Contents

What is Microsoft Defender for Cloud?

2/15/2022 • 10 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud is a tool for security posture management and threat protection. It strengthens the security posture of your cloud resources, and with its integrated Microsoft Defender plans, Defender for Cloud protects workloads running in Azure, hybrid, and other cloud platforms.

Defender for Cloud provides the tools needed to harden your resources, track your security posture, protect against cyber attacks, and streamline security management. Because it's natively integrated, deployment of Defender for Cloud is easy, providing you with simple auto provisioning to secure your resources by default.

Defender for Cloud fills three vital needs as you manage the security of your resources and workloads in the cloud and on-premises:

SECURITY REQUIREMENT	DEFENDER FOR CLOUD SOLUTION
Continuous assessment - Understand your current security posture.	**Secure score** - A single score so that you can tell, at a glance, your current security situation: the higher the score, the lower the identified risk level.
Secure - Harden all connected resources and services.	**Security recommendations** - Customized and prioritized hardening tasks to improve your posture. You implement a recommendation by following the detailed remediation steps provided in the recommendation. For many recommendations, Defender for Cloud offers a "Fix" button for automated implementation!
Defend - Detect and resolve threats to those resources and services.	**Security alerts** - With the enhanced security features enabled, Defender for Cloud detects threats to your resources and workloads. These alerts appear in the Azure portal and Defender for Cloud can also send them by email to the relevant personnel in your organization. Alerts can also be streamed to SIEM, SOAR, or IT Service Management solutions as required.

Posture management and workload protection

Microsoft Defender for Cloud's features cover the two broad pillars of cloud security: cloud security posture management and cloud workload protection.

Cloud security posture management (CSPM)

In Defender for Cloud, the posture management features provide:

- **Visibility** - to help you understand your current security situation
- **Hardening guidance** - to help you efficiently and effectively improve your security

The central feature in Defender for Cloud that enables you to achieve those goals is **secure score**. Defender for Cloud continually assesses your resources, subscriptions, and organization for security issues. It then aggregates all the findings into a single score so that you can tell, at a glance, your current security situation: the higher the score, the lower the identified risk level.

When you open Defender for Cloud for the first time, it will meet the visibility and strengthening goals as follows:

1. **Generate a secure score** for your subscriptions based on an assessment of your connected resources compared with the guidance in Azure Security Benchmark. Use the score to understand your security posture, and the compliance dashboard to review your compliance with the built-in benchmark. When you've enabled the enhanced security features, you can customize the standards used to assess your compliance, and add other regulations (such as NIST and Azure CIS) or organization-specific security requirements.
2. **Provide hardening recommendations** based on any identified security misconfigurations and weaknesses. Use these security recommendations to strengthen the security posture of your organization's Azure, hybrid, and multi-cloud resources.

Learn more about secure score.

Cloud workload protection (CWP)

Defender for Cloud offers security alerts that are powered by Microsoft Threat Intelligence. It also includes a range of advanced, intelligent, protections for your workloads. The workload protections are provided through Microsoft Defender plans specific to the types of resources in your subscriptions. For example, you can enable **Microsoft Defender for Storage** to get alerted about suspicious activities related to your Azure Storage accounts.

Azure, hybrid, and multi-cloud protections

Because Defender for Cloud is an Azure-native service, many Azure services are monitored and protected without needing any deployment.

When necessary, Defender for Cloud can automatically deploy a Log Analytics agent to gather security-related data. For Azure machines, deployment is handled directly. For hybrid and multi-cloud environments, Microsoft Defender plans are extended to non Azure machines with the help of Azure Arc. CSPM features are extended to multi-cloud machines without the need for any agents (see Defend resources running on other clouds).

Azure-native protections

Defender for Cloud helps you detect threats across:

- **Azure PaaS services** - Detect threats targeting Azure services including Azure App Service, Azure SQL, Azure Storage Account, and more data services. You can also perform anomaly detection on your Azure activity logs using the native integration with Microsoft Defender for Cloud Apps (formerly known as Microsoft Cloud App Security).

- **Azure data services** - Defender for Cloud includes capabilities that help you automatically classify your data in Azure SQL. You can also get assessments for potential vulnerabilities across Azure SQL and Storage services, and recommendations for how to mitigate them.
- **Networks** - Defender for Cloud helps you limit exposure to brute force attacks. By reducing access to virtual machine ports, using the just-in-time VM access, you can harden your network by preventing unnecessary access. You can set secure access policies on selected ports, for only authorized users, allowed source IP address ranges or IP addresses, and for a limited amount of time.

Defend your hybrid resources

In addition to defending your Azure environment, you can add Defender for Cloud capabilities to your hybrid cloud environment to protect your non-Azure servers. To help you focus on what matters the most, you'll get customized threat intelligence and prioritized alerts according to your specific environment.

To extend protection to on-premises machines, deploy Azure Arc and enable Defender for Cloud's enhanced security features. Learn more in Add non-Azure machines with Azure Arc.

Defend resources running on other clouds

Defender for Cloud can protect resources in other clouds (such as AWS and GCP).

For example, if you've connected an Amazon Web Services (AWS) account to an Azure subscription, you can enable any of these protections:

- **Defender for Cloud's CSPM features** extend to your AWS resources. This agentless plan assesses your AWS resources according to AWS-specific security recommendations and these are included in your secure score. The resources will also be assessed for compliance with built-in standards specific to AWS (AWS CIS, AWS PCI DSS, and AWS Foundational Security Best Practices). Defender for Cloud's asset inventory page is a multi-cloud enabled feature helping you manage your AWS resources alongside your Azure resources.
- **Microsoft Defender for Kubernetes** extends its container threat detection and advanced defenses to your **Amazon EKS Linux clusters**.
- **Microsoft Defender for servers** brings threat detection and advanced defenses to your Windows and Linux EC2 instances. This plan includes the integrated license for Microsoft Defender for Endpoint, security baselines and OS level assessments, vulnerability assessment scanning, adaptive application controls (AAC), file integrity monitoring (FIM), and more.

Learn more about connecting your AWS and GCP accounts to Microsoft Defender for Cloud.

Vulnerability assessment and management

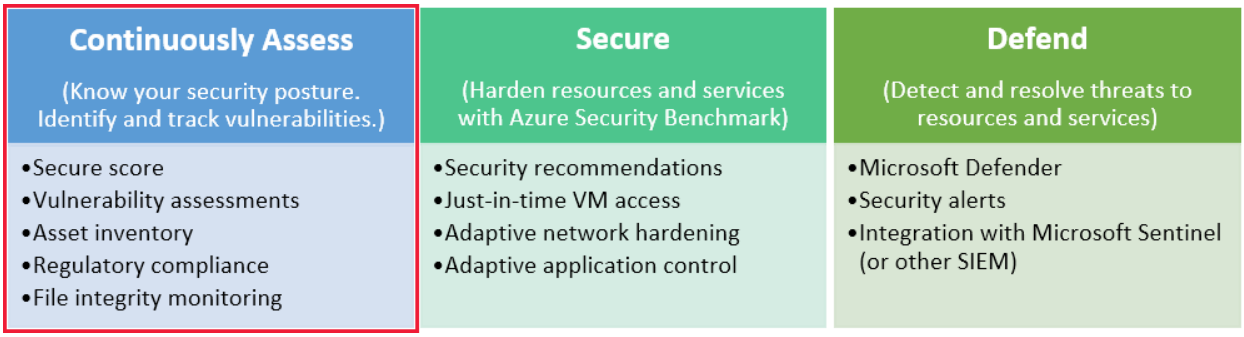

Defender for Cloud includes vulnerability assessment solutions for your virtual machines, container registries, and SQL servers as part of the enhanced security features. Some of the scanners are powered by Qualys. But you don't need a Qualys license, or even a Qualys account - everything's handled seamlessly inside Defender for Cloud.

Microsoft Defender for servers includes automatic, native integration with Microsoft Defender for Endpoint. Learn more, Protect your endpoints with Defender for Cloud's integrated EDR solution: Microsoft Defender for

Endpoint. With this integration enabled, you'll have access to the vulnerability findings from **Microsoft threat and vulnerability management**. Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management.

Review the findings from these vulnerability scanners and respond to them all from within Defender for Cloud. This broad approach brings Defender for Cloud closer to being the single pane of glass for all of your cloud security efforts.

Learn more on the following pages:

- Defender for Cloud's integrated Qualys scanner for Azure and hybrid machines
- Identify vulnerabilities in images in Azure container registries

Optimize and improve security by configuring recommended controls

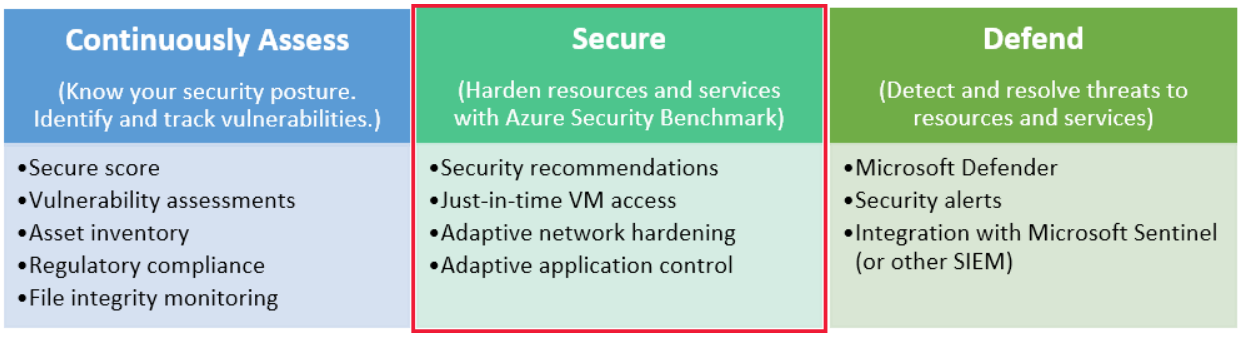

It's a security basic to know and make sure your workloads are secure, and it starts with having tailored security policies in place. Because policies in Defender for Cloud are built on top of Azure Policy controls, you're getting the full range and flexibility of a **world-class policy solution**. In Defender for Cloud, you can set your policies to run on management groups, across subscriptions, and even for a whole tenant.

Defender for Cloud continuously discovers new resources that are being deployed across your workloads and assesses whether they are configured according to security best practices. If not, they're flagged and you get a prioritized list of recommendations for what you need to fix. Recommendations help you reduce the attack surface across each of your resources.

The list of recommendations is enabled and supported by the Azure Security Benchmark. This Microsoft-authored, Azure-specific, benchmark provides a set of guidelines for security and compliance best practices based on common compliance frameworks. Learn more in Introduction to Azure Security Benchmark.

In this way, Defender for Cloud enables you not just to set security policies, but to *apply secure configuration standards across your resources*.

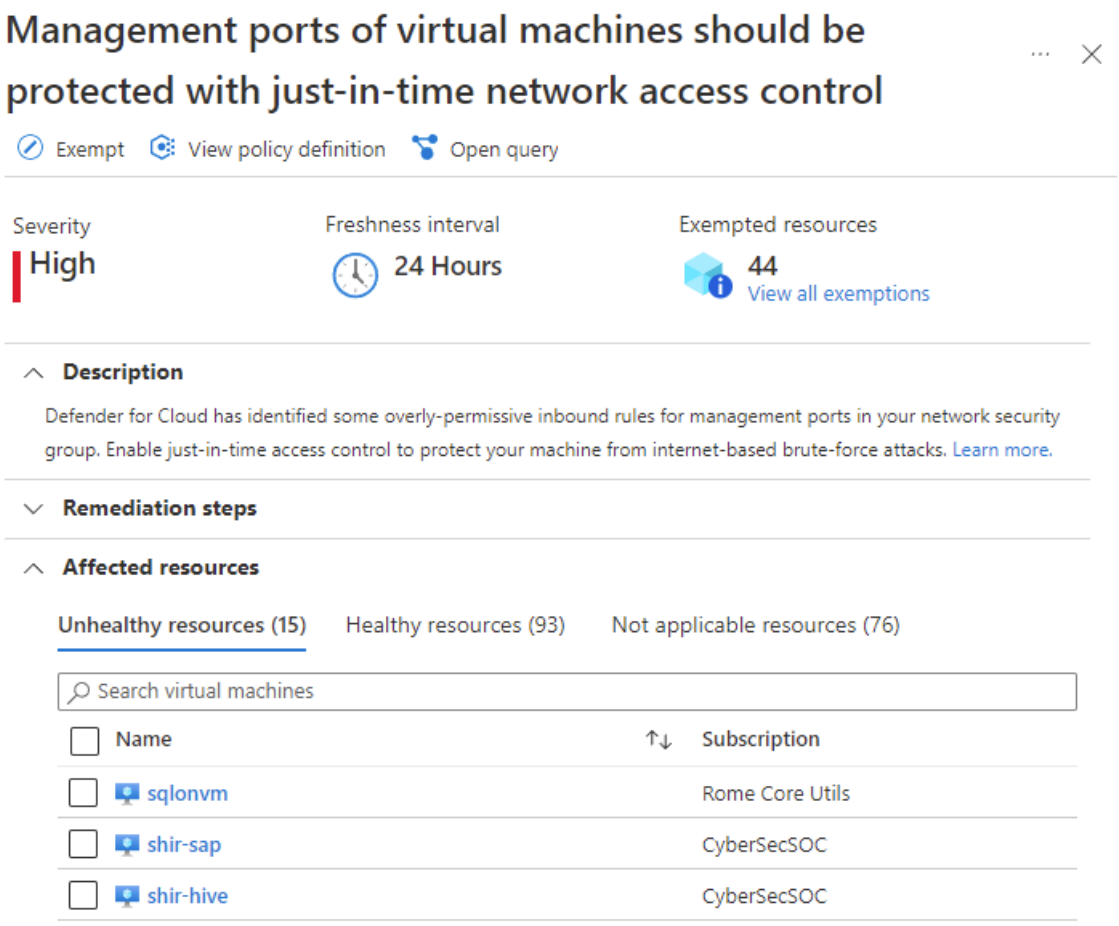

To help you understand how important each recommendation is to your overall security posture, Defender for Cloud groups the recommendations into security controls and adds a **secure score** value to each control. This is crucial in enabling you to **prioritize your security work**.

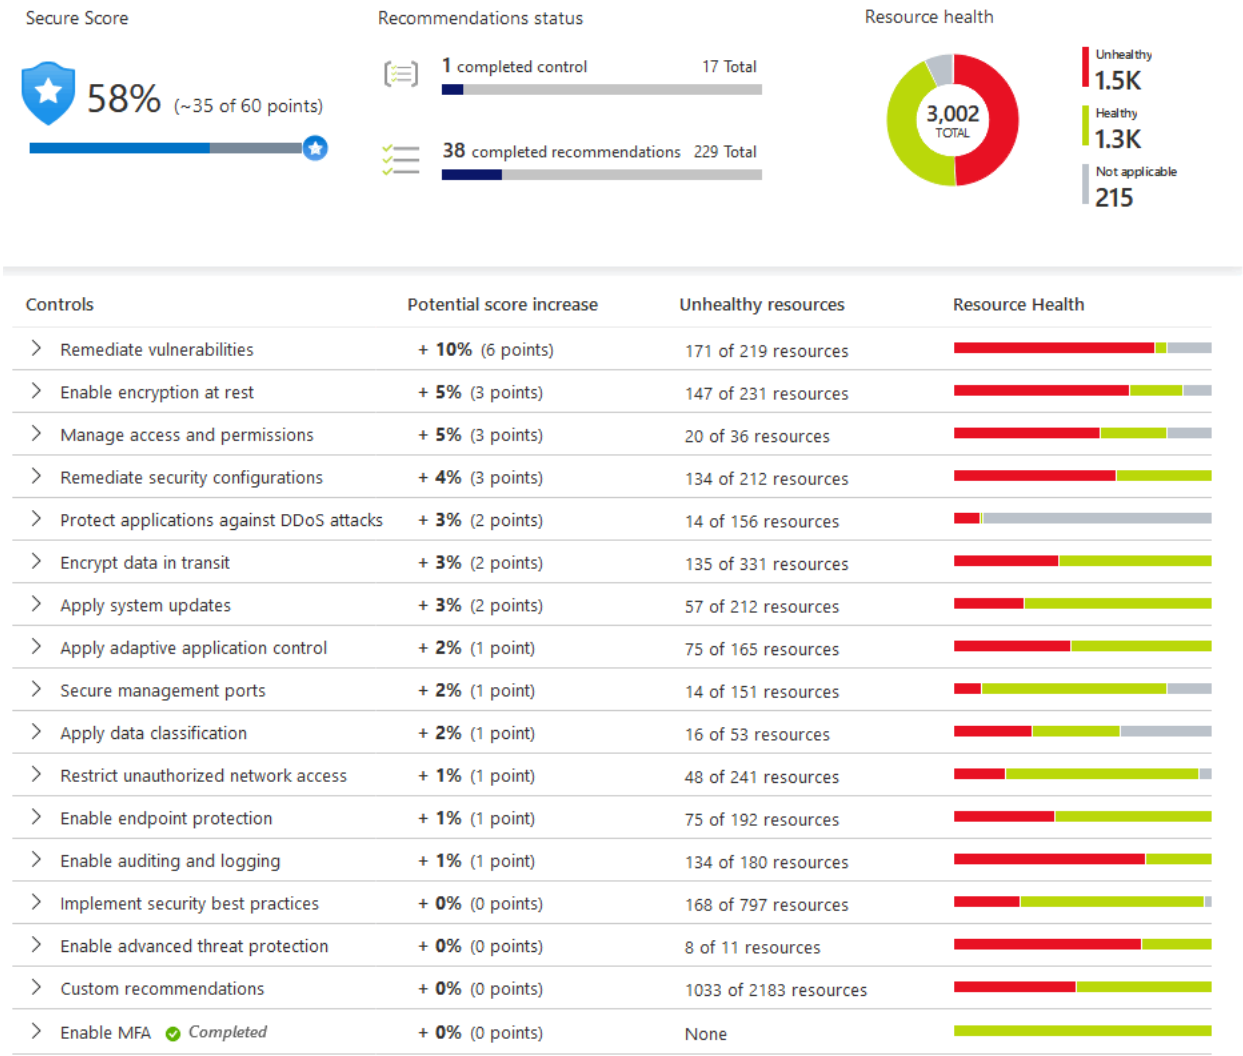

Defend against threats

Continuously Assess (Know your security posture. Identify and track vulnerabilities.)	Secure (Harden resources and services with Azure Security Benchmark)	Defend (Detect and resolve threats to resources and services)
•Secure score •Vulnerability assessments •Asset inventory •Regulatory compliance •File integrity monitoring	•Security recommendations •Just-in-time VM access •Adaptive network hardening •Adaptive application control	•Microsoft Defender •Security alerts •Integration with Microsoft Sentinel (or other SIEM)

Defender for Cloud provides:

- **Security alerts** - When Defender for Cloud detects a threat in any area of your environment, it generates a security alert. These alerts describe details of the affected resources, suggested remediation steps, and in some cases an option to trigger a logic app in response. Whether an alert is generated by Defender for Cloud, or received by Defender for Cloud from an integrated security product, you can export it. To export your alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Stream alerts to a SIEM, SOAR, or IT Service Management solution. Defender for Cloud's threat protection includes fusion kill-chain analysis, which automatically correlates alerts in your environment based on cyber kill-chain analysis, to help you better understand the full story of an attack campaign, where it started and what kind of impact it had on your resources. Defender for Cloud's supported kill chain intents are based on version 7 of the MITRE ATT&CK matrix.
- **Advanced threat protection features** for virtual machines, SQL databases, containers, web

applications, your network, and more - Protections include securing the management ports of your VMs with just-in-time access, and adaptive application controls to create allowlists for what apps should and shouldn't run on your machines.

The **Defender plans** page of Microsoft Defender for Cloud offers the following plans for comprehensive defenses for the compute, data, and service layers of your environment:

- Microsoft Defender for servers
- Microsoft Defender for Storage
- Microsoft Defender for SQL
- Microsoft Defender for Containers
- Microsoft Defender for App Service
- Microsoft Defender for Key Vault
- Microsoft Defender for Resource Manager
- Microsoft Defender for DNS
- Microsoft Defender for open-source relational databases

Use the advanced protection tiles in the workload protections dashboard to monitor and configure each of these protections.

> **TIP**
> Microsoft Defender for IoT is a separate product. You'll find all the details in Introducing Microsoft Defender for IoT.

Next steps

- To get started with Defender for Cloud, you need a subscription to Microsoft Azure. If you don't have a subscription, sign up for a free trial.
- Defender for Cloud's free plan is enabled on all your current Azure subscriptions when you visit the Defender for Cloud pages in the Azure portal for the first time, or if enabled programmatically via the REST API. To take advantage of advanced security management and threat detection capabilities, you must enable the enhanced security features. These features are free for the first 30 days. Learn more about the pricing.
- If you're ready to enable enhanced security features now, Quickstart: Enable enhanced security features walks you through the steps.

Enable Microsoft Defender plans

Microsoft Defender for Cloud's enhanced security features

2/15/2022 • 9 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The enhanced security features are free for the first 30 days. At the end of 30 days, if you decide to continue using the service, we'll automatically start charging for usage.

You can upgrade from the **Environment settings** page, as described in Quickstart: Enable enhanced security features. For pricing details in your local currency or region, see the pricing page.

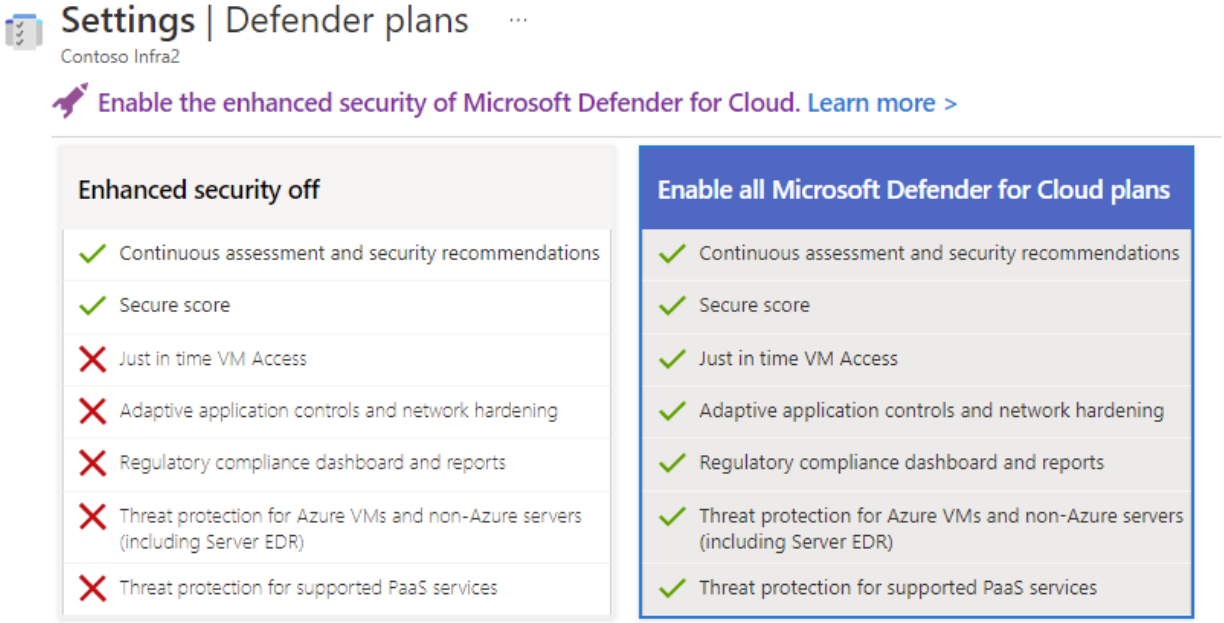

What are the benefits of enabling enhanced security features?

Defender for Cloud is offered in two modes:

- **Without enhanced security features** (Free) - Defender for Cloud is enabled for free on all your Azure subscriptions when you visit the workload protection dashboard in the Azure portal for the first time, or if enabled programmatically via API. Using this free mode provides the secure score and its related features: security policy, continuous security assessment, and actionable security recommendations to help you protect your Azure resources.
- **Defender for Cloud with all enhanced security features** - Enabling enhanced security extends the capabilities of the free mode to workloads running in private and other public clouds, providing unified security management and threat protection across your hybrid cloud workloads. Some of the major benefits include:
  - **Microsoft Defender for Endpoint** - Microsoft Defender for servers includes Microsoft Defender for

Endpoint for comprehensive endpoint detection and response (EDR). Learn more about the benefits of using Microsoft Defender for Endpoint together with Defender for Cloud in Use Defender for Cloud's integrated EDR solution.

- **Vulnerability assessment for virtual machines, container registries, and SQL resources** - Easily enable vulnerability assessment solutions to discover, manage, and resolve vulnerabilities. View, investigate, and remediate the findings directly from within Defender for Cloud.
- **Multi-cloud security** - Connect your accounts from Amazon Web Services (AWS) and Google Cloud Platform (GCP) to protect resources and workloads on those platforms with a range of Microsoft Defender for Cloud security features.
- **Hybrid security** – Get a unified view of security across all of your on-premises and cloud workloads. Apply security policies and continuously assess the security of your hybrid cloud workloads to ensure compliance with security standards. Collect, search, and analyze security data from multiple sources, including firewalls and other partner solutions.
- **Threat protection alerts** - Advanced behavioral analytics and the Microsoft Intelligent Security Graph provide an edge over evolving cyber-attacks. Built-in behavioral analytics and machine learning can identify attacks and zero-day exploits. Monitor networks, machines, data stores (SQL servers hosted inside and outside Azure, Azure SQL databases, Azure SQL Managed Instance, and Azure Storage) and cloud services for incoming attacks and post-breach activity. Streamline investigation with interactive tools and contextual threat intelligence.
- **Track compliance with a range of standards** - Defender for Cloud continuously assesses your hybrid cloud environment to analyze the risk factors according to the controls and best practices in Azure Security Benchmark. When you enable the enhanced security features, you can apply a range of other industry standards, regulatory standards, and benchmarks according to your organization's needs. Add standards and track your compliance with them from the regulatory compliance dashboard.
- **Access and application controls** - Block malware and other unwanted applications by applying machine learning powered recommendations adapted to your specific workloads to create allow and blocklists. Reduce the network attack surface with just-in-time, controlled access to management ports on Azure VMs. Access and application controls drastically reduce exposure to brute force and other network attacks.
- **Container security features** - Benefit from vulnerability management and real-time threat protection on your containerized environments. Charges are based on the number of unique container images pushed to your connected registry. After an image has been scanned once, you won't be charged for it again unless it's modified and pushed once more.
- **Breadth threat protection for resources connected to Azure** - Cloud-native threat protection for the Azure services common to all of your resources: Azure Resource Manager, Azure DNS, Azure network layer, and Azure Key Vault. Defender for Cloud has unique visibility into the Azure management layer and the Azure DNS layer, and can therefore protect cloud resources that are connected to those layers.

FAQ - Pricing and billing

- How can I track who in my organization enabled a Microsoft Defender plan in Defender for Cloud?
- What are the plans offered by Defender for Cloud?
- How do I enable Defender for Cloud's enhanced security for my subscription?
- Can I enable Microsoft Defender for servers on a subset of servers in my subscription?
- If I already have a license for Microsoft Defender for Endpoint can I get a discount for Defender for servers?
- My subscription has Microsoft Defender for servers enabled, do I pay for not-running servers?
- Will I be charged for machines without the Log Analytics agent installed?
- If a Log Analytics agent reports to multiple workspaces, will I be charged twice?

- If a Log Analytics agent reports to multiple workspaces, is the 500-MB free data ingestion available on all of them?
- Is the 500-MB free data ingestion calculated for an entire workspace or strictly per machine?
- What data types are included in the 500-MB data daily allowance?

How can I track who in my organization enabled a Microsoft Defender plan in Defender for Cloud?

Azure Subscriptions may have multiple administrators with permissions to change the pricing settings. To find out which user made a change, use the Azure Activity Log.

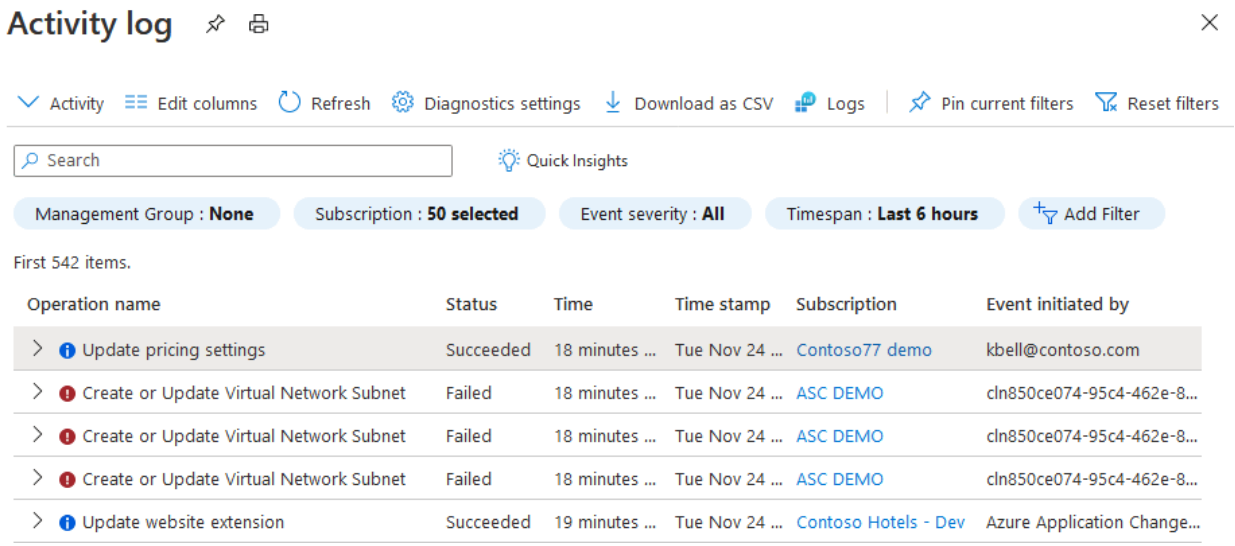

If the user's info isn't listed in the **Event initiated by** column, explore the event's JSON for the relevant details.

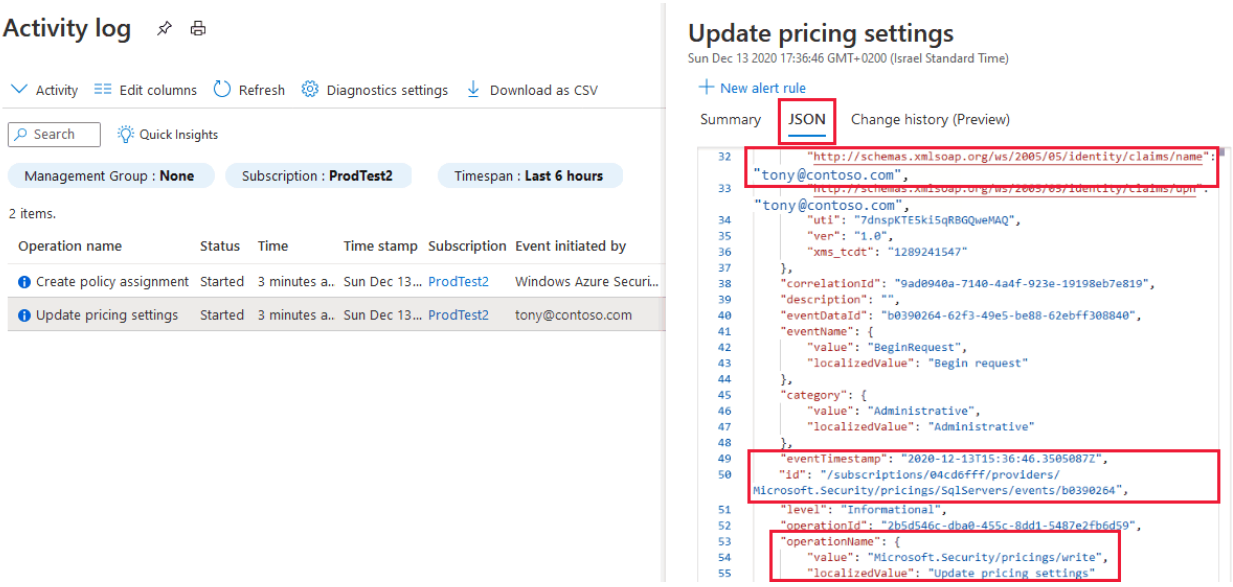

What are the plans offered by Defender for Cloud?

The free offering from Microsoft Defender for Cloud offers the secure score and related tools. Enabling enhanced security turns on all of the Microsoft Defender plans to provide a range of security benefits for all your resources in Azure, hybrid, and multi-cloud environments.

How do I enable Defender for Cloud's enhanced security for my subscription?

You can use any of the following ways to enable enhanced security for your subscription:

METHOD	INSTRUCTIONS
Defender for Cloud pages of the Azure portal	Enable enhanced protections

METHOD	INSTRUCTIONS
REST API	Pricings API
Azure CLI	az security pricing
PowerShell	Set-AzSecurityPricing
Azure Policy	Bundle Pricings

Can I enable Microsoft Defender for servers on a subset of servers in my subscription?

No. When you enable Microsoft Defender for servers on a subscription, all the machines in the subscription will be protected by Defender for servers.

An alternative is to enable Microsoft Defender for servers at the Log Analytics workspace level. If you do this, only servers reporting to that workspace will be protected and billed. However, several capabilities will be unavailable. These include just-in-time VM access, network detections, regulatory compliance, adaptive network hardening, adaptive application control, and more.

If I already have a license for Microsoft Defender for Endpoint can I get a discount for Defender for servers?

If you've already got a license for **Microsoft Defender for Endpoint for Servers**, you won't have to pay for that part of your Microsoft Defender for servers license. Learn more about this license.

To request your discount, contact Defender for Cloud's support team. You'll need to provide the relevant workspace ID, region, and number of Microsoft Defender for Endpoint for servers licenses applied for machines in the given workspace.

The discount will be effective starting from the approval date, and won't take place retroactively.

My subscription has Microsoft Defender for servers enabled, do I pay for not-running servers?

No. When you enable Microsoft Defender for servers on a subscription, you won't be charged for any machines that are in the deallocated power state while they're in that state. Machines are billed according to their power state as shown in the following table:

STATE	DESCRIPTION	INSTANCE USAGE BILLED
Starting	VM is starting up.	Not billed
Running	Normal working state for a VM	Billed
Stopping	This is a transitional state. When completed, it will show as Stopped.	Billed
Stopped	The VM has been shut down from within the guest OS or using the PowerOff APIs. Hardware is still allocated to the VM and it remains on the host.	Billed
Deallocating	Transitional state. When completed, the VM will show as Deallocated.	Not billed

STATE	DESCRIPTION	INSTANCE USAGE BILLED
Deallocated	The VM has been stopped successfully and removed from the host.	Not billed

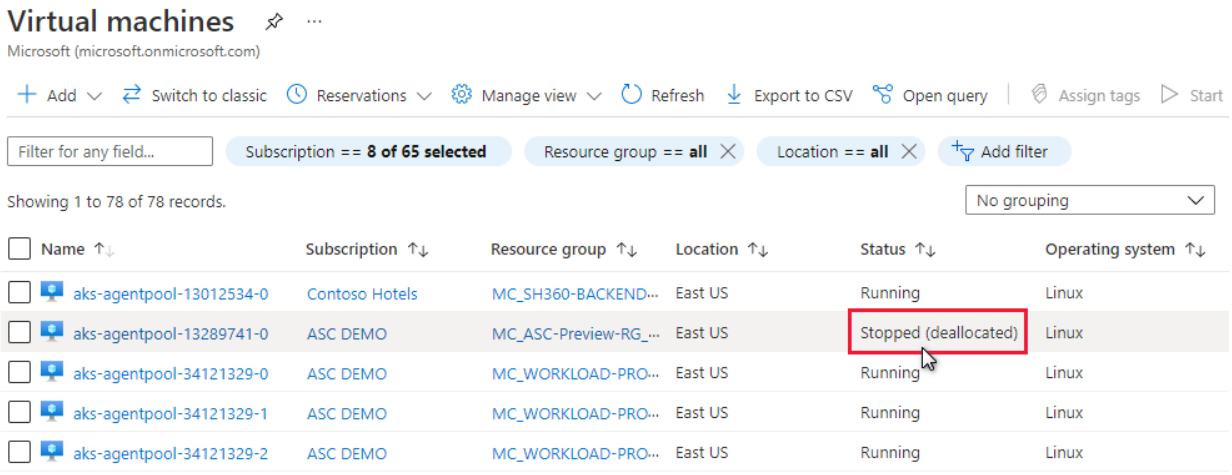

Will I be charged for machines without the Log Analytics agent installed?

Yes. When you enable Microsoft Defender for servers on a subscription, the machines in that subscription get a range of protections even if you haven't installed the Log Analytics agent. This is applicable for Azure virtual machines, Azure virtual machine scale sets instances, and Azure Arc-enabled servers.

If a Log Analytics agent reports to multiple workspaces, will I be charged twice?

Yes. If you've configured your Log Analytics agent to send data to two or more different Log Analytics workspaces (multi-homing), you'll be charged for every workspace that has a 'Security' or 'AntiMalware' solution installed.

If a Log Analytics agent reports to multiple workspaces, is the 500-MB free data ingestion available on all of them?

Yes. If you've configured your Log Analytics agent to send data to two or more different Log Analytics workspaces (multi-homing), you'll get 500-MB free data ingestion. It's calculated per node, per reported workspace, per day, and available for every workspace that has a 'Security' or 'AntiMalware' solution installed. You'll be charged for any data ingested over the 500-MB limit.

Is the 500-MB free data ingestion calculated for an entire workspace or strictly per machine?

You'll get 500-MB free data ingestion per day, for every Windows machine connected to the workspace. Specifically for security data types directly collected by Defender for Cloud.

This data is a daily rate averaged across all nodes. So even if some machines send 100-MB and others send 800-MB, if the total doesn't exceed the **[number of machines] x 500-MB** free limit, you won't be charged extra.

What data types are included in the 500-MB data daily allowance?

Defender for Cloud's billing is closely tied to the billing for Log Analytics. Microsoft Defender for servers provides a 500 MB/node/day allocation for Windows machines against the following subset of security data types:

- SecurityAlert
- SecurityBaseline
- SecurityBaselineSummary
- SecurityDetection
- SecurityEvent
- WindowsFirewall

- MaliciousIPCommunication
- SysmonEvent
- ProtectionStatus
- Update and UpdateSummary data types when the Update Management solution is not running on the workspace or solution targeting is enabled

If the workspace is in the legacy Per Node pricing tier, the Defender for Cloud and Log Analytics allocations are combined and applied jointly to all billable ingested data.

Next steps

This article explained Defender for Cloud's pricing options. For related material, see:

- How to optimize your Azure workload costs
- Pricing details according to currency or region
- You may want to manage your costs and limit the amount of data collected for a solution by limiting it to a particular set of agents. Use solution targeting to apply a scope to the solution and target a subset of computers in the workspace. If you're using solution targeting, Defender for Cloud lists the workspace as not having a solution.

What's new in Microsoft Defender for Cloud?

2/15/2022 • 35 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud is in active development and receives improvements on an ongoing basis. To stay up to date with the most recent developments, this page provides you with information about new features, bug fixes, and deprecated functionality.

This page is updated frequently, so revisit it often.

To learn about *planned* changes that are coming soon to Defender for Cloud, see Important upcoming changes to Microsoft Defender for Cloud.

> **TIP**
> If you're looking for items older than six months, you'll find them in the Archive for What's new in Microsoft Defender for Cloud.

January 2022

Updates in January include:

- Microsoft Defender for Resource Manager updated with new alerts and greater emphasis on high-risk operations mapped to MITRE ATT&CK® Matrix
- Recommendations to enable Microsoft Defender plans on workspaces (in preview)
- Auto provision Log Analytics agent to Azure Arc-enabled machines (preview)
- Deprecated the recommendation to classify sensitive data in SQL databases
- Communication with suspicious domain alert expanded to included known Log4Shell-related domains
- 'Copy alert JSON' button added to security alert details pane
- Renamed two recommendations
- Deprecate Kubernetes cluster containers should only listen on allowed ports policy
- Added 'Active Alerts' workbook
- 'System update' recommendation added to government cloud

Microsoft Defender for Resource Manager updated with new alerts and greater emphasis on high-risk operations mapped to MITRE ATT&CK® Matrix

The cloud management layer is a crucial service connected to all your cloud resources. Because of this, it is also a potential target for attackers. Consequently, we recommend security operations teams closely monitor the resource management layer.

Microsoft Defender for Resource Manager automatically monitors the resource management operations in your organization, whether they're performed through the Azure portal, Azure REST APIs, Azure CLI, or other Azure programmatic clients. Defender for Cloud runs advanced security analytics to detect threats and alerts you about suspicious activity.

The plan's protections greatly enhance an organization's resiliency against attacks from threat actors and significantly increase the number of Azure resources protected by Defender for Cloud.

In December 2020, we introduced the preview of Defender for Resource Manager, and in May 2021 the plan was release for general availability.

With this update, we've comprehensively revised the focus of the Microsoft Defender for Resource Manager plan. The updated plan includes many **new alerts focused on identifying suspicious invocation of high-risk operations**. These new alerts provide extensive monitoring for attacks across the *complete* MITRE ATT&CK® matrix for cloud-based techniques.

This matrix covers the following range of potential intentions of threat actors who may be targeting your organization's resources: *Initial Access, Execution, Persistence, Privilege Escalation, Defense Evasion, Credential Access, Discovery, Lateral Movement, Collection, Exfiltration, and Impact*.

The new alerts for this Defender plan cover these intentions as shown in the following table.

> **TIP**
> These alerts also appear in the alerts reference page.

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Suspicious invocation of a high-risk 'Initial Access' operation detected (Preview) (ARM_AnomalousOperation.InitialAccess)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to access restricted resources. The identified operations are designed to allow administrators to efficiently access their environments. While this activity may be legitimate, a threat actor might utilize such operations to gain initial access to restricted resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Suspicious invocation of a high-risk 'Execution' operation detected (Preview) (ARM_AnomalousOperation.Execution)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation on a machine in your subscription which might indicate an attempt to execute code. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Execution	Medium
Suspicious invocation of a high-risk 'Persistence' operation detected (Preview) (ARM_AnomalousOperation.Persistence)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to establish persistence. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to establish persistence in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Persistence	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Suspicious invocation of a high-risk 'Privilege Escalation' operation detected (Preview) (ARM_AnomalousOperation.PrivilegeEscalation)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to escalate privileges. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to escalate privileges while compromising resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Privilege Escalation	Medium
Suspicious invocation of a high-risk 'Defense Evasion' operation detected (Preview) (ARM_AnomalousOperation.DefenseEvasion)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to evade defenses. The identified operations are designed to allow administrators to efficiently manage the security posture of their environments. While this activity may be legitimate, a threat actor might utilize such operations to avoid being detected while compromising resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Suspicious invocation of a high-risk 'Credential Access' operation detected (Preview) (ARM_AnomalousOperation.CredentialAccess)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to access credentials. The identified operations are designed to allow administrators to efficiently access their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Credential Access	Medium
Suspicious invocation of a high-risk 'Lateral Movement' operation detected (Preview) (ARM_AnomalousOperation.LateralMovement)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to perform lateral movement. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to compromise additional resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Lateral Movement	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Suspicious invocation of a high-risk 'Data Collection' operation detected (Preview) (ARM_AnomalousOperation.Collection)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to collect data. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to collect sensitive data on resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Collection	Medium
Suspicious invocation of a high-risk 'Impact' operation detected (Preview) (ARM_AnomalousOperation.Impact)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempted configuration change. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Impact	Medium

In addition, these two alerts from this plan have come out of preview:

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Azure Resource Manager operation from suspicious IP address (ARM_OperationFromSuspiciousIP)	Microsoft Defender for Resource Manager detected an operation from an IP address that has been marked as suspicious in threat intelligence feeds.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (INTENTIONS)	SEVERITY
Azure Resource Manager operation from suspicious proxy IP address (ARM_OperationFromSuspiciousProxyIP)	Microsoft Defender for Resource Manager detected a resource management operation from an IP address that is associated with proxy services, such as TOR. While this behavior can be legitimate, it's often seen in malicious activities, when threat actors try to hide their source IP.	Defense Evasion	Medium

Recommendations to enable Microsoft Defender plans on workspaces (in preview)

To benefit from all of the security features available from Microsoft Defender for servers and Microsoft Defender for SQL on machines, the plans must be enabled on **both** the subscription and workspace levels.

When a machine is in a subscription with one of these plan enabled, you'll be billed for the full protections. However, if that machine is reporting to a workspace *without* the plan enabled, you won't actually receive those benefits.

We've added two recommendations that highlight workspaces without these plans enabled, that nevertheless have machines reporting to them from subscriptions that *do* have the plan enabled.

The two recommendations, which both offer automated remediation (the 'Fix' action), are:

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for servers should be enabled on workspaces	Microsoft Defender for servers brings threat detection and advanced defenses for your Windows and Linux machines. With this Defender plan enabled on your subscriptions but not on your workspaces, you're paying for the full capability of Microsoft Defender for servers but missing out on some of the benefits. When you enable Microsoft Defender for servers on a workspace, all machines reporting to that workspace will be billed for Microsoft Defender for servers - even if they're in subscriptions without Defender plans enabled. Unless you also enable Microsoft Defender for servers on the subscription, those machines won't be able to take advantage of just-in-time VM access, adaptive application controls, and network detections for Azure resources. Learn more in Introduction to Microsoft Defender for servers. (No related policy)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for SQL on machines should be enabled on workspaces	Microsoft Defender for servers brings threat detection and advanced defenses for your Windows and Linux machines. With this Defender plan enabled on your subscriptions but not on your workspaces, you're paying for the full capability of Microsoft Defender for servers but missing out on some of the benefits. When you enable Microsoft Defender for servers on a workspace, all machines reporting to that workspace will be billed for Microsoft Defender for servers - even if they're in subscriptions without Defender plans enabled. Unless you also enable Microsoft Defender for servers on the subscription, those machines won't be able to take advantage of just-in-time VM access, adaptive application controls, and network detections for Azure resources. Learn more in Introduction to Microsoft Defender for servers. (No related policy)	Medium

Auto provision Log Analytics agent to Azure Arc-enabled machines (preview)

Defender for Cloud uses the Log Analytics agent to gather security-related data from machines. The agent reads various security-related configurations and event logs and copies the data to your workspace for analysis.

Defender for Cloud's auto provisioning settings have a toggle for each type of supported extension, including the Log Analytics agent.

In a further expansion of our hybrid cloud features, we've added an option to auto provision the Log Analytics agent to machines connected to Azure Arc.

As with the other other auto provisioning options, this is configured at the subscription level.

When you enable this option, you'll be prompted for the workspace.

> **NOTE**
>
> For this preview, you can't select the default workspaces that was created by Defender for Cloud. To ensure you receive the full set of security features available for the Azure Arc-enabled servers, verify that you have the relevant security solution installed on the selected workspace.

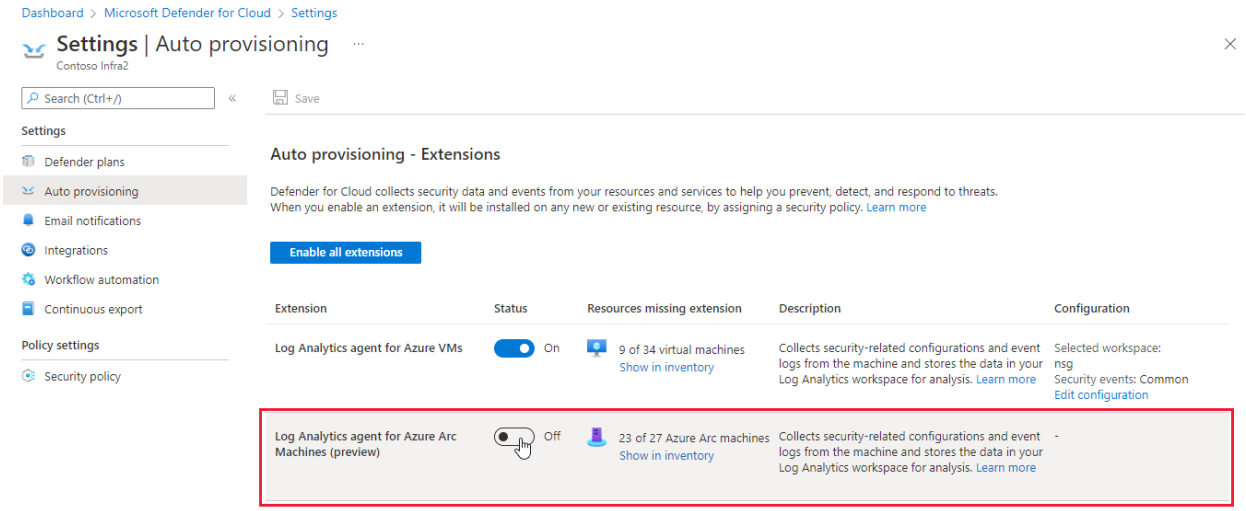

Deprecated the recommendation to classify sensitive data in SQL databases

We've removed the recommendation **Sensitive data in your SQL databases should be classified** as part of an overhaul of how Defender for Cloud identifies and protects sensitive date in your cloud resources.

Advance notice of this change appeared for the last six months in the Important upcoming changes to Microsoft Defender for Cloud page.

Communication with suspicious domain alert expanded to included known Log4Shell-related domains

The following alert was previously only available to organizations who'd enabled the Microsoft Defender for DNS plan.

With this update, the alert will also show for subscriptions with the Microsoft Defender for servers or Defender for App Service plan enabled.

In addition, Microsoft Threat Intelligence has expanded the list of known malicious domains to include domains associated with exploiting the widely publicised vulnerabilities associated with Log4j.

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS	SEVERITY
Communication with suspicious domain identified by threat intelligence (AzureDNS_ThreatIntelSuspectDomain)	Communication with suspicious domain was detected by analyzing DNS transactions from your resource and comparing against known malicious domains identified by threat intelligence feeds. Communication to malicious domains is frequently performed by attackers and could imply that your resource is compromised.	Initial Access / Persistence / Execution / Command And Control / Exploitation	Medium

'Copy alert JSON' button added to security alert details pane

To help our users quickly share an alert's details with others (for example, SOC analysts, resource owners, and developers) we've added the capability to easily extract all the details of a specific alert with one button from the security alert's details pane.

The new **Copy alert JSON** button puts the alert's details, in JSON format, into the user's clipboard.

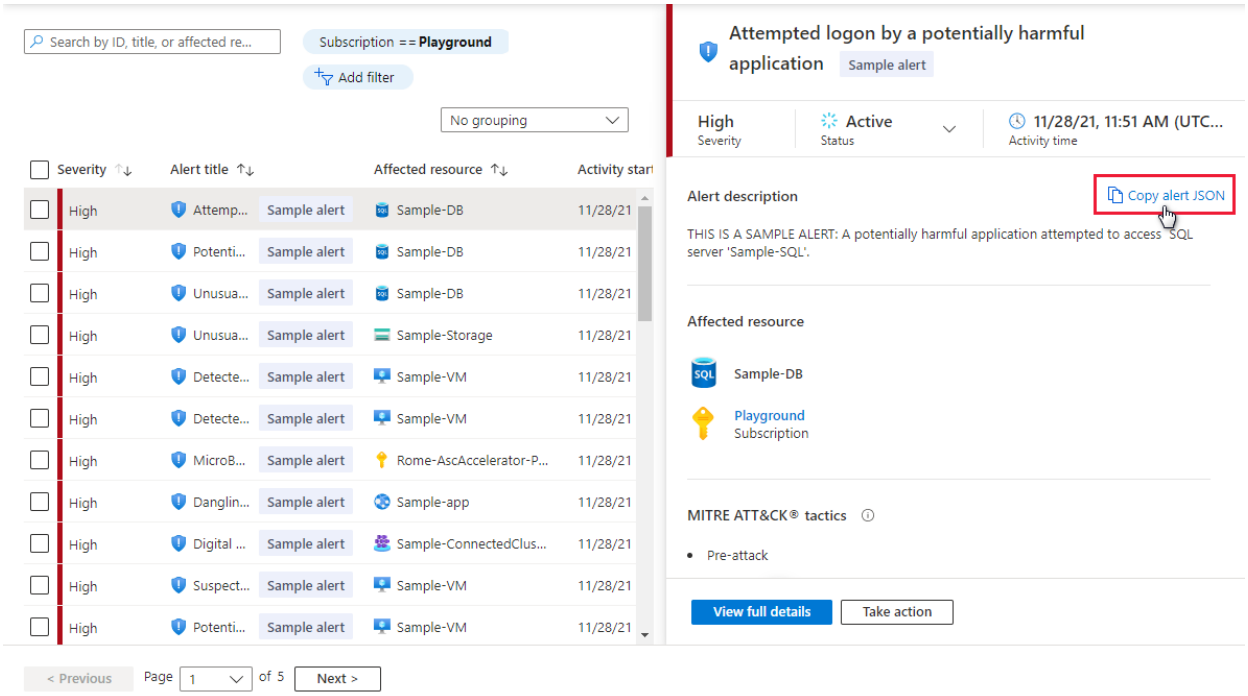

Renamed two recommendations

For consistency with other recommendation names, we've renamed the following two recommendations:

- Recommendation to resolve vulnerabilities discovered in running container images
  - Previous name: Vulnerabilities in running container images should be remediated (powered by Qualys)
  - New name: Running container images should have vulnerability findings resolved
- Recommendation to enable diagnostic logs for Azure App Service
  - Previous name: Diagnostic logs should be enabled in App Service
  - New name: Diagnostic logs in App Service should be enabled

Deprecate Kubernetes cluster containers should only listen on allowed ports policy

We have deprecated the **Kubernetes cluster containers should only listen on allowed ports** recommendation.

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes cluster containers should only listen on allowed ports	Restrict containers to listen only on allowed ports to secure access to the Kubernetes cluster. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	6.1.2

The **Services should listen on allowed ports only** recommendation should be used to limit ports that an application exposes to the internet.

Added 'Active Alert' workbook

To assist our users in their understanding of the active threats to their environments, and prioritize between

active alerts during the remediation process, we have added the Active Alerts workbook.

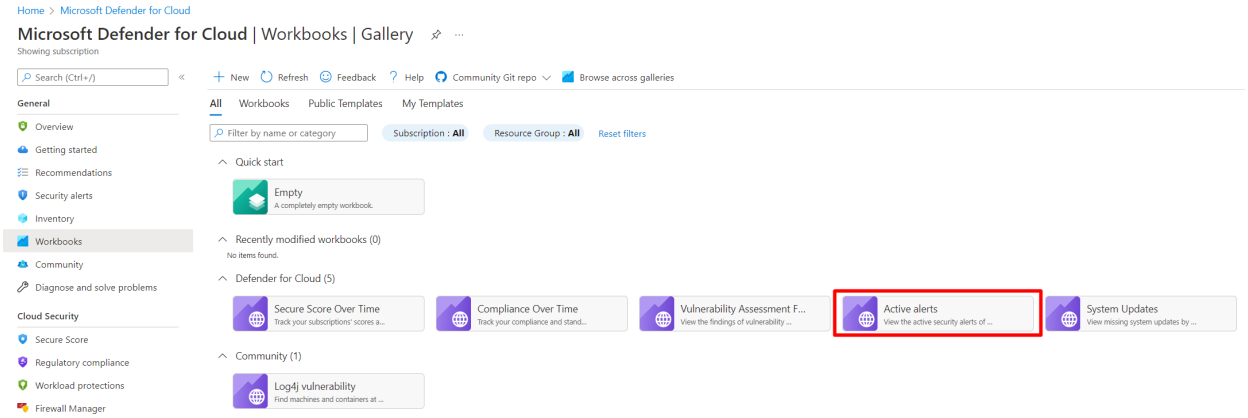

The active alerts workbook allows users to view a unified dashboard of their aggregated alerts by severity, type, tag, MITRE ATT&CK tactics, and location. Learn more in Use the 'Active Alerts' workbook.

'System update' recommendation added to government cloud

The 'System updates should be installed on your machines' recommendation is now available on all government clouds.

It's likely that this change will impact your government cloud subscription's secure score. We expect the change to lead to a decreased score, but it's possible the recommendation's inclusion might result in an increased score in some cases.

December 2021

Updates in December include:

- Microsoft Defender for Containers plan released for general availability (GA)
- New alerts for Microsoft Defender for Storage released for general availability (GA)
- Improvements to alerts for Microsoft Defender for Storage
- 'PortSweeping' alert removed from network layer alerts

Microsoft Defender for Containers plan released for general availability (GA)

Over two years ago, we introduced Defender for Kubernetes and Defender for container registries as part of the Azure Defender offering within Microsoft Defender for Cloud.

With the release of Microsoft Defender for Containers, we've merged these two existing Defender plans.

The new plan:

- **Combines the features of the two existing plans** - threat detection for Kubernetes clusters and vulnerability assessment for images stored in container registries
- **Brings new and improved features** - including multi-cloud support, host level threat detection with over **sixty** new Kubernetes-aware analytics, and vulnerability assessment for running images
- **Introduces Kubernetes-native at-scale onboarding** - by default, when you enable the plan all relevant components are configured to be deployed automatically

With this release, the availability and presentation of Defender for Kubernetes and Defender for container registries has changed as follows:

- New subscriptions - The two previous container plans are no longer available
- Existing subscriptions - Wherever they appear in the Azure portal, the plans are shown as **Deprecated** with instructions for how to upgrade to the newer plan

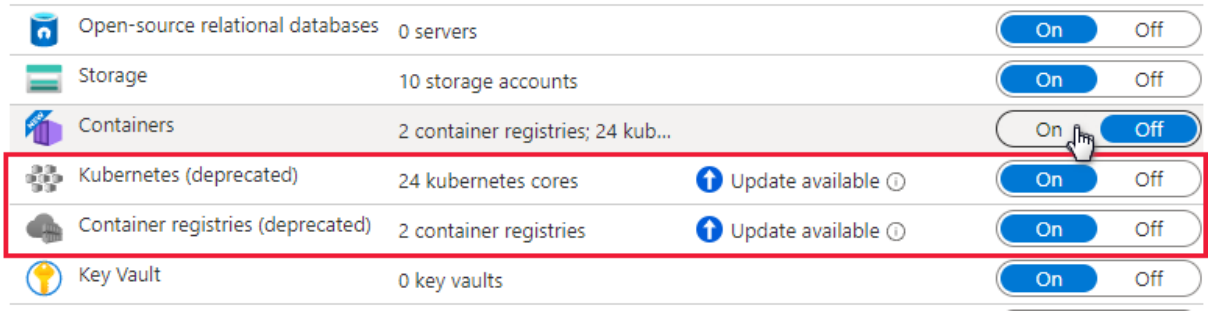

The new plan is free for the month of December 2021. For the potential changes to the billing from the old plans to Defender for Containers, and for more details on the benefits introduced with this plan, see Introducing Microsoft Defender for Containers.

For more information, see:

- Overview of Microsoft Defender for Containers
- Enable Microsoft Defender for Containers
- Introducing Microsoft Defender for Containers - Microsoft Tech Community
- Microsoft Defender for Containers | Defender for Cloud in the Field #3 - YouTube

New alerts for Microsoft Defender for Storage released for general availability (GA)

Threat actors use tools and scripts to scan for publicly open containers in the hope of finding misconfigured open storage containers with sensitive data.

Microsoft Defender for Storage detects these scanners so that you can block them and remediate your posture.

The preview alert that detected this was called **"Anonymous scan of public storage containers"**. To provide greater clarity about the suspicious events discovered, we've divided this into **two** new alerts. These alerts are relevant to Azure Blob Storage only.

We have improved the detection logic, updated the alert metadata, and changed the alert name and alert type.

These are the new alerts:

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTIC	SEVERITY

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTIC	SEVERITY
Publicly accessible storage containers successfully discovered (Storage.Blob_OpenContainersScanning.SuccessfulDiscovery)	A successful discovery of publicly open storage container(s) in your storage account was performed in the last hour by a scanning script or tool. This usually indicates a reconnaissance attack, where the threat actor tries to list blobs by guessing container names, in the hope of finding misconfigured open storage containers with sensitive data in them. The threat actor may use their own script or use known scanning tools like Microburst to scan for publicly open containers. ✔ Azure Blob Storage ✘ Azure Files ✘ Azure Data Lake Storage Gen2	Collection	Medium
Publicly accessible storage containers unsuccessfully scanned (Storage.Blob_OpenContainersScanning.FailedAttempt)	A series of failed attempts to scan for publicly open storage containers were performed in the last hour. This usually indicates a reconnaissance attack, where the threat actor tries to list blobs by guessing container names, in the hope of finding misconfigured open storage containers with sensitive data in them. The threat actor may use their own script or use known scanning tools like Microburst to scan for publicly open containers. ✔ Azure Blob Storage ✘ Azure Files ✘ Azure Data Lake Storage Gen2	Collection	Low

For more information, see:

- Threat matrix for storage services
- Introduction to Microsoft Defender for Storage
- List of alerts provided by Microsoft Defender for Storage

Improvements to alerts for Microsoft Defender for Storage

The initial access alerts now have improved accuracy and more data to support investigation.

Threat actors use various techniques in the initial access to gain a foothold within a network. Two of the Microsoft Defender for Storage alerts that detect behavioral anomalies in this stage now have improved detection logic and additional data to support investigations.

If you've configured automations or defined alert suppression rules for these alerts in the past, update them in accordance with these changes.

Detecting access from a Tor exit node

Access from a Tor exit node might indicate a threat actor trying to hide their identity.

The alert is now tuned to generate only for authenticated access, which results in higher accuracy and confidence that the activity is malicious. This enhancement reduces the benign positive rate.

An outlying pattern will have high severity, while less anomalous patterns will have medium severity.

The alert name and description have been updated. The AlertType remains unchanged.

- Alert name (old): Access from a Tor exit node to a storage account
- Alert name (new): Authenticated access from a Tor exit node
- Alert types: Storage.Blob_TorAnomaly / Storage.Files_TorAnomaly
- Description: One or more storage container(s) / file share(s) in your storage account were successfully accessed from an IP address known to be an active exit node of Tor (an anonymizing proxy). Threat actors use Tor to make it difficult to trace the activity back to them. Authenticated access from a Tor exit node is a likely indication that a threat actor is trying to hide their identity. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2
- MITRE tactic: Initial access
- Severity: High/Medium

Unusual unauthenticated access

A change in access patterns may indicate that a threat actor was able to exploit public read access to storage containers, either by exploiting a mistake in access configurations, or by changing the access permissions.

This medium severity alert is now tuned with improved behavioral logic, higher accuracy, and confidence that the activity is malicious. This enhancement reduces the benign positive rate.

The alert name and description have been updated. The AlertType remains unchanged.

- Alert name (old): Anonymous access to a storage account
- Alert name (new): Unusual unauthenticated access to a storage container
- Alert types: Storage.Blob_AnonymousAccessAnomaly
- Description: This storage account was accessed without authentication, which is a change in the common access pattern. Read access to this container is usually authenticated. This might indicate that a threat actor was able to exploit public read access to storage container(s) in this storage account(s). Applies to: Azure Blob Storage
- MITRE tactic: Collection
- Severity: Medium

For more information, see:

- Threat matrix for storage services
- Introduction to Microsoft Defender for Storage
- List of alerts provided by Microsoft Defender for Storage

'PortSweeping' alert removed from network layer alerts

The following alert was removed from our network layer alerts due to inefficiencies:

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS	SEVERITY
Possible outgoing port scanning activity detected (PortSweeping)	Network traffic analysis detected suspicious outgoing traffic from % {Compromised Host}. This traffic may be a result of a port scanning activity. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). If this behavior is intentional, please note that performing port scanning is against Azure Terms of service. If this behavior is unintentional, it may mean your resource has been compromised.	Discovery	Medium

November 2021

Our Ignite release includes:

- Azure Security Center and Azure Defender become Microsoft Defender for Cloud
- Native CSPM for AWS and threat protection for Amazon EKS, and AWS EC2
- Prioritize security actions by data sensitivity (powered by Azure Purview) (in preview)
- Expanded security control assessments with Azure Security Benchmark v3
- Microsoft Sentinel connector's optional bi-directional alert synchronization released for general availability (GA)
- New recommendation to push Azure Kubernetes Service (AKS) logs to Sentinel
- Recommendations mapped to the MITRE ATT&CK® framework - released for general availability (GA)

Other changes in November include:

- Microsoft Threat and Vulnerability Management added as vulnerability assessment solution - released for general availability (GA)
- Microsoft Defender for Endpoint for Linux now supported by Microsoft Defender for servers - released for general availability (GA)
- Snapshot export for recommendations and security findings (in preview)
- Auto provisioning of vulnerability assessment solutions released for general availability (GA)
- Software inventory filters in asset inventory released for general availability (GA)
- New AKS security policy added to default initiative – for use by private preview customers only
- Inventory display of on-premises machines applies different template for resource name

Azure Security Center and Azure Defender become Microsoft Defender for Cloud

According to the 2021 State of the Cloud report, 92% of organizations now have a multi-cloud strategy. At

Microsoft, our goal is to centralize security across these environments and help security teams work more effectively.

Microsoft Defender for Cloud (formerly known as Azure Security Center and Azure Defender) is a Cloud Security Posture Management (CSPM) and cloud workload protection (CWP) solution that discovers weaknesses across your cloud configuration, helps strengthen the overall security posture of your environment, and protects workloads across multi-cloud and hybrid environments.

At Ignite 2019, we shared our vision to create the most complete approach for securing your digital estate and integrating XDR technologies under the Microsoft Defender brand. Unifying Azure Security Center and Azure Defender under the new name **Microsoft Defender for Cloud**, reflects the integrated capabilities of our security offering and our ability to support any cloud platform.

Native CSPM for AWS and threat protection for Amazon EKS, and AWS EC2

A new **environment settings** page provides greater visibility and control over your management groups, subscriptions, and AWS accounts. The page is designed to onboard AWS accounts at scale: connect your AWS **management account**, and you'll automatically onboard existing and future accounts.

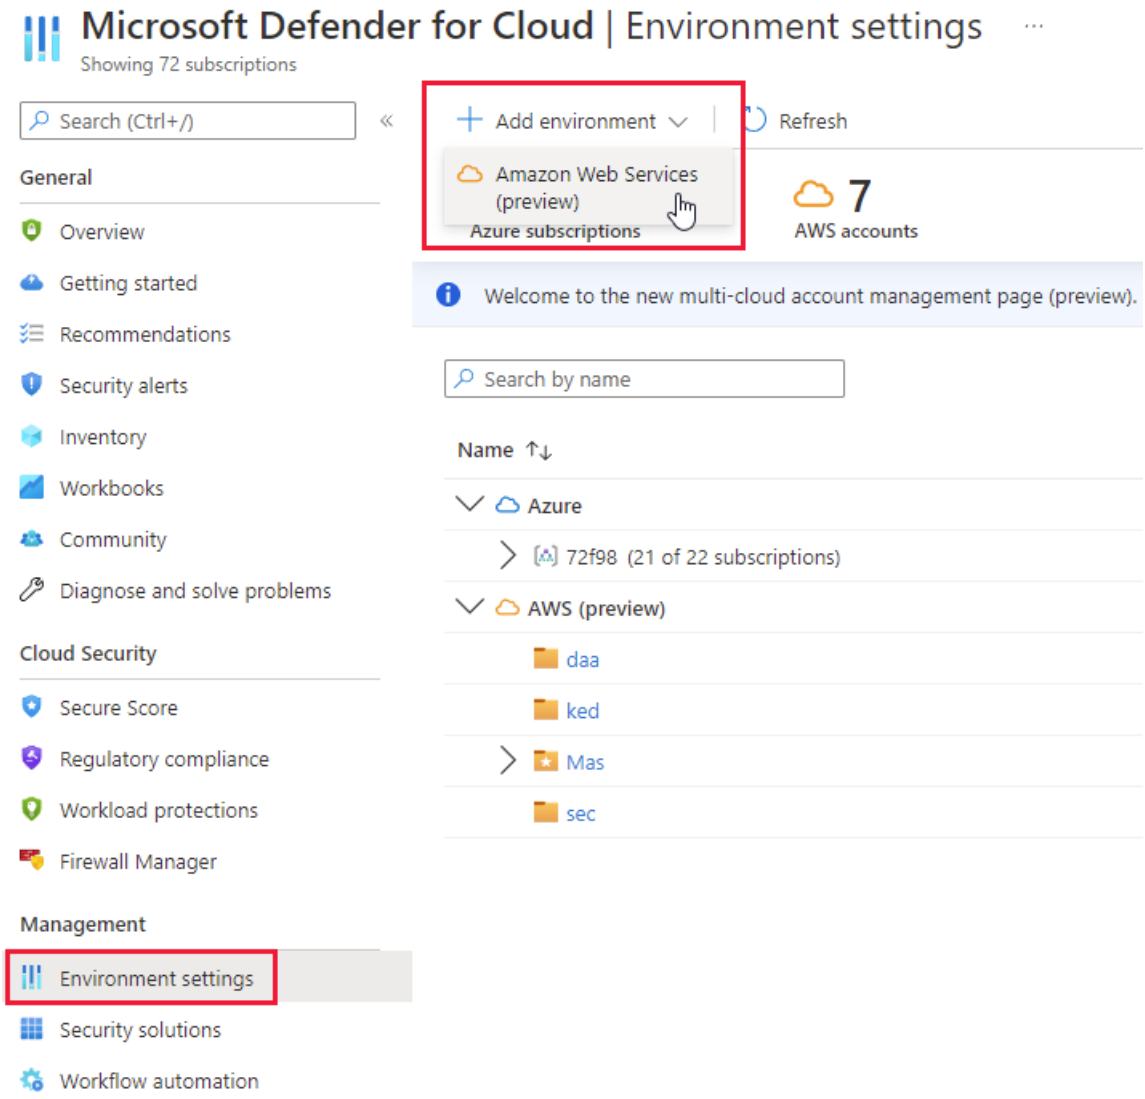

When you've added your AWS accounts, Defender for Cloud protects your AWS resources with any or all of the following plans:

- **Defender for Cloud's CSPM features** extend to your AWS resources. This agentless plan assesses your AWS resources according to AWS-specific security recommendations and these are included in your secure score. The resources will also be assessed for compliance with built-in standards specific to AWS (AWS CIS,

AWS PCI DSS, and AWS Foundational Security Best Practices). Defender for Cloud's asset inventory page is a multi-cloud enabled feature helping you manage your AWS resources alongside your Azure resources.

- **Microsoft Defender for Kubernetes** extends its container threat detection and advanced defenses to your **Amazon EKS Linux clusters**.
- **Microsoft Defender for servers** brings threat detection and advanced defenses to your Windows and Linux EC2 instances. This plan includes the integrated license for Microsoft Defender for Endpoint, security baselines and OS level assessments, vulnerability assessment scanning, adaptive application controls (AAC), file integrity monitoring (FIM), and more.

Learn more about connecting your AWS accounts to Microsoft Defender for Cloud.

Prioritize security actions by data sensitivity (powered by Azure Purview) (in preview)

Data resources remain a popular target for threat actors. So it's crucial for security teams to identify, prioritize, and secure sensitive data resources across their cloud environments.

To address this challenge, Microsoft Defender for Cloud now integrates sensitivity information from Azure Purview. Azure Purview is a unified data governance service that provides rich insights into the sensitivity of your data within multi-cloud, and on-premises workloads.

The integration with Azure Purview extends your security visibility in Defender for Cloud from the infrastructure level down to the data, enabling an entirely new way to prioritize resources and security activities for your security teams.

Learn more in Prioritize security actions by data sensitivity.

Expanded security control assessments with Azure Security Benchmark v3

Microsoft Defender for Cloud's security recommendations are enabled and supported by the Azure Security Benchmark.

Azure Security Benchmark is the Microsoft-authored, Azure-specific set of guidelines for security and compliance best practices based on common compliance frameworks. This widely respected benchmark builds on the controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security.

From Ignite 2021, Azure Security Benchmark **v3** is available in Defender for Cloud's regulatory compliance dashboard and enabled as the new default initiative for all Azure subscriptions protected with Microsoft Defender for Cloud.

Enhancements for v3 include:

- Additional mappings to industry frameworks PCI-DSS v3.2.1 and CIS Controls v8.
- More granular and actionable guidance for controls with the introduction of:
  - **Security Principles** - Providing insight into the overall security objectives that build the foundation for our recommendations.
  - **Azure Guidance** - The technical "how-to" for meeting these objectives.
- New controls include DevOps security for issues such as threat modeling and software supply chain security, as well as key and certificate management for best practices in Azure.

Learn more in Introduction to Azure Security Benchmark.

Microsoft Sentinel connector's optional bi-directional alert synchronization released for general availability (GA)

In July, we announced a preview feature, **bi-directional alert synchronization**, for the built-in connector in Microsoft Sentinel (Microsoft's cloud-native SIEM and SOAR solution). This feature is now released for general availability (GA).

When you connect Microsoft Defender for Cloud to Microsoft Sentinel, the status of security alerts is synchronized between the two services. So, for example, when an alert is closed in Defender for Cloud, that alert will display as closed in Microsoft Sentinel as well. Changing the status of an alert in Defender for Cloud won't affect the status of any Microsoft Sentinel **incidents** that contain the synchronized Microsoft Sentinel alert, only that of the synchronized alert itself.

When you enable **bi-directional alert synchronization** you'll automatically sync the status of the original Defender for Cloud alerts with Microsoft Sentinel incidents that contain the copies of those Defender for Cloud alerts. So, for example, when a Microsoft Sentinel incident containing a Defender for Cloud alert is closed, Defender for Cloud will automatically close the corresponding original alert.

Learn more in Connect Azure Defender alerts from Azure Security Center and Stream alerts to Azure Sentinel.

New recommendation to push Azure Kubernetes Service (AKS) logs to Sentinel

In a further enhancement to the combined value of Defender for Cloud and Microsoft Sentinel, we'll now highlight Azure Kubernetes Service instances that aren't sending log data to Microsoft Sentinel.

SecOps teams can choose the relevant Microsoft Sentinel workspace directly from the recommendation details page and immediately enable the streaming of raw logs. This seamless connection between the two products makes it easy for security teams to ensure complete logging coverage across their workloads to stay on top of their entire environment.

The new recommendation, "Diagnostic logs in Kubernetes services should be enabled" includes the 'Fix' option for faster remediation.

We've also enhanced the "Auditing on SQL server should be enabled" recommendation with the same Sentinel streaming capabilities.

Recommendations mapped to the MITRE ATT&CK® framework - released for general availability (GA)

We've enhanced Defender for Cloud's security recommendations to show their position on the MITRE ATT&CK® framework. This globally accessible knowledge base of threat actors' tactics and techniques based on real-world observations, provides more context to help you understand the associated risks of the recommendations for your environment.

You'll find these tactics wherever you access recommendation information:

- **Azure Resource Graph query results** for relevant recommendations include the MITRE ATT&CK® tactics and techniques.
- **Recommendation details pages** show the mapping for all relevant recommendations:

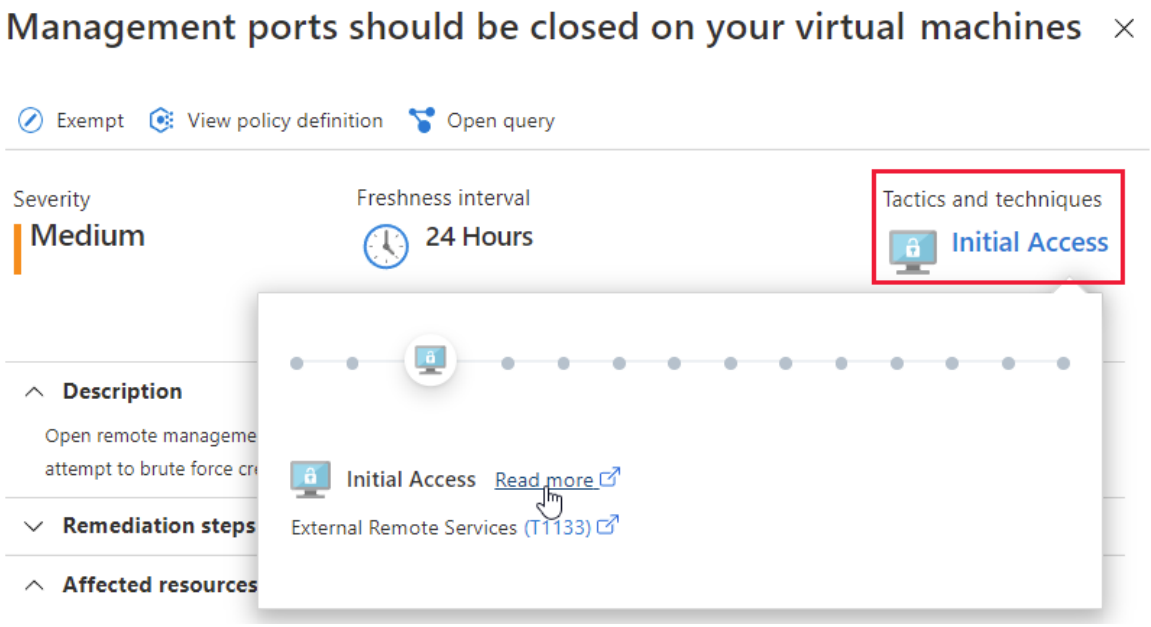

- **The recommendations page in Defender for Cloud** has a new 

filter to select recommendations according to their associated tactic:

Learn more in Review your security recommendations.

Microsoft Threat and Vulnerability Management added as vulnerability assessment solution - released for general availability (GA)

In October, we announced an extension to the integration between Microsoft Defender for servers and Microsoft Defender for Endpoint, to support a new vulnerability assessment provider for your machines: Microsoft threat and vulnerability management. This feature is now released for general availability (GA).

Use **threat and vulnerability management** to discover vulnerabilities and misconfigurations in near real time with the integration with Microsoft Defender for Endpoint enabled, and without the need for additional agents or periodic scans. Threat and vulnerability management prioritizes vulnerabilities based on the threat landscape and detections in your organization.

Use the security recommendation "A vulnerability assessment solution should be enabled on your virtual machines" to surface the vulnerabilities detected by threat and vulnerability management for your supported machines.

To automatically surface the vulnerabilities, on existing and new machines, without the need to manually remediate the recommendation, see Vulnerability assessment solutions can now be auto enabled (in preview).

Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management.

Microsoft Defender for Endpoint for Linux now supported by Microsoft Defender for servers - released for general availability (GA)

In August, we announced preview support for deploying the Defender for Endpoint for Linux sensor to supported Linux machines. This feature is now released for general availability (GA).

Microsoft Defender for servers includes an integrated license for Microsoft Defender for Endpoint. Together, they provide comprehensive endpoint detection and response (EDR) capabilities.

When Defender for Endpoint detects a threat, it triggers an alert. The alert is shown in Defender for Cloud. From Defender for Cloud, you can also pivot to the Defender for Endpoint console, and perform a detailed investigation to uncover the scope of the attack.

Learn more in Protect your endpoints with Security Center's integrated EDR solution: Microsoft Defender for

Endpoint.

Snapshot export for recommendations and security findings (in preview)

Defender for Cloud generates detailed security alerts and recommendations. You can view them in the portal or through programmatic tools. You might also need to export some or all of this information for tracking with other monitoring tools in your environment.

Defender for Cloud's **continuous export** feature lets you fully customize *what* will be exported, and *where* it will go. Learn more in Continuously export Microsoft Defender for Cloud data.

Even though the feature is called *continuous*, there's also an option to export weekly snapshots. Until now, these weekly snapshots were limited to secure score and regulatory compliance data. We've added the capability to export recommendations and security findings.

Auto provisioning of vulnerability assessment solutions released for general availability (GA)

In October, we announced the addition of vulnerability assessment solutions to Defender for Cloud's auto provisioning page. This is relevant to Azure virtual machines and Azure Arc machines on subscriptions protected by Azure Defender for servers. This feature is now released for general availability (GA).

If the integration with Microsoft Defender for Endpoint is enabled, Defender for Cloud presents a choice of vulnerability assessment solutions:

- (**NEW**) The Microsoft threat and vulnerability management module of Microsoft Defender for Endpoint (see the release note)
- The integrated Qualys agent

Your chosen solution will be automatically enabled on supported machines.

Learn more in Automatically configure vulnerability assessment for your machines.

Software inventory filters in asset inventory released for general availability (GA)

In October, we announced new filters for the asset inventory page to select machines running specific software - and even specify the versions of interest. This feature is now released for general availability (GA).

You can query the software inventory data in **Azure Resource Graph Explorer**.

To use these features, you'll need to enable the integration with Microsoft Defender for Endpoint.

For full details, including sample Kusto queries for Azure Resource Graph, see Access a software inventory.

New AKS security policy added to default initiative – for use by private preview customers only

To ensure that Kubernetes workloads are secure by default, Defender for Cloud includes Kubernetes level policies and hardening recommendations, including enforcement options with Kubernetes admission control.

As part of this project, we've added a policy and recommendation (disabled by default) for gating deployment on Kubernetes clusters. The policy is in the default initiative but is only relevant for organizations who register for the related private preview.

You can safely ignore the policies and recommendation ("Kubernetes clusters should gate deployment of vulnerable images") and there will be no impact on your environment.

If you'd like to participate in the private preview, you'll need to be a member of the private preview ring. If you're not already a member, submit a request here. Members will be notified when the preview begins.

Inventory display of on-premises machines applies different template for resource name

To improve the presentation of resources in the Asset inventory, we've removed the "source-computer-IP" element from the template for naming on-premises machines.

- **Previous format:** `machine-name_source-computer-id_VMUUID`

- From this update: `machine-name_VMUUID`

October 2021

Updates in October include:

- Microsoft Threat and Vulnerability Management added as vulnerability assessment solution (in preview)
- Vulnerability assessment solutions can now be auto enabled (in preview)
- Software inventory filters added to asset inventory (in preview)
- Changed prefix of some alert types from "ARM_" to "VM_"
- Changes to the logic of a security recommendation for Kubernetes clusters
- Recommendations details pages now show related recommendations
- New alerts for Azure Defender for Kubernetes (in preview)

Microsoft Threat and Vulnerability Management added as vulnerability assessment solution (in preview)

We've extended the integration between Azure Defender for servers and Microsoft Defender for Endpoint, to support a new vulnerability assessment provider for your machines: Microsoft threat and vulnerability management.

Use **threat and vulnerability management** to discover vulnerabilities and misconfigurations in near real time with the integration with Microsoft Defender for Endpoint enabled, and without the need for additional agents or periodic scans. Threat and vulnerability management prioritizes vulnerabilities based on the threat landscape and detections in your organization.

Use the security recommendation "A vulnerability assessment solution should be enabled on your virtual machines" to surface the vulnerabilities detected by threat and vulnerability management for your supported machines.

To automatically surface the vulnerabilities, on existing and new machines, without the need to manually remediate the recommendation, see Vulnerability assessment solutions can now be auto enabled (in preview).

Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management.

Vulnerability assessment solutions can now be auto enabled (in preview)

Security Center's auto provisioning page now includes the option to automatically enable a vulnerability assessment solution to Azure virtual machines and Azure Arc machines on subscriptions protected by Azure Defender for servers.

If the integration with Microsoft Defender for Endpoint is enabled, Defender for Cloud presents a choice of vulnerability assessment solutions:

- (**NEW**) The Microsoft threat and vulnerability management module of Microsoft Defender for Endpoint (see the release note)
- The integrated Qualys agent

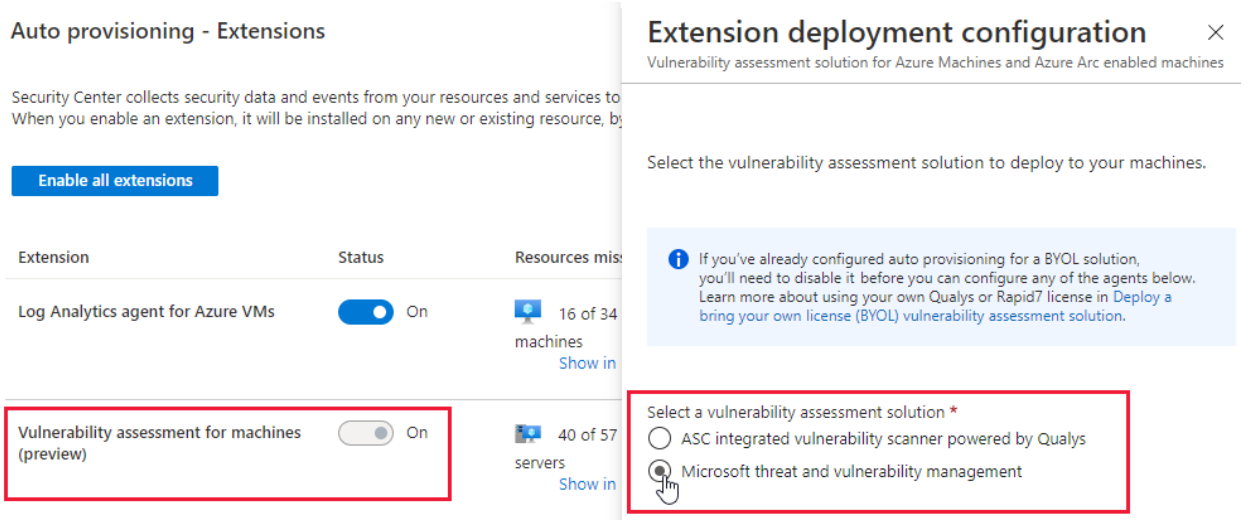

Your chosen solution will be automatically enabled on supported machines.

Learn more in Automatically configure vulnerability assessment for your machines.

Software inventory filters added to asset inventory (in preview)

The asset inventory page now includes a filter to select machines running specific software - and even specify the versions of interest.

Additionally, you can query the software inventory data in **Azure Resource Graph Explorer**.

To use these new features, you'll need to enable the integration with Microsoft Defender for Endpoint.

For full details, including sample Kusto queries for Azure Resource Graph, see Access a software inventory.

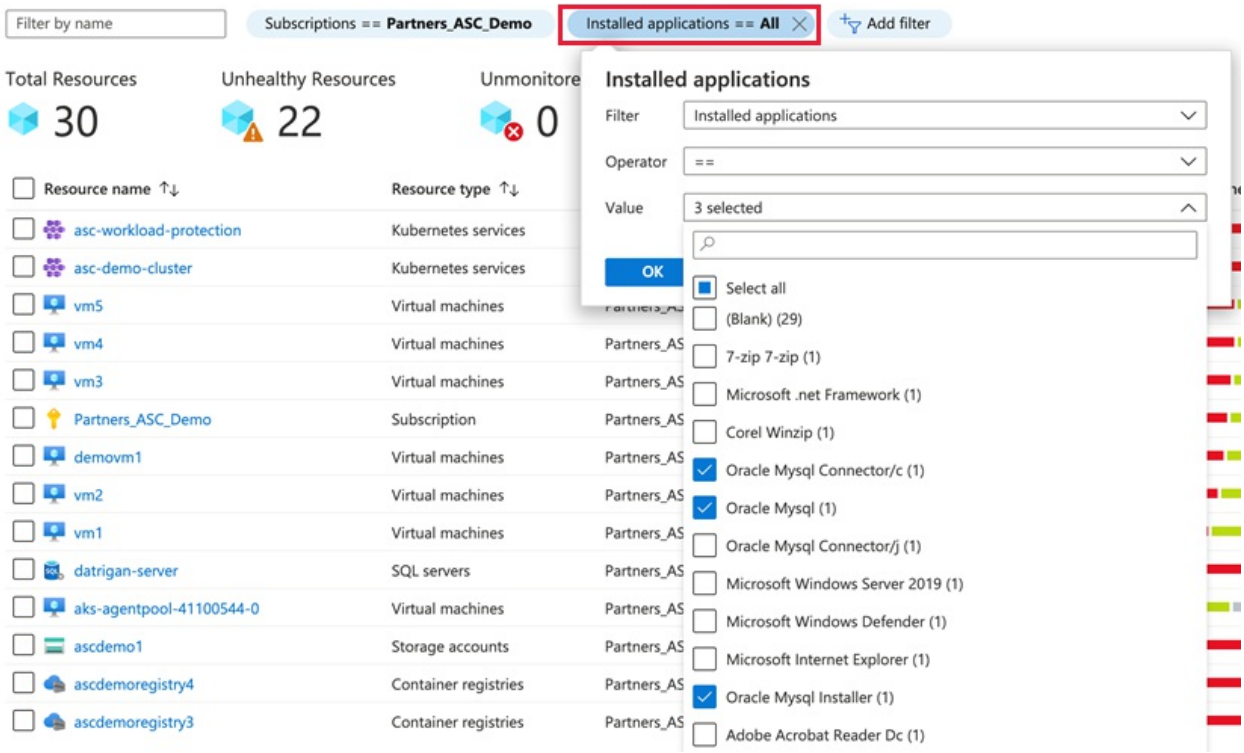

Changed prefix of some alert types from "ARM_" to "VM_"

In July 2021, we announced a logical reorganization of Azure Defender for Resource Manager alerts

As part of a logical reorganization of some of the Azure Defender plans, we moved twenty-one alerts from Azure Defender for Resource Manager to Azure Defender for servers.

With this update, we've changed the prefixes of these alerts to match this reassignment and replaced "ARM_" with "VM_" as shown in the following table:

ORIGINAL NAME	FROM THIS CHANGE
ARM_AmBroadFilesExclusion	VM_AmBroadFilesExclusion
ARM_AmDisablementAndCodeExecution	VM_AmDisablementAndCodeExecution
ARM_AmDisablement	VM_AmDisablement
ARM_AmFileExclusionAndCodeExecution	VM_AmFileExclusionAndCodeExecution
ARM_AmTempFileExclusionAndCodeExecution	VM_AmTempFileExclusionAndCodeExecution
ARM_AmTempFileExclusion	VM_AmTempFileExclusion
ARM_AmRealtimeProtectionDisabled	VM_AmRealtimeProtectionDisabled
ARM_AmTempRealtimeProtectionDisablement	VM_AmTempRealtimeProtectionDisablement
ARM_AmRealtimeProtectionDisablementAndCodeExec	VM_AmRealtimeProtectionDisablementAndCodeExec
ARM_AmMalwareCampaignRelatedExclusion	VM_AmMalwareCampaignRelatedExclusion
ARM_AmTemporarilyDisablement	VM_AmTemporarilyDisablement
ARM_UnusualAmFileExclusion	VM_UnusualAmFileExclusion
ARM_CustomScriptExtensionSuspiciousCmd	VM_CustomScriptExtensionSuspiciousCmd
ARM_CustomScriptExtensionSuspiciousEntryPoint	VM_CustomScriptExtensionSuspiciousEntryPoint
ARM_CustomScriptExtensionSuspiciousPayload	VM_CustomScriptExtensionSuspiciousPayload
ARM_CustomScriptExtensionSuspiciousFailure	VM_CustomScriptExtensionSuspiciousFailure
ARM_CustomScriptExtensionUnusualDeletion	VM_CustomScriptExtensionUnusualDeletion
ARM_CustomScriptExtensionUnusualExecution	VM_CustomScriptExtensionUnusualExecution
ARM_VMAccessUnusualConfigReset	VM_VMAccessUnusualConfigReset
ARM_VMAccessUnusualPasswordReset	VM_VMAccessUnusualPasswordReset
ARM_VMAccessUnusualSSHReset	VM_VMAccessUnusualSSHReset

Learn more about the Azure Defender for Resource Manager and Azure Defender for servers plans.

Changes to the logic of a security recommendation for Kubernetes clusters

The recommendation "Kubernetes clusters should not use the default namespace" prevents usage of the default namespace for a range of resource types. Two of the resource types that were included in this recommendation have been removed: ConfigMap and Secret.

Learn more about this recommendation and hardening your Kubernetes clusters in Understand Azure Policy for

Kubernetes clusters.

Recommendations details pages now show related recommendations

To clarify the relationships between different recommendations, we've added a **Related recommendations** area to the details pages of many recommendations.

The three relationship types that are shown on these pages are:

- **Prerequisite** - A recommendation that must be completed before the selected recommendation
- **Alternative** - A different recommendation which provides another way of achieving the goals of the selected recommendation
- **Dependent** - A recommendation for which the selected recommendation is a prerequisite

For each related recommendation, the number of unhealthy resources is shown in the "Affected resources" column.

> **TIP**
> If a related recommendation is grayed out, its dependency isn't yet completed and so isn't available.

An example of related recommendations:

1. Security Center checks your machines for supported vulnerability assessment solutions:
   A vulnerability assessment solution should be enabled on your virtual machines
2. If one is found, you'll get notified about discovered vulnerabilities:
   Vulnerabilities in your virtual machines should be remediated

Obviously, Security Center can't notify you about discovered vulnerabilities unless it finds a supported vulnerability assessment solution.

Therefore:

- Recommendation #1 is a prerequisite for recommendation #2
- Recommendation #2 depends upon recommendation #1

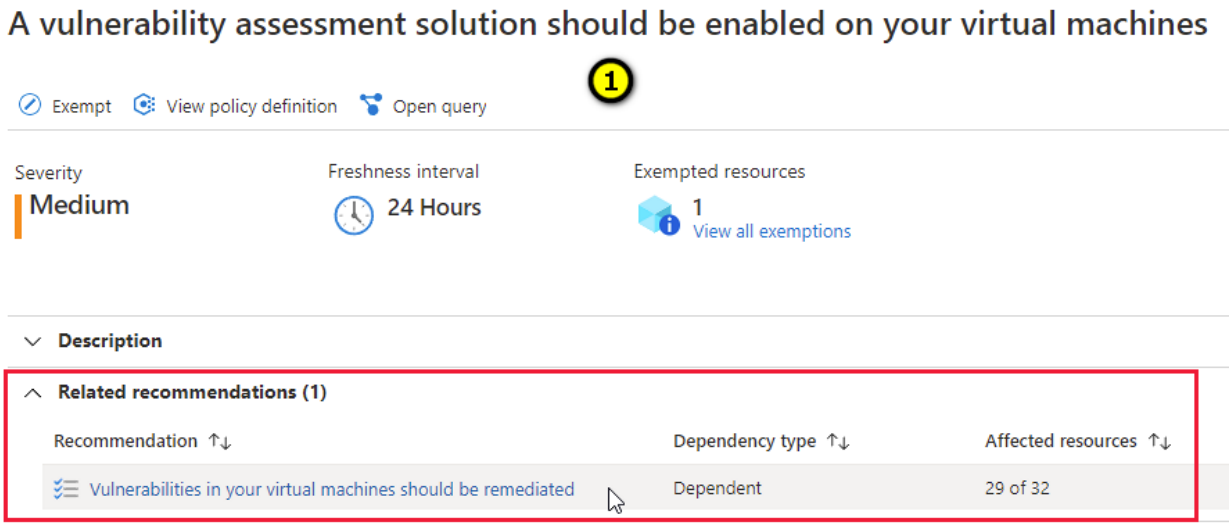

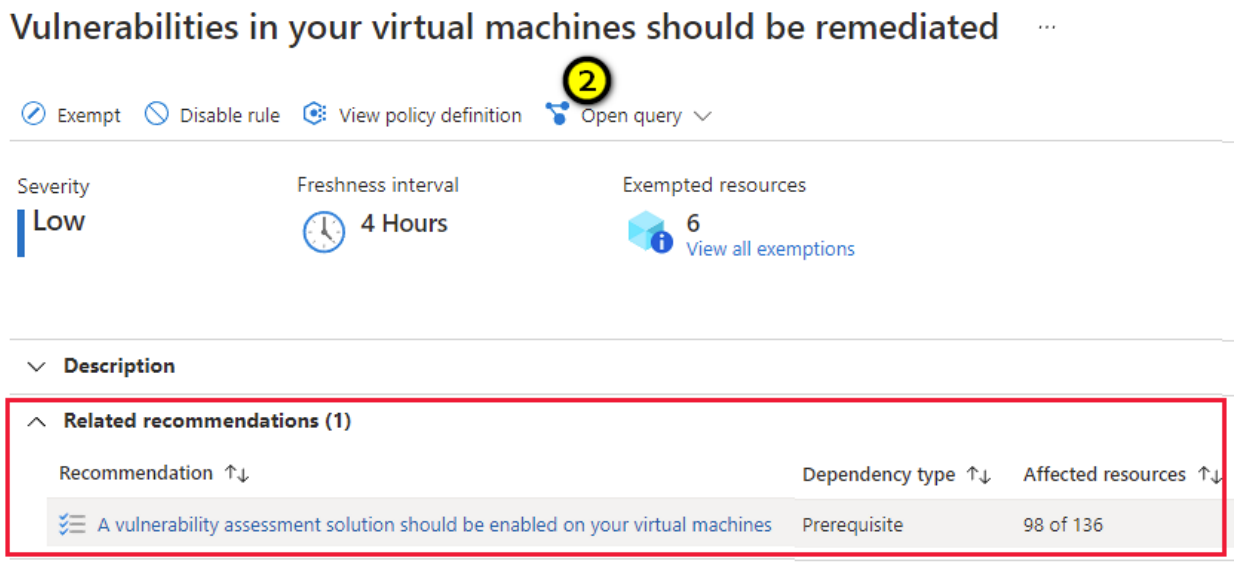

New alerts for Azure Defender for Kubernetes (in preview)

To expand the threat protections provided by Azure Defender for Kubernetes, we've added two preview alerts.

These alerts are generated based on a new machine learning model and Kubernetes advanced analytics, measuring multiple deployment and role assignment attributes against previous activities in the cluster and across all clusters monitored by Azure Defender.

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTIC	SEVERITY
Anomalous pod deployment (Preview) (K8S_AnomalousPodDeployment)	Kubernetes audit log analysis detected pod deployment that is anomalous based on previous pod deployment activity. This activity is considered an anomaly when taking into account how the different features seen in the deployment operation are in relations to one another. The features monitored by this analytics include the container image registry used, the account performing the deployment, day of the week, how often does this account performs pod deployments, user agent used in the operation, is this a namespace which is pod deployment occur to often, or other feature. Top contributing reasons for raising this alert as anomalous activity are detailed under the alert extended properties.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTIC	SEVERITY
Excessive role permissions assigned in Kubernetes cluster (Preview) (K8S_ServiceAcountPermissionAnomaly)	Analysis of the Kubernetes audit logs detected an excessive permissions role assignment to your cluster. From examining role assignments, the listed permissions are uncommon to the specific service account. This detection considers previous role assignments to the same service account across clusters monitored by Azure, volume per permission, and the impact of the specific permission. The anomaly detection model used for this alert takes into account how this permission is used across all clusters monitored by Azure Defender.	Privilege Escalation	Low

For a full list of the Kubernetes alerts, see Alerts for Kubernetes clusters.

September 2021

In September, the following update was released:

Two new recommendations to audit OS configurations for Azure security baseline compliance (in preview)

The following two recommendations have been released to assess your machines' compliance with the Windows security baseline and the Linux security baseline:

- For Windows machines, Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration)
- For Linux machines, Vulnerabilities in security configuration on your Linux machines should be remediated (powered by Guest Configuration)

These recommendations make use of the guest configuration feature of Azure Policy to compare the OS configuration of a machine with the baseline defined in the Azure Security Benchmark.

Learn more about using these recommendations in Harden a machine's OS configuration using guest configuration.

Important upcoming changes to Microsoft Defender for Cloud

2/15/2022 • 10 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

IMPORTANT

The information on this page relates to pre-release products or features, which may be substantially modified before they are commercially released, if ever. Microsoft makes no commitments or warranties, express or implied, with respect to the information provided here.

On this page, you'll learn about changes that are planned for Defender for Cloud. It describes planned modifications to the product that might impact things like your secure score or workflows.

If you're looking for the latest release notes, you'll find them in the What's new in Microsoft Defender for Cloud.

Planned changes

PLANNED CHANGE	ESTIMATED DATE FOR CHANGE
Deprecating a preview alert: ARM.MCAS_ActivityFromAnonymousIPAddresses	January 2022
Legacy implementation of ISO 27001 is being replaced with new ISO 27001:2013	January 2022
Multiple changes to identity recommendations	February 2022
Deprecating the recommendation to use service principals to protect your subscriptions	February 2022
Moving recommendation Vulnerabilities in container security configurations should be remediated from the secure score to best practices	February 2022
Deprecating the recommendations to install the network traffic data collection agent	February 2022
Changes to recommendations for managing endpoint protection solutions	March 2022
AWS recommendations to GA	March 2022

PLANNED CHANGE	ESTIMATED DATE FOR CHANGE
Relocation of custom recommendations	March 2022
Deprecating Microsoft Defender for IoT device recommendations	March 2022
Deprecating Microsoft Defender for IoT device alerts	March 2022

Deprecating a preview alert: ARM.MCAS_ActivityFromAnonymousIPAddresses

Estimated date for change: January 2022

We'll be deprecating the following preview alert:

ALERT NAME	DESCRIPTION
PREVIEW - Activity from a risky IP address (ARM.MCAS_ActivityFromAnonymousIPAddresses)	Users activity from an IP address that has been identified as an anonymous proxy IP address has been detected. These proxies are used by people who want to hide their device's IP address, and can be used for malicious intent. This detection uses a machine learning algorithm that reduces false positives, such as mis-tagged IP addresses that are widely used by users in the organization. Requires an active Microsoft Defender for Cloud Apps license.

We've created new alerts that provide this information and add to it. In addition, the newer alerts (ARM_OperationFromSuspiciousIP, ARM_OperationFromSuspiciousProxyIP) don't require a license for Microsoft Defender for Cloud Apps (formerly known as Microsoft Cloud App Security).

Legacy implementation of ISO 27001 is being replaced with new ISO 27001:2013

Estimated date for change: January 2022

The legacy implementation of ISO 27001 will be removed from Defender for Cloud's regulatory compliance dashboard. If you're tracking your ISO 27001 compliance with Defender for Cloud, onboard the new ISO 27001:2013 standard for all relevant management groups or subscriptions, and the current legacy ISO 27001 will soon be removed from the dashboard.

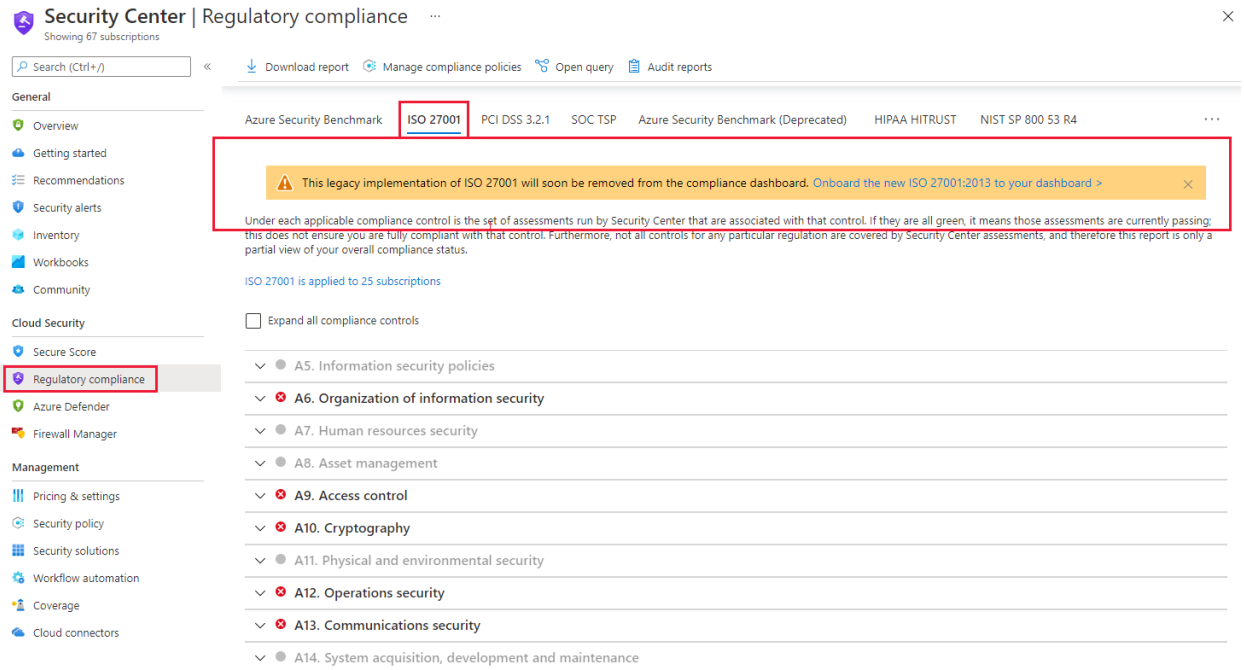

Multiple changes to identity recommendations

Estimated date for change: February 2022

Defender for Cloud includes multiple recommendations for improving the management of users and accounts. In December, we'll be making the changes outlined below.

- **Improved freshness interval** - Currently, the identity recommendations have a freshness interval of 24 hours. This update will reduce that interval to 12 hours.
- **Account exemption capability** - Defender for Cloud has many features for customizing the experience and making sure your secure score reflects your organization's security priorities. The exempt option on security recommendations is one such feature. For a full overview and instructions, see Exempting resources and recommendations from your secure score. With this update, you'll be able to exempt specific accounts from evaluation by the eight recommendations listed in the following table.

  Typically, you'd exempt emergency "break glass" accounts from MFA recommendations, because such accounts are often deliberately excluded from an organization's MFA requirements. Alternatively, you might have external accounts that you'd like to permit access to but which don't have MFA enabled.

  > **TIP**
  > When you exempt an account, it won't be shown as unhealthy and also won't cause a subscription to appear unhealthy.

RECOMMENDATION	ASSESSMENT KEY
MFA should be enabled on accounts with owner permissions on your subscription	94290b00-4d0c-d7b4-7cea-064a9554e681
MFA should be enabled on accounts with read permissions on your subscription	151e82c5-5341-a74b-1eb0-bc38d2c84bb5
MFA should be enabled on accounts with write permissions on your subscription	57e98606-6b1e-6193-0e3d-fe621387c16b

RECOMMENDATION	ASSESSMENT KEY
External accounts with owner permissions should be removed from your subscription	c3b6ae71-f1f0-31b4-e6c1-d5951285d03d
External accounts with read permissions should be removed from your subscription	a8c6a4ad-d51e-88fe-2979-d3ee3c864f8b
External accounts with write permissions should be removed from your subscription	04e7147b-0deb-9796-2e5c-0336343ceb3d
Deprecated accounts with owner permissions should be removed from your subscription	e52064aa-6853-e252-a11e-dffc675689c2
Deprecated accounts should be removed from your subscription	00c6d40b-e990-6acf-d4f3-471e747a27c4

- **Recommendations rename** - From this update, we're renaming two recommendations. We're also revising their descriptions. The assessment keys will remain unchanged.

PROPERTY	CURRENT VALUE	FROM THE UPDATE
Assessment key	e52064aa-6853-e252-a11e-dffc675689c2	Unchanged
Name	Deprecated accounts with owner permissions should be removed from your subscription	Subscriptions should be purged of accounts that are blocked in Active Directory and have owner permissions
Description	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed.	User accounts that have been blocked from signing into Active Directory, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. Learn more about securing the identity perimeter in Azure Identity Management and access control security best practices.
Related policy	Deprecated accounts with owner permissions should be removed from your subscription	Subscriptions should be purged of accounts that are blocked in Active Directory and have owner permissions

PROPERTY	CURRENT VALUE	FROM THE UPDATE
Assessment key	00c6d40b-e990-6acf-d4f3-471e747a27c4	Unchanged

PROPERTY	CURRENT VALUE	FROM THE UPDATE
Name	Deprecated accounts should be removed from your subscription	Subscriptions should be purged of accounts that are blocked in Active Directory and have read and write permissions
Description	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed.	User accounts that have been blocked from signing into Active Directory, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. Learn more about securing the identity perimeter in Azure Identity Management and access control security best practices.
Related policy	Deprecated accounts should be removed from your subscription	Subscriptions should be purged of accounts that are blocked in Active Directory and have read and write permissions

Deprecating the recommendation to use service principals to protect your subscriptions

Estimated date for change: February 2022

As organizations are moving away from using management certificates to manage their subscriptions, and our recent announcement that we're retiring the Cloud Services (classic) deployment model, we'll be deprecating the following Defender for Cloud recommendation and its related policy:

RECOMMENDATION	DESCRIPTION	SEVERITY
Service principals should be used to protect your subscriptions instead of Management Certificates	Management certificates allow anyone who authenticates with them to manage the subscription(s) they are associated with. To manage subscriptions more securely, using service principals with Resource Manager is recommended to limit the blast radius in the case of a certificate compromise. It also automates resource management. (Related policy: Service principals should be used to protect your subscriptions instead of management certificates)	Medium

Learn more:

- Cloud Services (classic) deployment model is retiring on 31 August 2024
- Overview of Azure Cloud Services (classic)
- Workflow of Windows Azure classic VM Architecture - including RDFE workflow basics

Deprecating the recommendations to install the network traffic data collection agent

Estimated date for change: February 2022

Changes in our roadmap and priorities have removed the need for the network traffic data collection agent. Consequently, we'll be deprecating the following two recommendations and their related policies.

RECOMMENDATION	DESCRIPTION	SEVERITY
Network traffic data collection agent should be installed on Linux virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations and specific network threats. (Related policy: Network traffic data collection agent should be installed on Linux virtual machines)	Medium
Network traffic data collection agent should be installed on Windows virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations, and specific network threats. (Related policy: Network traffic data collection agent should be installed on Windows virtual machines)	Medium

Moving recommendation Vulnerabilities in container security configurations should be remediated from the secure score to best practices

Estimated date for change: February 2022

The recommendation for 'Vulnerabilities in container security configurations should be remediated' is being moved from the secure score section to best practices section.

The current user experience only provides the score when all compliance checks have passed. Most customers have difficulties with meeting all the required checks. We are working on an improved experience for this recommendation, and once released the recommendation will be moved back to the secure score.

Changes to recommendations for managing endpoint protection solutions

Estimated date for change: March 2022

In August 2021, we added two new **preview** recommendations to deploy and maintain the endpoint protection solutions on your machines. For full details, see the release note.

When the recommendations are released to general availability, they will replace the following existing recommendations:

- **Endpoint protection should be installed on your machines** will replace:
  - Install endpoint protection solution on virtual machines (key: 83f577bd-a1b6-b7e1-0891-12ca19d1e6df)
  - Install endpoint protection solution on your machines (key: 383cf3bc-fdf9-4a02-120a-3e7e36c6bfee)
- **Endpoint protection health issues should be resolved on your machines** will replace the existing recommendation that has the same name. The two recommendations have different assessment keys:
  - Assessment key for the **preview** recommendation: 37a3689a-818e-4a0e-82ac-b1392b9bb000

- Assessment key for the **GA** recommendation: 3bcd234d-c9c7-c2a2-89e0-c01f419c1a8a

Learn more:

- Defender for Cloud's supported endpoint protection solutions
- How these recommendations assess the status of your deployed solutions

AWS recommendations to GA

Estimated date for change: March 2022

There are currently AWS recommendations in the preview stage. These recommendations come from the AWS Foundational Security Best Practices standard which is assigned by default. All of the recommendations will become Generally Available (GA) in March 2022.

When these recommendations go live, their impact will be included in the calculations of your secure score. Expect changes to your secure score.

To find these recommendations:

1. Navigate to **Environment settings** > `AWS connector` > **Standards (preview)**.
2. Right click on **AWS Foundational Security Best Practices (preview)**, and select **view assessments**.

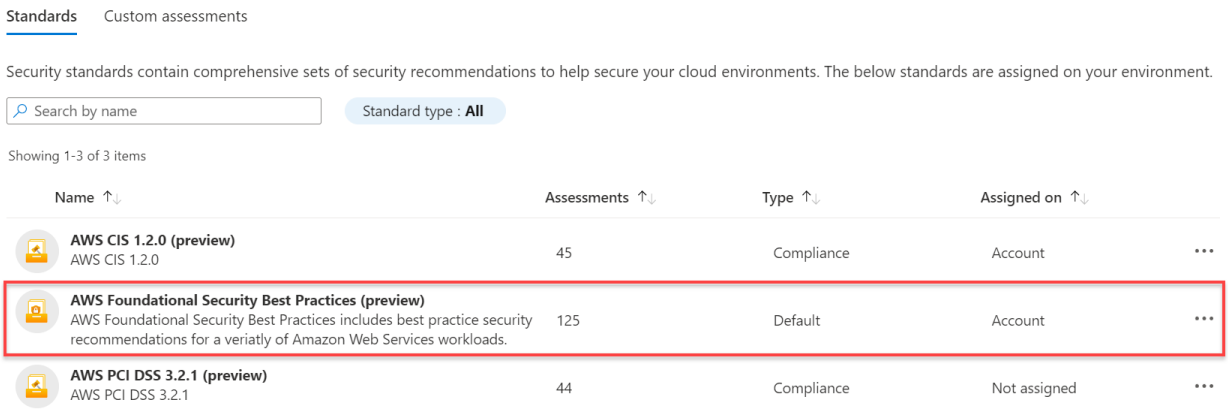

Relocation of custom recommendations

Estimated date for change: March 2022

Custom recommendation are those created by a user, and have no impact on the secure score. Therefore, the custom recommendations are being relocated from the Secure score recommendations tab to the All recommendations tab.

When the move occurs, the custom recommendations will be found via a new "recommendation type" filter.

Learn more in Create custom security initiatives and policies.

Deprecating Microsoft Defender for IoT device recommendations

Estimated date for change: March 2022

Microsoft Defender for IoT device recommendations will no longer be visible in Microsoft Defender for Cloud. These recommendations will still be available on Microsoft Defender for IoT's Recommendations page, and in Microsoft Sentinel.

The following recommendations will be deprecated:

ASSESSMENT KEY	RECOMMENDATIONS
1a36f14a-8bd8-45f5-abe5-eef88d76ab5b: IoT Devices	Open Ports On Device

ASSESSMENT KEY	RECOMMENDATIONS
ba975338-f956-41e7-a9f2-7614832d382d: IoT Devices	Permissive firewall rule in the input chain was found
beb62be3-5e78-49bd-ac5f-099250ef3c7c: IoT Devices	Permissive firewall policy in one of the chains was found
d5a8d84a-9ad0-42e2-80e0-d38e3d46028a: IoT Devices	Permissive firewall rule in the output chain was found
5f65e47f-7a00-4bf3-acae-90ee441ee876: IoT Devices	Operating system baseline validation failure
a9a59ebb-5d6f-42f5-92a1-036fd0fd1879: IoT Devices	Agent sending underutilized messages
2acc27c6-5fdb-405e-9080-cb66b850c8f5: IoT Devices	TLS cipher suite upgrade needed
d74d2738-2485-4103-9919-69c7e63776ec: IoT Devices	Auditd process stopped sending events

Deprecating Microsoft Defender for IoT device alerts

Estimated date for change: March 2022

All Microsoft Defender for IoT device alerts will no longer be visible in Microsoft Defender for Cloud. These alerts will still be available on Microsoft Defender for IoT's Alert page, and in Microsoft Sentinel.

Next steps

For all recent changes to Defender for Cloud, see What's new in Microsoft Defender for Cloud?

Supported platforms

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page shows the platforms and environments supported by Microsoft Defender for Cloud.

Combinations of environments

Microsoft Defender for Cloud supports virtual machines and servers on different types of hybrid environments:

- Only Azure
- Azure and on-premises
- Azure and other clouds
- Azure, other clouds, and on-premises

For an Azure environment activated on an Azure subscription, Microsoft Defender for Cloud will automatically discover IaaS resources that are deployed within the subscription.

Supported operating systems

Defender for Cloud depends on the Log Analytics agent. Ensure your machines are running one of the supported operating systems for this agent as described on the following pages:

- Log Analytics agent for Windows supported operating systems
- Log Analytics agent for Linux supported operating systems

Also ensure your Log Analytics agent is properly configured to send data to Defender for Cloud

To learn more about the specific Defender for Cloud features available on Windows and Linux, see Feature coverage for machines.

> **NOTE**
> Even though **Microsoft Defender for servers** is designed to protect servers, most of its features are supported for Windows 10 machines. One feature that isn't currently supported is Defender for Cloud's integrated EDR solution: Microsoft Defender for Endpoint.

Managed virtual machine services

Virtual machines are also created in a customer subscription as part of some Azure-managed services as well, such as Azure Kubernetes (AKS), Azure Databricks, and more. Defender for Cloud discovers these virtual machines too, and the Log Analytics agent can be installed and configured if a supported OS is available.

Cloud Services

Virtual machines that run in a cloud service are also supported. Only cloud services web and worker roles that run in production slots are monitored. To learn more about cloud services, see Overview of Azure Cloud Services.

Protection for VMs residing in Azure Stack Hub is also supported. For more information about Defender for Cloud's integration with Azure Stack Hub, see Onboard your Azure Stack Hub virtual machines to Defender for Cloud.

Next steps

- Learn how Defender for Cloud collects data using the Log Analytics Agent.
- Learn how Defender for Cloud manages and safeguards data.
- Learn how to plan and understand the design considerations to adopt Microsoft Defender for Cloud.

Feature coverage for machines

2/15/2022 • 6 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The **tabs** below show the features of Microsoft Defender for Cloud that are available for Windows and Linux machines.

Supported features for virtual machines and servers

- **Windows machines**
- **Linux machines**
- **Multi-cloud machines**

FEATURE	AZURE VIRTUAL MACHINES	AZURE VIRTUAL MACHINE SCALE SETS	AZURE ARC-ENABLED MACHINES	DEFENDER FOR SERVERS REQUIRED
Microsoft Defender for Endpoint integration	✔ (on supported versions)	✔ (on supported versions)	✔	Yes
Virtual machine behavioral analytics (and security alerts)	✔	✔	✔	Yes
Fileless security alerts	✔	✔	✔	Yes
Network-based security alerts	✔	✔	-	Yes
Just-in-time VM access	✔	-	-	Yes
Integrated Qualys vulnerability scanner	✔	-	✔	Yes
File integrity monitoring	✔	✔	✔	Yes
Adaptive application controls	✔	-	✔	Yes
Network map	✔	✔	-	Yes

FEATURE	AZURE VIRTUAL MACHINES	AZURE VIRTUAL MACHINE SCALE SETS	AZURE ARC-ENABLED MACHINES	DEFENDER FOR SERVERS REQUIRED
Adaptive network hardening	✔	-	-	Yes
Regulatory compliance dashboard & reports	✔	✔	✔	Yes
Docker host hardening	-	-	-	Yes
Missing OS patches assessment	✔	✔	✔	Azure: No Azure Arc-enabled: Yes
Security misconfigurations assessment	✔	✔	✔	Azure: No Azure Arc-enabled: Yes
Endpoint protection assessment	✔	✔	✔	Azure: No Azure Arc-enabled: Yes
Disk encryption assessment	✔ (for supported scenarios)	✔	-	No
Third-party vulnerability assessment	✔	-	✔	No
Network security assessment	✔	✔	-	No

TIP

To experiment with features that are only available with enhanced security features enabled, you can enroll in a 30-day trial. For more information, see the pricing page.

Supported endpoint protection solutions

The following table provides a matrix of supported endpoint protection solutions and whether you can use Microsoft Defender for Cloud to install each solution for you.

For information about when recommendations are generated for each of these solutions, see Endpoint Protection Assessment and Recommendations.

SOLUTION	SUPPORTED PLATFORMS	DEFENDER FOR CLOUD INSTALLATION
Microsoft Defender Antivirus	Windows Server 2016 or later	No (built into OS)
System Center Endpoint Protection (Microsoft Antimalware)	Windows Server 2012 R2	Via extension
Trend Micro – Deep Security	Windows Server (all)	No
Symantec v12.1.1100+	Windows Server (all)	No
McAfee v10+	Windows Server (all)	No
McAfee v10+	Linux (GA)	No
Microsoft Defender for Endpoint for Linux[1]	Linux (GA)	Via extension
Sophos V9+	Linux (GA)	No

[1] It's not enough to have Microsoft Defender for Endpoint on the Linux machine: the machine will only appear as healthy if the always-on scanning feature (also known as real-time protection (RTP)) is active. By default, the RTP feature is **disabled** to avoid clashes with other AV software.

Feature support in government and national clouds

FEATURE/SERVICE	AZURE	AZURE GOVERNMENT	AZURE CHINA 21VIANET
Defender for Cloud free features			
- Continuous export	GA	GA	GA
- Workflow automation	GA	GA	GA
- Recommendation exemption rules	Public Preview	Not Available	Not Available
- Alert suppression rules	GA	GA	GA
- Email notifications for security alerts	GA	GA	GA
- Auto provisioning for agents and extensions	GA	GA	GA
- Asset inventory	GA	GA	GA
- Azure Monitor Workbooks reports in Microsoft Defender for Cloud's workbooks gallery	GA	GA	GA

FEATURE/SERVICE	AZURE	AZURE GOVERNMENT	AZURE CHINA 21VIANET
- Integration with Microsoft Defender for Cloud Apps	GA	Not Available	Not Available
Microsoft Defender plans and extensions			
- Microsoft Defender for servers	GA	GA	GA
- Microsoft Defender for App Service	GA	Not Available	Not Available
- Microsoft Defender for DNS	GA	GA	GA
- Microsoft Defender for container registries [1]	GA	GA [2]	GA [2]
- Microsoft Defender for container registries scanning of images in CI/CD workflows [3]	Public Preview	Not Available	Not Available
- Microsoft Defender for Kubernetes [4]	GA	GA	GA
- Microsoft Defender for Containers [10]	GA	GA	GA
- Defender extension for Azure Arc-enabled Kubernetes clusters, servers or data services [5]	Public Preview	Not Available	Not Available
- Microsoft Defender for Azure SQL database servers	GA	GA	GA [9]
- Microsoft Defender for SQL servers on machines	GA	GA	Not Available
- Microsoft Defender for open-source relational databases	GA	Not Available	Not Available
- Microsoft Defender for Key Vault	GA	Not Available	Not Available
- Microsoft Defender for Resource Manager	GA	GA	GA
- Microsoft Defender for Storage [6]	GA	GA	Not Available

FEATURE/SERVICE	AZURE	AZURE GOVERNMENT	AZURE CHINA 21VIANET
- Threat protection for Cosmos DB	Public Preview	Not Available	Not Available
- Kubernetes workload protection	GA	GA	GA
- Bi-directional alert synchronization with Sentinel	Public Preview	Not Available	Not Available
Microsoft Defender for servers features [7]			
- Just-in-time VM access	GA	GA	GA
- File integrity monitoring	GA	GA	GA
- Adaptive application controls	GA	GA	GA
- Adaptive network hardening	GA	Not Available	Not Available
- Docker host hardening	GA	GA	GA
- Integrated Qualys vulnerability scanner	GA	Not Available	Not Available
- Regulatory compliance dashboard & reports [8]	GA	GA	GA
- Microsoft Defender for Endpoint deployment and integrated license	GA	GA	Not Available
- Connect AWS account	GA	Not Available	Not Available
- Connect GCP account	GA	Not Available	Not Available

[1] Partially GA: The ability to disable specific findings from vulnerability scans is in public preview.

[2] Vulnerability scans of container registries on the Azure Government cloud can only be performed with the scan on push feature.

[3] Requires Microsoft Defender for container registries.

[4] Partially GA: Support for Azure Arc-enabled clusters is in public preview and not available on Azure Government.

[5] Requires Microsoft Defender for Kubernetes or Microsoft Defender for Containers.

[6] Partially GA: Some of the threat protection alerts from Microsoft Defender for Storage are in public preview.

[7] These features all require Microsoft Defender for servers.

[8] There may be differences in the standards offered per cloud type.

[9] Partially GA: Subset of alerts and vulnerability assessment for SQL servers. Behavioral threat protections aren't available.

[10] Partially GA: Support for Arc-enabled Kubernetes clusters (and therefore AWS EKS too) is in public preview and not available on Azure Government. Run-time visibility of vulnerabilities in container images is also a preview feature.

Next steps

- Learn how Defender for Cloud collects data using the Log Analytics Agent.
- Learn how Defender for Cloud manages and safeguards data.
- Review the platforms that support Defender for Cloud.

Feature coverage for Azure PaaS services

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The table below shows the availability of Microsoft Defender for Cloud features for the supported Azure PaaS resources.

SERVICE	RECOMMENDATIONS (FREE)	SECURITY ALERTS	VULNERABILITY ASSESSMENT
Azure App Service	✔	✔	-
Azure Automation account	✔	-	-
Azure Batch account	✔	-	-
Azure Blob Storage	✔	✔	-
Azure Cache for Redis	✔	-	-
Azure Cloud Services	✔	-	-
Azure Cognitive Search	✔	-	-
Azure Container Registry	✔	✔	✔
Azure Cosmos DB*	✔	✔	-
Azure Data Lake Analytics	✔	-	-
Azure Data Lake Storage	✔	✔	-
Azure Database for MySQL*	-	✔	-
Azure Database for PostgreSQL*	-	✔	-
Azure Event Hubs namespace	✔	-	-
Azure Functions app	✔	-	-
Azure Key Vault	✔	✔	-

SERVICE	RECOMMENDATIONS (FREE)	SECURITY ALERTS	VULNERABILITY ASSESSMENT
Azure Kubernetes Service	✔	✔	-
Azure Load Balancer	✔	-	-
Azure Logic Apps	✔	-	-
Azure SQL Database	✔	✔	✔
Azure SQL Managed Instance	✔	✔	✔
Azure Service Bus namespace	✔	-	-
Azure Service Fabric account	✔	-	-
Azure Storage accounts	✔	✔	-
Azure Stream Analytics	✔	-	-
Azure Subscription	✔ **	✔	-
Azure Virtual Network (incl. subnets, NICs, and network security groups)	✔	-	-

* These features are currently supported in preview.

** Azure Active Directory (Azure AD) recommendations are available only for subscriptions with enhanced security features enabled.

Permissions in Microsoft Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud uses Azure role-based access control (Azure RBAC), which provides built-in roles that can be assigned to users, groups, and services in Azure.

Defender for Cloud assesses the configuration of your resources to identify security issues and vulnerabilities. In Defender for Cloud, you only see information related to a resource when you are assigned the role of Owner, Contributor, or Reader for the subscription or the resource's resource group.

In addition to the built-in roles, there are two roles specific to Defender for Cloud:

- **Security Reader**: A user that belongs to this role has viewing rights to Defender for Cloud. The user can view recommendations, alerts, a security policy, and security states, but cannot make changes.
- **Security Admin**: A user that belongs to this role has the same rights as the Security Reader and can also update the security policy and dismiss alerts and recommendations.

> **NOTE**
>
> The security roles, Security Reader and Security Admin, have access only in Defender for Cloud. The security roles do not have access to other Azure services such as Storage, Web & Mobile, or Internet of Things.

Roles and allowed actions

The following table displays roles and allowed actions in Defender for Cloud.

ACTION	SECURITY READER / READER	SECURITY ADMIN	CONTRIBUTOR / OWNER	CONTRIBUTOR	OWNER
			(Resource group level)	**(Subscription level)**	**(Subscription level)**
Add/assign initiatives (including) regulatory compliance standards)	-	-	-	✔	✔
Edit security policy	-	✔	-	✔	✔

ACTION	SECURITY READER / READER	SECURITY ADMIN	CONTRIBUTOR / OWNER	CONTRIBUTOR	OWNER
Enable / disable Microsoft Defender plans	-	✔	-	✔	✔
Dismiss alerts	-	✔	-	✔	✔
Apply security recommendations for a resource (and use Fix)	-	-	✔	✔	✔
View alerts and recommendations	✔	✔	✔	✔	✔

For **auto provisioning**, the specific role required depends on the extension you're deploying. For full details, check the tab for the specific extension in the availability table on the auto provisioning quick start page.

> **NOTE**
> We recommend that you assign the least permissive role needed for users to complete their tasks. For example, assign the Reader role to users who only need to view information about the security health of a resource but not take action, such as applying recommendations or editing policies.

Next steps

This article explained how Defender for Cloud uses Azure RBAC to assign permissions to users and identified the allowed actions for each role. Now that you're familiar with the role assignments needed to monitor the security state of your subscription, edit security policies, and apply recommendations, learn how to:

- Set security policies in Defender for Cloud
- Manage security recommendations in Defender for Cloud
- Manage and respond to security alerts in Defender for Cloud
- Monitor partner security solutions

Quickstart: Set up Microsoft Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud provides unified security management and threat protection across your hybrid and multi-cloud workloads. While the free features offer limited security for your Azure resources only, enabling enhanced security features extends these capabilities to on-premises and other clouds. Defender for Cloud helps you find and fix security vulnerabilities, apply access and application controls to block malicious activity, detect threats using analytics and intelligence, and respond quickly when under attack. You can try the enhanced security features at no cost. To learn more, see the pricing page.

This quickstart section will walk you through all the recommended steps to enable Microsoft Defender for Cloud and the enhanced security features. When you've completed all the quickstart steps, you'll have:

- Defender for Cloud enabled on your Azure subscriptions
- Enhanced security features enabled on your Azure subscriptions
- Automatic data collection set up
- Email notifications set up for security alerts
- Your hybrid and multi-cloud machines connected to Azure

Prerequisites

To get started with Defender for Cloud, you must have a subscription to Microsoft Azure. If you don't have a subscription, you can sign up for a free account.

To enable enhanced security features on a subscription, you must be assigned the role of Subscription Owner, Subscription Contributor, or Security Admin.

Enable Defender for Cloud on your Azure subscription

> **TIP**
> To enable Defender for Cloud on all subscriptions within a management group, see Enable Defender for Cloud on multiple Azure subscriptions.

1. Sign into the Azure portal.
2. From the portal's menu, select **Defender for Cloud**.

   Defender for Cloud's overview page opens.

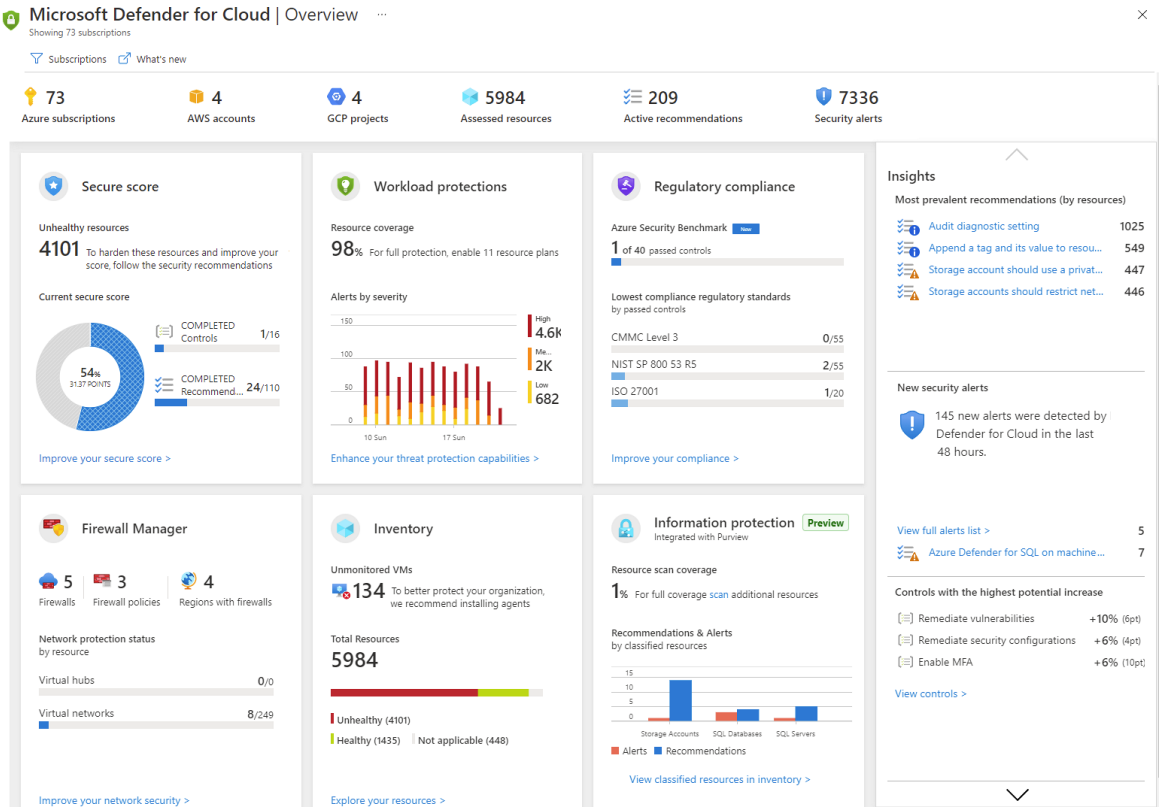

Defender for Cloud – Overview provides a unified view into the security posture of your hybrid cloud workloads, helping you discover and assess the security of your workloads and to identify and mitigate risks. Learn more in Microsoft Defender for Cloud's overview page.

Defender for Cloud automatically, at no cost, enables any of your Azure subscriptions not previously onboarded by you or another subscription user.

You can view and filter the list of subscriptions by selecting the **Subscriptions** menu item. Defender for Cloud will adjust the display to reflect the security posture of the selected subscriptions.

Within minutes of launching Defender for Cloud the first time, you might see:

- **Recommendations** for ways to improve the security of your connected resources.
- An inventory of your resources that are now being assessed by Defender for Cloud, along with the security posture of each.

To take full advantage of Defender for Cloud, continue with the next steps of the quickstart section.

Next steps

In this quickstart you enabled Defender for Cloud. The next step is to enable enhanced security features for unified security management and threat protection across your hybrid cloud workloads.

Quickstart: Enable enhanced security features

Quickstart: Enable enhanced security features

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

To learn about the benefits of enhanced security features, see Microsoft Defender for Cloud's enhanced security features.

Prerequisites

For the purpose of the Defender for Cloud quickstarts and tutorials you must enable the enhanced security features.

You can protect an entire Azure subscription with Defender for Cloud's enhanced security features and the protections will be inherited by all resources within the subscription.

A free 30-day trial is available. For pricing details in your local currency or region, see the pricing page.

Enable enhanced security features from the Azure portal

To enable all Defender for Cloud features including threat protection capabilities, you must enable enhanced security features on the subscription containing the applicable workloads. Enabling it at the workspace level doesn't enable just-in-time VM access, adaptive application controls, and network detections for Azure resources. In addition, the only Microsoft Defender plans available at the workspace level are Microsoft Defender for servers and Microsoft Defender for SQL servers on machines.

- You can enable **Microsoft Defender for Storage accounts** at either the subscription level or resource level
- You can enable **Microsoft Defender for SQL** at either the subscription level or resource level
- You can enable **Microsoft Defender for open-source relational databases** at the resource level only

To enable enhanced security features on your subscriptions and workspaces:

- To enable enhanced security features on one subscription:

  1. From Defender for Cloud's main menu, select **Environment settings**.
  2. Select the subscription or workspace that you want to protect.
  3. Select **Enable all Microsoft Defender plans** to upgrade.
  4. Select **Save**.

  > **TIP**
  > You'll notice that each Microsoft Defender plan is priced separately and can be individually set to on or off. For example, you might want to turn off Defender for App Service on subscriptions that don't have an associated Azure App Service plan.

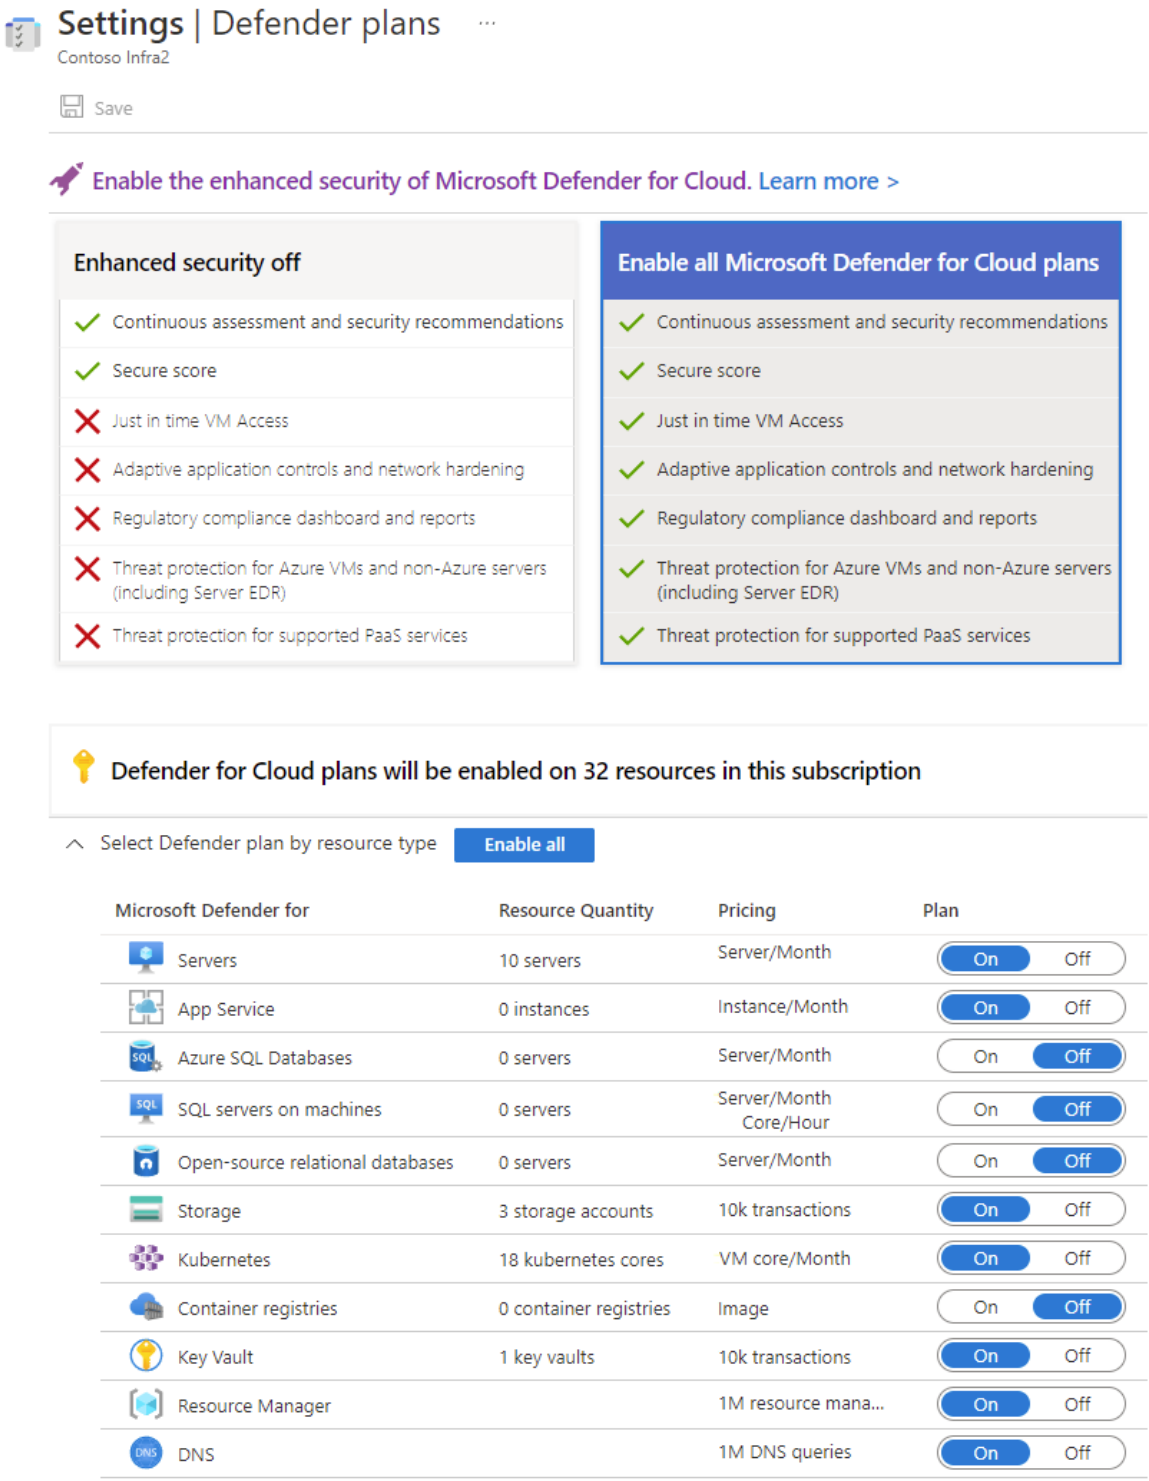

- To enable enhanced security on multiple subscriptions or workspaces:

  1. From Defender for Cloud's menu, select **Getting started**.

     The **Upgrade** tab lists subscriptions and workspaces eligible for onboarding.

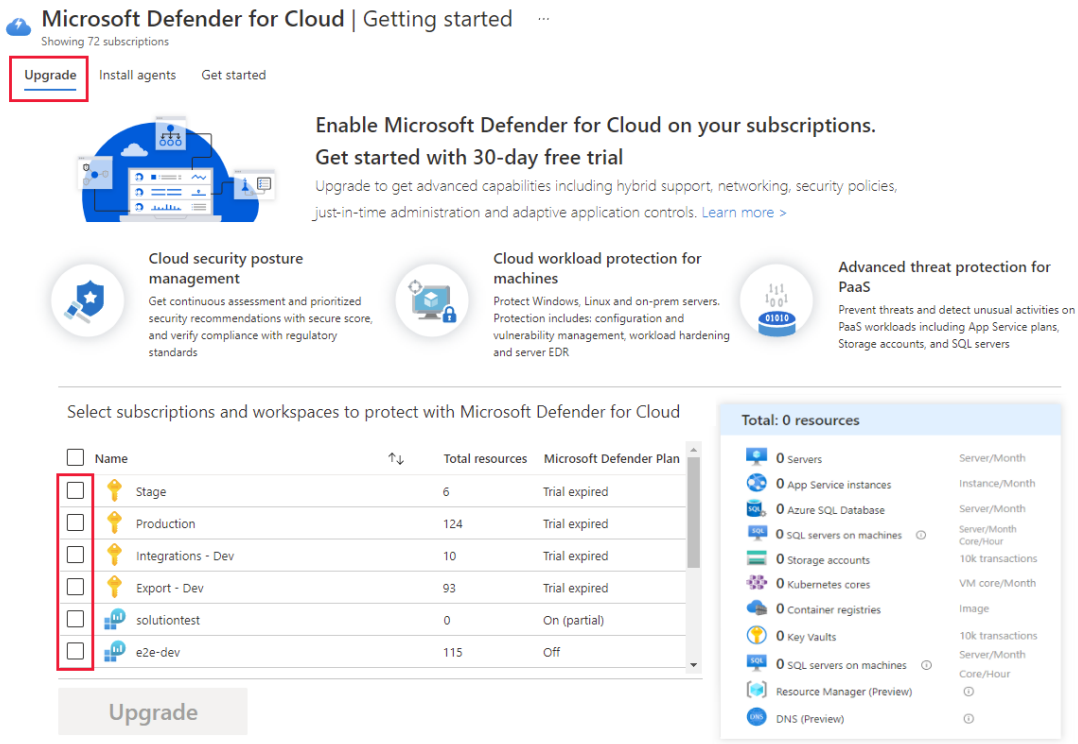

2. From the **Select subscriptions and workspaces to protect with Microsoft Defender for Cloud** list, select the subscriptions and workspaces to upgrade and select **Upgrade** to enable all Microsoft Defender for Cloud security features.

   - If you select subscriptions and workspaces that aren't eligible for trial, the next step will upgrade them and charges will begin.
   - If you select a workspace that's eligible for a free trial, the next step will begin a trial.

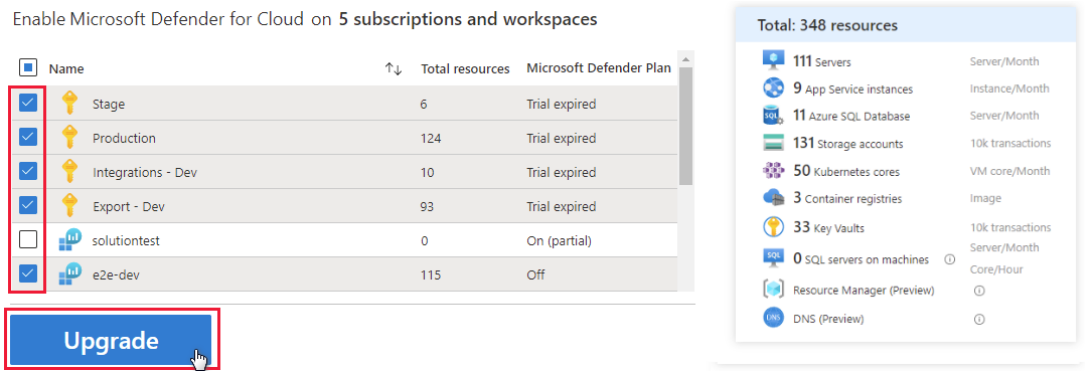

Disable enhanced security features

If you need to disable enhanced security features for a subscription, the procedure is the same but you select **Enhanced security off**:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.
3. Select **Defender plans** and select **Enhanced security off**.

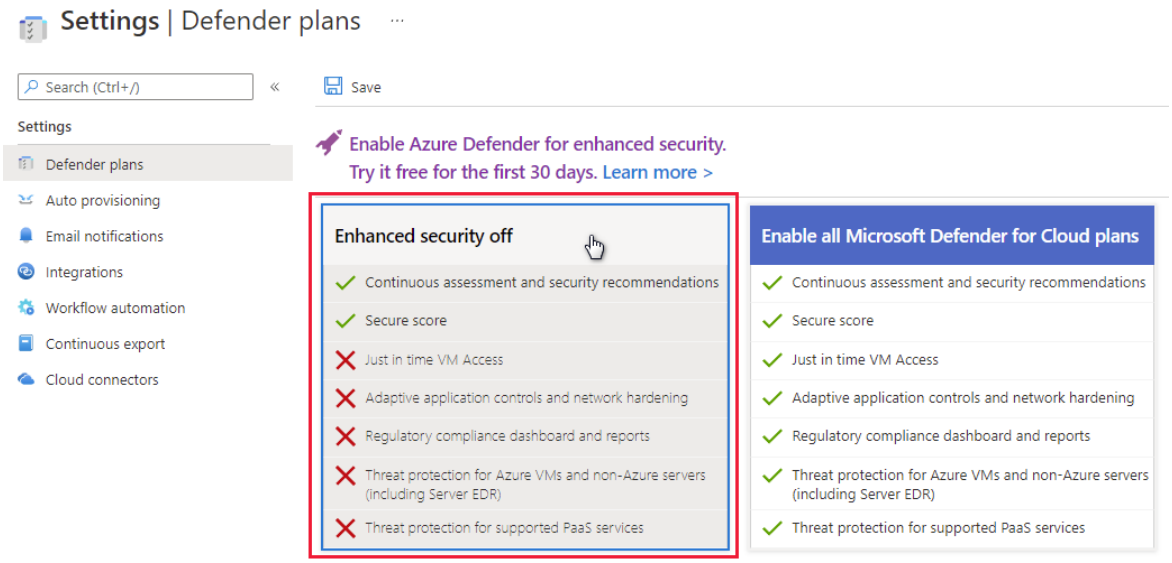

4. Select **Save**.

> **NOTE**
> After you disable enhanced security features - whether you disable a single plan or all plans at once - data collection may continue for a short period of time.

Next steps

Now that you've enabled enhanced security features, enable the necessary agents and extensions to perform automatic data collection as described in auto provisioning agents and extensions.

Connect your non-Azure machines to Microsoft Defender for Cloud

2/15/2022 • 5 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud can monitor the security posture of your non-Azure computers, but first you need to connect them to Azure.

You can connect your non-Azure computers in any of the following ways:

- Using Azure Arc-enabled servers (**recommended**)
- From Defender for Cloud's pages in the Azure portal (**Getting started** and **Inventory**)

Each of these is described on this page.

TIP

If you're connecting machines from other cloud providers, see Connect your AWS accounts or Connect your GCP accounts.

Add non-Azure machines with Azure Arc

The preferred way of adding your non-Azure machines to Microsoft Defender for Cloud is with Azure Arc-enabled servers.

A machine with Azure Arc-enabled servers becomes an Azure resource and - when you've installed the Log Analytics agent on it - appears in Defender for Cloud with recommendations like your other Azure resources.

In addition, Azure Arc-enabled servers provides enhanced capabilities such as the option to enable guest configuration policies on the machine, simplify deployment with other Azure services, and more. For an overview of the benefits, see Supported cloud operations.

NOTE

Defender for Cloud's auto-deploy tools for deploying the Log Analytics agent don't support machines running Azure Arc. When you've connected your machines using Azure Arc, use the relevant Defender for Cloud recommendation to deploy the agent and benefit from the full range of protections offered by Defender for Cloud:

- Log Analytics agent should be installed on your Linux-based Azure Arc machines
- Log Analytics agent should be installed on your Windows-based Azure Arc machines

Learn more about Azure Arc-enabled servers.

To deploy Azure Arc:

- For one machine, follow the instructions in Quickstart: Connect hybrid machines with Azure Arc-enabled servers.
- To connect multiple machines at scale to Azure Arc-enabled servers, see Connect hybrid machines to Azure at scale

> **TIP**
> If you're onboarding machines running on Amazon Web Services (AWS), Defender for Cloud's connector for AWS transparently handles the Azure Arc deployment for you. Learn more in Connect your AWS accounts to Microsoft Defender for Cloud.

Add non-Azure machines from the Azure portal

1. From Defender for Cloud's menu, open the **Getting started** page.
2. Select the **Get started** tab.
3. Below **Add non-Azure servers**, select **Configure** .

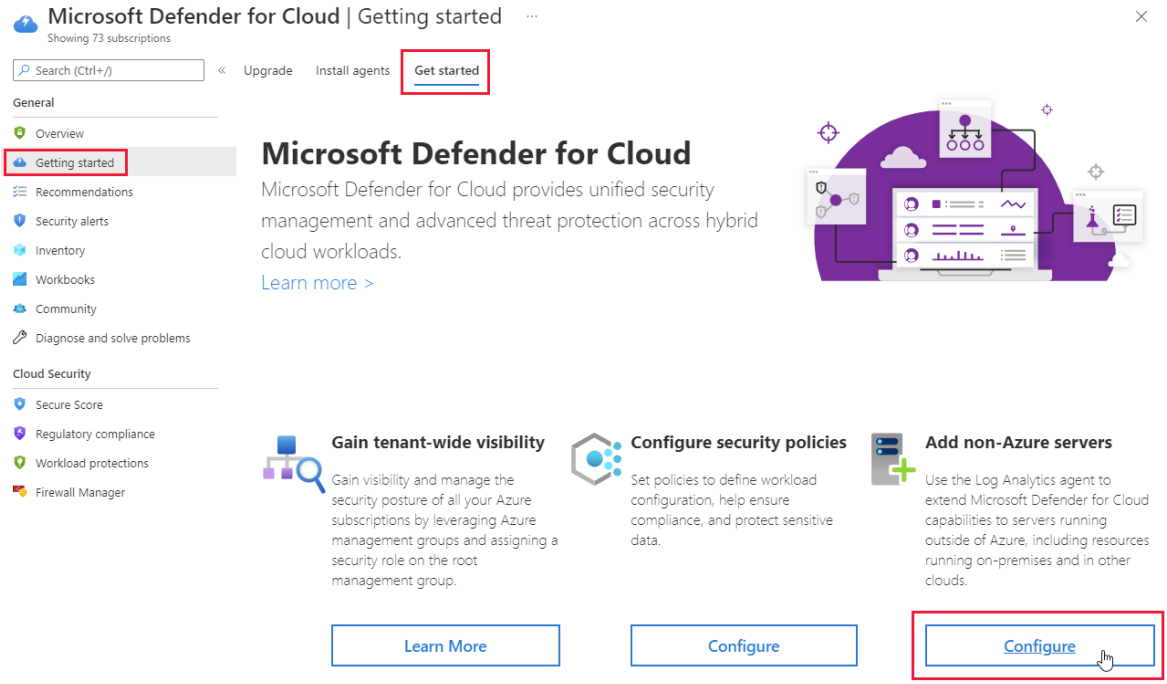

> **TIP**
> You can also open add machines from the **inventory** page's **Add non-Azure servers** button.
>
> 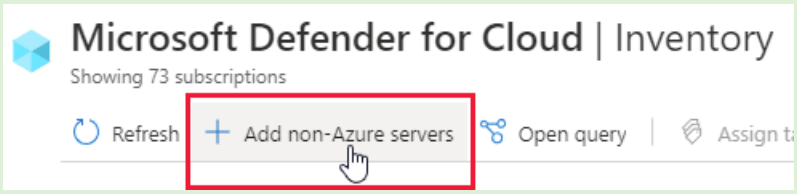
>

A list of your Log Analytics workspaces is shown. The list includes, if applicable, the default workspace created for you by Defender for Cloud when automatic provisioning was enabled. Select this workspace or another workspace you want to use.

You can add computers to an existing workspace or create a new workspace.

4. Optionally, to create a new workspace, select **Create new workspace**.

5. From the list of workspaces, select **Add Servers** for the relevant workspace.

   The **Agents management** page appears.

   From here, choose the relevant procedure below depending on the type of machines you're onboarding:

   - Onboard your Azure Stack Hub VMs
   - Onboard your Linux machines
   - Onboard your Windows machines

Onboard your Azure Stack Hub VMs

To add Azure Stack Hub VMs, you need the information on the **Agents management** page and to configure the **Azure Monitor, Update and Configuration Management** virtual machine extension on the virtual machines running on your Azure Stack Hub instance.

1. From the **Agents management** page, copy the **Workspace ID** and **Primary Key** into Notepad.
2. Log into your **Azure Stack Hub** portal and open the **Virtual machines** page.
3. Select the virtual machine that you want to protect with Defender for Cloud.

   > **TIP**
   > For information on how to create a virtual machine on Azure Stack Hub, see this quickstart for Windows virtual machines or this quickstart for Linux virtual machines.

4. Select **Extensions**. The list of virtual machine extensions installed on this virtual machine is shown.
5. Select the **Add** tab. The **New Resource** menu shows the list of available virtual machine extensions.
6. Select the **Azure Monitor, Update and Configuration Management** extension and select **Create**. The **Install extension** configuration page opens.

   > **NOTE**
   > If you do not see the **Azure Monitor, Update and Configuration Management** extension listed in your marketplace, please reach out to your Azure Stack Hub operator to make it available.

7. On the **Install extension** configuration page, paste the **Workspace ID** and **Workspace Key (Primary Key)** that you copied into Notepad in the previous step.
8. When you complete the configuration, select **OK**. The extension's status will show as **Provisioning Succeeded**. It might take up to one hour for the virtual machine to appear in Defender for Cloud.

Onboard your Linux machines

To add Linux machines, you need the WGET command from the **Agents management** page.

1. From the **Agents management** page, copy the **WGET** command into Notepad. Save this file to a location that can be accessible from your Linux computer.

2. On your Linux computer, open the file with the WGET command. Select the entire content and copy and paste it into a terminal console.

3. When the installation completes, you can validate that the `omsagent` is installed by running the `pgrep` command. The command will return the `omsagent` PID.

   The logs for the Agent can be found at: `/var/opt/microsoft/omsagent/\<workspace id>/log/`. It might take up to 30 minutes for the new Linux machine to appear in Defender for Cloud.

Onboard your Windows machines

To add Windows machines, you need the information on the **Agents management** page and to download the

appropriate agent file (32/64-bit).

1. Select the **Download Windows Agent** link applicable to your computer processor type to download the setup file.
2. From the **Agents management** page, copy the **Workspace ID** and **Primary Key** into Notepad.
3. Copy the downloaded setup file to the target computer and run it.
4. Follow the installation wizard (**Next**, **I Agree**, **Next**, **Next**).
    a. On the **Azure Log Analytics** page, paste the **Workspace ID** and **Workspace Key (Primary Key)** that you copied into Notepad.
    b. If the computer should report to a Log Analytics workspace in Azure Government cloud, select **Azure US Government** from the **Azure Cloud** dropdown list.
    c. If the computer needs to communicate through a proxy server to the Log Analytics service, select **Advanced** and provide the URL and port number of the proxy server.
    d. When you've entered all of the configuration settings, select **Next**.
    e. From the **Ready to Install** page, review the settings to be applied and select **Install**.
    f. On the **Configuration completed successfully** page, select **Finish**.

When complete, the **Microsoft Monitoring agent** appears in **Control Panel**. You can review your configuration there and verify that the agent is connected.

For further information on installing and configuring the agent, see Connect Windows machines.

Verifying

Congratulations! Now you can see your Azure and non-Azure machines together in one place. Open the asset inventory page and filter to the relevant resource types. These icons distinguish the types:

 Non-Azure machine

 Azure VM

 Azure Arc-enabled server

Next steps

This page showed you how to add your non-Azure machines to Microsoft Defender for Cloud. To monitor their status, use the inventory tools as explained in the following page:

- Explore and manage your resources with asset inventory

Connect your AWS accounts to Microsoft Defender for Cloud

2/15/2022 • 11 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

With cloud workloads commonly spanning multiple cloud platforms, cloud security services must do the same.

Microsoft Defender for Cloud protects workloads in Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP).

To protect your AWS-based resources, you can connect an account with one of two mechanisms:

- **Classic cloud connectors experience** - As part of the initial multi-cloud offering, we introduced these cloud connectors as a way to connect your AWS and GCP accounts. If you've already configured an AWS connector through the classic cloud connectors experience, we recommend deleting these connectors (as explained in Remove classic connectors), and connecting the account again using the newer mechanism. If you don't do this before creating the new connector through the environment settings page, do so afterwards to avoid seeing duplicate recommendations.
- **Environment settings page (in preview)** (recommended) - This preview page provides a greatly improved, simpler, onboarding experience (including auto provisioning). This mechanism also extends Defender for Cloud's enhanced security features to your AWS resources:
    - **Defender for Cloud's CSPM features** extend to your AWS resources. This agentless plan assesses your AWS resources according to AWS-specific security recommendations and these are included in your secure score. The resources will also be assessed for compliance with built-in standards specific to AWS (AWS CIS, AWS PCI DSS, and AWS Foundational Security Best Practices). Defender for Cloud's asset inventory page is a multi-cloud enabled feature helping you manage your AWS resources alongside your Azure resources.
    - **Microsoft Defender for Containers** extends Defender for Cloud's container threat detection and advanced defenses to your **Amazon EKS clusters**.
    - **Microsoft Defender for servers** brings threat detection and advanced defenses to your Windows and Linux EC2 instances. This plan includes the integrated license for Microsoft Defender for Endpoint, security baselines and OS level assessments, vulnerability assessment scanning, adaptive application controls (AAC), file integrity monitoring (FIM), and more.

For a reference list of all the recommendations Defender for Cloud can provide for AWS resources, see Security recommendations for AWS resources - a reference guide.

This screenshot shows AWS accounts displayed in Defender for Cloud's overview dashboard.

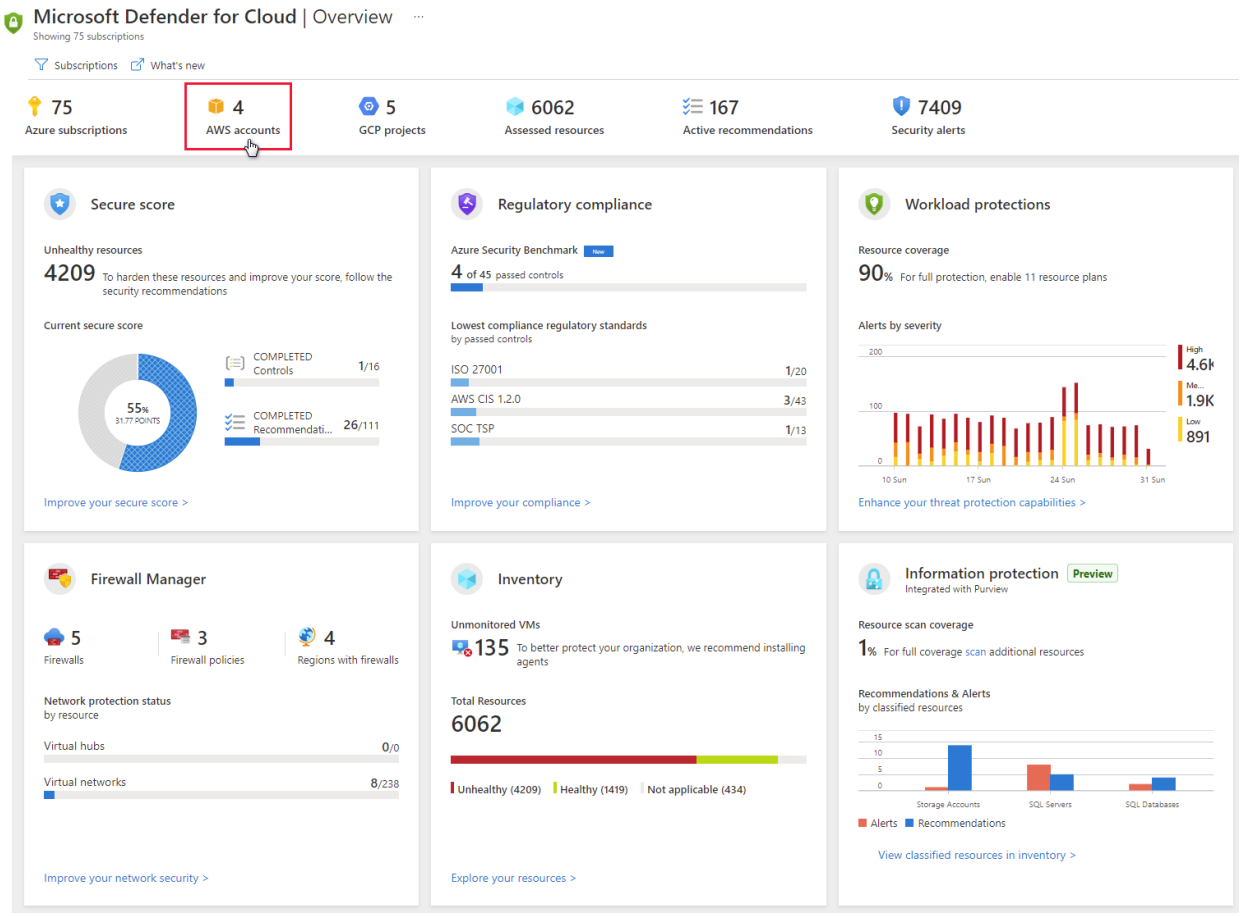

Availability

ASPECT	DETAILS
Release state:	Preview. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	The **CSPM plan** is free. The **Defender for Containers** plan is free during the preview. After which, it will be billed for AWS at the same price as for Azure resources. For every AWS machine connected to Azure with Azure Arc-enabled servers, the **Defender for servers** plan is billed at the same price as the Microsoft Defender for servers plan for Azure machines. If an AWS EC2 doesn't have the Azure Arc agent deployed, you won't be charged for that machine.
Required roles and permissions:	**Owner** on the relevant Azure subscription **Contributor** can also connect an AWS account if an owner provides the service principal details (required for the Defender for servers plan)
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

Prerequisites

- To connect an AWS account to your Azure subscription, you'll obviously need access to an AWS account.
- **To enable the Defender for Kubernetes plan**, you'll need:
  - At least one Amazon EKS cluster with permission to access to the EKS K8s API server. If you need to create a new EKS cluster, follow the instructions in Getting started with Amazon EKS – eksctl.
  - The resource capacity to create a new SQS queue, Kinesis Fire Hose delivery stream, and S3 bucket in the cluster's region.
- **To enable the Defender for servers plan**, you'll need:
  - Microsoft Defender for servers enabled (see Quickstart: Enable enhanced security features.
  - An active AWS account with EC2 instances managed by AWS Systems Manager (SSM) and using SSM agent. Some Amazon Machine Images (AMIs) have the SSM agent pre-installed, their AMIs are listed in AMIs with SSM Agent preinstalled. If your EC2 instances don't have the SSM Agent, follow the relevant instructions from Amazon:
    - Install SSM Agent for a hybrid environment (Windows)
    - Install SSM Agent for a hybrid environment (Linux)

Connect your AWS account

Follow the steps below to create your AWS cloud connector.

Remove 'classic' connectors

If you have any existing connectors created with the classic cloud connectors experience, remove them first:

1. From Defender for Cloud's menu, open **Environment settings** and select the option to switch back to the classic connectors experience.

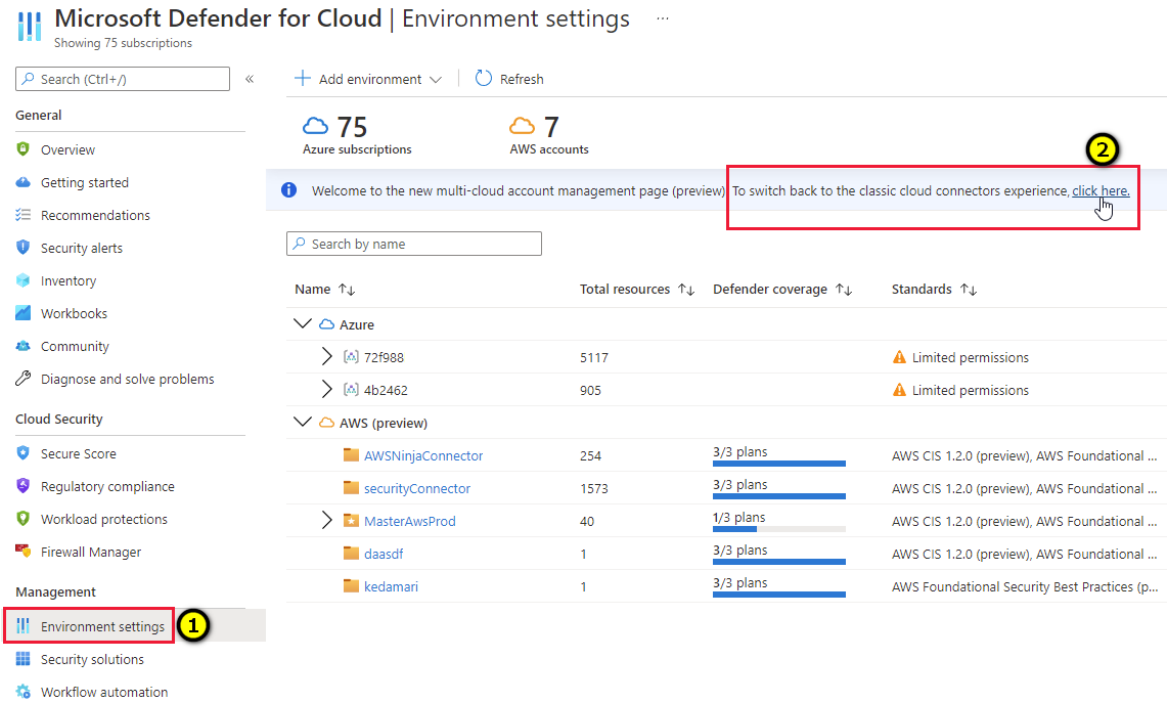

   From Defender for Cloud's menu, open **Environment settings**.

2. For each connector, select the "..." at the end of the row, and select **Delete**.
3. On AWS, delete the role ARN or the credentials created for the integration.

Create a new connector

1. From Defender for Cloud's menu, open **Environment settings**.

2. Select **Add environment** > **Amazon Web Services**.

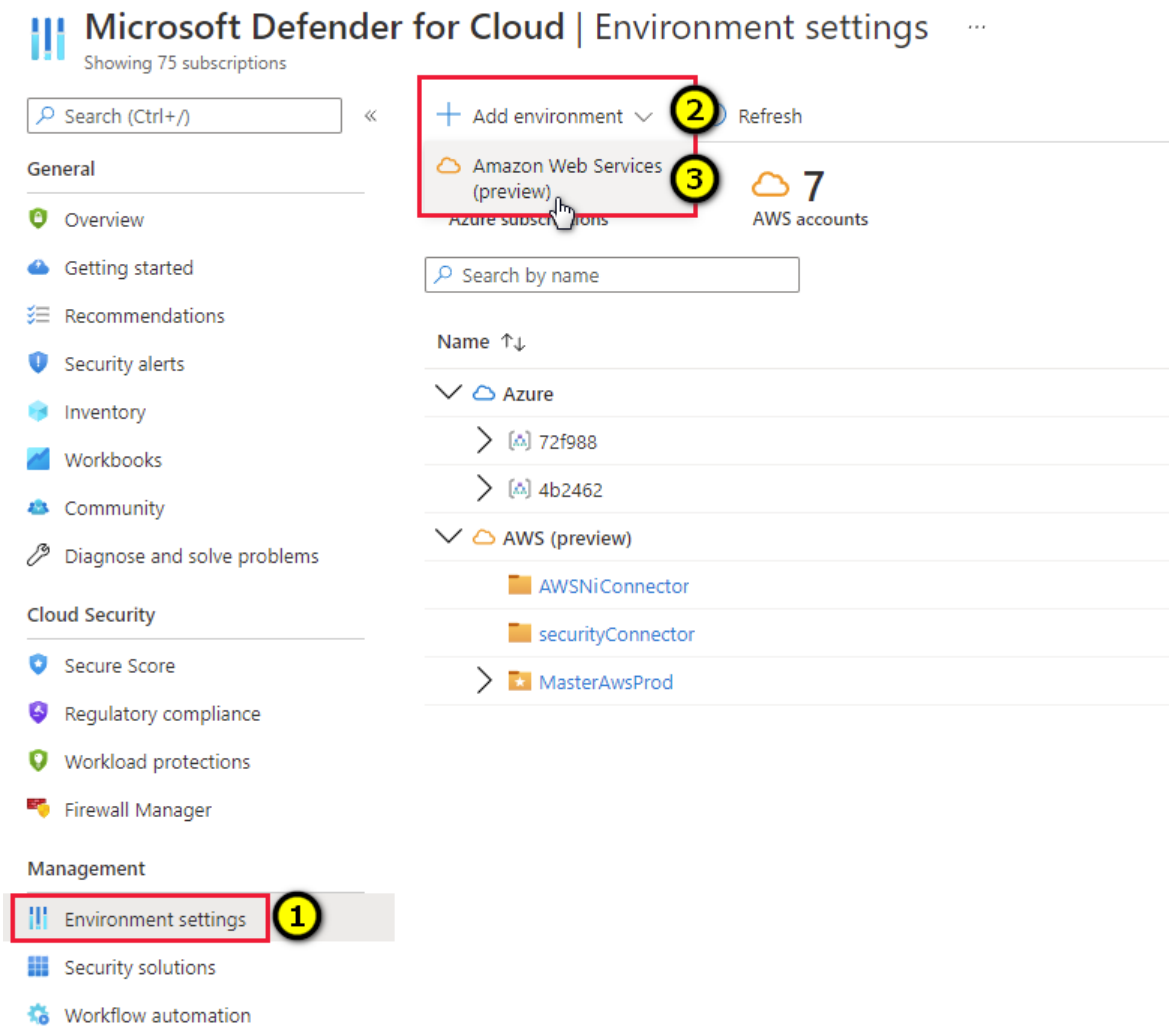

3. Enter the details of the AWS account, including the location where you'll store the connector resource, and select **Next: Select plans**.

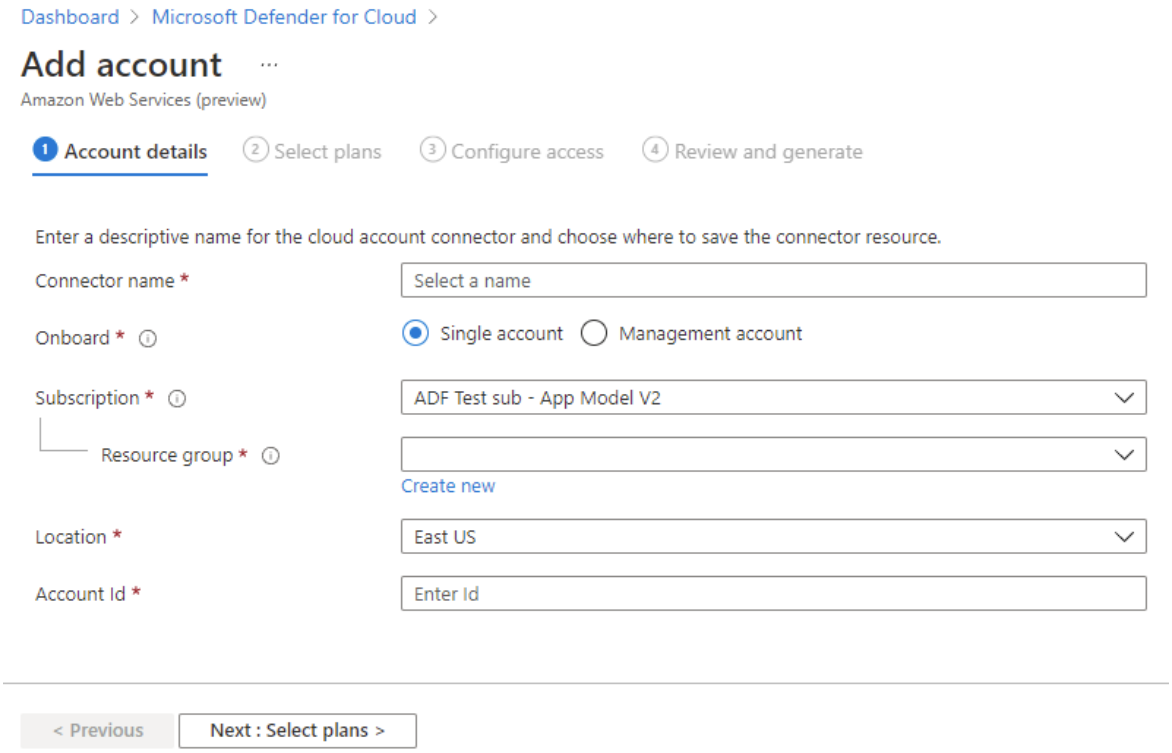

4. The select plans tab is where you choose which Defender for Cloud capabilities to enable for this AWS account.

> **NOTE**
>
> Each capability has its own requirements for permissions and might incur charges.

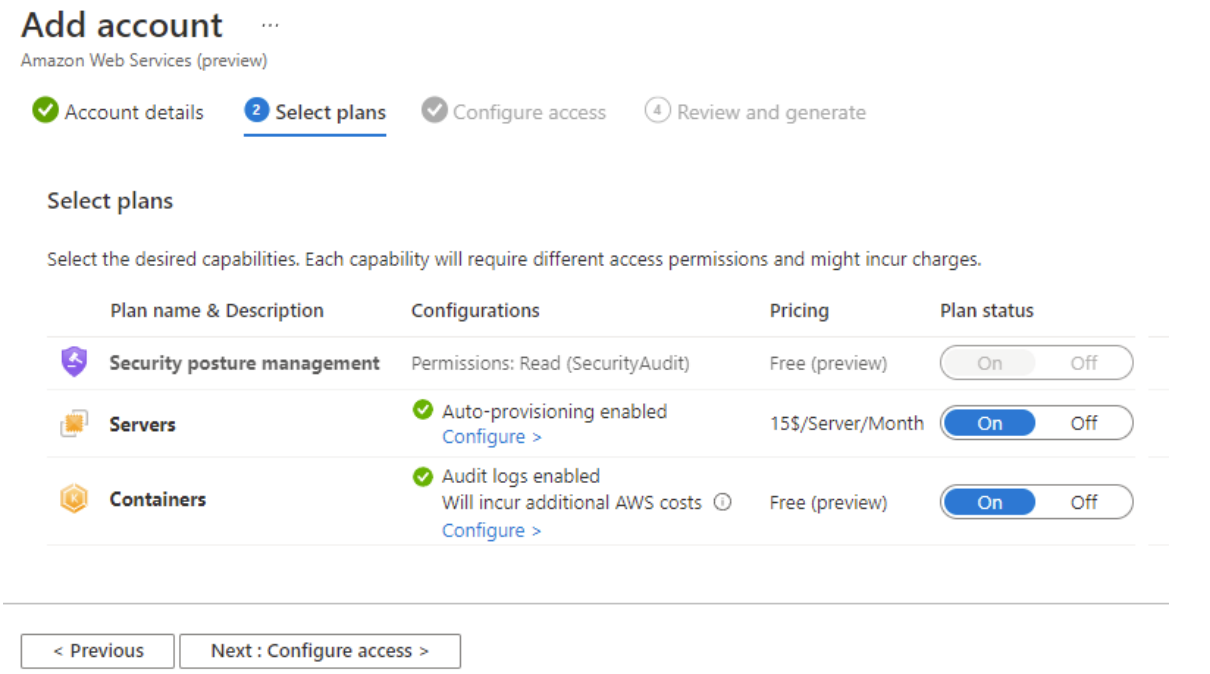

> **IMPORTANT**
>
> To present the current status of your recommendations, the CSPM plan queries the AWS resource APIs several times a day. These read-only API calls incur no charges, but they *are* registered in CloudTrail if you've enabled a trail for read events. As explained in the AWS documentation, there are no additional charges for keeping one trail. If you're exporting the data out of AWS (for example, to an external SIEM), this increased volume of calls might also increase ingestion costs. In such cases, We recommend filtering out the read-only calls from the Defender for Cloud user or role ARN: arn:aws:iam::[accountId]:role/CspmMonitorAws (this is the default role name, confirm the role name configured on your account).

- To extend Defender for Servers coverage to your AWS EC2, set the **Servers** plan to **On** and edit the configuration as required.
- For Defender for Kubernetes to protect your AWS EKS clusters, Azure Arc-enabled Kubernetes and the Defender extension should be installed. Set the **Containers** plan to **On**, and use the dedicated Defender for Cloud recommendation to deploy the extension (and Arc, if necessary) as explained in Protect Amazon Elastic Kubernetes Service clusters.

5. Complete the setup:

   a. Select **Next: Configure access**.
   b. Download the CloudFormation template.
   c. Using the downloaded CloudFormation template, create the stack in AWS as instructed on screen.
   d. Select **Next: Review and generate**.
   e. Select **Create**.

Defender for Cloud will immediately start scanning your AWS resources and you'll see security recommendations within a few hours. For a reference list of all the recommendations Defender for Cloud can

provide for AWS resources, see Security recommendations for AWS resources - a reference guide.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers
Required roles and permissions:	**Owner** on the relevant Azure subscription **Contributor** can also connect an AWS account if an owner provides the service principal details
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

Connect your AWS account

Follow the steps below to create your AWS cloud connector.

Step 1. Set up AWS Security Hub:

1. To view security recommendations for multiple regions, repeat the following steps for each relevant region.

   > **IMPORTANT**
   > If you're using an AWS management account, repeat the following three steps to configure the management account and all connected member accounts across all relevant regions

   a. Enable AWS Config.
   b. Enable AWS Security Hub.
   c. Verify that data is flowing to the Security Hub. When you first enable Security Hub, it might take several hours for data to be available.

Step 2. Set up authentication for Defender for Cloud in AWS

There are two ways to allow Defender for Cloud to authenticate to AWS:

- **Create an IAM role for Defender for Cloud** (Recommended) - The most secure method
- **AWS user for Defender for Cloud** - A less secure option if you don't have IAM enabled

Create an IAM role for Defender for Cloud

1. From your Amazon Web Services console, under **Security, Identity & Compliance**, select **IAM**.

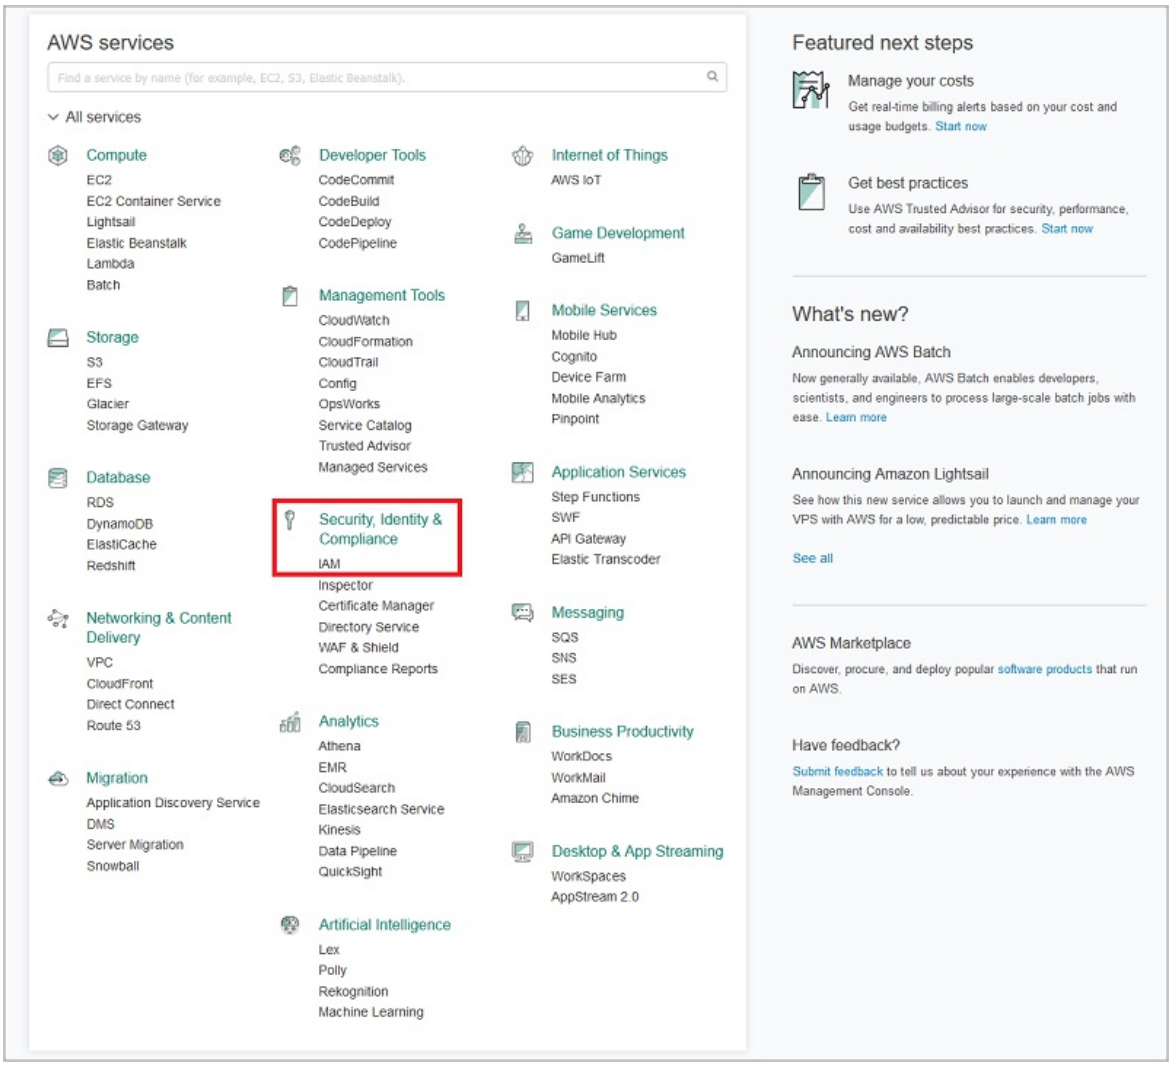

2. Select **Roles** and **Create role**.
3. Select **Another AWS account**.
4. Enter the following details:
   - **Account ID** - enter the Microsoft Account ID (**158177204117**) as shown in the AWS connector page in Defender for Cloud.
   - **Require External ID** - should be selected
   - **External ID** - enter the subscription ID as shown in the AWS connector page in Defender for Cloud
5. Select **Next**.
6. In the **Attach permission policies** section, select the following AWS managed policies:
   - SecurityAudit (`arn:aws:iam::aws:policy/SecurityAudit`)
   - AmazonSSMAutomationRole (`arn:aws:iam::aws:policy/service-role/AmazonSSMAutomationRole`)
   - AWSSecurityHubReadOnlyAccess (`arn:aws:iam::aws:policy/AWSSecurityHubReadOnlyAccess`)
7. Optionally add tags. Adding Tags to the user doesn't affect the connection.
8. Select **Next**.
9. In The Roles list, choose the role you created
10. Save the Amazon Resource Name (ARN) for later.

Create an AWS user for Defender for Cloud

1. Open the **Users** tab and select **Add user**.

2. In the **Details** step, enter a username for Defender for Cloud and ensure that you select **Programmatic access** for the AWS Access Type.
3. Select **Next Permissions**.
4. Select **Attach existing policies directly** and apply the following policies:
    - SecurityAudit
    - AmazonSSMAutomationRole
    - AWSSecurityHubReadOnlyAccess
5. Select **Next: Tags**. Optionally add tags. Adding Tags to the user doesn't affect the connection.
6. Select **Review**.
7. Save the automatically generated **Access key ID** and **Secret access key** CSV file for later.
8. Review the summary and select **Create user**.

Step 3. Configure the SSM Agent

AWS Systems Manager is required for automating tasks across your AWS resources. If your EC2 instances don't have the SSM Agent, follow the relevant instructions from Amazon:

- Installing and Configuring SSM Agent on Windows Instances
- Installing and Configuring SSM Agent on Amazon EC2 Linux Instances

Step 4. Complete Azure Arc prerequisites

1. Make sure the appropriate Azure resources providers are registered:
    - Microsoft.HybridCompute
    - Microsoft.GuestConfiguration
2. Create a Service Principal for onboarding at scale. As an **Owner** on the subscription you want to use for the onboarding, create a service principal for Azure Arc onboarding as described in Create a Service Principal for onboarding at scale.

Step 5. Connect AWS to Defender for Cloud

1. From Defender for Cloud's menu, open **Environment settings** and select the option to switch back to the classic connectors experience.

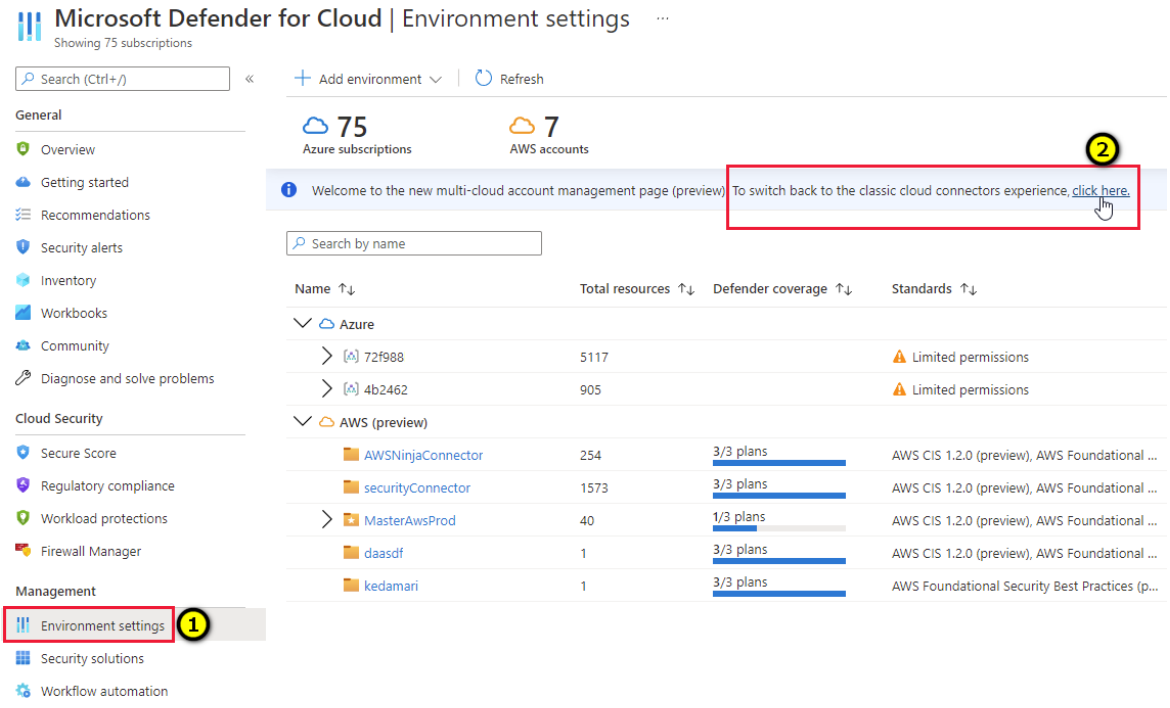

2. Select **Add AWS account**.

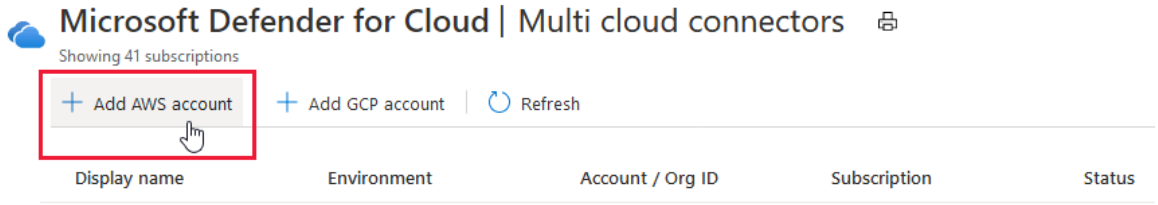

3. Configure the options in the **AWS authentication** tab:

   a. Enter a **Display name** for the connector.

   b. Confirm that the subscription is correct. It's the subscription that will include the connector and AWS Security Hub recommendations.

   c. Depending on the authentication option, you chose in Step 2. Set up authentication for Defender for Cloud in AWS:

      - Select **Assume Role** and paste the ARN from Create an IAM role for Defender for Cloud.

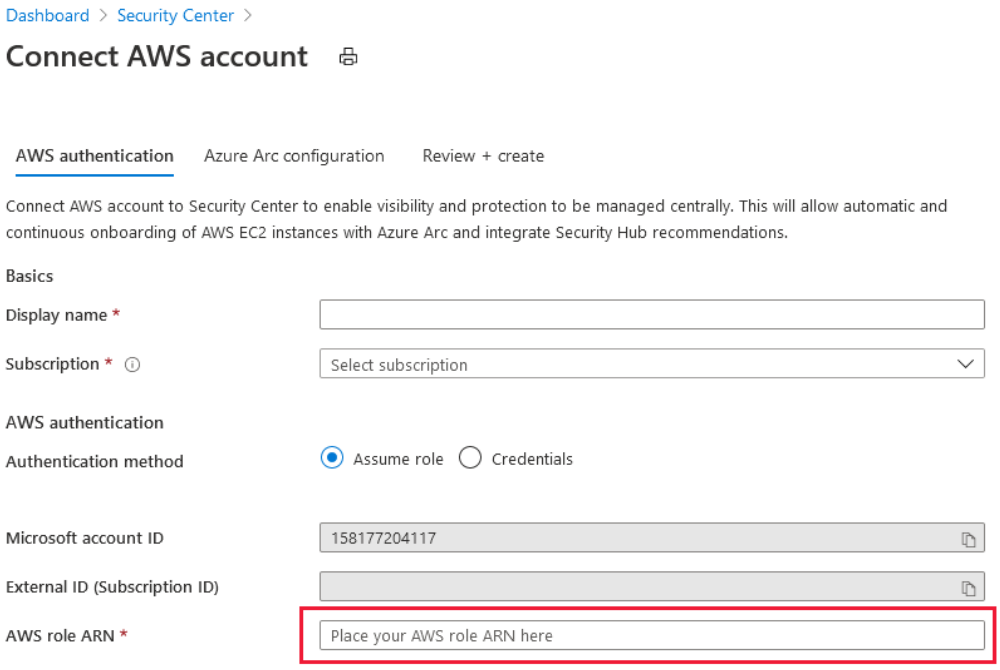

OR

- Select **Credentials** and paste the **access key** and **secret key** from the .csv file you saved in Create an AWS user for Defender for Cloud.

4. Select **Next**.

5. Configure the options in the **Azure Arc Configuration** tab:

   Defender for Cloud discovers the EC2 instances in the connected AWS account and uses SSM to onboard them to Azure Arc.

   > **TIP**
   > For the list of supported operating systems, see What operating systems for my EC2 instances are supported? in the FAQ.

   a. Select the **Resource Group** and **Azure Region** that the discovered AWS EC2s will be onboarded to in the selected subscription.

   b. Enter the **Service Principal ID** and **Service Principal Client Secret** for Azure Arc as described here Create a Service Principal for onboarding at scale

   c. If the machine is connecting to the internet via a proxy server, specify the proxy server IP address or the name and port number that the machine uses to communicate with the proxy server. Enter the value in the format `http://<proxyURL>:<proxyport>`

   d. Select **Review + create**.

      Review the summary information

      The Tags sections will list all Azure Tags that will be automatically created for each onboarded EC2 with its own relevant details to easily recognize it in Azure.

      Learn more about Azure Tags in Use tags to organize your Azure resources and management hierarchy.

Step 6. Confirmation

When the connector is successfully created, and AWS Security Hub has been configured properly:

- Defender for Cloud scans the environment for AWS EC2 instances, onboarding them to Azure Arc, enabling to install the Log Analytics agent and providing threat protection and security recommendations.
- The Defender for Cloud service scans for new AWS EC2 instances every 6 hours and onboards them according to the configuration.
- The AWS CIS standard will be shown in the Defender for Cloud's regulatory compliance dashboard.
- If Security Hub policy is enabled, recommendations will appear in the Defender for Cloud portal and the regulatory compliance dashboard 5-10 minutes after onboard completes.

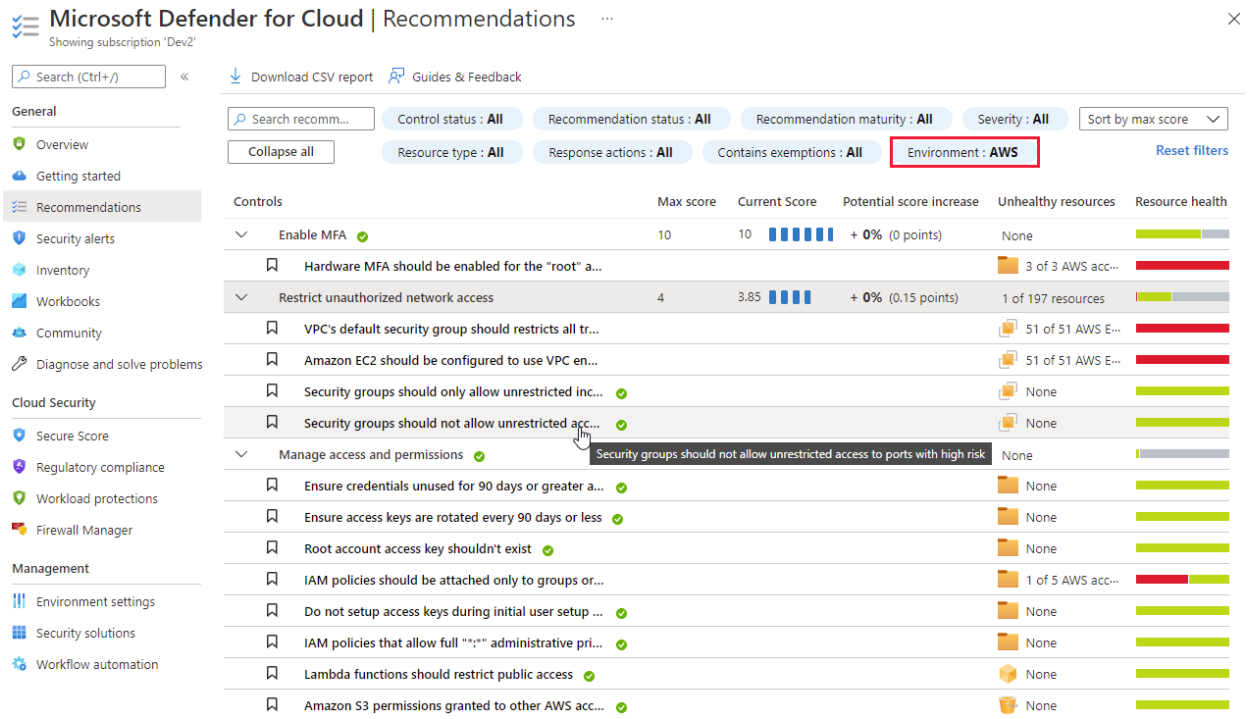

Monitoring your AWS resources

As you can see in the previous screenshot, Defender for Cloud's security recommendations page displays your AWS resources. You can use the environments filter to enjoy Defender for Cloud's multi-cloud capabilities: view the recommendations for Azure, AWS, and GCP resources together.

To view all the active recommendations for your resources by resource type, use Defender for Cloud's asset inventory page and filter to the AWS resource type in which you're interested:

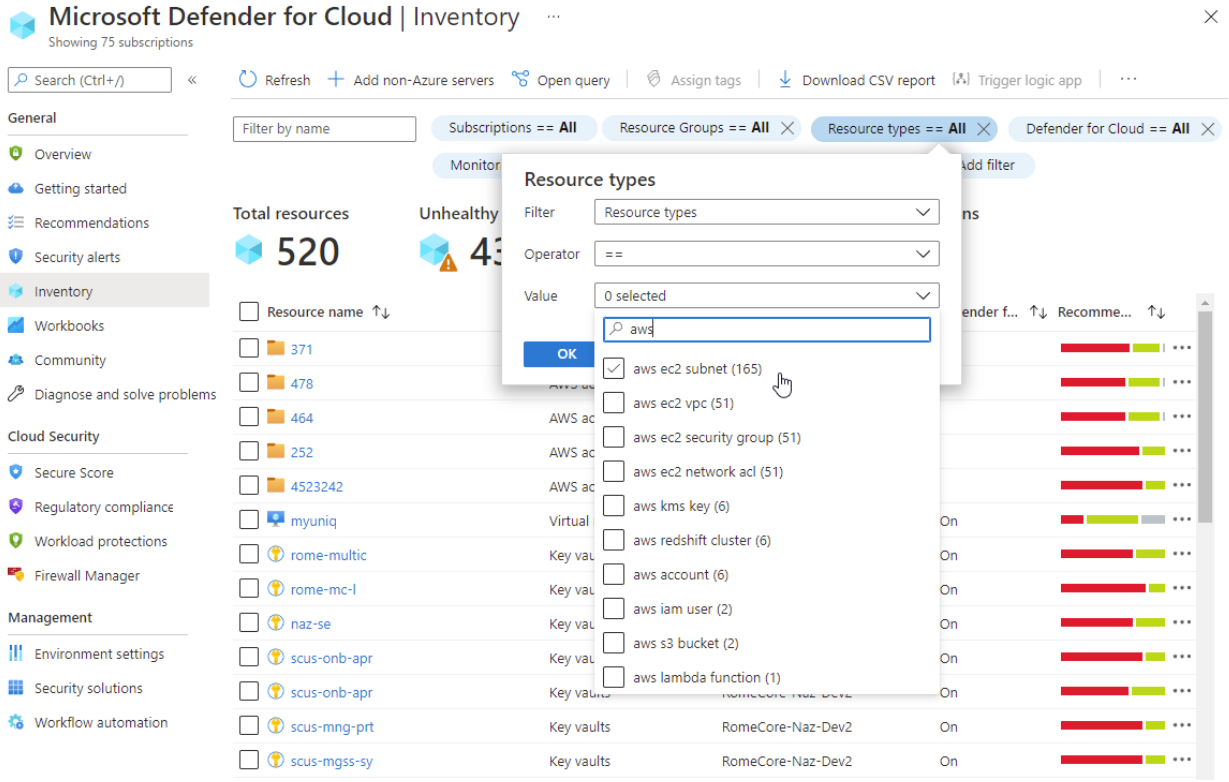

FAQ - AWS in Defender for Cloud

What operating systems for my EC2 instances are supported?

For a list of the AMIs with the SSM Agent preinstalled see this page in the AWS docs.

For other operating systems, the SSM Agent should be installed manually using the following instructions:

- Install SSM Agent for a hybrid environment (Windows)
- Install SSM Agent for a hybrid environment (Linux)

Next steps

Connecting your AWS account is part of the multi-cloud experience available in Microsoft Defender for Cloud. For related information, see the following page:

- Security recommendations for AWS resources - a reference guide.
- Connect your GCP accounts to Microsoft Defender for Cloud

Connect your GCP accounts to Microsoft Defender for Cloud

2/15/2022 • 5 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

With cloud workloads commonly spanning multiple cloud platforms, cloud security services must do the same.

Microsoft Defender for Cloud protects workloads in Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP).

Adding a GCP account to an Azure subscription connects Defender for Cloud with GCP Security Command. Defender for Cloud can then protect your resources across both of these cloud environments and provide:

- Detection of security misconfigurations
- A single view showing Defender for Cloud recommendations and GCP Security Command Center findings
- Incorporation of your GCP resources into Defender for Cloud's secure score calculations
- Integration of GCP Security Command Center recommendations based on the CIS standard into the Defender for Cloud's regulatory compliance dashboard

IMPORTANT

At Ignite Fall 2021, we announced an updated way of connecting your accounts from other cloud providers. This uses the new **Environment settings** page. GCP accounts aren't supported from that page. To connect a GCP account to your Azure subscription, you'll need to use the classic cloud connectors experience as described below.

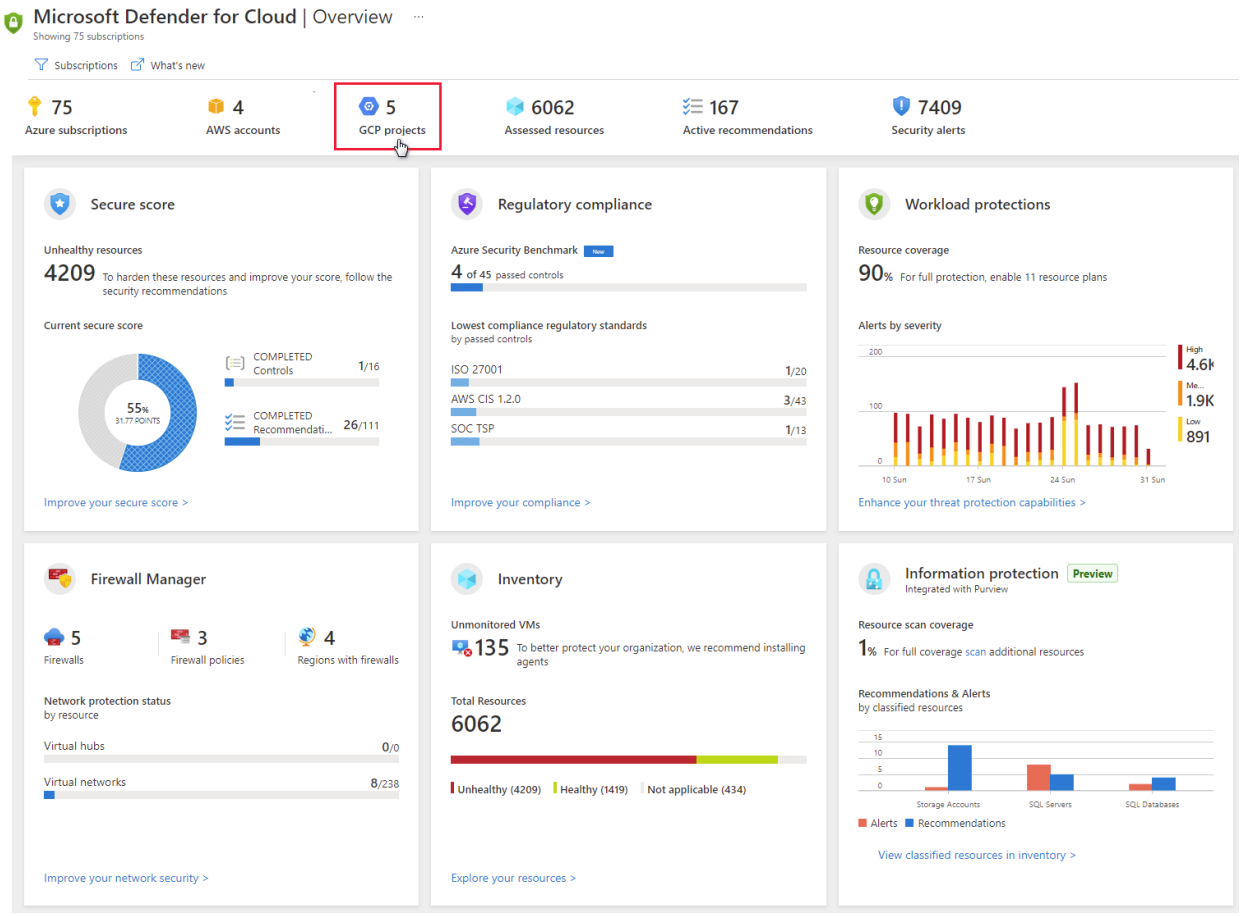

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers
Required roles and permissions:	**Owner** or **Contributor** on the relevant Azure Subscription
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

Connect your GCP account

Create a connector for every organization you want to monitor from Defender for Cloud.

When connecting your GCP accounts to specific Azure subscriptions, consider the Google Cloud resource hierarchy and these guidelines:

- You can connect your GCP accounts to Defender for Cloud in the *organization* level
- You can connect multiple organizations to one Azure subscription
- You can connect multiple organizations to multiple Azure subscriptions
- When you connect an organization, all *projects* within that organization are added to Defender for Cloud

Follow the steps below to create your GCP cloud connector.

Step 1. Set up GCP Security Command Center with Security Health Analytics

For all the GCP projects in your organization, you must also:

1. Set up **GCP Security Command Center** using these instructions from the GCP documentation.
2. Enable **Security Health Analytics** using these instructions from the GCP documentation.
3. Verify that there is data flowing to the Security Command Center.

The instructions for connecting your GCP environment for security configuration follow Google's recommendations for consuming security configuration recommendations. The integration leverages Google Security Command Center and will consume additional resources that might impact your billing.

When you first enable Security Health Analytics, it might take several hours for data to be available.

Step 2. Enable GCP Security Command Center API

1. From Google's **Cloud Console API Library**, select each project in the organization you want to connect to Microsoft Defender for Cloud.
2. In the API Library, find and select **Security Command Center API**.
3. On the API's page, select **ENABLE**.

Learn more about the Security Command Center API.

Step 3. Create a dedicated service account for the security configuration integration

1. In the **GCP Console**, select a project from the organization in which you're creating the required service account.

   > **NOTE**
   > When this service account is added at the organization level, it'll be used to access the data gathered by Security Command Center from all of the other enabled projects in the organization.

2. In the **Navigation menu**, Under **IAM & admin** options, select **Service accounts**.
3. Select **CREATE SERVICE ACCOUNT**.
4. Enter an account name, and select **Create**.
5. Specify the **Role** as **Defender for Cloud Admin Viewer**, and select **Continue**.
6. The **Grant users access to this service account** section is optional. Select **Done**.
7. Copy the **Email value** of the created service account, and save it for later use.
8. In the **Navigation menu**, Under **IAM & admin** options, select **IAM**
   a. Switch to organization level.
   b. Select **ADD**.
   c. In the **New members** field, paste the **Email value** you copied earlier.
   d. Specify the role as **Defender for Cloud Admin Viewer** and then select **Save**.

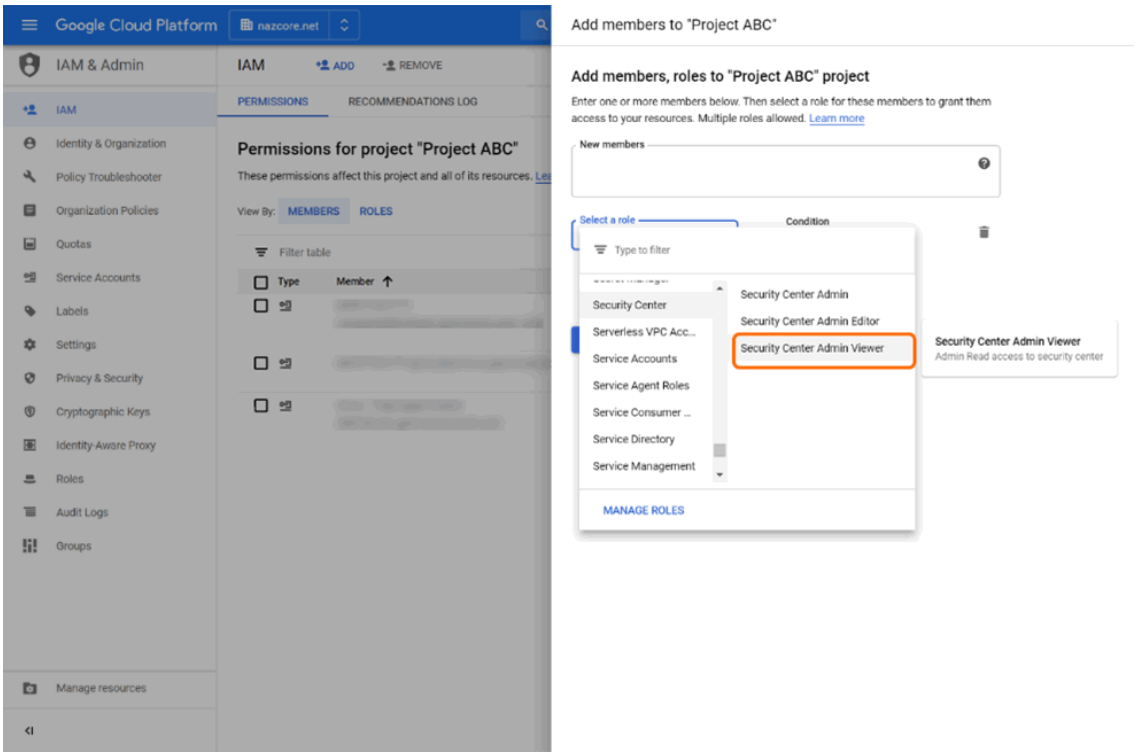

Step 4. Create a private key for the dedicated service account

1. Switch to project level.
2. In the **Navigation menu**, Under **IAM & admin** options, select **Service accounts**.
3. Open the dedicated service account and select Edit.
4. In the **Keys** section, select **ADD KEY** and then **Create new key**.
5. In the Create private key screen, select **JSON**, and then select **CREATE**.
6. Save this JSON file for later use.

Step 5. Connect GCP to Defender for Cloud

1. From Defender for Cloud's menu, open **Environment settings** and select the option to switch back to the classic connectors experience.

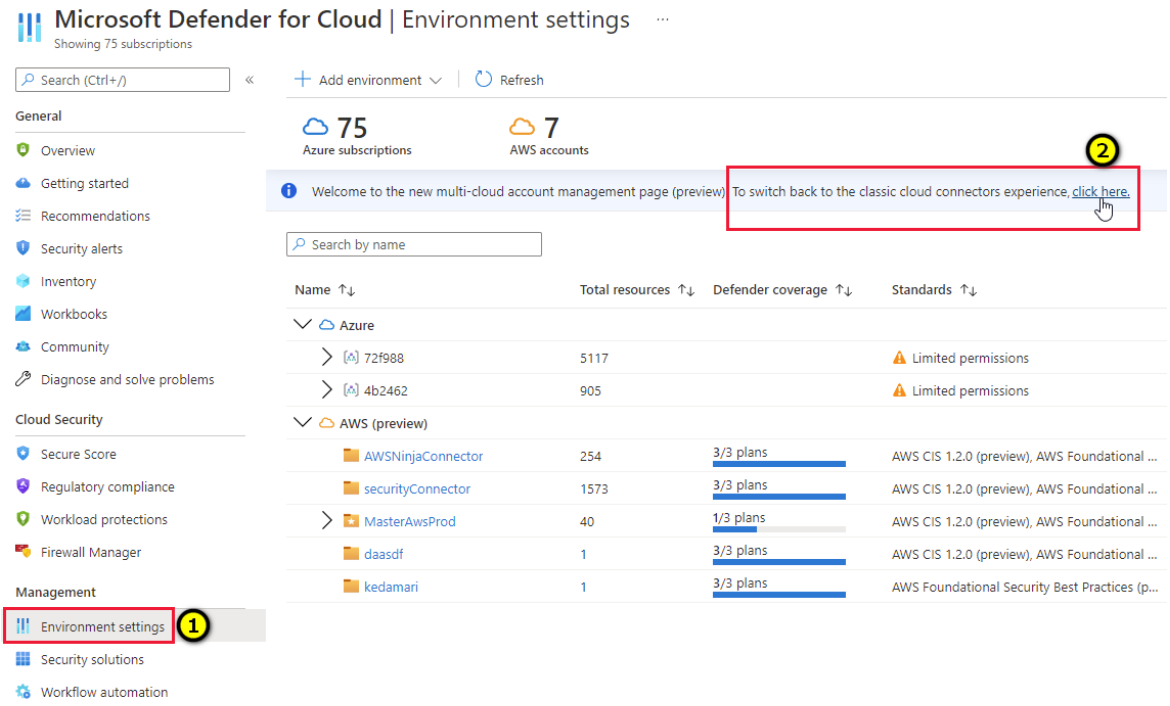

2. Select add GCP account.

3. In the onboarding page, do the following and then select **Next**.

    a. Validate the chosen subscription.
    b. In the **Display name** field, enter a display name for the connector.
    c. In the **Organization ID** field, enter your organization's ID. If you don't know it, see Creating and managing organizations.
    d. In the **Private key** file box, browse to the JSON file you downloaded in Step 4. Create a private key for the dedicated service account.

Step 6. Confirmation

When the connector is successfully created and GCP Security Command Center has been configured properly:

- The GCP CIS standard will be shown in the Defender for Cloud's regulatory compliance dashboard.
- Security recommendations for your GCP resources will appear in the Defender for Cloud portal and the regulatory compliance dashboard 5-10 minutes after onboard completes:

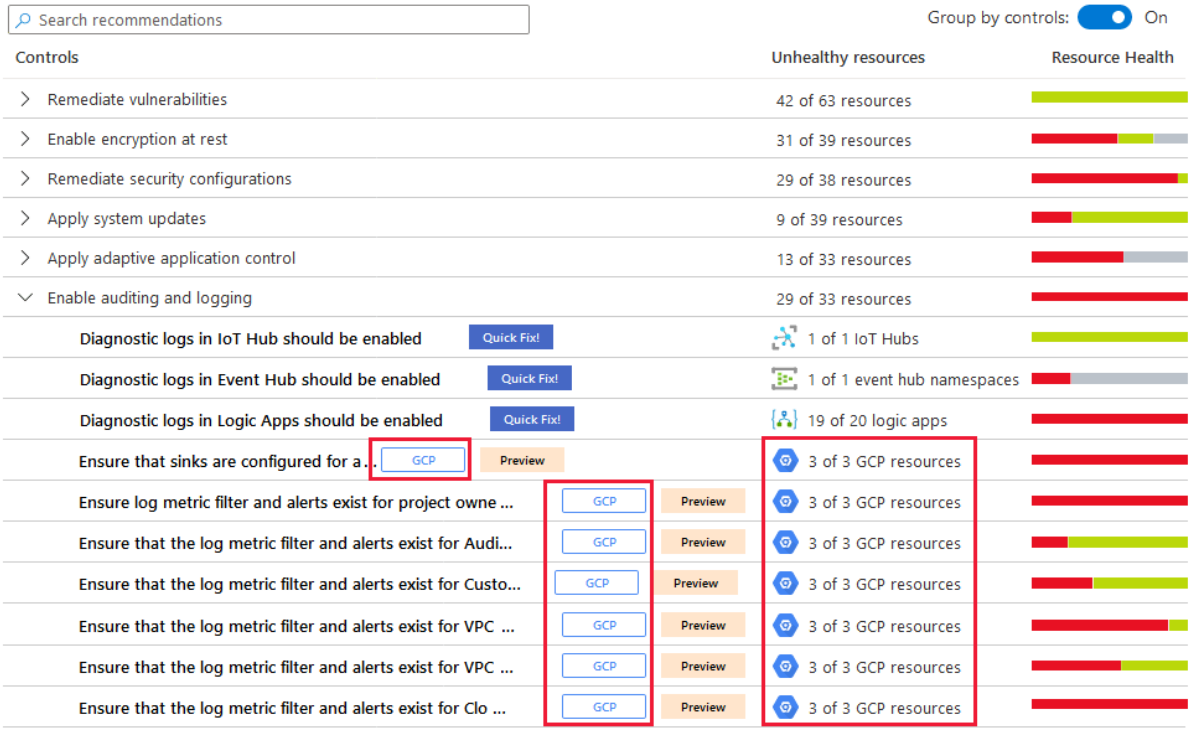

Monitoring your GCP resources

As shown above, Microsoft Defender for Cloud's security recommendations page displays your GCP resources together with your Azure and AWS resources for a true multi-cloud view.

To view all the active recommendations for your resources by resource type, use Defender for Cloud's asset inventory page and filter to the GCP resource type in which you're interested:

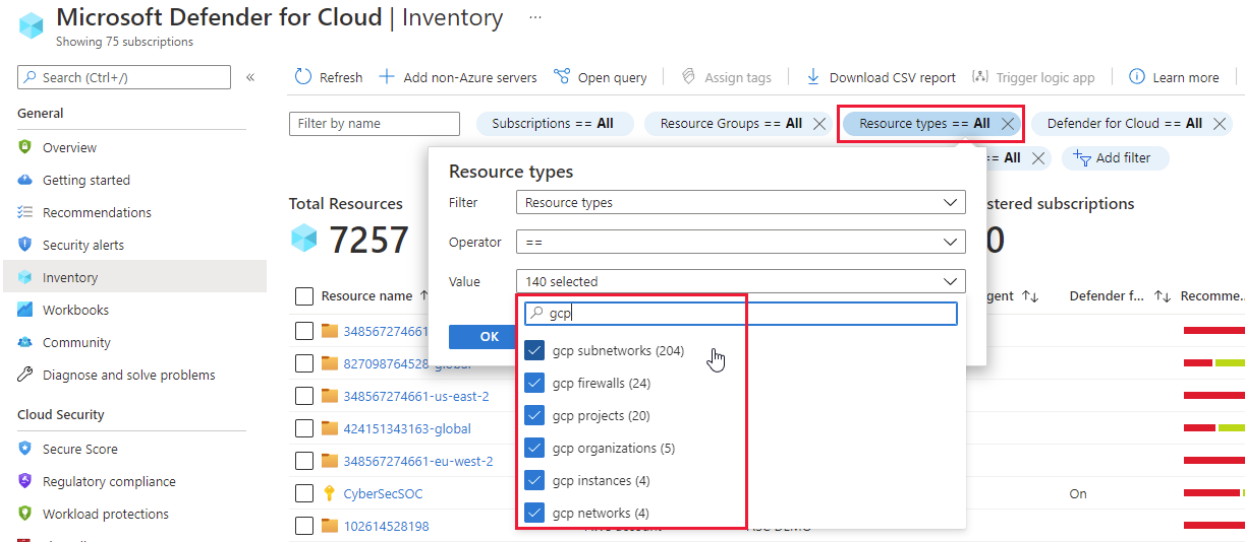

FAQ - Connecting GCP accounts to Microsoft Defender for Cloud

Can I connect multiple GCP organizations to Defender for Cloud?

Yes. Defender for Cloud's GCP connector connects your Google Cloud resources at the *organization* level.

Create a connector for every GCP organization you want to monitor from Defender for Cloud. When you connect an organization, all projects within that organization are added to Defender for Cloud.

Learn about the Google Cloud resource hierarchy in Google's online docs.

Is there an API for connecting my GCP resources to Defender for Cloud?

Yes. To create, edit, or delete Defender for Cloud cloud connectors with a REST API, see the details of the Connectors API.

Next steps

Connecting your GCP account is part of the multi-cloud experience available in Microsoft Defender for Cloud. For related information, see the following page:

- Connect your AWS accounts to Microsoft Defender for Cloud
- Google Cloud resource hierarchy--Learn about the Google Cloud resource hierarchy in Google's online docs

Configure auto provisioning for agents and extensions from Microsoft Defender for Cloud

2/15/2022 • 17 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud collects data from your resources using the relevant agent or extensions for that resource and the type of data collection you've enabled. Use the procedures below to ensure your resources have the necessary agents and extensions used by Defender for Cloud.

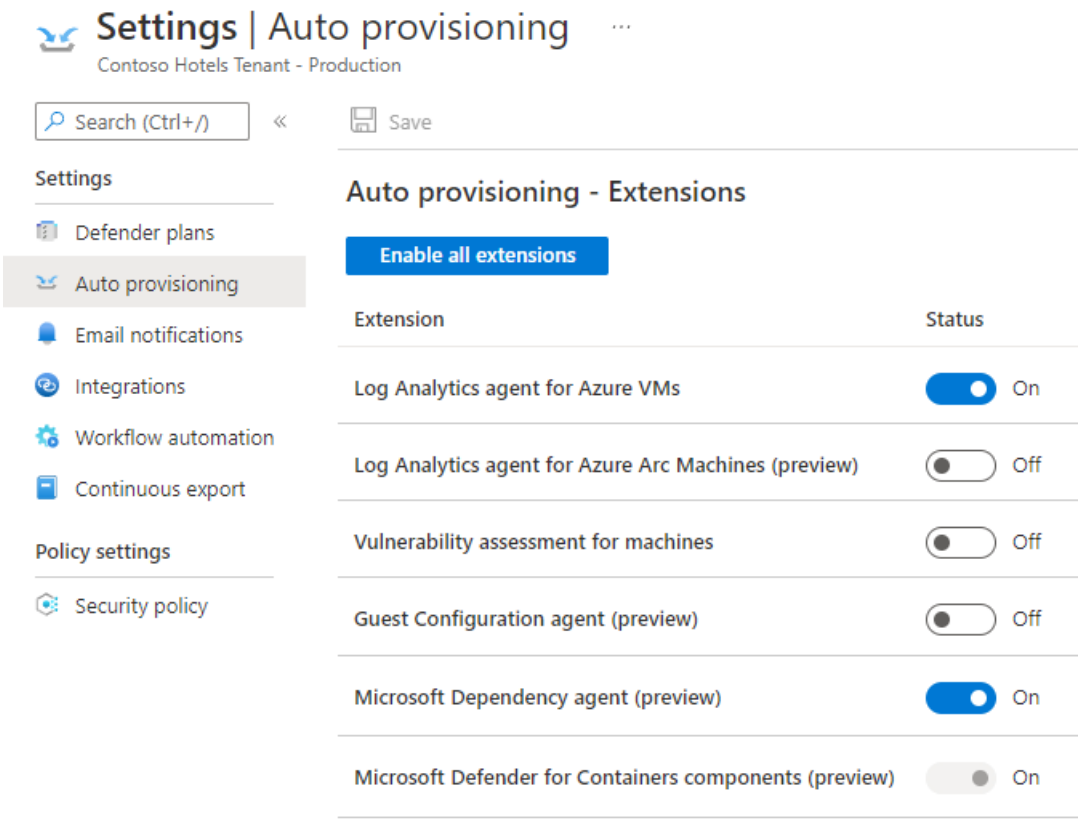

NOTE

When you enable auto provisioning of any of the supported extensions, you'll potentially impact *existing* and *future* machines. But when you **disable** auto provisioning for an extension, you'll only affect the *future* machines: nothing is uninstalled by disabling auto provisioning.

Prerequisites

To get started with Defender for Cloud, you must have a subscription to Microsoft Azure. If you don't have a subscription, you can sign up for a free account.

Availability

- Auto provisioning
- Log Analytics agent
- Vulnerability assessment
- Defender for Endpoint
- Guest Configuration
- Defender for Containers

This table shows the availability details for the auto provisioning **feature** itself.

ASPECT	DETAILS
Release state:	Auto provisioning is generally available (GA)
Pricing:	Auto provisioning is free to use
Required roles and permissions:	Depends on the specific extension - see relevant tab
Supported destinations:	Depends on the specific extension - see relevant tab
Clouds:	✅ Commercial clouds ✅ Azure Government, Azure China 21Vianet

TIP

For items marked in preview: The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

How does Defender for Cloud collect data?

Defender for Cloud collects data from your Azure virtual machines (VMs), virtual machine scale sets, IaaS containers, and non-Azure (including on-premises) machines to monitor for security vulnerabilities and threats.

Data collection is required to provide visibility into missing updates, misconfigured OS security settings, endpoint protection status, and health and threat protection. Data collection is only needed for compute resources such as VMs, virtual machine scale sets, IaaS containers, and non-Azure computers.

You can benefit from Microsoft Defender for Cloud even if you don't provision agents. However, you'll have limited security and the capabilities listed above aren't supported.

Data is collected using:

- The **Log Analytics agent**, which reads various security-related configurations and event logs from the machine and copies the data to your workspace for analysis. Examples of such data are: operating system type and version, operating system logs (Windows event logs), running processes, machine name, IP addresses, and logged in user.
- **Security extensions**, such as the Azure Policy Add-on for Kubernetes, which can also provide data to Defender for Cloud regarding specialized resource types.

TIP

As Defender for Cloud has grown, the types of resources that can be monitored has also grown. The number of extensions has also grown. Auto provisioning has expanded to support additional resource types by leveraging the capabilities of Azure Policy.

Why use auto provisioning?

Any of the agents and extensions described on this page *can* be installed manually (see Manual installation of the Log Analytics agent). However, **auto provisioning** reduces management overhead by installing all required agents and extensions on existing - and new - machines to ensure faster security coverage for all supported resources.

We recommend enabling auto provisioning, but it's disabled by default.

How does auto provisioning work?

Defender for Cloud's auto provisioning settings have a toggle for each type of supported extension. When you enable auto provisioning of an extension, you assign the appropriate **Deploy if not exists** policy. This policy type ensures the extension is provisioned on all existing and future resources of that type.

TIP

Learn more about Azure Policy effects including deploy if not exists in Understand Azure Policy effects.

Enable auto provisioning of the Log Analytics agent and extensions

When automatic provisioning is on for the Log Analytics agent, Defender for Cloud deploys the agent on all supported Azure VMs and any new ones created. For the list of supported platforms, see Supported platforms in Microsoft Defender for Cloud.

To enable auto provisioning of the Log Analytics agent:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.
3. In the **Auto provisioning** page, set the status of auto provisioning for the Log Analytics agent to **On**.

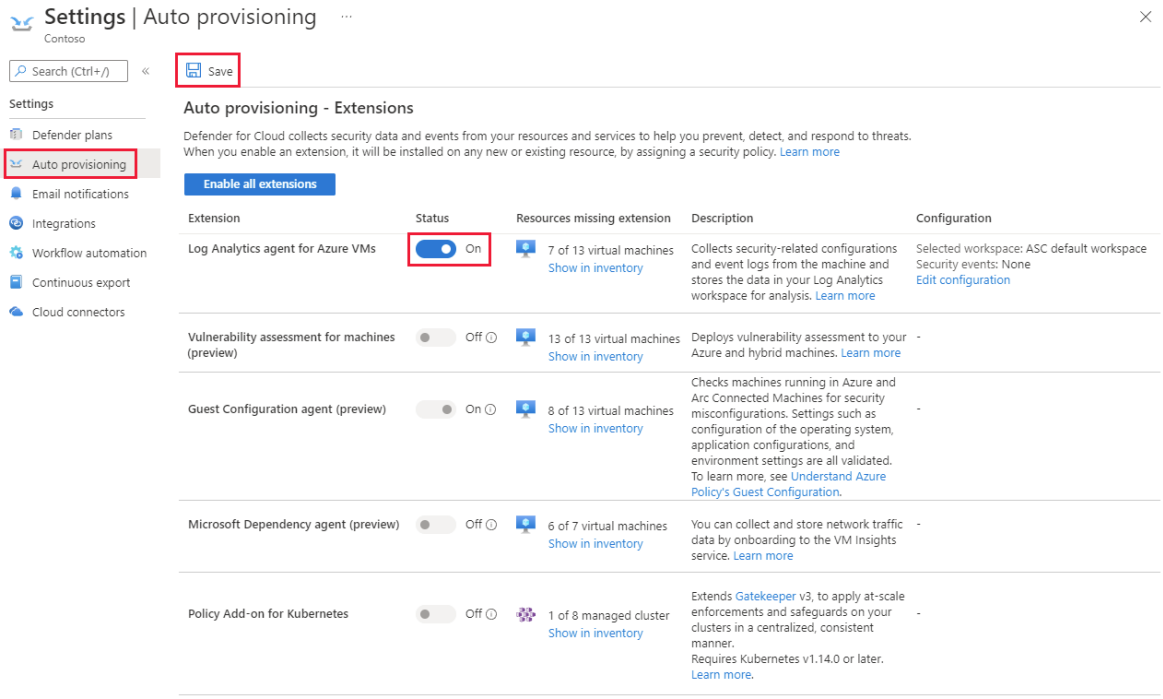

4. From the configuration options pane, define the workspace to use.

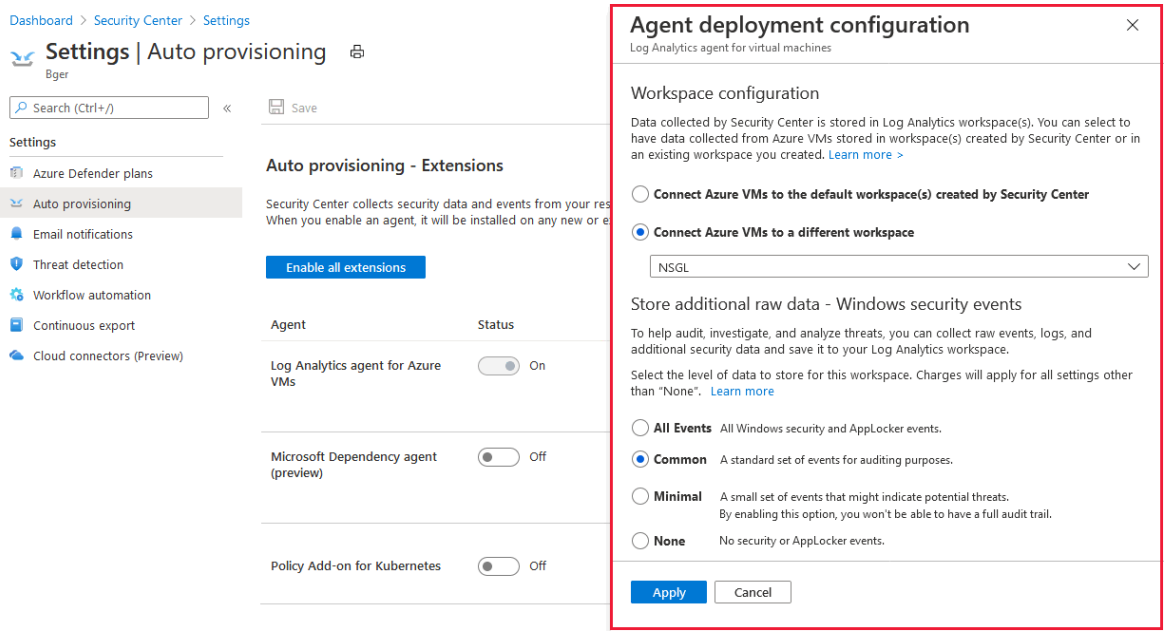

- **Connect Azure VMs to the default workspace(s) created by Defender for Cloud** - Defender for Cloud creates a new resource group and default workspace in the same geolocation, and connects the agent to that workspace. If a subscription contains VMs from multiple geolocations, Defender for Cloud creates multiple workspaces to ensure compliance with data privacy requirements.

  The naming convention for the workspace and resource group is:

  - Workspace: DefaultWorkspace-[subscription-ID]-[geo]
  - Resource Group: DefaultResourceGroup-[geo]

  A Defender for Cloud solution is automatically enabled on the workspace per the pricing tier set for the subscription.

> **TIP**
>
> For questions regarding default workspaces, see:
>
> - Am I billed for Azure Monitor logs on the workspaces created by Defender for Cloud?
> - Where is the default Log Analytics workspace created?
> - Can I delete the default workspaces created by Defender for Cloud?

- **Connect Azure VMs to a different workspace** - From the dropdown list, select the workspace to store collected data. The dropdown list includes all workspaces across all of your subscriptions. You can use this option to collect data from virtual machines running in different subscriptions and store it all in your selected workspace.

  If you already have an existing Log Analytics workspace, you might want to use the same workspace (requires read and write permissions on the workspace). This option is useful if you're using a centralized workspace in your organization and want to use it for security data collection. Learn more in Manage access to log data and workspaces in Azure Monitor.

  If your selected workspace already has a "Security" or "SecurityCenterFree" solution enabled, the pricing will be set automatically. If not, install a Defender for Cloud solution on the workspace:

  a. From Defender for Cloud's menu, open **Environment settings**.
  b. Select the workspace to which you'll be connecting the agents.
  c. Select **Enhanced security off** or **Enable all Microsoft Defender plans**.

5. From the **Windows security events** configuration, select the amount of raw event data to store:

   - **None** – Disable security event storage. This is the default setting.
   - **Minimal** – A small set of events for when you want to minimize the event volume.
   - **Common** – A set of events that satisfies most customers and provides a full audit trail.
   - **All events** – For customers who want to make sure all events are stored.

   > **TIP**
   >
   > To set these options at the workspace level, see Setting the security event option at the workspace level.
   >
   > For more information of these options, see Windows security event options for the Log Analytics agent.

6. Select **Apply** in the configuration pane.

7. To enable automatic provisioning of an extension other than the Log Analytics agent:

   a. If you're enabling auto provisioning for the Microsoft Dependency agent, ensure the Log Analytics agent is set to auto deploy.

   b. Toggle the status to **On** for the relevant extension.

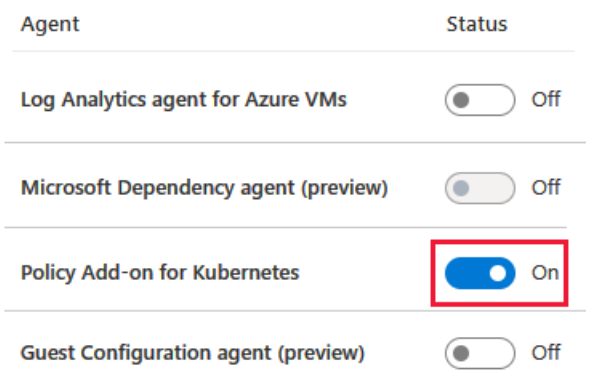

c. Select **Save**. The Azure Policy definition is assigned and a remediation task is created.

EXTENSION	POLICY
Policy Add-on for Kubernetes	Deploy Azure Policy Add-on to Azure Kubernetes Service clusters
Microsoft Dependency agent (preview) (Windows VMs)	Deploy Dependency agent for Windows virtual machines
Microsoft Dependency agent (preview) (Linux VMs)	Deploy Dependency agent for Linux virtual machines
Guest Configuration agent (preview)	Deploy prerequisites to enable Guest Configuration policies on virtual machines

8. Select **Save**. If a workspace needs to be provisioned, agent installation might take up to 25 minutes.

9. You'll be asked if you want to reconfigure monitored VMs that were previously connected to a default workspace:

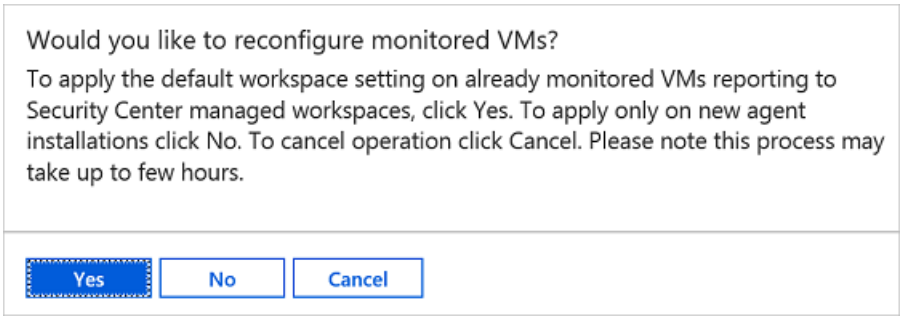

- **No** - your new workspace settings will only be applied to newly discovered VMs that don't have the Log Analytics agent installed.
- **Yes** - your new workspace settings will apply to all VMs and every VM currently connected to a Defender for Cloud created workspace will be reconnected to the new target workspace.

> **NOTE**
> If you select **Yes**, don't delete the workspace(s) created by Defender for Cloud until all VMs have been reconnected to the new target workspace. This operation fails if a workspace is deleted too early.

Windows security event options for the Log Analytics agent

Selecting a data collection tier in Microsoft Defender for Cloud only affects the *storage* of security events in your

Log Analytics workspace. The Log Analytics agent will still collect and analyze the security events required for Defender for Cloud's threat protection, regardless of the level of security events you choose to store in your workspace. Choosing to store security events enables investigation, search, and auditing of those events in your workspace.

Requirements

The enhanced security protections of Defender for Cloud are required for storing Windows security event data. Learn more about the enhanced protection plans.

Storing data in Log Analytics might incur additional charges for data storage. For more information, see the pricing page.

Information for Microsoft Sentinel users

Users of Microsoft Sentinel: note that security events collection within the context of a single workspace can be configured from either Microsoft Defender for Cloud or Microsoft Sentinel, but not both. If you're planning to add Microsoft Sentinel to a workspace that is already getting alerts from Microsoft Defender for Cloud, and is set to collect Security Events, you have two options:

- Leave the Security Events collection in Microsoft Defender for Cloud as is. You will be able to query and analyze these events in Microsoft Sentinel as well as in Defender for Cloud. You will not, however, be able to monitor the connector's connectivity status or change its configuration in Microsoft Sentinel. If this is important to you, consider the second option.
- Disable Security Events collection in Microsoft Defender for Cloud (by setting **Windows security events** to **None** in the configuration of your Log Analytics agent). Then add the Security Events connector in Microsoft Sentinel. As with the first option, you will be able to query and analyze events in both Microsoft Sentinel and Defender for Cloud, but you will now be able to monitor the connector's connectivity status or change its configuration in - and only in - Microsoft Sentinel.

What event types are stored for "Common" and "Minimal"?

These sets were designed to address typical scenarios. Make sure to evaluate which one fits your needs before implementing it.

To determine the events for the **Common** and **Minimal** options, we worked with customers and industry standards to learn about the unfiltered frequency of each event and their usage. We used the following guidelines in this process:

- **Minimal** - Make sure that this set covers only events that might indicate a successful breach and important events that have a very low volume. For example, this set contains user successful and failed login (event IDs 4624, 4625), but it doesn't contain sign out which is important for auditing but not meaningful for detection and has relatively high volume. Most of the data volume of this set is the login events and process creation event (event ID 4688).
- **Common** - Provide a full user audit trail in this set. For example, this set contains both user logins and user sign outs (event ID 4634). We include auditing actions like security group changes, key domain controller Kerberos operations, and other events that are recommended by industry organizations.

Events that have very low volume were included in the common set as the main motivation to choose it over all the events is to reduce the volume and not to filter out specific events.

Here is a complete breakdown of the Security and App Locker event IDs for each set:

DATA TIER	COLLECTED EVENT INDICATORS
Minimal	1102,4624,4625,4657,4663,4688,4700,4702,4719,4720,4722,4723,4724,4727,4728,4732,4735,4737,4739,4740,4754,4755,

DATA TIER	COLLECTED EVENT INDICATORS
	4756,4767,4799,4825,4946,4948,4956,5024,5033,8001,8002,8003,8004,8005,8006,8007,8222
Common	1,299,300,324,340,403,404,410,411,412,413,431,500,501,1100,1102,1107,1108,4608,4610,4611,4614,4622,
	4624,4625,4634,4647,4648,4649,4657,4661,4662,4663,4665,4666,4667,4688,4670,4672,4673,4674,4675,4689,4697,
	4700,4702,4704,4705,4716,4717,4718,4719,4720,4722,4723,4724,4725,4726,4727,4728,4729,4733,4732,4735,4737,
	4738,4739,4740,4742,4744,4745,4746,4750,4751,4752,4754,4755,4756,4757,4760,4761,4762,4764,4767,4768,4771,
	4774,4778,4779,4781,4793,4797,4798,4799,4800,4801,4802,4803,4825,4826,4870,4886,4887,4888,4893,4898,4902,
	4904,4905,4907,4931,4932,4933,4946,4948,4956,4985,5024,5033,5059,5136,5137,5140,5145,5632,6144,6145,6272,
	6273,6278,6416,6423,6424,8001,8002,8003,8004,8005,8006,8007,8222,26401,30004

NOTE

- If you are using Group Policy Object (GPO), it is recommended that you enable audit policies Process Creation Event 4688 and the *CommandLine* field inside event 4688. For more information about Process Creation Event 4688, see Defender for Cloud's FAQ. For more information about these audit policies, see Audit Policy Recommendations.
- To enable data collection for Adaptive Application Controls, Defender for Cloud configures a local AppLocker policy in Audit mode to allow all applications. This will cause AppLocker to generate events which are then collected and leveraged by Defender for Cloud. It is important to note that this policy will not be configured on any machines on which there is already a configured AppLocker policy.
- To collect Windows Filtering Platform Event ID 5156, you need to enable Audit Filtering Platform Connection (Auditpol /set /subcategory:"Filtering Platform Connection" /Success:Enable)

Setting the security event option at the workspace level

You can define the level of security event data to store at the workspace level.

1. From Defender for Cloud's menu in the Azure portal, select **Environment settings**.
2. Select the relevant workspace. The only data collection events for a workspace are the Windows security events described on this page.

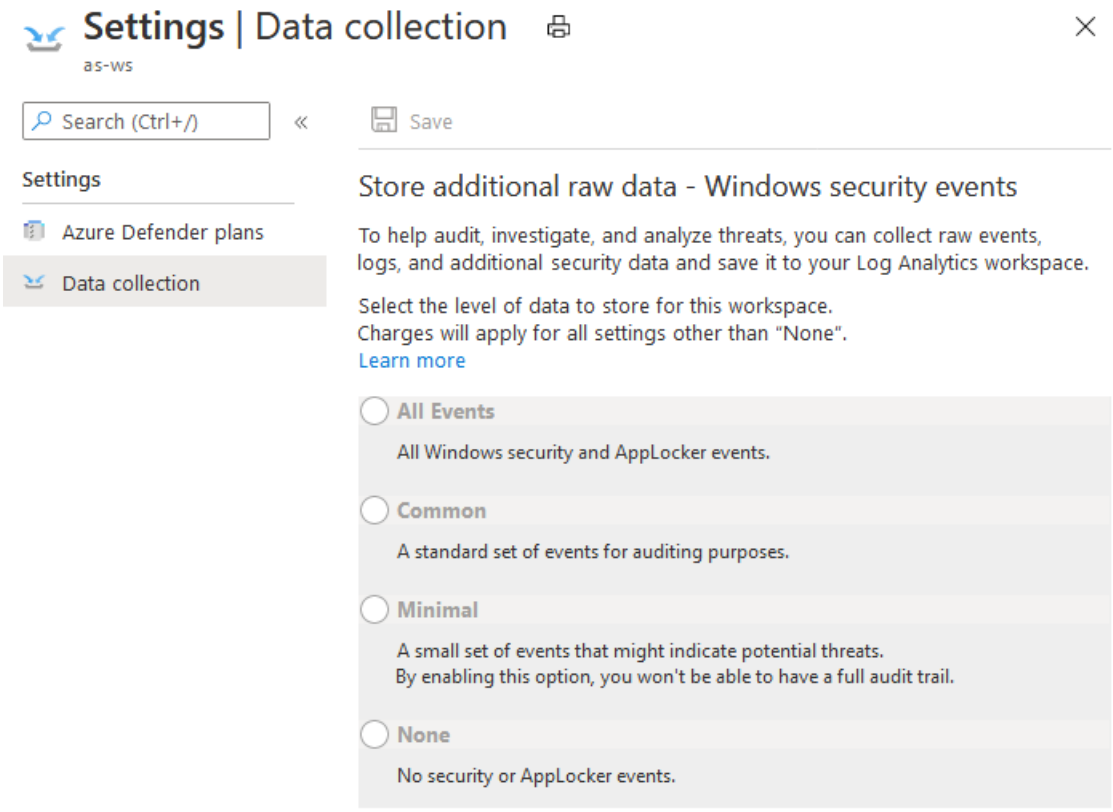

3. Select the amount of raw event data to store and select **Save**.

Manual agent provisioning

To manually install the Log Analytics agent:

1. Disable auto provisioning.
2. Optionally, create a workspace.
3. Enable Microsoft Defender for Cloud on the workspace on which you're installing the Log Analytics agent:
   a. From Defender for Cloud's menu, open **Environment settings**.
   b. Set the workspace on which you're installing the agent. Make sure the workspace is in the same subscription you use in Defender for Cloud and that you have read/write permissions for the workspace.
   c. Select **Microsoft Defender for Cloud on**, and **Save**.

   > **NOTE**
   > If the workspace already has a **Security** or **SecurityCenterFree** solution enabled, the pricing will be set automatically.

4. To deploy agents on new VMs using a Resource Manager template, install the Log Analytics agent:
   - Install the Log Analytics agent for Windows
   - Install the Log Analytics agent for Linux
5. To deploy agents on your existing VMs, follow the instructions in Collect data about Azure Virtual Machines (the section **Collect event and performance data** is optional).
6. To use PowerShell to deploy the agents, use the instructions from the virtual machines documentation:

- For Windows machines
- For Linux machines

> **TIP**
> For more information about onboarding, see Automate onboarding of Microsoft Defender for Cloud using PowerShell.

Automatic provisioning in cases of a pre-existing agent installation

The following use cases specify how automatic provision works in cases when there is already an agent or extension installed.

- **Log Analytics agent is installed on the machine, but not as an extension (Direct agent)** - If the Log Analytics agent is installed directly on the VM (not as an Azure extension), Defender for Cloud will install the Log Analytics agent extension, and might upgrade the Log Analytics agent to the latest version. The agent installed will continue to report to its already configured workspace(s), and additionally will report to the workspace configured in Defender for Cloud (Multi-homing is supported on Windows machines).

  If the configured workspace is a user workspace (not Defender for Cloud's default workspace), then you will need to install the "Security" or "SecurityCenterFree" solution on it for Defender for Cloud to start processing events from VMs and computers reporting to that workspace.

  For Linux machines, Agent multi-homing is not yet supported - hence, if an existing agent installation is detected, automatic provisioning will not occur and the machine's configuration will not be altered.

  For existing machines on subscriptions onboarded to Defender for Cloud before 17 March 2019, when an existing agent will be detected, the Log Analytics agent extension will not be installed and the machine will not be affected. For these machines, see to the "Resolve monitoring agent health issues on your machines" recommendation to resolve the agent installation issues on these machines.

- **System Center Operations Manager agent is installed on the machine** - Defender for Cloud will install the Log Analytics agent extension side by side to the existing Operations Manager. The existing Operations Manager agent will continue to report to the Operations Manager server normally. The Operations Manager agent and Log Analytics agent share common run-time libraries, which will be updated to the latest version during this process. If Operations Manager agent version 2012 is installed, **do not** enable automatic provisioning.

- **A pre-existing VM extension is present**:

  - When the Monitoring Agent is installed as an extension, the extension configuration allows reporting to only a single workspace. Defender for Cloud does not override existing connections to user workspaces. Defender for Cloud will store security data from the VM in the workspace already connected, provided that the "Security" or "SecurityCenterFree" solution has been installed on it. Defender for Cloud may upgrade the extension version to the latest version in this process.
  - To see to which workspace the existing extension is sending data to, run the test to Validate connectivity with Microsoft Defender for Cloud. Alternatively, you can open Log Analytics workspaces, select a workspace, select the VM, and look at the Log Analytics agent connection.
  - If you have an environment where the Log Analytics agent is installed on client workstations and reporting to an existing Log Analytics workspace, review the list of operating systems supported by Microsoft Defender for Cloud to make sure your operating system is supported. For more information, see Existing log analytics customers.

Disable auto provisioning

When you disable auto provisioning, agents will not be provisioned on new VMs.

To turn off automatic provisioning of an agent:

1. From Defender for Cloud's menu in the portal, select **Environment settings**.
2. Select the relevant subscription.
3. Select **Auto provisioning**.
4. Toggle the status to **Off** for the relevant agent.

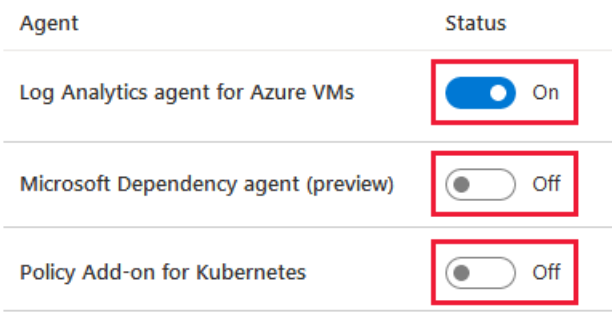

5. Select **Save**. When auto provisioning is disabled, the default workspace configuration section is not displayed:

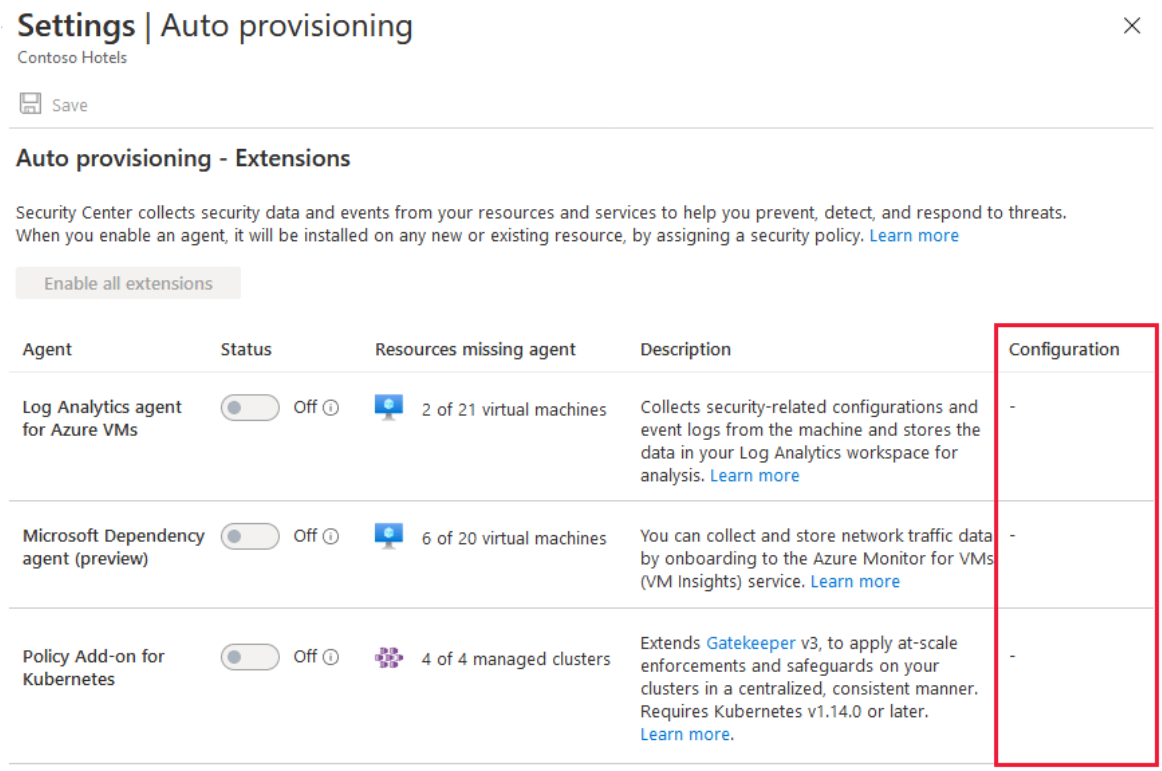

NOTE

Disabling automatic provisioning does not remove the Log Analytics agent from Azure VMs where the agent was provisioned. For information on removing the OMS extension, see How do I remove OMS extensions installed by Defender for Cloud.

Troubleshooting

- To identify monitoring agent network requirements, see Troubleshooting monitoring agent network requirements.

- To identify manual onboarding issues, see How to troubleshoot Operations Management Suite onboarding issues.

Next steps

This page explained how to enable auto provisioning for the Log Analytics agent and other Defender for Cloud extensions. It also described how to define a Log Analytics workspace in which to store the collected data. Both operations are required to enable data collection. Storing data in Log Analytics, whether you use a new or existing workspace, might incur more charges for data storage. For pricing details in your local currency or region, see the pricing page.

Configure email notifications for security alerts

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Security alerts need to reach the right people in your organization. By default, Microsoft Defender for Cloud emails subscription owners whenever a high-severity alert is triggered for their subscription. This page explains how to customize these notifications.

Use Defender for Cloud's **Email notifications** settings page to define preferences for notification emails including:

- ***who* should be notified** - Emails can be sent to select individuals or to anyone with a specified Azure role for a subscription.
- ***what* they should be notified about** - Modify the severity levels for which Defender for Cloud should send out notifications.

To avoid alert fatigue, Defender for Cloud limits the volume of outgoing mails. For each subscription, Defender for Cloud sends:

- a maximum of one email per **6 hours** (4 emails per day) for **high-severity** alerts
- a maximum of one email per **12 hours** (2 emails per day) for **medium-severity** alerts
- a maximum of one email per **24 hours** for **low-severity** alerts

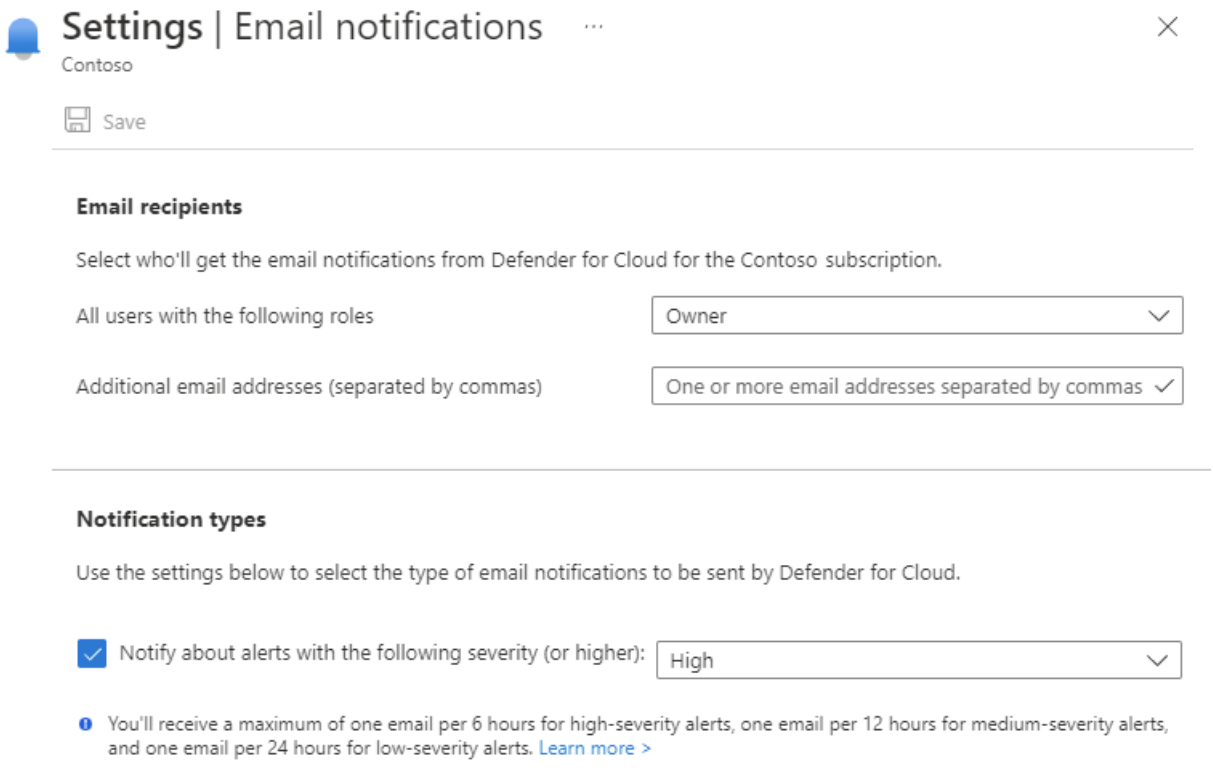

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free
Required roles and permissions:	**Security Admin** **Subscription Owner**
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

Customize the security alerts email notifications via the portal

You can send email notifications to individuals or to all users with specific Azure roles.

1. From Defender for Cloud's **Environment settings** area, select the relevant subscription, and open **Email notifications**.
2. Define the recipients for your notifications with one or both of these options:
   - From the dropdown list, select from the available roles.
   - Enter specific email addresses separated by commas. There's no limit to the number of email addresses that you can enter.
3. To apply the security contact information to your subscription, select **Save**.

Customize the alerts email notifications through the API

You can also manage your email notifications through the supplied REST API. For full details see the SecurityContacts API documentation.

This is an example request body for the PUT request when creating a security contact configuration:

URI:

```
https://management.azure.com/subscriptions/<SubscriptionId>/providers/Microsoft.Security/securityContacts/default?api-version=2020-01-01-preview
```

```
{
    "properties": {
        "emails": admin@contoso.com;admin2@contoso.com,
        "notificationsByRole": {
            "state": "On",
            "roles": ["AccountAdmin", "Owner"]
        },
        "alertNotifications": {
            "state": "On",
            "minimalSeverity": "Medium"
        },
        "phone": ""
    }
}
```

See also

To learn more about security alerts, see the following pages:

- Security alerts - a reference guide--Learn about the security alerts you might see in Microsoft Defender for Cloud's Threat Protection module
- Manage and respond to security alerts in Microsoft Defender for Cloud--Learn how to manage and respond to security alerts
- Workflow automation--Automate responses to alerts with custom notification logic

Quickstart: Create an automatic response to a specific security alert using an ARM template

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This quickstart describes how to use an Azure Resource Manager template (ARM template) to create a workflow automation that triggers a logic app when specific security alerts are received by Microsoft Defender for Cloud.

An ARM template is a JavaScript Object Notation (JSON) file that defines the infrastructure and configuration for your project. The template uses declarative syntax. In declarative syntax, you describe your intended deployment without writing the sequence of programming commands to create the deployment.

If your environment meets the prerequisites and you're familiar with using ARM templates, select the **Deploy to Azure** button. The template will open in the Azure portal.

Prerequisites

If you don't have an Azure subscription, create a free account before you begin.

For a list of the roles and permissions required to work with Microsoft Defender for Cloud's workflow automation feature, see workflow automation.

Review the template

The template used in this quickstart is from Azure Quickstart Templates.

```
{
  "$schema": "https://schema.management.azure.com/schemas/2019-04-01/deploymentTemplate.json#",
  "contentVersion": "1.0.0.0",
  "parameters": {
    "automationName": {
      "type": "string",
      "minLength": 3,
      "maxLength": 24
    },
    "location": {
      "type": "string",
      "defaultValue": "[resourceGroup().location]",
      "metadata": {
        "description": "Location for the automation"
      }
    },
    "logicAppName": {
      "type": "string",
      "minLength": 3
    },
    "logicAppResourceGroupName": {
```

```
      "type": "string",
      "minLength": 3
    },
    "subscriptionId": {
      "type": "string",
      "defaultValue": "[subscription().subscriptionId]",
      "metadata": {
        "description": "The Azure resource GUID id of the subscription"
      }
    },
    "alertSettings": {
      "type": "object",
      "metadata": {
        "description": "The alert settings object used for deploying the automation"
      }
    }
  },
  "variables": {
    "automationDescription": "automation description for subscription {0}",
    "scopeDescription": "automation scope for subscription {0}"
  },
  "resources": [
    {
      "type": "Microsoft.Security/automations",
      "apiVersion": "2019-01-01-preview",
      "name": "[parameters('automationName')]",
      "location": "[parameters('location')]",
      "properties": {
        "description": "[format(variables('automationDescription'),'{0}', parameters('subscriptionId'))]",
        "isEnabled": true,
        "actions": [
          {
            "actionType": "LogicApp",
            "logicAppResourceId": "[resourceId('Microsoft.Logic/workflows', parameters('logicAppName'))]",
            "uri": "[listCallbackURL(resourceId(parameters('subscriptionId'),
parameters('logicAppResourceGroupName'), 'Microsoft.Logic/workflows/triggers', parameters('logicAppName'),
'manual'), '2019-05-01').value]"
          }
        ],
        "scopes": [
          {
            "description": "[format(variables('scopeDescription'),'{0}', parameters('subscriptionId'))]",
            "scopePath": "[subscription().id]"
          }
        ],
        "sources": [
          {
            "eventSource": "Alerts",
            "eventVersionType": "Api",
            "copy": [
              {
                "name": "ruleSets",
                "count": "[length(parameters('alertSettings').alertSeverityMapping)]",
                "input": {
                  "rules": [
                    {
                      "propertyJPath": "
[parameters('alertSettings').alertSeverityMapping[copyIndex('ruleSets')].jpath]",
                      "propertyType": "string",
                      "expectedValue": "
[parameters('alertSettings').alertSeverityMapping[copyIndex('ruleSets')].expectedValue]",
                      "operator": "
[parameters('alertSettings').alertSeverityMapping[copyIndex('ruleSets')].operator]"
                    },
                    {
                      "propertyJPath": "Severity",
                      "propertyType": "string",
                      "expectedValue": "
[parameters('alertSettings').alertSeverityMapping[copyIndex('ruleSets')].severity]"
```

```
                    "operator": "Equals"
                  }
                ]
              }
            }
          ]
        }
      ]
    }
  ]
}
```

Relevant resources

- **Microsoft.Security/automations**: The automation that will trigger the logic app, upon receiving a Microsoft Defender for Cloud alert that contains a specific string.
- **Microsoft.Logic/workflows**: An empty triggerable Logic App.

For other Defender for Cloud quickstart templates, see these community contributed templates.

Deploy the template

- **PowerShell**:

```
New-AzResourceGroup -Name <resource-group-name> -Location <resource-group-location> #use this command
when you need to create a new resource group for your deployment
New-AzResourceGroupDeployment -ResourceGroupName <resource-group-name> -TemplateUri
https://raw.githubusercontent.com/Azure/azure-quickstart-
templates/master/quickstarts/microsoft.security/securitycenter-create-automation-for-
alertnamecontains/azuredeploy.json
```

- **CLI**:

```
az group create --name <resource-group-name> --location <resource-group-location> #use this command
when you need to create a new resource group for your deployment
az deployment group create --resource-group <my-resource-group> --template-uri
https://raw.githubusercontent.com/Azure/azure-quickstart-
templates/master/quickstarts/microsoft.security/securitycenter-create-automation-for-
alertnamecontains/azuredeploy.json
```

- **Portal**:

  Deploy to Azure

  To find more information about this deployment option, see Use a deployment button to deploy templates from GitHub repository.

Review deployed resources

Use the Azure portal to check the workflow automation has been deployed.

1. From the Azure portal, open **Microsoft Defender for Cloud**.
2. From the top menu bar, select the filter icon, and select the specific subscription on which you deployed the new workflow automation.
3. From Microsoft Defender for Cloud's menu, open **workflow automation** and check for your new automation.

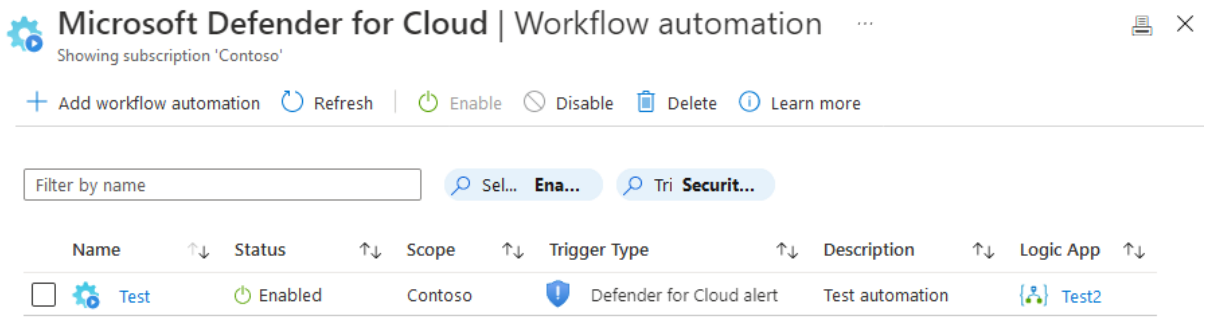

> **TIP**
> If you have many workflow automations on your subscription, use the **filter by name** option.

Clean up resources

When no longer needed, delete the workflow automation using the Azure portal.

1. From the Azure portal, open **Microsoft Defender for Cloud**.
2. From the top menu bar, select the filter icon, and select the specific subscription on which you deployed the new workflow automation.
3. From Microsoft Defender for Cloud's menu, open **workflow automation** and find the automation to be deleted.

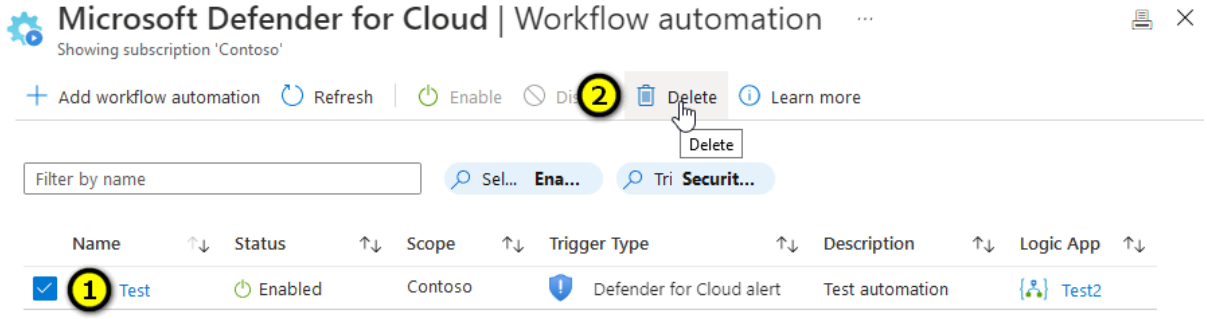

4. Select the checkbox for the item to be deleted.
5. From the toolbar, select **Delete**.

Next steps

For a step-by-step tutorial that guides you through the process of creating a template, see:

Tutorial: Create and deploy your first ARM template

Customize the set of standards in your regulatory compliance dashboard

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud continually compares the configuration of your resources with requirements in industry standards, regulations, and benchmarks. The **regulatory compliance dashboard** provides insights into your compliance posture based on how you're meeting specific compliance requirements.

> **TIP**
> Learn more about Defender for Cloud's regulatory compliance dashboard in the frequently asked questions.

How are regulatory compliance standards represented in Defender for Cloud?

Industry standards, regulatory standards, and benchmarks are represented in Defender for Cloud's regulatory compliance dashboard. Each standard is an initiative defined in Azure Policy.

To see compliance data mapped as assessments in your dashboard, add a compliance standard to your management group or subscription from within the **Security policy** page. To learn more about Azure Policy and initiatives, see Working with security policies.

When you've assigned a standard or benchmark to your selected scope, the standard appears in your regulatory compliance dashboard with all associated compliance data mapped as assessments. You can also download summary reports for any of the standards that have been assigned.

Microsoft tracks the regulatory standards themselves and automatically improves its coverage in some of the packages over time. When Microsoft releases new content for the initiative, it will appear automatically in your dashboard as new policies mapped to controls in the standard.

What regulatory compliance standards are available in Defender for Cloud?

By default, every subscription has the **Azure Security Benchmark** assigned. This is the Microsoft-authored, Azure-specific guidelines for security and compliance best practices based on common compliance frameworks. Learn more about Azure Security Benchmark.

Available regulatory standards:

- PCI-DSS v3.2.1:2018
- SOC TSP
- NIST SP 800-53 R4

- NIST SP 800 171 R2
- UK OFFICIAL and UK NHS
- Canada Federal PBMM
- Azure CIS 1.1.0
- HIPAA/HITRUST
- SWIFT CSP CSCF v2020
- ISO 27001:2013
- New Zealand ISM Restricted
- CMMC Level 3
- Azure CIS 1.3.0
- NIST SP 800-53 R5
- FedRAMP H
- FedRAMP M

> **TIP**
> Standards are added to the dashboard as they become available. The preceding list might not contain recently added standards.

Add a regulatory standard to your dashboard

The following steps explain how to add a package to monitor your compliance with one of the supported regulatory standards.

Prerequisites

To add standards to your dashboard:

- The subscription must have Defender for Cloud's enhanced security features enabled
- The user must have owner or policy contributor permissions

Add a standard

1. From Defender for Cloud's menu, select **Regulatory compliance** to open the regulatory compliance dashboard. Here you can see the compliance standards currently assigned to the currently selected subscriptions.
2. From the top of the page, select **Manage compliance policies**. The Policy Management page appears.
3. Select the subscription or management group for which you want to manage the regulatory compliance posture.

   > **TIP**
   > We recommend selecting the highest scope for which the standard is applicable so that compliance data is aggregated and tracked for all nested resources.

4. To add the standards relevant to your organization, expand the **Industry & regulatory standards** section and select **Add more standards**.
5. From the **Add regulatory compliance standards** page, you can search for any of the available standards:

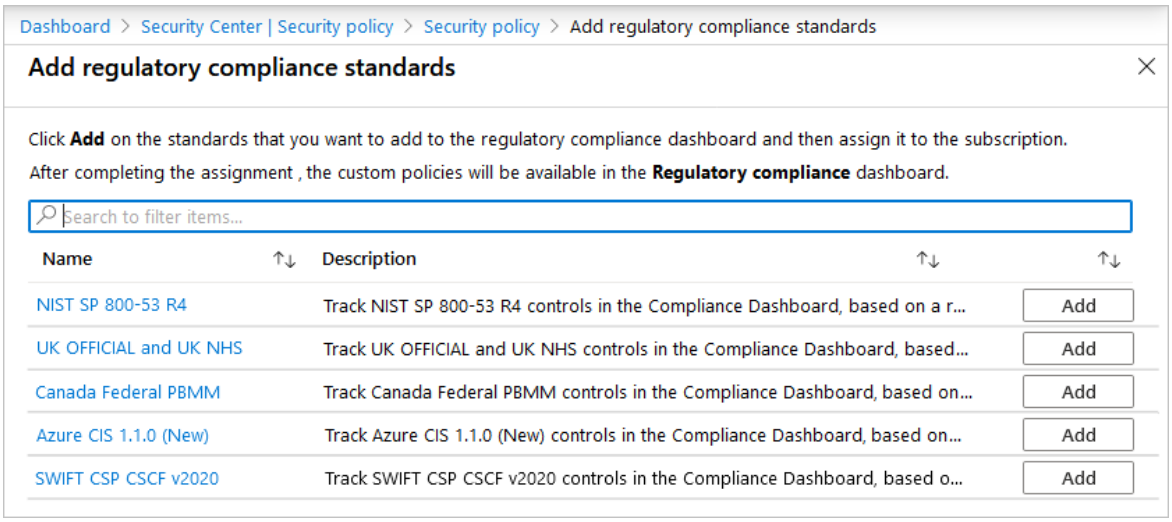

6. Select **Add** and enter all the necessary details for the specific initiative such as scope, parameters, and remediation.

7. From Defender for Cloud's menu, select **Regulatory compliance** again to go back to the regulatory compliance dashboard.

   Your new standard appears in your list of Industry & regulatory standards.

> **NOTE**
> It may take a few hours for a newly added standard to appear in the compliance dashboard.

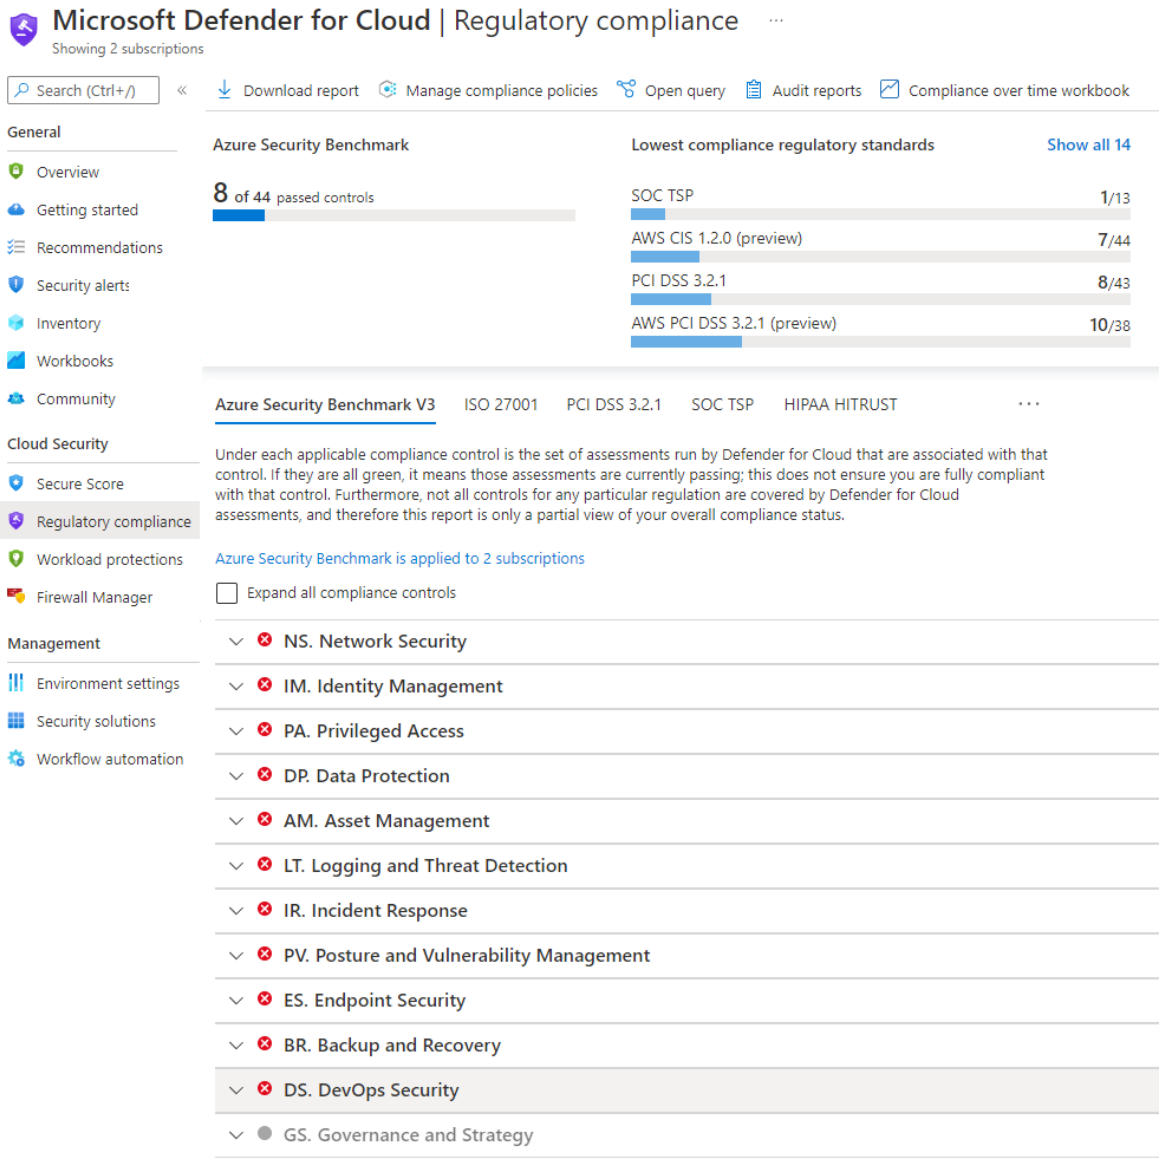

Remove a standard from your dashboard

You can continue to customize the regulatory compliance dashboard, to focus only on the standards that are applicable to you, by removing any of the supplied regulatory standards that aren't relevant to your organization.

To remove a standard:

1. From Defender for Cloud's menu, select **Security policy**.
2. Select the relevant subscription from which you want to remove a standard.

> **NOTE**
> You can remove a standard from a subscription, but not from a management group.

The security policy page opens. For the selected subscription, it shows the default policy, the industry and regulatory standards, and any custom initiatives you've created.

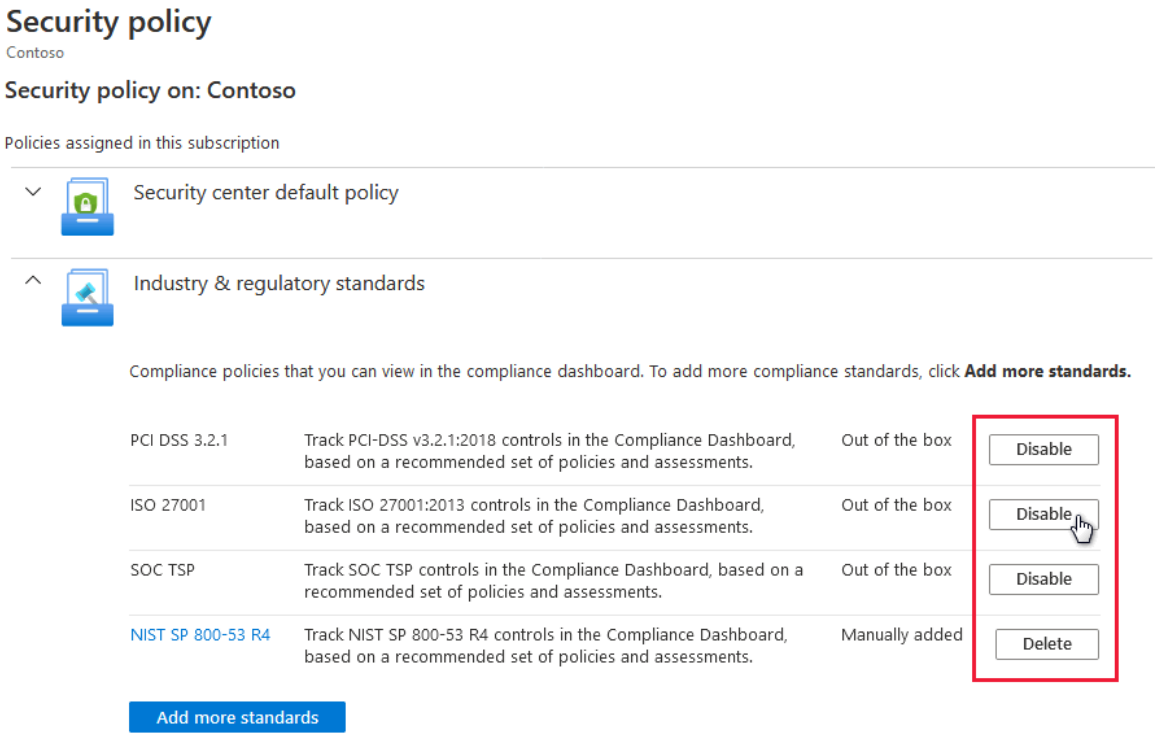

3. For the standard you want to remove, select **Disable**. A confirmation window appears.

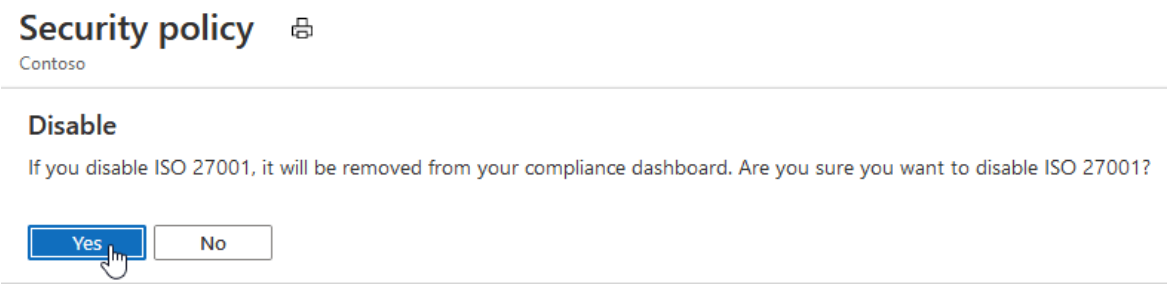

4. Select **Yes**. The standard will be removed.

Next steps

In this article, you learned how to **add compliance standards** to monitor your compliance with regulatory and industry standards.

For related material, see the following pages:

- Azure Security Benchmark
- Defender for Cloud regulatory compliance dashboard - Learn how to track and export your regulatory compliance data with Defender for Cloud and external tools
- Working with security policies

Planning and operations guide

2/15/2022 • 11 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This guide is for information technology (IT) professionals, IT architects, information security analysts, and cloud administrators planning to use Defender for Cloud.

Planning guide

This guide covers tasks that you can follow to optimize your use of Defender for Cloud based on your organization's security requirements and cloud management model. To take full advantage of Defender for Cloud, it is important to understand how different individuals or teams in your organization use the service to meet secure development and operations, monitoring, governance, and incident response needs. The key areas to consider when planning to use Defender for Cloud are:

- Security Roles and Access Controls
- Security Policies and Recommendations
- Data Collection and Storage
- Onboarding non-Azure resources
- Ongoing Security Monitoring
- Incident Response

In the next section, you will learn how to plan for each one of those areas and apply those recommendations based on your requirements.

> **NOTE**
> Read Defender for Cloud frequently asked questions (FAQ) for a list of common questions that can also be useful during the designing and planning phase.

Security roles and access controls

Depending on the size and structure of your organization, multiple individuals and teams may use Defender for Cloud to perform different security-related tasks. In the following diagram, you have an example of fictitious personas and their respective roles and security responsibilities:

Defender for Cloud enables these individuals to meet these various responsibilities. For example:

Jeff (Workload Owner)

- Manage a cloud workload and its related resources
- Responsible for implementing and maintaining protections in accordance with company security policy

Ellen (CISO/CIO)

- Responsible for all aspects of security for the company
- Wants to understand the company's security posture across cloud workloads
- Needs to be informed of major attacks and risks

David (IT Security)

- Sets company security policies to ensure the appropriate protections are in place
- Monitors compliance with policies
- Generates reports for leadership or auditors

Judy (Security Operations)

- Monitors and responds to security alerts 24/7
- Escalates to Cloud Workload Owner or IT Security Analyst

Sam (Security Analyst)

- Investigate attacks
- Work with Cloud Workload Owner to apply remediation

Defender for Cloud uses Azure role-based access control (Azure RBAC), which provides built-in roles that can be assigned to users, groups, and services in Azure. When a user opens Defender for Cloud, they only see information related to resources they have access to. Which means the user is assigned the role of Owner, Contributor, or Reader to the subscription or resource group that a resource belongs to. In addition to these roles, there are two specific Defender for Cloud roles:

- **Security reader**: a user that belongs to this role is able to view only Defender for Cloud configurations, which include recommendations, alerts, policy, and health, but it won't be able to make changes.
- **Security admin**: same as security reader but it can also update the security policy, dismiss recommendations and alerts.

The Defender for Cloud roles described above do not have access to other service areas of Azure such as Storage, Web & Mobile, or Internet of Things.

Using the personas explained in the previous diagram, the following Azure RBAC would be needed:

Jeff (Workload Owner)

- Resource Group Owner/Contributor

Ellen (CISO/CIO)

- Subscription Owner/Contributor or Security Admin

David (IT Security)

- Subscription Owner/Contributor or Security Admin

Judy (Security Operations)

- Subscription Reader or Security Reader to view Alerts
- Subscription Owner/Contributor or Security Admin required to dismiss Alerts

Sam (Security Analyst)

- Subscription Reader to view Alerts
- Subscription Owner/Contributor required to dismiss Alerts
- Access to the workspace may be required

Some other important information to consider:

- Only subscription Owners/Contributors and Security Admins can edit a security policy.
- Only subscription and resource group Owners and Contributors can apply security recommendations for a resource.

When planning access control using Azure RBAC for Defender for Cloud, be sure to understand who in your organization will be using Defender for Cloud. Also, what types of tasks they will be performing and then configure Azure RBAC accordingly.

> **NOTE**
> We recommend that you assign the least permissive role needed for users to complete their tasks. For example, users who only need to view information about the security state of resources but not take action, such as applying recommendations or editing policies, should be assigned the Reader role.

Security policies and recommendations

A security policy defines the desired configuration of your workloads and helps ensure compliance with company or regulatory security requirements. In Defender for Cloud, you can define policies for your Azure subscriptions, which can be tailored to the type of workload or the sensitivity of data.

Defender for Cloud policies contain the following components:

- Data collection: agent provisioning and data collection settings.
- Security policy: an Azure Policy that determines which controls are monitored and recommended by Defender for Cloud, or use Azure Policy to create new definitions, define additional policies, and assign policies across management groups.
- Email notifications: security contacts and notification settings.

- Pricing tier: with or without Microsoft Defender for Cloud's enhanced security features, which determine which Defender for Cloud features are available for resources in scope (can be specified for subscriptions and workspaces using the API).

> **NOTE**
> Specifying a security contact will ensure that Azure can reach the right person in your organization if a security incident occurs. Read Provide security contact details in Defender for Cloud for more information on how to enable this recommendation.

Security policies definitions and recommendations

Defender for Cloud automatically creates a default security policy for each of your Azure subscriptions. You can edit the policy in Defender for Cloud or use Azure Policy to create new definitions, define additional policies, and assign policies across Management Groups (which can represent the entire organization, a business unit in it etc.), and monitor compliance to these policies across these scopes.

Before configuring security policies, review each of the security recommendations, and determine whether these policies are appropriate for your various subscriptions and resource groups. It is also important to understand what action should be taken to address Security Recommendations and who in your organization will be responsible for monitoring for new recommendations and taking the needed steps.

Data collection and storage

Defender for Cloud uses the Log Analytics agent – this is the same agent used by the Azure Monitor service – to collect security data from your virtual machines. Data collected from this agent will be stored in your Log Analytics workspace(s).

Agent

When automatic provisioning is enabled in the security policy, the Log Analytics agent (for Windows or Linux) is installed on all supported Azure VMs, and any new ones that are created. If the VM or computer already has the Log Analytics agent installed, Defender for Cloud will leverage the current installed agent. The agent's process is designed to be non-invasive and have very minimal impact on VM performance.

The Log Analytics agent for Windows requires use TCP port 443. See the Troubleshooting article for additional details.

If at some point you want to disable Data Collection, you can turn it off in the security policy. However, because the Log Analytics agent may be used by other Azure management and monitoring services, the agent will not be uninstalled automatically when you turn off data collection in Defender for Cloud. You can manually uninstall the agent if needed.

> **NOTE**
> To find a list of supported VMs, read the Defender for Cloud frequently asked questions (FAQ).

Workspace

A workspace is an Azure resource that serves as a container for data. You or other members of your organization might use multiple workspaces to manage different sets of data that is collected from all or portions of your IT infrastructure.

Data collected from the Log Analytics agent (on behalf of Defender for Cloud) will be stored in either an existing Log Analytics workspace(s) associated with your Azure subscription or a new workspace(s), taking into account the Geo of the VM.

In the Azure portal, you can browse to see a list of your Log Analytics workspaces, including any created by Defender for Cloud. A related resource group will be created for new workspaces. Both will follow this naming convention:

- Workspace: *DefaultWorkspace-[subscription-ID]-[geo]*
- Resource Group: *DefaultResourceGroup-[geo]*

For workspaces created by Defender for Cloud, data is retained for 30 days. For existing workspaces, retention is based on the workspace pricing tier. If you want, you can also use an existing workspace.

If your agent reports to a workspace other than the **default** workspace, any Microsoft Defender plans providing enhanced security features that you've enabled on the subscription should also be enabled on the workspace.

> **NOTE**
> Microsoft makes strong commitments to protect the privacy and security of this data. Microsoft adheres to strict compliance and security guidelines—from coding to operating a service. For more information about data handling and privacy, read Defender for Cloud Data Security.

Onboard non-Azure resources

Defender for Cloud can monitor the security posture of your non-Azure computers but you need to first onboard these resources. Read Onboard non-Azure computers for more information on how to onboard non-Azure resources.

Ongoing security monitoring

After initial configuration and application of Defender for Cloud recommendations, the next step is considering Defender for Cloud operational processes.

The Defender for Cloud Overview provides a unified view of security across all your Azure resources and any non-Azure resources you have connected. The example below shows an environment with many issues to be addressed:

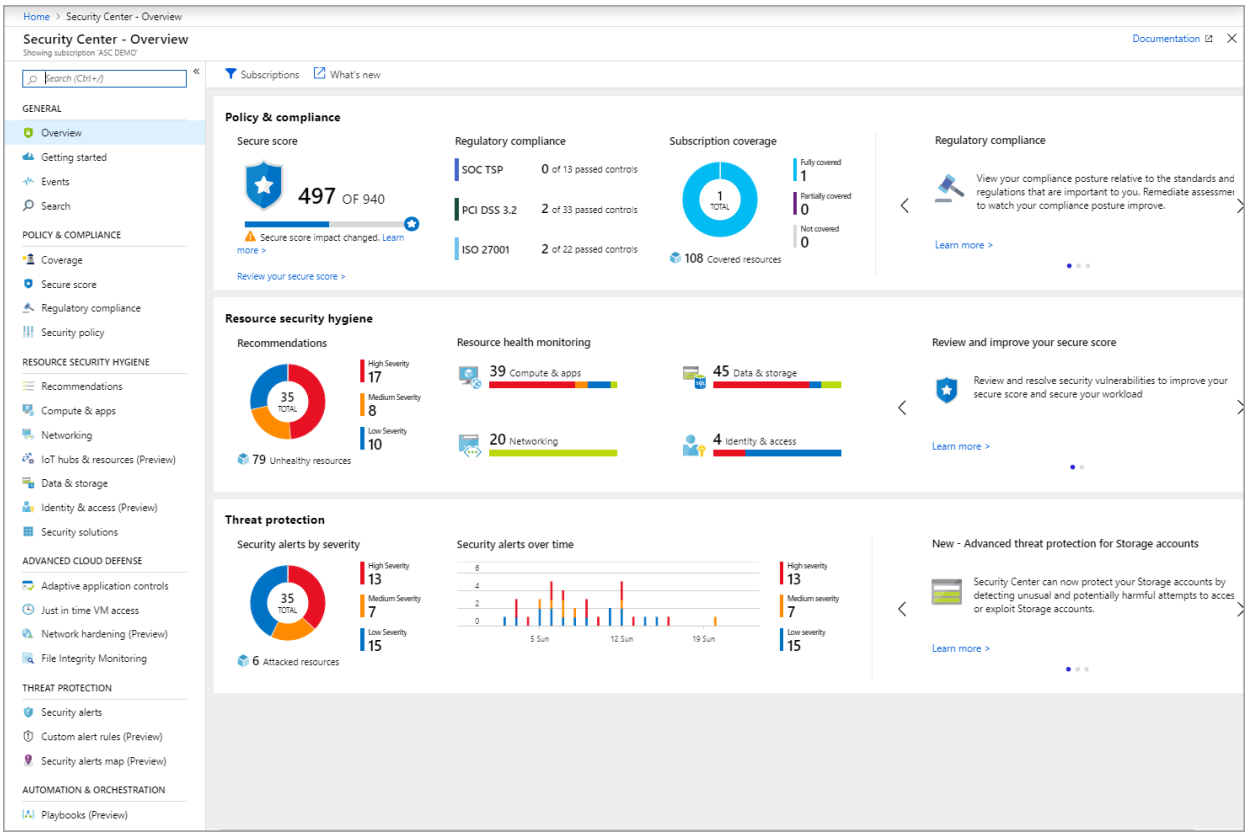

> **NOTE**
>
> Defender for Cloud will not interfere with your normal operational procedures, it will passively monitor your deployments and provide recommendations based on the security policies you enabled.

When you first opt in to use Defender for Cloud for your current Azure environment, make sure that you review all recommendations, which can be done in the **Recommendations** page.

Plan to visit the threat intelligence option as part of your daily security operations. There you can identify security threats against the environment, such as identify if a particular computer is part of a botnet.

Monitoring for new or changed resources

Most Azure environments are dynamic, with resources regularly being created, spun up or down, reconfigured, and changed. Defender for Cloud helps ensure that you have visibility into the security state of these new resources.

When you add new resources (VMs, SQL DBs) to your Azure Environment, Defender for Cloud will automatically discover these resources and begin to monitor their security. This also includes PaaS web roles and worker roles. If Data Collection is enabled in the Security Policy, additional monitoring capabilities will be enabled automatically for your virtual machines.

You should also regularly monitor existing resources for configuration changes that could have created security risks, drift from recommended baselines, and security alerts.

Hardening access and applications

As part of your security operations, you should also adopt preventative measures to restrict access to VMs, and control the applications that are running on VMs. By locking down inbound traffic to your Azure VMs, you are reducing the exposure to attacks, and at the same time providing easy access to connect to VMs when needed. Use just-in-time VM access access feature to hardening access to your VMs.

You can use adaptive application controls to limit which applications can run on your VMs located in Azure. Among other benefits, this helps harden your VMs against malware. Using machine learning, Defender for

Cloud analyzes processes running in the VM to help you create allow listing rules.

Incident response

Defender for Cloud detects and alerts you to threats as they occur. Organizations should monitor for new security alerts and take action as needed to investigate further or remediate the attack. For more information on how Defender for Cloud threat protection works, read How Defender for Cloud detects and responds to threats.

While this article doesn't have the intent to assist you creating your own Incident Response plan, we are going to use Microsoft Azure Security Response in the Cloud lifecycle as the foundation for incident response stages. The stages are shown in the following diagram:

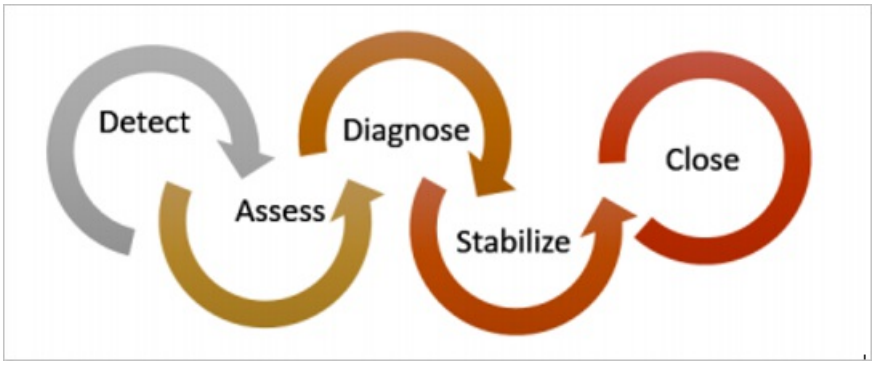

> **NOTE**
> You can use the National Institute of Standards and Technology (NIST) Computer Security Incident Handling Guide as a reference to assist you building your own.

You can use Defender for Cloud Alerts during the following stages:

- **Detect**: identify a suspicious activity in one or more resources.
- **Assess**: perform the initial assessment to obtain more information about the suspicious activity.
- **Diagnose**: use the remediation steps to conduct the technical procedure to address the issue.

Each Security Alert provides information that can be used to better understand the nature of the attack and suggest possible mitigations. Some alerts also provide links to either more information or to other sources of information within Azure. You can use the information provided for further research and to begin mitigation, and you can also search security-related data that is stored in your workspace.

The following example shows a suspicious RDP activity taking place:

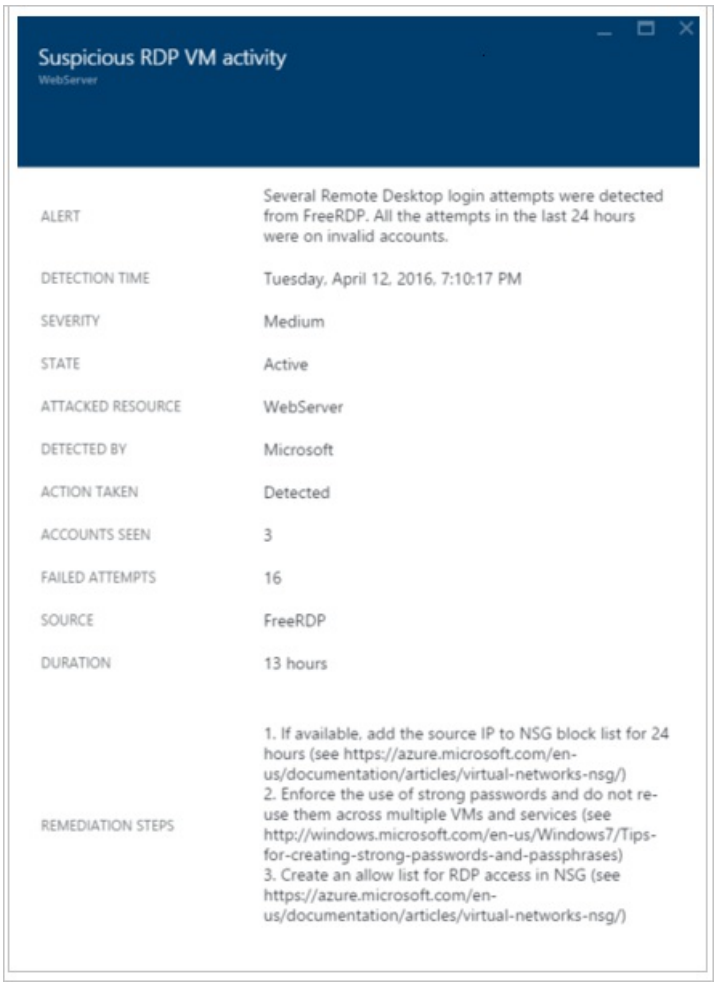

This page shows the details regarding the time that the attack took place, the source hostname, the target VM and also gives recommendation steps. In some circumstances, the source information of the attack may be empty. Read Missing Source Information in Defender for Cloud Alerts for more information about this type of behavior.

Once you identify the compromised system, you can run a workflow automation that was previously created. These are a collection of procedures that can be executed from Defender for Cloud once triggered by an alert.

In the How to Leverage the Defender for Cloud & Microsoft Operations Management Suite for an Incident Response video, you can see some demonstrations that show how Defender for Cloud can be used in each one of those stages.

> **NOTE**
> Read Managing and responding to security alerts in Defender for Cloud for more information on how to use Defender for Cloud capabilities to assist you during your Incident Response process.

Next steps

In this document, you learned how to plan for Defender for Cloud adoption. To learn more about Defender for Cloud, see the following:

- Managing and responding to security alerts in Defender for Cloud
- Monitoring partner solutions with Defender for Cloud — Learn how to monitor the health status of your partner solutions.
- Defender for Cloud FAQ — Find frequently asked questions about using the service.
- Azure Security blog — Find blog posts about Azure security and compliance.

Tutorial: Protect your resources with Microsoft Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud limits your exposure to threats by using access and application controls to block malicious activity. Just-in-time (JIT) virtual machine (VM) access reduces your exposure to attacks by enabling you to deny persistent access to VMs. Instead, you provide controlled and audited access to VMs only when needed. Adaptive application controls help harden VMs against malware by controlling which applications can run on your VMs. Defender for Cloud uses machine learning to analyze the processes running in the VM and helps you apply allow listing rules using this intelligence.

In this tutorial you'll learn how to:

- Configure a just-in-time VM access policy
- Configure an application control policy

Prerequisites

To step through the features covered in this tutorial, you must have Defender for Cloud's enhanced security features enabled. A free trial is available. To upgrade, see Enable enhanced protections.

Manage VM access

JIT VM access can be used to lock down inbound traffic to your Azure VMs, reducing exposure to attacks while providing easy access to connect to VMs when needed.

Management ports do not need to be open at all times. They only need to be open while you are connected to the VM, for example to perform management or maintenance tasks. When just-in-time is enabled, Defender for Cloud uses Network Security Group (NSG) rules, which restrict access to management ports so they cannot be targeted by attackers.

Follow the guidance in Secure your management ports with just-in-time access.

Harden VMs against malware

Adaptive application controls help you define a set of applications that are allowed to run on configured resource groups, which among other benefits helps harden your VMs against malware. Defender for Cloud uses machine learning to analyze the processes running in the VM and helps you apply allow listing rules using this intelligence.

Follow the guidance in Use adaptive application controls to reduce your machines' attack surfaces.

Next steps

In this tutorial, you learned how to limit your exposure to threats by:

- Configuring a just-in-time VM access policy to provide controlled and audited access to VMs only when needed
- Configuring an adaptive application controls policy to control which applications can run on your VMs

Advance to the next tutorial to learn about responding to security incidents.

Tutorial: Respond to security incidents

Tutorial: Triage, investigate, and respond to security alerts

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud continuously analyzes your hybrid cloud workloads using advanced analytics and threat intelligence to alert you about potentially malicious activities in your cloud resources. You can also integrate alerts from other security products and services into Defender for Cloud. Once an alert is raised, swift action is needed to investigate and remediate the potential security issue.

In this tutorial, you will learn how to:

- Triage security alerts
- Investigate a security alert to determine the root cause
- Respond to a security alert and mitigate that root cause

If you don't have an Azure subscription, create a free account before you begin.

Prerequisites

To step through the features covered in this tutorial, you must have Defender for Cloud's enhanced security features enabled. You can try these at no cost. To learn more, see the pricing page. The quickstart Get started with Defender for Cloud walks you through how to upgrade.

Triage security alerts

Defender for Cloud provides a unified view of all security alerts. Security alerts are ranked based on the severity of the detected activity.

Triage your alerts from the **Security alerts** page:

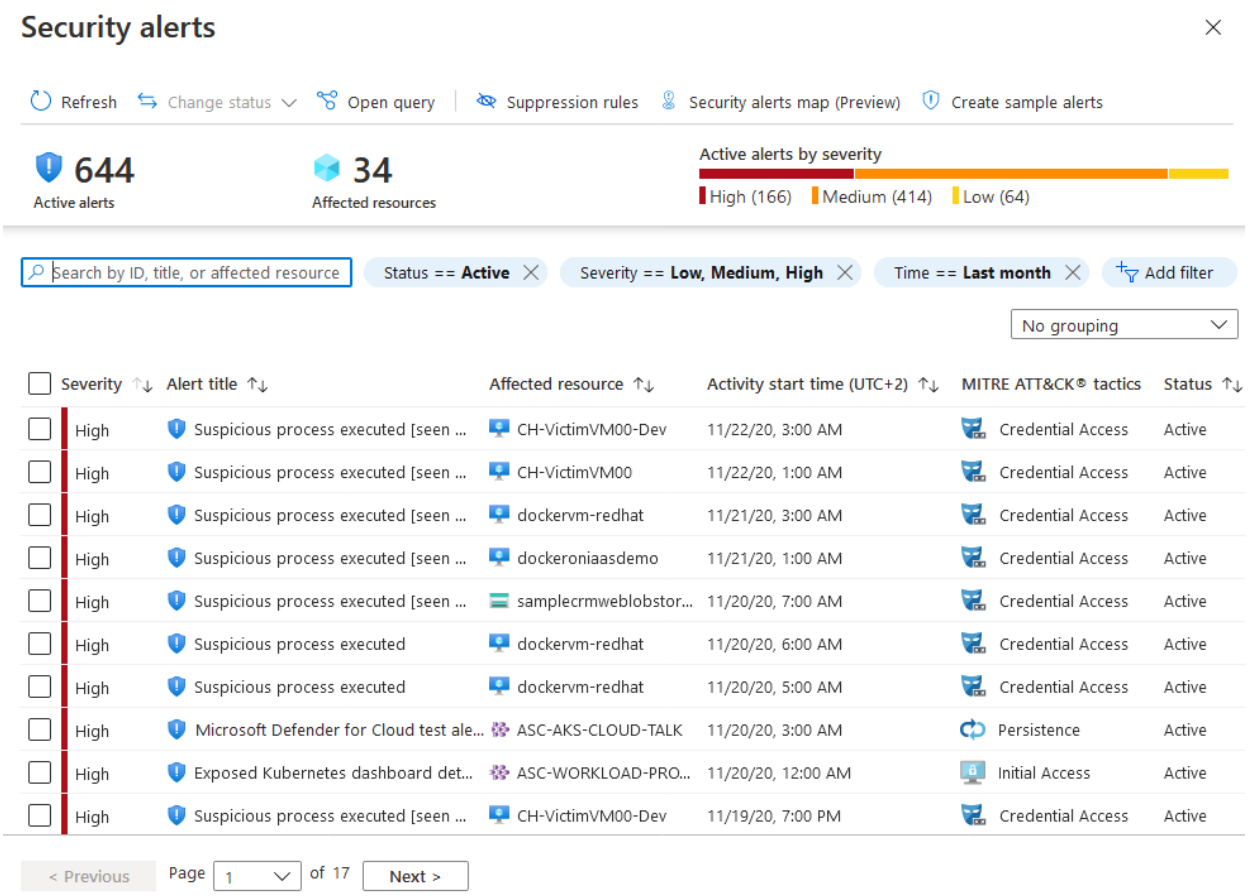

Use this page to review the active security alerts in your environment to decide which alert to investigate first.

When triaging security alerts, prioritize alerts based on the alert severity by addressing alerts with higher severity first. Learn more about alerts severity in How are alerts classified?.

> **TIP**
> You can connect Microsoft Defender for Cloud to most popular SIEM solutions including Microsoft Sentinel and consume the alerts from your tool of choice. Learn more in Stream alerts to a SIEM, SOAR, or IT Service Management solution.

Investigate a security alert

When you've decided which alert to investigate first:

1. Select the desired alert.
2. From the alert overview page, select the resource to investigate first.
3. Begin your investigation from the left pane, which shows the high-level information about the security alert.

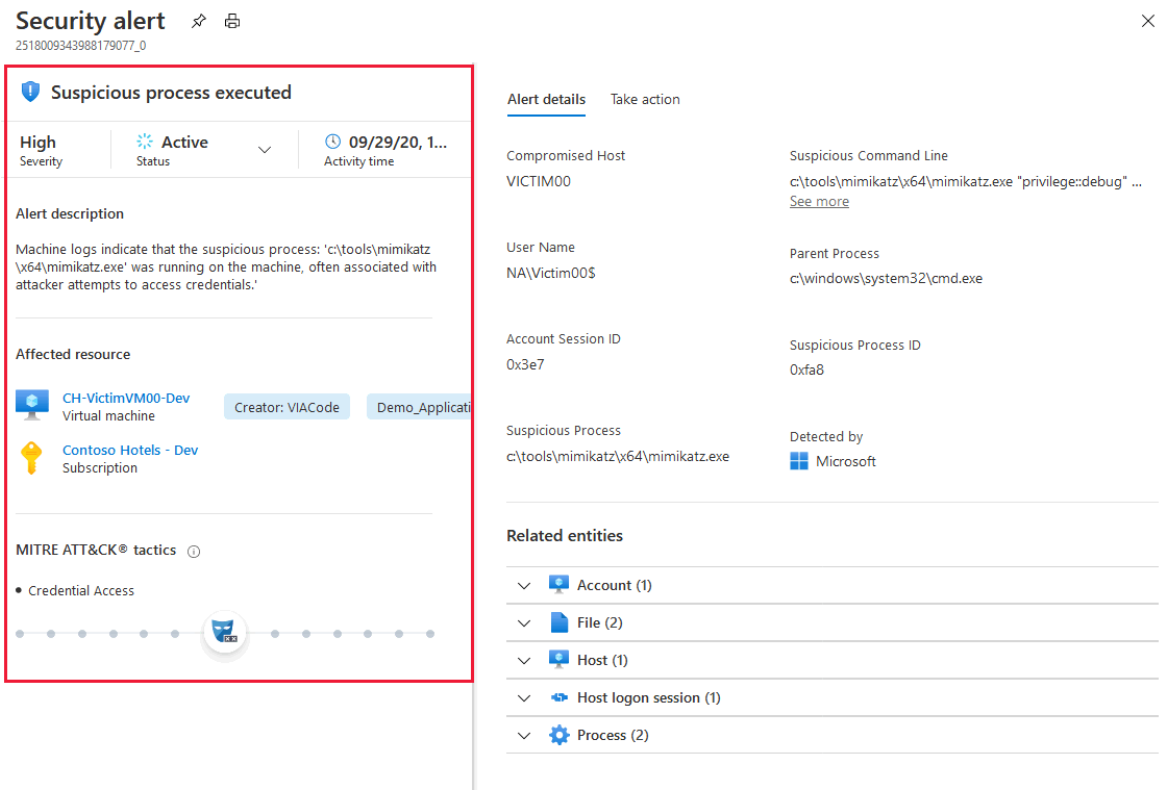

This pane shows:

- Alert severity, status, and activity time
- Description that explains the precise activity that was detected
- Affected resources
- Kill chain intent of the activity on the MITRE ATT&CK matrix

4. For more detailed information that can help you investigate the suspicious activity, examine the **Alert details** tab.

5. When you've reviewed the information on this page, you may have enough to proceed with a response. If you need further details:

   - Contact the resource owner to verify whether the detected activity is a false positive.
   - Investigate the raw logs generated by the attacked resource

Respond to a security alert

After you've investigated a security alert and understood its scope, you can respond to the alert from within Microsoft Defender for Cloud:

1. Open the **Take action** tab to see the recommended responses.

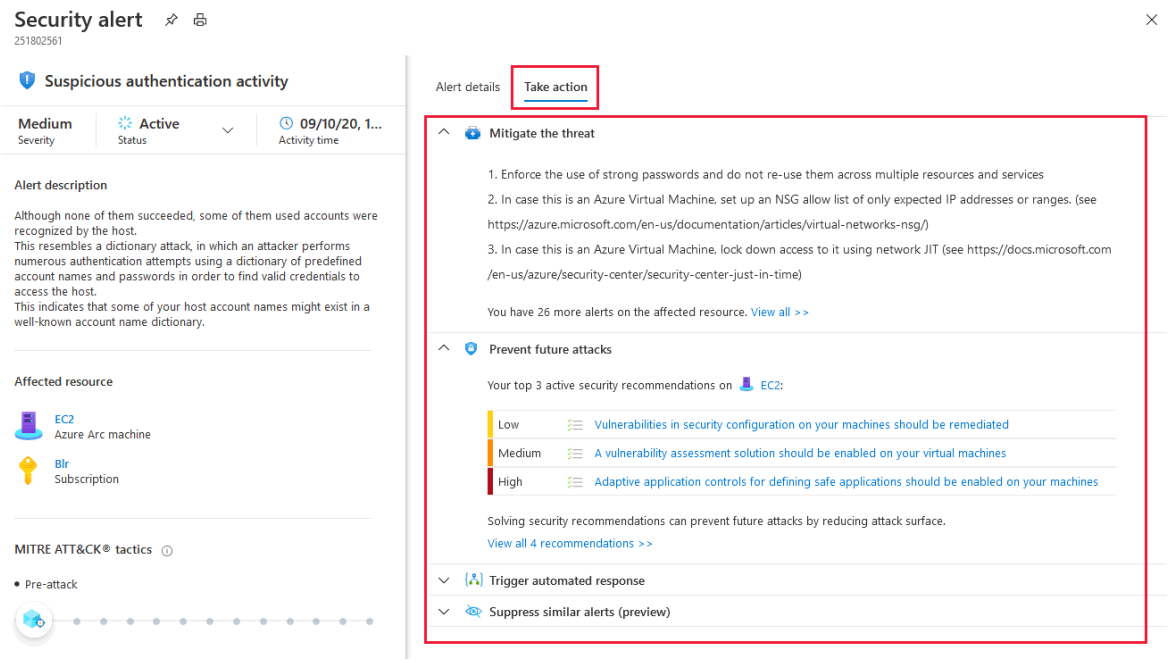

2. Review the **Mitigate the threat** section for the manual investigation steps necessary to mitigate the issue.
3. To harden your resources and prevent future attacks of this kind, remediate the security recommendations in the **Prevent future attacks** section.
4. To trigger a logic app with automated response steps, use the **Trigger automated response** section.
5. If the detected activity *isn't* malicious, you can suppress future alerts of this kind using the **Suppress similar alerts** section.
6. When you've completed the investigation into the alert and responded in the appropriate way, change the status to **Dismissed**.

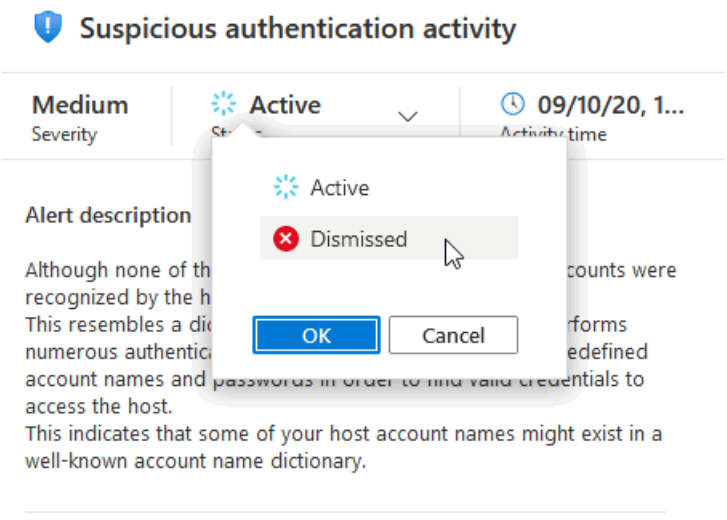

This removes the alert from the main alerts list. You can use the filter from the alerts list page to view all alerts with **Dismissed** status.

7. We encourage you to provide feedback about the alert to Microsoft:
   a. Marking the alert as **Useful** or **Not useful**.
   b. Select a reason and add a comment.

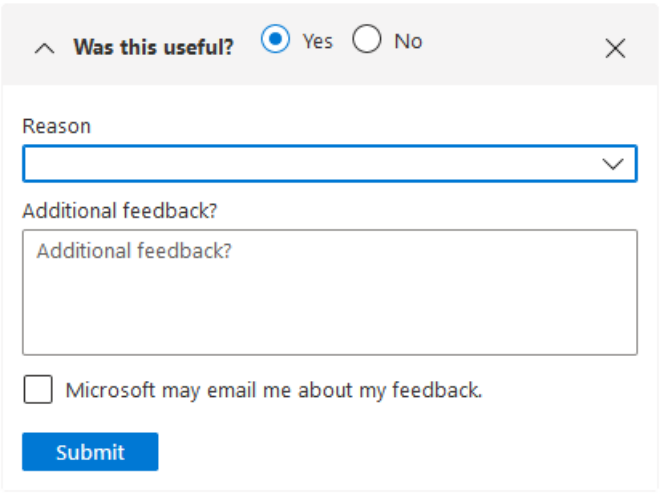

> **TIP**
> We review your feedback to improve our algorithms and provide better security alerts.

End the tutorial

Other quickstarts and tutorials in this collection build upon this quickstart. If you plan to continue to work with subsequent quickstarts and tutorials, keep automatic provisioning and Defender for Cloud's enhanced security features enabled.

If you don't plan to continue, or you want to disable either of these features:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.
3. Select **Defender plans** and select **Enhanced security off**.

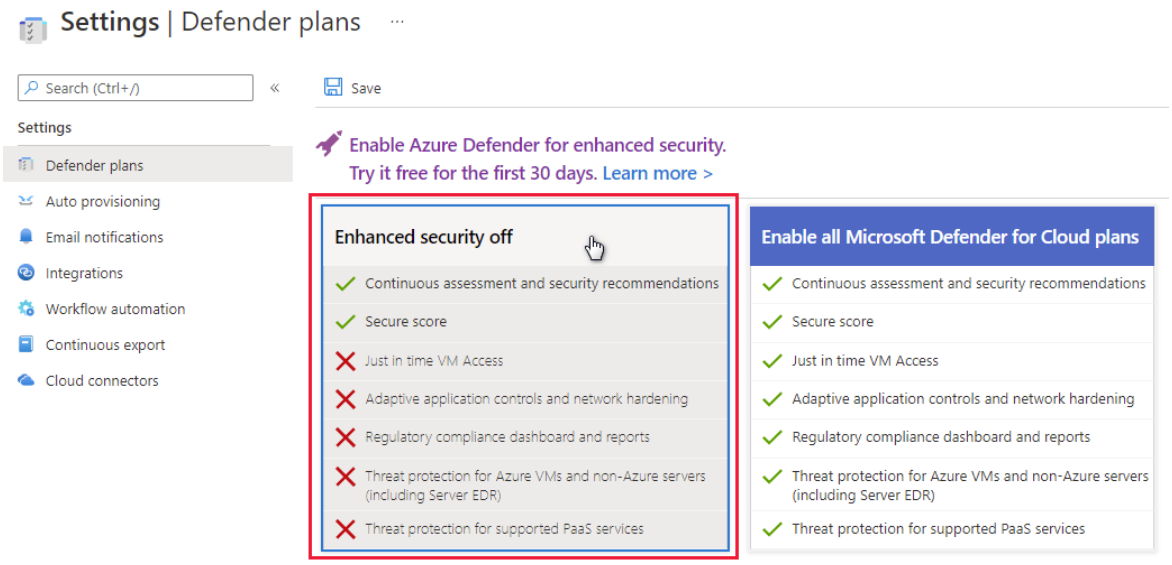

4. Select **Save**.

> **NOTE**
> After you disable enhanced security features - whether you disable a single plan or all plans at once - data collection may continue for a short period of time.

5. From Defender for Cloud's menu, open **Environment settings**.
6. Select the relevant subscription.
7. Select **Auto provisioning**.
8. Disable the relevant extensions.

> **NOTE**
> Disabling automatic provisioning does not remove the Log Analytics agent from Azure VMs that already have the agent. Disabling automatic provisioning limits security monitoring for your resources.

Next steps

In this tutorial, you learned about Defender for Cloud features to be used when responding to a security alert. For related material see:

- Respond to Microsoft Defender for Key Vault alerts
- Security alerts - a reference guide
- Introduction to Defender for Cloud

Tutorial: Improve your regulatory compliance

2/15/2022 • 11 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud helps streamline the process for meeting regulatory compliance requirements, using the **regulatory compliance dashboard**.

Defender for Cloud continuously assesses your hybrid cloud environment to analyze the risk factors according to the controls and best practices in the standards that you've applied to your subscriptions. The dashboard reflects the status of your compliance with these standards.

When you enable Defender for Cloud on an Azure subscription, the Azure Security Benchmark is automatically assigned to that subscription. This widely respected benchmark builds on the controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security.

The regulatory compliance dashboard shows the status of all the assessments within your environment for your chosen standards and regulations. As you act on the recommendations and reduce risk factors in your environment, your compliance posture improves.

In this tutorial you'll learn how to:

- Evaluate your regulatory compliance using the regulatory compliance dashboard
- Improve your compliance posture by taking action on recommendations
- Download PDF/CSV reports as well as certification reports of your compliance status
- Setup alerts on changes to your compliance status
- Export your compliance data as a continuous stream and as weekly snapshots

If you don't have an Azure subscription, create a free account before you begin.

Prerequisites

To step through the features covered in this tutorial:

- Enable enhanced security features. You can enable these for free for 30 days.
- You must be signed in with an account that has reader access to the policy compliance data (**Security Reader** is insufficient). The role of **Global reader** for the subscription will work. At a minimum, you'll need to have **Resource Policy Contributor** and **Security Admin** roles assigned.

Assess your regulatory compliance

The regulatory compliance dashboard shows your selected compliance standards with all their requirements, where supported requirements are mapped to applicable security assessments. The status of these assessments reflects your compliance with the standard.

Use the regulatory compliance dashboard to help focus your attention on the gaps in compliance with your

chosen standards and regulations. This focused view also enables you to continuously monitor your compliance over time within dynamic cloud and hybrid environments.

1. From Defender for Cloud's menu, select **Regulatory compliance**.

   At the top of the screen, is a dashboard with an overview of your compliance status and the set of supported compliance regulations. You'll see your overall compliance score, and the number of passing vs. failing assessments associated with each standard.

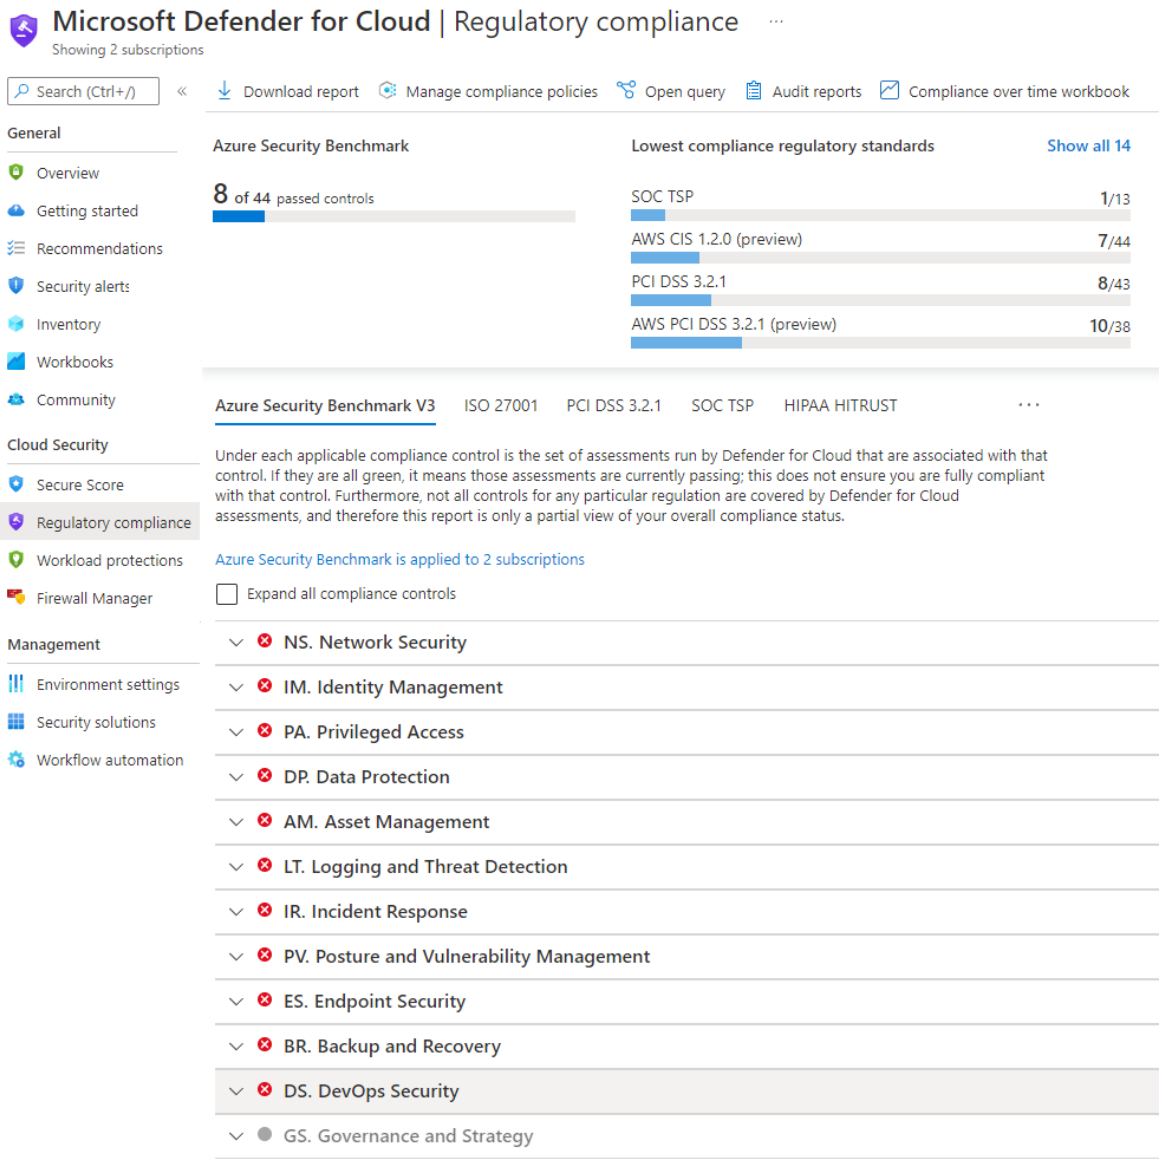

2. Select a tab for a compliance standard that is relevant to you (1). You'll see which subscriptions the standard is applied on (2), and the list of all controls for that standard (3). For the applicable controls, you can view the details of passing and failing assessments associated with that control (4), and the number of affected resources (5). Some controls are grayed out. These controls don't have any Defender for Cloud assessments associated with them. Check their requirements and assess them in your environment. Some of these might be process-related and not technical.

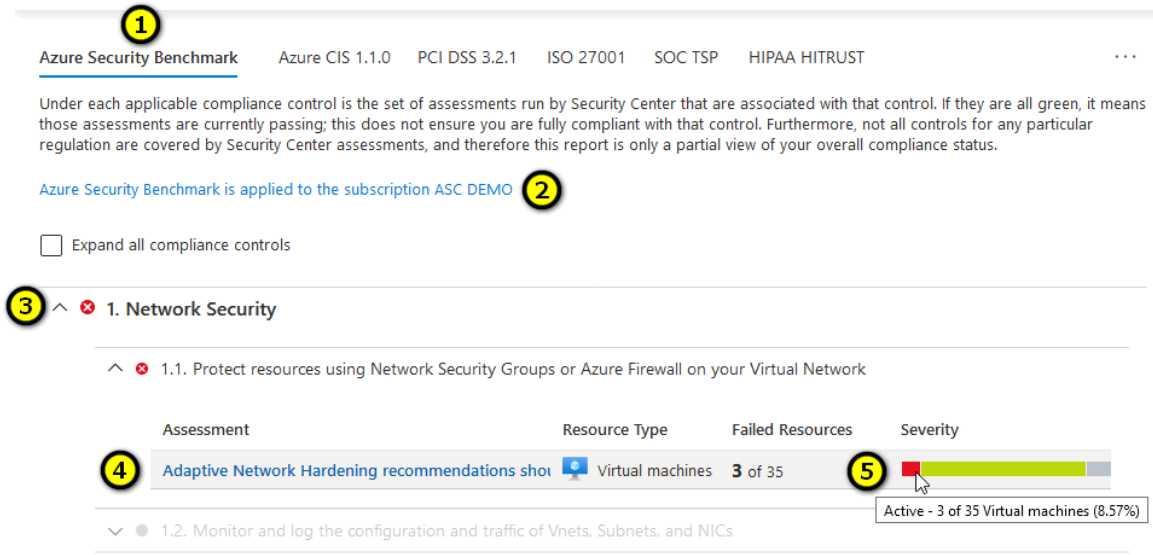

Improve your compliance posture

Using the information in the regulatory compliance dashboard, improve your compliance posture by resolving recommendations directly within the dashboard.

1. Select any of the failing assessments that appear in the dashboard to view the details for that recommendation. Each recommendation includes a set of remediation steps to resolve the issue.

2. Select a particular resource to view more details and resolve the recommendation for that resource. For example, in the **Azure CIS 1.1.0** standard, select the recommendation **Disk encryption should be applied on virtual machines**.

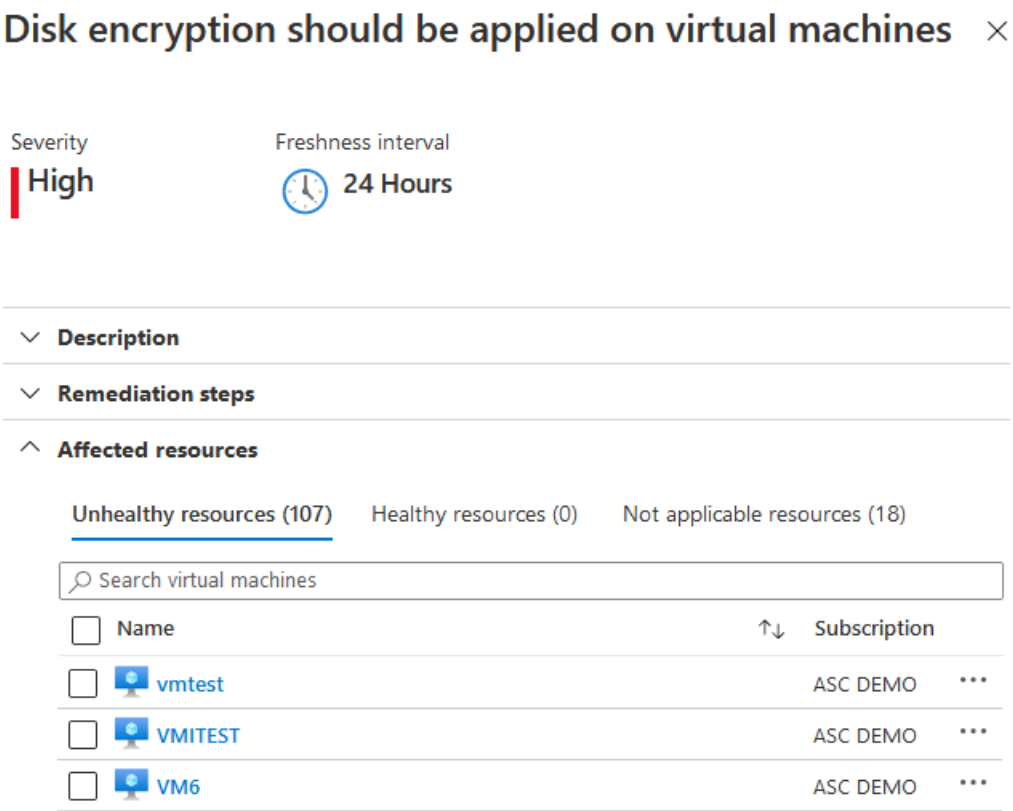

3. In this example, when you select **Take action** from the recommendation details page, you arrive in the Azure Virtual Machine pages of the Azure portal, where you can enable encryption from the **Security** tab:

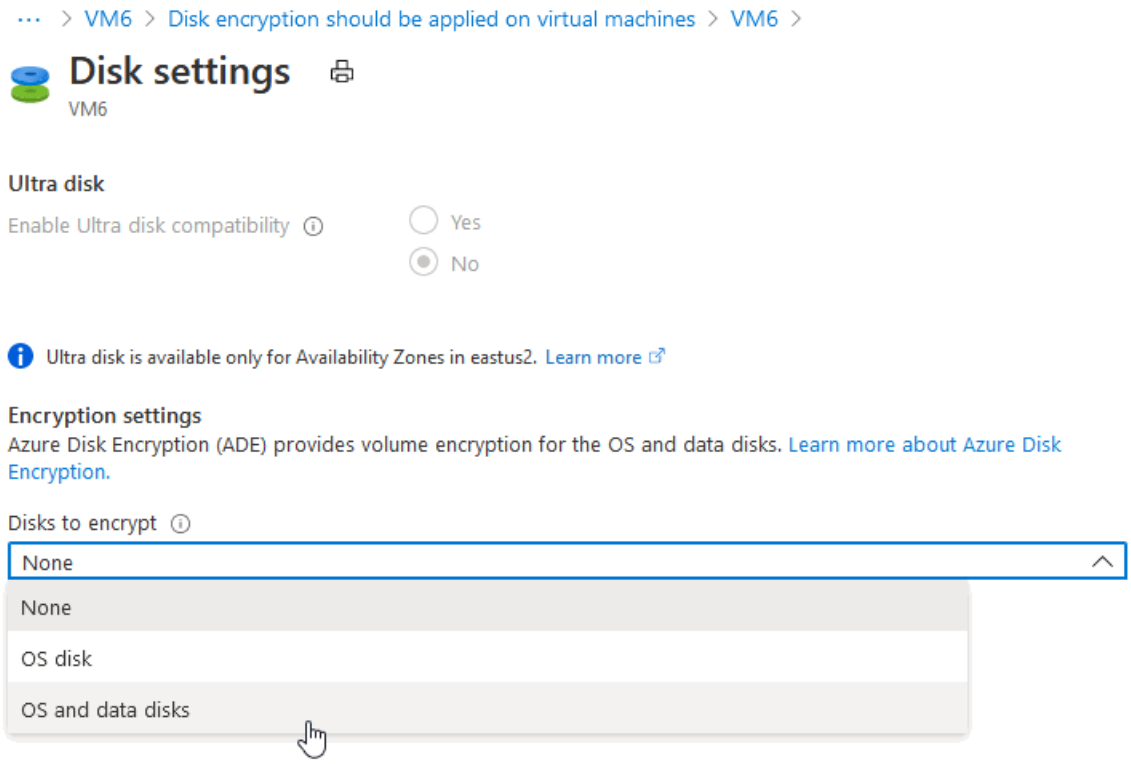

For more information about how to apply recommendations, see Implementing security recommendations in Microsoft Defender for Cloud.

4. After you take action to resolve recommendations, you'll see the result in the compliance dashboard report because your compliance score improves.

> **NOTE**
> Assessments run approximately every 12 hours, so you will see the impact on your compliance data only after the next run of the relevant assessment.

Generate compliance status reports and certificates

- To generate a PDF report with a summary of your current compliance status for a particular standard, select **Download report**.

  The report provides a high-level summary of your compliance status for the selected standard based on Defender for Cloud assessments data. The report's organized according to the controls of that particular standard. The report can be shared with relevant stakeholders, and might provide evidence to internal and external auditors.

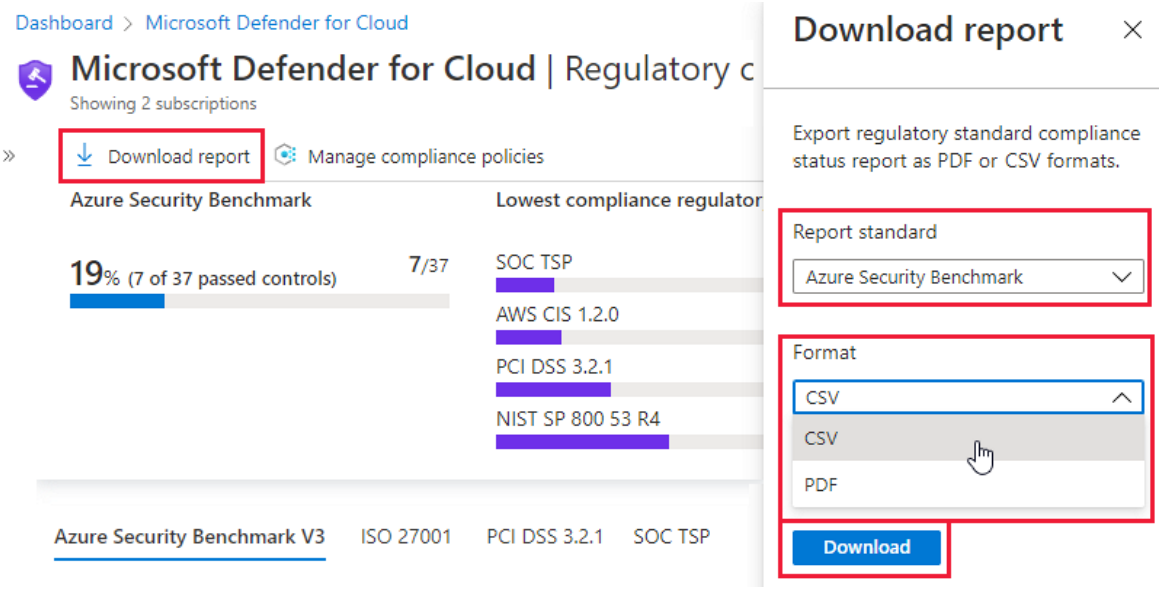

- To download Azure and Dynamics **certification reports** for the standards applied to your subscriptions, use the **Audit reports** option.

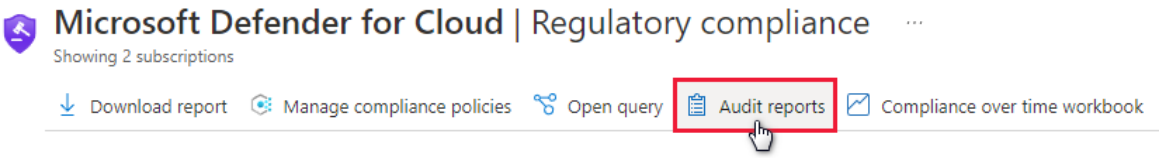

Select the tab for the relevant reports types (PCI, SOC, ISO, and others) and use filters to find the specific reports you need:

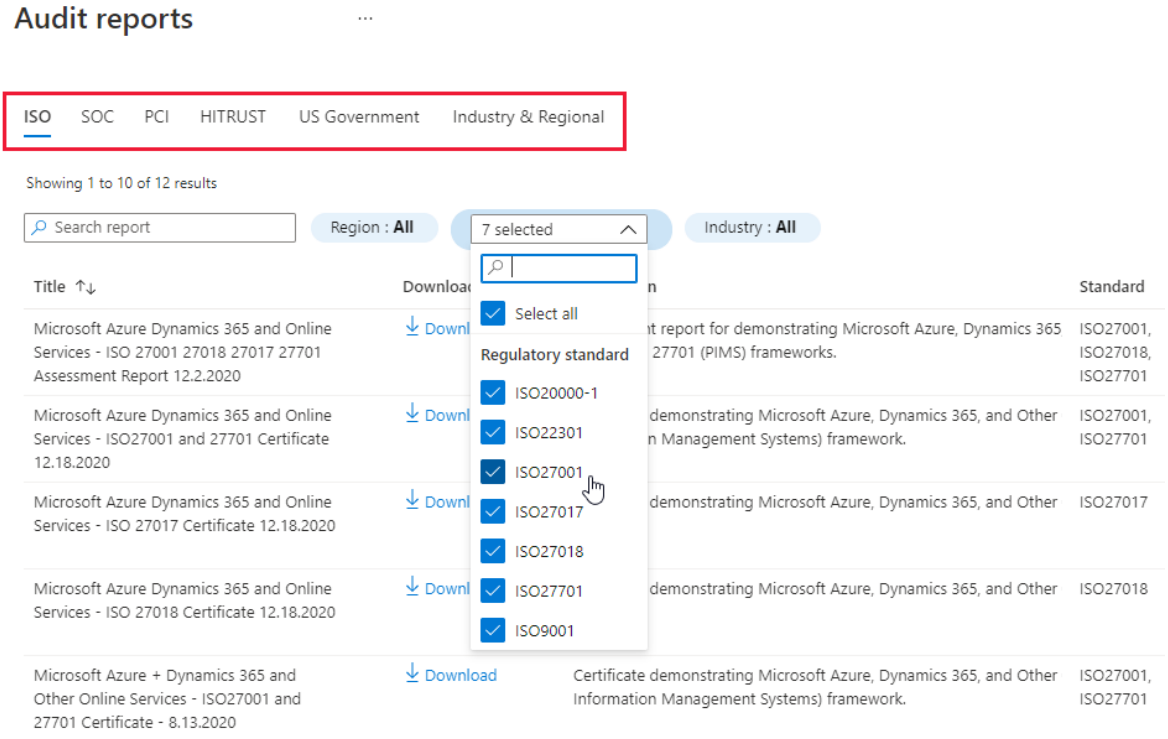

For example, from the PCI tab you can download a ZIP file containing a digitally signed certificate demonstrating Microsoft Azure, Dynamics 365, and Other Online Services' compliance with ISO22301 framework, together with the necessary collateral to interpret and present the certificate.

> **NOTE**
>
> When you download one of these certification reports, you'll be shown the following privacy notice:
>
> *By downloading this file, you are giving consent to Microsoft to store the current user and the selected subscriptions at the time of download. This data is used in order to notify you in case of changes or updates to the downloaded audit report. This data is used by Microsoft and the audit firms that produce the certification/reports only when notification is required.*

Configure frequent exports of your compliance status data

If you want to track your compliance status with other monitoring tools in your environment, Defender for Cloud includes an export mechanism to make this straightforward. Configure **continuous export** to send select data to an Azure Event Hub or a Log Analytics workspace. Learn more in continuously export Defender for Cloud data.

Use continuous export data to an Azure Event Hub or a Log Analytics workspace:

- Export all regulatory compliance data in a **continuous stream**:

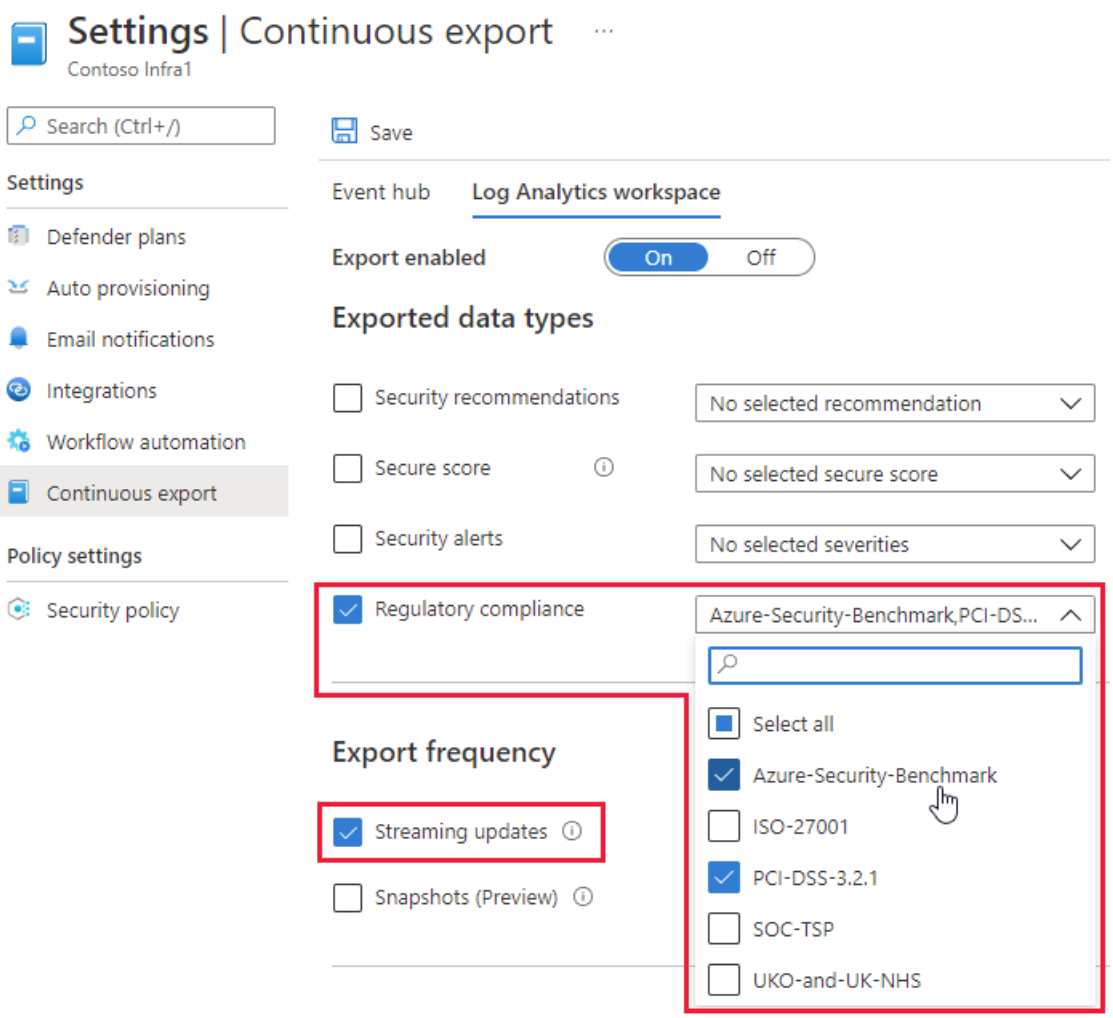

- Export **weekly snapshots** of your regulatory compliance data:

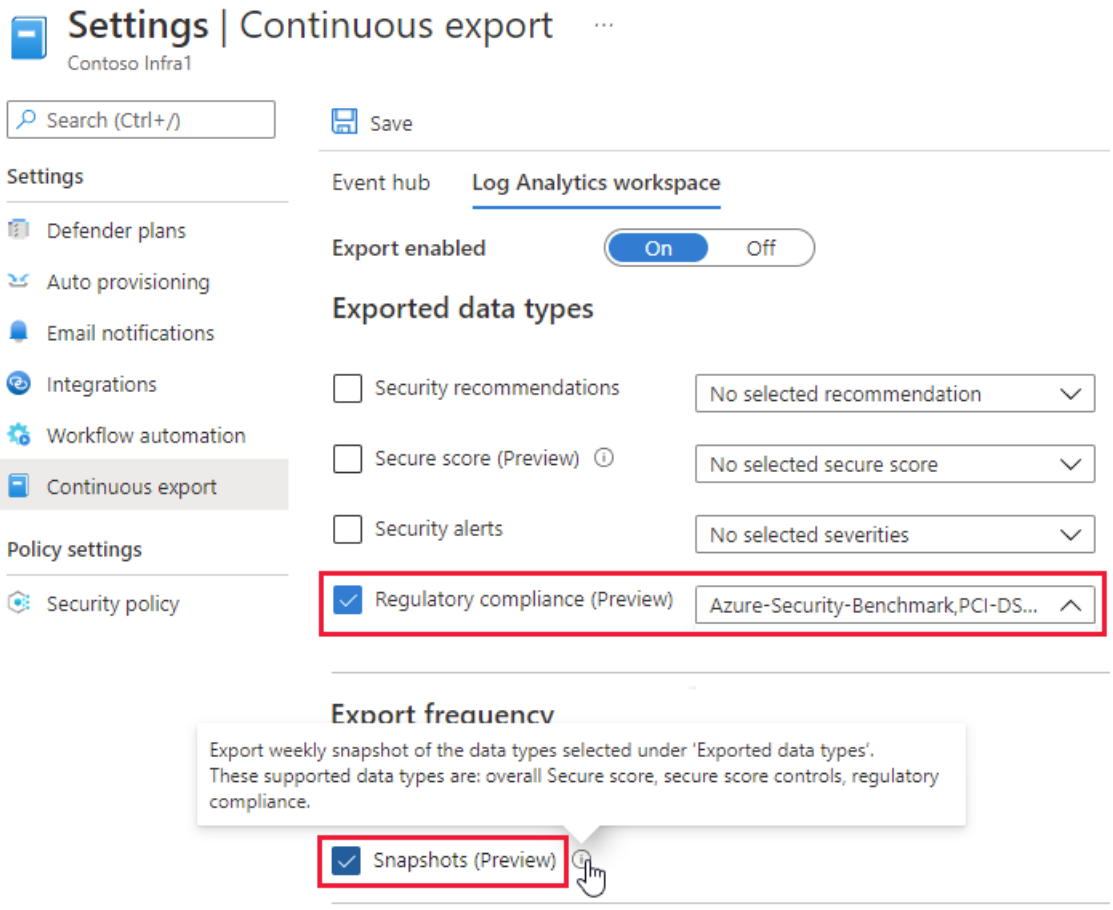

TIP

You can also manually export reports about a single point in time directly from the regulatory compliance dashboard. Generate these **PDF/CSV reports** or **Azure and Dynamics certification reports** using the **Download report** or **Audit reports** toolbar options. See Assess your regulatory compliance

Run workflow automations when there are changes to your compliance

Defender for Cloud's workflow automation feature can trigger Logic Apps whenever one of your regulatory compliance assessments change state.

For example, you might want Defender for Cloud to email a specific user when a compliance assessment fails. You'll need to create the logic app first (using Azure Logic Apps) and then set up the trigger in a new workflow automation as explained in Automate responses to Defender for Cloud triggers.

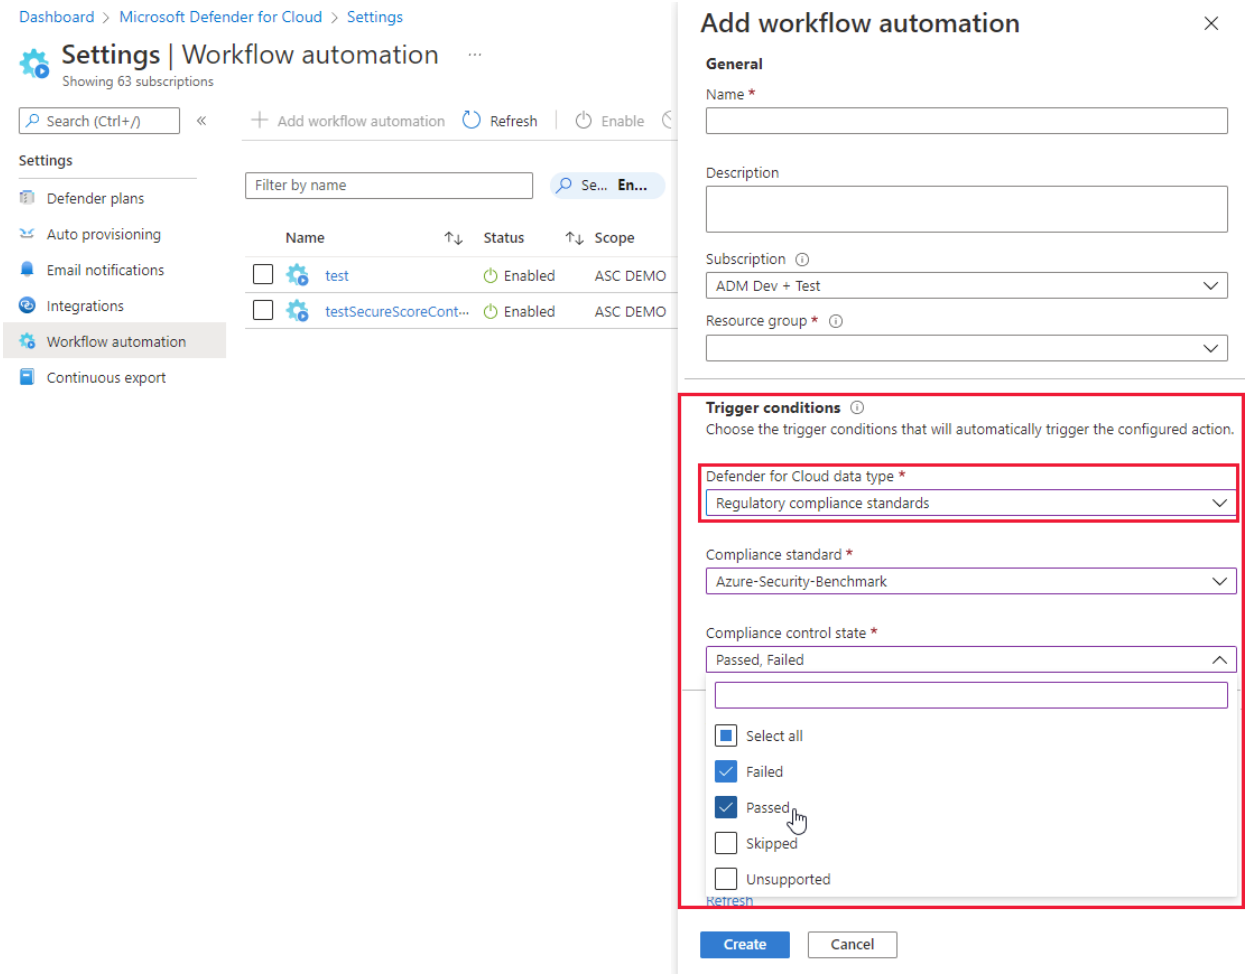

FAQ - Regulatory compliance dashboard

What standards are supported in the compliance dashboard?

By default, the regulatory compliance dashboard shows you the Azure Security Benchmark. The Azure Security Benchmark is the Microsoft-authored, Azure-specific guidelines for security, and compliance best practices based on common compliance frameworks. Learn more in the Azure Security Benchmark introduction.

To track your compliance with any other standard, you'll need to explicitly add them to your dashboard.

You can add other standards such as Azure CIS 1.3.0, NIST SP 800-53, NIST SP 800-171, SWIFT CSP CSCF-v2020, UK Official and UK NHS, HIPAA, Canada Federal PBMM, ISO 27001, SOC2-TSP, and PCI-DSS 3.2.1.

More standards will be added to the dashboard and included in the information on Customize the set of standards in your regulatory compliance dashboard.

Why do some controls appear grayed out?

For each compliance standard in the dashboard, there's a list of the standard's controls. For the applicable controls, you can view the details of passing and failing assessments.

Some controls are grayed out. These controls don't have any Defender for Cloud assessments associated with them. Some may be procedure or process-related, and so can't be verified by Defender for Cloud. Some don't have any automated policies or assessments implemented yet, but will have in the future. And some controls may be the platform's responsibility as explained in Shared responsibility in the cloud.

How can I remove a built-in standard, like PCI-DSS, ISO 27001, or SOC2 TSP from the dashboard?

To customize the regulatory compliance dashboard, and focus only on the standards that are applicable to you, you can remove any of the displayed regulatory standards that aren't relevant to your organization. To remove a standard, follow the instructions in Remove a standard from your dashboard.

I made the suggested changed based on the recommendation, yet it isn't being reflected in the dashboard

After you take action to resolve recommendations, wait 12 hours to see the changes to your compliance data. Assessments are run approximately every 12 hours, so you'll see the effect on your compliance data only after the assessments run.

What permissions do I need to access the compliance dashboard?

To view compliance data, you need to have at least **Reader** access to the policy compliance data as well; so Security Reader alone won't suffice. If you're a global reader on the subscription, that will be enough too.

The minimum set of roles for accessing the dashboard and managing standards is **Resource Policy Contributor** and **Security Admin**.

The regulatory compliance dashboard isn't loading for me

To use the regulatory compliance dashboard, Defender for Cloud must be enabled at the subscription level. If the dashboard isn't loading correctly, try the following steps:

1. Clear your browser's cache.
2. Try a different browser.
3. Try opening the dashboard from different network location.

How can I view a report of passing and failing controls per standard in my dashboard?

On the main dashboard, you can see a report of passing and failing controls for (1) the 'top 4' lowest compliance standards in the dashboard. To see all the passing/failing controls status, select (2) **Show all *x*** (where x is the number of standards you're tracking). A context plane displays the compliance status for every one of your tracked standards.

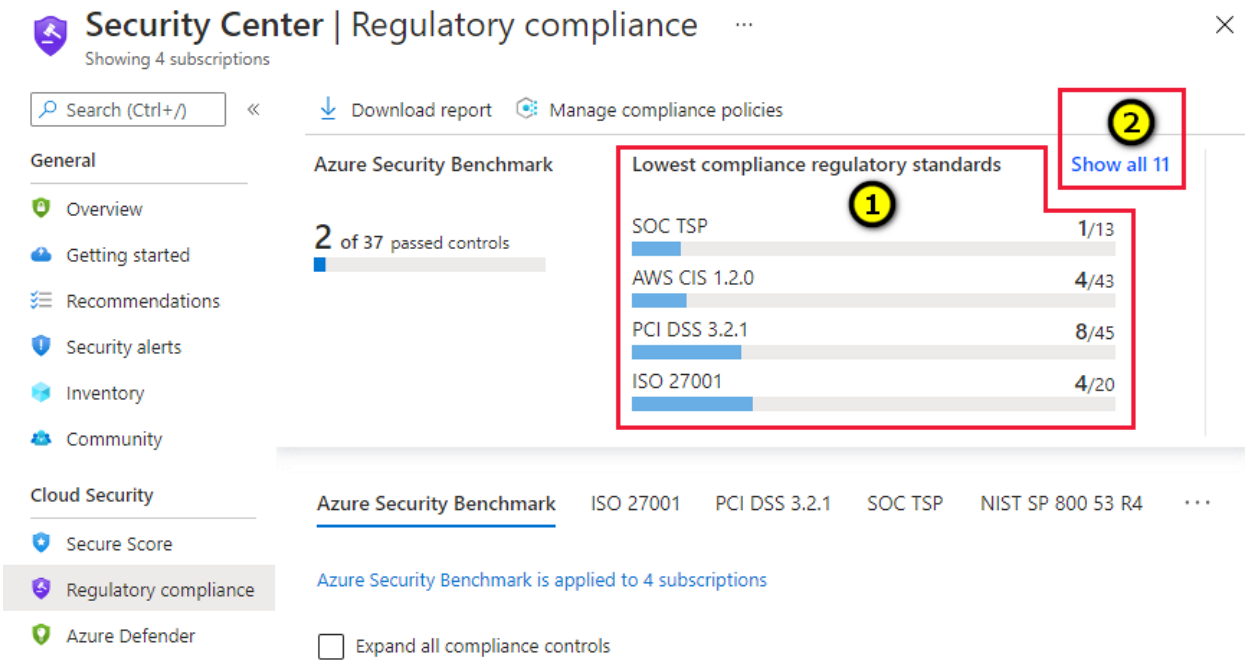

How can I download a report with compliance data in a format other than PDF?

When you select **Download report**, select the standard and the format (PDF or CSV). The resulting report will reflect the current set of subscriptions you've selected in the portal's filter.

- The PDF report shows a summary status for the standard you selected
- The CSV report provides detailed results per resource, as it relates to policies associated with each control

Currently, there's no support for downloading a report for a custom policy; only for the supplied regulatory standards.

How can I create exceptions for some of the policies in the regulatory compliance dashboard?

For policies that are built into Defender for Cloud and included in the secure score, you can create exemptions for one or more resources directly in the portal as explained in Exempting resources and recommendations from your secure score.

For other policies, you can create an exemption directly in the policy itself, by following the instructions in Azure Policy exemption structure.

What Microsoft Defender plans or licenses do I need to use the regulatory compliance dashboard?

If you've got *any* of the Microsoft Defender plan enabled on *any* of your Azure resources, you can access Defender for Cloud's regulatory compliance dashboard and all of its data.

How do I know which benchmark or standard to use?

Azure Security Benchmark (ASB) is the canonical set of security recommendations and best practices defined by Microsoft, aligned with common compliance control frameworks including CIS Microsoft Azure Foundations Benchmark and NIST SP 800-53. ASB is a comprehensive benchmark, and is designed to recommend the most up-to-date security capabilities of a wide range of Azure services. We recommend ASB to customers who want to maximize their security posture and align their compliance status with industry standards.

The CIS Benchmark is authored by an independent entity – Center for Internet Security (CIS) – and contains recommendations on a subset of core Azure services. We work with CIS to try to ensure that their recommendations are up to date with the latest enhancements in Azure, but they do sometimes fall behind and become outdated. Nonetheless, some customers like to use this objective, third-party assessment from CIS as their initial and primary security baseline.

Since we've released the Azure Security Benchmark, many customers have chosen to migrate to it as a replacement for CIS benchmarks.

Next steps

In this tutorial, you learned about using Defender for Cloud's regulatory compliance dashboard to:

- View and monitor your compliance posture regarding the standards and regulations that are important to you.
- Improve your compliance status by resolving relevant recommendations and watching the compliance score improve.

The regulatory compliance dashboard can greatly simplify the compliance process, and significantly cut the time required for gathering compliance evidence for your Azure, hybrid, and multi-cloud environment.

To learn more, see these related pages:

- Customize the set of standards in your regulatory compliance dashboard - Learn how to select which standards appear in your regulatory compliance dashboard.
- Managing security recommendations in Defender for Cloud - Learn how to use recommendations in Defender for Cloud to help protect your Azure resources.

Manage security policies

2/15/2022 • 5 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how security policies are configured, and how to view them in Microsoft Defender for Cloud.

To understand the relationships between initiatives, policies, and recommendations, see What are security policies, initiatives, and recommendations?

Who can edit security policies?

Defender for Cloud uses Azure role-based access control (Azure RBAC), which provides built-in roles you can assign to Azure users, groups, and services. When users open Defender for Cloud, they see only information related to the resources they can access. Which means users are assigned the role of *owner*, *contributor*, or *reader* to the resource's subscription. There are also two specific Defender for Cloud roles:

- **Security reader**: Has rights to view Defender for Cloud items such as recommendations, alerts, policy, and health. Can't make changes.
- **Security admin**: Has the same view rights as *security reader*. Can also update the security policy and dismiss alerts.

You can edit security policies through the Azure Policy portal, via REST API or using Windows PowerShell.

Manage your security policies

To view your security policies in Defender for Cloud:

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the relevant subscription or management group whose policies you want to view.
3. Open the **Security policy** page.
4. The security policy page for that subscription or management group appears. It shows the available and assigned policies.

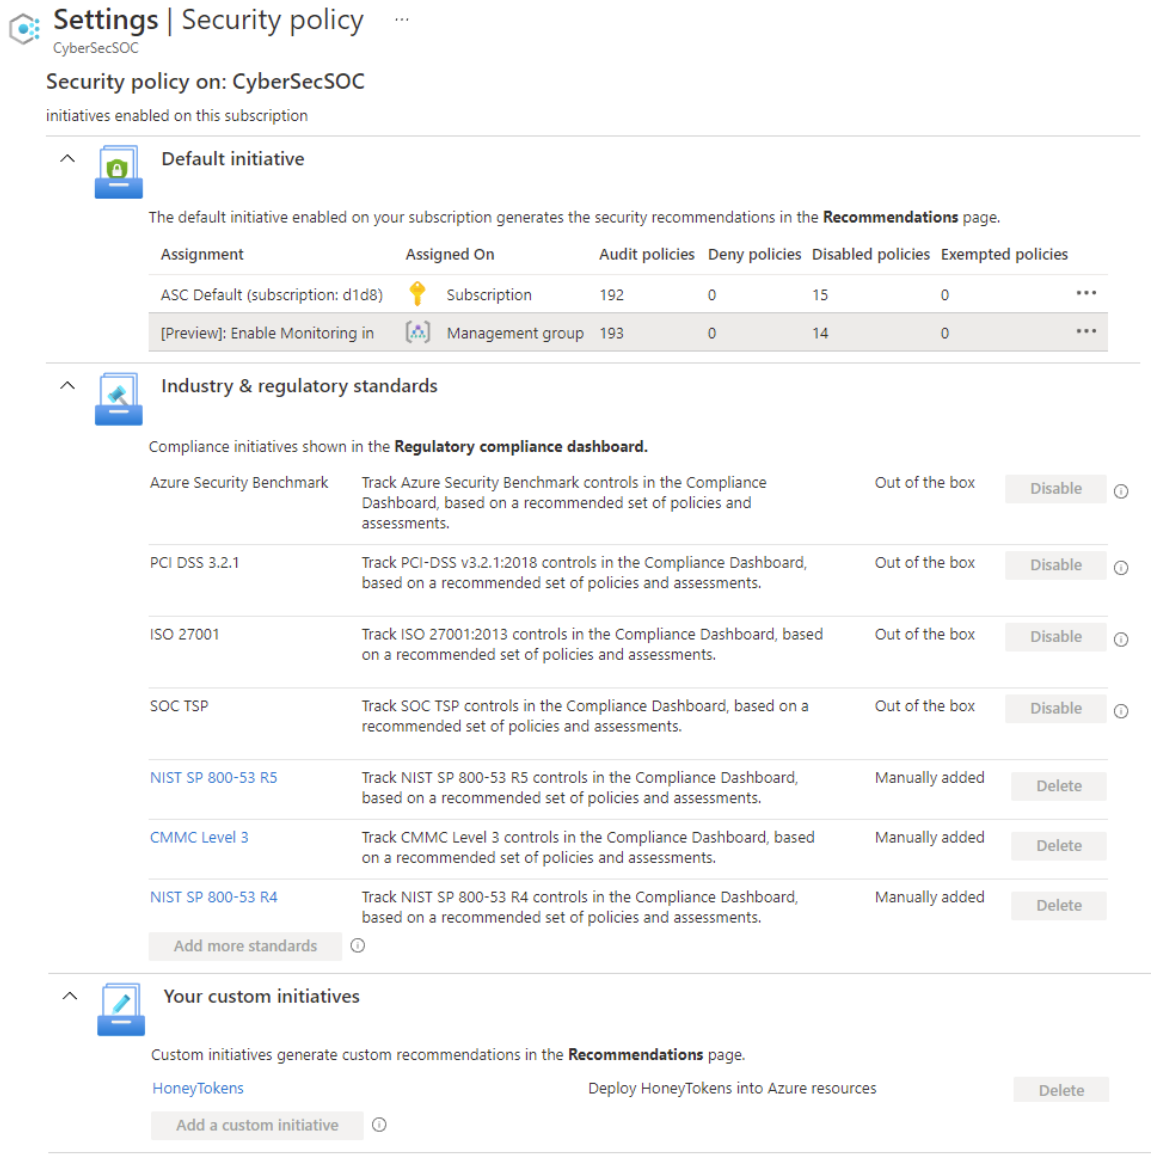

> **NOTE**
>
> If there is a label "MG Inherited" alongside your default initiative, it means that the initiative has been assigned to a management group and inherited by the subscription you're viewing.

5. Choose from the available options on this page:

   a. To work with industry standards, select **Add more standards**. For more information, see Customize the set of standards in your regulatory compliance dashboard.

   b. To assign and manage custom initiatives, select **Add custom initiatives**. For more information, see Using custom security initiatives and policies.

   c. To view and edit the default initiative, select it and proceed as described below.

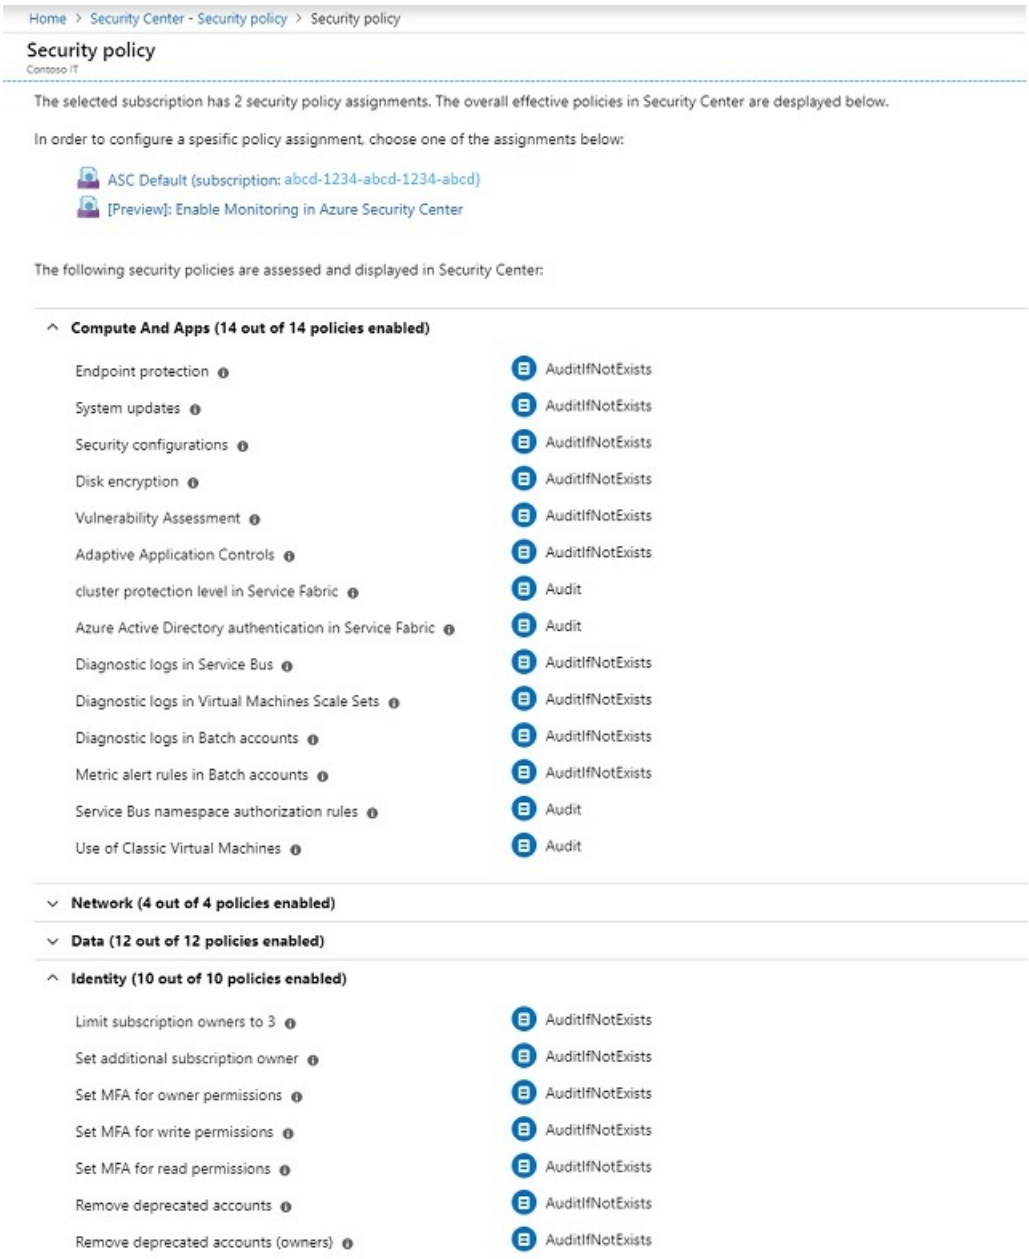

This **Security policy** screen reflects the action taken by the policies assigned on the subscription or management group you selected.

- Use the links at the top to open a policy **assignment** that applies on the subscription or management group. These links let you access the assignment and edit or disable the policy. For example, if you see that a particular policy assignment is effectively denying endpoint protection, use the link to edit or disable the policy.
- In the list of policies, you can see the effective application of the policy on your subscription or management group. The settings of each policy that apply to the scope are taken into consideration and the cumulative outcome of actions taken by the policy is shown. For example, if in one assignment of the policy is disabled, but in another it's set to AuditIfNotExist, then the cumulative effect applies AuditIfNotExist. The more active effect always takes precedence.
- The policies' effect can be: Append, Audit, AuditIfNotExists, Deny, DeployIfNotExists, Disabled. For more information on how effects are applied, see Understand Policy effects.

> **NOTE**
> When you view assigned policies, you can see multiple assignments and you can see how each assignment is configured on its own.

Disable security policies and disable recommendations

When your security initiative triggers a recommendation that's irrelevant for your environment, you can prevent that recommendation from appearing again. To disable a recommendation, disable the specific policy that generates the recommendation.

The recommendation you want to disable will still appear if it's required for a regulatory standard you've applied with Defender for Cloud's regulatory compliance tools. Even if you've disabled a policy in the built-in initiative, a policy in the regulatory standard's initiative will still trigger the recommendation if it's necessary for compliance. You can't disable policies from regulatory standard initiatives.

For more information about recommendations, see Managing security recommendations.

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the subscription or management group for which you want to disable the recommendation (and policy).

   > **NOTE**
   > Remember that a management group applies its policies to its subscriptions. Therefore, if you disable a subscription's policy, and the subscription belongs to a management group that still uses the same policy, then you will continue to receive the policy recommendations. The policy will still be applied from the management level and the recommendations will still be generated.

3. Open the **Security policy** page.
4. From the **Default initiative** or **Your custom initiatives** sections, select the relevant initiative containing the policy you want to disable.
5. Open the **Parameters** section and search for the policy that invokes the recommendation that you want to disable.
6. From the dropdown list, change the value for the corresponding policy to **Disabled**.

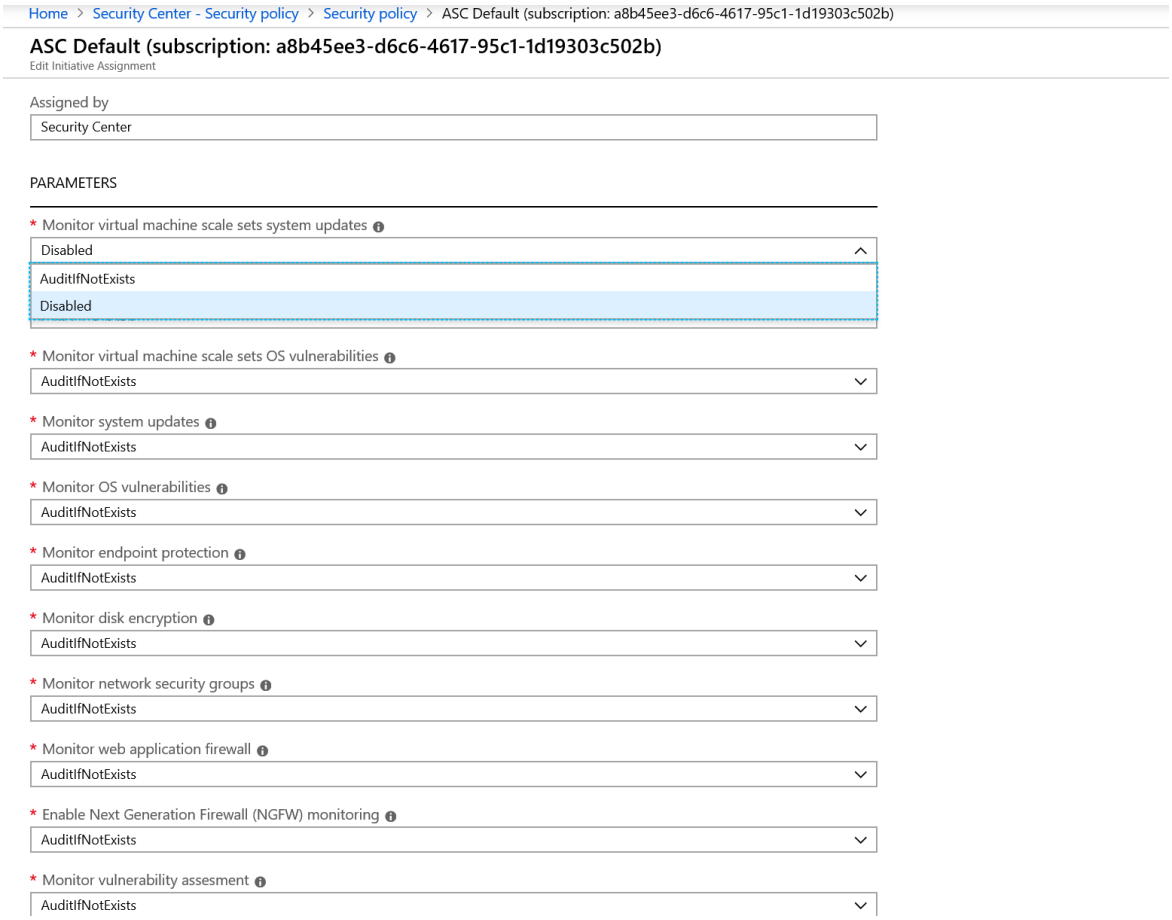

7. Select **Save**.

> **NOTE**
>
> The change might take up to 12 hours to take effect.

Enable a security policy

Some policies in your initiatives might be disabled by default. For example, in the Azure Security Benchmark initiative, some policies are provided for you to enable only if they meet a specific regulatory or compliance requirement for your organization. Such policies include recommendations to encrypt data at rest with customer-managed keys, such as "Container registries should be encrypted with a customer-managed key (CMK)".

To enable a disabled policy and ensure it's assessed for your resources:

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the subscription or management group for which you want to enable the recommendation (and policy).
3. Open the **Security policy** page.
4. From the **Default initiative**, **Industry & regulatory standards**, or **Your custom initiatives** sections, select the relevant initiative with the policy you want to enable.
5. Open the **Parameters** section and search for the policy that invokes the recommendation that you want to disable.
6. From the dropdown list, change the value for the corresponding policy to **AuditIfNotExists** or **Enforce**.

7. Select **Save**.

> **NOTE**
>
> The change might take up to 12 hours to take effect.

Next steps

This page explained security policies. For related information, see the following pages:

- Learn how to set policies using PowerShell
- Learn how to edit a security policy in Azure Policy
- Learn how to set a policy across subscriptions or on Management groups using Azure Policy
- Learn how to enable Defender for Cloud on all subscriptions in a management group

Tutorial: Investigate the health of your resources

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

NOTE

The resource health page described in this tutorial is a preview release.

The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability. | |Pricing:|**Microsoft Defender for DNS** is billed as shown on the pricing page.

The resource health page provides a snapshot view of the overall health of a single resource. You can review detailed information about the resource and all recommendations that apply to that resource. Also, if you're using any of the advanced protection plans of Microsoft Defender for Cloud, you can see outstanding security alerts for that specific resource too.

This single page, currently in preview, in Defender for Cloud's portal pages shows:

1. **Resource information** - The resource group and subscription it's attached to, the geographic location, and more.
2. **Applied security feature** - Whether a Microsoft Defender plan is enabled for the resource.
3. **Counts of outstanding recommendations and alerts** - The number of outstanding security recommendations and Defender for Cloud alerts.
4. **Actionable recommendations and alerts** - Two tabs list the recommendations and alerts that apply to the resource.

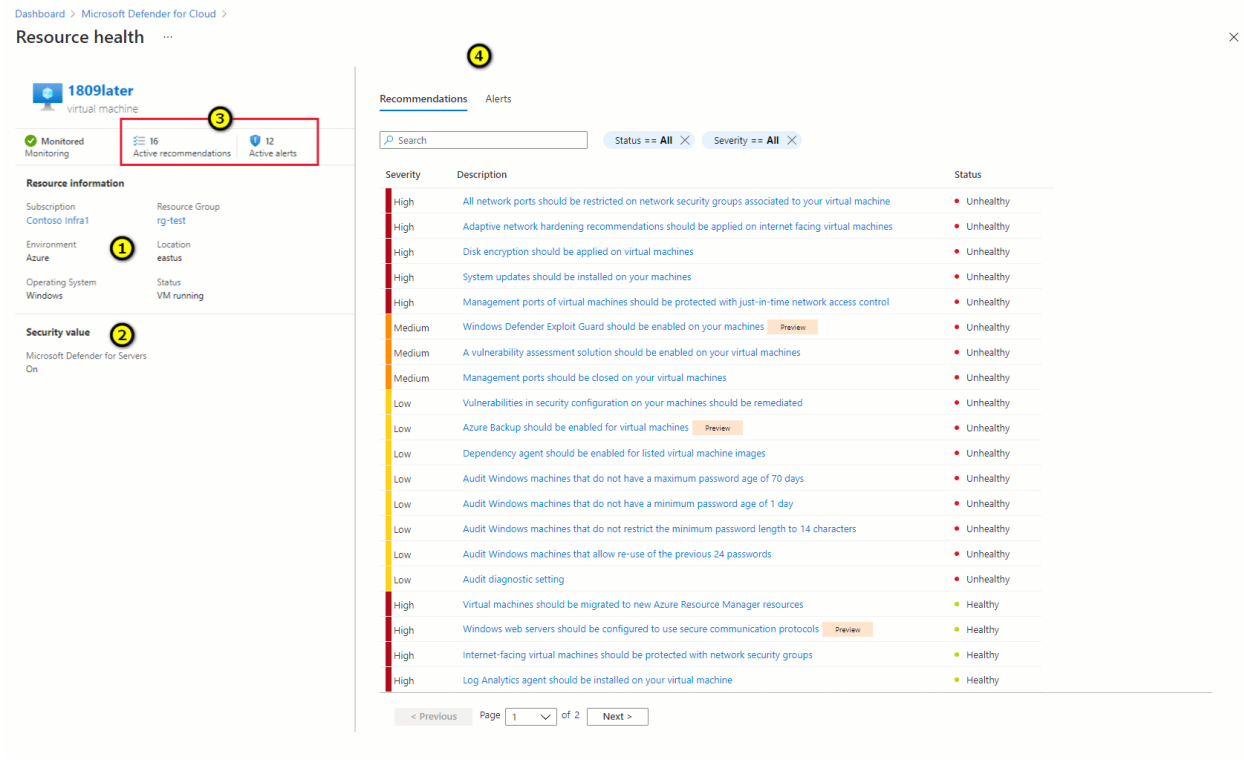

In this tutorial you'll learn how to:

- Access the resource health page for all resource types
- Evaluate the outstanding security issues for a resource
- Improve the security posture for the resource

Prerequisites

To step through the features covered in this tutorial:

- An Azure subscription If you don't have an Azure subscription, create a free account before you begin.
- To apply security recommendations, you must be signed in with an account that has the relevant permissions (Resource Group Contributor, Resource Group Owner, Subscription Contributor, or Subscription Owner)
- To dismiss alerts, you must be signed in with an account that has the relevant permissions (Security Admin, Subscription Contributor, or Subscription Owner)

Access the health information for a resource

> **TIP**
> In the screenshots below, we're opening a virtual machine, but the resource health page can show you the details for all resource types.

To open the resource health page for a resource:

1. Select any resource from the asset inventory page.

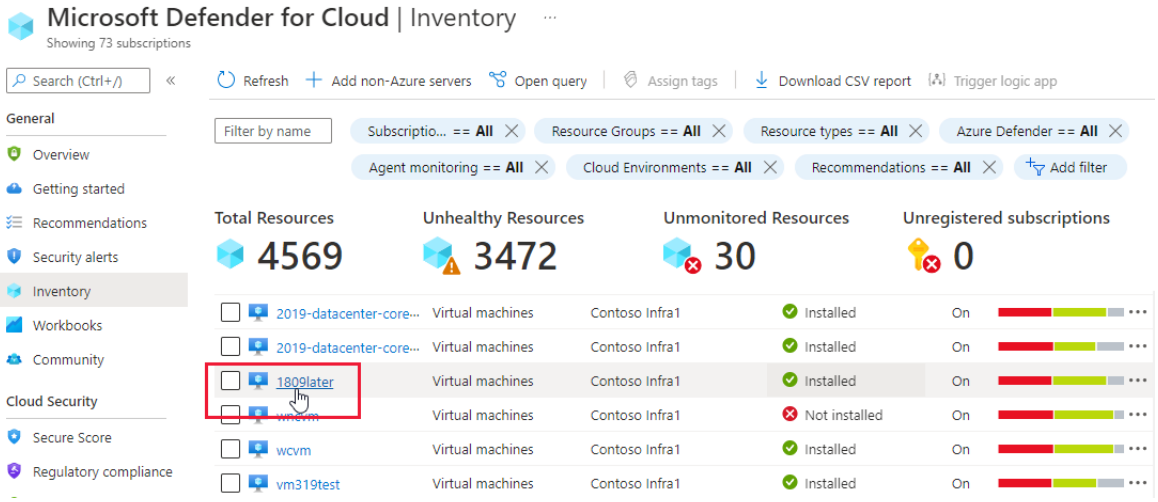

2. Use the left pane of the resource health page for an overview of the subscription, status, and monitoring information about the resource. You can also see whether enhanced security features are enabled for the resource:

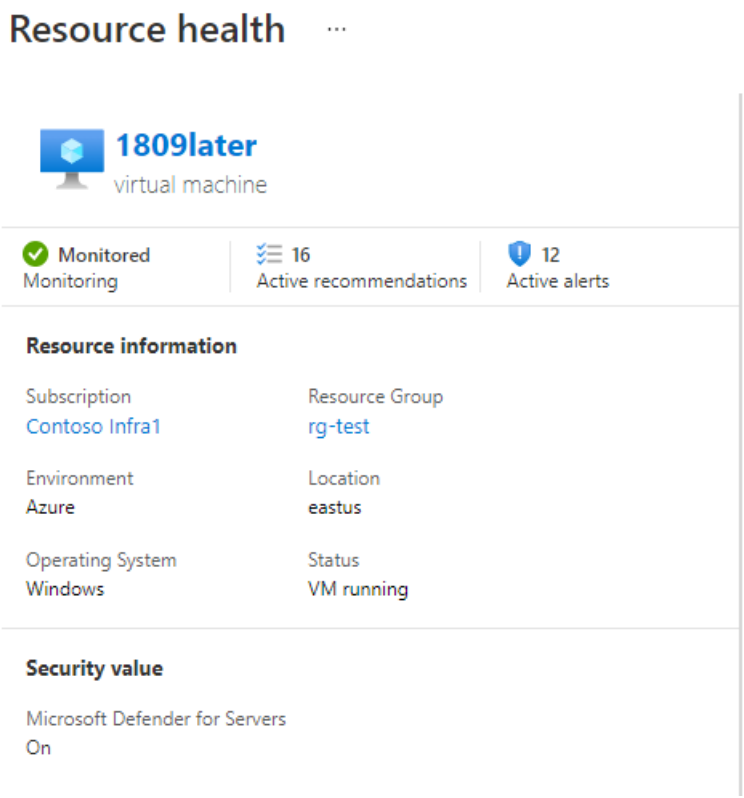

3. Use the two tabs on the right pane to review the lists of security recommendations and alerts that apply to this resource:

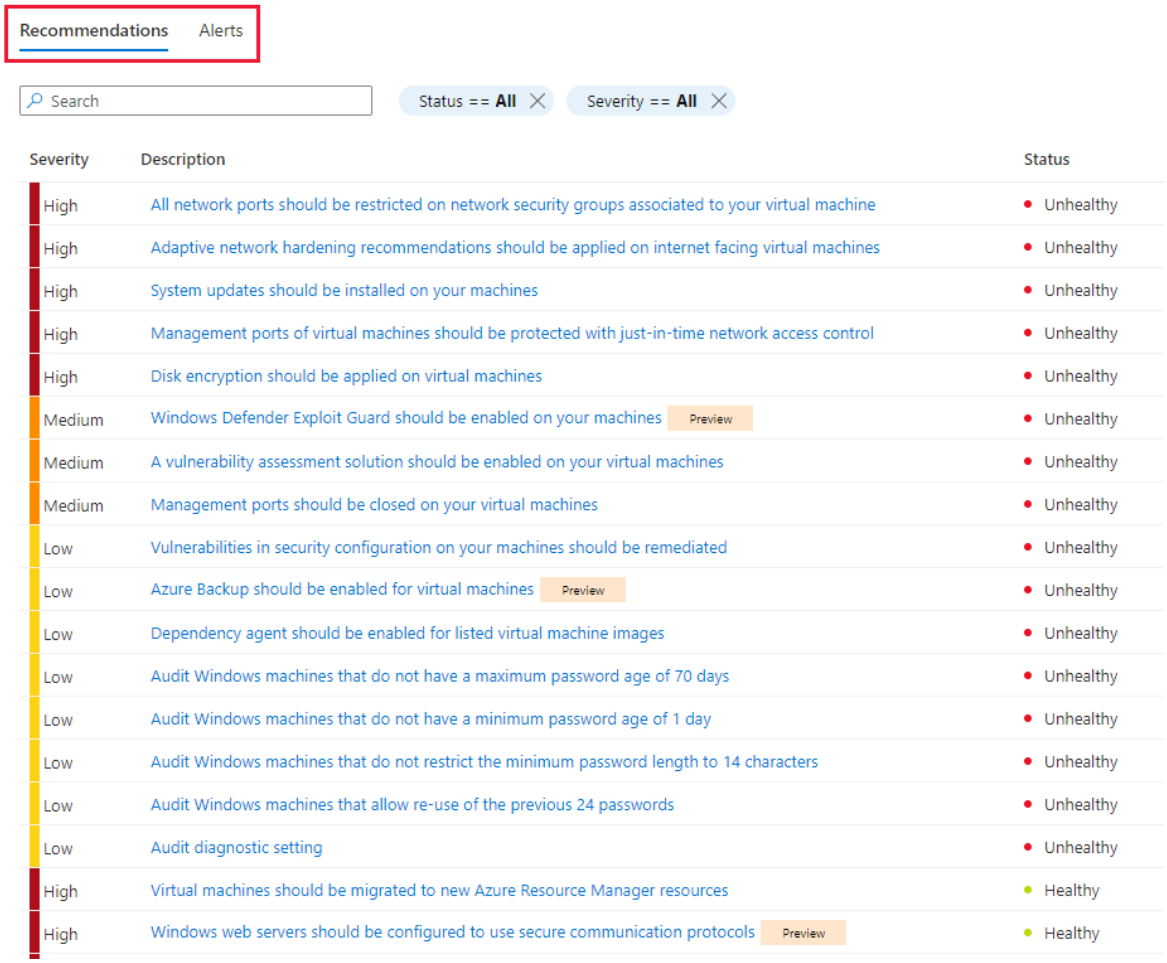

> **NOTE**
>
> Microsoft Defender for Cloud uses the terms "healthy" and "unhealthy" to describe the security status of a resource. These terms relate to whether the resource is compliant with a specific security recommendation.
>
> In the screenshot above, you can see that recommendations are listed even when this resource is "healthy". One advantage of the resource health page is that all recommendations are listed so you can get a complete picture of your resources' health.

Evaluate the outstanding security issues for a resource

The resource health page lists the recommendations for which your resource is "unhealthy" and the alerts that are active.

- To ensure your resource is hardened according to the policies applied to your subscriptions, fix the issues described in the recommendations:

  1. From the right pane, select a recommendation.
  2. Continue as instructed on screen.

  > **TIP**
  >
  > The instructions for fixing issues raised by security recommendations differ for each of Defender for Cloud's recommendations.
  >
  > To decide which recommendations to resolve first, look at the severity of each one and its potential impact on your secure score.

- To investigate a security alert:

    1. From the right pane, select an alert.
    2. Follow the instructions in Respond to security alerts.

Next steps

In this tutorial, you learned about using Defender for Cloud's resource health page.

To learn more, see these related pages:

- Respond to security alerts
- Review your security recommendations

Azure Resource Graph sample queries for Microsoft Defender for Cloud

2/15/2022 • 11 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page is a collection of Azure Resource Graph sample queries for Microsoft Defender for Cloud. For a complete list of Azure Resource Graph samples, see Resource Graph samples by Category and Resource Graph samples by Table.

Sample queries

Controls secure score per subscription

Returns controls secure score per subscription.

```
SecurityResources
| where type == 'microsoft.security/securescores/securescorecontrols'
| extend controlName=properties.displayName,
 controlId=properties.definition.name,
 notApplicableResourceCount=properties.notApplicableResourceCount,
 unhealthyResourceCount=properties.unhealthyResourceCount,
 healthyResourceCount=properties.healthyResourceCount,
 percentageScore=properties.score.percentage,
 currentScore=properties.score.current,
 maxScore=properties.definition.properties.maxScore,
 weight=properties.weight,
 controlType=properties.definition.properties.source.sourceType,
 controlRecommendationIds=properties.definition.properties.assessmentDefinitions
| project tenantId, subscriptionId, controlName, controlId, unhealthyResourceCount, healthyResourceCount,
notApplicableResourceCount, percentageScore, currentScore, maxScore, weight, controlType,
controlRecommendationIds
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/securescores/securescorecontrols' |
extend controlName=properties.displayName, controlId=properties.definition.name,
notApplicableResourceCount=properties.notApplicableResourceCount,
unhealthyResourceCount=properties.unhealthyResourceCount,
healthyResourceCount=properties.healthyResourceCount, percentageScore=properties.score.percentage,
currentScore=properties.score.current, maxScore=properties.definition.properties.maxScore,
weight=properties.weight, controlType=properties.definition.properties.source.sourceType,
controlRecommendationIds=properties.definition.properties.assessmentDefinitions | project tenantId,
subscriptionId, controlName, controlId, unhealthyResourceCount, healthyResourceCount,
notApplicableResourceCount, percentageScore, currentScore, maxScore, weight, controlType,
controlRecommendationIds"
```

Count healthy, unhealthy, and not applicable resources per recommendation

Returns count of healthy, unhealthy, and not applicable resources per recommendation. Use `summarize` and `count` to define how to group and aggregate the values by property.

```
SecurityResources
| where type == 'microsoft.security/assessments'
| extend resourceId=id,
 recommendationId=name,
 resourceType=type,
 recommendationName=properties.displayName,
 source=properties.resourceDetails.Source,
 recommendationState=properties.status.code,
 description=properties.metadata.description,
 assessmentType=properties.metadata.assessmentType,
 remediationDescription=properties.metadata.remediationDescription,
 policyDefinitionId=properties.metadata.policyDefinitionId,
 implementationEffort=properties.metadata.implementationEffort,
 recommendationSeverity=properties.metadata.severity,
 category=properties.metadata.categories,
 userImpact=properties.metadata.userImpact,
 threats=properties.metadata.threats,
 portalLink=properties.links.azurePortal
| summarize numberOfResources=count(resourceId) by tostring(recommendationName),
tostring(recommendationState)
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/assessments' | extend
resourceId=id, recommendationId=name, resourceType=type, recommendationName=properties.displayName,
source=properties.resourceDetails.Source, recommendationState=properties.status.code,
description=properties.metadata.description, assessmentType=properties.metadata.assessmentType,
remediationDescription=properties.metadata.remediationDescription,
policyDefinitionId=properties.metadata.policyDefinitionId,
implementationEffort=properties.metadata.implementationEffort,
recommendationSeverity=properties.metadata.severity, category=properties.metadata.categories,
userImpact=properties.metadata.userImpact, threats=properties.metadata.threats,
portalLink=properties.links.azurePortal | summarize numberOfResources=count(resourceId) by
tostring(recommendationName), tostring(recommendationState)"
```

Get all IoT alerts on hub, filtered by type

Returns all IoT alerts for a specific hub (replace placeholder `{hub_id}`) and alert type (replace placeholder `{alert_type}`).

```
SecurityResources
| where type =~ 'microsoft.security/iotalerts' and id contains '{hub_id}' and properties.alertType contains
'{alert_type}'
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type =~ 'microsoft.security/iotalerts' and id contains
'{hub_id}' and properties.alertType contains '{alert_type}'"
```

Get sensitivity insight of a specific resource

Returns sensitivity insight of a specific resource (replace placeholder {resource_id}).

```
SecurityResources
| where type == 'microsoft.security/insights/classification'
| where properties.associatedResource contains '$resource_id'
| project SensitivityInsight = properties.insightProperties.purviewCatalogs[0].sensitivity
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/insights/classification' | where
properties.associatedResource contains '\$resource_id' | project SensitivityInsight =
properties.insightProperties.purviewCatalogs[0].sensitivity"
```

Get specific IoT alert

Returns specific IoT alert by a provided system alert ID (replace placeholder `{system_Alert_Id}`).

```
SecurityResources
| where type =~ 'microsoft.security/iotalerts' and properties.systemAlertId contains '{system_Alert_Id}'
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type =~ 'microsoft.security/iotalerts' and
properties.systemAlertId contains '{system_Alert_Id}'"
```

List Azure Security Center recommendations

Returns all Azure Security Center assessments, organized in tabular manner with field per property.

```
SecurityResources
| where type == 'microsoft.security/assessments'
| extend resourceId=id,
 recommendationId=name,
 recommendationName=properties.displayName,
 source=properties.resourceDetails.Source,
 recommendationState=properties.status.code,
 description=properties.metadata.description,
 assessmentType=properties.metadata.assessmentType,
 remediationDescription=properties.metadata.remediationDescription,
 policyDefinitionId=properties.metadata.policyDefinitionId,
 implementationEffort=properties.metadata.implementationEffort,
 recommendationSeverity=properties.metadata.severity,
 category=properties.metadata.categories,
 userImpact=properties.metadata.userImpact,
 threats=properties.metadata.threats,
 portalLink=properties.links.azurePortal
| project tenantId, subscriptionId, resourceId, recommendationName, recommendationId, recommendationState,
recommendationSeverity, description, remediationDescription, assessmentType, policyDefinitionId,
implementationEffort, userImpact, category, threats, source, portalLink
```

- Azure CLI

- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/assessments' | extend
resourceId=id, recommendationId=name, recommendationName=properties.displayName,
source=properties.resourceDetails.Source, recommendationState=properties.status.code,
description=properties.metadata.description, assessmentType=properties.metadata.assessmentType,
remediationDescription=properties.metadata.remediationDescription,
policyDefinitionId=properties.metadata.policyDefinitionId,
implementationEffort=properties.metadata.implementationEffort,
recommendationSeverity=properties.metadata.severity, category=properties.metadata.categories,
userImpact=properties.metadata.userImpact, threats=properties.metadata.threats,
portalLink=properties.links.azurePortal | project tenantId, subscriptionId, resourceId, recommendationName,
recommendationId, recommendationState, recommendationSeverity, description, remediationDescription,
assessmentType, policyDefinitionId, implementationEffort, userImpact, category, threats, source, portalLink"
```

List Container Registry vulnerability assessment results

Returns all the all the vulnerabilities found on container images. Azure Defender for Containers has to be enabled in order to view these security findings.

```
SecurityResources
| where type == 'microsoft.security/assessments'
| where properties.displayName contains 'Vulnerabilities in Azure Container Registry images should be
remediated'
| summarize by assessmentKey=name //the ID of the assessment
| join kind=inner (
 securityresources
 | where type == 'microsoft.security/assessments/subassessments'
 | extend assessmentKey = extract('.*assessments/(.+?)/.*',1,  id)
) on assessmentKey
| project assessmentKey, subassessmentKey=name, id, parse_json(properties), resourceGroup, subscriptionId,
tenantId
| extend description = properties.description,
 displayName = properties.displayName,
 resourceId = properties.resourceDetails.id,
 resourceSource = properties.resourceDetails.source,
 category = properties.category,
 severity = properties.status.severity,
 code = properties.status.code,
 timeGenerated = properties.timeGenerated,
 remediation = properties.remediation,
 impact = properties.impact,
 vulnId = properties.id,
 additionalData = properties.additionalData
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/assessments' | where
properties.displayName contains 'Vulnerabilities in Azure Container Registry images should be remediated' |
summarize by assessmentKey=name //the ID of the assessment | join kind=inner ( securityresources | where
type == 'microsoft.security/assessments/subassessments' | extend assessmentKey =
extract('.*assessments/(.+?)/.*',1, id) ) on assessmentKey | project assessmentKey, subassessmentKey=name,
id, parse_json(properties), resourceGroup, subscriptionId, tenantId | extend description =
properties.description, displayName = properties.displayName, resourceId = properties.resourceDetails.id,
resourceSource = properties.resourceDetails.source, category = properties.category, severity =
properties.status.severity, code = properties.status.code, timeGenerated = properties.timeGenerated,
remediation = properties.remediation, impact = properties.impact, vulnId = properties.id, additionalData =
properties.additionalData"
```

List Qualys vulnerability assessment results

Returns all the vulnerabilities found on virtual machines that have a Qualys agent installed.

```
SecurityResources
| where type == 'microsoft.security/assessments'
| where * contains 'vulnerabilities in your virtual machines'
| summarize by assessmentKey=name //the ID of the assessment
| join kind=inner (
 securityresources
 | where type == 'microsoft.security/assessments/subassessments'
 | extend assessmentKey = extract('.*assessments/(.+?)/.*',1,  id)
) on assessmentKey
| project assessmentKey, subassessmentKey=name, id, parse_json(properties), resourceGroup, subscriptionId,
tenantId
| extend description = properties.description,
 displayName = properties.displayName,
 resourceId = properties.resourceDetails.id,
 resourceSource = properties.resourceDetails.source,
 category = properties.category,
 severity = properties.status.severity,
 code = properties.status.code,
 timeGenerated = properties.timeGenerated,
 remediation = properties.remediation,
 impact = properties.impact,
 vulnId = properties.id,
 additionalData = properties.additionalData
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/assessments' | where * contains
'vulnerabilities in your virtual machines' | summarize by assessmentKey=name //the ID of the assessment |
join kind=inner ( securityresources | where type == 'microsoft.security/assessments/subassessments' | extend
assessmentKey = extract('.*assessments/(.+?)/.*',1, id) ) on assessmentKey | project assessmentKey,
subassessmentKey=name, id, parse_json(properties), resourceGroup, subscriptionId, tenantId | extend
description = properties.description, displayName = properties.displayName, resourceId =
properties.resourceDetails.id, resourceSource = properties.resourceDetails.source, category =
properties.category, severity = properties.status.severity, code = properties.status.code, timeGenerated =
properties.timeGenerated, remediation = properties.remediation, impact = properties.impact, vulnId =
properties.id, additionalData = properties.additionalData"
```

Regulatory compliance assessments state

Returns regulatory compliance assessments state per compliance standard and control.

```
SecurityResources
| where type ==
'microsoft.security/regulatorycompliancestandards/regulatorycompliancecontrols/regulatorycomplianceassessmen
ts'
| extend assessmentName=properties.description,
 complianceStandard=extract(@'/regulatoryComplianceStandards/(.+)/regulatoryComplianceControls',1,id),
 complianceControl=extract(@'/regulatoryComplianceControls/(.+)/regulatoryComplianceAssessments',1,id),
 skippedResources=properties.skippedResources,
 passedResources=properties.passedResources,
 failedResources=properties.failedResources,
 state=properties.state
| project tenantId, subscriptionId, id, complianceStandard, complianceControl, assessmentName, state,
skippedResources, passedResources, failedResources
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type ==
'microsoft.security/regulatorycompliancestandards/regulatorycompliancecontrols/regulatorycomplianceassessmen
ts' | extend assessmentName=properties.description,
complianceStandard=extract(@'/regulatoryComplianceStandards/(.+)/regulatoryComplianceControls',1,id),
complianceControl=extract(@'/regulatoryComplianceControls/(.+)/regulatoryComplianceAssessments',1,id),
skippedResources=properties.skippedResources, passedResources=properties.passedResources,
failedResources=properties.failedResources, state=properties.state | project tenantId, subscriptionId, id,
complianceStandard, complianceControl, assessmentName, state, skippedResources, passedResources,
failedResources"
```

Regulatory compliance state per compliance standard

Returns regulatory compliance state per compliance standard per subscription.

```
SecurityResources
| where type == 'microsoft.security/regulatorycompliancestandards'
| extend complianceStandard=name,
 state=properties.state,
 passedControls=properties.passedControls,
 failedControls=properties.failedControls,
 skippedControls=properties.skippedControls,
 unsupportedControls=properties.unsupportedControls
| project tenantId, subscriptionId, complianceStandard, state, passedControls, failedControls,
skippedControls, unsupportedControls
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/regulatorycompliancestandards' |
extend complianceStandard=name, state=properties.state, passedControls=properties.passedControls,
failedControls=properties.failedControls, skippedControls=properties.skippedControls,
unsupportedControls=properties.unsupportedControls | project tenantId, subscriptionId, complianceStandard,
state, passedControls, failedControls, skippedControls, unsupportedControls"
```

Secure score per management group

Returns secure score per management group.

```
SecurityResources
| where type == 'microsoft.security/securescores'
| project subscriptionId,
 subscriptionTotal = iff(properties.score.max == 0, 0.00, round(tolong(properties.weight) *
todouble(properties.score.current)/tolong(properties.score.max),2)),
 weight = tolong(iff(properties.weight == 0, 1, properties.weight))
| join kind=leftouter (
 ResourceContainers
 | where type == 'microsoft.resources/subscriptions' and properties.state == 'Enabled'
 | project subscriptionId, mgChain=properties.managementGroupAncestorsChain )
 on subscriptionId
| mv-expand mg=mgChain
| summarize sumSubs = sum(subscriptionTotal), sumWeight = sum(weight), resultsNum = count() by
tostring(mg.displayName), mgId = tostring(mg.name)
| extend secureScore = iff(tolong(resultsNum) == 0, 404.00, round(sumSubs/sumWeight*100,2))
| project mgName=mg_displayName, mgId, sumSubs, sumWeight, resultsNum, secureScore
| order by mgName asc
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/securescores' | project
subscriptionId, subscriptionTotal = iff(properties.score.max == 0, 0.00, round(tolong(properties.weight) *
todouble(properties.score.current)/tolong(properties.score.max),2)), weight = tolong(iff(properties.weight
== 0, 1, properties.weight)) | join kind=leftouter ( ResourceContainers | where type ==
'microsoft.resources/subscriptions' and properties.state == 'Enabled' | project subscriptionId,
mgChain=properties.managementGroupAncestorsChain ) on subscriptionId | mv-expand mg=mgChain | summarize
sumSubs = sum(subscriptionTotal), sumWeight = sum(weight), resultsNum = count() by tostring(mg.displayName),
mgId = tostring(mg.name) | extend secureScore = iff(tolong(resultsNum) == 0, 404.00,
round(sumSubs/sumWeight*100,2)) | project mgName=mg_displayName, mgId, sumSubs, sumWeight, resultsNum,
secureScore | order by mgName asc"
```

Secure score per subscription

Returns secure score per subscription.

```
SecurityResources
| where type == 'microsoft.security/securescores'
| extend percentageScore=properties.score.percentage,
 currentScore=properties.score.current,
 maxScore=properties.score.max,
 weight=properties.weight
| project tenantId, subscriptionId, percentageScore, currentScore, maxScore, weight
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/securescores' | extend
percentageScore=properties.score.percentage, currentScore=properties.score.current,
maxScore=properties.score.max, weight=properties.weight | project tenantId, subscriptionId, percentageScore,
currentScore, maxScore, weight"
```

Show Azure Defender pricing tier per subscription

Returns Azure Defender pricing tier plan per subscription.

```
SecurityResources
| where type == 'microsoft.security/pricings'
| project Subscription= subscriptionId, Azure_Defender_plan= name, Status= properties.pricingTier
```

- Azure CLI
- Azure PowerShell
- Portal

```
az graph query -q "SecurityResources | where type == 'microsoft.security/pricings' | project Subscription=
subscriptionId, Azure_Defender_plan= name, Status= properties.pricingTier"
```

Next steps

- Learn more about the query language.
- Learn more about how to explore resources.
- See samples of Starter language queries.
- See samples of Advanced language queries.

What are security policies, initiatives, and recommendations?

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud applies security initiatives to your subscriptions. These initiatives contain one or more security policies. Each of those policies results in a security recommendation for improving your security posture. This page explains each of these ideas in detail.

What is a security policy?

An Azure Policy definition, created in Azure Policy, is a rule about specific security conditions that you want controlled. Built in definitions include things like controlling what type of resources can be deployed or enforcing the use of tags on all resources. You can also create your own custom policy definitions.

To implement these policy definitions (whether built-in or custom), you'll need to assign them. You can assign any of these policies through the Azure portal, PowerShell, or Azure CLI. Policies can be disabled or enabled from Azure Policy.

There are different types of policies in Azure Policy. Defender for Cloud mainly uses 'Audit' policies that check specific conditions and configurations then report on compliance. There are also "Enforce' policies that can be used to apply secure settings.

What is a security initiative?

An Azure Policy initiative is a collection of Azure Policy definitions, or rules, that are grouped together towards a specific goal or purpose. Azure initiatives simplify management of your policies by grouping a set of policies together, logically, as a single item.

A security initiative defines the desired configuration of your workloads and helps ensure you're complying with the security requirements of your company or regulators.

Like security policies, Defender for Cloud initiatives are also created in Azure Policy. You can use Azure Policy to manage your policies, build initiatives, and assign initiatives to multiple subscriptions or for entire management groups.

The default initiative automatically assigned to every subscription in Microsoft Defender for Cloud is Azure Security Benchmark. This benchmark is the Microsoft-authored, Azure-specific set of guidelines for security and compliance best practices based on common compliance frameworks. This widely respected benchmark builds on the controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security. Learn more about Azure Security Benchmark.

Defender for Cloud offers the following options for working with security initiatives and policies:

- **View and edit the built-in default initiative** - When you enable Defender for Cloud, the initiative named 'Azure Security Benchmark' is automatically assigned to all Defender for Cloud registered

subscriptions. To customize this initiative, you can enable or disable individual policies within it by editing a policy's parameters. See the list of built-in security policies to understand the options available out-of-the-box.

- **Add your own custom initiatives** - If you want to customize the security initiatives applied to your subscription, you can do so within Defender for Cloud. You'll then receive recommendations if your machines don't follow the policies you create. For instructions on building and assigning custom policies, see Using custom security initiatives and policies.
- **Add regulatory compliance standards as initiatives** - Defender for Cloud's regulatory compliance dashboard shows the status of all the assessments within your environment in the context of a particular standard or regulation (such as Azure CIS, NIST SP 800-53 R4, SWIFT CSP CSCF-v2020). For more information, see Improve your regulatory compliance.

What is a security recommendation?

Using the policies, Defender for Cloud periodically analyzes the compliance status of your resources to identify potential security misconfigurations and weaknesses. It then provides you with recommendations on how to remediate those issues. Recommendations are the result of assessing your resources against the relevant policies and identifying resources that aren't meeting your defined requirements.

Defender for Cloud makes its security recommendations based on your chosen initiatives. When a policy from your initiative is compared against your resources and finds one or more that aren't compliant it is presented as a recommendation in Defender for Cloud.

Recommendations are actions for you to take to secure and harden your resources. Each recommendation provides you with the following information:

- A short description of the issue
- The remediation steps to carry out in order to implement the recommendation
- The affected resources

In practice, it works like this:

1. Azure Security Benchmark is an ***initiative*** that contains requirements.

   For example, Azure Storage accounts must restrict network access to reduce their attack surface.

2. The initiative includes multiple ***policies***, each with a requirement of a specific resource type. These policies enforce the requirements in the initiative.

   To continue the example, the storage requirement is enforced with the policy "Storage accounts should restrict network access using virtual network rules".

3. Microsoft Defender for Cloud continually assesses your connected subscriptions. If it finds a resource that doesn't satisfy a policy, it displays a ***recommendation*** to fix that situation and harden the security of resources that aren't meeting your security requirements.

   So, for example, if an Azure Storage account on any of your protected subscriptions isn't protected with virtual network rules, you'll see the recommendation to harden those resources.

So, (1) an initiative includes (2) policies that generate (3) environment-specific recommendations.

Viewing the relationship between a recommendation and a policy

As mentioned above, Defender for Cloud's built in recommendations are based on the Azure Security Benchmark. Almost every recommendation has an underlying policy that is derived from a requirement in the benchmark.

When you're reviewing the details of a recommendation, it's often helpful to be able to see the underlying policy. For every recommendation supported by a policy, use the **View policy definition** link from the recommendation details page to go directly to the Azure Policy entry for the relevant policy:

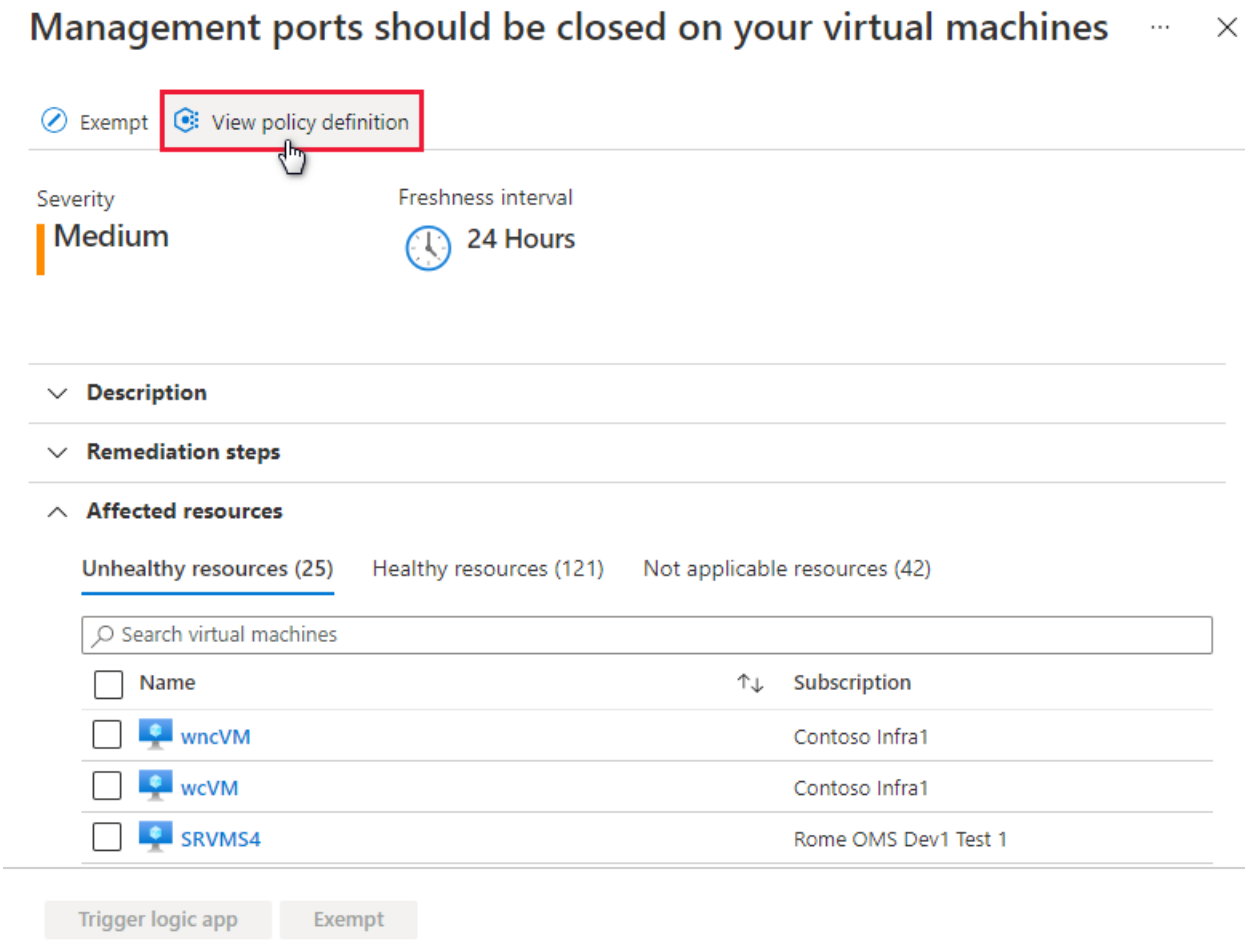

Use this link to view the policy definition and review the evaluation logic.

If you're reviewing the list of recommendations on our Security recommendations reference guide, you'll also see links to the policy definition pages:

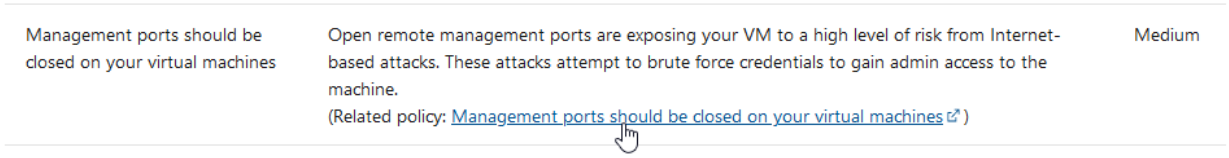

Next steps

This page explained, at a high level, the basic concepts and relationships between policies, initiatives, and recommendations. For related information, see:

- Create custom initiatives
- Disable security policies to disable recommendations
- Learn how to edit a security policy in Azure Policy

Secure score in Microsoft Defender for Cloud

2/15/2022 • 24 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Introduction to secure score

Microsoft Defender for Cloud has two main goals:

- to help you understand your current security situation
- to help you efficiently and effectively improve your security

The central feature in Defender for Cloud that enables you to achieve those goals is **secure score**.

Defender for Cloud continually assesses your resources, subscriptions, and organization for security issues. It then aggregates all the findings into a single score so that you can tell, at a glance, your current security situation: the higher the score, the lower the identified risk level.

The secure score is shown in the Azure portal pages as a percentage value, but the underlying values are also clearly presented:

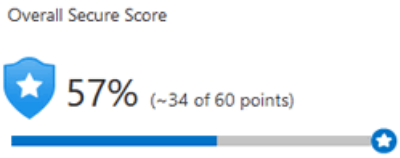

To increase your security, review Defender for Cloud's recommendations page for the outstanding actions necessary to raise your score. Each recommendation includes instructions to help you remediate the specific issue.

Recommendations are grouped into **security controls**. Each control is a logical group of related security recommendations, and reflects your vulnerable attack surfaces. Your score only improves when you remediate *all* of the recommendations for a single resource within a control. To see how well your organization is securing each individual attack surface, review the scores for each security control.

For more information, see How your secure score is calculated below.

How your secure score is calculated

The contribution of each security control towards the overall secure score is shown clearly on the recommendations page.

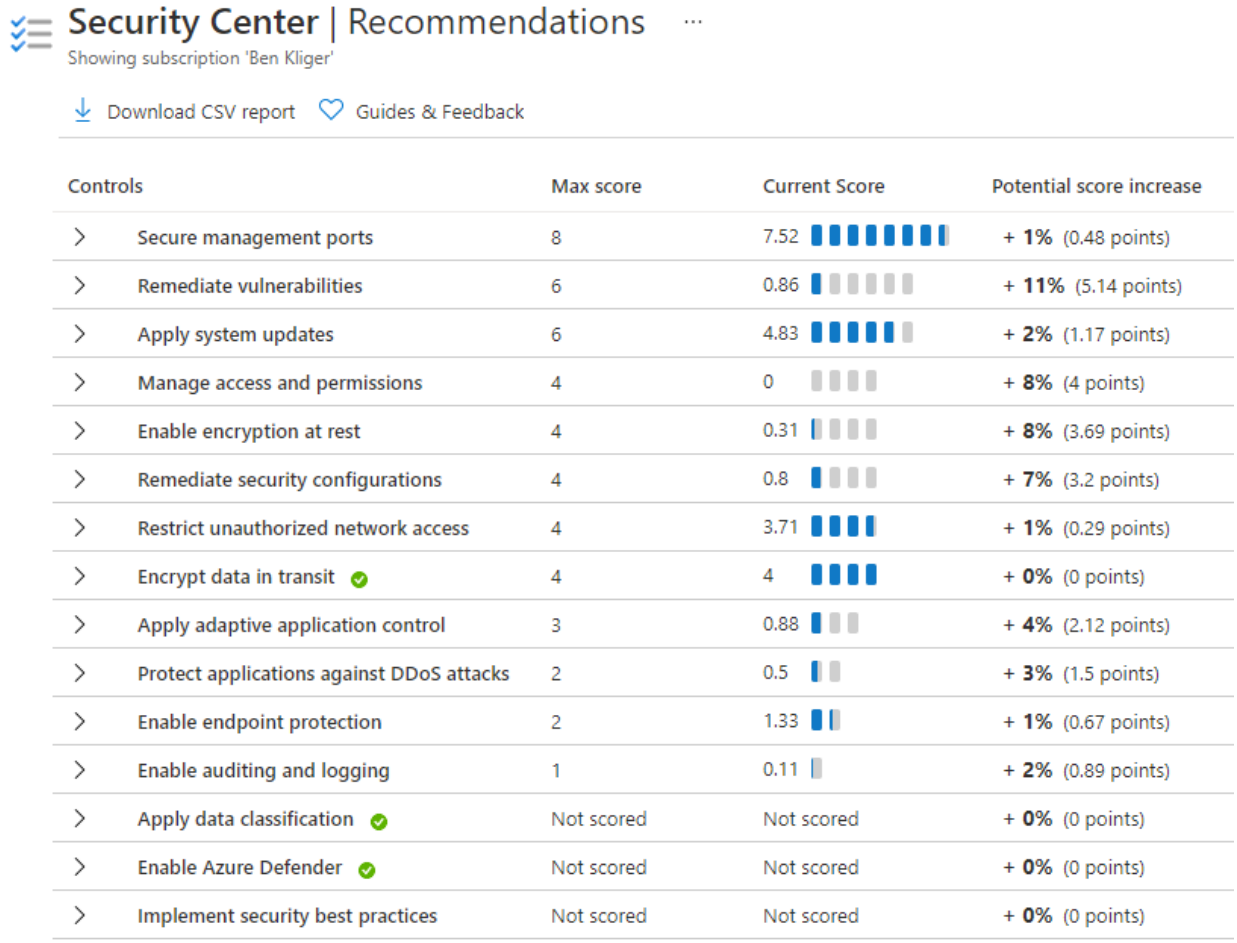

To get all the possible points for a security control, all your resources must comply with all of the security recommendations within the security control. For example, Defender for Cloud has multiple recommendations regarding how to secure your management ports. You'll need to remediate them all to make a difference to your secure score.

Example scores for a control

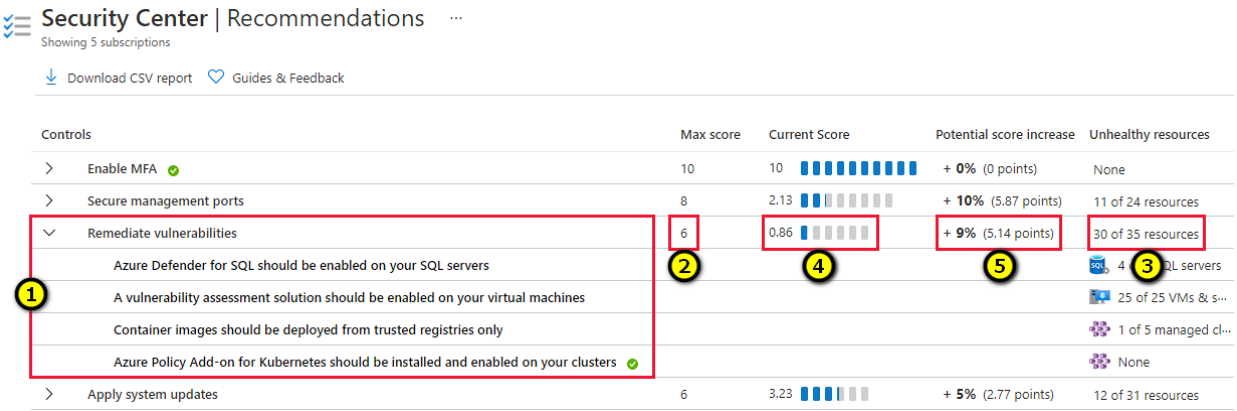

In this example:

#	NAME	DESCRIPTION
1	**Remediate vulnerabilities security control**	This control groups multiple recommendations related to discovering and resolving known vulnerabilities.

#	NAME	DESCRIPTION
2	**Max score**	The maximum number of points you can gain by completing all recommendations within a control. The maximum score for a control indicates the relative significance of that control and is fixed for every environment. Use the max score values to triage the issues to work on first. For a list of all controls and their max scores, see Security controls and their recommendations.
3	**Number of resources**	There are 35 resources affected by this control. To understand the possible contribution of every resource, divide the max score by the number of resources. For this example, 6/35=0.1714 **Every resource contributes 0.1714 points.**
4	**Current score**	The current score for this control. Current score=[Score per resource]*[Number of healthy resources] 0.1714 x 5 healthy resources = 0.86 Each control contributes towards the total score. In this example, the control is contributing 0.86 points to current total secure score.
5	**Potential score increase**	The remaining points available to you within the control. If you remediate all the recommendations in this control, your score will increase by 9%. Potential score increase=[Score per resource]*[Number of unhealthy resources] 0.1714 x 30 unhealthy resources = 5.14

Calculations - understanding your score

METRIC	FORMULA AND EXAMPLE

METRIC	FORMULA AND EXAMPLE
Security control's current score	$$Secure\ score\ for\ a\ single\ security\ control = \frac{Max\ score}{Healthy + Unhealthy}\ x\ Healthy$$ Each individual security control contributes towards the Security Score. Each resource affected by a recommendation within the control, contributes towards the control's current score. The current score for each control is a measure of the status of the resources *within* the control. 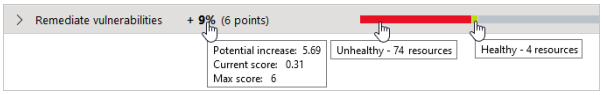 In this example, the max score of 6 would be divided by 78 because that's the sum of the healthy and unhealthy resources. 6 / 78 = 0.0769 Multiplying that by the number of healthy resources (4) results in the current score: 0.0769 * 4 = **0.31**
Secure score Single subscription	$$Secure\ score\ for\ a\ subscription = \frac{\sum current\ scores\ for\ all\ controls}{\sum maximum\ scores\ for\ all\ controls}\ x\ 100$$ **Subscription** ↑↓ **Secure score** ASC DEMO — 47% (28 of 60) In this example, there is a single subscription with all security controls available (a potential maximum score of 60 points). The score shows 28 points out of a possible 60 and the remaining 32 points are reflected in the "Potential score increase" figures of the security controls. **Controls** — **Potential score increase** Remediate vulnerabilities — + **9%** (6 points) Secure management ports — + **9%** (5 points) Enable encryption at rest — + **6%** (3 points) Restrict unauthorized network access — + **4%** (2 points) Enable DDoS protection on Vnet — + **3%** (2 points) Apply system updates — + **3%** (2 points) Manage access and permissions — + **3%** (2 points) Remediate security configurations — + **3%** (2 points) Apply data classification — + **2%** (1 point) Encrypt data in transit — + **2%** (1 point) Adaptive application control — + **1%** (1 point) Enable auditing and logging — + **1%** (1 point) Enable endpoint protection — + **1%** (1 point) Enable MFA — *Completed* — + **0%** (0 points) Additional best practices — + **0%** (0 points)

METRIC	FORMULA AND EXAMPLE
Secure score Multiple subscriptions	$$Secure\ score\ multiple\ subscriptions = \frac{\sum (subscription\ score\ x\ subscription\ weight)}{\sum Weights\ for\ all\ subscriptions}\ x\ 100$$ When calculating the combined score for multiple subscriptions, Defender for Cloud includes a *weight* for each subscription. The relative weights for your subscriptions are determined by Defender for Cloud based on factors such as the number of resources. The current score for each subscription is calculated in the same way as for a single subscription, but then the weight is applied as shown in the equation. When viewing multiple subscriptions, secure score evaluates all resources within all enabled policies and groups their combined impact on each security control's maximum score. 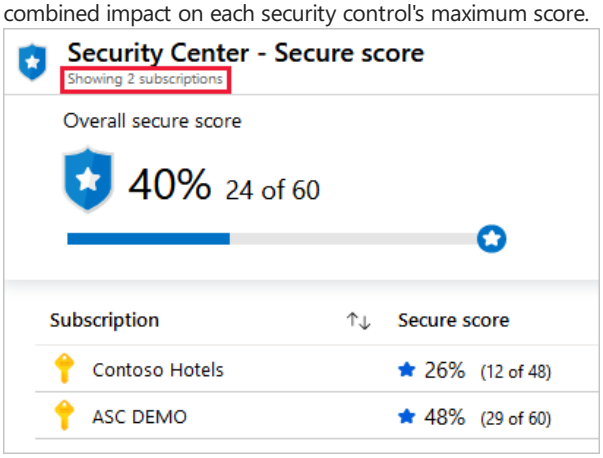 The combined score is **not** an average; rather it's the evaluated posture of the status of all resources across all subscriptions. Here too, if you go to the recommendations page and add up the potential points available, you will find that it's the difference between the current score (24) and the maximum score available (60).

Which recommendations are included in the secure score calculations?

Only built-in recommendations have an impact on the secure score.

Recommendations flagged as **Preview** aren't included in the calculations of your secure score. They should still be remediated wherever possible, so that when the preview period ends they'll contribute towards your score.

An example of a preview recommendation:

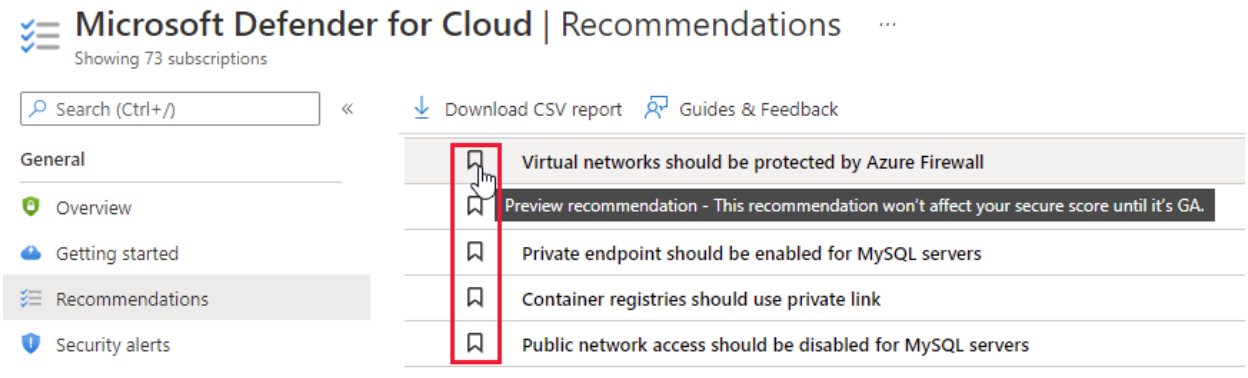

Improve your secure score

To improve your secure score, remediate security recommendations from your recommendations list. You can

remediate each recommendation manually for each resource, or by using the **Fix** option (when available) to resolve an issue on multiple resources quickly. For more information, see Remediate recommendations.

Another way to improve your score and ensure your users don't create resources that negatively impact your score is to configure the Enforce and Deny options on the relevant recommendations. Learn more in Prevent misconfigurations with Enforce/Deny recommendations.

Security controls and their recommendations

The table below lists the security controls in Microsoft Defender for Cloud. For each control, you can see the maximum number of points you can add to your secure score if you remediate *all* of the recommendations listed in the control, for *all* of your resources.

The set of security recommendations provided with Defender for Cloud is tailored to the available resources in each organization's environment. The recommendations can be further customized by disabling policies and exempting specific resources from a recommendation.

We recommend every organization carefully review their assigned Azure Policy initiatives.

> **TIP**
> For details of reviewing and editing your initiatives, see Working with security policies.

Even though Defender for Cloud's default security initiative is based on industry best practices and standards, there are scenarios in which the built-in recommendations listed below might not completely fit your organization. Consequently, it'll sometimes be necessary to adjust the default initiative - without compromising security - to ensure it's aligned with your organization's own policies, industry standards, regulatory standards, and benchmarks you're obligated to meet.

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
10	**Enable MFA** - Defender for Cloud places a high value on multi-factor authentication (MFA). Use these recommendations to secure the users of your subscriptions. There are three ways to enable MFA and be compliant with the recommendations: security defaults, per-user assignment, conditional access policy. Learn more about these options in Manage MFA enforcement on your subscriptions.	- MFA should be enabled on accounts with owner permissions on subscriptions - MFA should be enabled on accounts with owner permissions on your subscription - MFA should be enabled on accounts with write permissions on subscriptions - MFA should be enabled on accounts with write permissions on your subscription
8	**Secure management ports** - Brute force attacks often target management ports. Use these recommendations to reduce your exposure with tools like just-in-time VM access and network security groups.	- Internet-facing virtual machines should be protected with network security groups - Management ports of virtual machines should be protected with just-in-time network access control - Management ports should be closed on your virtual machines

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
6	**Apply system updates** - Not applying updates leaves unpatched vulnerabilities and results in environments that are susceptible to attacks. Use these recommendations to maintain operational efficiency, reduce security vulnerabilities, and provide a more stable environment for your end users. To deploy system updates, you can use the Update Management solution to manage patches and updates for your machines.	- Log Analytics agent should be installed on Linux-based Azure Arc-enabled machines - Log Analytics agent should be installed on virtual machine scale sets - Log Analytics agent should be installed on virtual machines - Log Analytics agent should be installed on Windows-based Azure Arc-enabled machines - System updates on virtual machine scale sets should be installed - System updates should be installed on your machines - System updates should be installed on your machines (powered by Update Center)
6	**Remediate vulnerabilities** - Defender for Cloud includes multiple vulnerability assessment scanners to check your machines, databases, and container registries for weaknesses that threat actors might leverage. Use these recommendations to enable these scanners and review their findings. Learn more about scanning machines, SQL servers, and container registries.	- [Preview] Kubernetes clusters should gate deployment of vulnerable images - Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed - Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed - Container images should be deployed from trusted registries only - Container registry images should have vulnerability findings resolved - Machines should have a vulnerability assessment solution - Machines should have vulnerability findings resolved - Vulnerabilities in running container images should be remediated (powered by Qualys)

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
4	**Encrypt data in transit** - Use these recommendations to secure data that's moving between components, locations, or programs. Such data is susceptible to man-in-the-middle attacks, eavesdropping, and session hijacking.	- API App should only be accessible over HTTPS - Enforce SSL connection should be enabled for MySQL database servers - Enforce SSL connection should be enabled for PostgreSQL database servers - FTPS should be required in API apps - FTPS should be required in function apps - FTPS should be required in web apps - Function App should only be accessible over HTTPS - Redis Cache should allow access only via SSL - Secure transfer to storage accounts should be enabled - TLS should be updated to the latest version for API apps - TLS should be updated to the latest version for function apps - TLS should be updated to the latest version for web apps - Web Application should only be accessible over HTTPS
4	**Restrict unauthorized network access** - Azure offers a suite of tools designed to ensure accesses across your network meet the highest security standards. Use these recommendations to manage Defender for Cloud's adaptive network hardening settings, ensure you've configured Azure Private Link for all relevant PaaS services, enable Azure Firewall on your virtual networks, and more.	- Adaptive network hardening recommendations should be applied on internet facing virtual machines - All network ports should be restricted on network security groups associated to your virtual machine - App Configuration should use private link - Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed - Azure Cache for Redis should reside within a virtual network - Azure Event Grid domains should use private link - Azure Event Grid topics should use private link - Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed - Azure Machine Learning workspaces should use private link - Azure SignalR Service should use private link - Azure Spring Cloud should use network injection - Container registries should not allow unrestricted network access - Container registries should use private link - Containers should listen on allowed ports only - CORS should not allow every resource to access API Apps - CORS should not allow every

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
		resource to access Function Apps - CORS should not allow every resource to access Web Applications - Firewall should be enabled on Key Vault - Internet-facing virtual machines should be protected with network security groups - IP forwarding on your virtual machine should be disabled - Kubernetes API server should be configured with restricted access - Private endpoint should be configured for Key Vault - Private endpoint should be enabled for MariaDB servers - Private endpoint should be enabled for MySQL servers - Private endpoint should be enabled for PostgreSQL servers - Public network access should be disabled for MariaDB servers - Public network access should be disabled for MySQL servers - Public network access should be disabled for PostgreSQL servers - Services should listen on allowed ports only - Storage account should use a private link connection - Storage accounts should restrict network access using virtual network rules - Usage of host networking and ports should be restricted - Virtual networks should be protected by Azure Firewall - VM Image Builder templates should use private link
4	**Enable encryption at rest** - Use these recommendations to ensure you mitigate misconfigurations around the protection of your stored data.	- Service Fabric clusters should have the ClusterProtectionLevel property set to EncryptAndSign - Transparent Data Encryption on SQL databases should be enabled - Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources
4	**Manage access and permissions** - A core part of a security program is ensuring your users have the necessary access to do their jobs but no more than that: the least privilege access model. Use these recommendations to manage your identity and access requirements.	- Authentication to Linux machines should require SSH keys - Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed - Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed - Container with privilege escalation should be avoided - Containers sharing sensitive host namespaces should be avoided - Deprecated accounts should be removed from subscriptions

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
		- Deprecated accounts should be removed from your subscription - Deprecated accounts with owner permissions should be removed from subscriptions - Deprecated accounts with owner permissions should be removed from your subscription - External accounts with owner permissions should be removed from subscriptions - External accounts with owner permissions should be removed from your subscription - External accounts with write permissions should be removed from subscriptions - External accounts with write permissions should be removed from your subscription - Function apps should have Client Certificates (Incoming client certificates) enabled - Guest Configuration extension should be installed on machines - Immutable (read-only) root filesystem should be enforced for containers - Least privileged Linux capabilities should be enforced for containers - Managed identity should be used in API apps - Managed identity should be used in function apps - Managed identity should be used in web apps - Privileged containers should be avoided - Role-Based Access Control should be used on Kubernetes Services - Running containers as root user should be avoided - Service Fabric clusters should only use Azure Active Directory for client authentication - Service principals should be used to protect your subscriptions instead of Management Certificates - Storage account public access should be disallowed - Usage of pod HostPath volume mounts should be restricted to a known list to restrict node access from compromised containers - Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
4	**Remediate security configurations** - Misconfigured IT assets have a higher risk of being attacked. Use these recommendations to harden the identified misconfigurations across your infrastructure.	- Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed - Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed - Container hosts should be configured securely - Log Analytics agent should be installed on Linux-based Azure Arc-enabled machines - Log Analytics agent should be installed on virtual machine scale sets - Log Analytics agent should be installed on virtual machines - Log Analytics agent should be installed on Windows-based Azure Arc-enabled machines - Machines should be configured securely - Overriding or disabling of containers AppArmor profile should be restricted - Pod Security Policies should be defined on Kubernetes Services (Deprecated) - SQL databases should have vulnerability findings resolved - SQL managed instances should have vulnerability assessment configured - SQL servers on machines should have vulnerability findings resolved - SQL servers should have vulnerability assessment configured - Virtual machine scale sets should be configured securely - Vulnerabilities in security configuration on your Linux machines should be remediated (powered by Guest Configuration) - Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration)
3	**Apply adaptive application control** - Adaptive application control is an intelligent, automated, end-to-end solution to control which applications can run on your machines. It also helps to harden your machines against malware.	- Adaptive application controls for defining safe applications should be enabled on your machines - Allowlist rules in your adaptive application control policy should be updated - Log Analytics agent should be installed on Linux-based Azure Arc-enabled machines - Log Analytics agent should be installed on virtual machines - Log Analytics agent should be installed on Windows-based Azure Arc-enabled machines

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
2	**Protect your applications with Azure advanced networking solutions** -	- Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed - Azure DDoS Protection Standard should be enabled - Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed - Container CPU and memory limits should be enforced - Web Application Firewall (WAF) should be enabled for Application Gateway - Web Application Firewall (WAF) should be enabled for Azure Front Door Service service
2	**Enable endpoint protection** - Defender for Cloud checks your organization's endpoints for active threat detection and response solutions such as Microsoft Defender for Endpoint or any of the major solutions shown in this list. When an Endpoint Detection and Response (EDR) solution isn't found, you can use these recommendations to deploy Microsoft Defender for Endpoint (included as part of Microsoft Defender for servers). Other recommendations in this control help you deploy the Log Analytics agent and configure file integrity monitoring.	- Endpoint protection health issues on machines should be resolved - Endpoint protection health issues on machines should be resolved - Endpoint protection health issues on virtual machine scale sets should be resolved - Endpoint protection should be installed on machines - Endpoint protection should be installed on machines - Endpoint protection should be installed on virtual machine scale sets - File integrity monitoring should be enabled on servers - Install endpoint protection solution on virtual machines - Log Analytics agent should be installed on Linux-based Azure Arc-enabled machines - Log Analytics agent should be installed on virtual machine scale sets - Log Analytics agent should be installed on virtual machines - Log Analytics agent should be installed on Windows-based Azure Arc-enabled machines

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
1	**Enable auditing and logging** - Detailed logs are a crucial part of incident investigations and many other troubleshooting operations. The recommendations in this control focus on ensuring you've enabled diagnostic logs wherever relevant.	- Auditing on SQL server should be enabled - Diagnostic logs in App Service should be enabled - Diagnostic logs in Azure Data Lake Store should be enabled - Diagnostic logs in Azure Stream Analytics should be enabled - Diagnostic logs in Batch accounts should be enabled - Diagnostic logs in Data Lake Analytics should be enabled - Diagnostic logs in Event Hub should be enabled - Diagnostic logs in Key Vault should be enabled - Diagnostic logs in Kubernetes services should be enabled - Diagnostic logs in Logic Apps should be enabled - Diagnostic logs in Search services should be enabled - Diagnostic logs in Service Bus should be enabled - Diagnostic logs in Virtual Machine Scale Sets should be enabled
0	**Implement security best practices** - This control has no impact on your secure score. For that reason, it's a collection of recommendations which are important to fulfil for the sake of your organization's security, but which we feel shouldn't be a part of how you assess your overall score.	- [Enable if required] Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest - [Enable if required] Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK) - [Enable if required] Cognitive Services accounts should enable data encryption with a customer-managed key (CMK) - [Enable if required] Container registries should be encrypted with a customer-managed key (CMK) - [Enable if required] MySQL servers should use customer-managed keys to encrypt data at rest - [Enable if required] PostgreSQL servers should use customer-managed keys to encrypt data at rest - [Enable if required] SQL managed instances should use customer-managed keys to encrypt data at rest - [Enable if required] SQL servers should use customer-managed keys to encrypt data at rest - [Enable if required] Storage accounts should use customer-managed key (CMK) for encryption - A maximum of 3 owners should be designated for subscriptions - Access to storage accounts with firewall and virtual network configurations should be restricted

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
		- All advanced threat protection types should be enabled in SQL managed instance advanced data security settings - All advanced threat protection types should be enabled in SQL server advanced data security settings - API Management services should use a virtual network - Audit retention for SQL servers should be set to at least 90 days - Auto provisioning of the Log Analytics agent should be enabled on subscriptions - Automation account variables should be encrypted - Azure Backup should be enabled for virtual machines - Azure Cosmos DB accounts should have firewall rules - Cognitive Services accounts should enable data encryption - Cognitive Services accounts should restrict network access - Cognitive Services accounts should use customer owned storage or enable data encryption - Default IP Filter Policy should be Deny - Diagnostic logs in IoT Hub should be enabled - Email notification for high severity alerts should be enabled - Email notification to subscription owner for high severity alerts should be enabled - Ensure API app has Client Certificates Incoming client certificates set to On - External accounts with read permissions should be removed from subscriptions - External accounts with read permissions should be removed from your subscription - Geo-redundant backup should be enabled for Azure Database for MariaDB - Geo-redundant backup should be enabled for Azure Database for MySQL - Geo-redundant backup should be enabled for Azure Database for PostgreSQL - Guest Attestation extension should be installed on supported Linux virtual machine scale sets - Guest Attestation extension should be installed on supported Linux virtual machines - Guest Attestation extension should be installed on supported Windows virtual machine scale sets - Guest Attestation extension should be installed on supported Windows virtual machines - Guest Configuration extension

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
		should be installed on machines - Identical Authentication Credentials - IoT Devices - Agent sending underutilized messages - IoT Devices - Auditd process stopped sending events - IoT Devices - Open Ports On Device - IoT Devices - Operating system baseline validation failure - IoT Devices - Permissive firewall policy in one of the chains was found - IoT Devices - Permissive firewall rule in the input chain was found - IoT Devices - Permissive firewall rule in the output chain was found - IoT Devices - TLS cipher suite upgrade needed - IP Filter rule large IP range - Java should be updated to the latest version for API apps - Java should be updated to the latest version for function apps - Java should be updated to the latest version for web apps - Key Vault keys should have an expiration date - Key Vault secrets should have an expiration date - Key vaults should have purge protection enabled - Key vaults should have soft delete enabled - Kubernetes clusters should be accessible only over HTTPS - Kubernetes clusters should disable automounting API credentials - Kubernetes clusters should not grant CAPSYSADMIN security capabilities - Kubernetes clusters should not use the default namespace - Linux virtual machines should enforce kernel module signature validation - Linux virtual machines should use only signed and trusted boot components - Linux virtual machines should use Secure Boot - Machines should be restarted to apply security configuration updates - MFA should be enabled on accounts with read permissions on subscriptions - MFA should be enabled on accounts with read permissions on your subscription - Microsoft Defender for SQL should be enabled for unprotected Azure SQL servers - Microsoft Defender for SQL should be enabled for unprotected SQL Managed Instances - Network traffic data collection agent should be installed on Linux virtual machines - Network traffic data collection agent

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
		should be installed on Windows virtual machines - Network Watcher should be enabled - Non-internet-facing virtual machines should be protected with network security groups - PHP should be updated to the latest version for API apps - PHP should be updated to the latest version for web apps - Private endpoint connections on Azure SQL Database should be enabled - Public network access on Azure SQL Database should be disabled - Public network access should be disabled for Cognitive Services accounts - Python should be updated to the latest version for API apps - Python should be updated to the latest version for function apps - Python should be updated to the latest version for web apps - Remote debugging should be turned off for API App - Remote debugging should be turned off for Function App - Remote debugging should be turned off for Web Applications - Secure Boot should be enabled on supported Windows virtual machines - SQL servers should have an Azure Active Directory administrator provisioned - Storage accounts should be migrated to new Azure Resource Manager resources - Subnets should be associated with a network security group - Subscriptions should have a contact email address for security issues - There should be more than one owner assigned to subscriptions - Validity period of certificates stored in Azure Key Vault should not exceed 12 months - Virtual machines guest attestation status should be healthy - Virtual machines should be migrated to new Azure Resource Manager resources - Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity - vTPM should be enabled on supported virtual machines - Web apps should request an SSL certificate for all incoming requests - Windows Defender Exploit Guard should be enabled on machines - Windows web servers should be configured to use secure communication protocols

SECURE SCORE	SECURITY CONTROL AND DESCRIPTION	RECOMMENDATIONS
0	**Enable enhanced security features** - Use these recommendations to enable any of the enhanced security features plans.	- Azure Arc-enabled Kubernetes clusters should have the Defender extension installed - Azure Kubernetes Service clusters should have Defender profile enabled - Microsoft Defender for App Service should be enabled - Microsoft Defender for Azure SQL Database servers should be enabled - Microsoft Defender for Containers should be enabled - Microsoft Defender for DNS should be enabled - Microsoft Defender for Key Vault should be enabled - Microsoft Defender for open-source relational databases should be enabled - Microsoft Defender for Resource Manager should be enabled - Microsoft Defender for servers should be enabled - Microsoft Defender for servers should be enabled on workspaces - Microsoft Defender for SQL on machines should be enabled on workspaces - Microsoft Defender for SQL servers on machines should be enabled - Microsoft Defender for Storage should be enabled

FAQ - Secure score

If I address only three out of four recommendations in a security control, will my secure score change?

No. It won't change until you remediate all of the recommendations for a single resource. To get the maximum score for a control, you must remediate all recommendations, for all resources.

If a recommendation isn't applicable to me, and I disable it in the policy, will my security control be fulfilled and my secure score updated?

Yes. We recommend disabling recommendations when they're inapplicable in your environment. For instructions on how to disable a specific recommendation, see Disable security policies.

If a security control offers me zero points towards my secure score, should I ignore it?

In some cases, you'll see a control max score greater than zero, but the impact is zero. When the incremental score for fixing resources is negligible, it's rounded to zero. Don't ignore these recommendations as they still bring security improvements. The only exception is the "Additional Best Practice" control. Remediating these recommendations won't increase your score, but it will enhance your overall security.

Next steps

This article described the secure score and the included security controls.

Access and track your secure score

For related material, see the following articles:

- Learn about the different elements of a recommendation
- Learn how to remediate recommendations
- View the GitHub-based tools for working programmatically with secure score

Security recommendations - a reference guide

2/15/2022 • 83 minutes to read • Edit Online

This article lists the recommendations you might see in Microsoft Defender for Cloud. The recommendations shown in your environment depend on the resources you're protecting and your customized configuration.

Defender for Cloud's recommendations are based on the Azure Security Benchmark. Azure Security Benchmark is the Microsoft-authored, Azure-specific set of guidelines for security and compliance best practices based on common compliance frameworks. This widely respected benchmark builds on the controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security.

To learn about how to respond to these recommendations, see Remediate recommendations in Defender for Cloud.

Your secure score is based on the number of security recommendations you've completed. To decide which recommendations to resolve first, look at the severity of each one and its potential impact on your secure score.

TIP

If a recommendation's description says "No related policy", it's usually because that recommendation is dependent on a different recommendation and *its* policy. For example, the recommendation "Endpoint protection health failures should be remediated...", relies on the recommendation that checks whether an endpoint protection solution is even *installed* ("Endpoint protection solution should be installed..."). The underlying recommendation *does* have a policy. Limiting the policies to only the foundational recommendation simplifies policy management.

AppServices recommendations

There are **31** recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
API App should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks. (Related policy: API App should only be accessible over HTTPS)	Medium
CORS should not allow every resource to access API Apps	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your API app. Allow only required domains to interact with your API app. (Related policy: CORS should not allow every resource to access your API App)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
CORS should not allow every resource to access Function Apps	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your Function app. Allow only required domains to interact with your Function app. (Related policy: CORS should not allow every resource to access your Function Apps)	Low
CORS should not allow every resource to access Web Applications	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your web application. Allow only required domains to interact with your web app. (Related policy: CORS should not allow every resource to access your Web Applications)	Low
Diagnostic logs in App Service should be enabled	Audit enabling of diagnostic logs on the app. This enables you to recreate activity trails for investigation purposes if a security incident occurs or your network is compromised (No related policy)	Medium
Ensure API app has Client Certificates Incoming client certificates set to On	Client certificates allow for the app to request a certificate for incoming requests. Only clients that have a valid certificate will be able to reach the app. (Related policy: Ensure API app has 'Client Certificates (Incoming client certificates)' set to 'On')	Medium
FTPS should be required in API apps	Enable FTPS enforcement for enhanced security (Related policy: FTPS only should be required in your API App)	High
FTPS should be required in function apps	Enable FTPS enforcement for enhanced security (Related policy: FTPS only should be required in your Function App)	High
FTPS should be required in web apps	Enable FTPS enforcement for enhanced security (Related policy: FTPS should be required in your Web App)	High
Function App should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks. (Related policy: Function App should only be accessible over HTTPS)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Function apps should have Client Certificates (Incoming client certificates) enabled	Client certificates allow for the app to request a certificate for incoming requests. Only clients with valid certificates will be able to reach the app. (Related policy: Function apps should have 'Client Certificates (Incoming client certificates)' enabled)	Medium
Java should be updated to the latest version for API apps	Periodically, newer versions are released for Java either due to security flaws or to include additional functionality. Using the latest Python version for API apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Java version' is the latest, if used as a part of the API app)	Medium
Java should be updated to the latest version for function apps	Periodically, newer versions are released for Java software either due to security flaws or to include additional functionality. Using the latest Java version for function apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Java version' is the latest, if used as a part of the Function app)	Medium
Java should be updated to the latest version for web apps	Periodically, newer versions are released for Java software either due to security flaws or to include additional functionality. Using the latest Java version for web apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Java version' is the latest, if used as a part of the Web app)	Medium
Managed identity should be used in API apps	For enhanced authentication security, use a managed identity. On Azure, managed identities eliminate the need for developers to have to manage credentials by providing an identity for the Azure resource in Azure AD and using it to obtain Azure Active Directory (Azure AD) tokens. (Related policy: Managed identity should be used in your API App)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Managed identity should be used in function apps	For enhanced authentication security, use a managed identity. On Azure, managed identities eliminate the need for developers to have to manage credentials by providing an identity for the Azure resource in Azure AD and using it to obtain Azure Active Directory (Azure AD) tokens. (Related policy: Managed identity should be used in your Function App)	Medium
Managed identity should be used in web apps	For enhanced authentication security, use a managed identity. On Azure, managed identities eliminate the need for developers to have to manage credentials by providing an identity for the Azure resource in Azure AD and using it to obtain Azure Active Directory (Azure AD) tokens. (Related policy: Managed identity should be used in your Web App)	Medium
Microsoft Defender for App Service should be enabled	Microsoft Defender for App Service leverages the scale of the cloud, and the visibility that Azure has as a cloud provider, to monitor for common web app attacks. Microsoft Defender for App Service can discover attacks on your applications and identify emerging attacks. Important: Remediating this recommendation will result in charges for protecting your App Service plans. If you don't have any App Service plans in this subscription, no charges will be incurred. If you create any App Service plans on this subscription in the future, they will automatically be protected and charges will begin at that time. Learn more in Protect your web apps and APIs. (Related policy: Azure Defender for App Service should be enabled)	High
PHP should be updated to the latest version for API apps	Periodically, newer versions are released for PHP software either due to security flaws or to include additional functionality. Using the latest PHP version for API apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'PHP version' is the latest, if used as a part of the API app)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
PHP should be updated to the latest version for web apps	Periodically, newer versions are released for PHP software either due to security flaws or to include additional functionality. Using the latest PHP version for web apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'PHP version' is the latest, if used as a part of the WEB app)	Medium
Python should be updated to the latest version for API apps	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for API apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Python version' is the latest, if used as a part of the API app)	Medium
Python should be updated to the latest version for function apps	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for function apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Python version' is the latest, if used as a part of the Function app)	Medium
Python should be updated to the latest version for web apps	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for web apps is recommended to benefit from security fixes, if any, and/or new functionalities of the latest version. (Related policy: Ensure that 'Python version' is the latest, if used as a part of the Web app)	Medium
Remote debugging should be turned off for API App	Remote debugging requires inbound ports to be opened on an API app. Remote debugging should be turned off. (Related policy: Remote debugging should be turned off for API Apps)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Remote debugging should be turned off for Function App	Remote debugging requires inbound ports to be opened on an Azure Function app. Remote debugging should be turned off. (Related policy: Remote debugging should be turned off for Function Apps)	Low
Remote debugging should be turned off for Web Applications	Remote debugging requires inbound ports to be opened on a web application. Remote debugging is currently enabled. If you no longer need to use remote debugging, it should be turned off. (Related policy: Remote debugging should be turned off for Web Applications)	Low
TLS should be updated to the latest version for API apps	Upgrade to the latest TLS version (Related policy: Latest TLS version should be used in your API App)	High
TLS should be updated to the latest version for function apps	Upgrade to the latest TLS version (Related policy: Latest TLS version should be used in your Function App)	High
TLS should be updated to the latest version for web apps	Upgrade to the latest TLS version (Related policy: Latest TLS version should be used in your Web App)	High
Web Application should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks. (Related policy: Web Application should only be accessible over HTTPS)	Medium
Web apps should request an SSL certificate for all incoming requests	Client certificates allow for the app to request a certificate for incoming requests. Only clients that have a valid certificate will be able to reach the app. (Related policy: Ensure WEB app has 'Client Certificates (Incoming client certificates)' set to 'On')	Medium

Compute recommendations

There are **61** recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY

RECOMMENDATION	DESCRIPTION	SEVERITY
Adaptive application controls for defining safe applications should be enabled on your machines	Enable application controls to define the list of known-safe applications running on your machines, and alert you when other applications run. This helps harden your machines against malware. To simplify the process of configuring and maintaining your rules, Defender for Cloud uses machine learning to analyze the applications running on each machine and suggest the list of known-safe applications. (Related policy: Adaptive application controls for defining safe applications should be enabled on your machines)	High
Allowlist rules in your adaptive application control policy should be updated	Monitor for changes in behavior on groups of machines configured for auditing by Defender for Cloud's adaptive application controls. Defender for Cloud uses machine learning to analyze the running processes on your machines and suggest a list of known-safe applications. These are presented as recommended apps to allow in adaptive application control policies. (Related policy: Allowlist rules in your adaptive application control policy should be updated)	High
Authentication to Linux machines should require SSH keys	Although SSH itself provides an encrypted connection, using passwords with SSH still leaves the VM vulnerable to brute-force attacks. The most secure option for authenticating to an Azure Linux virtual machine over SSH is with a public-private key pair, also known as SSH keys. Learn more in Detailed steps: Create and manage SSH keys for authentication to a Linux VM in Azure. (Related policy: Audit Linux machines that are not using SSH key for authentication)	Medium
Automation account variables should be encrypted	It is important to enable encryption of Automation account variable assets when storing sensitive data. (Related policy: Automation account variables should be encrypted)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Azure Backup should be enabled for virtual machines	Protect the data on your Azure virtual machines with Azure Backup. Azure Backup is an Azure-native, cost-effective, data protection solution. It creates recovery points that are stored in geo-redundant recovery vaults. When you restore from a recovery point, you can restore the whole VM or specific files. (Related policy: Azure Backup should be enabled for Virtual Machines)	Low
Container hosts should be configured securely	Remediate vulnerabilities in security configuration on machines with Docker installed to protect them from attacks. (Related policy: Vulnerabilities in container security configurations should be remediated)	High
Diagnostic logs in Azure Stream Analytics should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Azure Stream Analytics should be enabled)	Low
Diagnostic logs in Batch accounts should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Batch accounts should be enabled)	Low
Diagnostic logs in Event Hub should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Event Hub should be enabled)	Low
Diagnostic logs in Kubernetes services should be enabled	Enable diagnostic logs in your Kubernetes services and retain them up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs. (No related policy)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Diagnostic logs in Logic Apps should be enabled	To ensure you can recreate activity trails for investigation purposes when a security incident occurs or your network is compromised, enable logging. If your diagnostic logs aren't being sent to a Log Analytics workspace, Azure Storage account, or Azure Event Hub, ensure you've configured diagnostic settings to send platform metrics and platform logs to the relevant destinations. Learn more in Create diagnostic settings to send platform logs and metrics to different destinations. (Related policy: Diagnostic logs in Logic Apps should be enabled)	Low
Diagnostic logs in Search services should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Search services should be enabled)	Low
Diagnostic logs in Service Bus should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Service Bus should be enabled)	Low
Diagnostic logs in Virtual Machine Scale Sets should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Virtual Machine Scale Sets should be enabled)	Low
Endpoint protection health issues on machines should be resolved	For full Defender for Cloud protection, resolve monitoring agent issues on your machines by following the instructions in the Troubleshooting guide. (Related policy: Monitor missing Endpoint Protection in Azure Security Center)	Medium
Endpoint protection health issues on machines should be resolved	Resolve endpoint protection health issues on your virtual machines to protect them from latest threats and vulnerabilities. See the documentation for the endpoint protection solutions supported by Defender for Cloud and the endpoint protection assessments. (No related policy)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Endpoint protection health issues on virtual machine scale sets should be resolved	Remediate endpoint protection health failures on your virtual machine scale sets to protect them from threats and vulnerabilities. (Related policy: Endpoint protection solution should be installed on virtual machine scale sets)	Low
Endpoint protection should be installed on machines	To protect machines from threats and vulnerabilities, install a supported endpoint protection solution. Learn more about how endpoint protection for machines is evaluated in Endpoint protection assessment and recommendations in Microsoft Defender for Cloud. (No related policy)	High
Endpoint protection should be installed on machines	Install an endpoint protection solution on your Windows and Linux machines, to protect them from threats and vulnerabilities. (No related policy)	Medium
Endpoint protection should be installed on virtual machine scale sets	Install an endpoint protection solution on your virtual machines scale sets, to protect them from threats and vulnerabilities. (Related policy: Endpoint protection solution should be installed on virtual machine scale sets)	High
File integrity monitoring should be enabled on servers	Defender for Cloud has identified virtual machines that are missing a file integrity monitoring solution. To monitor changes to critical files, registry keys, and more on your servers, enable file integrity monitoring. Learn more > (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Guest Attestation extension should be installed on supported Linux virtual machine scale sets	Install Guest Attestation extension on supported Linux virtual machine scale sets to allow Microsoft Defender for Cloud to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machine scale sets. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines. (No related policy)	Low
Guest Attestation extension should be installed on supported Linux virtual machines	Install Guest Attestation extension on supported Linux virtual machines to allow Microsoft Defender for Cloud to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machines. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines. (No related policy)	Low
Guest Attestation extension should be installed on supported Windows virtual machine scale sets	Install Guest Attestation extension on supported virtual machine scale sets to allow Microsoft Defender for Cloud to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machine scale sets. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines. (No related policy)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Guest Attestation extension should be installed on supported Windows virtual machines	Install Guest Attestation extension on supported virtual machines to allow Microsoft Defender for Cloud to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machines. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines. (No related policy)	Low
Guest Configuration extension should be installed on machines	To ensure secure configurations of in-guest settings of your machine, install the Guest Configuration extension. In-guest settings that the extension monitors include the configuration of the operating system, application configuration or presence, and environment settings. Once installed, in-guest policies will be available such as 'Windows Exploit guard should be enabled'. Learn more. (Related policy: Virtual machines should have the Guest Configuration extension)	Medium
Install endpoint protection solution on virtual machines	Install an endpoint protection solution on your virtual machines, to protect them from threats and vulnerabilities. (Related policy: Monitor missing Endpoint Protection in Azure Security Center)	High
Linux virtual machines should enforce kernel module signature validation	To help mitigate against the execution of malicious or unauthorized code in kernel mode, enforce kernel module signature validation on supported Linux virtual machines. Kernel module signature validation ensures that only trusted kernel modules will be allowed to run. This assessment only applies to Linux virtual machines that have the Azure Monitor Agent installed. (No related policy)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Linux virtual machines should use only signed and trusted boot components	With Secure Boot enabled, all OS boot components (boot loader, kernel, kernel drivers) must be signed by trusted publishers. Defender for Cloud has identified untrusted OS boot components on one or more of your Linux machines. To protect your machines from potentially malicious components, add them to your allow list or remove the identified components. (No related policy)	Low
Linux virtual machines should use Secure Boot	To protect against the installation of malware-based rootkits and boot kits, enable Secure Boot on supported Linux virtual machines. Secure Boot ensures that only signed operating systems and drivers will be allowed to run. This assessment only applies to Linux virtual machines that have the Azure Monitor Agent installed. (No related policy)	Low
Log Analytics agent should be installed on Linux-based Azure Arc-enabled machines	Defender for Cloud uses the Log Analytics agent (also known as OMS) to collect security events from your Azure Arc machines. To deploy the agent on all your Azure Arc machines, follow the remediation steps. (No related policy)	High
Log Analytics agent should be installed on virtual machine scale sets	Defender for Cloud collects data from your Azure virtual machines (VMs) to monitor for security vulnerabilities and threats. Data is collected using the Log Analytics agent, formerly known as the Microsoft Monitoring Agent (MMA), which reads various security-related configurations and event logs from the machine and copies the data to your workspace for analysis. You'll also need to follow that procedure if your VMs are used by an Azure managed service such as Azure Kubernetes Service or Azure Service Fabric. You cannot configure auto-provisioning of the agent for Azure virtual machine scale sets. To deploy the agent on virtual machine scale sets (including those used by Azure managed services such as Azure Kubernetes Service and Azure Service Fabric), follow the procedure in the remediation steps. (Related policy: Log Analytics agent should be installed on your virtual machine scale sets for Azure Security Center monitoring)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Log Analytics agent should be installed on virtual machines	Defender for Cloud collects data from your Azure virtual machines (VMs) to monitor for security vulnerabilities and threats. Data is collected using the Log Analytics agent, formerly known as the Microsoft Monitoring Agent (MMA), which reads various security-related configurations and event logs from the machine and copies the data to your Log Analytics workspace for analysis. This agent is also required if your VMs are used by an Azure managed service such as Azure Kubernetes Service or Azure Service Fabric. We recommend configuring auto-provisioning to automatically deploy the agent. If you choose not to use auto-provisioning, manually deploy the agent to your VMs using the instructions in the remediation steps. (Related policy: Log Analytics agent should be installed on your virtual machine for Azure Security Center monitoring)	High
Log Analytics agent should be installed on Windows-based Azure Arc-enabled machines	Defender for Cloud uses the Log Analytics agent (also known as MMA) to collect security events from your Azure Arc machines. To deploy the agent on all your Azure Arc machines, follow the remediation steps. (No related policy)	High
Machines should be configured securely	Remediate vulnerabilities in security configuration on your machines to protect them from attacks. (Related policy: Vulnerabilities in security configuration on your machines should be remediated)	Low
Machines should be restarted to apply security configuration updates	To apply security configuration updates and protect against vulnerabilities, restart your machines. This assessment only applies to Linux virtual machines that have the Azure Monitor Agent installed. (No related policy)	Low
Machines should have a vulnerability assessment solution	Defender for Cloud regularly checks your connected machines to ensure they're running vulnerability assessment tools. Use this recommendation to deploy a vulnerability assessment solution. (Related policy: A vulnerability assessment solution should be enabled on your virtual machines)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Machines should have vulnerability findings resolved	Resolve the findings from the vulnerability assessment solutions on your virtual machines. (Related policy: A vulnerability assessment solution should be enabled on your virtual machines)	Low
Management ports of virtual machines should be protected with just-in-time network access control	Defender for Cloud has identified some overly-permissive inbound rules for management ports in your Network Security Group. Enable just-in-time access control to protect your VM from internet-based brute-force attacks. Learn more in Understanding just-in-time (JIT) VM access. (Related policy: Management ports of virtual machines should be protected with just-in-time network access control)	High
Microsoft Defender for servers should be enabled	Microsoft Defender for servers provides real-time threat protection for your server workloads and generates hardening recommendations as well as alerts about suspicious activities. You can use this information to quickly remediate security issues and improve the security of your servers. Important: Remediating this recommendation will result in charges for protecting your servers. If you don't have any servers in this subscription, no charges will be incurred. If you create any servers on this subscription in the future, they will automatically be protected and charges will begin at that time. Learn more in Introduction to Microsoft Defender for servers. (Related policy: Azure Defender for servers should be enabled)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for servers should be enabled on workspaces	Microsoft Defender for servers brings threat detection and advanced defenses for your Windows and Linux machines. With this Defender plan enabled on your subscriptions but not on your workspaces, you're paying for the full capability of Microsoft Defender for servers but missing out on some of the benefits. When you enable Microsoft Defender for servers on a workspace, all machines reporting to that workspace will be billed for Microsoft Defender for servers - even if they're in subscriptions without Defender plans enabled. Unless you also enable Microsoft Defender for servers on the subscription, those machines won't be able to take advantage of just-in-time VM access, adaptive application controls, and network detections for Azure resources. Learn more in Introduction to Microsoft Defender for servers. (No related policy)	Medium
Network traffic data collection agent should be installed on Linux virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations and specific network threats. (Related policy: Network traffic data collection agent should be installed on Linux virtual machines)	Medium
Network traffic data collection agent should be installed on Windows virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations, and specific network threats. (Related policy: Network traffic data collection agent should be installed on Windows virtual machines)	Medium
Pod Security Policies should be defined on Kubernetes Services (Deprecated)	(Deprecated) Define Pod Security Policies to reduce the attack vector by removing unnecessary application privileges. It is recommended to configure Pod Security Policies to ensure that Pods that request resources you don't allow can't run in the AKS cluster. (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Secure Boot should be enabled on supported Windows virtual machines	Enable Secure Boot on supported Windows virtual machines to mitigate against malicious and unauthorized changes to the boot chain. Once enabled, only trusted bootloaders, kernel and kernel drivers will be allowed to run. This assessment only applies to trusted launch enabled Windows virtual machines. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines.. (No related policy)	Low
Service Fabric clusters should have the ClusterProtectionLevel property set to EncryptAndSign	Service Fabric provides three levels of protection (None, Sign and EncryptAndSign) for node-to-node communication using a primary cluster certificate. Set the protection level to ensure that all node-to-node messages are encrypted and digitally signed. (Related policy: Service Fabric clusters should have the ClusterProtectionLevel property set to EncryptAndSign)	High
Service Fabric clusters should only use Azure Active Directory for client authentication	Perform Client authentication only via Azure Active Directory in Service Fabric (Related policy: Service Fabric clusters should only use Azure Active Directory for client authentication)	High
System updates on virtual machine scale sets should be installed	Install missing system security and critical updates to secure your Windows and Linux virtual machine scale sets. (Related policy: System updates on virtual machine scale sets should be installed)	High
System updates should be installed on your machines	Install missing system security and critical updates to secure your Windows and Linux virtual machines and computers (Related policy: System updates should be installed on your machines)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
System updates should be installed on your machines (powered by Update Center)	Your machines are missing system, security, and critical updates. Software updates often include critical patches to security holes. Such holes are frequently exploited in malware attacks so it's vital to keep your software updated. To install all outstanding patches and secure your machines, follow the remediation steps. (No related policy)	High
Virtual machine scale sets should be configured securely	Remediate vulnerabilities in security configuration on your virtual machine scale sets to protect them from attacks. (Related policy: Vulnerabilities in security configuration on your virtual machine scale sets should be remediated)	High
Virtual machines guest attestation status should be healthy	Guest attestation is performed by sending a trusted log (TCGLog) to an attestation server. The server uses these logs to determine whether boot components are trustworthy. This assessment is intended to detect compromises of the boot chain which might be the result of a bootkit or rootkit infection. This assessment only applies to Trusted Launch enabled virtual machines that have the Guest Attestation extension installed. (No related policy)	Medium
Virtual machines should be migrated to new Azure Resource Manager resources	Virtual Machines (classic) was deprecated and these VMs should be migrated to Azure Resource Manager. Because Azure Resource Manager now has full IaaS capabilities and other advancements, we deprecated the management of IaaS virtual machines (VMs) through Azure Service Manager (ASM) on February 28, 2020. This functionality will be fully retired on March 1, 2023. Available resources and information about this tool & migration: Overview of Virtual machines (classic) deprecation, step by step process for migration & available Microsoft resources. Details about Migrate to Azure Resource Manager migration tool. Migrate to Azure Resource Manager migration tool using PowerShell. (Related policy: Virtual machines should be migrated to new Azure Resource Manager resources)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources	By default, a virtual machine's OS and data disks are encrypted-at-rest using platform-managed keys; temp disks and data caches aren't encrypted, and data isn't encrypted when flowing between compute and storage resources. For a comparison of different disk encryption technologies in Azure, see https://aka.ms/diskencryptioncomparison. Use Azure Disk Encryption to encrypt all this data. Disregard this recommendation if: 1. You're using the encryption-at-host feature, or 2. Server-side encryption on Managed Disks meets your security requirements. Learn more in Server-side encryption of Azure Disk Storage. (Related policy: Disk encryption should be applied on virtual machines)	High
Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity	The Guest Configuration extension requires a system assigned managed identity. Azure virtual machines in the scope of this policy will be non-compliant when they have the Guest Configuration extension installed but do not have a system assigned managed identity. Learn more (Related policy: Guest Configuration extension should be deployed to Azure virtual machines with system assigned managed identity)	Medium
vTPM should be enabled on supported virtual machines	Enable virtual TPM device on supported virtual machines to facilitate Measured Boot and other OS security features that require a TPM. Once enabled, vTPM can be used to attest boot integrity. This assessment only applies to trusted launch enabled virtual machines. Important: Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it. Learn more about Trusted launch for Azure virtual machines. (No related policy)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Vulnerabilities in running container images should be remediated (powered by Qualys)	Container image vulnerability assessment scans container images running on your Kubernetes clusters for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks. (No related policy)	High
Vulnerabilities in security configuration on your Linux machines should be remediated (powered by Guest Configuration)	Remediate vulnerabilities in security configuration on your Linux machines to protect them from attacks. (Related policy: Linux machines should meet requirements for the Azure security baseline)	Low
Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration)	Remediate vulnerabilities in security configuration on your Windows machines to protect them from attacks. (No related policy)	Low
Windows Defender Exploit Guard should be enabled on machines	Windows Defender Exploit Guard uses the Azure Policy Guest Configuration agent. Exploit Guard has four components that are designed to lock down devices against a wide variety of attack vectors and block behaviors commonly used in malware attacks while enabling enterprises to balance their security risk and productivity requirements (Windows only). (Related policy: Audit Windows machines on which Windows Defender Exploit Guard is not enabled)	Medium
Windows web servers should be configured to use secure communication protocols	To protect the privacy of information communicated over the Internet, your web servers should use the latest version of the industry-standard cryptographic protocol, Transport Layer Security (TLS). TLS secures communications over a network by using security certificates to encrypt a connection between machines. (Related policy: Audit Windows web servers that are not using secure communication protocols)	High

Container recommendations

There are 24 recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] Container registries should be encrypted with a customer-managed key (CMK)	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Use customer-managed keys to manage the encryption at rest of the contents of your registries. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more about CMK encryption at https://aka.ms/acr/CMK. (Related policy: Container registries should be encrypted with a customer-managed key (CMK))	Low
[Preview] Kubernetes clusters should gate deployment of vulnerable images	Protect your Kubernetes clusters and container workloads from potential threats by restricting deployment of container images with vulnerable software components. Use Defender for Cloud's CI/CD scanning and Microsoft Defender for container registries to identify and patch vulnerabilities prior to deployment. Evaluation prerequisite: Azure policy add-on/extension and the Defender profile/extension. Applicable only for private preview customers. (No related policy)	High
Azure Arc-enabled Kubernetes clusters should have the Azure Policy extension installed	Azure Policy extension for Kubernetes extends Gatekeeper v3, an admission controller webhook for Open Policy Agent (OPA), to apply at-scale enforcements and safeguards on your clusters in a centralized, consistent manner. (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Azure Kubernetes Service clusters should have the Azure Policy add-on for Kubernetes installed	Azure Policy add-on for Kubernetes extends Gatekeeper v3, an admission controller webhook for Open Policy Agent (OPA), to apply at-scale enforcements and safeguards on your clusters in a centralized, consistent manner. Defender for Cloud requires the Add-on to audit and enforce security capabilities and compliance inside your clusters. Learn more. Requires Kubernetes v1.14.0 or later. (Related policy: Azure Policy Add-on for Kubernetes service (AKS) should be installed and enabled on your clusters)	High
Container CPU and memory limits should be enforced	Enforcing CPU and memory limits prevents resource exhaustion attacks (a form of denial of service attack). We recommend setting limits for containers to ensure the runtime prevents the container from using more than the configured resource limit. (Related policy: Ensure container CPU and memory resource limits do not exceed the specified limits in Kubernetes cluster)	Medium
Container images should be deployed from trusted registries only	Images running on your Kubernetes cluster should come from known and monitored container image registries. Trusted registries reduce your cluster's exposure risk by limiting the potential for the introduction of unknown vulnerabilities, security issues and malicious images. (Related policy: Ensure only allowed container images in Kubernetes cluster)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Container registries should not allow unrestricted network access	Azure container registries by default accept connections over the internet from hosts on any network. To protect your registries from potential threats, allow access from only specific public IP addresses or address ranges. If your registry doesn't have an IP/firewall rule or a configured virtual network, it will appear in the unhealthy resources. Learn more about Container Registry network rules here: https://aka.ms/acr/portal/public-network and here https://aka.ms/acr/vnet. (Related policy: Container registries should not allow unrestricted network access)	Medium
Container registries should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your container registries instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/acr/private-link. (Related policy: Container registries should use private link)	Medium
Container registry images should have vulnerability findings resolved	Container image vulnerability assessment scans your registry for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks. (Related policy: Vulnerabilities in Azure Container Registry images should be remediated)	High
Container with privilege escalation should be avoided	Containers shouldn't run with privilege escalation to root in your Kubernetes cluster. The AllowPrivilegeEscalation attribute controls whether a process can gain more privileges than its parent process. (Related policy: Kubernetes clusters should not allow container privilege escalation)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Containers sharing sensitive host namespaces should be avoided	To protect against privilege escalation outside the container, avoid pod access to sensitive host namespaces (host process ID and host IPC) in a Kubernetes cluster. (Related policy: Kubernetes cluster containers should not share host process ID or host IPC namespace)	Medium
Containers should listen on allowed ports only	To reduce the attack surface of your Kubernetes cluster, restrict access to the cluster by limiting containers access to the configured ports. (Related policy: Ensure containers listen only on allowed ports in Kubernetes cluster)	Medium
Immutable (read-only) root filesystem should be enforced for containers	Containers should run with a read only root file system in your Kubernetes cluster. Immutable filesystem protects containers from changes at run-time with malicious binaries being added to PATH. (Related policy: Kubernetes cluster containers should run with a read only root file system)	Medium
Kubernetes API server should be configured with restricted access	To ensure that only applications from allowed networks, machines, or subnets can access your cluster, restrict access to your Kubernetes API server. You can restrict access by defining authorized IP ranges, or by setting up your API servers as private clusters as explained inCreate a private Azure Kubernetes Service cluster. (Related policy: Authorized IP ranges should be defined on Kubernetes Services)	High
Kubernetes clusters should be accessible only over HTTPS	Use of HTTPS ensures authentication and protects data in transit from network layer eavesdropping attacks. This capability is currently generally available for Kubernetes Service (AKS), and in preview for AKS Engine and Azure Arc-enabled Kubernetes. For more info, visit https://aka.ms/kubepolicydoc (Related policy: Enforce HTTPS ingress in Kubernetes cluster)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Least privileged Linux capabilities should be enforced for containers	To reduce attack surface of your container, restrict Linux capabilities and grant specific privileges to containers without granting all the privileges of the root user. We recommend dropping all capabilities, then adding those that are required (Related policy: Kubernetes cluster containers should only use allowed capabilities)	Medium
Microsoft Defender for Containers should be enabled	Microsoft Defender for Containers provides hardening, vulnerability assessment and run-time protections for your Azure, hybrid, and multi-cloud Kubernetes environments. You can use this information to quickly remediate security issues and improve the security of your containers. Important: Remediating this recommendation will result in charges for protecting your Kubernetes clusters. If you don't have any Kubernetes clusters in this subscription, no charges will be incurred. If you create any Kubernetes clusters on this subscription in the future, they will automatically be protected and charges will begin at that time. Learn more in Introduction to Microsoft Defender for Containers. (No related policy)	High
Overriding or disabling of containers AppArmor profile should be restricted	Containers running on your Kubernetes cluster should be limited to allowed AppArmor profiles only. ;AppArmor (Application Armor) is a Linux security module that protects an operating system and its applications from security threats. To use it, a system administrator associates an AppArmor security profile with each program. (Related policy: Kubernetes cluster containers should only use allowed AppArmor profiles)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Privileged containers should be avoided	To prevent unrestricted host access, avoid privileged containers whenever possible. Privileged containers have all of the root capabilities of a host machine. They can be used as entry points for attacks and to spread malicious code or malware to compromised applications, hosts and networks. (Related policy: Do not allow privileged containers in Kubernetes cluster)	Medium
Role-Based Access Control should be used on Kubernetes Services	To provide granular filtering on the actions that users can perform, use Role-Based Access Control (RBAC) to manage permissions in Kubernetes Service Clusters and configure relevant authorization policies. For more information, see Azure role-based access control. (Related policy: Role-Based Access Control (RBAC) should be used on Kubernetes Services)	High
Running containers as root user should be avoided	Containers shouldn't run as root users in your Kubernetes cluster. Running a process as the root user inside a container runs it as root on the host. If there's a compromise, an attacker has root in the container, and any misconfigurations become easier to exploit. (Related policy: Kubernetes cluster pods and containers should only run with approved user and group IDs)	High
Services should listen on allowed ports only	To reduce the attack surface of your Kubernetes cluster, restrict access to the cluster by limiting services access to the configured ports. (Related policy: Ensure services listen only on allowed ports in Kubernetes cluster)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Usage of host networking and ports should be restricted	Restrict pod access to the host network and the allowable host port range in a Kubernetes cluster. Pods created with the hostNetwork attribute enabled will share the node's network space. To avoid compromised container from sniffing network traffic, we recommend not putting your pods on the host network. If you need to expose a container port on the node's network, and using a Kubernetes Service node port does not meet your needs, another possibility is to specify a hostPort for the container in the pod spec. (Related policy: Kubernetes cluster pods should only use approved host network and port range)	Medium
Usage of pod HostPath volume mounts should be restricted to a known list to restrict node access from compromised containers	We recommend limiting pod HostPath volume mounts in your Kubernetes cluster to the configured allowed host paths. If there's a compromise, the container node access from the containers should be restricted. (Related policy: Kubernetes cluster pod hostPath volumes should only use allowed host paths)	Medium

Data recommendations

There are 72 recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Use customer-managed keys to manage the encryption at rest of your Azure Cosmos DB. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more about CMK encryption at https://aka.ms/cosmosdb-cmk. (Related policy: Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK)	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Manage encryption at rest of your Azure Machine Learning workspace data with customer-managed keys (CMK). By default, customer data is encrypted with service-managed keys, but CMKs are commonly required to meet regulatory compliance standards. CMKs enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more about CMK encryption at https://aka.ms/azureml-workspaces-cmk. (Related policy: Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK))	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] Cognitive Services accounts should enable data encryption with a customer-managed key (CMK)	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs enable the data stored in Cognitive Services to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more about CMK encryption at https://aka.ms/cosmosdb-cmk. (Related policy: Cognitive Services accounts should enable data encryption with a customer-managed key(CMK))	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] MySQL servers should use customer-managed keys to encrypt data at rest	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Use customer-managed keys to manage the encryption at rest of your MySQL servers. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. (Related policy: Bring your own key data protection should be enabled for MySQL servers)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] PostgreSQL servers should use customer-managed keys to encrypt data at rest	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Use customer-managed keys to manage the encryption at rest of your PostgreSQL servers. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. (Related policy: Bring your own key data protection should be enabled for PostgreSQL servers)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] SQL managed instances should use customer-managed keys to encrypt data at rest	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Implementing Transparent Data Encryption (TDE) with your own key provides you with increased transparency and control over the TDE Protector, increased security with an HSM-backed external service, and promotion of separation of duties. This recommendation applies to organizations with a related compliance requirement. (Related policy: SQL managed instances should use customer-managed keys to encrypt data at rest)	Low
[Enable if required] SQL servers should use customer-managed keys to encrypt data at rest	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Implementing Transparent Data Encryption (TDE) with your own key provides increased transparency and control over the TDE Protector, increased security with an HSM-backed external service, and promotion of separation of duties. This recommendation applies to organizations with a related compliance requirement. (Related policy: SQL servers should use customer-managed keys to encrypt data at rest)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
[Enable if required] Storage accounts should use customer-managed key (CMK) for encryption	Recommendations to use customer-managed keys for encryption of data at rest are not assessed by default, but are available to enable for applicable scenarios. Data is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when obligated by compliance or restrictive policy requirements. To enable this recommendation, navigate to your Security Policy for the applicable scope, and update the *Effect* parameter for the corresponding policy to audit or enforce the use of customer-managed keys. Learn more in Manage security policies. Secure your storage account with greater flexibility using customer-managed keys (CMKs). When you specify a CMK, that key is used to protect and control access to the key that encrypts your data. Using CMKs provides additional capabilities to control rotation of the key encryption key or cryptographically erase data. (Related policy: Storage accounts should use customer-managed key (CMK) for encryption)	Low
All advanced threat protection types should be enabled in SQL managed instance advanced data security settings	It is recommended to enable all advanced threat protection types on your SQL managed instances. Enabling all types protects against SQL injection, database vulnerabilities, and any other anomalous activities. (No related policy)	Medium
All advanced threat protection types should be enabled in SQL server advanced data security settings	It is recommended to enable all advanced threat protection types on your SQL servers. Enabling all types protects against SQL injection, database vulnerabilities, and any other anomalous activities. (No related policy)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
API Management services should use a virtual network	Azure Virtual Network deployment provides enhanced security, isolation and allows you to place your API Management service in a non-internet routable network that you control access to. These networks can then be connected to your on-premises networks using various VPN technologies, which enables access to your backend services within the network and/or on-premises. The developer portal and API gateway, can be configured to be accessible either from the Internet or only within the virtual network. (Related policy: API Management services should use a virtual network)	Medium
App Configuration should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your app configuration instances instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/appconfig/private-endpoint. (Related policy: App Configuration should use private link)	Medium
Audit retention for SQL servers should be set to at least 90 days	Audit SQL servers configured with an auditing retention period of less than 90 days. (Related policy: SQL servers should be configured with 90 days auditing retention or higher.)	Low
Auditing on SQL server should be enabled	Enable auditing on your SQL Server to track database activities across all databases on the server and save them in an audit log. (Related policy: Auditing on SQL server should be enabled)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Auto provisioning of the Log Analytics agent should be enabled on subscriptions	To monitor for security vulnerabilities and threats, Microsoft Defender for Cloud collects data from your Azure virtual machines. Data is collected by the Log Analytics agent, formerly known as the Microsoft Monitoring Agent (MMA), which reads various security-related configurations and event logs from the machine and copies the data to your Log Analytics workspace for analysis. We recommend enabling auto provisioning to automatically deploy the agent to all supported Azure VMs and any new ones that are created. (Related policy: Auto provisioning of the Log Analytics agent should be enabled on your subscription)	Low
Azure Arc-enabled Kubernetes clusters should have the Defender extension installed	Defender's extension for Azure Arc provides threat protection for your Arc-enabled Kubernetes clusters. The extension collects data from all control plane (master) nodes in the cluster and sends it to the Microsoft Defender for Kubernetes backend in the cloud for further analysis. Learn more in /azure/defender-for-cloud/defender-for-kubernetes-azure-arc?wt.mc_id=defenderforcloud_inproduct_portal_recoremediation. (No related policy)	High
Azure Cache for Redis should reside within a virtual network	Azure Virtual Network (VNet) deployment provides enhanced security and isolation for your Azure Cache for Redis, as well as subnets, access control policies, and other features to further restrict access. When an Azure Cache for Redis instance is configured with a VNet, it is not publicly addressable and can only be accessed from virtual machines and applications within the VNet. (Related policy: Azure Cache for Redis should reside within a virtual network)	Medium
Azure Cosmos DB accounts should have firewall rules	Firewall rules should be defined on your Azure Cosmos DB accounts to prevent traffic from unauthorized sources. Accounts that have at least one IP rule defined with the virtual network filter enabled are deemed compliant. Accounts disabling public access are also deemed compliant. (Related policy: Azure Cosmos DB accounts should have firewall rules)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Azure Event Grid domains should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your Event Grid domains instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/privateendpoints. (Related policy: Azure Event Grid domains should use private link)	Medium
Azure Event Grid topics should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your topics instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/privateendpoints. (Related policy: Azure Event Grid topics should use private link)	Medium
Azure Kubernetes Service clusters should have Defender profile enabled	Microsoft Defender for Containers provides cloud-native Kubernetes security capabilities including environment hardening, workload protection, and run-time protection. When you enable the SecurityProfile.AzureDefender profile on your Azure Kubernetes Service cluster, an agent is deployed to your cluster to collect security event data. Learn more about Microsoft Defender for Containers. (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Azure Machine Learning workspaces should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your Azure Machine Learning workspaces instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/azureml-workspaces-privatelink. (Related policy: Azure Machine Learning workspaces should use private link)	Medium
Azure SignalR Service should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your SignalR resources instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/asrs/privatelink. (Related policy: Azure SignalR Service should use private link)	Medium
Azure Spring Cloud should use network injection	Azure Spring Cloud instances should use virtual network injection for the following purposes: 1. Isolate Azure Spring Cloud from Internet. 2. Enable Azure Spring Cloud to interact with systems in either on premises data centers or Azure service in other virtual networks. 3. Empower customers to control inbound and outbound network communications for Azure Spring Cloud. (Related policy: Azure Spring Cloud should use network injection)	Medium
Cognitive Services accounts should enable data encryption	This policy audits any Cognitive Services account not using data encryption. For each Cognitive Services account with storage, should enable data encryption with either customer managed or Microsoft managed key. (Related policy: Cognitive Services accounts should enable data encryption)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Cognitive Services accounts should restrict network access	Network access to Cognitive Services accounts should be restricted. Configure network rules so only applications from allowed networks can access the Cognitive Services account. To allow connections from specific internet or on-premises clients, access can be granted to traffic from specific Azure virtual networks or to public internet IP address ranges. (Related policy: Cognitive Services accounts should restrict network access)	Medium
Cognitive Services accounts should use customer owned storage or enable data encryption	This policy audits any Cognitive Services account not using customer owned storage nor data encryption. For each Cognitive Services account with storage, use either customer owned storage or enable data encryption. (Related policy: Cognitive Services accounts should use customer owned storage or enable data encryption.)	Low
Diagnostic logs in Azure Data Lake Store should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Azure Data Lake Store should be enabled)	Low
Diagnostic logs in Data Lake Analytics should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Data Lake Analytics should be enabled)	Low
Email notification for high severity alerts should be enabled	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, enable email notifications for high severity alerts in Defender for Cloud. (Related policy: Email notification for high severity alerts should be enabled)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Email notification to subscription owner for high severity alerts should be enabled	To ensure your subscription owners are notified when there is a potential security breach in their subscription, set email notifications to subscription owners for high severity alerts in Defender for Cloud. (Related policy: Email notification to subscription owner for high severity alerts should be enabled)	Medium
Enforce SSL connection should be enabled for MySQL database servers	Azure Database for MySQL supports connecting your Azure Database for MySQL server to client applications using Secure Sockets Layer (SSL). Enforcing SSL connections between your database server and your client applications helps protect against 'man in the middle' attacks by encrypting the data stream between the server and your application. This configuration enforces that SSL is always enabled for accessing your database server. (Related policy: Enforce SSL connection should be enabled for MySQL database servers)	Medium
Enforce SSL connection should be enabled for PostgreSQL database servers	Azure Database for PostgreSQL supports connecting your Azure Database for PostgreSQL server to client applications using Secure Sockets Layer (SSL). Enforcing SSL connections between your database server and your client applications helps protect against 'man in the middle' attacks by encrypting the data stream between the server and your application. This configuration enforces that SSL is always enabled for accessing your database server. (Related policy: Enforce SSL connection should be enabled for PostgreSQL database servers)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Geo-redundant backup should be enabled for Azure Database for MariaDB	Azure Database for MariaDB allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery options in case of a region failure. Configuring geo-redundant storage for backup is only allowed when creating a server. (Related policy: Geo-redundant backup should be enabled for Azure Database for MariaDB)	Low
Geo-redundant backup should be enabled for Azure Database for MySQL	Azure Database for MySQL allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery options in case of a region failure. Configuring geo-redundant storage for backup is only allowed when creating a server. (Related policy: Geo-redundant backup should be enabled for Azure Database for MySQL)	Low
Geo-redundant backup should be enabled for Azure Database for PostgreSQL	Azure Database for PostgreSQL allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery options in case of a region failure. Configuring geo-redundant storage for backup is only allowed when creating a server. (Related policy: Geo-redundant backup should be enabled for Azure Database for PostgreSQL)	Low
Kubernetes clusters should disable automounting API credentials	Disable automounting API credentials to prevent a potentially compromised Pod resource to run API commands against Kubernetes clusters. For more information, see https://aka.ms/kubepolicydoc. (Related policy: Kubernetes clusters should disable automounting API credentials)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Kubernetes clusters should not grant CAPSYSADMIN security capabilities	To reduce the attack surface of your containers, restrict CAP_SYS_ADMIN Linux capabilities. For more information, see https://aka.ms/kubepolicydoc. (No related policy)	High
Kubernetes clusters should not use the default namespace	Prevent usage of the default namespace in Kubernetes clusters to protect against unauthorized access for ConfigMap, Pod, Secret, Service, and ServiceAccount resource types. For more information, see https://aka.ms/kubepolicydoc. (Related policy: Kubernetes clusters should not use the default namespace)	Low
Microsoft Defender for Azure SQL Database servers should be enabled	Microsoft Defender for SQL is a unified package that provides advanced SQL security capabilities. It includes functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate a threat to your database, and discovering and classifying sensitive data. Important: Protections from this plan are charged as shown on the **Defender plans** page. If you don't have any Azure SQL Database servers in this subscription, you won't be charged. If you later create Azure SQL Database servers on this subscription, they'll automatically be protected and charges will begin. Learn about the pricing details per region. Learn more in Introduction to Microsoft Defender for SQL. (Related policy: Azure Defender for Azure SQL Database servers should be enabled)	High
Microsoft Defender for DNS should be enabled	Microsoft Defender for DNS provides an additional layer of protection for your cloud resources by continuously monitoring all DNS queries from your Azure resources. Defender for DNS alerts you about suspicious activity at the DNS layer. Learn more in Introduction to Microsoft Defender for DNS. Enabling this Defender plan results in charges. Learn about the pricing details per region on Defender for Cloud's pricing page: https://azure.microsoft.com/services/defender-for-cloud/#pricing. (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for open-source relational databases should be enabled	Microsoft Defender for open-source relational databases detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Learn more in Introduction to Microsoft Defender for open-source relational databases. Important: Enabling this plan will result in charges for protecting your open-source relational databases. If you don't have any open-source relational databases in this subscription, no charges will be incurred. If you create any open-source relational databases on this subscription in the future, they will automatically be protected and charges will begin at that time. (No related policy)	High
Microsoft Defender for Resource Manager should be enabled	Microsoft Defender for Resource Manager automatically monitors the resource management operations in your organization. Defender for Cloud detects threats and alerts you about suspicious activity. Learn more in Introduction to Microsoft Defender for Resource Manager. Enabling this Defender plan results in charges. Learn about the pricing details per region on Defender for Cloud's pricing page: https://azure.microsoft.com/services/defender-for-cloud/#pricing. (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for SQL on machines should be enabled on workspaces	Microsoft Defender for servers brings threat detection and advanced defenses for your Windows and Linux machines. With this Defender plan enabled on your subscriptions but not on your workspaces, you're paying for the full capability of Microsoft Defender for servers but missing out on some of the benefits. When you enable Microsoft Defender for servers on a workspace, all machines reporting to that workspace will be billed for Microsoft Defender for servers - even if they're in subscriptions without Defender plans enabled. Unless you also enable Microsoft Defender for servers on the subscription, those machines won't be able to take advantage of just-in-time VM access, adaptive application controls, and network detections for Azure resources. Learn more in Introduction to Microsoft Defender for servers. (No related policy)	Medium
Microsoft Defender for SQL servers on machines should be enabled	Microsoft Defender for SQL is a unified package that provides advanced SQL security capabilities. It includes functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate a threat to your database, and discovering and classifying sensitive data. Important: Remediating this recommendation will result in charges for protecting your SQL servers on machines. If you don't have any SQL servers on machines in this subscription, no charges will be incurred. If you create any SQL servers on machines on this subscription in the future, they will automatically be protected and charges will begin at that time. Learn more about Microsoft Defender for SQL servers on machines. (Related policy: Azure Defender for SQL servers on machines should be enabled)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for SQL should be enabled for unprotected Azure SQL servers	Microsoft Defender for SQL is a unified package that provides advanced SQL security capabilities. It surfaces and mitigates potential database vulnerabilities, and detects anomalous activities that could indicate a threat to your database. Microsoft Defender for SQL is billed as shown on pricing details per region. (Related policy: Advanced data security should be enabled on your SQL servers)	High
Microsoft Defender for SQL should be enabled for unprotected SQL Managed Instances	Microsoft Defender for SQL is a unified package that provides advanced SQL security capabilities. It surfaces and mitigates potential database vulnerabilities, and detects anomalous activities that could indicate a threat to your database. Microsoft Defender for SQL is billed as shown on pricing details per region. (Related policy: Advanced data security should be enabled on SQL Managed Instance)	High
Microsoft Defender for Storage should be enabled	Microsoft Defender for storage detects unusual and potentially harmful attempts to access or exploit storage accounts. Important: Protections from this plan are charged as shown on the **Defender plans** page. If you don't have any Azure Storage accounts in this subscription, you won't be charged. If you later create Azure Storage accounts on this subscription, they'll automatically be protected and charges will begin. Learn about the pricing details per region. Learn more in Introduction to Microsoft Defender for Storage. (Related policy: Azure Defender for Storage should be enabled)	High
Network Watcher should be enabled	Network Watcher is a regional service that enables you to monitor and diagnose conditions at a network scenario level in, to, and from Azure. Scenario level monitoring enables you to diagnose problems at an end-to-end network level view. Network diagnostic and visualization tools available with Network Watcher help you understand, diagnose, and gain insights to your network in Azure. (Related policy: Network Watcher should be enabled)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Private endpoint connections on Azure SQL Database should be enabled	Private endpoint connections enforce secure communication by enabling private connectivity to Azure SQL Database. (Related policy: Private endpoint connections on Azure SQL Database should be enabled)	Medium
Private endpoint should be enabled for MariaDB servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for MariaDB. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure. (Related policy: Private endpoint should be enabled for MariaDB servers)	Medium
Private endpoint should be enabled for MySQL servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for MySQL. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure. (Related policy: Private endpoint should be enabled for MySQL servers)	Medium
Private endpoint should be enabled for PostgreSQL servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for PostgreSQL. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure. (Related policy: Private endpoint should be enabled for PostgreSQL servers)	Medium
Public network access on Azure SQL Database should be disabled	Disabling the public network access property improves security by ensuring your Azure SQL Database can only be accessed from a private endpoint. This configuration denies all logins that match IP or virtual network based firewall rules. (Related policy: Public network access on Azure SQL Database should be disabled)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Public network access should be disabled for Cognitive Services accounts	This policy audits any Cognitive Services account in your environment with public network access enabled. Public network access should be disabled so that only connections from private endpoints are allowed. (Related policy: Public network access should be disabled for Cognitive Services accounts)	Medium
Public network access should be disabled for MariaDB servers	Disable the public network access property to improve security and ensure your Azure Database for MariaDB can only be accessed from a private endpoint. This configuration strictly disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules. (Related policy: Public network access should be disabled for MariaDB servers)	Medium
Public network access should be disabled for MySQL servers	Disable the public network access property to improve security and ensure your Azure Database for MySQL can only be accessed from a private endpoint. This configuration strictly disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules. (Related policy: Public network access should be disabled for MySQL servers)	Medium
Public network access should be disabled for PostgreSQL servers	Disable the public network access property to improve security and ensure your Azure Database for PostgreSQL can only be accessed from a private endpoint. This configuration disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules. (Related policy: Public network access should be disabled for PostgreSQL servers)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Redis Cache should allow access only via SSL	Enable only connections via SSL to Redis Cache. Use of secure connections ensures authentication between the server and the service and protects data in transit from network layer attacks such as man-in-the-middle, eavesdropping, and session-hijacking. (Related policy: Only secure connections to your Azure Cache for Redis should be enabled)	High
SQL databases should have vulnerability findings resolved	SQL Vulnerability assessment scans your database for security vulnerabilities, and exposes any deviations from best practices such as misconfigurations, excessive permissions, and unprotected sensitive data. Resolving the vulnerabilities found can greatly improve your database security posture. Learn more (Related policy: Vulnerabilities on your SQL databases should be remediated)	High
SQL managed instances should have vulnerability assessment configured	Vulnerability assessment can discover, track, and help you remediate potential database vulnerabilities. (Related policy: Vulnerability assessment should be enabled on SQL Managed Instance)	High
SQL servers on machines should have vulnerability findings resolved	SQL Vulnerability assessment scans your database for security vulnerabilities, and exposes any deviations from best practices such as misconfigurations, excessive permissions, and unprotected sensitive data. Resolving the vulnerabilities found can greatly improve your database security posture. Learn more (Related policy: Vulnerabilities on your SQL servers on machine should be remediated)	High
SQL servers should have an Azure Active Directory administrator provisioned	Provision an Azure AD administrator for your SQL server to enable Azure AD authentication. Azure AD authentication enables simplified permission management and centralized identity management of database users and other Microsoft services. (Related policy: An Azure Active Directory administrator should be provisioned for SQL servers)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
SQL servers should have vulnerability assessment configured	Vulnerability assessment can discover, track, and help you remediate potential database vulnerabilities. (Related policy: Vulnerability assessment should be enabled on your SQL servers)	High
Storage account should use a private link connection	Private links enforce secure communication, by providing private connectivity to the storage account (Related policy: Storage account should use a private link connection)	Medium
Storage accounts should be migrated to new Azure Resource Manager resources	To benefit from new capabilities in Azure Resource Manager, you can migrate existing deployments from the Classic deployment model. Resource Manager enables security enhancements such as: stronger access control (RBAC), better auditing, ARM-based deployment and governance, access to managed identities, access to key vault for secrets, Azure AD-based authentication and support for tags and resource groups for easier security management. Learn more (Related policy: Storage accounts should be migrated to new Azure Resource Manager resources)	Low
Storage accounts should restrict network access using virtual network rules	Protect your storage accounts from potential threats using virtual network rules as a preferred method instead of IP-based filtering. Disabling IP-based filtering prevents public IPs from accessing your storage accounts. (Related policy: Storage accounts should restrict network access using virtual network rules)	Medium
Subscriptions should have a contact email address for security issues	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, set a security contact to receive email notifications from Defender for Cloud. (Related policy: Subscriptions should have a contact email address for security issues)	Low
Transparent Data Encryption on SQL databases should be enabled	Enable transparent data encryption to protect data-at-rest and meet compliance requirements (Related policy: Transparent Data Encryption on SQL databases should be enabled)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
VM Image Builder templates should use private link	Audit VM Image Builder templates that do not have a virtual network configured. When a virtual network is not configured, a public IP is created and used instead, which may directly expose resources to the internet and increase the potential attack surface. (Related policy: VM Image Builder templates should use private link)	Medium
Web Application Firewall (WAF) should be enabled for Application Gateway	Deploy Azure Web Application Firewall (WAF) in front of public facing web applications for additional inspection of incoming traffic. Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities such as SQL injections, Cross-Site Scripting, local and remote file executions. You can also restrict access to your web applications by countries, IP address ranges, and other http(s) parameters via custom rules. (Related policy: Web Application Firewall (WAF) should be enabled for Application Gateway)	Low
Web Application Firewall (WAF) should be enabled for Azure Front Door Service service	Deploy Azure Web Application Firewall (WAF) in front of public facing web applications for additional inspection of incoming traffic. Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities such as SQL injections, Cross-Site Scripting, local and remote file executions. You can also restrict access to your web applications by countries, IP address ranges, and other http(s) parameters via custom rules. (Related policy: Web Application Firewall (WAF) should be enabled for Azure Front Door Serviceservice)	Low

IdentityAndAccess recommendations

There are 29 recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
A maximum of 3 owners should be designated for subscriptions	To reduce the potential for breaches by compromised owner accounts, we recommend limiting the number of owner accounts to a maximum of 3 (Related policy: A maximum of 3 owners should be designated for your subscription)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Deprecated accounts should be removed from subscriptions	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: Deprecated accounts should be removed from your subscription)	High
Deprecated accounts should be removed from your subscription	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: Deprecated accounts should be removed from your subscription)	High
Deprecated accounts with owner permissions should be removed from subscriptions	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: Deprecated accounts with owner permissions should be removed from your subscription)	High
Deprecated accounts with owner permissions should be removed from your subscription	User accounts that have been blocked from signing in, should be removed from your subscriptions. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: Deprecated accounts with owner permissions should be removed from your subscription)	High
Diagnostic logs in Key Vault should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in Key Vault should be enabled)	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
External accounts with owner permissions should be removed from subscriptions	Accounts with owner permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with owner permissions should be removed from your subscription)	High
External accounts with owner permissions should be removed from your subscription	Accounts with owner permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with owner permissions should be removed from your subscription)	High
External accounts with read permissions should be removed from subscriptions	Accounts with read permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with read permissions should be removed from your subscription)	High
External accounts with read permissions should be removed from your subscription	Accounts with read permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with read permissions should be removed from your subscription)	High
External accounts with write permissions should be removed from subscriptions	Accounts with write permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with write permissions should be removed from your subscription)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
External accounts with write permissions should be removed from your subscription	Accounts with write permissions that have different domain names (external accounts), should be removed from your subscription. This prevents unmonitored access. These accounts can be targets for attackers looking to find ways to access your data without being noticed. (Related policy: External accounts with write permissions should be removed from your subscription)	High
Firewall should be enabled on Key Vault	Key vault's firewall prevents unauthorized traffic from reaching your key vault and provides an additional layer of protection for your secrets. Enable the firewall to make sure that only traffic from allowed networks can access your key vault. (Related policy: Firewall should be enabled on Key Vault)	Medium
Key Vault keys should have an expiration date	Cryptographic keys should have a defined expiration date and not be permanent. Keys that are valid forever provide a potential attacker with more time to compromise the key. It is a recommended security practice to set expiration dates on cryptographic keys. (Related policy: Key Vault keys should have an expiration date)	High
Key Vault secrets should have an expiration date	Secrets should have a defined expiration date and not be permanent. Secrets that are valid forever provide a potential attacker with more time to compromise them. It is a recommended security practice to set expiration dates on secrets. (Related policy: Key Vault secrets should have an expiration date)	High
Key vaults should have purge protection enabled	Malicious deletion of a key vault can lead to permanent data loss. A malicious insider in your organization can potentially delete and purge key vaults. Purge protection protects you from insider attacks by enforcing a mandatory retention period for soft deleted key vaults. No one inside your organization or Microsoft will be able to purge your key vaults during the soft delete retention period. (Related policy: Key vaults should have purge protection enabled)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Key vaults should have soft delete enabled	Deleting a key vault without soft delete enabled permanently deletes all secrets, keys, and certificates stored in the key vault. Accidental deletion of a key vault can lead to permanent data loss. Soft delete allows you to recover an accidentally deleted key vault for a configurable retention period. (Related policy: Key vaults should have soft delete enabled)	High
MFA should be enabled on accounts with owner permissions on subscriptions	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with owner permissions to prevent a breach of accounts or resources. (Related policy: MFA should be enabled on accounts with owner permissions on your subscription)	High
MFA should be enabled on accounts with owner permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with owner permissions to prevent a breach of accounts or resources. (Related policy: MFA should be enabled on accounts with owner permissions on your subscription)	High
MFA should be enabled on accounts with read permissions on subscriptions	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with read privileges to prevent a breach of accounts or resources. (Related policy: MFA should be enabled on accounts with read permissions on your subscription)	High
MFA should be enabled on accounts with read permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with read privileges to prevent a breach of accounts or resources. (Related policy: MFA should be enabled on accounts with read permissions on your subscription)	High
MFA should be enabled on accounts with write permissions on subscriptions	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with write privileges to prevent a breach of accounts or resources. (Related policy: MFA should be enabled accounts with write permissions on your subscription)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
MFA should be enabled on accounts with write permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with write privileges to prevent a breach of accounts or resources. (Related policy: MFA should be enabled accounts with write permissions on your subscription)	High
Microsoft Defender for Key Vault should be enabled	Microsoft Defender for Cloud includes Microsoft Defender for Key Vault, providing an additional layer of security intelligence. Microsoft Defender for Key Vault detects unusual and potentially harmful attempts to access or exploit Key Vault accounts. Important: Protections from this plan are charged as shown on the **Defender plans** page. If you don't have any key vaults in this subscription, you won't be charged. If you later create key vaults on this subscription, they'll automatically be protected and charges will begin. Learn about the pricing details per region. Learn more in Introduction to Microsoft Defender for Key Vault. (Related policy: Azure Defender for Key Vault should be enabled)	High
Private endpoint should be configured for Key Vault	Private link provides a way to connect Key Vault to your Azure resources without sending traffic over the public internet. Private link provides defense in depth protection against data exfiltration. (Related policy: Private endpoint should be configured for Key Vault)	Medium
Service principals should be used to protect your subscriptions instead of Management Certificates	Management certificates allow anyone who authenticates with them to manage the subscription(s) they are associated with. To manage subscriptions more securely, using service principals with Resource Manager is recommended to limit the blast radius in the case of a certificate compromise. It also automates resource management. (Related policy: Service principals should be used to protect your subscriptions instead of management certificates)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Storage account public access should be disallowed	Anonymous public read access to containers and blobs in Azure Storage is a convenient way to share data, but might present security risks. To prevent data breaches caused by undesired anonymous access, Microsoft recommends preventing public access to a storage account unless your scenario requires it. (Related policy: Storage account public access should be disallowed)	Medium
There should be more than one owner assigned to subscriptions	Designate more than one subscription owner in order to have administrator access redundancy. (Related policy: There should be more than one owner assigned to your subscription)	High
Validity period of certificates stored in Azure Key Vault should not exceed 12 months	Ensure your certificates do not have a validity period that exceeds 12 months. (Related policy: Certificates should have the specified maximum validity period)	Medium

IoT recommendations

There are 12 recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Default IP Filter Policy should be Deny	IP Filter Configuration should have rules defined for allowed traffic and should deny all other traffic by default (No related policy)	Medium
Diagnostic logs in IoT Hub should be enabled	Enable logs and retain them for up to a year. This enables you to recreate activity trails for investigation purposes when a security incident occurs or your network is compromised. (Related policy: Diagnostic logs in IoT Hub should be enabled)	Low
Identical Authentication Credentials	Identical authentication credentials to the IoT Hub used by multiple devices. This could indicate an illegitimate device impersonating a legitimate device. It also exposes the risk of device impersonation by an attacker (No related policy)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
IoT Devices - Agent sending underutilized messages	IoT agent message size capacity is currently underutilized, causing an increase in the number of sent messages. Adjust message intervals for better utilization (No related policy)	Low
IoT Devices - Auditd process stopped sending events	Security events originated from Auditd process are no longer received from this device (No related policy)	High
IoT Devices - Open Ports On Device	A listening endpoint was found on the device (No related policy)	Medium
IoT Devices - Operating system baseline validation failure	Security-related system configuration issues identified (No related policy)	Medium
IoT Devices - Permissive firewall policy in one of the chains was found	An allowed firewall policy was found in main firewall Chains (INPUT/OUTPUT). The policy should Deny all traffic by default define rules to allow necessary communication to/from the device (No related policy)	Medium
IoT Devices - Permissive firewall rule in the input chain was found	A rule in the firewall has been found that contains a permissive pattern for a wide range of IP addresses or Ports (No related policy)	Medium
IoT Devices - Permissive firewall rule in the output chain was found	A rule in the firewall has been found that contains a permissive pattern for a wide range of IP addresses or ports (No related policy)	Medium
IoT Devices - TLS cipher suite upgrade needed	Unsecure TLS configurations detected. Immediate TLS cipher suite upgrade recommended (No related policy)	Medium
IP Filter rule large IP range	An Allow IP Filter rule's source IP range is too large. Overly permissive rules might expose your IoT hub to malicious intenders (No related policy)	Medium

Networking recommendations

There are **14** recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Access to storage accounts with firewall and virtual network configurations should be restricted	Review the settings of network access in your storage account firewall settings. We recommended configuring network rules so that only applications from allowed networks can access the storage account. To allow connections from specific internet or on-premise clients, access can be granted to traffic from specific Azure virtual networks or to public internet IP address ranges. (Related policy: Storage accounts should restrict network access)	Low
Adaptive network hardening recommendations should be applied on internet facing virtual machines	Defender for Cloud has analyzed the internet traffic communication patterns of the virtual machines listed below, and determined that the existing rules in the NSGs associated to them are overly-permissive, resulting in an increased potential attack surface. This typically occurs when this IP address doesn't communicate regularly with this resource. Alternatively, the IP address has been flagged as malicious by Defender for Cloud's threat intelligence sources. Learn more in Improve your network security posture with adaptive network hardening. (Related policy: Adaptive network hardening recommendations should be applied on internet facing virtual machines)	High
All network ports should be restricted on network security groups associated to your virtual machine	Defender for Cloud has identified some of your network security groups' inbound rules to be too permissive. Inbound rules should not allow access from 'Any' or 'Internet' ranges. This can potentially enable attackers to target your resources. (Related policy: All network ports should be restricted on network security groups associated to your virtual machine)	High
Azure DDoS Protection Standard should be enabled	Defender for Cloud has discovered virtual networks with Application Gateway resources unprotected by the DDoS protection service. These resources contain public IPs. Enable mitigation of network volumetric and protocol attacks. (Related policy: Azure DDoS Protection Standard should be enabled)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Internet-facing virtual machines should be protected with network security groups	Protect your VM from potential threats by restricting access to it with a network security group (NSG). NSGs contain a list of Access Control List (ACL) rules that allow or deny network traffic to your VM from other instances, in or outside the same subnet. To keep your machine as secure as possible, the VM access to the internet must be restricted and an NSG should be enabled on the subnet. VMs with 'High' severity are internet-facing VMs. (Related policy: Internet-facing virtual machines should be protected with network security groups)	High
IP forwarding on your virtual machine should be disabled	Defender for Cloud has discovered that IP forwarding is enabled on some of your virtual machines. Enabling IP forwarding on a virtual machine's NIC allows the machine to receive traffic addressed to other destinations. IP forwarding is rarely required (e.g., when using the VM as a network virtual appliance), and therefore, this should be reviewed by the network security team. (Related policy: IP Forwarding on your virtual machine should be disabled)	Medium
Management ports of virtual machines should be protected with just-in-time network access control	Defender for Cloud has identified some overly-permissive inbound rules for management ports in your Network Security Group. Enable just-in-time access control to protect your VM from internet-based brute-force attacks. Learn more in Understanding just-in-time (JIT) VM access. (Related policy: Management ports of virtual machines should be protected with just-in-time network access control)	High
Management ports should be closed on your virtual machines	Open remote management ports are exposing your VM to a high level of risk from Internet-based attacks. These attacks attempt to brute force credentials to gain admin access to the machine. (Related policy: Management ports should be closed on your virtual machines)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Network traffic data collection agent should be installed on Linux virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations and specific network threats. (Related policy: Network traffic data collection agent should be installed on Linux virtual machines)	Medium
Network traffic data collection agent should be installed on Windows virtual machines	Defender for Cloud uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations, and specific network threats. (Related policy: Network traffic data collection agent should be installed on Windows virtual machines)	Medium
Non-internet-facing virtual machines should be protected with network security groups	Protect your non-internet-facing virtual machine from potential threats by restricting access to it with a network security group (NSG). NSGs contain a list of Access Control List (ACL) rules that allow or deny network traffic to your VM from other instances, whether or not they're on the same subnet. Note that to keep your machine as secure as possible, the VM's access to the internet must be restricted and an NSG should be enabled on the subnet. (Related policy: Non-internet-facing virtual machines should be protected with network security groups)	Low
Secure transfer to storage accounts should be enabled	Secure transfer is an option that forces your storage account to accept requests only from secure connections (HTTPS). Use of HTTPS ensures authentication between the server and the service and protects data in transit from network layer attacks such as man-in-the-middle, eavesdropping, and session-hijacking. (Related policy: Secure transfer to storage accounts should be enabled)	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Subnets should be associated with a network security group	Protect your subnet from potential threats by restricting access to it with a network security group (NSG). NSGs contain a list of Access Control List (ACL) rules that allow or deny network traffic to your subnet. When an NSG is associated with a subnet, the ACL rules apply to all the VM instances and integrated services in that subnet, but don't apply to internal traffic inside the subnet. To secure resources in the same subnet from one another, enable NSG directly on the resources as well. (Related policy: Subnets should be associated with a Network Security Group)	Low
Virtual networks should be protected by Azure Firewall	Some of your virtual networks aren't protected with a firewall. Use Azure Firewall to restrict access to your virtual networks and prevent potential threats. Learn more about Azure Firewall. (Related policy: All Internet traffic should be routed via your deployed Azure Firewall)	Low

Deprecated recommendations

RECOMMENDATION	DESCRIPTION & RELATED POLICY	SEVERITY
Access to App Services should be restricted	Restrict access to your App Services by changing the networking configuration, to deny inbound traffic from ranges that are too broad. (Related policy: [Preview]: Access to App Services should be restricted)	High
The rules for web applications on IaaS NSGs should be hardened	Harden the network security group (NSG) of your virtual machines that are running web applications, with NSG rules that are overly permissive with regard to web application ports. (Related policy: The NSGs rules for web applications on IaaS should be hardened)	High

RECOMMENDATION	DESCRIPTION & RELATED POLICY	SEVERITY
Pod Security Policies should be defined to reduce the attack vector by removing unnecessary application privileges (Preview)	Define Pod Security Policies to reduce the attack vector by removing unnecessary application privileges. It is recommended to configure pod security policies so pods can only access resources which they are allowed to access. (Related policy: [Preview]: Pod Security Policies should be defined on Kubernetes Services)	Medium
Install Azure Security Center for IoT security module to get more visibility into your IoT devices	Install Azure Security Center for IoT security module to get more visibility into your IoT devices.	Low
Your machines should be restarted to apply system updates	Restart your machines to apply the system updates and secure the machine from vulnerabilities. (Related policy: System updates should be installed on your machines)	Medium
Monitoring agent should be installed on your machines	This action installs a monitoring agent on the selected virtual machines. Select a workspace for the agent to report to. (No related policy)	High

Next steps

To learn more about recommendations, see the following:

- What are security policies, initiatives, and recommendations?
- Review your security recommendations

Security recommendations for AWS resources - a reference guide

2/15/2022 • 74 minutes to read • Edit Online

This article lists the recommendations you might see in Microsoft Defender for Cloud if you've connected an AWS account from the **Environment settings** page. The recommendations shown in your environment depend on the resources you're protecting and your customized configuration.

To learn about how to respond to these recommendations, see Remediate recommendations in Defender for Cloud.

Your secure score is based on the number of security recommendations you've completed. To decide which recommendations to resolve first, look at the severity of each one and its potential impact on your secure score.

AWS Compute recommendations

There are **18** AWS recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon EC2 instances managed by Systems Manager should have a patch compliance status of COMPLIANT after a patch installation	This control checks whether the compliance status of the Amazon EC2 Systems Manager patch compliance is COMPLIANT or NON_COMPLIANT after the patch installation on the instance. It only checks instances that are managed by AWS Systems Manager Patch Manager. It does not check whether the patch was applied within the 30-day limit prescribed by PCI DSS requirement '6.2'. It also does not validate whether the patches applied were classified as security patches. You should create patching groups with the appropriate baseline settings and ensure in-scope systems are managed by those patch groups in Systems Manager. For more information about patch groups, see the AWS Systems Manager User Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon EFS should be configured to encrypt file data at rest using AWS KMS	This control checks whether Amazon Elastic File System is configured to encrypt the file data using AWS KMS. The check fails in the following cases: *"Encrypted" is set to "false" in the DescribeFileSystems response. The "KmsKeyId" key in the DescribeFileSystems response does not match the KmsKeyId parameter for efs-encrypted-check. Note that this control does not use the "KmsKeyId" parameter for efs-encrypted-check. It only checks the value of "Encrypted". For an added layer of security for your sensitive data in Amazon EFS, you should create encrypted file systems. Amazon EFS supports encryption for file systems at-rest. You can enable encryption of data at rest when you create an Amazon EFS file system. To learn more about Amazon EFS encryption, see Data encryption in Amazon EFS in the Amazon Elastic File System User Guide.	Medium
Amazon EFS volumes should be in backup plans	This control checks whether Amazon Elastic File System (Amazon EFS) file systems are added to the backup plans in AWS Backup. The control fails if Amazon EFS file systems are not included in the backup plans. Including EFS file systems in the backup plans helps you to protect your data from deletion and data loss.	Medium
Application Load Balancer deletion protection should be enabled	This control checks whether an Application Load Balancer has deletion protection enabled. The control fails if deletion protection is not configured. Enable deletion protection to protect your Application Load Balancer from deletion.	Medium
Auto Scaling groups associated with a load balancer should use health checks	Auto Scaling groups that are associated with a load balancer are using Elastic Load Balancing health checks. PCI DSS does not require load balancing or highly available configurations. This is recommended by AWS best practices.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
AWS accounts should have Azure Arc auto provisioning enabled	For full visibility of the security content from Microsoft Defender for servers, EC2 instances should be connected to Azure Arc. To ensure that all eligible EC2 instances automatically receive Azure Arc, enable auto-provisioning from Defender for Cloud at the AWS account level. Learn more about Azure Arc, and Microsoft Defender for Servers.	High
CloudFront distributions should have origin failover configured	This control checks whether an Amazon CloudFront distribution is configured with an origin group that has two or more origins. CloudFront origin failover can increase availability. Origin failover automatically redirects traffic to a secondary origin if the primary origin is unavailable or if it returns specific HTTP response status codes.	Medium
CodeBuild GitHub or Bitbucket source repository URLs should use OAuth	This control checks whether the GitHub or Bitbucket source repository URL contains either personal access tokens or a user name and password. Authentication credentials should never be stored or transmitted in clear text or appear in the repository URL. Instead of personal access tokens or user name and password, you should use OAuth to grant authorization for accessing GitHub or Bitbucket repositories. Using personal access tokens or a user name and password could expose your credentials to unintended data exposure and unauthorized access.	High
CodeBuild project environment variables should not contain credentials	This control checks whether the project contains the environment variables `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY`. Authentication credentials `AWS_ACCESS_KEY_ID` and `AWS_SECRET_ACCESS_KEY` should never be stored in clear text, as this could lead to unintended data exposure and unauthorized access.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
DynamoDB Accelerator (DAX) clusters should be encrypted at rest	This control checks whether a DAX cluster is encrypted at rest. Encrypting data at rest reduces the risk of data stored on disk being accessed by a user not authenticated to AWS. The encryption adds another set of access controls to limit the ability of unauthorized users to access to the data. For example, API permissions are required to decrypt the data before it can be read.	Medium
DynamoDB tables should automatically scale capacity with demand	This control checks whether an Amazon DynamoDB table can scale its read and write capacity as needed. This control passes if the table uses either on-demand capacity mode or provisioned mode with auto scaling configured. Scaling capacity with demand avoids throttling exceptions, which helps to maintain availability of your applications.	Medium
EC2 instances should be connected to Azure Arc	Connect your EC2 instances to Azure Arc in order to have full visibility to Microsoft Defender for Servers security content. Learn more about Azure Arc, and about Microsoft Defender for Servers on hybrid-cloud environment.	High
EC2 instances should be managed by AWS Systems Manager	Status of the Amazon EC2 Systems Manager patch compliance is 'COMPLIANT' or 'NON_COMPLIANT' after the patch installation on the instance. Only instances that are managed by AWS Systems Manager Patch Manager are checked. Patches that were applied within the 30-day limit prescribed by PCI DSS requirement '6' are not checked.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Instances managed by Systems Manager should have an association compliance status of COMPLIANT	This control checks whether the status of the AWS Systems Manager association compliance is COMPLIANT or NON_COMPLIANT after the association is run on an instance. The control passes if the association compliance status is COMPLIANT. A State Manager association is a configuration that is assigned to your managed instances. The configuration defines the state that you want to maintain on your instances. For example, an association can specify that antivirus software must be installed and running on your instances, or that certain ports must be closed. After you create one or more State Manager associations, compliance status information is immediately available to you in the console or in response to AWS CLI commands or corresponding Systems Manager API operations. For associations, "Configuration" Compliance shows statuses of Compliant or Non-compliant and the severity level assigned to the association, such as "Critical" or "Medium". To learn more about State Manager association compliance, see About About State Manager association compliance in the AWS Systems Manager User Guide. You must configure your in-scope EC2 instances for Systems Manager association. You must also configure the patch baseline for the security rating of the vendor of patches, and set the autoapproval date to meet PCI DSS '3.2.1' requirement '6.2'. For additional guidance on how to Create an association, see Create an association in the AWS Systems Manager User Guide. For additional information on working with patching in Systems Manager, see AWS Systems Manager Patch Manager in the AWS Systems Manager User Guide.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Lambda functions should have a dead-letter queue configured	This control checks whether a Lambda function is configured with a dead-letter queue. The control fails if the Lambda function is not configured with a dead-letter queue. As an alternative to an on-failure destination, you can configure your function with a dead-letter queue to save discarded events for further processing. A dead-letter queue acts the same as an on-failure destination. It is used when an event fails all processing attempts or expires without being processed. A dead-letter queue allows you to look back at errors or failed requests to your Lambda function to debug or identify unusual behavior. From a security perspective, it is important to understand why your function failed and to ensure that your function does not drop data or compromise data security as a result. For example, if your function cannot communicate to an underlying resource, that could be a symptom of a denial of service (DoS) attack elsewhere in the network.	Medium
Lambda functions should use supported runtimes	This control checks that the Lambda function settings for runtimes match the expected values set for the supported runtimes for each language. This control checks for the following runtimes: **nodejs14.x**, **nodejs12.x**, **nodejs10.x**, **python3.8**, **python3.7**, **python3.6**, **ruby2.7**, **ruby2.5**, **java11**, **java8**, **java8.al2**, **go1.x**, **dotnetcore3.1**, **dotnetcore2.1** Lambda runtimes are built around a combination of operating system, programming language, and software libraries that are subject to maintenance and security updates. When a runtime component is no longer supported for security updates, Lambda deprecates the runtime. Even though you cannot create functions that use the deprecated runtime, the function is still available to process invocation events. Make sure that your Lambda functions are current and do not use out-of-date runtime environments. To learn more about the supported runtimes that this control checks for the supported languages, see AWS Lambda runtimes in the AWS Lambda Developer Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Management ports of EC2 instances should be protected with just-in-time network access control	Microsoft Defender for Cloud has identified some overly-permissive inbound rules for management ports in your network. Enable just-in-time access control to protect your Instances from internet-based brute-force attacks. Learn more.	High
Unused EC2 security groups should be removed	Security groups should be attached to Amazon EC2 instances or to an ENI. healthy finding can indicate there are unused Amazon EC2 security groups.	Low

AWS Container recommendations

There are **3** AWS recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
EKS clusters should grant the required AWS permissions to Microsoft Defender for Cloud	Microsoft Defender for Containers provides protections for your EKS clusters. To monitor your cluster for security vulnerabilities and threats, Defender for Containers needs permissions for your AWS account. These permissions will be used to enable Kubernetes control plane logging on your cluster and establish a reliable pipeline between your cluster and Defender for Cloud's backend in the cloud. Learn more about Microsoft Defender for Cloud's security features for containerized environments.	High
EKS clusters should have Microsoft Defender's extension for Azure Arc installed	Microsoft Defender's cluster extension provides security capabilities for your EKS clusters. The extension collects data from a cluster and its nodes to identify security vulnerabilities and threats. The extension works with Azure Arc-enabled Kubernetes. If your cluster isn't connected to Azure Arc-enabled Kubernetes, connect it as described in the remediation steps. Learn more about Microsoft Defender for Cloud's security features for containerized environments.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Microsoft Defender for Containers should be enabled on AWS connectors	Microsoft Defender for Containers provides real-time threat protection for containerized environments and generates alerts about suspicious activities. Use this information to harden the security of Kubernetes clusters and remediate security issues. Important: When you've enabled Microsoft Defender for Containers and deployed Azure Arc to your EKS clusters, the protections - and charges - will begin. If you don't deploy Azure Arc on a cluster, Defender for Containers will not protect it and no charges will be incurred for this Microsoft Defender plan for that cluster.	High

AWS Data recommendations

There are 61 AWS recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon Aurora clusters should have backtracking enabled	This control checks whether Amazon Aurora clusters have backtracking enabled. Backups help you to recover more quickly from a security incident. They also strengthens the resilience of your systems. Aurora backtracking reduces the time to recover a database to a point in time. It does not require a database restore to do so. For more information about backtracking in Aurora, see Backtracking an Aurora DB cluster in the Amazon Aurora User Guide.	Medium
Amazon EBS snapshots should not be publicly restorable	Amazon EBS snapshots should not be publicly restorable by everyone unless explicitly allowed, to avoid accidental exposure of data. Additionally, permission to change Amazon EBS configurations should be restricted to authorized AWS accounts only.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon ECS task definitions should have secure networking modes and user definitions	This control checks whether an active Amazon ECS task definition that has host networking mode also has privileged or user container definitions. The control fails for task definitions that have host network mode and container definitions where privileged=false or is empty and user=root or is empty. If a task definition has elevated privileges, it is because the customer has specifically opted in to that configuration. This control checks for unexpected privilege escalation when a task definition has host networking enabled but the customer has not opted in to elevated privileges.	High
Amazon Elasticsearch Service domains should encrypt data sent between nodes	This control checks whether Amazon ES domains have node-to-node encryption enabled. HTTPS (TLS) can be used to help prevent potential attackers from eavesdropping on or manipulating network traffic using person-in-the-middle or similar attacks. Only encrypted connections over HTTPS (TLS) should be allowed. Enabling node-to-node encryption for Amazon ES domains ensures that intra-cluster communications are encrypted in transit. There can be a performance penalty associated with this configuration. You should be aware of and test the performance trade-off before enabling this option.	Medium
Amazon Elasticsearch Service domains should have encryption at rest enabled	It is important to enable encryptions rest of Amazon ES domains to protect sensitive data	Medium
Amazon Redshift clusters should have audit logging enabled	This control checks whether an Amazon Redshift cluster has audit logging enabled. Amazon Redshift audit logging provides additional information about connections and user activities in your cluster. This data can be stored and secured in Amazon S3 and can be helpful in security audits and investigations. For more information, see Database audit logging in the *Amazon Redshift Cluster Management Guide*.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon Redshift clusters should have automatic snapshots enabled	This control checks whether Amazon Redshift clusters have automated snapshots enabled. It also checks whether the snapshot retention period is greater than or equal to seven. Backups help you to recover more quickly from a security incident. They strengthen the resilience of your systems. Amazon Redshift takes periodic snapshots by default. This control checks whether automatic snapshots are enabled and retained for at least seven days. For more details on Amazon Redshift automated snapshots, see Automated snapshots in the *Amazon Redshift Cluster Management Guide*.	Medium
Amazon Redshift clusters should prohibit public access	We recommend Amazon Redshift clusters to avoid public accessibility by evaluating the 'publiclyAccessible' field in the cluster configuration item.	High
Amazon Redshift should have automatic upgrades to major versions enabled	This control checks whether automatic major version upgrades are enabled for the Amazon Redshift cluster. Enabling automatic major version upgrades ensures that the latest major version updates to Amazon Redshift clusters are installed during the maintenance window. These updates might include security patches and bug fixes. Keeping up to date with patch installation is an important step in securing systems.	Medium
Amazon SQS queues should be encrypted at rest	This control checks whether Amazon SQS queues are encrypted at rest. Server-side encryption (SSE) allows you to transmit sensitive data in encrypted queues. To protect the content of messages in queues, SSE uses keys managed in AWS KMS. For more information, see Encryption at rest in the Amazon Simple Queue Service Developer Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
An RDS event notifications subscription should be configured for critical cluster events	This control checks whether an Amazon RDS event subscription exists that has notifications enabled for the following source type, event category key-value pairs. DBCluster: ["maintenance" and "failure"]. RDS event notifications uses Amazon SNS to make you aware of changes in the availability or configuration of your RDS resources. These notifications allow for rapid response. For additional information about RDS event notifications, see Using Amazon RDS event notification in the Amazon RDS User Guide.	Low
An RDS event notifications subscription should be configured for critical database instance events	This control checks whether an Amazon RDS event subscription exists with notifications enabled for the following source type. event category key-value pairs. DBInstance: ["maintenance", "configuration change" and "failure"]. RDS event notifications use Amazon SNS to make you aware of changes in the availability or configuration of your RDS resources. These notifications allow for rapid response. For additional information about RDS event notifications, see Using Amazon RDS event notification in the Amazon RDS User Guide.	Low
An RDS event notifications subscription should be configured for critical database parameter group events	This control checks whether an Amazon RDS event subscription exists with notifications enabled for the following source type. event category key-value pairs. DBParameterGroup: ["configuration","change"]. RDS event notifications use Amazon SNS to make you aware of changes in the availability or configuration of your RDS resources. These notifications allow for rapid response. For additional information about RDS event notifications, see Using Amazon RDS event notification in the Amazon RDS User Guide.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
An RDS event notifications subscription should be configured for critical database security group events	This control checks whether an Amazon RDS event subscription exists with notifications enabled for the following source type, event category key-value pairs.DBSecurityGroup: ["configuration","change","failure"]. RDS event notifications use Amazon SNS to make you aware of changes in the availability or configuration of your RDS resources. These notifications allow for a rapid response. For additional information about RDS event notifications , see Using Amazon RDS event notification in the Amazon RDS User Guide.	Low
API Gateway REST and WebSocket API logging should be enabled	This control checks whether all stages of an Amazon API Gateway REST or WebSocket API have logging enabled. The control fails if logging is not enabled for all methods of a stage or if logging Level is neither ERROR nor INFO. API Gateway REST or WebSocket API stages should have relevant logs enabled. API Gateway REST and WebSocket API execution logging provides detailed records of requests made to API Gateway REST and WebSocket API stages. The stages include API integration backend responses, Lambda authorizer responses, and the requestId for AWS integration endpoints.	Medium
API Gateway REST API cache data should be encrypted at rest	This control checks whether all methods in API Gateway REST API stages that have cache enabled are encrypted. The control fails if any method in an API Gateway REST API stage is configured to cache and the cache is not encrypted. Encrypting data at rest reduces the risk of data stored on disk being accessed by a user not authenticated to AWS. It adds another set of access controls to limit unauthorized users ability access the data. For example, API permissions are required to decrypt the data before it can be read. API Gateway REST API caches should be encrypted at rest for an added layer of security.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
API Gateway REST API stages should be configured to use SSL certificates for backend authentication	This control checks whether Amazon API Gateway REST API stages have SSL certificates configured. Backend systems use these certificates to authenticate that incoming requests are from API Gateway. API Gateway REST API stages should be configured with SSL certificates to allow backend systems to authenticate that requests originate from API Gateway.	Medium
API Gateway REST API stages should have AWS X-Ray tracing enabled	This control checks whether AWS X-Ray active tracing is enabled for your Amazon API Gateway REST API stages. X-Ray active tracing enables a more rapid response to performance changes in the underlying infrastructure. Changes in performance could result in a lack of availability of the API. X-Ray active tracing provides real-time metrics of user requests that flow through your API Gateway REST API operations and connected services.	Low
API Gateway should be associated with an AWS WAF web ACL	This control checks whether an API Gateway stage uses an AWS WAF web access control list (ACL). This control fails if an AWS WAF web ACL is not attached to a REST API Gateway stage. AWS WAF is a web application firewall that helps protect web applications and APIs from attacks. It enables you to configure an ACL, which is a set of rules that allow, block, or count web requests based on customizable web security rules and conditions that you define. Ensure that your API Gateway stage is associated with an AWS WAF web ACL to help protect it from malicious attacks.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Application and Classic Load Balancers logging should be enabled	This control checks whether the Application Load Balancer and the Classic Load Balancer have logging enabled. The control fails if access_logs.s3.enabled is false. Elastic Load Balancing provides access logs that capture detailed information about requests sent to your load balancer. Each log contains information such as the time the request was received, the client's IP address, latencies, request paths, and server responses. You can use these access logs to analyze traffic patterns and to troubleshoot issues. To learn more, see Access logs for your Classic Load Balancer in User Guide for Classic Load Balancers.	Medium
Attached EBS volumes should be encrypted at-rest	This control checks whether the EBS volumes that are in an attached state are encrypted. To pass this check, EBS volumes must be in use and encrypted. If the EBS volume is not attached, then it is not subject to this check. For an added layer of security of your sensitive data in EBS volumes, you should enable EBS encryption at rest. Amazon EBS encryption offers a straightforward encryption solution for your EBS resources that doesn't require you to build, maintain, and secure your own key management infrastructure. It uses AWS KMS customer master keys (CMK) when creating encrypted volumes and snapshots. To learn more about Amazon EBS encryption, see Amazon EBS encryption in the Amazon EC2 User Guide for Linux Instances.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
AWS Database Migration Service replication instances should not be public	To protect your replicated instances from threats. A private replication instance should have a private IP address that you cannot access outside of the replication network. A replication instance should have a private IP address when the source and target databases are in the same network, and the network is connected to the replication instance's VPC using a VPN, AWS Direct Connect, or VPC peering. You should also ensure that access to your AWS DMS instance configuration is limited to only authorized users. To do this, restrict users' IAM permissions to modify AWS DMS settings and resources.	High
Classic Load Balancer listeners should be configured with HTTPS or TLS termination	This control checks whether your Classic Load Balancer listeners are configured with HTTPS or TLS protocol for front-end (client to load balancer) connections. The control is applicable if a Classic Load Balancer has listeners. If your Classic Load Balancer does not have a listener configured, then the control does not report any findings. The control passes if the Classic Load Balancer listeners are configured with TLS or HTTPS for front-end connections. The control fails if the listener is not configured with TLS or HTTPS for front-end connections. Before you start to use a load balancer, you must add one or more listeners. A listener is a process that uses the configured protocol and port to check for connection requests. Listeners can support both HTTP and HTTPS/TLS protocols. You should always use an HTTPS or TLS listener, so that the load balancer does the work of encryption and decryption in transit.	Medium
Classic Load Balancers should have connection draining enabled	This control checks whether Classic Load Balancers have connection draining enabled. Enabling connection draining on Classic Load Balancers ensures that the load balancer stops sending requests to instances that are de-registering or unhealthy. It keeps the existing connections open. This is particularly useful for instances in Auto Scaling groups, to ensure that connections aren't severed abruptly.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
CloudFront distributions should have AWS WAF enabled	This control checks whether CloudFront distributions are associated with either AWS WAF or AWS WAFv2 web ACLs. The control fails if the distribution is not associated with a web ACL. AWS WAF is a web application firewall that helps protect web applications and APIs from attacks. It allows you to configure a set of rules, called a web access control list (web ACL), that allow, block, or count web requests based on customizable web security rules and conditions that you define. Ensure your CloudFront distribution is associated with an AWS WAF web ACL to help protect it from malicious attacks.	Medium
CloudFront distributions should have logging enabled	This control checks whether server access logging is enabled on CloudFront distributions. The control fails if access logging is not enabled for a distribution. CloudFront access logs provide detailed information about every user request that CloudFront receives. Each log contains information such as the date and time the request was received, the IP address of the viewer that made the request, the source of the request, and the port number of the request from the viewer. These logs are useful for applications such as security and access audits and forensics investigation. For additional guidance on how to analyze access logs, see Querying Amazon CloudFront logs in the Amazon Athena User Guide.	Medium
CloudFront distributions should require encryption in transit	This control checks whether an Amazon CloudFront distribution requires viewers to use HTTPS directly or whether it uses redirection. The control fails if ViewerProtocolPolicy is set to allow-all for defaultCacheBehavior or for cacheBehaviors. HTTPS (TLS) can be used to help prevent potential attackers from using person-in-the-middle or similar attacks to eavesdrop on or manipulate network traffic. Only encrypted connections over HTTPS (TLS) should be allowed. Encrypting data in transit can affect performance. You should test your application with this feature to understand the performance profile and the impact of TLS.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
CloudTrail logs should be encrypted at rest using KMS CMKs	We recommended to configure CloudTrail use SSE-KMS. Configuring CloudTrail to use SSE-KMS provides additional confidentiality controls on log data as a given user must have S3 read permission on the corresponding log bucket and must be granted decrypt permission by the CMK policy.	Medium
Connections to Amazon Redshift clusters should be encrypted in transit	This control checks whether connections to Amazon Redshift clusters are required to use encryption in transit. The check fails if the Amazon Redshift cluster parameter require_SSL is not set to '1'. TLS can be used to help prevent potential attackers from using person-in-the-middle or similar attacks to eavesdrop on or manipulate network traffic. Only encrypted connections over TLS should be allowed. Encrypting data in transit can affect performance. You should test your application with this feature to understand the performance profile and the impact of TLS.	Medium
Connections to Elasticsearch domains should be encrypted using TLS 1.2	This control checks whether connections to Elasticsearch domains are required to use TLS 1.2. The check fails if the Elasticsearch domain TLSSecurityPolicy is not Policy-Min-TLS-1-2-2019-07. HTTPS (TLS) can be used to help prevent potential attackers from using person-in-the-middle or similar attacks to eavesdrop on or manipulate network traffic. Only encrypted connections over HTTPS (TLS) should be allowed. Encrypting data in transit can affect performance. You should test your application with this feature to understand the performance profile and the impact of TLS. TLS 1.2 provides several security enhancements over previous versions of TLS.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
DynamoDB tables should have point-in-time recovery enabled	This control checks whether point-in-time recovery (PITR) is enabled for an Amazon DynamoDB table. Backups help you to recover more quickly from a security incident. They also strengthen the resilience of your systems. DynamoDB point-in-time recovery automates backups for DynamoDB tables. It reduces the time to recover from accidental delete or write operations. DynamoDB tables that have PITR enabled can be restored to any point in time in the last 35 days.	Medium
EBS default encryption should be enabled	This control checks whether account-level encryption is enabled by default for Amazon Elastic Block Store(Amazon EBS). The control fails if the account level encryption is not enabled. When encryption is enabled for your account, Amazon EBS volumes and snapshot copies are encrypted at rest. This adds an additional layer of protection for your data. For more information, see Encryption by default in the Amazon EC2 User Guide for Linux Instances. Note that following instance types do not support encryption: R1, C1, and M1.	Medium
Elastic Beanstalk environments should have enhanced health reporting enabled	This control checks whether enhanced health reporting is enabled for your AWS Elastic Beanstalk environments. Elastic Beanstalk enhanced health reporting enables a more rapid response to changes in the health of the underlying infrastructure. These changes could result in a lack of availability of the application. Elastic Beanstalk enhanced health reporting provides a status descriptor to gauge the severity of the identified issues and identify possible causes to investigate. The Elastic Beanstalk health agent, included in supported Amazon Machine Images (AMIs), evaluates logs and metrics of environment EC2 instances.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Elastic Beanstalk managed platform updates should be enabled	his control checks whether managed platform updates are enabled for the Elastic Beanstalk environment. Enabling managed platform updates ensures that the latest available platform fixes, updates, and features for the environment are installed. Keeping up to date with patch installation is an important step in securing systems.	High
Elasticsearch domain error logging to CloudWatch Logs should be enabled	This control checks whether Elasticsearch domains are configured to send error logs to CloudWatch Logs. You should enable error logs for Elasticsearch domains and send those logs to CloudWatch Logs for retention and response. Domain error logs can assist with security and access audits, and can help to diagnose availability issues.	Medium
Elasticsearch domains should be configured with at least three dedicated master nodes	This control checks whether Elasticsearch domains are configured with at least three dedicated master nodes. This control fails if the domain does not use dedicated master nodes. This control passes if Elasticsearch domains have five dedicated master nodes. However, using more than three master nodes might be unnecessary to mitigate the availability risk, and will result in additional cost. An Elasticsearch domain requires at least three dedicated master nodes for high availability and fault-tolerance. Dedicated master node resources can be strained during data node blue/green deployments because there are additional nodes to manage. Deploying an Elasticsearch domain with at least three dedicated master nodes ensures sufficient master node resource capacity and cluster operations if a node fails.	Medium
Elasticsearch domains should have at least three data nodes	This control checks whether Elasticsearch domains are configured with at least three data nodes and zoneAwarenessEnabled is true. An Elasticsearch domain requires at least three data nodes for high availability and fault-tolerance. Deploying an Elasticsearch domain with at least three data nodes ensures cluster operations if a node fails.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Elasticsearch domains should have audit logging enabled	This control checks whether Elasticsearch domains have audit logging enabled. This control fails if an Elasticsearch domain does not have audit logging enabled. Audit logs are highly customizable. They allow you to track user activity on your Elasticsearch clusters, including authentication successes and failures, requests to OpenSearch, index changes, and incoming search queries.	Medium
Enhanced monitoring should be configured for RDS DB instances and clusters	This control checks whether enhanced monitoring is enabled for your RDS DB instances. In Amazon RDS, Enhanced Monitoring enables a more rapid response to performance changes in underlying infrastructure. These performance changes could result in a lack of availability of the data. Enhanced Monitoring provides real-time metrics of the operating system that your RDS DB instance runs on. An agent is installed on the instance. The agent can obtain metrics more accurately than is possible from the hypervisor layer. Enhanced Monitoring metrics are useful when you want to see how different processes or threads on a DB instance use the CPU. For more information, see Enhanced Monitoring in the *Amazon RDS User Guide*.	Low
Ensure rotation for customer created CMKs is enabled	AWS Key Management Service (KMS) allows customers to rotate the backing key which is key material stored within the KMS which is tied to the key ID of the Customer Created customer master key (CMK). It is the backing key that is used to perform cryptographic operations such as encryption and decryption. Automated key rotation currently retains all prior backing keys so that decryption of encrypted data can take place transparently. It is recommended that CMK key rotation be enabled. Rotating encryption keys helps reduce the potential impact of a compromised key as data encrypted with a new key cannot be accessed with a previous key that may have been exposed.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure S3 bucket access logging is enabled on the CloudTrail S3 bucket	S3 Bucket Access Logging generates a log that contains access records Ensure S3 bucket access logging is enabled on the CloudTrail S3 bucket for each request made to your S3 bucket. An access log record contains details about the request, such as the request type, the resources specified in the request worked, and the time and date the request was processed. It is recommended that bucket access logging be enabled on the CloudTrail S3 bucket. By enabling S3 bucket logging on target S3 buckets, it is possible to capture all events which may affect objects within an target buckets. Configuring logs to be placed in a separate bucket allows access to log information which can be useful in security and incident response workflows.	Low
Ensure the S3 bucket used to store CloudTrail logs is not publicly accessible	CloudTrail logs a record of every API call made in your AWS account. These log files are stored in an S3 bucket. It is recommended that the bucket policy, or access control list (ACL), applied to the S3 bucket that CloudTrail logs to prevents public access to the CloudTrail logs. Allowing public access to CloudTrail log content may aid an adversary in identifying weaknesses in the affected account's use or configuration.	High
Imported ACM certificates should be renewed after a specified time period	This control checks whether ACM certificates in your account are marked for expiration within 30 days. It checks both imported certificates and certificates provided by AWS Certificate Manager. ACM can automatically renew certificates that use DNS validation. For certificates that use email validation, you must respond to a domain validation email. ACM also does not automatically renew certificates that you import. You must renew imported certificates manually. For more information about managed renewal for ACM certificates, see Managed renewal for ACM certificates in the AWS Certificate Manager User Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
RDS automatic minor version upgrades should be enabled	This control checks whether automatic minor version upgrades are enabled for the RDS database instance. Enabling automatic minor version upgrades ensures that the latest minor version updates to the relational database management system (RDBMS) are installed. These upgrades might include security patches and bug fixes. Keeping up to date with patch installation is an important step in securing systems.	High
RDS cluster snapshots and database snapshots should be encrypted at rest	This control checks whether RDS DB snapshots are encrypted. This control is intended for RDS DB instances. However, it can also generate findings for snapshots of Aurora DB instances, Neptune DB instances, and Amazon DocumentDB clusters. If these findings are not useful, then you can suppress them. Encrypting data at rest reduces the risk that an unauthenticated user gets access to data that is stored on disk. Data in RDS snapshots should be encrypted at rest for an added layer of security.	Medium
RDS clusters should have deletion protection enabled	This control checks whether RDS clusters have deletion protection enabled. This control is intended for RDS DB instances. However, it can also generate findings for Aurora DB instances, Neptune DB instances, and Amazon DocumentDB clusters. If these findings are not useful, then you can suppress them. Enabling cluster deletion protection is an additional layer of protection against accidental database deletion or deletion by an unauthorized entity. When deletion protection is enabled, an RDS cluster cannot be deleted. Before a deletion request can succeed, deletion protection must be disabled.	Low
RDS DB clusters should be configured for multiple Availability Zones	RDS DB clusters should be configured for multiple the data that is stored. Deployment to multiple Availability Zones allows for automate Availability Zones to ensure availability of ed failover in the event of an Availability Zone availability issue and during regular RDS maintenance events.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
RDS DB clusters should be configured to copy tags to snapshots	Identification and inventory of your IT assets is a crucial aspect of governance and security. You need to have visibility of all your RDS DB clusters so that you can assess their security posture and act on potential areas of weakness. Snapshots should be tagged in the same way as their parent RDS database clusters. Enabling this setting ensures that snapshots inherit the tags of their parent database clusters.	Low
RDS DB instances should be configured to copy tags to snapshots	This control checks whether RDS DB instances are configured to copy all tags to snapshots when the snapshots are created. Identification and inventory of your IT assets is a crucial aspect of governance and security. You need to have visibility of all your RDS DB instances so that you can assess their security posture and take action on potential areas of weakness. Snapshots should be tagged in the same way as their parent RDS database instances. Enabling this setting ensures that snapshots inherit the tags of their parent database instances.	Low
RDS DB instances should be configured with multiple Availability Zones	This control checks whether high availability is enabled for your RDS DB instances. RDS DB instances should be configured for multiple Availability Zones (AZs). This ensures the availability of the data stored. Multi-AZ deployments allow for automated failover if there is an issue with Availability Zone availability and during regular RDS maintenance.	Medium
RDS DB instances should have deletion protection enabled	This control checks whether your RDS DB instances that use one of the listed database engines have deletion protection enabled. Enabling instance deletion protection is an additional layer of protection against accidental database deletion or deletion by an unauthorized entity. While deletion protection is enabled, an RDS DB instance cannot be deleted. Before a deletion request can succeed, deletion protection must be disabled.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
RDS DB instances should have encryption at rest enabled	This control checks whether storage encryption is enabled for your Amazon RDS DB instances. This control is intended for RDS DB instances. However, it can also generate findings for Aurora DB instances, Neptune DB instances, and Amazon DocumentDB clusters. If these findings are not useful, then you can suppress them. For an added layer of security for your sensitive data in RDS DB instances, you should configure your RDS DB instances to be encrypted at rest. To encrypt your RDS DB instances and snapshots at rest, enable the encryption option for your RDS DB instances. Data that is encrypted at rest includes the underlying storage for DB instances, its automated backups, read replicas, and snapshots. RDS encrypted DB instances use the open standard AES-256 encryption algorithm to encrypt your data on the server that hosts your RDS DB instances. After your data is encrypted, Amazon RDS handles authentication of access and decryption of your data transparently with a minimal impact on performance. You do not need to modify your database client applications to use encryption. Amazon RDS encryption is currently available for all database engines and storage types. Amazon RDS encryption is available for most DB instance classes. To learn about DB instance classes that do not support Amazon RDS encryption, see Encrypting Amazon RDS resources in the *Amazon RDS User Guide*.	Medium
RDS DB Instances should prohibit public access	We recommend that you also ensure that access to your RDS instance's configuration is limited to authorized users only, by restricting users' IAM permissions to modify RDS instances' settings and resources.	High
RDS snapshots should prohibit public access	We recommend only allowing authorized principals to access the snapshot and change Amazon RDS configuration.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Remove unused Secrets Manager secrets	This control checks whether your secrets have been accessed within a specified number of days. The default value is 90 days. If a secret was not accessed within the defined number of days, this control fails. Deleting unused secrets is as important as rotating secrets. Unused secrets can be abused by their former users, who no longer need access to these secrets. Also, as more users get access to a secret, someone might have mishandled and leaked it to an unauthorized entity, which increases the risk of abuse. Deleting unused secrets helps revoke secret access from users who no longer need it. It also helps to reduce the cost of using Secrets Manager. Therefore, it is essential to routinely delete unused secrets.	Medium
S3 buckets should have cross-region replication enabled	Enabling S3 cross-region replication ensures that multiple versions of the data are available in different distinct Regions. This allows you to protect your S3 bucket against DDoS attacks and data corruption events.	Low
S3 buckets should have server-side encryption enabled	Enable server-side encryption to protect data in your S3 buckets. Encrypting the data can prevent access to sensitive data in the event of a data breach.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Secrets Manager secrets configured with automatic rotation should rotate successfully	This control checks whether an AWS Secrets Manager secret rotated successfully based on the rotation schedule. The control fails if **RotationOccurringAsScheduled** is **false**. The control does not evaluate secrets that do not have rotation configured. Secrets Manager helps you improve the security posture of your organization. Secrets include database credentials, passwords, and third-party API keys. You can use Secrets Manager to store secrets centrally, encrypt secrets automatically, control access to secrets, and rotate secrets safely and automatically. Secrets Manager can rotate secrets. You can use rotation to replace long-term secrets with short-term ones. Rotating your secrets limits how long an unauthorized user can use a compromised secret. For this reason, you should rotate your secrets frequently. In addition to configuring secrets to rotate automatically, you should ensure that those secrets rotate successfully based on the rotation schedule. To learn more about rotation, see Rotating your AWS Secrets Manager secrets in the AWS Secrets Manager User Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Secrets Manager secrets should be rotated within a specified number of days	This control checks whether your secrets have been rotated at least once within 90 days. Rotating secrets can help you to reduce the risk of an unauthorized use of your secrets in your AWS account. Examples include database credentials, passwords, third-party API keys, and even arbitrary text. If you do not change your secrets for a long period of time, the secrets are more likely to be compromised. As more users get access to a secret, it can become more likely that someone mishandled and leaked it to an unauthorized entity. Secrets can be leaked through logs and cache data. They can be shared for debugging purposes and not changed or revoked once the debugging completes. For all these reasons, secrets should be rotated frequently. You can configure your secrets for automatic rotation in AWS Secrets Manager. With automatic rotation, you can replace long-term secrets with short-term ones, significantly reducing the risk of compromise. Security Hub recommends that you enable rotation for your Secrets Manager secrets. To learn more about rotation, see Rotating your AWS Secrets Manager secrets in the AWS Secrets Manager User Guide.	Medium
SNS topics should be encrypted at rest using AWS KMS	This control checks whether an SNS topic is encrypted at rest using AWS KMS. Encrypting data at rest reduces the risk of data stored on disk being accessed by a user not authenticated to AWS. It also adds another set of access controls to limit the ability of unauthorized users to access the data. For example, API permissions are required to decrypt the data before it can be read. SNS topics should be encrypted at-rest for an added layer of security. For more information, see Encryption at rest in the Amazon Simple Notification Service Developer Guide.	Medium
VPC flow logging should be enabled in all VPCs	VPC Flow Logs provide visibility into network traffic that passes through the VPC and can be used to detect anomalous traffic or insight during security events.	Medium

AWS IdentityAndAccess recommendations

There are **46** AWS recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon Elasticsearch Service domains should be in a VPC	VPC cannot contain domains with a public endpoint. Note: this does not evaluate the VPC subnet routing configuration to determine public reachability.	High
Amazon S3 permissions granted to other AWS accounts in bucket policies should be restricted	Implementing least privilege access is fundamental to reducing security risk and the impact of errors or malicious intent. If an S3 bucket policy allows access from external accounts, it could result in data exfiltration by an insider threat or an attacker. The 'blacklistedactionpatterns' parameter allows for successful evaluation of the rule for S3 buckets. The parameter grants access to external accounts for action patterns that are not included in the 'blacklistedactionpatterns' list.	High
Avoid the use of the "root" account	The "root" account has unrestricted access to all resources in the AWS account. It is highly recommend that the use of this account be avoided. The "root" account is the most privileged AWS account. Minimizing the use of this account and adopting the principle of least privilege for access management will reduce the risk of accidental changes and unintended disclosure of highly privileged credentials.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
AWS KMS keys should not be unintentionally deleted	This control checks whether KMS keys are scheduled for deletion. The control fails if a KMS key is scheduled for deletion. KMS keys cannot be recovered once deleted. Data encrypted under a KMS key is also permanently unrecoverable if the KMS key is deleted. If meaningful data has been encrypted under a KMS key scheduled for deletion, consider decrypting the data or re-encrypting the data under a new KMS key unless you are intentionally performing a cryptographic erasure. When a KMS key is scheduled for deletion, a mandatory waiting period is enforced to allow time to reverse the deletion, if it was scheduled in error. The default waiting period is 30 days, but it can be reduced to as short as 7 days when the KMS key is scheduled for deletion. During the waiting period, the scheduled deletion can be canceled and the KMS key will not be deleted. For additional information regarding deleting KMS keys, see Deleting KMS keys in the AWS Key Management Service Developer Guide.	High
AWS WAF Classic global web ACL logging should be enabled	This control checks whether logging is enabled for an AWS WAF global Web ACL. This control fails if logging is not enabled for the web ACL. Logging is an important part of maintaining the reliability, availability, and performance of AWS WAF globally. It is a business and compliance requirement in many organizations, and allows you to troubleshoot application behavior. It also provides detailed information about the traffic that is analyzed by the web ACL that is attached to AWS WAF.	Medium
CloudFront distributions should have a default root object configured	This control checks whether an Amazon CloudFront distribution is configured to return a specific object that is the default root object. The control fails if the CloudFront distribution does not have a default root object configured. A user might sometimes request the distributions root URL instead of an object in the distribution. When this happens, specifying a default root object can help you to avoid exposing the contents of your web distribution.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
CloudFront distributions should have origin access identity enabled	This control checks whether an Amazon CloudFront distribution with Amazon S3 Origin type has Origin Access Identity (OAI) configured. The control fails if OAI is not configured. CloudFront OAI prevents users from accessing S3 bucket content directly. When users access an S3 bucket directly, they effectively bypass the CloudFront distribution and any permissions that are applied to the underlying S3 bucket content.	Medium
CloudTrail log file validation should be enabled	To ensure additional integrity checking of CloudTrail logs, we recommend enabling file validation on all CloudTrails.	Low
CloudTrail should be enabled	AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. Not all services enable logging by default for all APIs and events. You should implement any additional audit trails other than CloudTrail and review the documentation for each service in CloudTrail Supported Services and Integrations.	High
CloudTrail trails should be integrated with CloudWatch Logs	In addition to capturing CloudTrail logs within a specified S3 bucket for long term analysis, real-time analysis can be performed by configuring CloudTrail to send logs to CloudWatch Logs. For a trail that is enabled in all regions in an account, CloudTrail sends log files from all those regions to a CloudWatch Logs log group. We recommended that CloudTrail logs will be sent to CloudWatch Logs to ensure AWS account activity is being captured, monitored, and appropriately alarmed on. Sending CloudTrail logs to CloudWatch Logs facilitates real-time and historic activity logging based on user, API, resource, and IP address, and provides opportunity to establish alarms and notifications for anomalous or sensitivity account activity.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Database logging should be enabled	This control checks whether the following logs of Amazon RDS are enabled and sent to CloudWatch Logs: - Oracle: (Alert, Audit, Trace, Listener) - PostgreSQL: (Postgresql, Upgrade) - MySQL: (Audit, Error, General, SlowQuery) - MariaDB: (Audit, Error, General, SlowQuery) - SQL Server: (Error, Agent) - Aurora: (Audit, Error, General, SlowQuery) - Aurora-MySQL: (Audit, Error, General, SlowQuery) - Aurora-PostgreSQL: (Postgresql, Upgrade). RDS databases should have relevant logs enabled. Database logging provides detailed records of requests made to RDS. Database logs can assist with security and access audits and can help to diagnose availability issues.	Medium
Disable direct internet access for Amazon SageMaker notebook instances	Direct internet access should be disabled for an SageMaker notebook instance. This checks whether the 'DirectInternetAccess' field is disabled for the notebook instance. Your instance should be configured with a VPC and the default setting should be Disable - Access the internet through a VPC. In order to enable internet access to train or host models from a notebook, make sure that your VPC has a NAT gateway and your security group allows outbound connections. Ensure access to your SageMaker configuration is limited to only authorized users, and restrict users' IAM permissions to modify SageMaker settings and resources.	High
Do not setup access keys during initial user setup for all IAM users that have a console password	AWS console defaults the checkbox for creating access keys to enabled. This results in many access keys being generated unnecessarily. In addition to unnecessary credentials, it also generates unnecessary management work in auditing and rotating these keys. Requiring that additional steps be taken by the user after their profile has been created will give a stronger indication of intent that access keys are [a] necessary for their work and [b] once the access key is established on an account that the keys may be in use somewhere in the organization	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure a support role has been created to manage incidents with AWS Support	AWS provides a support center that can be used for incident notification and response, as well as technical support and customer services. Create an IAM Role to allow authorized users to manage incidents with AWS Support. By implementing least privilege for access control, an IAM Role will require an appropriate IAM Policy to allow Support Center Access in order to manage Incidents with AWS Support.	Low
Ensure access keys are rotated every 90 days or less	Access keys consist of an access key ID and secret access key, which are used to sign programmatic requests that you make to AWS. AWS users need their own access keys to make programmatic calls to AWS from the AWS Command Line Interface (AWS CLI), Tools for Windows PowerShell, the AWS SDKs, or direct HTTP calls using the APIs for individual AWS services. It is recommended that all access keys be regularly rotated. Rotating access keys will reduce the window of opportunity for an access key that is associated with a compromised or terminated account to be used. Access keys should be rotated to ensure that data cannot be accessed with an old key which might have been lost, cracked, or stolen.	Medium
Ensure AWS Config is enabled in all regions	AWS Config is a web service that performs configuration management of supported AWS resources within your account and delivers log files to you. The recorded information includes the configuration item (AWS resource), relationships between configuration items (AWS resources), any configuration changes between resources. It is recommended to enable AWS Config be enabled in all regions. The AWS configuration item history captured by AWS Config enables security analysis, resource change tracking, and compliance auditing.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure CloudTrail is enabled in all regions	AWS CloudTrail is a web service that records AWS API calls for your account and delivers log files to you. The recorded information includes the identity of the API caller, the time of the API call, the source IP address of the API caller, the request parameters, and the response elements returned by the AWS service. CloudTrail provides a history of AWS API calls for an account, including API calls made via the Management Console, SDKs, command line tools, and higher-level AWS services (such as CloudFormation). The AWS API call history produced by CloudTrail enables security analysis, resource change tracking, and compliance auditing. Additionally, * ensuring that a multi-regions trail exists will ensure that unexpected activity occurring in otherwise unused regions is detected * ensuring that a multi-regions trail exists will ensure that "Global Service Logging" is enabled for a trail by default to capture recording of events generated on AWS global services * for a multi-regions trail, ensuring that management events configured for all type of Read/Writes ensures recording of management operations that are performed on all resources in an AWS account.	High
Ensure credentials unused for 90 days or greater are disabled	AWS IAM users can access AWS resources using different types of credentials, such as passwords or access keys. It is recommended that all credentials that have been unused in 90 or greater days be removed or deactivated. Disabling or removing unnecessary credentials will reduce the window of opportunity for credentials associated with a compromised or abandoned account to be used.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure IAM password policy expires passwords within 90 days or less	IAM password policies can require passwords to be rotated or expired after a given number of days. It is recommended that the password policy expire passwords after 90 days or less. Reducing the password lifetime increases account resiliency against brute force login attempts. Additionally, requiring regular password changes help in the following scenarios: * Passwords can be stolen or compromised sometimes without your knowledge. This can happen via a system compromise, software vulnerability, or internal threat. * Certain corporate and government web filters or proxy servers have the ability to intercept and record traffic even if it's encrypted. * Many people use the same password for many systems such as work, email, and personal. * Compromised end user workstations might have a keystroke logger.	Low
Ensure IAM password policy prevents password reuse	IAM password policies can prevent the reuse of a given password by the same user. It is recommended that the password policy prevent the reuse of passwords. Preventing password reuse increases account resiliency against brute force login attempts.	Low
Ensure IAM password policy requires at least one lowercase letter	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are comprised of different character sets. It is recommended that the password policy require at least one lowercase letter. Setting a password complexity policy increases account resiliency against brute force login attempts	Medium
Ensure IAM password policy requires at least one number	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are comprised of different character sets. It is recommended that the password policy require at least one number. Setting a password complexity policy increases account resiliency against brute force login attempts.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure IAM password policy requires at least one symbol	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are comprised of different character sets. It is recommended that the password policy require at least one symbol. Setting a password complexity policy increases account resiliency against brute force login attempts.	Medium
Ensure IAM password policy requires at least one uppercase letter	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are comprised of different character sets. It is recommended that the password policy require at least one uppercase letter. Setting a password complexity policy increases account resiliency against brute force login attempts.	Medium
Ensure IAM password policy requires minimum length of 14 or greater	Password policies are, in part, used to enforce password complexity requirements. IAM password policies can be used to ensure password are at least a given length. It is recommended that the password policy require a minimum password length '14'. Setting a password complexity policy increases account resiliency against brute force login attempts.	Medium
Ensure multi-factor authentication (MFA) is enabled for all IAM users that have a console password	Multi-Factor Authentication (MFA) adds an extra layer of protection on top of a user name and password. With MFA enabled, when a user signs in to an AWS website, they will be prompted for their user name and password as well as for an authentication code from their AWS MFA device. It is recommended that MFA be enabled for all accounts that have a console password. Enabling MFA provides increased security for console access as it requires the authenticating principal to possess a device that emits a time-sensitive key and have knowledge of a credential.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
GuardDuty should be enabled	To provide additional protection against intrusions, GuardDuty should be enabled on your AWS account and region. Note: GuardDuty might not be a complete solution for every environment	Medium
Hardware MFA should be enabled for the "root" account	The root account is the most privileged user in an account. MFA adds an extra layer of protection on top of a user name and password. With MFA enabled, when a user signs in to an AWS website, they're prompted for their user name and password and for an authentication code from their AWS MFA device. For Level 2, it is recommended that you protect the root account with a hardware MFA. A hardware MFA has a smaller attack surface than a virtual MFA. For example, a hardware MFA doesn't suffer the attack surface introduced by the mobile smartphone that a virtual MFA resides on. Using hardware MFA for many, many accounts might create a logistical device management issue. If this occurs, consider implementing this Level 2 recommendation selectively to the highest security accounts. You can then apply the Level 1 recommendation to the remaining accounts.	Low
IAM authentication should be configured for RDS clusters	This control checks whether an RDS DB cluster has IAM database authentication enabled. IAM database authentication allows for password-free authentication to database instances. The authentication uses an authentication token. Network traffic to and from the database is encrypted using SSL. For more information, see IAM database authentication in the Amazon Aurora User Guide.	Medium
IAM authentication should be configured for RDS instances	This control checks whether an RDS DB instance has IAM database authentication enabled. IAM database authentication allows authentication to database instances with an authentication token instead of a password. Network traffic to and from the database is encrypted using SSL. For more information, see IAM database authentication in the Amazon Aurora User Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
IAM customer managed policies should not allow decryption actions on all KMS keys	Checks whether the default version of IAM customer managed policies allow principals to use the AWS KMS decryption actions on all resources. This control uses Zelkova, an automated reasoning engine, to validate and warn you about policies that may grant broad access to your secrets across AWS accounts.This control fails if the "kms:Decrypt" or "kms:ReEncryptFrom" actions are allowed on all KMS keys. The control evaluates both attached and unattached customer managed policies. It does not check inline policies or AWS managed policies. With AWS KMS, you control who can use your KMS keys and gain access to your encrypted data. IAM policies define which actions an identity (user, group, or role) can perform on which resources. Following security best practices, AWS recommends that you allow least privilege. In other words, you should grant to identities only the "kms:Decrypt" or "kms:ReEncryptFrom" permissions and only for the keys that are required to perform a task. Otherwise, the user might use keys that are not appropriate for your data. Instead of granting permissions for all keys, determine the minimum set of keys that users need to access encrypted data. Then design policies that allow users to use only those keys. For example, do not allow "kms:Decrypt" permission on all KMS keys. Instead, allow "kms:Decrypt" only on keys in a particular Region for your account. By adopting the principle of least privilege, you can reduce the risk of unintended disclosure of your data.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY

IAM customer managed policies that you create should not allow wildcard actions for services	This control checks whether the IAM identity-based policies that you create have Allow statements that use the * wildcard to grant permissions for all actions on any service. The control fails if any policy statement includes 'Effect': 'Allow' with 'Action': 'Service:*'.
	For example, the following statement in a policy results in a failed finding.
	`'Statement': [ { 'Sid': 'EC2-Wildcard', 'Effect': 'Allow', 'Action': 'ec2:*', 'Resource': '*' }`
	The control also fails if you use 'Effect': 'Allow' with 'NotAction': 'service:'*. In that case, the NotAction element provides access to all of the actions in an AWS service, except for the actions specified in NotAction.
	This control only applies to customer managed IAM policies. It does not apply to IAM policies that are managed by AWS.
	When you assign permissions to AWS services, it is important to scope the allowed IAM actions in your IAM policies. You should restrict IAM actions to only those actions that are needed. This helps you to provision least privilege permissions. Overly permissive policies might lead to privilege escalation if the policies are attached to an IAM principal that might not require the permission.
	In some cases, you might want to allow IAM actions that have a similar prefix, such as DescribeFlowLogs and DescribeAvailabilityZones. In these authorized cases, you can add a suffixed wildcard to the common prefix. For example, ec2:Describe*.
	
	This control passes if you use a prefixed IAM action with a suffixed wildcard. For example, the following statement in a policy results in a passed finding.
	`'Statement': [ { 'Sid': 'EC2-Wildcard', 'Effect': 'Allow', 'Action': 'ec2:Describe*', 'Resource': '*' }`
	When you group related IAM actions in this way, you can also avoid exceeding the IAM policy size limits.
		Low

RECOMMENDATION	DESCRIPTION	SEVERITY
IAM policies should be attached only to groups or roles	By default, IAM users, groups, and roles have no access to AWS resources. IAM policies are the means by which privileges are granted to users, groups, or roles. It is recommended that IAM policies be applied directly to groups and roles but not users. Assigning privileges at the group or role level reduces the complexity of access management as the number of users grow. Reducing access management complexity may in-turn reduce opportunity for a principal to inadvertently receive or retain excessive privileges.	Low
IAM policies that allow full "*:*" administrative privileges should not be created	IAM policies are the means by which privileges are granted to users, groups, or roles. It is recommended and considered a standard security advice to grant least privilege-that is, granting only the permissions required to perform a task. Determine what users need to do and then craft policies for them that let the users perform only those tasks, instead of allowing full administrative privileges. It's more secure to start with a minimum set of permissions and grant additional permissions as necessary, rather than starting with permissions that are too lenient and then trying to tighten them later. Providing full administrative privileges instead of restricting to the minimum set of permissions that the user is required to do exposes the resources to potentially unwanted actions. IAM policies that have a statement with "Effect": "Allow" with "Action": "*" *over "Resource": "*" should be removed.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
IAM principals should not have IAM inline policies that allow decryption actions on all KMS keys	Checks whether the inline policies that are embedded in your IAM identities (role, user, or group) allow the AWS KMS decryption actions on all KMS keys. This control uses Zelkova, an automated reasoning engine, to validate and warn you about policies that may grant broad access to your secrets across AWS accounts. This control fails if "kms:Decrypt" or "kms:ReEncryptFrom" actions are allowed on all KMS keys in an inline policy. With AWS KMS, you control who can use your KMS keys and gain access to your encrypted data. IAM policies define which actions an identity (user, group, or role) can perform on which resources. Following security best practices, AWS recommends that you allow least privilege. In other words, you should grant to identities only the permissions they need and only for keys that are required to perform a task. Otherwise, the user might use keys that are not appropriate for your data. Instead of granting permission for all keys, determine the minimum set of keys that users need to access encrypted data. Then design policies that allow the users to use only those keys. For example, do not allow "kms:Decrypt" permission on all KMS keys. Instead, allow them only on keys in a particular Region for your account. By adopting the principle of least privilege, you can reduce the risk of unintended disclosure of your data.	Medium
Lambda functions should restrict public access	Lambda function resource-based policy should restrict public access. This recommendation does not check access by internal principals. Ensure access to the function is restricted to authorized principals only by using least privilege resource-based policies.	High
MFA should be enabled for all IAM users	All IAM users should have multi-factor authentication (MFA) enabled.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
MFA should be enabled for the "root" account	The root account is the most privileged user in an account. MFA adds an extra layer of protection on top of a user name and password. With MFA enabled, when a user signs in to an AWS website, they're prompted for their user name and password and for an authentication code from their AWS MFA device. When you use virtual MFA for root accounts, it is recommended that the device used is not a personal device. Instead, use a dedicated mobile device (tablet or phone) that you manage to keep charged and secured independent of any individual personal devices. This lessens the risks of losing access to the MFA due to device loss, device trade-in, or if the individual owning the device is no longer employed at the company.	Low
Password policies for IAM users should have strong configurations	Checks whether the account password policy for IAM users uses the following minimum configurations. * RequireUppercaseCharacters- Require at least one uppercase character in password. (Default = true) * RequireLowercaseCharacters- Require at least one lowercase character in password. (Default = true) * RequireNumbers- Require at least one number in password. (Default = true) * MinimumPasswordLength- Password minimum length. (Default = 7 or longer) * PasswordReusePrevention- Number of passwords before allowing reuse. (Default = 4) * MaxPasswordAge- Number of days before password expiration. (Default = 90)	Medium
Root account access key shouldn't exist	The root account is the most privileged user in an AWS account. AWS Access Keys provide programmatic access to a given AWS account. It is recommended that all access keys associated with the root account be removed. Removing access keys associated with the root account limits vectors by which the account can be compromised. Additionally, removing the root access keys encourages the creation and use of role based accounts that are least privileged.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
S3 Block Public Access setting should be enabled	Enabling Block Public Access setting for your S3 bucket can help prevent sensitive data leaks and protect your bucket from malicious actions.	Medium
S3 Block Public Access setting should be enabled at the bucket level	This control checks whether S3 buckets have bucket-level public access blocks applied. This control fails is if any of the following settings are set to false: * ignorePublicAcls * blockPublicPolicy * blockPublicAcls * restrictPublicBuckets Block Public Access at the S3 bucket level provides controls to ensure that objects never have public access. Public access is granted to buckets and objects through access control lists (ACLs), bucket policies, or both. Unless you intend to have your S3 buckets publicly accessible, you should configure the bucket level Amazon S3 Block Public Access feature.	High
S3 buckets public read access should be removed	Removing public read access to your S3 bucket can help protect your data and prevent a data breach.	High
S3 buckets public write access should be removed	Allowing public write access to your S3 bucket can leave you vulnerable to malicious actions such as storing data at your expense, encrypting your files for ransom, or using your bucket to operate malware.	High
Secrets Manager secrets should have automatic rotation enabled	This control checks whether a secret stored in AWS Secrets Manager is configured with automatic rotation. Secrets Manager helps you improve the security posture of your organization. Secrets include database credentials, passwords, and third-party API keys. You can use Secrets Manager to store secrets centrally, encrypt secrets automatically, control access to secrets, and rotate secrets safely and automatically. Secrets Manager can rotate secrets. You can use rotation to replace long-term secrets with short-term ones. Rotating your secrets limits how long an unauthorized user can use a compromised secret. For this reason, you should rotate your secrets frequently. To learn more about rotation, see Rotating your AWS Secrets Manager secrets in the AWS Secrets Manager User Guide.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Stopped EC2 instances should be removed after a specified time period	This control checks whether any EC2 instances have been stopped for more than the allowed number of days. An EC2 instance fails this check if it is stopped for longer than the maximum allowed time period, which by default is 30 days. A failed finding indicates that an EC2 instance has not run for a significant period of time. This creates a security risk because the EC2 instance is not being actively maintained (analyzed, patched, updated). If it is later launched, the lack of proper maintenance could result in unexpected issues in your AWS environment. To safely maintain an EC2 instance over time in a nonrunning state, start it periodically for maintenance and then stop it after maintenance. Ideally this is an automated process.	Medium

AWS Networking recommendations

There are 36 AWS recommendations in this category.

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon EC2 should be configured to use VPC endpoints	This control checks whether a service endpoint for Amazon EC2 is created for each VPC. The control fails if a VPC does not have a VPC endpoint created for the Amazon EC2 service. To improve the security posture of your VPC, you can configure Amazon EC2 to use an interface VPC endpoint. Interface endpoints are powered by AWS PrivateLink, a technology that enables you to access Amazon EC2 API operations privately. It restricts all network traffic between your VPC and Amazon EC2 to the Amazon network. Because endpoints are supported within the same Region only, you cannot create an endpoint between a VPC and a service in a different Region. This prevents unintended Amazon EC2 API calls to other Regions. To learn more about creating VPC endpoints for Amazon EC2, see Amazon EC2 and interface VPC endpoints in the Amazon EC2 User Guide for Linux Instances.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Amazon ECS services should not have public IP addresses assigned to them automatically	A public IP address is an IP address that is reachable from the internet. If you launch your Amazon ECS instances with a public IP address, then your Amazon ECS instances are reachable from the internet. Amazon ECS services should not be publicly accessible, as this may allow unintended access to your container application servers.	High
Amazon EMR cluster master nodes should not have public IP addresses	This control checks whether master nodes on Amazon EMR clusters have public IP addresses. The control fails if the master node has public IP addresses that are associated with any of its instances. Public IP addresses are designated in the PublicIp field of the NetworkInterfaces configuration for the instance. This control only checks Amazon EMR clusters that are in a RUNNING or WAITING state.	High
Amazon Redshift clusters should use enhanced VPC routing	This control checks whether an Amazon Redshift cluster has EnhancedVpcRouting enabled. Enhanced VPC routing forces all COPY and UNLOAD traffic between the cluster and data repositories to go through your VPC. You can then use VPC features such as security groups and network access control lists to secure network traffic. You can also use VPC Flow Logs to monitor network traffic.	High
Application Load Balancer should be configured to redirect all HTTP requests to HTTPS	To enforce encryption in transit, you should use redirect actions with Application Load Balancers to redirect client HTTP requests to an HTTPS request on port 443.	Medium
Application load balancers should be configured to drop HTTP headers	This control evaluates AWS Application Load Balancers (ALB) to ensure they are configured to drop invalid HTTP headers. The control fails if the value of routing.http.drop_invalid_header_fields. enabled is set to false. By default, ALBs are not configured to drop invalid HTTP header values. Removing these header values prevents HTTP desync attacks.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Configure Lambda functions to a VPC	This control checks whether a Lambda function is in a VPC. It does not evaluate the VPC subnet routing configuration to determine public reachability. Note that if Lambda@Edge is found in the account, then this control generates failed findings. To prevent these findings, you can disable this control.	Low
EC2 instances should not have a public IP address	This control checks whether EC2 instances have a public IP address. The control fails if the "publicIp" field is present in the EC2 instance configuration item. This control applies to IPv4 addresses only. A public IPv4 address is an IP address that is reachable from the internet. If you launch your instance with a public IP address, then your EC2 instance is reachable from the internet. A private IPv4 address is an IP address that is not reachable from the internet. You can use private IPv4 addresses for communication between EC2 instances in the same VPC or in your connected private network. IPv6 addresses are globally unique, and therefore are reachable from the internet. However, by default all subnets have the IPv6 addressing attribute set to false. For more information about IPv6, see IP addressing in your VPC in the Amazon VPC User Guide. If you have a legitimate use case to maintain EC2 instances with public IP addresses, then you can suppress the findings from this control. For more information about front-end architecture options, see the AWS Architecture Blog or the This Is My Architecture series.	High
EC2 instances should not use multiple ENIs	This control checks whether an EC2 instance uses multiple Elastic Network Interfaces (ENIs) or Elastic Fabric Adapters (EFAs).This control passes if a single network adapter is used. The control includes an optional parameter list to identify the allowed ENIs. Multiple ENIs can cause dual-homed instances, meaning instances that have multiple subnets. This can add network security complexity and introduce unintended network paths and access.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
EC2 instances should use IMDSv2	This control checks whether your EC2 instance metadata version is configured with Instance Metadata Service Version 2 (IMDSv2). The control passes if "HttpTokens" is set to "required" for IMDSv2. The control fails if "HttpTokens" is set to "optional". You use instance metadata to configure or manage the running instance. The IMDS provides access to temporary, frequently rotated credentials. These credentials remove the need to hard code or distribute sensitive credentials to instances manually or programmatically. The IMDS is attached locally to every EC2 instance. It runs on a special 'link local' IP address of 169.254.169.254. This IP address is only accessible by software that runs on the instance. Version 2 of the IMDS adds new protections for the following types of vulnerabilities. These vulnerabilities could be used to try to access the IMDS. * Open website application firewalls * Open reverse proxies * Server-side request forgery (SSRF) vulnerabilities * Open Layer 3 firewalls and network address translation (NAT) Security Hub recommends that you configure your EC2 instances with IMDSv2.	High
EC2 subnets should not automatically assign public IP addresses	This control checks whether the assignment of public IPs in Amazon Virtual Private Cloud (Amazon VPC) subnets have "MapPublicIpOnLaunch" set to "FALSE". The control passes if the flag is set to "FALSE". All subnets have an attribute that determines whether a network interface created in the subnet automatically receives a public IPv4 address. Instances that are launched into subnets that have this attribute enabled have a public IP address assigned to their primary network interface.	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure a log metric filter and alarm exist for AWS Config configuration changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for detecting changes to CloudTrail's configurations. Monitoring changes to AWS Config configuration will help ensure sustained visibility of configuration items within the AWS account.	Low
Ensure a log metric filter and alarm exist for AWS Management Console authentication failures	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for failed console authentication attempts. Monitoring failed console logins may decrease lead time to detect an attempt to brute force a credential, which may provide an indicator, such as source IP, that can be used in other event correlation.	Low
Ensure a log metric filter and alarm exist for changes to Network Access Control Lists (NACL)	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. NACLs are used as a stateless packet filter to control ingress and egress traffic for subnets within a VPC. It is recommended that a metric filter and alarm be established for changes made to NACLs. Monitoring changes to NACLs will help ensure that AWS resources and services are not unintentionally exposed.	Low
Ensure a log metric filter and alarm exist for changes to network gateways	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. Network gateways are required to send/receive traffic to a destination outside of a VPC. It is recommended that a metric filter and alarm be established for changes to network gateways. Monitoring changes to network gateways will help ensure that all ingress/egress traffic traverses the VPC border via a controlled path.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure a log metric filter and alarm exist for CloudTrail configuration changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for detecting changes to CloudTrail's configurations. Monitoring changes to CloudTrail's configuration will help ensure sustained visibility to activities performed in the AWS account.	Low
Ensure a log metric filter and alarm exist for disabling or scheduled deletion of customer created CMKs	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for customer created CMKs which have changed state to disabled or scheduled deletion. Data encrypted with disabled or deleted keys will no longer be accessible.	Low
Ensure a log metric filter and alarm exist for IAM policy changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established changes made to Identity and Access Management (IAM) policies. Monitoring changes to IAM policies will help ensure authentication and authorization controls remain intact.	Low
Ensure a log metric filter and alarm exist for Management Console sign-in without MFA	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for console logins that are not protected by multi-factor authentication (MFA). Monitoring for single-factor console logins will increase visibility into accounts that are not protected by MFA.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure a log metric filter and alarm exist for route table changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. Routing tables are used to route network traffic between subnets and to network gateways. It is recommended that a metric filter and alarm be established for changes to route tables. Monitoring changes to route tables will help ensure that all VPC traffic flows through an expected path.	Low
Ensure a log metric filter and alarm exist for S3 bucket policy changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for changes to S3 bucket policies. Monitoring changes to S3 bucket policies may reduce time to detect and correct permissive policies on sensitive S3 buckets.	Low
Ensure a log metric filter and alarm exist for security group changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. Security Groups are a stateful packet filter that controls ingress and egress traffic within a VPC. It is recommended that a metric filter and alarm be established changes to Security Groups. Monitoring changes to security group will help ensure that resources and services are not unintentionally exposed.	Low
Ensure a log metric filter and alarm exist for unauthorized API calls	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for unauthorized API calls. Monitoring unauthorized API calls will help reveal application errors and may reduce time to detect malicious activity.	Low

RECOMMENDATION	DESCRIPTION	SEVERITY
Ensure a log metric filter and alarm exist for usage of 'root' account	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is recommended that a metric filter and alarm be established for root login attempts. Monitoring for root account logins will provide visibility into the use of a fully privileged account and an opportunity to reduce the use of it.	Low
Ensure a log metric filter and alarm exist for VPC changes	Real-time monitoring of API calls can be achieved by directing CloudTrail Logs to CloudWatch Logs and establishing corresponding metric filters and alarms. It is possible to have more than 1 VPC within an account, in addition it is also possible to create a peer connection between 2 VPCs enabling network traffic to route between VPCs. It is recommended that a metric filter and alarm be established for changes made to VPCs. Monitoring changes to IAM policies will help ensure authentication and authorization controls remain intact.	Low
Ensure no security groups allow ingress from 0.0.0.0/0 to port 3389	Security groups provide stateful filtering of ingress/egress network traffic to AWS resources. It is recommended that no security group allows unrestricted ingress access to port 3389. Removing unfettered connectivity to remote console services, such as RDP, reduces a server's exposure to risk.	High
Management ports of EC2 instances should be protected with just-in-time network access control	Microsoft Defender for Cloud has identified some overly-permissive inbound rules for management ports in your network. Enable just-in-time access control to protect your Instances from internet-based brute-force attacks. Learn more.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
RDS databases and clusters should not use a database engine default port	This control checks whether the RDS cluster or instance uses a port other than the default port of the database engine. If you use a known port to deploy an RDS cluster or instance, an attacker can guess information about the cluster or instance. The attacker can use this information in conjunction with other information to connect to an RDS cluster or instance or gain additional information about your application. When you change the port, you must also update the existing connection strings that were used to connect to the old port. You should also check the security group of the DB instance to ensure that it includes an ingress rule that allows connectivity on the new port.	Low
RDS instances should be deployed in a VPC	VPCs provide a number of network controls to secure access to RDS resources. These controls include VPC Endpoints, network ACLs, and security groups. To take advantage of these controls, we recommend that you move EC2-Classic RDS instances to EC2-VPC.	Low
S3 buckets should require requests to use Secure Socket Layer	We recommend to require requests to use Secure Socket Layer (SSL) on all Amazon S3 bucket. S3 buckets should have policies that require all requests ('Action: S3:*') to only accept transmission of data over HTTPS in the S3 resource policy, indicated by the condition key 'aws:SecureTransport'.	Medium
Security groups should not allow ingress from 0.0.0.0/0 to port 22	To reduce the server's exposure, it is recommended not to allow unrestricted ingress access to port '22'.	High

RECOMMENDATION	DESCRIPTION	SEVERITY
Security groups should not allow unrestricted access to ports with high risk	This control checks whether unrestricted incoming traffic for the security groups is accessible to the specified ports that have the highest risk. This control passes when none of the rules in a security group allow ingress traffic from 0.0.0.0/0 for those ports. Unrestricted access (0.0.0.0/0) increases opportunities for malicious activity, such as hacking, denial-of-service attacks, and loss of data. Security groups provide stateful filtering of ingress and egress network traffic to AWS resources. No security group should allow unrestricted ingress access to the following ports: - 3389 (RDP) - 20, 21 (FTP) - 22 (SSH) - 23 (Telnet) - 110 (POP3) - 143 (IMAP) - 3306 (mySQL) - 8080 (proxy) - 1433, 1434 (MSSQL) - 9200 or 9300 (Elasticsearch) - 5601 (Kibana) - 25 (SMTP) - 445 (CIFS) - 135 (RPC) - 4333 (ahsp) - 5432 (postgresql) - 5500 (fcp-addr-srvr1)	Medium

RECOMMENDATION	DESCRIPTION	SEVERITY
Security groups should only allow unrestricted incoming traffic for authorized ports	This control checks whether the security groups that are in use allow unrestricted incoming traffic. Optionally the rule checks whether the port numbers are listed in the "authorizedTcpPorts" parameter. - If the security group rule port number allows unrestricted incoming traffic, but the port number is specified in "authorizedTcpPorts", then the control passes. The default value for "authorizedTcpPorts" is **80, 443**. - If the security group rule port number allows unrestricted incoming traffic, but the port number is not specified in authorizedTcpPorts input parameter, then the control fails. - If the parameter is not used, then the control fails for any security group that has an unrestricted inbound rule. Security groups provide stateful filtering of ingress and egress network traffic to AWS. Security group rules should follow the principal of least privileged access. Unrestricted access (IP address with a /0 suffix) increases the opportunity for malicious activity such as hacking, denial-of-service attacks, and loss of data. Unless a port is specifically allowed, the port should deny unrestricted access.	High
Unused EC2 EIPs should be removed	Elastic IP addresses that are allocated to a VPC should be attached to Amazon EC2 instances or in-use elastic network interfaces (ENIs).	Low
Unused network access control lists should be removed	This control checks whether there are any unused network access control lists (ACLs). The control checks the item configuration of the resource "AWS::EC2::NetworkAcl" and determines the relationships of the network ACL. If the only relationship is the VPC of the network ACL, then the control fails. If other relationships are listed, then the control passes.	Low
VPC's default security group should restricts all traffic	Security group should restrict all traffic to reduce resource exposure.	Low

Next steps

For related information, see the following:

- Connect your AWS accounts to Microsoft Defender for Cloud

- What are security policies, initiatives, and recommendations?
- Review your security recommendations

Security alerts and incidents in Microsoft Defender for Cloud

2/15/2022 • 10 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud generates alerts for resources deployed on your Azure, on-premises, and hybrid cloud environments.

Security alerts are triggered by advanced detections and are available only with enhanced security features enabled. You can upgrade from the **Environment settings** page, as described in Quickstart: Enable enhanced security features. A free 30-day trial is available. For pricing details in your currency of choice and according to your region, see the pricing page.

What are security alerts and security incidents?

Alerts are the notifications that Defender for Cloud generates when it detects threats on your resources. Defender for Cloud prioritizes and lists the alerts, along with the information needed for you to quickly investigate the problem. Defender for Cloud also provides detailed steps to help you remediate attacks. Alerts data is retained for 90 days.

A security incident is a collection of related alerts, instead of listing each alert individually. Defender for Cloud uses Cloud smart alert correlation (incidents) to correlate different alerts and low fidelity signals into security incidents.

Using incidents, Defender for Cloud provides you with a single view of an attack campaign and all of the related alerts. This view enables you to quickly understand what actions the attacker took, and what resources were affected.

Respond to today's threats

There have been significant changes in the threat landscape over the last 20 years. In the past, companies typically only had to worry about web site defacement by individual attackers who were mostly interested in seeing "what they could do". Today's attackers are much more sophisticated and organized. They often have specific financial and strategic goals. They also have more resources available to them, as they might be funded by nation states or organized crime.

These changing realities have led to an unprecedented level of professionalism in the attacker ranks. No longer are they interested in web defacement. They are now interested in stealing information, financial accounts, and private data – all of which they can use to generate cash on the open market or to leverage a particular business, political, or military position. Even more concerning than those attackers with a financial objective are the attackers who breach networks to do harm to infrastructure and people.

In response, organizations often deploy various point solutions, which focus on defending either the enterprise perimeter or endpoints by looking for known attack signatures. These solutions tend to generate a high volume of low fidelity alerts, which require a security analyst to triage and investigate. Most organizations lack the time

and expertise required to respond to these alerts – so many go unaddressed.

In addition, attackers have evolved their methods to subvert many signature-based defenses and adapt to cloud environments. New approaches are required to more quickly identify emerging threats and expedite detection and response.

Continuous monitoring and assessments

Microsoft Defender for Cloud benefits from having security research and data science teams throughout Microsoft who continuously monitor for changes in the threat landscape. This includes the following initiatives:

- **Threat intelligence monitoring**: Threat intelligence includes mechanisms, indicators, implications, and actionable advice about existing or emerging threats. This information is shared in the security community and Microsoft continuously monitors threat intelligence feeds from internal and external sources.
- **Signal sharing**: Insights from security teams across Microsoft's broad portfolio of cloud and on-premises services, servers, and client endpoint devices are shared and analyzed.
- **Microsoft security specialists**: Ongoing engagement with teams across Microsoft that work in specialized security fields, like forensics and web attack detection.
- **Detection tuning**: Algorithms are run against real customer data sets and security researchers work with customers to validate the results. True and false positives are used to refine machine learning algorithms.

These combined efforts culminate in new and improved detections, which you can benefit from instantly – there's no action for you to take.

How does Defender for Cloud detect threats?

Microsoft security researchers are constantly on the lookout for threats. Because of our global presence in the cloud and on-premises, we have access to an expansive set of telemetry. The wide-reaching and diverse collection of datasets enables us to discover new attack patterns and trends across our on-premises consumer and enterprise products, as well as our online services. As a result, Defender for Cloud can rapidly update its detection algorithms as attackers release new and increasingly sophisticated exploits. This approach helps you keep pace with a fast moving threat environment.

To detect real threats and reduce false positives, Defender for Cloud collects, analyzes, and integrates log data from your Azure resources and the network. It also works with connected partner solutions, like firewall and endpoint protection solutions. Defender for Cloud analyzes this information, often correlating information from multiple sources, to identify threats.

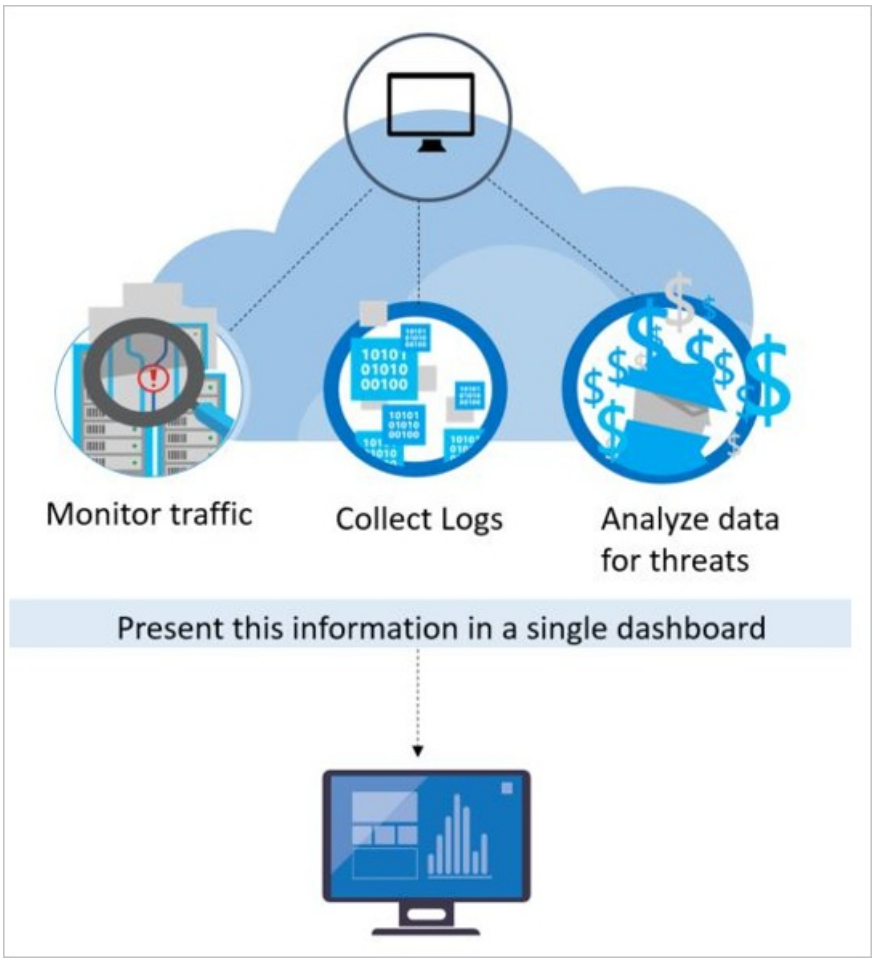

Defender for Cloud employs advanced security analytics, which go far beyond signature-based approaches. Breakthroughs in big data and machine learning technologies are leveraged to evaluate events across the entire cloud fabric – detecting threats that would be impossible to identify using manual approaches and predicting the evolution of attacks. These security analytics include:

- **Integrated threat intelligence**: Microsoft has an immense amount of global threat intelligence. Telemetry flows in from multiple sources, such as Azure, Microsoft 365, Microsoft CRM online, Microsoft Dynamics AX, outlook.com, MSN.com, the Microsoft Digital Crimes Unit (DCU), and Microsoft Security Response Center (MSRC). Researchers also receive threat intelligence information that is shared among major cloud service providers and feeds from other third parties. Microsoft Defender for Cloud can use this information to alert you to threats from known bad actors.
- **Behavioral analytics**: Behavioral analytics is a technique that analyzes and compares data to a collection of known patterns. However, these patterns are not simple signatures. They are determined through complex machine learning algorithms that are applied to massive datasets. They are also determined through careful analysis of malicious behaviors by expert analysts. Microsoft Defender for Cloud can use behavioral analytics to identify compromised resources based on analysis of virtual machine logs, virtual network device logs, fabric logs, and other sources.
- **Anomaly detection**: Microsoft Defender for Cloud also uses anomaly detection to identify threats. In contrast to behavioral analytics (which depends on known patterns derived from large data sets), anomaly detection is more "personalized" and focuses on baselines that are specific to your deployments. Machine learning is applied to determine normal activity for your deployments and then rules are generated to define outlier conditions that could represent a security event.

How are alerts classified?

Defender for Cloud assigns a severity to alerts, to help you prioritize the order in which you attend to each alert, so that when a resource is compromised, you can get to it right away. The severity is based on how confident

Defender for Cloud is in the finding or the analytic used to issue the alert as well as the confidence level that there was malicious intent behind the activity that led to the alert.

> **NOTE**
> Alert severity is displayed differently in the portal and versions of the REST API that predate 01-01-2019. If you're using an older version of the API, upgrade for the consistent experience described below.

SEVERITY	RECOMMENDED RESPONSE
High	There is a high probability that your resource is compromised. You should look into it right away. Defender for Cloud has high confidence in both the malicious intent and in the findings used to issue the alert. For example, an alert that detects the execution of a known malicious tool such as Mimikatz, a common tool used for credential theft.
Medium	This is probably a suspicious activity might indicate that a resource is compromised. Defender for Cloud's confidence in the analytic or finding is medium and the confidence of the malicious intent is medium to high. These would usually be machine learning or anomaly-based detections. For example, a sign-in attempt from an anomalous location.
Low	This might be a benign positive or a blocked attack. Defender for Cloud isn't confident enough that the intent is malicious and the activity might be innocent. For example, log clear is an action that might happen when an attacker tries to hide their tracks, but in many cases is a routine operation performed by admins. Defender for Cloud doesn't usually tell you when attacks were blocked, unless it's an interesting case that we suggest you look into.
Informational	An incident is typically made up of a number of alerts, some of which might appear on their own to be only informational, but in the context of the other alerts might be worthy of a closer look.

Export alerts

You have a range of options for viewing your alerts outside of Defender for Cloud, including:

- **Download CSV report** on the alerts dashboard provides a one-time export to CSV.
- **Continuous export** from Environment settings allows you to configure streams of security alerts and recommendations to Log Analytics workspaces and Event Hubs. Learn more about continuous export.
- **Microsoft Sentinel connector** streams security alerts from Microsoft Defender for Cloud into Microsoft Sentinel. Learn more about connecting Microsoft Defender for Cloud with Microsoft Sentinel.

Learn about all of the export options in Stream alerts to a SIEM, SOAR, or IT Service Management solution and Continuously export Defender for Cloud data.

Cloud smart alert correlation (incidents)

Microsoft Defender for Cloud continuously analyzes hybrid cloud workloads by using advanced analytics and threat intelligence to alert you about malicious activity.

The breadth of threat coverage is growing. The need to detect even the slightest compromise is important, and it can be challenging for security analysts to triage the different alerts and identify an actual attack. Defender for Cloud helps analysts cope with this alert fatigue. It helps diagnose attacks as they occur, by correlating different alerts and low fidelity signals into security incidents.

Fusion analytics is the technology and analytic back end that powers Defender for Cloud incidents, enabling it to correlate different alerts and contextual signals together. Fusion looks at the different signals reported on a subscription across the resources. Fusion finds patterns that reveal attack progression or signals with shared contextual information, indicating that you should use a unified response procedure for them.

Fusion analytics combines security domain knowledge with AI to analyze alerts, discovering new attack patterns as they occur.

Defender for Cloud leverages MITRE Attack Matrix to associate alerts with their perceived intent, helping formalize security domain knowledge. In addition, by using the information gathered for each step of an attack, Defender for Cloud can rule out activity that appears to be steps of an attack, but actually isn't.

Because attacks often occur across different tenants, Defender for Cloud can combine AI algorithms to analyze attack sequences that are reported on each subscription. This technique identifies the attack sequences as prevalent alert patterns, instead of just being incidentally associated with each other.

During an investigation of an incident, analysts often need extra context to reach a verdict about the nature of the threat and how to mitigate it. For example, even when a network anomaly is detected, without understanding what else is happening on the network or with regard to the targeted resource, it's difficult to understand what actions to take next. To help, a security incident can include artifacts, related events, and information. The additional information available for security incidents varies, depending on the type of threat detected and the configuration of your environment.

> **TIP**
> For a list of security incident alerts that can be produced by the fusion analytics, see the Reference table of alerts.

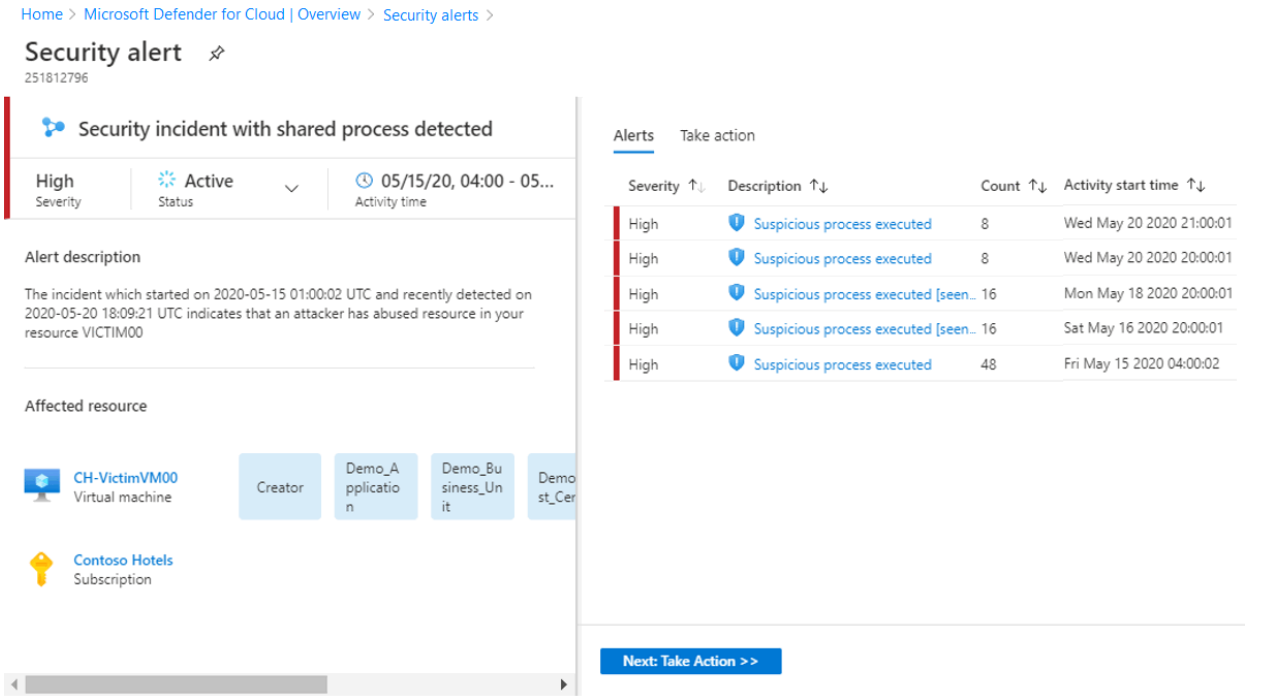

To manage your security incidents, see How to manage security incidents in Microsoft Defender for Cloud.

Next steps

In this article, you learned about the different types of alerts available in Defender for Cloud. For more information, see:

- Security alerts in Azure Activity log - In addition to being available in the Azure portal or programmatically, Security alerts and incidents are audited as events in Azure Activity Log
- Reference table of Defender for Cloud alerts
- Respond to security alerts

Security alerts - a reference guide

2/15/2022 • 127 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This article lists the security alerts you might get from Microsoft Defender for Cloud and any Microsoft Defender plans you've enabled. The alerts shown in your environment depend on the resources and services you're protecting, as well as your customized configuration.

At the bottom of this page, there's a table describing the Microsoft Defender for Cloud kill chain aligned with version 7 of the MITRE ATT&CK matrix.

Learn how to respond to these alerts.

Learn how to export alerts.

> **NOTE**
> Alerts from different sources might take different amounts of time to appear. For example, alerts that require analysis of network traffic might take longer to appear than alerts related to suspicious processes running on virtual machines.

Alerts for Windows machines

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
A logon from a malicious IP has been detected. [seen multiple times]	A successful remote authentication for the account [account] and process [process] occurred, however the logon IP address (x.x.x.x) has previously been reported as malicious or highly unusual. A successful attack has probably occurred. Files with the .scr extensions are screen saver files and are normally reside and execute from the Windows system directory.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Addition of Guest account to Local Administrators group	Analysis of host data has detected the addition of the built-in Guest account to the Local Administrators group on %{Compromised Host}, which is strongly associated with attacker activity.	-	Medium
An event log was cleared	Machine logs indicate a suspicious event log clearing operation by user: '%{user name}' in Machine: '%{CompromisedEntity}'. The %{log channel} log was cleared.	-	Informational
Antimalware Action Failed	Microsoft Antimalware has encountered an error when taking an action on malware or other potentially unwanted software.	-	Medium
Antimalware Action Taken	Microsoft Antimalware for Azure has taken an action to protect this machine from malware or other potentially unwanted software.	-	Medium
Antimalware broad files exclusion in your virtual machine (VM_AmBroadFilesExclusion)	Files exclusion from antimalware extension with broad exclusion rule was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Such exclusion practically disabling the Antimalware protection. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware disabled and code execution in your virtual machine (VM_AmDisablementAndCodeExecution)	Antimalware disabled at the same time as code execution on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers disable antimalware scanners to prevent detection while running unauthorized tools or infecting the machine with malware.	-	High
Antimalware disabled in your virtual machine (VM_AmDisablement)	Antimalware disabled in your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might disable the antimalware on your virtual machine to prevent detection.	Defense Evasion	Medium
Antimalware file exclusion and code execution in your virtual machine (VM_AmFileExclusionAndCodeExecution)	File excluded from your antimalware scanner at the same time as code was executed via a custom script extension on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running unauthorized tools or infecting the machine with malware.	Defense Evasion, Execution	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware file exclusion and code execution in your virtual machine (VM_AmTempFileExclusionAndCodeExecution)	Temporary file exclusion from antimalware extension in parallel to execution of code via custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion, Execution	High
Antimalware file exclusion in your virtual machine (VM_AmTempFileExclusion)	File excluded from your antimalware scanner on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running unauthorized tools or infecting the machine with malware.	Defense Evasion	Medium
Antimalware real-time protection was disabled in your virtual machine (VM_AmRealtimeProtectionDisabled)	Real-time protection disablement of the antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware real-time protection was disabled temporarily in your virtual machine (VM_AmTempRealtimeProtectionDisablement)	Real-time protection temporary disablement of the antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium
Antimalware real-time protection was disabled temporarily while code was executed in your virtual machine (VM_AmRealtimeProtectionDisablementAndCodeExec)	Real-time protection temporary disablement of the antimalware extension in parallel to code execution via custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	-	High
Antimalware scans blocked for files potentially related to malware campaigns on your virtual machine (Preview) (VM_AmMalwareCampaignRelatedExclusion)	An exclusion rule was detected in your virtual machine to prevent your antimalware extension scanning certain files that are suspected of being related to a malware campaign. The rule was detected by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from antimalware scans to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware temporarily disabled in your virtual machine (VM_AmTemporarilyDisablement)	Antimalware temporarily disabled in your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might disable the antimalware on your virtual machine to prevent detection.	-	Medium
Antimalware unusual file exclusion in your virtual machine (VM_UnusualAmFileExclusion)	Unusual file exclusion from antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium
Communication with suspicious domain identified by threat intelligence (AzureDNS_ThreatIntelSuspectDomain)	Communication with suspicious domain was detected by analyzing DNS transactions from your resource and comparing against known malicious domains identified by threat intelligence feeds. Communication to malicious domains is frequently performed by attackers and could imply that your resource is compromised.	Initial Access, Persistence, Execution, Command And Control, Exploitation	Medium
Custom script extension with suspicious command in your virtual machine (VM_CustomScriptExtensionSuspiciousCmd)	Custom script extension with suspicious command was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extension to execute a malicious code on your virtual machine via the Azure Resource Manager.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Custom script extension with suspicious entry-point in your virtual machine (VM_CustomScriptExtensionSuspiciousEntryPoint)	Custom script extension with a suspicious entry-point was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. The entry-point refers to a suspicious GitHub repository. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Custom script extension with suspicious payload in your virtual machine (VM_CustomScriptExtensionSuspiciousPayload)	Custom script extension with a payload from a suspicious GitHub repository was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Detected actions indicative of disabling and deleting IIS log files	Analysis of host data detected actions that show IIS log files being disabled and/or deleted.	-	Medium
Detected anomalous mix of upper and lower case characters in command-line	Analysis of host data on % {Compromised Host} detected a command line with anomalous mix of upper and lower case characters. This kind of pattern, while possibly benign, is also typical of attackers trying to hide from case-sensitive or hash-based rule matching when performing administrative tasks on a compromised host.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected change to a registry key that can be abused to bypass UAC	Analysis of host data on % {Compromised Host} detected that a registry key that can be abused to bypass UAC (User Account Control) was changed. This kind of configuration, while possibly benign, is also typical of attacker activity when trying to move from unprivileged (standard user) to privileged (for example administrator) access on a compromised host.	-	Medium
Detected decoding of an executable using built-in certutil.exe tool	Analysis of host data on % {Compromised Host} detected that certutil.exe, a built-in administrator utility, was being used to decode an executable instead of its mainstream purpose that relates to manipulating certificates and certificate data. Attackers are known to abuse functionality of legitimate administrator tools to perform malicious actions, for example using a tool such as certutil.exe to decode a malicious executable that will then be subsequently executed.	-	High
Detected enabling of the WDigest UseLogonCredential registry key	Analysis of host data detected a change in the registry key HKLM\SYSTEM\ CurrentControlSet\Control\ SecurityProviders\WDigest\ "UseLogonCredential". Specifically this key has been updated to allow logon credentials to be stored in clear text in LSA memory. Once enabled an attacker can dump clear text passwords from LSA memory with credential harvesting tools such as Mimikatz.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected encoded executable in command line data	Analysis of host data on % {Compromised Host} detected a base-64 encoded executable. This has previously been associated with attackers attempting to construct executables on-the-fly through a sequence of commands, and attempting to evade intrusion detection systems by ensuring that no individual command would trigger an alert. This could be legitimate activity, or an indication of a compromised host.	-	High
Detected obfuscated command line	Attackers use increasingly complex obfuscation techniques to evade detections that run against the underlying data. Analysis of host data on % {Compromised Host} detected suspicious indicators of obfuscation on the commandline.	-	Informational
Detected Petya ransomware indicators	Analysis of host data on % {Compromised Host} detected indicators associated with Petya ransomware. See https://aka.ms/petya-blog for more information. Review the command line associated in this alert and escalate this alert to your security team.	-	High
Detected possible execution of keygen executable	Analysis of host data on % {Compromised Host} detected execution of a process whose name is indicative of a keygen tool; such tools are typically used to defeat software licensing mechanisms but their download is often bundled with other malicious software. Activity group GOLD has been known to make use of such keygens to covertly gain back door access to hosts that they compromise.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected possible execution of malware dropper	Analysis of host data on %{Compromised Host} detected a filename that has previously been associated with one of activity group GOLD's methods of installing malware on a victim host.	-	High
Detected possible local reconnaissance activity	Analysis of host data on %{Compromised Host} detected a combination of systeminfo commands that has previously been associated with one of activity group GOLD's methods of performing reconnaissance activity. While 'systeminfo.exe' is a legitimate Windows tool, executing it twice in succession in the way that has occurred here is rare.	-	
Detected potentially suspicious use of Telegram tool	Analysis of host data shows installation of Telegram, a free cloud-based instant messaging service that exists both for mobile and desktop system. Attackers are known to abuse this service to transfer malicious binaries to any other computer, phone, or tablet.	-	Medium
Detected suppression of legal notice displayed to users at logon	Analysis of host data on %{Compromised Host} detected changes to the registry key that controls whether a legal notice is displayed to users when they log on. Microsoft security analysis has determined that this is a common activity undertaken by attackers after having compromised a host.	-	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious combination of HTA and PowerShell	mshta.exe (Microsoft HTML Application Host) which is a signed Microsoft binary is being used by the attackers to launch malicious PowerShell commands. Attackers often resort to having an HTA file with inline VBScript. When a victim browses to the HTA file and chooses to run it, the PowerShell commands and scripts that it contains are executed. Analysis of host data on % {Compromised Host} detected mshta.exe launching PowerShell commands.	-	Medium
Detected suspicious commandline arguments	Analysis of host data on % {Compromised Host} detected suspicious commandline arguments that have been used in conjunction with a reverse shell used by activity group HYDROGEN.	-	High
Detected suspicious commandline used to start all executables in a directory	Analysis of host data has detected a suspicious process running on % {Compromised Host}. The commandline indicates an attempt to start all executables (*.exe) that may reside in a directory. This could be an indication of a compromised host.	-	Medium
Detected suspicious credentials in commandline	Analysis of host data on % {Compromised Host} detected a suspicious password being used to execute a file by activity group BORON. This activity group has been known to use this password to execute Pirpi malware on a victim host.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious document credentials	Analysis of host data on % {Compromised Host} detected a suspicious, common precomputed password hash used by malware being used to execute a file. Activity group HYDROGEN has been known to use this password to execute malware on a victim host.	-	High
Detected suspicious execution of VBScript.Encode command	Analysis of host data on % {Compromised Host} detected the execution of VBScript.Encode command. This encodes the scripts into unreadable text, making it more difficult for users to examine the code. Microsoft threat research shows that attackers often use encoded VBscript files as part of their attack to evade detection systems. This could be legitimate activity, or an indication of a compromised host.	-	Medium
Detected suspicious execution via rundll32.exe	Analysis of host data on % {Compromised Host} detected rundll32.exe being used to execute a process with an uncommon name, consistent with the process naming scheme previously seen used by activity group GOLD when installing their first stage implant on a compromised host.	-	High
Detected suspicious file cleanup commands	Analysis of host data on % {Compromised Host} detected a combination of systeminfo commands that has previously been associated with one of activity group GOLD's methods of performing post-compromise self-cleanup activity. While 'systeminfo.exe' is a legitimate Windows tool, executing it twice in succession, followed by a delete command in the way that has occurred here is rare.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious file creation	Analysis of host data on %{Compromised Host} detected creation or execution of a process which has previously indicated post-compromise action taken on a victim host by activity group BARIUM. This activity group has been known to use this technique to download additional malware to a compromised host after an attachment in a phishing doc has been opened.	-	High
Detected suspicious named pipe communications	Analysis of host data on %{Compromised Host} detected data being written to a local named pipe from a Windows console command. Named pipes are known to be a channel used by attackers to task and communicate with a malicious implant. This could be legitimate activity, or an indication of a compromised host.	-	High
Detected suspicious network activity	Analysis of network traffic from %{Compromised Host} detected suspicious network activity. Such traffic, while possibly benign, is typically used by an attacker to communicate with malicious servers for downloading of tools, command-and-control and exfiltration of data. Typical related attacker activity includes copying remote administration tools to a compromised host and exfiltrating user data from it.	-	Low
Detected suspicious new firewall rule	Analysis of host data detected a new firewall rule has been added via netsh.exe to allow traffic from an executable in a suspicious location.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious use of Cacls to lower the security state of the system	Attackers use myriad ways like brute force, spear phishing etc. to achieve initial compromise and get a foothold on the network. Once initial compromise is achieved they often take steps to lower the security settings of a system. Cacls —short for change access control list is Microsoft Windows native command-line utility often used for modifying the security permission on folders and files. A lot of time the binary is used by the attackers to lower the security settings of a system. This is done by giving Everyone full access to some of the system binaries like ftp.exe, net.exe, wscript.exe etc. Analysis of host data on % {Compromised Host} detected suspicious use of Cacls to lower the security of a system.	-	Medium
Detected suspicious use of FTP -s Switch	Analysis of process creation data from the % {Compromised Host} detected the use of the FTP "-s:filename" switch. This switch is used to specify an FTP script file for the client to run. Malware or malicious processes are known to use this FTP switch (-s:filename) to point to a script file which is configured to connect to a remote FTP server and download additional malicious binaries.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious use of Pcalua.exe to launch executable code	Analysis of host data on %{Compromised Host} detected the use of pcalua.exe to launch executable code. Pcalua.exe is component of the Microsoft Windows "Program Compatibility Assistant" which detects compatibility issues during the installation or execution of a program. Attackers are known to abuse functionality of legitimate Windows system tools to perform malicious actions, for example using pcalua.exe with the -a switch to launch malicious executables either locally or from remote shares.	-	Medium
Detected the disabling of critical services	The analysis of host data on %{Compromised Host} detected execution of "net.exe stop" command being used to stop critical services like SharedAccess or the Windows Security app. The stopping of either of these services can be indication of a malicious behavior.	-	Medium
Digital currency mining related behavior detected	Analysis of host data on %{Compromised Host} detected the execution of a process or command normally associated with digital currency mining.	-	High
Dynamic PS script construction	Analysis of host data on %{Compromised Host} detected a PowerShell script being constructed dynamically. Attackers sometimes use this approach of progressively building up a script in order to evade IDS systems. This could be legitimate activity, or an indication that one of your machines has been compromised.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Executable found running from a suspicious location	Analysis of host data detected an executable file on %{Compromised Host} that is running from a location in common with known suspicious files. This executable could either be legitimate activity, or an indication of a compromised host.	-	High
Fileless attack behavior detected (VM_FilelessAttackBehavior.Windows)	The memory of the process specified contains behaviors commonly used by fileless attacks. Specific behaviors include: 1) Shellcode, which is a small piece of code typically used as the payload in the exploitation of a software vulnerability. 2) Active network connections. See NetworkConnections below for details. 3) Function calls to security sensitive operating system interfaces. See Capabilities below for referenced OS capabilities. 4) Contains a thread that was started in a dynamically allocated code segment. This is a common pattern for process injection attacks.	Defense Evasion	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Fileless attack technique detected (VM_FilelessAttackTechnique.Windows)	The memory of the process specified below contains evidence of a fileless attack technique. Fileless attacks are used by attackers to execute code while evading detection by security software. Specific behaviors include: 1) Shellcode, which is a small piece of code typically used as the payload in the exploitation of a software vulnerability. 2) Executable image injected into the process, such as in a code injection attack. 3) Active network connections. See NetworkConnections below for details. 4) Function calls to security sensitive operating system interfaces. See Capabilities below for referenced OS capabilities. 5) Process hollowing, which is a technique used by malware in which a legitimate process is loaded on the system to act as a container for hostile code. 6) Contains a thread that was started in a dynamically allocated code segment. This is a common pattern for process injection attacks.	Defense Evasion, Execution	High
Fileless attack toolkit detected (VM_FilelessAttackToolkit.Windows)	The memory of the process specified contains a fileless attack toolkit: [toolkit name]. Fileless attack toolkits use techniques that minimize or eliminate traces of malware on disk, and greatly reduce the chances of detection by disk-based malware scanning solutions. Specific behaviors include: 1) Well-known toolkits and crypto mining software. 2) Shellcode, which is a small piece of code typically used as the payload in the exploitation of a software vulnerability. 3) Injected malicious executable in process memory.	Defense Evasion, Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
High risk software detected	Analysis of host data from %{Compromised Host} detected the usage of software that has been associated with the installation of malware in the past. A common technique utilized in the distribution of malicious software is to package it within otherwise benign tools such as the one seen in this alert. Upon using these tools, the malware can be silently installed in the background.	-	Medium
Local Administrators group members were enumerated	Machine logs indicate a successful enumeration on group %{Enumerated Group Domain Name}% {Enumerated Group Name}. Specifically, %{Enumerating User Domain Name}% {Enumerating User Name} remotely enumerated the members of the % {Enumerated Group Domain Name}% {Enumerated Group Name} group. This activity could either be legitimate activity, or an indication that a machine in your organization has been compromised and used to reconnaissance %{vmname}.	-	Informational
Malicious firewall rule created by ZINC server implant [seen multiple times]	A firewall rule was created using techniques that match a known actor, ZINC. The rule was possibly used to open a port on % {Compromised Host} to allow for Command & Control communications. This behavior was seen [x] times today on the following machines: [Machine names]	-	High
Malicious SQL activity	Machine logs indicate that '%{process name}' was executed by account: % {user name}. This activity is considered malicious.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Multiple Domain Accounts Queried	Analysis of host data has determined that an unusual number of distinct domain accounts are being queried within a short time period from %{Compromised Host}. This kind of activity could be legitimate, but can also be an indication of compromise.	-	Medium
Possible credential dumping detected [seen multiple times]	Analysis of host data has detected use of native windows tool (e.g. sqldumper.exe) being used in a way that allows to extract credentials from memory. Attackers often use these techniques to extract credentials that they then further use for lateral movement and privilege escalation. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Potential attempt to bypass AppLocker detected	Analysis of host data on % {Compromised Host} detected a potential attempt to bypass AppLocker restrictions. AppLocker can be configured to implement a policy that limits what executables are allowed to run on a Windows system. The command-line pattern similar to that identified in this alert has been previously associated with attacker attempts to circumvent AppLocker policy by using trusted executables (allowed by AppLocker policy) to execute untrusted code. This could be legitimate activity, or an indication of a compromised host.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PsExec execution detected (VM_RunByPsExec)	Analysis of host data indicates that the process %{Process Name} was executed by PsExec utility. PsExec can be used for running processes remotely. This technique might be used for malicious purposes.	Lateral Movement, Execution	Informational
Ransomware indicators detected [seen multiple times]	Analysis of host data indicates suspicious activity traditionally associated with lock-screen and encryption ransomware. Lock screen ransomware displays a full-screen message preventing interactive use of the host and access to its files. Encryption ransomware prevents access by encrypting data files. In both cases a ransom message is typically displayed, requesting payment in order to restore file access. This behavior was seen [x] times today on the following machines: [Machine names]	-	High
Ransomware indicators detected	Analysis of host data indicates suspicious activity traditionally associated with lock-screen and encryption ransomware. Lock screen ransomware displays a full-screen message preventing interactive use of the host and access to its files. Encryption ransomware prevents access by encrypting data files. In both cases a ransom message is typically displayed, requesting payment in order to restore file access.	-	High
Rare SVCHOST service group executed (VM_SvcHostRunInRareServiceGroup)	The system process SVCHOST was observed running a rare service group. Malware often uses SVCHOST to masquerade its malicious activity.	Defense Evasion, Execution	Informational

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Sticky keys attack detected	Analysis of host data indicates that an attacker may be subverting an accessibility binary (for example sticky keys, onscreen keyboard, narrator) in order to provide backdoor access to the host %{Compromised Host}.	-	Medium
Successful brute force attack (VM_LoginBruteForceSuccess)	Several sign in attempts were detected from the same source. Some successfully authenticated to the host. This resembles a burst attack, in which an attacker performs numerous authentication attempts to find valid account credentials.	Exploitation	Medium/High
Suspect integrity level indicative of RDP hijacking	Analysis of host data has detected the tscon.exe running with SYSTEM privileges - this can be indicative of an attacker abusing this binary in order to switch context to any other logged on user on this host; it is a known attacker technique to compromise additional user accounts and move laterally across a network.	-	Medium
Suspect service installation	Analysis of host data has detected the installation of tscon.exe as a service: this binary being started as a service potentially allows an attacker to trivially switch to any other logged on user on this host by hijacking RDP connections; it is a known attacker technique to compromise additional user accounts and move laterally across a network.	-	Medium
Suspected Kerberos Golden Ticket attack parameters observed	Analysis of host data detected commandline parameters consistent with a Kerberos Golden Ticket attack.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious Account Creation Detected	Analysis of host data on %{Compromised Host} detected creation or use of a local account %{Suspicious account name} : this account name closely resembles a standard Windows account or group name '%{Similar To Account Name}'. This is potentially a rogue account created by an attacker, so named in order to avoid being noticed by a human administrator.	-	Medium
Suspicious Activity Detected (VM_SuspiciousActivity)	Analysis of host data has detected a sequence of one or more processes running on %{machine name} that have historically been associated with malicious activity. While individual commands may appear benign the alert is scored based on an aggregation of these commands. This could either be legitimate activity, or an indication of a compromised host.	Execution	Medium
Suspicious authentication activity (VM_LoginBruteForceValidUserFailed)	Although none of them succeeded, some of them used accounts were recognized by the host. This resembles a dictionary attack, in which an attacker performs numerous authentication attempts using a dictionary of predefined account names and passwords in order to find valid credentials to access the host. This indicates that some of your host account names might exist in a well-known account name dictionary.	Probing	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious code segment detected	Indicates that a code segment has been allocated by using non-standard methods, such as reflective injection and process hollowing. The alert provides additional characteristics of the code segment that have been processed to provide context for the capabilities and behaviors of the reported code segment.	-	Medium
Suspicious command execution (VM_SuspiciousCommandLineExecution)	Machine logs indicate a suspicious command-line execution by user %{user name}.	Execution	High
Suspicious double extension file executed	Analysis of host data indicates an execution of a process with a suspicious double extension. This extension may trick users into thinking files are safe to be opened and might indicate the presence of malware on the system.	-	High
Suspicious download using Certutil detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the use of certutil.exe, a built-in administrator utility, for the download of a binary instead of its mainstream purpose that relates to manipulating certificates and certificate data. Attackers are known to abuse functionality of legitimate administrator tools to perform malicious actions, for example using certutil.exe to download and decode a malicious executable that will then be subsequently executed. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious download using Certutil detected	Analysis of host data on % {Compromised Host} detected the use of certutil.exe, a built-in administrator utility, for the download of a binary instead of its mainstream purpose that relates to manipulating certificates and certificate data. Attackers are known to abuse functionality of legitimate administrator tools to perform malicious actions, for example using certutil.exe to download and decode a malicious executable that will then be subsequently executed.	-	Medium
Suspicious failed execution of custom script extension in your virtual machine (VM_CustomScriptExtension SuspiciousFailure)	Suspicious failure of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Such failures may be associated with malicious scripts run by this extension.	Execution	Medium
Suspicious PowerShell Activity Detected	Analysis of host data detected a PowerShell script running on % {Compromised Host} that has features in common with known suspicious scripts. This script could either be legitimate activity, or an indication of a compromised host.	-	High
Suspicious PowerShell cmdlets executed	Analysis of host data indicates execution of known malicious PowerShell PowerSploit cmdlets.	-	Medium
Suspicious process executed [seen multiple times]	Machine logs indicate that the suspicious process: '% {Suspicious Process}' was running on the machine, often associated with attacker attempts to access credentials. This behavior was seen [x] times today on the following machines: [Machine names]	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious process executed	Machine logs indicate that the suspicious process: '% {Suspicious Process}' was running on the machine, often associated with attacker attempts to access credentials.	-	High
Suspicious process name detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected a process whose name is suspicious, for example corresponding to a known attacker tool or named in a way that is suggestive of attacker tools that try to hide in plain sight. This process could be legitimate activity, or an indication that one of your machines has been compromised. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Suspicious process name detected	Analysis of host data on % {Compromised Host} detected a process whose name is suspicious, for example corresponding to a known attacker tool or named in a way that is suggestive of attacker tools that try to hide in plain sight. This process could be legitimate activity, or an indication that one of your machines has been compromised.	-	Medium
Suspicious process termination burst (VM_TaskkillBurst)	Analysis of host data indicates a suspicious process termination burst in %{Machine Name}. Specifically, % {NumberOfCommands} processes were killed between %{Begin} and % {Ending}.	Defense Evasion	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious Screensaver process executed (VM_SuspiciousScreenSaverExecution)	The process '%{process name}' was observed executing from an uncommon location. Files with the .scr extensions are screen saver files and are normally reside and execute from the Windows system directory.	Defense Evasion, Execution	Medium
Suspicious SQL activity	Machine logs indicate that '%{process name}' was executed by account: %{user name}. This activity is uncommon with this account.	-	Medium
Suspicious SVCHOST process executed	The system process SVCHOST was observed running in an abnormal context. Malware often uses SVCHOST to masquerade its malicious activity.	-	High
Suspicious system process executed (VM_SystemProcessInAbnormalContext)	The system process %{process name} was observed running in an abnormal context. Malware often uses this process name to masquerade its malicious activity.	Defense Evasion, Execution	High
Suspicious Volume Shadow Copy Activity	Analysis of host data has detected a shadow copy deletion activity on the resource. Volume Shadow Copy (VSC) is an important artifact that stores data snapshots. Some malware and specifically Ransomware, targets VSC to sabotage backup strategies.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious WindowPosition registry value detected	Analysis of host data on % {Compromised Host} detected an attempted WindowPosition registry configuration change that could be indicative of hiding application windows in non-visible sections of the desktop. This could be legitimate activity, or an indication of a compromised machine: this type of activity has been previously associated with known adware (or unwanted software) such as Win32/OneSystemCare and Win32/SystemHealer and malware such as Win32/Creprote. When the WindowPosition value is set to 201329664, (Hex: 0x0c00 0c00, corresponding to X-axis=0c00 and the Y-axis=0c00) this places the console app's window in a non-visible section of the user's screen in an area that is hidden from view below the visible start menu/taskbar. Known suspect Hex value includes, but not limited to c000c000	-	Low
Suspiciously named process detected	Analysis of host data on % {Compromised Host} detected a process whose name is very similar to but different from a very commonly run process (% {Similar To Process Name}). While this process could be benign attackers are known to sometimes hide in plain sight by naming their malicious tools to resemble legitimate process names.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual config reset in your virtual machine (VM_VMAccessUnusualConfigReset)	An unusual config reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing VM Access extension to reset the configuration in your virtual machine and compromise it.	Credential Access	Medium
Unusual deletion of custom script extension in your virtual machine (VM_CustomScriptExtensionUnusualDeletion)	Unusual deletion of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Unusual execution of custom script extension in your virtual machine (VM_CustomScriptExtensionUnusualExecution)	Unusual execution of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Unusual process execution detected	Analysis of host data on %{Compromised Host} detected the execution of a process by %{User Name} that was unusual. Accounts such as %{User Name} tend to perform a limited set of operations, this execution was determined to be out of character and may be suspicious.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual user password reset in your virtual machine (VM_VMAccessUnusualPasswordReset)	An unusual user password reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing the VM Access extension to reset the credentials of a local user in your virtual machine and compromise it.	Credential Access	Medium
Unusual user SSH key reset in your virtual machine (VM_VMAccessUnusualSSHReset)	An unusual user SSH key reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing VM Access extension to reset SSH key of a user account in your virtual machine and compromise it.	Credential Access	Medium
VBScript HTTP object allocation detected	Creation of a VBScript file using Command Prompt has been detected. The following script contains HTTP object allocation command. This action can be used to download malicious files.	-	High
Windows registry persistence method detected (VM_RegistryPersistencyKey)	Analysis of host data has detected an attempt to persist an executable in the Windows registry. Malware often uses such a technique to survive a boot.	Persistence	Low

Alerts for Linux machines

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
a history file has been cleared	Analysis of host data indicates that the command history log file has been cleared. Attackers may do this to cover their traces. The operation was performed by user: '%{user name}'.	-	Medium
Access of htaccess file detected (VM_SuspectHtaccessFileAccess)	Analysis of host data on % {Compromised Host} detected possible manipulation of a htaccess file. Htaccess is a powerful configuration file that allows you to make multiple changes to a web server running the Apache Web software including basic redirect functionality, or for more advanced functions such as basic password protection. Attackers will often modify htaccess files on machines they have compromised to gain persistence.	Persistence, Defense Evasion, Execution	Medium
Antimalware broad files exclusion in your virtual machine (VM_AmBroadFilesExclusion)	Files exclusion from antimalware extension with broad exclusion rule was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Such exclusion practically disabling the Antimalware protection. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware disabled and code execution in your virtual machine (VM_AmDisablementAndCodeExecution)	Antimalware disabled at the same time as code execution on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers disable antimalware scanners to prevent detection while running unauthorized tools or infecting the machine with malware.	-	High
Antimalware disabled in your virtual machine (VM_AmDisablement)	Antimalware disabled in your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might disable the antimalware on your virtual machine to prevent detection.	Defense Evasion	Medium
Antimalware file exclusion and code execution in your virtual machine (VM_AmFileExclusionAndCodeExecution)	File excluded from your antimalware scanner at the same time as code was executed via a custom script extension on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running unauthorized tools or infecting the machine with malware.	Defense Evasion, Execution	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware file exclusion and code execution in your virtual machine (VM_AmTempFileExclusionAndCodeExecution)	Temporary file exclusion from antimalware extension in parallel to execution of code via custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion, Execution	High
Antimalware file exclusion in your virtual machine (VM_AmTempFileExclusion)	File excluded from your antimalware scanner on your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running unauthorized tools or infecting the machine with malware.	Defense Evasion	Medium
Antimalware real-time protection was disabled in your virtual machine (VM_AmRealtimeProtectionDisabled)	Real-time protection disablement of the antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware real-time protection was disabled temporarily in your virtual machine (VM_AmTempRealtimeProtectionDisablement)	Real-time protection temporary disablement of the antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium
Antimalware real-time protection was disabled temporarily while code was executed in your virtual machine (VM_AmRealtimeProtectionDisablementAndCodeExec)	Real-time protection temporary disablement of the antimalware extension in parallel to code execution via custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might disable real-time protection from the antimalware scan on your virtual machine to avoid detection while running arbitrary code or infecting the machine with malware.	-	High
Antimalware scans blocked for files potentially related to malware campaigns on your virtual machine (Preview) (VM_AmMalwareCampaignRelatedExclusion)	An exclusion rule was detected in your virtual machine to prevent your antimalware extension scanning certain files that are suspected of being related to a malware campaign. The rule was detected by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from antimalware scans to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Antimalware temporarily disabled in your virtual machine (VM_AmTemporarilyDisablement)	Antimalware temporarily disabled in your virtual machine. This was detected by analyzing Azure Resource Manager operations in your subscription. Attackers might disable the antimalware on your virtual machine to prevent detection.	-	Medium
Antimalware unusual file exclusion in your virtual machine (VM_UnusualAmFileExclusion)	Unusual file exclusion from antimalware extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers might exclude files from the antimalware scan on your virtual machine to prevent detection while running arbitrary code or infecting the machine with malware.	Defense Evasion	Medium
Attempt to stop apt-daily-upgrade.timer service detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected an attempt to stop apt-daily-upgrade.timer service. In some recent attacks, attackers have been observed stopping this service, to download malicious files and granting execution privileges for their attack. This behavior was seen [x] times today on the following machines: [Machine names]	-	Low
Attempt to stop apt-daily-upgrade.timer service detected (VM_TimerServiceDisabled)	Analysis of host data on % {Compromised Host} detected an attempt to stop apt-daily-upgrade.timer service. In some recent attacks, attackers have been observed stopping this service, to download malicious files and granting execution privileges for their attack.	Defense Evasion	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Behavior similar to common Linux bots detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of a process normally associated with common Linux botnets. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Behavior similar to common Linux bots detected (VM_CommonBot)	Analysis of host data on % {Compromised Host} detected the execution of a process normally associated with common Linux botnets.	Execution, Collection, Command and Control	Medium
Behavior similar to Fairware ransomware detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of rm -rf commands applied to suspicious locations. As rm -rf will recursively delete files, it is normally used on discrete folders. In this case, it is being used in a location that could remove a lot of data. Fairware ransomware is known to execute rm -rf commands in this folder. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Behavior similar to Fairware ransomware detected (VM_FairwareMalware)	Analysis of host data on % {Compromised Host} detected the execution of rm -rf commands applied to suspicious locations. As rm -rf will recursively delete files, it is normally used on discrete folders. In this case, it is being used in a location that could remove a lot of data. Fairware ransomware is known to execute rm -rf commands in this folder.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Behavior similar to ransomware detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of files that have resemblance of known ransomware that can prevent users from accessing their system or personal files, and demands ransom payment in order to regain access. This behavior was seen [x] times today on the following machines: [Machine names]	-	High
Communication with suspicious domain identified by threat intelligence (AzureDNS_ThreatIntelSuspectDomain)	Communication with suspicious domain was detected by analyzing DNS transactions from your resource and comparing against known malicious domains identified by threat intelligence feeds. Communication to malicious domains is frequently performed by attackers and could imply that your resource is compromised.	Initial Access, Persistence, Execution, Command And Control, Exploitation	Medium
Container with a miner image detected (VM_MinerInContainerImage)	Machine logs indicate execution of a Docker container that run an image associated with a digital currency mining.	Execution	High
Crypto coin miner execution (VM_CryptoCoinMinerExecution)	Analysis of host/device data detected a process being started in a way very similar to a coin mining process.	Execution	Medium
Custom script extension with suspicious command in your virtual machine (VM_CustomScriptExtensionSuspiciousCmd)	Custom script extension with suspicious command was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extension to execute a malicious code on your virtual machine via the Azure Resource Manager.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Custom script extension with suspicious entry-point in your virtual machine (VM_CustomScriptExtensionSuspiciousEntryPoint)	Custom script extension with a suspicious entry-point was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. The entry-point refers to a suspicious GitHub repository. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Custom script extension with suspicious payload in your virtual machine (VM_CustomScriptExtensionSuspiciousPayload)	Custom script extension with a payload from a suspicious GitHub repository was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Detected anomalous mix of upper and lower case characters in command line	Analysis of host data on %{Compromised Host} detected a command line with anomalous mix of upper and lower case characters. This kind of pattern, while possibly benign, is also typical of attackers trying to hide from case-sensitive or hash-based rule matching when performing administrative tasks on a compromised host.	-	Medium
Detected file download from a known malicious source [seen multiple times] (VM_SuspectDownload)	Analysis of host data has detected the download of a file from a known malware source on %{Compromised Host}. This behavior was seen over [x] times today on the following machines: [Machine names]	Privilege Escalation, Execution, Exfiltration, Command and Control	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected file download from a known malicious source	Analysis of host data has detected the download of a file from a known malware source on %{Compromised Host}.	-	Medium
Detected persistence attempt [seen multiple times]	Analysis of host data on %{Compromised Host} has detected installation of a startup script for single-user mode. It is extremely rare that any legitimate process needs to execute in that mode, so this may indicate that an attacker has added a malicious process to every run-level to guarantee persistence. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Detected persistence attempt (VM_NewSingleUserModeStartupScript)	Host data analysis has detected that a startup script for single-user mode has been installed. Because it's rare that any legitimate process would be required to run in that mode, this might indicate that an attacker has added a malicious process to every run-level to guarantee persistence.	Persistence	Medium
Detected suspicious file download [seen multiple times]	Analysis of host data has detected suspicious download of remote file on %{Compromised Host}. This behavior was seen 10 times today on the following machines: [Machine name]	-	Low
Detected suspicious file download (VM_SuspectDownloadArtifacts)	Analysis of host data has detected suspicious download of remote file on %{Compromised Host}.	Persistence	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected suspicious network activity	Analysis of network traffic from %{Compromised Host} detected suspicious network activity. Such traffic, while possibly benign, is typically used by an attacker to communicate with malicious servers for downloading of tools, command-and-control and exfiltration of data. Typical related attacker activity includes copying remote administration tools to a compromised host and exfiltrating user data from it.	-	Low
Detected suspicious use of the useradd command [seen multiple times]	Analysis of host data has detected suspicious use of the useradd command on %{Compromised Host}. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Detected suspicious use of the useradd command (VM_SuspectUserAddition)	Analysis of host data has detected suspicious use of the useradd command on %{Compromised Host}.	Persistence	Medium
Digital currency mining related behavior detected	Analysis of host data on % {Compromised Host} detected the execution of a process or command normally associated with digital currency mining.	-	High
Disabling of auditd logging [seen multiple times]	The Linux Audit system provides a way to track security-relevant information on the system. It records as much information about the events that are happening on your system as possible. Disabling auditd logging could hamper discovering violations of security policies used on the system. This behavior was seen [x] times today on the following machines: [Machine names]	-	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Docker build operation detected on a Kubernetes node (VM_ImageBuildOnNode)	Machine logs indicate a build operation of a container image on a Kubernetes node. While this behavior might be legitimate, attackers might build their malicious images locally to avoid detection.	Defense Evasion	Low
Executable found running from a suspicious location (VM_SuspectExecutablePath)	Analysis of host data detected an executable file on %{Compromised Host} that is running from a location in common with known suspicious files. This executable could either be legitimate activity, or an indication of a compromised host.	Execution	High
Exploitation of Xorg vulnerability [seen multiple times]	Analysis of host data on % {Compromised Host} detected the user of Xorg with suspicious arguments. Attackers may use this technique in privilege escalation attempts. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Exposed Docker daemon on TCP socket (VM_ExposedDocker)	Machine logs indicate that your Docker daemon (dockerd) exposes a TCP socket. By default, Docker configuration, does not use encryption or authentication when a TCP socket is enabled. This enables full access to the Docker daemon, by anyone with access to the relevant port.	Execution, Exploitation	Medium
Failed SSH brute force attack (VM_SshBruteForceFailed)	Failed brute force attacks were detected from the following attackers: % {Attackers}. Attackers were trying to access the host with the following user names: %{Accounts used on failed sign in to host attempts}.	Probing	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Fileless Attack Behavior Detected (VM_FilelessAttackBehavior.Linux)	The memory of the process specified below contains behaviors commonly used by fileless attacks. Specific behaviors include: {list of observed behaviors}	Execution	Low
Fileless Attack Technique Detected (VM_FilelessAttackTechnique.Linux)	The memory of the process specified below contains evidence of a fileless attack technique. Fileless attacks are used by attackers to execute code while evading detection by security software. Specific behaviors include: {list of observed behaviors}	Execution	High
Fileless Attack Toolkit Detected (VM_FilelessAttackToolkit.Linux)	The memory of the process specified below contains a fileless attack toolkit: {ToolKitName}. Fileless attack toolkits typically do not have a presence on the filesystem, making detection by traditional anti-virus software difficult. Specific behaviors include: {list of observed behaviors}	Defense Evasion, Execution	High
Hidden file execution detected	Analysis of host data indicates that a hidden file was executed by %{user name}. This activity could either be legitimate activity, or an indication of a compromised host.	-	Informational
Indicators associated with DDOS toolkit detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected file names that are part of a toolkit associated with malware capable of launching DDoS attacks, opening ports and services and taking full control over the infected system. This could also possibly be legitimate activity. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Indicators associated with DDOS toolkit detected (VM_KnownLinuxDDoSToolkit)	Analysis of host data on % {Compromised Host} detected file names that are part of a toolkit associated with malware capable of launching DDoS attacks, opening ports and services and taking full control over the infected system. This could also possibly be legitimate activity.	Persistence, Lateral Movement, Execution, Exploitation	Medium
Local host reconnaissance detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of a command normally associated with common Linux bot reconnaissance. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Local host reconnaissance detected (VM_LinuxReconnaissance)	Analysis of host data on % {Compromised Host} detected the execution of a command normally associated with common Linux bot reconnaissance.	Discovery	Medium
Manipulation of host firewall detected [seen multiple times] (VM_FirewallDisabled)	Analysis of host data on % {Compromised Host} detected possible manipulation of the on-host firewall. Attackers will often disable this to exfiltrate data. This behavior was seen [x] times today on the following machines: [Machine names]	Defense Evasion, Exfiltration	Medium
Manipulation of host firewall detected	Analysis of host data on % {Compromised Host} detected possible manipulation of the on-host firewall. Attackers will often disable this to exfiltrate data.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
MITRE Caldera agent detected (VM_MitreCalderaTools)	Machine logs indicate that the suspicious process: '% {Suspicious Process}' was running on % {Compromised Host}. This is often associated with the MITRE 54ndc47 agent which could be used maliciously to attack other machines in some way.	All	Medium
New SSH key added [seen multiple times] (VM_SshKeyAddition)	A new SSH key was added to the authorized keys file. This behavior was seen [x] times today on the following machines: [Machine names]	Persistence	Low
New SSH key added	A new SSH key was added to the authorized keys file	-	Low
Possible attack tool detected [seen multiple times]	Machine logs indicate that the suspicious process: '% {Suspicious Process}' was running on % {Compromised Host}. This tool is often associated with malicious users attacking other machines in some way. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Possible attack tool detected (VM_KnownLinuxAttackTool)	Machine logs indicate that the suspicious process: '% {Suspicious Process}' was running on % {Compromised Host}. This tool is often associated with malicious users attacking other machines in some way.	Execution, Collection, Command and Control, Probing	Medium
Possible backdoor detected [seen multiple times]	Analysis of host data has detected a suspicious file being downloaded then run on %{Compromised Host} in your subscription. This activity has previously been associated with installation of a backdoor. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible credential access tool detected [seen multiple times]	Machine logs indicate a possible known credential access tool was running on %{Compromised Host} launched by process: '%{Suspicious Process}'. This tool is often associated with attacker attempts to access credentials. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Possible credential access tool detected (VM_KnownLinuxCredentialAccessTool)	Machine logs indicate a possible known credential access tool was running on %{Compromised Host} launched by process: '%{Suspicious Process}'. This tool is often associated with attacker attempts to access credentials.	Credential Access	Medium
Possible data exfiltration [seen multiple times]	Analysis of host data on %{Compromised Host} detected a possible data egress condition. Attackers will often egress data from machines they have compromised. This behavior was seen [x]] times today on the following machines: [Machine names]	-	Medium
Possible data exfiltration (VM_DataEgressArtifacts)	Analysis of host data on %{Compromised Host} detected a possible data egress condition. Attackers will often egress data from machines they have compromised.	Collection, Exfiltration	Medium
Possible exploitation of Hadoop Yarn (VM_HadoopYarnExploit)	Analysis of host data on %{Compromised Host} detected the possible exploitation of the Hadoop Yarn service.	Exploitation	Medium
Possible exploitation of the mailserver detected (VM_MailserverExploitation)	Analysis of host data on %{Compromised Host} detected an unusual execution under the mail server account	Exploitation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible Log Tampering Activity Detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected possible removal of files that tracks user's activity during the course of its operation. Attackers often try to evade detection and leave no trace of malicious activities by deleting such log files. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Possible Log Tampering Activity Detected (VM_SystemLogRemoval)	Analysis of host data on % {Compromised Host} detected possible removal of files that tracks user's activity during the course of its operation. Attackers often try to evade detection and leave no trace of malicious activities by deleting such log files.	Defense Evasion	Medium
Possible malicious web shell detected [seen multiple times] (VM_Webshell)	Analysis of host data on % {Compromised Host} detected a possible web shell. Attackers will often upload a web shell to a machine they have compromised to gain persistence or for further exploitation. This behavior was seen [x] times today on the following machines: [Machine names]	Persistence, Exploitation	Medium
Possible malicious web shell detected	Analysis of host data on % {Compromised Host} detected a possible web shell. Attackers will often upload a web shell to a machine they have compromised to gain persistence or for further exploitation.	-	Medium
Possible password change using crypt-method detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected password change using crypt method. Attackers can make this change to continue access and gaining persistence after compromise. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Potential overriding of common files [seen multiple times]	Analysis of host data has detected common executables being overwritten on % {Compromised Host}. Attackers will overwrite common files as a way to obfuscate their actions or for persistence. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Potential overriding of common files (VM_OverridingCommonFiles)	Analysis of host data has detected common executables being overwritten on % {Compromised Host}. Attackers will overwrite common files as a way to obfuscate their actions or for persistence.	Persistence	Medium
Potential port forwarding to external IP address [seen multiple times]	Analysis of host data on % {Compromised Host} detected the initiation of port forwarding to an external IP address. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Potential port forwarding to external IP address (VM_SuspectPortForwarding)	Host data analysis detected the initiation of port forwarding to an external IP address.	Exfiltration, Command and Control	Medium
Potential reverse shell detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected a potential reverse shell. These are used to get a compromised machine to call back into a machine an attacker owns. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Potential reverse shell detected (VM_ReverseShell)	Analysis of host data on % {Compromised Host} detected a potential reverse shell. These are used to get a compromised machine to call back into a machine an attacker owns.	Exfiltration, Exploitation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Privileged command run in container (VM_PrivilegedExecutionInContainer)	Machine logs indicate that a privileged command was run in a Docker container. A privileged command has extended privileges on the host machine.	Privilege Escalation	Low
Privileged Container Detected (VM_PrivilegedContainerArtifacts)	Machine logs indicate that a privileged Docker container is running. A privileged container has a full access to the host's resources. If compromised, an attacker can use the privileged container to gain access to the host machine.	Privilege Escalation, Execution	Low
Process associated with digital currency mining detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of a process normally associated with digital currency mining. This behavior was seen over 100 times today on the following machines: [Machine name]	-	Medium
Process associated with digital currency mining detected	Host data analysis detected the execution of a process that is normally associated with digital currency mining.	Exploitation, Execution	Medium
Process seen accessing the SSH authorized keys file in an unusual way (VM_SshKeyAccess)	An SSH authorized keys file has been accessed in a method similar to known malware campaigns. This access can indicate that an attacker is attempting to gain persistent access to a machine.	-	Low
Python encoded downloader detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected the execution of encoded Python that downloads and runs code from a remote location. This may be an indication of malicious activity. This behavior was seen [x] times today on the following machines: [Machine names]	-	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Screenshot taken on host [seen multiple times]	Analysis of host data on % {Compromised Host} detected the user of a screen capture tool. Attackers may use these tools to access private data. This behavior was seen [x] times today on the following machines: [Machine names]	-	Low
Script extension mismatch detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected a mismatch between the script interpreter and the extension of the script file provided as input. This has frequently been associated with attacker script executions. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Script extension mismatch detected (VM_MismatchedScriptFeatures)	Analysis of host data on % {Compromised Host} detected a mismatch between the script interpreter and the extension of the script file provided as input. This has frequently been associated with attacker script executions.	Defense Evasion	Medium
Shellcode detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected shellcode being generated from the command line. This process could be legitimate activity, or an indication that one of your machines has been compromised. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
SSH server is running inside a container (VM_ContainerSSH)	Machine logs indicate that an SSH server is running inside a Docker container. While this behavior can be intentional, it frequently indicates that a container is misconfigured or breached.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Successful SSH brute force attack (VM_SshBruteForceSuccess)	Analysis of host data has detected a successful brute force attack. The IP %{Attacker source IP} was seen making multiple login attempts. Successful logins were made from that IP with the following user(s): %{Accounts used to successfully sign in to host}. This means that the host may be compromised and controlled by a malicious actor.	Exploitation	High
Suspect Password File Access (VM_SuspectPasswordFileAccess)	Analysis of host data has detected suspicious access to encrypted user passwords.	Persistence	Informational
Suspicious Account Creation Detected	Analysis of host data on %{Compromised Host} detected creation or use of a local account %{Suspicious account name} : this account name closely resembles a standard Windows account or group name '%{Similar To Account Name}'. This is potentially a rogue account created by an attacker, so named in order to avoid being noticed by a human administrator.	-	Medium
Suspicious compilation detected [seen multiple times]	Analysis of host data on %{Compromised Host} detected suspicious compilation. Attackers will often compile exploits on a machine they have compromised to escalate privileges. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Suspicious compilation detected (VM_SuspectCompilation)	Analysis of host data on %{Compromised Host} detected suspicious compilation. Attackers will often compile exploits on a machine they have compromised to escalate privileges.	Privilege Escalation, Exploitation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious DNS Over Https (VM_SuspiciousDNSOverHttps)	Analysis of host data indicates the use of a DNS call over HTTPS in an uncommon fashion. This technique is used by attackers to hide calls out to suspect or malicious sites.	DefenseEvasion, Exfiltration	Medium
Suspicious failed execution of custom script extension in your virtual machine (VM_CustomScriptExtensionSuspiciousFailure)	Suspicious failure of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Such failures may be associated with malicious scripts run by this extension.	Execution	Medium
Suspicious kernel module detected [seen multiple times]	Analysis of host data on % {Compromised Host} detected a shared object file being loaded as a kernel module. This could be legitimate activity, or an indication that one of your machines has been compromised. This behavior was seen [x] times today on the following machines: [Machine names]	-	Medium
Suspicious password access [seen multiple times]	Analysis of host data has detected suspicious access to encrypted user passwords on % {Compromised Host}. This behavior was seen [x] times today on the following machines: [Machine names]	-	Informational
Suspicious password access	Analysis of host data has detected suspicious access to encrypted user passwords on % {Compromised Host}.	-	Informational

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious PHP execution detected (VM_SuspectPhp)	Machine logs indicate that a suspicious PHP process is running. The action included an attempt to run OS commands or PHP code from the command line using the PHP process. While this behavior can be legitimate, in web applications this behavior is also observed in malicious activities such as attempts to infect websites with web shells.	Execution	Medium
Suspicious request to Kubernetes API (VM_KubernetesAPI)	Machine logs indicate that a suspicious request was made to the Kubernetes API. The request was sent from a Kubernetes node, possibly from one of the containers running in the node. Although this behavior can be intentional, it might indicate that the node is running a compromised container.	LateralMovement	Medium
Suspicious request to the Kubernetes Dashboard (VM_KubernetesDashboard)	Machine logs indicate that a suspicious request was made to the Kubernetes Dashboard. The request was sent from a Kubernetes node, possibly from one of the containers running in the node. Although this behavior can be intentional, it might indicate that the node is running a compromised container.	LateralMovement	Medium
Threat Intel Command Line Suspect Domain (VM_ThreatIntelCommandLineSuspectDomain)	The process 'PROCESSNAME' on 'HOST' connected to a location that has been reported to be malicious or unusual. This is an indicator that a compromise may have occurred.	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual config reset in your virtual machine (VM_VMAccessUnusualConfigReset)	An unusual config reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing VM Access extension to reset the configuration in your virtual machine and compromise it.	Credential Access	Medium
Unusual deletion of custom script extension in your virtual machine (VM_CustomScriptExtensionUnusualDeletion)	Unusual deletion of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Unusual execution of custom script extension in your virtual machine (VM_CustomScriptExtensionUnusualExecution)	Unusual execution of a custom script extension was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. Attackers may use custom script extensions to execute malicious code on your virtual machines via the Azure Resource Manager.	Execution	Medium
Unusual user password reset in your virtual machine (VM_VMAccessUnusualPasswordReset)	An unusual user password reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing the VM Access extension to reset the credentials of a local user in your virtual machine and compromise it.	Credential Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual user SSH key reset in your virtual machine (VM_VMAccessUnusualSSHReset)	An unusual user SSH key reset was detected in your virtual machine by analyzing the Azure Resource Manager operations in your subscription. While this action may be legitimate, attackers can try utilizing VM Access extension to reset SSH key of a user account in your virtual machine and compromise it.	Credential Access	Medium

Alerts for Azure App Service

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
An attempt to run Linux commands on a Windows App Service (AppServices_LinuxCommandOnWindows)	Analysis of App Service processes detected an attempt to run a Linux command on a Windows App Service. This action was running by the web application. This behavior is often seen during campaigns that exploit a vulnerability in a common web application. (Applies to: App Service on Windows)	-	Medium
An IP that connected to your Azure App Service FTP Interface was found in Threat Intelligence (AppServices_IncomingTiClientIpFtp)	Azure App Service FTP log indicates a connection from a source address that was found in the threat intelligence feed. During this connection, a user accessed the pages listed. (Applies to: App Service on Windows and App Service on Linux)	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Attempt to run high privilege command detected (AppServices_HighPrivilegeCommand)	Analysis of App Service processes detected an attempt to run a command that requires high privileges. The command ran in the web application context. While this behavior can be legitimate, in web applications this behavior is also observed in malicious activities. (Applies to: App Service on Windows)	-	Medium
Communication with suspicious domain identified by threat intelligence (AzureDNS_ThreatIntelSuspectDomain)	Communication with suspicious domain was detected by analyzing DNS transactions from your resource and comparing against known malicious domains identified by threat intelligence feeds. Communication to malicious domains is frequently performed by attackers and could imply that your resource is compromised.	Initial Access, Persistence, Execution, Command And Control, Exploitation	Medium
Connection to web page from anomalous IP address detected (AppServices_AnomalousPageAccess)	Azure App Service activity log indicates an anomalous connection to a sensitive web page from the listed source IP address. This might indicate that someone is attempting a brute force attack into your web app administration pages. It might also be the result of a new IP address being used by a legitimate user. If the source IP address is trusted, you can safely suppress this alert for this resource. To learn how to suppress security alerts, see Suppress alerts from Microsoft Defender for Cloud. (Applies to: App Service on Windows and App Service on Linux)	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Dangling DNS record for an App Service resource detected (AppServices_DanglingDomain)	A DNS record that points to a recently deleted App Service resource (also known as "dangling DNS" entry) has been detected. This leaves you susceptible to a subdomain takeover. Subdomain takeovers enable malicious actors to redirect traffic intended for an organization's domain to a site performing malicious activity. (Applies to: App Service on Windows and App Service on Linux)	-	High
Detected encoded executable in command line data (AppServices_Base64EncodedExecutableInCommandLineParams)	Analysis of host data on {Compromised host} detected a base-64 encoded executable. This has previously been associated with attackers attempting to construct executables on-the-fly through a sequence of commands, and attempting to evade intrusion detection systems by ensuring that no individual command would trigger an alert. This could be legitimate activity, or an indication of a compromised host. (Applies to: App Service on Windows)	Defense Evasion, Execution	High
Detected file download from a known malicious source (AppServices_SuspectDownload)	Analysis of host data has detected the download of a file from a known malware source on your host. (Applies to: App Service on Linux)	Privilege Escalation, Execution, Exfiltration, Command and Control	Medium
Detected suspicious file download (AppServices_SuspectDownloadArtifacts)	Analysis of host data has detected suspicious download of remote file. (Applies to: App Service on Linux)	Persistence	Medium
Digital currency mining related behavior detected (AppServices_DigitalCurrencyMining)	Analysis of host data on Inn-Flow-WebJobs detected the execution of a process or command normally associated with digital currency mining. (Applies to: App Service on Windows and App Service on Linux)	Execution	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Executable decoded using certutil (AppServices_ExecutableDecodedUsingCertutil)	Analysis of host data on [Compromised entity] detected that certutil.exe, a built-in administrator utility, was being used to decode an executable instead of its mainstream purpose that relates to manipulating certificates and certificate data. Attackers are known to abuse functionality of legitimate administrator tools to perform malicious actions, for example using a tool such as certutil.exe to decode a malicious executable that will then be subsequently executed. (Applies to: App Service on Windows)	Defense Evasion, Execution	High
Fileless Attack Behavior Detected (AppServices_FilelessAttackBehaviorDetection)	The memory of the process specified below contains behaviors commonly used by fileless attacks. Specific behaviors include: {list of observed behaviors} (Applies to: App Service on Windows and App Service on Linux)	Execution	Medium
Fileless Attack Technique Detected (AppServices_FilelessAttackTechniqueDetection)	The memory of the process specified below contains evidence of a fileless attack technique. Fileless attacks are used by attackers to execute code while evading detection by security software. Specific behaviors include: {list of observed behaviors} (Applies to: App Service on Windows and App Service on Linux)	Execution	High
Fileless Attack Toolkit Detected (AppServices_FilelessAttackToolkitDetection)	The memory of the process specified below contains a fileless attack toolkit: {ToolKitName}. Fileless attack toolkits typically do not have a presence on the filesystem, making detection by traditional anti-virus software difficult. Specific behaviors include: {list of observed behaviors} (Applies to: App Service on Windows and App Service on Linux)	Defense Evasion, Execution	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Microsoft Defender for Cloud test alert for App Service (not a threat) (AppServices_EICAR)	This is a test alert generated by Microsoft Defender for Cloud. No further action is needed. (Applies to: App Service on Windows and App Service on Linux)	-	High
NMap scanning detected (AppServices_Nmap)	Azure App Service activity log indicates a possible web fingerprinting activity on your App Service resource. The suspicious activity detected is associated with NMAP. Attackers often use this tool for probing the web application to find vulnerabilities. (Applies to: App Service on Windows and App Service on Linux)	PreAttack	Medium
Phishing content hosted on Azure Webapps (AppServices_PhishingContent)	URL used for phishing attack found on the Azure AppServices website. This URL was part of a phishing attack sent to Microsoft 365 customers. The content typically lures visitors into entering their corporate credentials or financial information into a legitimate looking website. (Applies to: App Service on Windows and App Service on Linux)	Collection	High
PHP file in upload folder (AppServices_PhpInUploadFolder)	Azure App Service activity log indicates an access to a suspicious PHP page located in the upload folder. This type of folder does not usually contain PHP files. The existence of this type of file might indicate an exploitation taking advantage of arbitrary file upload vulnerabilities. (Applies to: App Service on Windows and App Service on Linux)	Execution	Medium
Possible Cryptocoinminer download detected (AppServices_CryptoCoinMinerDownload)	Analysis of host data has detected the download of a file normally associated with digital currency mining. (Applies to: App Service on Linux)	Defense Evasion, Command and Control, Exploitation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible data exfiltration detected (AppServices_DataEgressArtifacts)	Analysis of host/device data detected a possible data egress condition. Attackers will often egress data from machines they have compromised. (Applies to: App Service on Linux)	Collection, Exfiltration	Medium
Potential dangling DNS record for an App Service resource detected (AppServices_PotentialDanglingDomain)	A DNS record that points to a recently deleted App Service resource (also known as "dangling DNS" entry) has been detected. This might leave you susceptible to a subdomain takeover. Subdomain takeovers enable malicious actors to redirect traffic intended for an organization's domain to a site performing malicious activity. In this case, a text record with the Domain Verification ID was found. Such text records prevent subdomain takeover but we still recommend removing the dangling domain. If you leave the DNS record pointing at the subdomain you're at risk if anyone in your organization deletes the TXT file or record in the future. (Applies to: App Service on Windows and App Service on Linux)	-	Low
Potential reverse shell detected (AppServices_ReverseShell)	Analysis of host data detected a potential reverse shell. These are used to get a compromised machine to call back into a machine an attacker owns. (Applies to: App Service on Linux)	Exfiltration, Exploitation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Raw data download detected (AppServices_DownloadCodeFromWebsite)	Analysis of App Service processes detected an attempt to download code from raw-data websites such as Pastebin. This action was run by a PHP process. This behavior is associated with attempts to download web shells or other malicious components to the App Service. (Applies to: App Service on Windows)	Execution	Medium
Saving curl output to disk detected (AppServices_CurlToDisk)	Analysis of App Service processes detected the running of a curl command in which the output was saved to the disk. While this behavior can be legitimate, in web applications this behavior is also observed in malicious activities such as attempts to infect websites with web shells. (Applies to: App Service on Windows)	-	Low
Spam folder referrer detected (AppServices_SpamReferrer)	Azure App Service activity log indicates web activity that was identified as originating from a web site associated with spam activity. This can occur if your website is compromised and used for spam activity. (Applies to: App Service on Windows and App Service on Linux)	-	Low
Suspicious access to possibly vulnerable web page detected (AppServices_ScanSensitivePage)	Azure App Service activity log indicates a web page that seems to be sensitive was accessed. This suspicious activity originated from a source IP address whose access pattern resembles that of a web scanner. This activity is often associated with an attempt by an attacker to scan your network to try and gain access to sensitive or vulnerable web pages. (Applies to: App Service on Windows and App Service on Linux)	-	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious domain name reference (AppServices_CommandlineSuspectDomain)	Analysis of host data detected reference to suspicious domain name. Such activity, while possibly legitimate user behavior, is frequently an indication of the download or execution of malicious software. Typical related attacker activity is likely to include the download and execution of further malicious software or remote administration tools. (Applies to: App Service on Linux)	Exfiltration	Low
Suspicious download using Certutil detected (AppServices_DownloadUsingCertutil)	Analysis of host data on {NAME} detected the use of certutil.exe, a built-in administrator utility, for the download of a binary instead of its mainstream purpose that relates to manipulating certificates and certificate data. Attackers are known to abuse functionality of legitimate administrator tools to perform malicious actions, for example using certutil.exe to download and decode a malicious executable that will then be subsequently executed. (Applies to: App Service on Windows)	Execution	Medium
Suspicious PHP execution detected (AppServices_SuspectPhp)	Machine logs indicate that a suspicious PHP process is running. The action included an attempt to run operating system commands or PHP code from the command line, by using the PHP process. While this behavior can be legitimate, in web applications this behavior might indicate malicious activities, such as attempts to infect websites with web shells. (Applies to: App Service on Windows and App Service on Linux)	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious PowerShell cmdlets executed (AppServices_PowerShellPowerSploitScriptExecution)	Analysis of host data indicates execution of known malicious PowerShell PowerSploit cmdlets. (Applies to: App Service on Windows)	Execution	Medium
Suspicious process executed (AppServices_KnownCredential AccessTools)	Machine logs indicate that the suspicious process: '% {process path}' was running on the machine, often associated with attacker attempts to access credentials. (Applies to: App Service on Windows)	Credential Access	High
Suspicious process name detected (AppServices_ProcessWithKnownSuspiciousExtension)	Analysis of host data on {NAME} detected a process whose name is suspicious, for example corresponding to a known attacker tool or named in a way that is suggestive of attacker tools that try to hide in plain sight. This process could be legitimate activity, or an indication that one of your machines has been compromised. (Applies to: App Service on Windows)	Persistence, Defense Evasion	Medium
Suspicious SVCHOST process executed (AppServices_SVCHostFromInvalidPath)	The system process SVCHOST was observed running in an abnormal context. Malware often use SVCHOST to mask its malicious activity. (Applies to: App Service on Windows)	Defense Evasion, Execution	High
Suspicious User Agent detected (AppServices_UserAgentInjection)	Azure App Service activity log indicates requests with suspicious user agent. This behavior can indicate on attempts to exploit a vulnerability in your App Service application. (Applies to: App Service on Windows and App Service on Linux)	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious WordPress theme invocation detected (AppServices_WpThemeInjection)	Azure App Service activity log indicates a possible code injection activity on your App Service resource. The suspicious activity detected resembles that of a manipulation of WordPress theme to support server side execution of code, followed by a direct web request to invoke the manipulated theme file. This type of activity was seen in the past as part of an attack campaign over WordPress. If your App Service resource isn't hosting a WordPress site, it isn't vulnerable to this specific code injection exploit and you can safely suppress this alert for the resource. To learn how to suppress security alerts, see Suppress alerts from Microsoft Defender for Cloud. (Applies to: App Service on Windows and App Service on Linux)	Execution	High
Vulnerability scanner detected (AppServices_DrupalScanner)	Azure App Service activity log indicates that a possible vulnerability scanner was used on your App Service resource. The suspicious activity detected resembles that of tools targeting a content management system (CMS). If your App Service resource isn't hosting a Drupal site, it isn't vulnerable to this specific code injection exploit and you can safely suppress this alert for the resource. To learn how to suppress security alerts, see Suppress alerts from Microsoft Defender for Cloud. (Applies to: App Service on Windows)	PreAttack	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Vulnerability scanner detected (AppServices_JoomlaScanner)	Azure App Service activity log indicates that a possible vulnerability scanner was used on your App Service resource. The suspicious activity detected resembles that of tools targeting Joomla applications. If your App Service resource isn't hosting a Joomla site, it isn't vulnerable to this specific code injection exploit and you can safely suppress this alert for the resource. To learn how to suppress security alerts, see Suppress alerts from Microsoft Defender for Cloud. (Applies to: App Service on Windows and App Service on Linux)	PreAttack	Medium
Vulnerability scanner detected (AppServices_WpScanner)	Azure App Service activity log indicates that a possible vulnerability scanner was used on your App Service resource. The suspicious activity detected resembles that of tools targeting WordPress applications. If your App Service resource isn't hosting a WordPress site, it isn't vulnerable to this specific code injection exploit and you can safely suppress this alert for the resource. To learn how to suppress security alerts, see Suppress alerts from Microsoft Defender for Cloud. (Applies to: App Service on Windows and App Service on Linux)	PreAttack	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Web fingerprinting detected (AppServices_WebFingerprinting)	Azure App Service activity log indicates a possible web fingerprinting activity on your App Service resource. The suspicious activity detected is associated with a tool called Blind Elephant. The tool fingerprint web servers and tries to detect the installed applications and version. Attackers often use this tool for probing the web application to find vulnerabilities. (Applies to: App Service on Windows and App Service on Linux)	PreAttack	Medium
Website is tagged as malicious in threat intelligence feed (AppServices_SmartScreen)	Your website as described below is marked as a malicious site by Windows SmartScreen. If you think this is a false positive, contact Windows SmartScreen via report feedback link provided. (Applies to: App Service on Windows and App Service on Linux)	Collection	Medium

Alerts for containers - Kubernetes clusters

Microsoft Defender for Containers provides security alerts on the cluster level and on the underlying cluster nodes by monitoring both control plane (API server) and the containerized workload itself. Control plane security alerts can be recognized by a prefix of `K8S_` of the alert type. Security alerts for runtime workload in the clusters can be recognized by the `K8S.NODE_` prefix of the alert type.

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
A file was downloaded and executed (Preview) (K8S.NODE_LinuxSuspiciousActivity)	Analysis of processes running within a container indicates that a file has been downloaded to the container, given execution privileges and then executed.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
A history file has been cleared (Preview) (K8S.NODE_HistoryFileCleared)	Analysis of processes running within a container indicates that the command history log file has been cleared. Attackers may do this to cover their tracks. The operation was performed by the specified user account.	DefenseEvasion	Medium
An uncommon connection attempt detected (Preview) (K8S.NODE_SuspectConnection)	Analysis of processes running within a container detected an uncommon connection attempt utilizing a socks protocol. This is very rare in normal operations, but a known technique for attackers attempting to bypass network-layer detections.	Execution, Exfiltration, Exploitation	Medium
Anomalous pod deployment (Preview) (K8S_AnomalousPodDeployment)	Kubernetes audit log analysis detected pod deployment which is anomalous based on previous pod deployment activity. This activity is considered an anomaly when taking into account how the different features seen in the deployment operation are in relations to one another. The features monitored include the container image registry used, the account performing the deployment, day of the week, how often this account performs pod deployments, user agent used in the operation, whether this is a namespace to which pod deployments often occur, and other features. Top contributing reasons for raising this alert as anomalous activity are detailed under the alert's extended properties.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Attempt to stop apt-daily-upgrade.timer service detected (Preview) (K8S.NODE_TimerServiceDisabled)	Analysis of host/device data detected an attempt to stop apt-daily-upgrade.timer service. Attackers have been observed stopping this service to download malicious files and grant execution privileges for their attacks. This activity can also happen if the service is updated through normal administrative actions.	DefenseEvasion	Informational
Behavior similar to common Linux bots detected (Preview) (K8S.NODE_CommonBot)	Analysis of processes running within a container detected execution of a process normally associated with common Linux botnets.	Execution, Collection, Command And Control	Medium
Behavior similar to Fairware ransomware detected (Preview) (K8S.NODE_FairwareMalware)	Analysis of processes running within a container detected the execution of rm -rf commands applied to suspicious locations. As rm -rf will recursively delete files, it is normally used on discrete folders. In this case, it is being used in a location that could remove a lot of data. Fairware ransomware is known to execute rm -rf commands in this folder.	Execution	Medium
Command within a container running with high privileges (Preview) (K8S.NODE_PrivilegedExecutionInContainer)	Machine logs indicate that a privileged command was run in a Docker container. A privileged command has extended privileges on the host machine.	PrivilegeEscalation	Low
Container running in privileged mode (Preview) (K8S.NODE_PrivilegedContainerArtifacts)	Machine logs indicate that a privileged Docker container is running. A privileged container has full access to the host's resources. If compromised, an attacker can use the privileged container to gain access to the host machine.	PrivilegeEscalation, Execution	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Container with a sensitive volume mount detected (K8S_SensitiveMount)	Kubernetes audit log analysis detected a new container with a sensitive volume mount. The volume that was detected is a hostPath type which mounts a sensitive file or folder from the node to the container. If the container gets compromised, the attacker can use this mount for gaining access to the node.	Privilege Escalation	Medium
CoreDNS modification in Kubernetes detected (K8S_CoreDnsModification)	Kubernetes audit log analysis detected a modification of the CoreDNS configuration. The configuration of CoreDNS can be modified by overriding its configmap. While this activity can be legitimate, if attackers have permissions to modify the configmap, they can change the behavior of the cluster's DNS server and poison it.	Lateral Movement	Low
Creation of admission webhook configuration detected (K8S_AdmissionController)	Kubernetes audit log analysis detected a new admission webhook configuration. Kubernetes has two built-in generic admission controllers: MutatingAdmissionWebhook and ValidatingAdmissionWebhook. The behavior of these admission controllers is determined by an admission webhook that the user deploys to the cluster. The usage of such admission controllers can be legitimate, however attackers can use such webhooks for modifying the requests (in case of MutatingAdmissionWebhook) or inspecting the requests and gain sensitive information (in case of ValidatingAdmissionWebhook).	Credential Access, Persistence	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Detected file download from a known malicious source (Preview) (K8S.NODE_SuspectDownload)	Analysis of processes running within a container detected download of a file from a source frequently used to distribute malware.	PrivilegeEscalation, Execution, Exfiltration, Command And Control	Medium
Detected Persistence Attempt (Preview) (K8S.NODE_NewSingleUserModeStartupScript)	Analysis of processes running within a container detected installation of a startup script for single-user mode. It is extremely rare that any legitimate process needs to execute in that mode so it may indicate an attacker has added a malicious process to every run-level to guarantee persistence.	Persistence	Medium
Detected suspicious file download (Preview) (K8S.NODE_SuspectDownloadArtifacts)	Analysis of processes running within a container detected suspicious download of a remote file.	Persistence	Low
Detected suspicious use of the nohup command (Preview) (K8S.NODE_SuspectNohup)	Analysis of processes running within a container detected suspicious use of the nohup command. Attackers have been seen using the command nohup to run hidden files from a temporary directory to allow their executables to run in the background. It is rare to see this command run on hidden files located in a temporary directory.	Persistence, DefenseEvasion	Medium
Detected suspicious use of the useradd command (Preview) (K8S.NODE_SuspectUserAddition)	Analysis of processes running within a container detected suspicious use of the useradd command.	Persistence	Medium
Digital currency mining container detected (K8S_MaliciousContainerImage)	Kubernetes audit log analysis detected a container that has an image associated with a digital currency mining tool.	Execution	High
Digital currency mining related behavior detected (Preview) (K8S.NODE_DigitalCurrencyMining)	Analysis of host data detected the execution of a process or command normally associated with digital currency mining.	Execution	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Docker build operation detected on a Kubernetes node (Preview) (K8S.NODE_ImageBuildOnNode)	Analysis of processes running within a container indicates a build operation of a container image on a Kubernetes node. While this behavior might be legitimate, attackers might build their malicious images locally to avoid detection.	DefenseEvasion	Low
Excessive role permissions assigned in Kubernetes cluster (Preview) (K8S_ServiceAcountPermissionAnomaly)	Analysis of the Kubernetes audit logs detected an excessive permissions role assignment to your cluster. The listed permissions for the assigned roles are uncommon to the specific service account. This detection considers previous role assignments to the same service account across clusters monitored by Azure, volume per permission, and the impact of the specific permission. The anomaly detection model used for this alert takes into account how this permission is used across all clusters monitored by Microsoft Defender for Cloud.	Privilege Escalation	Low
Executable found running from a suspicious location (Preview) (K8S.NODE_SuspectExecutablePath)	Analysis of host data detected an executable file that is running from a location associated with known suspicious files. This executable could either be legitimate activity, or an indication of a compromised host.	Execution	Medium
Execution of hidden file (Preview) (K8S.NODE_ExecuteHiddenFile)	Analysis of host data indicates that a hidden file was executed by the specified user account.	Persistence, DefenseEvasion	Informational

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Exposed Docker daemon on TCP socket (Preview) (K8S.NODE_ExposedDocker)	Machine logs indicate that your Docker daemon (dockerd) exposes a TCP socket. By default, Docker configuration, does not use encryption or authentication when a TCP socket is enabled. This enables full access to the Docker daemon, by anyone with access to the relevant port.	Execution, Exploitation	Medium
Exposed Kubeflow dashboard detected (K8S_ExposedKubeflow)	The Kubernetes audit log analysis detected exposure of the Istio Ingress by a load balancer in a cluster that runs Kubeflow. This action might expose the Kubeflow dashboard to the internet. If the dashboard is exposed to the internet, attackers can access it and run malicious containers or code on the cluster. Find more details in the following article: https://aka.ms/exposedkubeflow-blog	Initial Access	Medium
Exposed Kubernetes dashboard detected (K8S_ExposedDashboard)	Kubernetes audit log analysis detected exposure of the Kubernetes Dashboard by a LoadBalancer service. Exposed dashboard allows an unauthenticated access to the cluster management and poses a security threat.	Initial Access	High
Exposed Kubernetes service detected (K8S_ExposedService)	The Kubernetes audit log analysis detected exposure of a service by a load balancer. This service is related to a sensitive application that allows high impact operations in the cluster such as running processes on the node or creating new containers. In some cases, this service doesn't require authentication. If the service doesn't require authentication, exposing it to the internet poses a security risk.	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Exposed Redis service in AKS detected (K8S_ExposedRedis)	The Kubernetes audit log analysis detected exposure of a Redis service by a load balancer. If the service doesn't require authentication, exposing it to the internet poses a security risk.	Initial Access	Low
Indicators associated with DDOS toolkit detected (Preview) (K8S.NODE_KnownLinuxDDoSToolkit)	Analysis of processes running within a container detected file names that are part of a toolkit associated with malware capable of launching DDoS attacks, opening ports and services, and taking full control over the infected system. This could also possibly be legitimate activity.	Persistence, LateralMovement, Execution, Exploitation	Medium
K8S API requests from proxy IP address detected (K8S_TI_Proxy)	Kubernetes audit log analysis detected API requests to your cluster from an IP address that is associated with proxy services, such as TOR. While this behavior can be legitimate, it's often seen in malicious activities, when attackers try to hide their source IP.	Execution	Low
Kubernetes events deleted (K8S_DeleteEvents)	Defender for Cloud detected that some Kubernetes events have been deleted. Kubernetes events are objects in Kubernetes which contain information about changes in the cluster. Attackers might delete those events for hiding their operations in the cluster.	Defense Evasion	Medium
Kubernetes penetration testing tool detected (K8S_PenTestToolsKubeHunter)	Kubernetes audit log analysis detected usage of Kubernetes penetration testing tool in the AKS cluster. While this behavior can be legitimate, attackers might use such public tools for malicious purposes.	Execution	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Local host reconnaissance detected (Preview) (K8S.NODE_LinuxReconnaissance)	Analysis of processes running within a container detected the execution of a command normally associated with common Linux bot reconnaissance.	Discovery	Medium
Manipulation of host firewall detected (Preview) (K8S.NODE_FirewallDisabled)	Analysis of processes running within a container detected possible manipulation of the on-host firewall. Attackers will often disable this to exfiltrate data.	DefenseEvasion, Exfiltration	Medium
Microsoft Defender for Cloud test alert (not a threat). (Preview) (K8S.NODE_EICAR)	This is a test alert generated by Microsoft Defender for Cloud. No further action is needed.	Execution	High
MITRE Caldera agent detected (Preview) (K8S.NODE_MitreCalderaTools)	Analysis of processes running within a container indicate that a suspicious process was running. This is often associated with the MITRE 54ndc47 agent which could be used maliciously to attack other machines.	Persistence, PrivilegeEscalation, DefenseEvasion, CredentialAccess, Discovery, LateralMovement, Execution, Collection, Exfiltration, Command And Control, Probing, Exploitation	Medium
New container in the kube-system namespace detected (K8S_KubeSystemContainer)	Kubernetes audit log analysis detected a new container in the kube-system namespace that isn't among the containers that normally run in this namespace. The kube-system namespaces should not contain user resources. Attackers can use this namespace for hiding malicious components.	Persistence	Low
New high privileges role detected (K8S_HighPrivilegesRole)	Kubernetes audit log analysis detected a new role with high privileges. A binding to a role with high privileges gives the user\group high privileges in the cluster. Unnecessary privileges might cause privilege escalation in the cluster.	Persistence	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible attack tool detected (Preview) (K8S.NODE_KnownLinuxAttackTool)	Analysis of processes running within a container indicates a suspicious tool ran. This tool is often associated with malicious users attacking others.	Execution, Collection, Command And Control, Probing	Medium
Possible backdoor detected (Preview) (K8S.NODE_LinuxBackdoorArtifact)	Analysis of processes running within a container detected a suspicious file being downloaded and run. This activity has previously been associated with installation of a backdoor.	Persistence, DefenseEvasion, Execution, Exploitation	Medium
Possible command line exploitation attempt (Preview) (K8S.NODE_ExploitAttempt)	Analysis of processes running within a container detected a possible exploitation attempt against a known vulnerability.	Exploitation	Medium
Possible credential access tool detected (Preview) (K8S.NODE_KnownLinuxCredentialAccessTool)	Analysis of processes running within a container indicates a possible known credential access tool was running on the container, as identified by the specified process and commandline history item. This tool is often associated with attacker attempts to access credentials.	CredentialAccess	Medium
Possible Cryptocoinminer download detected (Preview) (K8S.NODE_CryptoCoinMinerDownload)	Analysis of processes running within a container detected the download of a file normally associated with digital currency mining.	DefenseEvasion, Command And Control, Exploitation	Medium
Possible data exfiltration detected (Preview) (K8S.NODE_DataEgressArtifacts)	Analysis of host/device data detected a possible data egress condition. Attackers will often egress data from machines they have compromised.	Collection, Exfiltration	Medium
Possible Log Tampering Activity Detected (Preview) (K8S.NODE_SystemLogRemoval)	Analysis of processes running within a container detected possible removal of files that tracks user's activity during the course of its operation. Attackers often try to evade detection and leave no trace of malicious activities by deleting such log files.	DefenseEvasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible password change using crypt-method detected (Preview) (K8S.NODE_SuspectPasswordChange)	Analysis of processes running within a container detected a password change using the crypt method. Attackers can make this change to continue access and gain persistence after compromise.	CredentialAccess	Medium
Potential overriding of common files (Preview) (K8S.NODE_OverridingCommonFiles)	Analysis of processes running within a container detected common files as a way to obfuscate their actions or for persistence.	Persistence	Medium
Potential port forwarding to external IP address (Preview) (K8S.NODE_SuspectPortForwarding)	Analysis of processes running within a container detected the initiation of port forwarding to an external IP address.	Exfiltration, Command And Control	Medium
Potential reverse shell detected (Preview) (K8S.NODE_ReverseShell)	Analysis of processes running within a container detected a potential reverse shell. These are used to get a compromised machine to call back into a machine an attacker owns.	Exfiltration, Exploitation	Medium
Privileged container detected (K8S_PrivilegedContainer)	Kubernetes audit log analysis detected a new privileged container. A privileged container has access to the node's resources and breaks the isolation between containers. If compromised, an attacker can use the privileged container to gain access to the node.	Privilege Escalation	Low
Process associated with digital currency mining detected (Preview) (K8S.NODE_CryptoCoinMinerArtifacts)	Analysis of processes running within a container detected the execution of a process normally associated with digital currency mining.	Execution, Exploitation	Medium
Process seen accessing the SSH authorized keys file in an unusual way (Preview) (K8S.NODE_SshKeyAccess)	An SSH authorized_keys file was accessed in a method similar to known malware campaigns. This access could signify that an actor is attempting to gain persistent access to a machine.	Unknown	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Role binding to the cluster-admin role detected (K8S_ClusterAdminBinding)	Kubernetes audit log analysis detected a new binding to the cluster-admin role which gives administrator privileges. Unnecessary administrator privileges might cause privilege escalation in the cluster.	Persistence	Low
Screenshot taken on host (Preview) (K8S.NODE_KnownLinuxScreenshotTool)	Analysis of host/device data detected the use of a screen capture tool. Attackers may use these tools to access private data.	Collection	Low
Script extension mismatch detected (Preview) (K8S.NODE_MismatchedScriptFeatures)	Analysis of processes running within a container detected a mismatch between the script interpreter and the extension of the script file provided as input. This has frequently been associated with attacker script executions.	DefenseEvasion	Medium
Security-related process termination detected (Preview) (K8S.NODE_SuspectProcessTermination)	Analysis of processes running within a container detected attempt to terminate processes related to security monitoring on the container. Attackers will often try to terminate such processes using predefined scripts post-compromise.	Persistence	Low
SSH server is running inside a container (Preview) (Preview) (K8S.NODE_ContainerSSH)	Analysis of processes running within a container detected an SSH server running inside the container.	Execution	Medium
Suspicious compilation detected (Preview) (K8S.NODE_SuspectCompilation)	Analysis of processes running within a container detected suspicious compilation. Attackers will often compile exploits to escalate privileges.	PrivilegeEscalation, Exploitation	Medium
Suspicious file timestamp modification (Preview) (K8S.NODE_TimestampTampering)	Analysis of host/device data detected a suspicious timestamp modification. Attackers will often copy timestamps from existing legitimate files to new tools to avoid detection of these newly dropped files.	Persistence, DefenseEvasion	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious request to Kubernetes API (Preview) (K8S.NODE_KubernetesAPI)	Analysis of processes running within a container indicates that a suspicious request was made to the Kubernetes API. The request was sent from a container in the cluster. Although this behavior can be intentional, it might indicate that a compromised container is running in the cluster.	LateralMovement	Medium
Suspicious request to the Kubernetes Dashboard (Preview) (K8S.NODE_KubernetesDashboard)	Analysis of processes running within a container indicates that a suspicious request was made to the Kubernetes Dashboard. The request was sent from a container in the cluster. Although this behavior can be intentional, it might indicate that a compromised container is running in the cluster.	Execution	Medium
Potential crypto coin miner started (Preview) (K8S.NODE_CryptoCoinMinerExecution)	Analysis of processes running within a container detected a process being started in a way normally associated with digital currency mining.	Execution	Medium
Suspicious password access (Preview) (K8S.NODE_SuspectPasswordFileAccess)	Analysis of processes running within a container detected suspicious access to encrypted user passwords.	Persistence	Informational
Suspicious use of DNS over HTTPS (Preview) (K8S.NODE_SuspiciousDNSOverHttps)	Analysis of processes running within a container indicates the use of a DNS call over HTTPS in an uncommon fashion. This technique is used by attackers to hide calls out to suspect or malicious sites.	DefenseEvasion, Exfiltration	Medium
A possible connection to malicious location has been detected. (Preview) (K8S.NODE_ThreatIntelCommandLineSuspectDomain)	Analysis of processes running within a container detected a connection to a location that has been reported to be malicious or unusual. This is an indicator that a compromise may have occured.	InitialAccess	Medium

Alerts for SQL Database and Azure Synapse Analytics

Further details and notes

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
A possible vulnerability to SQL Injection (SQL.VM_VulnerabilityToSqlInjection SQL.DB_VulnerabilityToSqlInjection SQL.MI_VulnerabilityToSqlInjection SQL.DW_VulnerabilityToSqlInjection)	An application has generated a faulty SQL statement in the database. This can indicate a possible vulnerability to SQL injection attacks. There are two possible reasons for a faulty statement. A defect in application code might have constructed the faulty SQL statement. Or, application code or stored procedures didn't sanitize user input when constructing the faulty SQL statement, which can be exploited for SQL injection.	PreAttack	Medium
Attempted logon by a potentially harmful application (SQL.DB_HarmfulApplication SQL.VM_HarmfulApplication SQL.MI_HarmfulApplication SQL.DW_HarmfulApplication)	A potentially harmful application attempted to access SQL server '{name}'.	PreAttack	High
Log on from an unusual Azure Data Center (SQL.DB_DataCenterAnomaly SQL.VM_DataCenterAnomaly SQL.DW_DataCenterAnomaly SQL.MI_DataCenterAnomaly)	There has been a change in the access pattern to an SQL Server, where someone has signed in to the server from an unusual Azure Data Center. In some cases, the alert detects a legitimate action (a new application or Azure service). In other cases, the alert detects a malicious action (attacker operating from breached resource in Azure).	Probing	Low

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Log on from an unusual location (SQL.DB_GeoAnomaly SQL.VM_GeoAnomaly SQL.DW_GeoAnomaly SQL.MI_GeoAnomaly)	There has been a change in the access pattern to SQL Server, where someone has signed in to the server from an unusual geographical location. In some cases, the alert detects a legitimate action (a new application or developer maintenance). In other cases, the alert detects a malicious action (a former employee or external attacker).	Exploitation	Medium
Login from a principal user not seen in 60 days (SQL.DB_PrincipalAnomaly SQL.VM_PrincipalAnomaly SQL.DW_PrincipalAnomaly SQL.MI_PrincipalAnomaly)	A principal user not seen in the last 60 days has logged into your database. If this database is new or this is expected behavior caused by recent changes in the users accessing the database, Defender for Cloud will identify significant changes to the access patterns and attempt to prevent future false positives.	Exploitation	Medium
Login from a suspicious IP (SQL.VM_SuspiciousIpAnomaly)	Your resource has been accessed successfully from an IP address that Microsoft Threat Intelligence has associated with suspicious activity.	PreAttack	Medium
Potential SQL Brute Force attempt	An abnormally high number of failed sign in attempts with different credentials have occurred. In some cases, the alert detects penetration testing in action. In other cases, the alert detects a brute force attack.	Probing	High
Potential SQL injection (SQL.DB_PotentialSqlInjection SQL.VM_PotentialSqlInjection SQL.MI_PotentialSqlInjection SQL.DW_PotentialSqlInjection Synapse.SQLPool_PotentialSqlInjection)	An active exploit has occurred against an identified application vulnerable to SQL injection. This means an attacker is trying to inject malicious SQL statements by using the vulnerable application code or stored procedures.	PreAttack	High

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Potentially Unsafe Action (SQL.DB_UnsafeCommands SQL.MI_UnsafeCommands SQL.DW_UnsafeCommands)	A potentially unsafe action was attempted on your database '{name}' on server '{name}'.	-	High
Suspected brute force attack using a valid user	A potential brute force attack has been detected on your resource. The attacker is using the valid user sa, which has permissions to login.	PreAttack	High
Suspected brute force attack	A potential brute force attack has been detected on your SQL server '{name}'.	PreAttack	High
Suspected successful brute force attack (SQL.DB_BruteForce SQL.VM_BruteForce SQL.DW_BruteForce SQL.MI_BruteForce)	A successful login occurred after an apparent brute force attack on your resource	PreAttack	High
Unusual export location	Someone has extracted a massive amount of data from your SQL Server '{name}' to an unusual location.	Exfiltration	High

Alerts for open-source relational databases

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspected brute force attack using a valid user (SQL.PostgreSQL_BruteForce SQL.MariaDB_BruteForce SQL.MySQL_BruteForce)	A potential brute force attack has been detected on your resource. The attacker is using the valid user (username), which has permissions to login.	PreAttack	High
Suspected successful brute force attack (SQL.PostgreSQL_BruteForce SQL.MySQL_BruteForce SQL.MariaDB_BruteForce)	A successful login occurred after an apparent brute force attack on your resource.	PreAttack	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspected brute force attack ("SQL.MySQL_BruteForce")	A potential brute force attack has been detected on your SQL server '{name}'.	PreAttack	High
Attempted logon by a potentially harmful application (SQL.PostgreSQL_HarmfulApplication SQL.MariaDB_HarmfulApplication SQL.MySQL_HarmfulApplication)	A potentially harmful application attempted to access your resource.	PreAttack	High
Login from a principal user not seen in 60 days (SQL.PostgreSQL_PrincipalAnomaly SQL.MariaDB_PrincipalAnomaly SQL.MySQL_PrincipalAnomaly)	A principal user not seen in the last 60 days has logged into your database. If this database is new or this is expected behavior caused by recent changes in the users accessing the database, Defender for Cloud will identify significant changes to the access patterns and attempt to prevent future false positives.	Exploitation	Medium
Login from a domain not seen in 60 days (SQL.MariaDB_DomainAnomaly SQL.PostgreSQL_DomainAnomaly SQL.MySQL_DomainAnomaly)	A user has logged in to your resource from a domain no other users have connected from in the last 60 days. If this resource is new or this is expected behavior caused by recent changes in the users accessing the resource, Defender for Cloud will identify significant changes to the access patterns and attempt to prevent future false positives.	Exploitation	Medium
Log on from an unusual Azure Data Center (SQL.PostgreSQL_DataCenterAnomaly SQL.MariaDB_DataCenterAnomaly SQL.MySQL_DataCenterAnomaly)	Someone logged on to your resource from an unusual Azure Data Center.	Probing	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Logon from an unusual cloud provider (SQL.PostgreSQL_CloudProviderAnomaly SQL.MariaDB_CloudProviderAnomaly SQL.MySQL_CloudProviderAnomaly)	Someone logged on to your resource from a cloud provider not seen in the last 60 days. It's quick and easy for threat actors to obtain disposable compute power for use in their campaigns. If this is expected behavior caused by the recent adoption of a new cloud provider, Defender for Cloud will learn over time and attempt to prevent future false positives.	Exploitation	Medium
Log on from an unusual location (SQL.MariaDB_GeoAnomaly SQL.PostgreSQL_GeoAnomaly SQL.MySQL_GeoAnomaly)	Someone logged on to your resource from an unusual Azure Data Center.	Exploitation	Medium
Login from a suspicious IP (SQL.PostgreSQL_SuspiciousIpAnomaly SQL.MariaDB_SuspiciousIpAnomaly SQL.MySQL_SuspiciousIpAnomaly)	Your resource has been accessed successfully from an IP address that Microsoft Threat Intelligence has associated with suspicious activity.	PreAttack	Medium

Alerts for Resource Manager

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Azure Resource Manager operation from suspicious IP address (ARM_OperationFromSuspiciousIP)	Microsoft Defender for Resource Manager detected an operation from an IP address that has been marked as suspicious in threat intelligence feeds.	Execution	Medium
Azure Resource Manager operation from suspicious proxy IP address (ARM_OperationFromSuspiciousProxyIP)	Microsoft Defender for Resource Manager detected a resource management operation from an IP address that is associated with proxy services, such as TOR. While this behavior can be legitimate, it's often seen in malicious activities, when threat actors try to hide their source IP.	Defense Evasion	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
MicroBurst exploitation toolkit used to enumerate resources in your subscriptions (ARM_MicroBurst.AzDomainInfo)	MicroBurst's Information Gathering module was run on your subscription. This tool can be used to discover resources, permissions and network structures. This was detected by analyzing the Azure Activity logs and resource management operations in your subscription	-	High
MicroBurst exploitation toolkit used to enumerate resources in your subscriptions (ARM_MicroBurst.AzureDomainInfo)	MicroBurst's Information Gathering module was run on your subscription. This tool can be used to discover resources, permissions and network structures. This was detected by analyzing the Azure Activity logs and resource management operations in your subscription	-	High
MicroBurst exploitation toolkit used to execute code on your virtual machine (ARM_MicroBurst.AzVMBulkCMD)	MicroBurst's exploitation toolkit was used to execute code on your virtual machines. This was detected by analyzing Azure Resource Manager operations in your subscription.	Execution	High
MicroBurst exploitation toolkit used to execute code on your virtual machine (RM_MicroBurst.AzureRmVMBulkCMD)	MicroBurst's exploitation toolkit was used to execute code on your virtual machines. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High
MicroBurst exploitation toolkit used to extract keys from your Azure key vaults (ARM_MicroBurst.AzKeyVaultKeysREST)	MicroBurst's exploitation toolkit was used to extract keys from your Azure key vaults. This was detected by analyzing Azure Activity logs and resource management operations in your subscription.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
MicroBurst exploitation toolkit used to extract keys to your storage accounts (ARM_MicroBurst.AZStorageKeysREST)	MicroBurst's exploitation toolkit was used to extract keys to your storage accounts. This was detected by analyzing Azure Activity logs and resource management operations in your subscription.	Collection	High
MicroBurst exploitation toolkit used to extract secrets from your Azure key vaults (ARM_MicroBurst.AzKeyVaultSecretsREST)	MicroBurst's exploitation toolkit was used to extract secrets from your Azure key vaults. This was detected by analyzing Azure Activity logs and resource management operations in your subscription.	-	High
Permissions granted for an RBAC role in an unusual way for your Azure environment (Preview) (ARM_AnomalousRBACRoleAssignment)	Microsoft Defender for Resource Manager detected an RBAC role assignment that's unusual when compared with other assignments performed by the same assigner / performed for the same assignee / in your tenant due to the following anomalies: assignment time, assigner location, assigner, authentication method, assigned entities, client software used, assignment extent. This operation might have been performed by a legitimate user in your organization. Alternatively, it might indicate that an account in your organization was breached, and that the threat actor is trying to grant permissions to an additional user account they own.	Lateral Movement, Defense Evasion	Medium
PowerZure exploitation toolkit used to elevate access from Azure AD to Azure (ARM_PowerZure.AzureElevatedPrivileges)	PowerZure exploitation toolkit was used to elevate access from AzureAD to Azure. This was detected by analyzing Azure Resource Manager operations in your tenant.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PowerZure exploitation toolkit used to enumerate resources (ARM_PowerZure.GetAzureTargets)	PowerZure exploitation toolkit was used to enumerate resources on behalf of a legitimate user account in your organization. This was detected by analyzing Azure Resource Manager operations in your subscription.	Collection	High
PowerZure exploitation toolkit used to enumerate storage containers, shares, and tables (ARM_PowerZure.ShowStorageContent)	PowerZure exploitation toolkit was used to enumerate storage shares, tables, and containers. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High
PowerZure exploitation toolkit used to execute a Runbook in your subscription (ARM_PowerZure.StartRunbook)	PowerZure exploitation toolkit was used to execute a Runbook. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High
PowerZure exploitation toolkit used to extract Runbooks content (ARM_PowerZure.AzureRunbookContent)	PowerZure exploitation toolkit was used to extract Runbook content. This was detected by analyzing Azure Resource Manager operations in your subscription.	Collection	High
PREVIEW - Activity from a risky IP address (ARM.MCAS_ActivityFromAnonymousIPAddresses)	Users activity from an IP address that has been identified as an anonymous proxy IP address has been detected. These proxies are used by people who want to hide their device's IP address, and can be used for malicious intent. This detection uses a machine learning algorithm that reduces false positives, such as mis-tagged IP addresses that are widely used by users in the organization. Requires an active Microsoft Defender for Cloud Apps license.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW - Activity from infrequent country (ARM.MCAS_ActivityFromInfrequentCountry)	Activity from a location that wasn't recently or ever visited by any user in the organization has occurred. This detection considers past activity locations to determine new and infrequent locations. The anomaly detection engine stores information about previous locations used by users in the organization. Requires an active Microsoft Defender for Cloud Apps license.	-	Medium
PREVIEW - Azurite toolkit run detected (ARM_Azurite)	A known cloud-environment reconnaissance toolkit run has been detected in your environment. The tool Azurite can be used by an attacker (or penetration tester) to map your subscriptions' resources and identify insecure configurations.	Collection	High
PREVIEW - Impossible travel activity (ARM.MCAS_ImpossibleTravelActivity)	Two user activities (in a single or multiple sessions) have occurred, originating from geographically distant locations. This occurs within a time period shorter than the time it would have taken the user to travel from the first location to the second. This indicates that a different user is using the same credentials. This detection uses a machine learning algorithm that ignores obvious false positives contributing to the impossible travel conditions, such as VPNs and locations regularly used by other users in the organization. The detection has an initial learning period of seven days, during which it learns a new user's activity pattern. Requires an active Microsoft Defender for Cloud Apps license.	-	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW - Suspicious management session using an inactive account detected (ARM_UnusedAccountPersistence)	Subscription activity logs analysis has detected suspicious behavior. A principal not in use for a long period of time is now performing actions that can secure persistence for an attacker.	Persistence	Medium
PREVIEW - Suspicious management session using PowerShell detected (ARM_UnusedAppPowershellPersistence)	Subscription activity logs analysis has detected suspicious behavior. A principal that doesn't regularly use PowerShell to manage the subscription environment is now using PowerShell, and performing actions that can secure persistence for an attacker.	Persistence	Medium
PREVIEW – Suspicious management session using Azure portal detected (ARM_UnusedAppIbizaPersistence)	Analysis of your subscription activity logs has detected a suspicious behavior. A principal that doesn't regularly use the Azure portal (Ibiza) to manage the subscription environment (hasn't used Azure portal to manage for the last 45 days, or a subscription that it is actively managing), is now using the Azure portal and performing actions that can secure persistence for an attacker.	Persistence	Medium
Privileged custom role created for your subscription in a suspicious way (Preview) (ARM_PrivilegedRoleDefinitionCreation)	Microsoft Defender for Resource Manager detected a suspicious creation of privileged custom role definition in your subscription. This operation might have been performed by a legitimate user in your organization. Alternatively, it might indicate that an account in your organization was breached, and that the threat actor is trying to create a privileged role to use in the future to evade detection.	Privilege Escalation, Defense Evasion	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious invocation of a high-risk 'Credential Access' operation detected (Preview) (ARM_AnomalousOperation.CredentialAccess)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to access credentials. The identified operations are designed to allow administrators to efficiently access their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Credential Access	Medium
Suspicious invocation of a high-risk 'Data Collection' operation detected (Preview) (ARM_AnomalousOperation.Collection)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to collect data. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to collect sensitive data on resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Collection	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious invocation of a high-risk 'Defense Evasion' operation detected (Preview) (ARM_AnomalousOperation.DefenseEvasion)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to evade defenses. The identified operations are designed to allow administrators to efficiently manage the security posture of their environments. While this activity may be legitimate, a threat actor might utilize such operations to avoid being detected while compromising resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Defense Evasion	Medium
Suspicious invocation of a high-risk 'Execution' operation detected (Preview) (ARM_AnomalousOperation.Execution)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation on a machine in your subscription which might indicate an attempt to execute code. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious invocation of a high-risk 'Impact' operation detected (Preview) (ARM_AnomalousOperation.Impact)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempted configuration change. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to access restricted credentials and compromise resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Impact	Medium
Suspicious invocation of a high-risk 'Initial Access' operation detected (Preview) (ARM_AnomalousOperation.InitialAccess)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to access restricted resources. The identified operations are designed to allow administrators to efficiently access their environments. While this activity may be legitimate, a threat actor might utilize such operations to gain initial access to restricted resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious invocation of a high-risk 'Lateral Movement' operation detected (Preview) (ARM_AnomalousOperation.LateralMovement)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to perform lateral movement. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to compromise additional resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Lateral Movement	Medium
Suspicious invocation of a high-risk 'Persistence' operation detected (Preview) (ARM_AnomalousOperation.Persistence)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to establish persistence. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to establish persistence in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Persistence	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious invocation of a high-risk 'Privilege Escalation' operation detected (Preview) (ARM_AnomalousOperation.PrivilegeEscalation)	Microsoft Defender for Resource Manager identified a suspicious invocation of a high-risk operation in your subscription which might indicate an attempt to escalate privileges. The identified operations are designed to allow administrators to efficiently manage their environments. While this activity may be legitimate, a threat actor might utilize such operations to escalate privileges while compromising resources in your environment. This can indicate that the account is compromised and is being used with malicious intent.	Privilege Escalation	Medium
Usage of MicroBurst exploitation toolkit to run an arbitrary code or exfiltrate Azure Automation account credentials (ARM_MicroBurst.RunCodeOnBehalf)	Usage of MicroBurst exploitation toolkit to run an arbitrary code or exfiltrate Azure Automation account credentials. This was detected by analyzing Azure Resource Manager operations in your subscription.	Persistence, Credential Access	High
Usage of NetSPI techniques to maintain persistence in your Azure environment (ARM_NetSPI.MaintainPersistence)	Usage of NetSPI persistence technique to create a webhook backdoor and maintain persistence in your Azure environment. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High
Usage of PowerZure exploitation toolkit to run an arbitrary code or exfiltrate Azure Automation account credentials (ARM_PowerZure.RunCodeOnBehalf)	PowerZure exploitation toolkit detected attempting to run code or exfiltrate Azure Automation account credentials. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Usage of PowerZure function to maintain persistence in your Azure environment (ARM_PowerZure.MaintainPersistence)	PowerZure exploitation toolkit detected creating a webhook backdoor to maintain persistence in your Azure environment. This was detected by analyzing Azure Resource Manager operations in your subscription.	-	High

Alerts for DNS

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Anomalous network protocol usage (AzureDNS_ProtocolAnomaly)	Analysis of DNS transactions from % {CompromisedEntity} detected anomalous protocol usage. Such traffic, while possibly benign, may indicate abuse of this common protocol to bypass network traffic filtering. Typical related attacker activity includes copying remote administration tools to a compromised host and exfiltrating user data from it.	Exfiltration	-
Anonymity network activity (AzureDNS_DarkWeb)	Analysis of DNS transactions from % {CompromisedEntity} detected anonymity network activity. Such activity, while possibly legitimate user behavior, is frequently employed by attackers to evade tracking and fingerprinting of network communications. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Anonymity network activity using web proxy (AzureDNS_DarkWebProxy)	Analysis of DNS transactions from % {CompromisedEntity} detected anonymity network activity. Such activity, while possibly legitimate user behavior, is frequently employed by attackers to evade tracking and fingerprinting of network communications. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-
Attempted communication with suspicious sinkholed domain (AzureDNS_SinkholedDomain)	Analysis of DNS transactions from % {CompromisedEntity} detected request for sinkholed domain. Such activity, while possibly legitimate user behavior, is frequently an indication of the download or execution of malicious software. Typical related attacker activity is likely to include the download and execution of further malicious software or remote administration tools.	Exfiltration	-
Communication with possible phishing domain (AzureDNS_PhishingDomain)	Analysis of DNS transactions from % {CompromisedEntity} detected a request for a possible phishing domain. Such activity, while possibly benign, is frequently performed by attackers to harvest credentials to remote services. Typical related attacker activity is likely to include the exploitation of any credentials on the legitimate service.	Exfiltration	-

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Communication with suspicious algorithmically generated domain (AzureDNS_DomainGenerationAlgorithm)	Analysis of DNS transactions from % {CompromisedEntity} detected possible usage of a domain generation algorithm. Such activity, while possibly benign, is frequently performed by attackers to evade network monitoring and filtering. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-
Communication with suspicious domain identified by threat intelligence (AzureDNS_ThreatIntelSuspectDomain)	Communication with suspicious domain was detected by analyzing DNS transactions from your resource and comparing against known malicious domains identified by threat intelligence feeds. Communication to malicious domains is frequently performed by attackers and could imply that your resource is compromised.	Initial Access	Medium
Communication with suspicious random domain name (AzureDNS_RandomizedDomain)	Analysis of DNS transactions from % {CompromisedEntity} detected usage of a suspicious randomly generated domain name. Such activity, while possibly benign, is frequently performed by attackers to evade network monitoring and filtering. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Digital currency mining activity (AzureDNS_CurrencyMining)	Analysis of DNS transactions from % {CompromisedEntity} detected digital currency mining activity. Such activity, while possibly legitimate user behavior, is frequently performed by attackers following compromise of resources. Typical related attacker activity is likely to include the download and execution of common mining tools.	Exfiltration	-
Network intrusion detection signature activation (AzureDNS_SuspiciousDomain)	Analysis of DNS transactions from % {CompromisedEntity} detected a known malicious network signature. Such activity, while possibly legitimate user behavior, is frequently an indication of the download or execution of malicious software. Typical related attacker activity is likely to include the download and execution of further malicious software or remote administration tools.	Exfiltration	-
Possible data download via DNS tunnel (AzureDNS_DataInfiltration)	Analysis of DNS transactions from % {CompromisedEntity} detected a possible DNS tunnel. Such activity, while possibly legitimate user behavior, is frequently performed by attackers to evade network monitoring and filtering. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible data exfiltration via DNS tunnel (AzureDNS_DataExfiltration)	Analysis of DNS transactions from % {CompromisedEntity} detected a possible DNS tunnel. Such activity, while possibly legitimate user behavior, is frequently performed by attackers to evade network monitoring and filtering. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-
Possible data transfer via DNS tunnel (AzureDNS_DataObfuscation)	Analysis of DNS transactions from % {CompromisedEntity} detected a possible DNS tunnel. Such activity, while possibly legitimate user behavior, is frequently performed by attackers to evade network monitoring and filtering. Typical related attacker activity is likely to include the download and execution of malicious software or remote administration tools.	Exfiltration	-

Alerts for Azure Storage

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Access from a suspicious IP address (Storage.Blob_SuspiciousIp Storage.Files_SuspiciousIp)	Indicates that this storage account has been successfully accessed from an IP address that is considered suspicious. This alert is powered by Microsoft Threat Intelligence. Learn more about Microsoft's threat intelligence capabilities. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Initial Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW – Phishing content hosted on a storage account (Storage.Blob_PhishingContent Storage.Files_PhishingContent)	A URL used in a phishing attack points to your Azure Storage account. This URL was part of a phishing attack affecting users of Microsoft 365. Typically, content hosted on such pages is designed to trick visitors into entering their corporate credentials or financial information into a web form that looks legitimate. This alert is powered by Microsoft Threat Intelligence. Learn more about Microsoft's threat intelligence capabilities. Applies to: Azure Blob Storage, Azure Files	Collection	High
PREVIEW - Storage account identified as source for distribution of malware (Storage.Files_WidespreadeAm)	Antimalware alerts indicate that an infected file(s) is stored in an Azure file share that is mounted to multiple VMs. If attackers gain access to a VM with a mounted Azure file share, they can use it to spread malware to other VMs that mount the same share. Applies to: Azure Files	Lateral Movement, Execution	High
PREVIEW - Storage account with potentially sensitive data has been detected with a publicly exposed container (Storage.Blob_OpenACL)	The access policy of a container in your storage account was modified to allow anonymous access. This might lead to a data breach if the container holds any sensitive data. This alert is based on analysis of Azure activity log. Applies to: Azure Blob Storage, Azure Data Lake Storage Gen2	Privilege Escalation	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Authenticated access from a Tor exit node (Storage.Blob_TorAnomaly Storage.Files_TorAnomaly)	One or more storage container(s) / file share(s) in your storage account were successfully accessed from an IP address known to be an active exit node of Tor (an anonymizing proxy). Threat actors use Tor to make it difficult to trace the activity back to them. Authenticated access from a Tor exit node is a likely indication that a threat actor is trying to hide their identity. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Initial access	High/Medium
Access from an unusual location to a storage account (Storage.Blob_GeoAnomaly Storage.Files_GeoAnomaly)	Indicates that there was a change in the access pattern to an Azure Storage account. Someone has accessed this account from an IP address considered unfamiliar when compared with recent activity. Either an attacker has gained access to the account, or a legitimate user has connected from a new or unusual geographic location. An example of the latter is remote maintenance from a new application or developer. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Exploitation	Low
Unusual unauthenticated access to a storage container (Storage.Blob_AnonymousAccessAnomaly)	This storage account was accessed without authentication, which is a change in the common access pattern. Read access to this container is usually authenticated. This might indicate that a threat actor was able to exploit public read access to storage container(s) in this storage account(s). Applies to: Azure Blob Storage	Collection	Low

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Potential malware uploaded to a storage account (Storage.Blob_MalwareHashReputation Storage.Files_MalwareHashReputation)	Indicates that a blob containing potential malware has been uploaded to a blob container or a file share in a storage account. This alert is based on hash reputation analysis leveraging the power of Microsoft threat intelligence, which includes hashes for viruses, trojans, spyware and ransomware. Potential causes may include an intentional malware upload by an attacker, or an unintentional upload of a potentially malicious blob by a legitimate user. Applies to: Azure Blob Storage, Azure Files (Only for transactions over REST API) Learn more about Azure's hash reputation analysis for malware. Learn more about Microsoft's threat intelligence capabilities.	Lateral Movement	High
Publicly accessible storage containers successfully discovered (Storage.Blob_OpenContainersScanning.SuccessfulDiscovery)	A successful discovery of publicly open storage container(s) in your storage account was performed in the last hour by a scanning script or tool. This usually indicates a reconnaissance attack, where the threat actor tries to list blobs by guessing container names, in the hope of finding misconfigured open storage containers with sensitive data in them. The threat actor may use their own script or use known scanning tools like Microburst to scan for publicly open containers. ✔ Azure Blob Storage ✖ Azure Files ✖ Azure Data Lake Storage Gen2	Collection	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Publicly accessible storage containers unsuccessfully scanned (Storage.Blob_OpenContainersScanning.FailedAttempt)	A series of failed attempts to scan for publicly open storage containers were performed in the last hour. This usually indicates a reconnaissance attack, where the threat actor tries to list blobs by guessing container names, in the hope of finding misconfigured open storage containers with sensitive data in them. The threat actor may use their own script or use known scanning tools like Microburst to scan for publicly open containers. ✔ Azure Blob Storage ❌ Azure Files ❌ Azure Data Lake Storage Gen2	Collection	Low
Unusual access inspection in a storage account (Storage.Blob_AccessInspectionAnomaly Storage.Files_AccessInspectionAnomaly)	Indicates that the access permissions of a storage account have been inspected in an unusual way, compared to recent activity on this account. A potential cause is that an attacker has performed reconnaissance for a future attack. Applies to: Azure Blob Storage, Azure Files	Collection	Medium
Unusual amount of data extracted from a storage account (Storage.Blob_DataExfiltration.AmountOfDataAnomaly Storage.Blob_DataExfiltration.NumberOfBlobsAnomaly Storage.Files_DataExfiltration.AmountOfDataAnomaly Storage.Files_DataExfiltration.NumberOfFilesAnomaly)	Indicates that an unusually large amount of data has been extracted compared to recent activity on this storage container. A potential cause is that an attacker has extracted a large amount of data from a container that holds blob storage. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Exfiltration	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual application accessed a storage account (Storage.Blob_ApplicationAnomaly Storage.Files_ApplicationAnomaly)	Indicates that an unusual application has accessed this storage account. A potential cause is that an attacker has accessed your storage account by using a new application. Applies to: Azure Blob Storage, Azure Files	Exploitation	Medium
Unusual change of access permissions in a storage account (Storage.Blob_PermissionsChangeAnomaly Storage.Files_PermissionsChangeAnomaly)	Indicates that the access permissions of this storage container have been changed in an unusual way. A potential cause is that an attacker has changed container permissions to weaken its security posture or to gain persistence. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Persistence	Medium
Unusual data exploration in a storage account (Storage.Blob_DataExplorationAnomaly Storage.Files_DataExplorationAnomaly)	Indicates that blobs or containers in a storage account have been enumerated in an abnormal way, compared to recent activity on this account. A potential cause is that an attacker has performed reconnaissance for a future attack. Applies to: Azure Blob Storage, Azure Files	Collection	Medium
Unusual deletion in a storage account (Storage.Blob_DeletionAnomaly Storage.Files_DeletionAnomaly)	Indicates that one or more unexpected delete operations has occurred in a storage account, compared to recent activity on this account. A potential cause is that an attacker has deleted data from your storage account. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Exfiltration	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual upload of .cspkg to a storage account (Storage.Blob_CspkgUploadAnomaly)	Indicates that an Azure Cloud Services package (.cspkg file) has been uploaded to a storage account in an unusual way, compared to recent activity on this account. A potential cause is that an attacker has been preparing to deploy malicious code from your storage account to an Azure cloud service. Applies to: Azure Blob Storage, Azure Data Lake Storage Gen2	Lateral Movement, Execution	Medium
Unusual upload of .exe to a storage account (Storage.Blob_ExeUploadAnomaly Storage.Files_ExeUploadAnomaly)	Indicates that an .exe file has been uploaded to a storage account in an unusual way, compared to recent activity on this account. A potential cause is that an attacker has uploaded a malicious executable file to your storage account, or that a legitimate user has uploaded an executable file. Applies to: Azure Blob Storage, Azure Files, Azure Data Lake Storage Gen2	Lateral Movement, Execution	Medium

Alerts for Azure Cosmos DB (Preview)

Further details and notes

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW - Access from a Tor exit node	This Cosmos DB account was successfully accessed from an IP address known to be an active exit node of Tor, an anonymizing proxy. Authenticated access from a Tor exit node is a likely indication that a threat actor is trying to hide their identity.	Initial Access	High/Medium
PREVIEW - Access from a suspicious IP	This Cosmos DB account was successfully accessed from an IP address that was identified as a threat by Microsoft Threat Intelligence.	Initial Access	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW - Access from an unusual location	This Cosmos DB account was accessed from a location considered unfamiliar, based on the usual access pattern. Either a threat actor has gained access to the account, or a legitimate user has connected from a new or unusual geographic location	Initial Access	Low
PREVIEW - Unusual volume of data extracted	An unusually large volume of data has been extracted from this Cosmos DB account. This might indicate that a threat actor exfiltrated data.	Exfiltration	Medium
PREVIEW - Extraction of Cosmos DB accounts keys via a potentially malicious script	A PowerShell script was run in your subscription and performed a suspicious pattern of key-listing operations to get the keys of Cosmos DB accounts in your subscription. Threat actors use automated scripts, like Microburst, to list keys and find Cosmos DB accounts they can access. This operation might indicate that an identity in your organization was breached, and that the threat actor is trying to compromise Cosmos DB accounts in your environment for malicious intentions. Alternatively, a malicious insider could be trying to access sensitive data and perform lateral movement.	Collection	High

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
PREVIEW - SQL injection: potential data exfiltration	A suspicious SQL statement was used to query a container in this Cosmos DB account. The injected statement might have succeeded in exfiltrating data that the threat actor isn't authorized to access. Due to the structure and capabilities of Cosmos DB queries, many known SQL injection attacks on Cosmos DB accounts cannot work. However, the variation used in this attack may work and threat actors can exfiltrate data.	Exfiltration	Medium
PREVIEW - SQL injection: fuzzing attempt	A suspicious SQL statement was used to query a container in this Cosmos DB account. Like other well-known SQL injection attacks, this attack won't succeed in compromising the Cosmos DB account. Nevertheless, it's an indication that a threat actor is trying to attack the resources in this account, and your application may be compromised. Some SQL injection attacks can succeed and be used to exfiltrate data. This means that if the attacker continues performing SQL injection attempts, they may be able to compromise your Cosmos DB account and exfiltrate data. You can prevent this threat by using parameterized queries.	Pre-attack	Low

Alerts for Azure network layer

Further details and notes

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Network communication with a malicious machine detected (Network_CommunicationWithC2)	Network traffic analysis indicates that your machine (IP %{Victim IP}) has communicated with what is possibly a Command and Control center. When the compromised resource is a load balancer or an application gateway, the suspected activity might indicate that one or more of the resources in the backend pool (of the load balancer or application gateway) has communicated with what is possibly a Command and Control center.	Command and Control	Medium
Possible compromised machine detected (Network_ResourceIpIndicatedAsMalicious)	Threat intelligence indicates that your machine (at IP %{Machine IP}) may have been compromised by a malware of type Conficker. Conficker was a computer worm that targets the Microsoft Windows operating system and was first detected in November 2008. Conficker infected millions of computers including government, business and home computers in over 200 countries/regions, making it the largest known computer worm infection since the 2003 Welchia worm.	Command and Control	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible incoming %{Service Name} brute force attempts detected (Generic_Incoming_BF_OneToOne)	Network traffic analysis detected incoming %{Service Name} communication to %{Victim IP}, associated with your resource %{Compromised Host} from %{Attacker IP}. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows suspicious activity between %{Start Time} and %{End Time} on port %{Victim Port}. This activity is consistent with brute force attempts against %{Service Name} servers.	PreAttack	Medium
Possible incoming SQL brute force attempts detected (SQL_Incoming_BF_OneToOne)	Network traffic analysis detected incoming SQL communication to %{Victim IP}, associated with your resource %{Compromised Host}, from %{Attacker IP}. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows suspicious activity between %{Start Time} and %{End Time} on port %{Port Number} (%{SQL Service Type}). This activity is consistent with brute force attempts against SQL servers.	PreAttack	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Possible outgoing denial-of-service attack detected (DDOS)	Network traffic analysis detected anomalous outgoing activity originating from %{Compromised Host}, a resource in your deployment. This activity may indicate that your resource was compromised and is now engaged in denial-of-service attacks against external endpoints. When the compromised resource is a load balancer or an application gateway, the suspected activity might indicate that one or more of the resources in the backend pool (of the load balancer or application gateway) was compromised. Based on the volume of connections, we believe that the following IPs are possibly the targets of the DOS attack: %{Possible Victims}. Note that it is possible that the communication to some of these IPs is legitimate.	Impact	Medium
Suspicious incoming RDP network activity from multiple sources (RDP_Incoming_BF_ManyToOne)	Network traffic analysis detected anomalous incoming Remote Desktop Protocol (RDP) communication to %{Victim IP}, associated with your resource %{Compromised Host}, from multiple sources. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Attacking IPs} unique IPs connecting to your resource, which is considered abnormal for this environment. This activity may indicate an attempt to brute force your RDP end point from multiple hosts (Botnet)	PreAttack	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious incoming RDP network activity (RDP_Incoming_BF_OneToOne)	Network traffic analysis detected anomalous incoming Remote Desktop Protocol (RDP) communication to %{Victim IP}, associated with your resource %{Compromised Host}, from %{Attacker IP}. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Connections} incoming connections to your resource, which is considered abnormal for this environment. This activity may indicate an attempt to brute force your RDP end point	PreAttack	Medium
Suspicious incoming SSH network activity from multiple sources (SSH_Incoming_BF_ManyToOne)	Network traffic analysis detected anomalous incoming SSH communication to %{Victim IP}, associated with your resource %{Compromised Host}, from multiple sources. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Attacking IPs} unique IPs connecting to your resource, which is considered abnormal for this environment. This activity may indicate an attempt to brute force your SSH end point from multiple hosts (Botnet)	PreAttack	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious incoming SSH network activity (SSH_Incoming_BF_OneToOne)	Network traffic analysis detected anomalous incoming SSH communication to %{Victim IP}, associated with your resource %{Compromised Host}, from %{Attacker IP}. When the compromised resource is a load balancer or an application gateway, the suspected incoming traffic has been forwarded to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Connections} incoming connections to your resource, which is considered abnormal for this environment. This activity may indicate an attempt to brute force your SSH end point	PreAttack	Medium
Suspicious outgoing %{Attacked Protocol} traffic detected (PortScanning)	Network traffic analysis detected suspicious outgoing traffic from %{Compromised Host} to destination port %{Most Common Port}. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). This behavior may indicate that your resource is taking part in %{Attacked Protocol} brute force attempts or port sweeping attacks.	Discovery	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious outgoing RDP network activity to multiple destinations (RDP_Outgoing_BF_OneToMany)	Network traffic analysis detected anomalous outgoing Remote Desktop Protocol (RDP) communication to multiple destinations originating from %{Compromised Host} (%{Attacker IP}), a resource in your deployment. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows your machine connecting to %{Number of Attacked IPs} unique IPs, which is considered abnormal for this environment. This activity may indicate that your resource was compromised and is now used to brute force external RDP end points. Note that this type of activity could possibly cause your IP to be flagged as malicious by external entities.	Discovery	High

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious outgoing RDP network activity (RDP_Outgoing_BF_OneToOne)	Network traffic analysis detected anomalous outgoing Remote Desktop Protocol (RDP) communication to %{Victim IP} originating from %{Compromised Host} (%{Attacker IP}), a resource in your deployment. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Connections} outgoing connections from your resource, which is considered abnormal for this environment. This activity may indicate that your machine was compromised and is now used to brute force external RDP end points. Note that this type of activity could possibly cause your IP to be flagged as malicious by external entities.	Lateral Movement	High

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious outgoing SSH network activity to multiple destinations (SSH_Outgoing_BF_OneToMany)	Network traffic analysis detected anomalous outgoing SSH communication to multiple destinations originating from %{Compromised Host} (%{Attacker IP}), a resource in your deployment. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows your resource connecting to % {Number of Attacked IPs} unique IPs, which is considered abnormal for this environment. This activity may indicate that your resource was compromised and is now used to brute force external SSH end points. Note that this type of activity could possibly cause your IP to be flagged as malicious by external entities.	Discovery	Medium

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious outgoing SSH network activity (SSH_Outgoing_BF_OneToOne)	Network traffic analysis detected anomalous outgoing SSH communication to %{Victim IP} originating from %{Compromised Host} (%{Attacker IP}), a resource in your deployment. When the compromised resource is a load balancer or an application gateway, the suspected outgoing traffic has been originated from to one or more of the resources in the backend pool (of the load balancer or application gateway). Specifically, sampled network data shows %{Number of Connections} outgoing connections from your resource, which is considered abnormal for this environment. This activity may indicate that your resource was compromised and is now used to brute force external SSH end points. Note that this type of activity could possibly cause your IP to be flagged as malicious by external entities.	Lateral Movement	Medium
Traffic detected from IP addresses recommended for blocking	Microsoft Defender for Cloud detected inbound traffic from IP addresses that are recommended to be blocked. This typically occurs when this IP address doesn't communicate regularly with this resource. Alternatively, the IP address has been flagged as malicious by Defender for Cloud's threat intelligence sources.	Probing	Low

Alerts for Azure Key Vault

Further details and notes

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Access from a suspicious IP address to a key vault (KV_SuspiciousIPAccess)	A key vault has been successfully accessed by an IP that has been identified by Microsoft Threat Intelligence as a suspicious IP address. This may indicate that your infrastructure has been compromised. We recommend further investigation. Learn more about Microsoft's threat intelligence capabilities.	Credential Access	Medium
Access from a TOR exit node to a key vault (KV_TORAccess)	A key vault has been accessed from a known TOR exit node. This could be an indication that a threat actor has accessed the key vault and is using the TOR network to hide their source location. We recommend further investigations.	Credential Access	Medium
High volume of operations in a key vault (KV_OperationVolumeAnomaly)	An anomalous number of key vault operations were performed by a user, service principal, and/or a specific key vault. This anomalous activity pattern may be legitimate, but it could be an indication that a threat actor has gained access to the key vault and the secrets contained within it. We recommend further investigations.	Credential Access	Medium
Suspicious policy change and secret query in a key vault (KV_PutGetAnomaly)	A user or service principal has performed an anomalous Vault Put policy change operation followed by one or more Secret Get operations. This pattern is not normally performed by the specified user or service principal. This may be legitimate activity, but it could be an indication that a threat actor has updated the key vault policy to access previously inaccessible secrets. We recommend further investigations.	Credential Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Suspicious secret listing and query in a key vault (KV_ListGetAnomaly)	A user or service principal has performed an anomalous Secret List operation followed by one or more Secret Get operations. This pattern is not normally performed by the specified user or service principal and is typically associated with secret dumping. This may be legitimate activity, but it could be an indication that a threat actor has gained access to the key vault and is trying to discover secrets that can be used to move laterally through your network and/or gain access to sensitive resources. We recommend further investigations.	Credential Access	Medium
Unusual application accessed a key vault (KV_AppAnomaly)	A key vault has been accessed by a service principal that does not normally access it. This anomalous access pattern may be legitimate activity, but it could be an indication that a threat actor has gained access to the key vault in an attempt to access the secrets contained within it. We recommend further investigations.	Credential Access	Medium
Unusual operation pattern in a key vault KV_OperationPatternAnomaly)	An anomalous pattern of key vault operations was performed by a user, service principal, and/or a specific key vault. This anomalous activity pattern may be legitimate, but it could be an indication that a threat actor has gained access to the key vault and the secrets contained within it. We recommend further investigations.	Credential Access	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Unusual user accessed a key vault (KV_UserAnomaly)	A key vault has been accessed by a user that does not normally access it. This anomalous access pattern may be legitimate activity, but it could be an indication that a threat actor has gained access to the key vault in an attempt to access the secrets contained within it. We recommend further investigations.	Credential Access	Medium
Unusual user-application pair accessed a key vault (KV_UserAppAnomaly)	A key vault has been accessed by a user-service principal pair that does not normally access it. This anomalous access pattern may be legitimate activity, but it could be an indication that a threat actor has gained access to the key vault in an attempt to access the secrets contained within it. We recommend further investigations.	Credential Access	Medium
User accessed high volume of key vaults (KV_AccountVolumeAnomaly)	A user or service principal has accessed an anomalously high volume of key vaults. This anomalous access pattern may be legitimate activity, but it could be an indication that a threat actor has gained access to multiple key vaults in an attempt to access the secrets contained within them. We recommend further investigations.	Credential Access	Medium

Alerts for Azure DDoS Protection

Further details and notes

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
DDoS Attack detected for Public IP	DDoS Attack detected for Public IP (IP address) and being mitigated.	Probing	High
DDoS Attack mitigated for Public IP	DDoS Attack mitigated for Public IP (IP address).	Probing	Low

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY

Security incident alerts

Further details and notes

ALERT	DESCRIPTION	MITRE TACTICS (LEARN MORE)	SEVERITY
Security incident with shared process detected	The incident which started on {Start Time (UTC)} and recently detected on {Detected Time (UTC)} indicates that an attacker has {Action taken} your resource {Host}	-	High
Security incident detected on multiple resources	The incident which started on {Start Time (UTC)} and recently detected on {Detected Time (UTC)} indicates that similar attack methods were performed on your cloud resources {Host}	-	Medium
Security incident detected from same source	The incident which started on {Start Time (UTC)} and recently detected on {Detected Time (UTC)} indicates that an attacker has {Action taken} your resource {Host}	-	High
Security incident detected on multiple machines	The incident which started on {Start Time (UTC)} and recently detected on {Detected Time (UTC)} indicates that an attacker has {Action taken} your resources {Host}	-	Medium

MITRE ATT&CK tactics

Understanding the intention of an attack can help you investigate and report the event more easily. To help with these efforts, Microsoft Defender for Cloud alerts include the MITRE tactics with many alerts.

The series of steps that describe the progression of a cyberattack from reconnaissance to data exfiltration is often referred to as a "kill chain".

Defender for Cloud's supported kill chain intents are based on version 7 of the MITRE ATT&CK matrix and described in the table below.

TACTIC	DESCRIPTION
PreAttack	PreAttack could be either an attempt to access a certain resource regardless of a malicious intent, or a failed attempt to gain access to a target system to gather information prior to exploitation. This step is usually detected as an attempt, originating from outside the network, to scan the target system and identify an entry point.
Initial Access	Initial Access is the stage where an attacker manages to get a foothold on the attacked resource. This stage is relevant for compute hosts and resources such as user accounts, certificates etc. Threat actors will often be able to control the resource after this stage.
Persistence	Persistence is any access, action, or configuration change to a system that gives a threat actor a persistent presence on that system. Threat actors will often need to maintain access to systems through interruptions such as system restarts, loss of credentials, or other failures that would require a remote access tool to restart or provide an alternate backdoor for them to regain access.
Privilege Escalation	Privilege escalation is the result of actions that allow an adversary to obtain a higher level of permissions on a system or network. Certain tools or actions require a higher level of privilege to work and are likely necessary at many points throughout an operation. User accounts with permissions to access specific systems or perform specific functions necessary for adversaries to achieve their objective may also be considered an escalation of privilege.
Defense Evasion	Defense evasion consists of techniques an adversary may use to evade detection or avoid other defenses. Sometimes these actions are the same as (or variations of) techniques in other categories that have the added benefit of subverting a particular defense or mitigation.
Credential Access	Credential access represents techniques resulting in access to or control over system, domain, or service credentials that are used within an enterprise environment. Adversaries will likely attempt to obtain legitimate credentials from users or administrator accounts (local system administrator or domain users with administrator access) to use within the network. With sufficient access within a network, an adversary can create accounts for later use within the environment.
Discovery	Discovery consists of techniques that allow the adversary to gain knowledge about the system and internal network. When adversaries gain access to a new system, they must orient themselves to what they now have control of and what benefits operating from that system give to their current objective or overall goals during the intrusion. The operating system provides many native tools that aid in this post-compromise information-gathering phase.

TACTIC	DESCRIPTION
LateralMovement	Lateral movement consists of techniques that enable an adversary to access and control remote systems on a network and could, but does not necessarily, include execution of tools on remote systems. The lateral movement techniques could allow an adversary to gather information from a system without needing additional tools, such as a remote access tool. An adversary can use lateral movement for many purposes, including remote Execution of tools, pivoting to additional systems, access to specific information or files, access to additional credentials, or to cause an effect.
Execution	The execution tactic represents techniques that result in execution of adversary-controlled code on a local or remote system. This tactic is often used in conjunction with lateral movement to expand access to remote systems on a network.
Collection	Collection consists of techniques used to identify and gather information, such as sensitive files, from a target network prior to exfiltration. This category also covers locations on a system or network where the adversary may look for information to exfiltrate.
Exfiltration	Exfiltration refers to techniques and attributes that result or aid in the adversary removing files and information from a target network. This category also covers locations on a system or network where the adversary may look for information to exfiltrate.
Command and Control	The command and control tactic represents how adversaries communicate with systems under their control within a target network.
Impact	Impact events primarily try to directly reduce the availability or integrity of a system, service, or network; including manipulation of data to impact a business or operational process. This would often refer to techniques such as ransomware, defacement, data manipulation, and others.

NOTE

For alerts that are in preview: The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

Next steps

To learn more about Microsoft Defender for Cloud security alerts, see the following:

- Security alerts in Microsoft Defender for Cloud
- Manage and respond to security alerts in Microsoft Defender for Cloud
- Continuously export Defender for Cloud data

Microsoft Defender for Cloud's overview page

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

When you open Microsoft Defender for Cloud, the first page to appear is the overview page.

This interactive dashboard provides a unified view into the security posture of your hybrid cloud workloads. Additionally, it shows security alerts, coverage information, and more.

You can select any element on the page to get more detailed information.

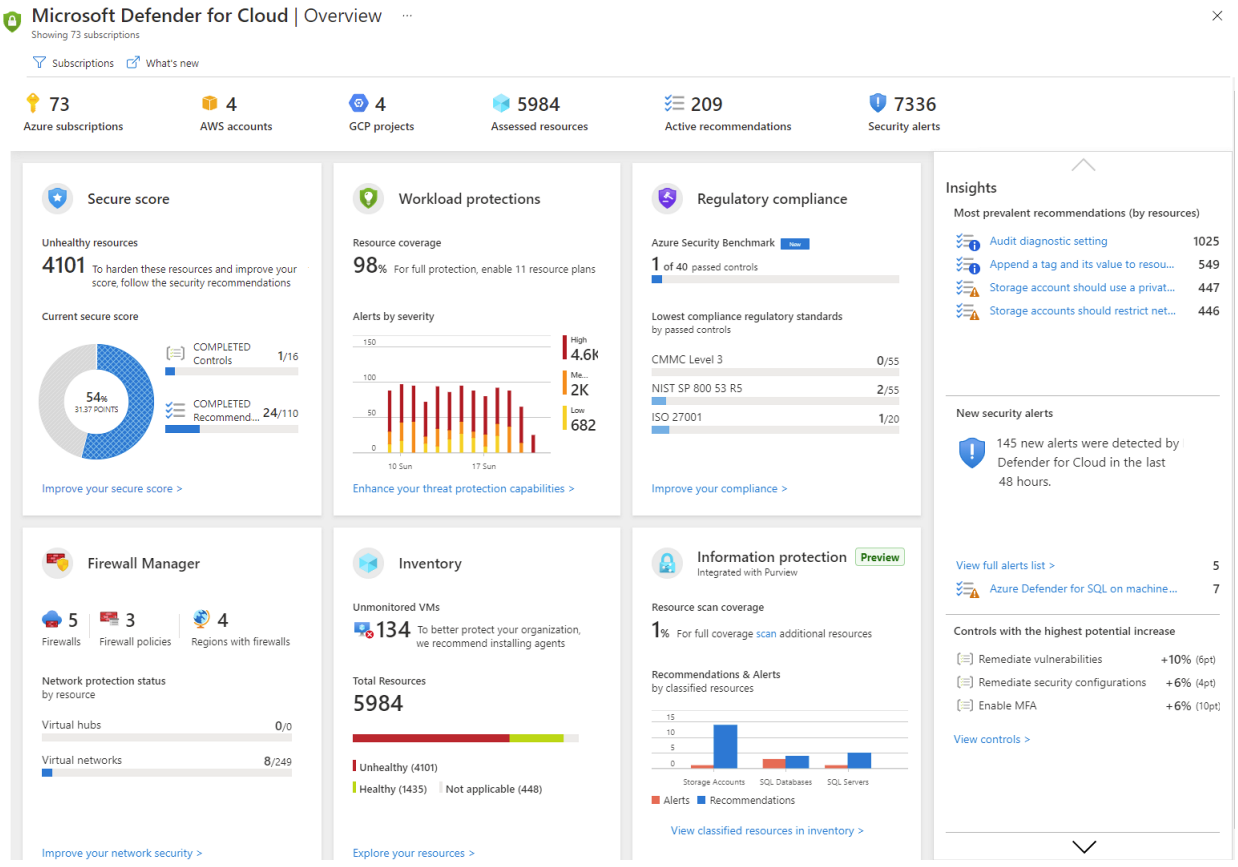

Features of the overview page

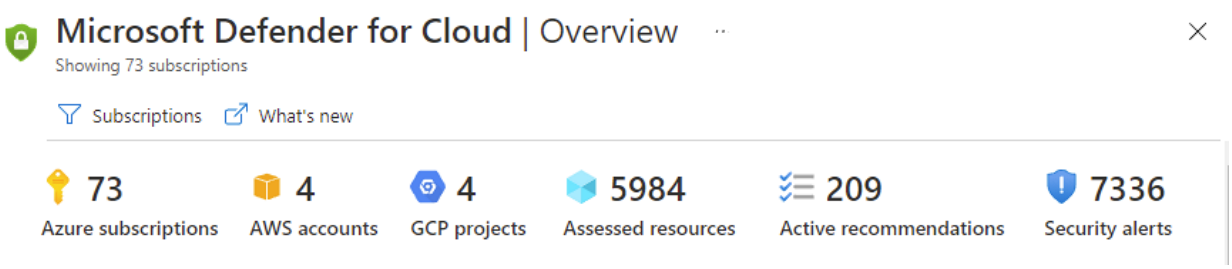

Metrics

The **top menu bar** offers:

- **Subscriptions** - You can view and filter the list of subscriptions by selecting this button. Defender for Cloud will adjust the display to reflect the security posture of the selected subscriptions.
- **What's new** - Opens the release notes so you can keep up to date with new features, bug fixes, and deprecated functionality.
- **High-level numbers** for the connected cloud accounts, to show the context of the information in the main tiles below. As well as the number of assessed resources, active recommendations, and security alerts. Select the assessed resources number to access Asset inventory. Learn more about connecting your AWS accounts and your GCP projects.

Feature tiles

In the center of the page are the **feature tiles**, each linking to a high profile feature or dedicated dashboard:

- **Secure score** - Defender for Cloud continually assesses your resources, subscriptions, and organization for security issues. It then aggregates all the findings into a single score so that you can tell, at a glance, your current security situation: the higher the score, the lower the identified risk level. Learn more.
- **Workload protections** - This is the cloud workload protection platform (CWPP) integrated within Defender for Cloud for advanced, intelligent protection of your workloads running on Azure, on-premises machines, or other cloud providers. For each resource type, there's a corresponding Microsoft Defender plan. The tile shows the coverage of your connected resources (for the currently selected subscriptions) and the recent alerts, color-coded by severity. Learn more about the enhanced security features.
- **Regulatory compliance** - Defender for Cloud provides insights into your compliance posture based on continuous assessments of your Azure environment. Defender for Cloud analyzes risk factors in your environment according to security best practices. These assessments are mapped to compliance controls from a supported set of standards. Learn more.
- **Firewall Manager** - This tile shows the status of your hubs and networks from Azure Firewall Manager.
- **Inventory** - The asset inventory page of Microsoft Defender for Cloud provides a single page for viewing the security posture of the resources you've connected to Microsoft Defender for Cloud. All resources with unresolved security recommendations are shown in the inventory. If you've enabled the integration with Microsoft Defender for Endpoint and enabled Microsoft Defender for servers, you'll also have access to a software inventory. The tile on the overview page shows you at a glance the total healthy and unhealthy resources (for the currently selected subscriptions). Learn more.
- **Information protection** - A graph on this tile shows the resource types that have been scanned by Azure Purview, found to contain sensitive data, and have outstanding recommendations and alerts. Follow the **scan** link to access the Azure Purview accounts and configure new scans, or select any other part of the tile to open the asset inventory and view your resources according to your Azure Purview data sensitivity classifications. Learn more.

Insights

The **Insights** pane offers customized items for your environment including:

- Your most attacked resources
- Your security controls that have the highest potential to increase your secure score
- The active recommendations with the most resources impacted
- Recent blog posts by Microsoft Defender for Cloud experts

Next steps

This page introduced the Defender for Cloud overview page. For related information, see:

- Explore and manage your resources with asset inventory and management tools
- Secure score in Microsoft Defender for Cloud

The workload protections dashboard

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This dashboard provides:

- Visibility into your Microsoft Defender for Cloud coverage across your different resource types
- Links to configure advanced threat protection capabilities
- The onboarding state and agent installation
- Threat detection alerts

To access the workload protections dashboard, select **Workload protections** from the Cloud Security section of Defender for Cloud's menu.

What's shown on the dashboard?

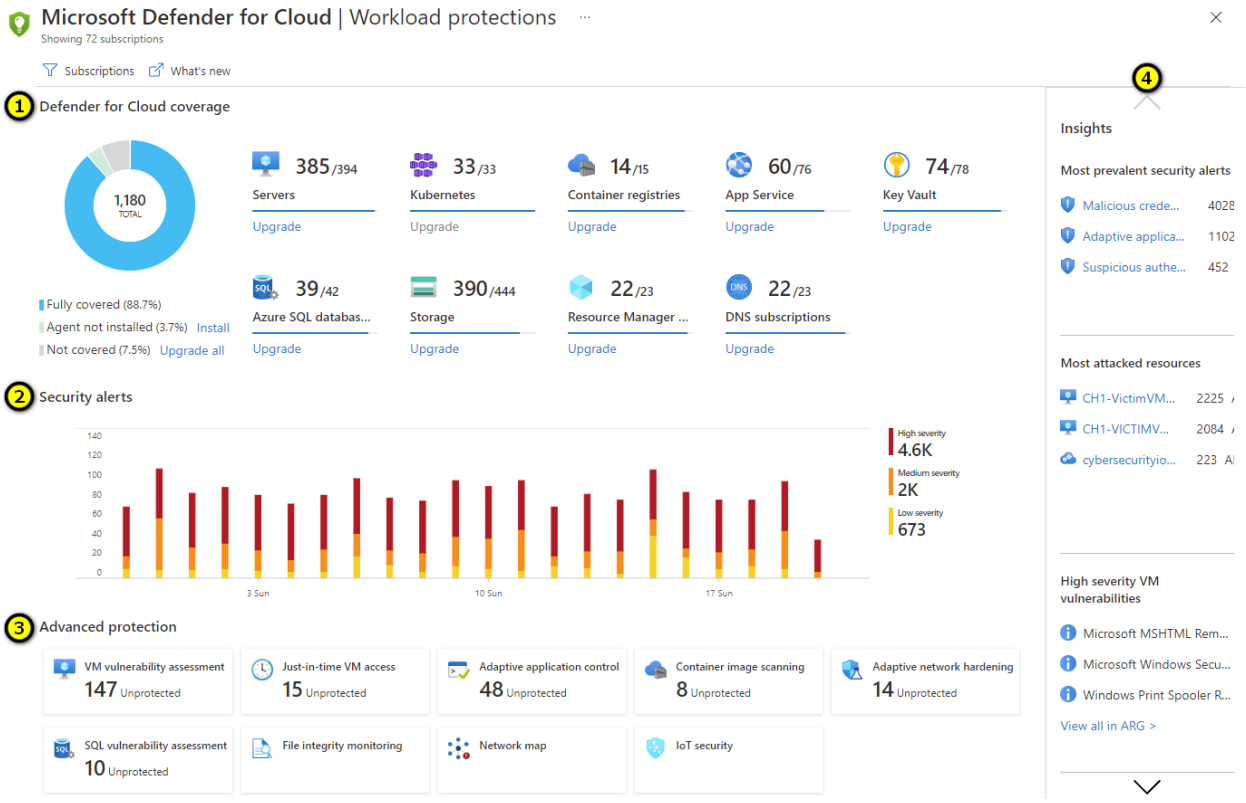

The dashboard includes the following sections:

1. **Microsoft Defender for Cloud coverage** - Here you can see the resources types that are in your subscription and eligible for protection by Defender for Cloud. Wherever relevant, you'll have the option to upgrade too. If you want to upgrade all possible eligible resources, select **Upgrade all**.

2. **Security alerts** - When Defender for Cloud detects a threat in any area of your environment, it generates an alert. These alerts describe details of the affected resources, suggested remediation steps,

and in some cases an option to trigger a logic app in response. Selecting anywhere in this graph opens the **Security alerts page**.

3. **Advanced protection** - Defender for Cloud includes many advanced threat protection capabilities for virtual machines, SQL databases, containers, web applications, your network, and more. In this advanced protection section, you can see the status of the resources in your selected subscriptions for each of these protections. Select any of them to go directly to the configuration area for that protection type.

4. **Insights** - This rolling pane of news, suggested reading, and high priority alerts gives Defender for Cloud's insights into pressing security matters that are relevant to you and your subscription. Whether it's a list of high severity CVEs discovered on your VMs by a vulnerability analysis tool, or a new blog post by a member of the Defender for Cloud team, you'll find it here in the Insights panel.

Next steps

In this article, you learned about the workload protections dashboard.

Enable enhanced protections

For more on the advanced protection plans of Microsoft Defender for Cloud, see Introduction to Microsoft Defender for Cloud

Access and track your secure score

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

You can find your overall secure score, as well as your score per subscription, through the Azure portal or programmatically as described in the following sections:

> **TIP**
>
> For a detailed explanation of how your scores are calculated, see Calculations - understanding your score.

Get your secure score from the portal

Defender for Cloud displays your score prominently in the portal: it's the first main tile the Defender for Cloud overview page. Selecting this tile, takes you to the dedicated secure score page, where you'll see the score broken down by subscription. Select a single subscription to see the detailed list of prioritized recommendations and the potential impact that remediating them will have on the subscription's score.

To recap, your secure score is shown in the following locations in Defender for Cloud's portal pages.

- In a tile on Defender for Cloud's **Overview** (main dashboard):

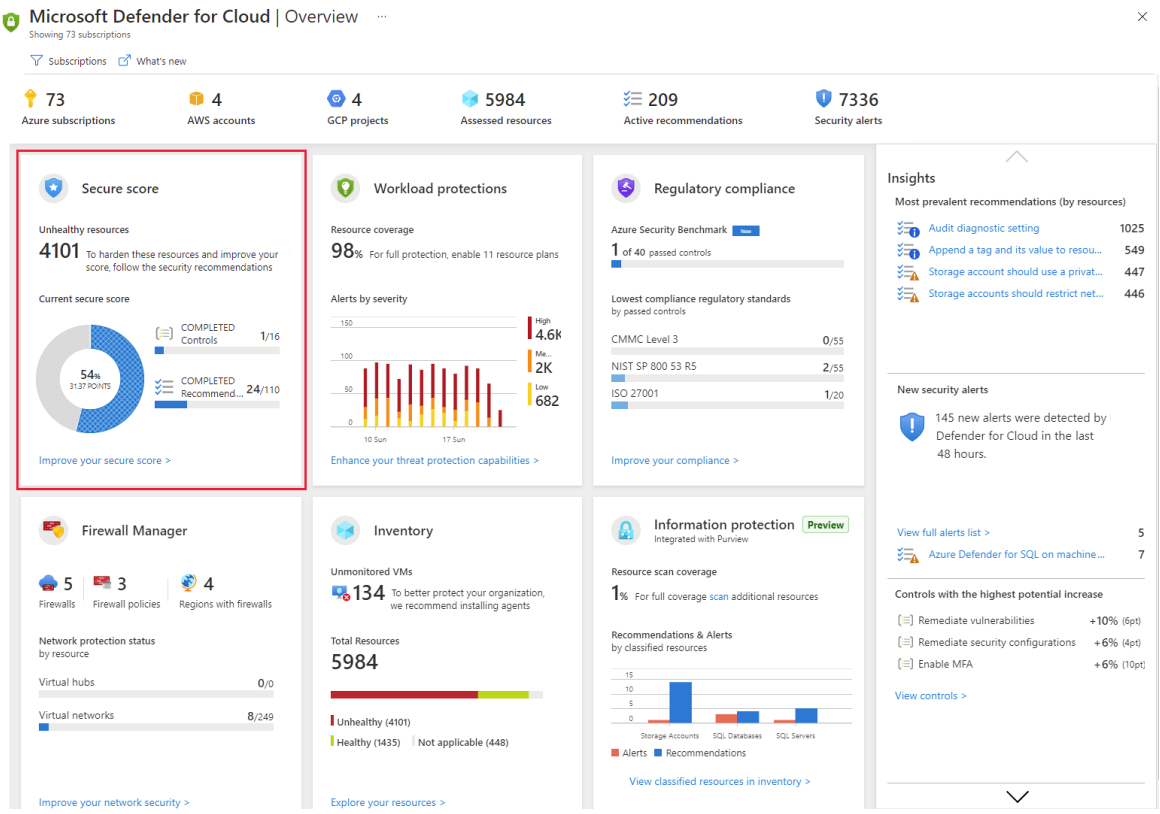

- In the dedicated **Secure score** page you can see the secure score for your subscription and your

management groups:

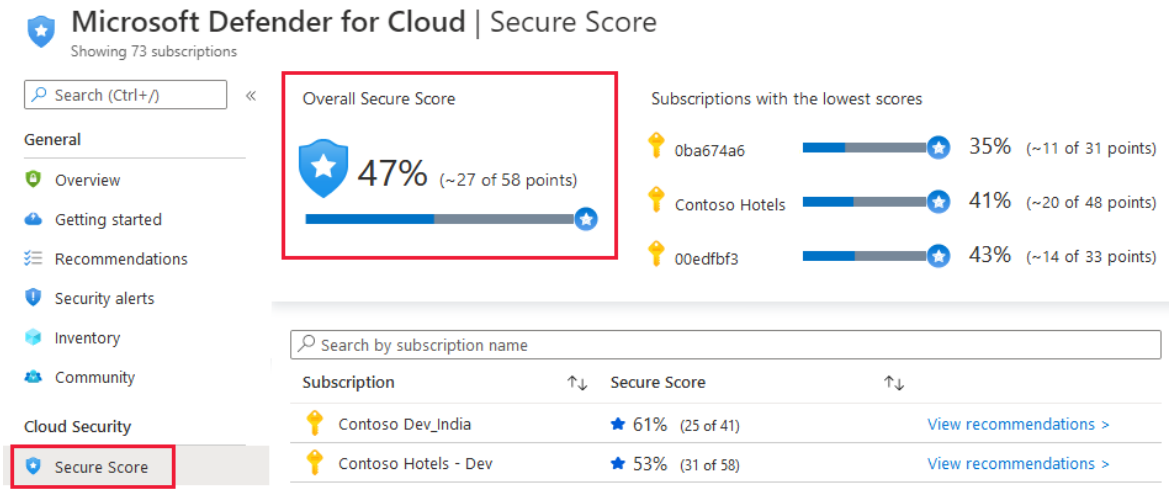

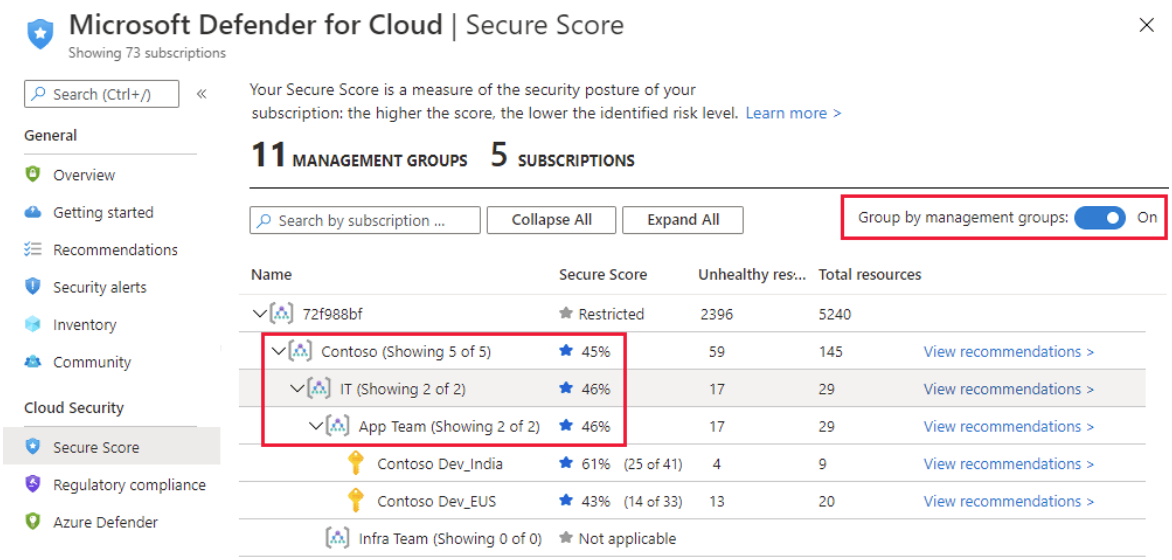

> **NOTE**
>
> Any management groups for which you don't have sufficient permissions, will show their score as "Restricted."

- At the top of the **Recommendations** page:

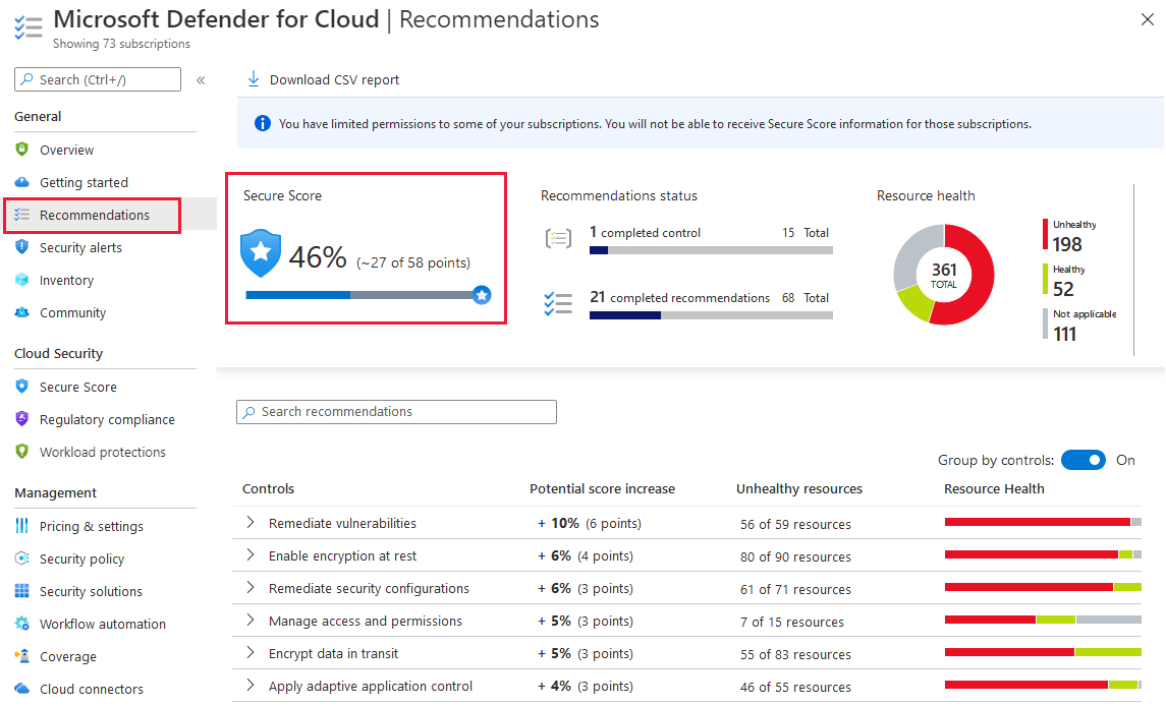

Get your secure score from the REST API

You can access your score via the secure score API. The API methods provide the flexibility to query the data and build your own reporting mechanism of your secure scores over time. For example, you can use the Secure Scores API to get the score for a specific subscription. In addition, you can use the Secure Score Controls API to list the security controls and the current score of your subscriptions.

Get single secure score

Sample Request

HTTP Copy

```
GET https://management.azure.com/subscriptions/20ff7fc3-e762-44dd-bd96-b71116dcdc23/providers/Microsoft.Security
/secureScores/ascScore?api-version=2020-01-01-preview
```

Sample Response

Status code: 200

JSON Copy

```
{
  "id": "/subscriptions/20ff7fc3-e762-44dd-bd96-b71116dcdc23/providers/Microsoft.Security/secureScores/ascScore",
  "name": "ascScore",
  "type": "Microsoft.Security/secureScores",
  "properties": {
    "displayName": "ASC score",
    "score": {
      "max": 13,
      "current": 3
    }
  }
}
```

For examples of tools built on top of the secure score API, see the secure score area of our GitHub community.

Get your secure score from Azure Resource Graph

Azure Resource Graph provides instant access to resource information across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal. Learn more about Azure Resource Graph.

To access the secure score for multiple subscriptions with Azure Resource Graph:

1. From the Azure portal, open **Azure Resource Graph Explorer**.

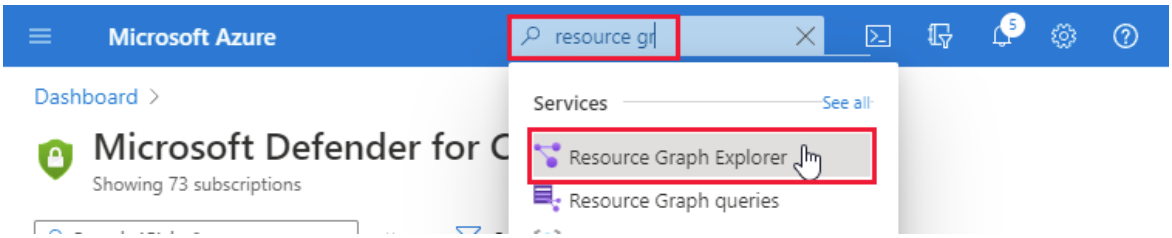

2. Enter your Kusto query (using the examples below for guidance).

   - This query returns the subscription ID, the current score in points and as a percentage, and the maximum score for the subscription.

     ```
     SecurityResources
     | where type == 'microsoft.security/securescores'
     | extend current = properties.score.current, max = todouble(properties.score.max)
     | project subscriptionId, current, max, percentage = ((current / max)*100)
     ```

   - This query returns the status of all the security controls. For each control, you'll get the number of unhealthy resources, the current score, and the maximum score.

     ```
     SecurityResources
     | where type == 'microsoft.security/securescores/securescorecontrols'
     | extend SecureControl = properties.displayName, unhealthy =
     properties.unhealthyResourceCount, currentscore = properties.score.current, maxscore =
     properties.score.max
     | project SecureControl , unhealthy, currentscore, maxscore
     ```

3. Select **Run query**.

Tracking your secure score over time

Secure Score Over Time report in workbooks page

Defender for Cloud's workbooks page includes a ready-made report for visually tracking the scores of your subscriptions, security controls, and more. Learn more in Create rich, interactive reports of Defender for Cloud data.

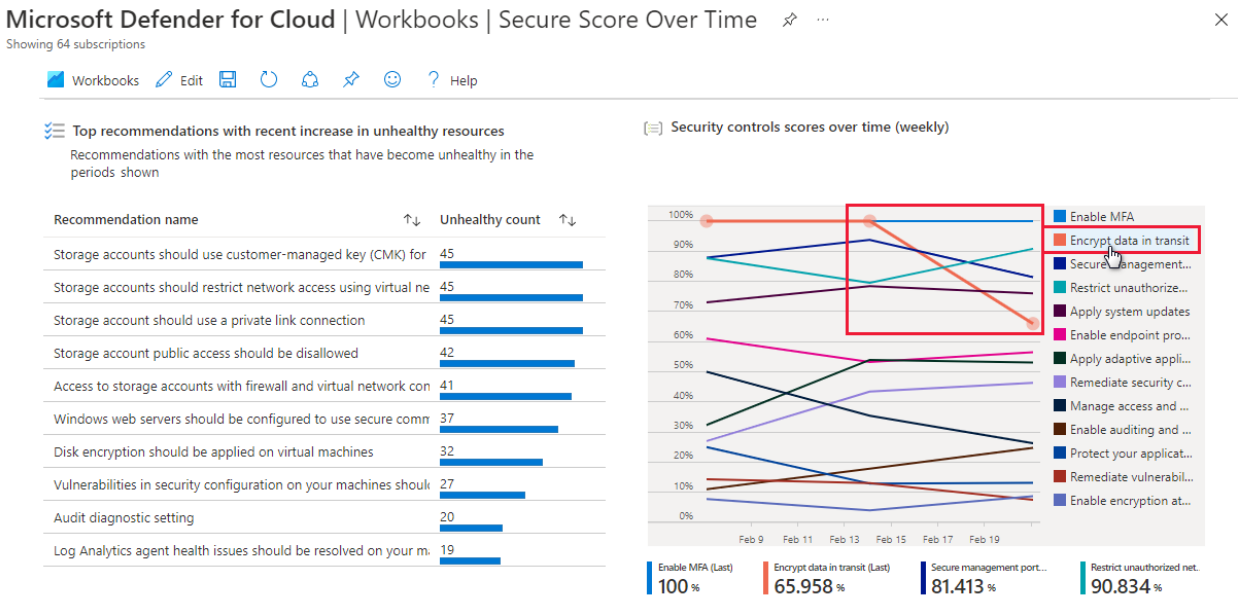

Power BI Pro dashboards

If you're a Power BI user with a Pro account, you can use the **Secure Score Over Time** Power BI dashboard to track your secure score over time and investigate any changes.

> **TIP**
>
> You can find this dashboard, as well as other tools for working programmatically with secure score, in the dedicated area of the Microsoft Defender for Cloud community on GitHub: https://github.com/Azure/Azure-Security-Center/tree/master/Secure%20Score

The dashboard contains the following two reports to help you analyze your security status:

- **Resources Summary** - provides summarized data regarding your resources' health.
- **Secure Score Summary** - provides summarized data regarding your score progress. Use the "Secure score over time per subscription" chart to view changes in the score. If you notice a dramatic change in your score, check the "detected changes that may affect your secure score" table for possible changes that could have caused the change. This table presents deleted resources, newly deployed resources, or resources that their security status changed for one of the recommendations.

Next steps

This article described how to access and track your secure score. For related material, see the following articles:

- Learn about the different elements of a recommendation
- Learn how to remediate recommendations
- View the GitHub-based tools for working programmatically with secure score

Prioritize security actions by data sensitivity

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Purview, Microsoft's data governance service, provides rich insights into the *sensitivity of your data*. With automated data discovery, sensitive data classification, and end-to-end data lineage, Azure Purview helps organizations manage and govern data in hybrid and multi-cloud environments.

Microsoft Defender for Cloud customers using Azure Purview can benefit from an additional vital layer of metadata in alerts and recommendations: information about any potentially sensitive data involved. This knowledge helps solve the triage challenge and ensures security professionals can focus their attention on threats to sensitive data.

This page explains the integration of Azure Purview's data sensitivity classification labels within Defender for Cloud.

Availability

ASPECT	DETAILS
Release state:	Preview. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	You'll need an Azure Purview account to create the data sensitivity classifications and run the scans. Viewing the scan results and using the output is free for Defender for Cloud users
Required roles and permissions:	**Security admin** and **Security contributor**
Clouds:	✅ Commercial clouds ❌ Azure Government ❌ Azure China 21Vianet (**Partial**: Subset of alerts and vulnerability assessment for SQL servers. Behavioral threat protections aren't available.)

The triage problem and Defender for Cloud's solution

Security teams regularly face the challenge of how to triage incoming issues.

Defender for Cloud includes two mechanisms to help prioritize recommendations and security alerts:

- For recommendations, we've provided **security controls** to help you understand how important each recommendation is to your overall security posture. Defender for Cloud includes a **secure score** value for each control to help you prioritize your security work. Learn more in Security controls and their recommendations.
- For alerts, we've assigned **severity labels** to each alert to help you prioritize the order in which you attend to each alert. Learn more in How are alerts classified?.

However, where possible, you'd want to focus the security team's efforts on risks to the organization's **data**. If two recommendations have equal impact on your secure score, but one relates to a resource with sensitive data, ideally you'd include that knowledge when determining prioritization.

Azure Purview's data sensitivity classifications and data sensitivity labels provide that knowledge.

Discover resources with sensitive data

To provide the vital information about discovered sensitive data, and help ensure you have that information when you need it, Defender for Cloud displays information from Azure Purview in multiple locations.

> **TIP**
> If a resource is scanned by multiple Azure Purview accounts, the information shown in Defender for Cloud relates to the most recent scan.

Alerts and recommendations pages

When you're reviewing a recommendation or investigating an alert, the information about any potentially sensitive data involved is included on the page.

This vital additional layer of metadata helps solve the triage challenge and ensures your security team can focus its attention on the threats to sensitive data.

Inventory filters

The asset inventory page has a collection of powerful filters to group your resources with outstanding alerts and recommendations according to the criteria relevant for any scenario. These filters include **Data sensitivity classifications** and **Data sensitivity labels**. Use these filters to evaluate the security posture of resources on which Azure Purview has discovered sensitive data.

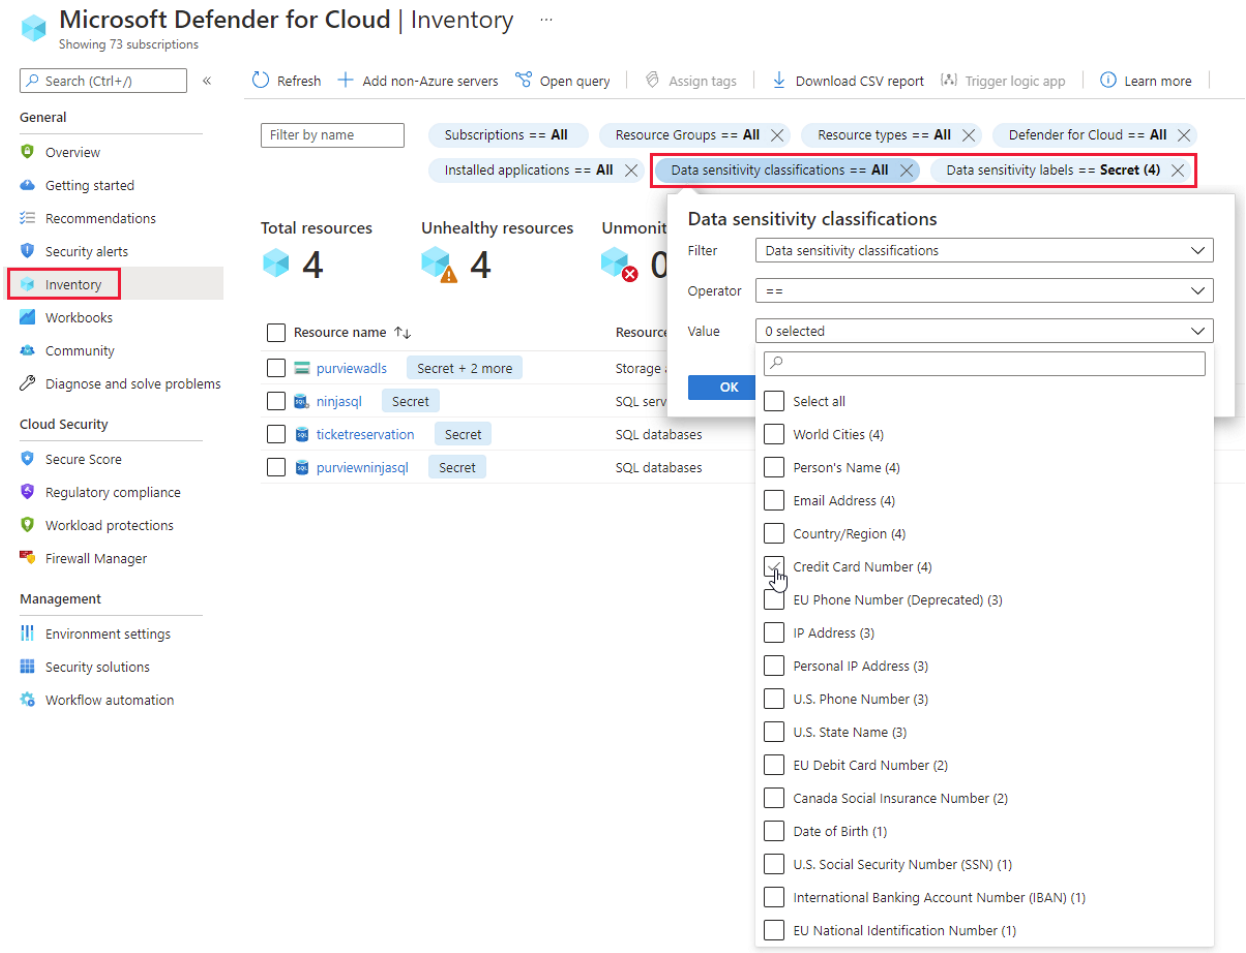

Resource health

When you select a single resource - whether from an alert, recommendation, or the inventory page - you reach a detailed health page showing a resource-centric view with the important security information related to that resource.

The resource health page provides a snapshot view of the overall health of a single resource. You can review detailed information about the resource and all recommendations that apply to that resource. Also, if you're using any of the Microsoft Defender plans, you can see outstanding security alerts for that specific resource too.

When reviewing the health of a specific resource, you'll see the Azure Purview information on this page and can use it determine what data has been discovered on this resource alongside the Azure Purview account used to scan the resource.

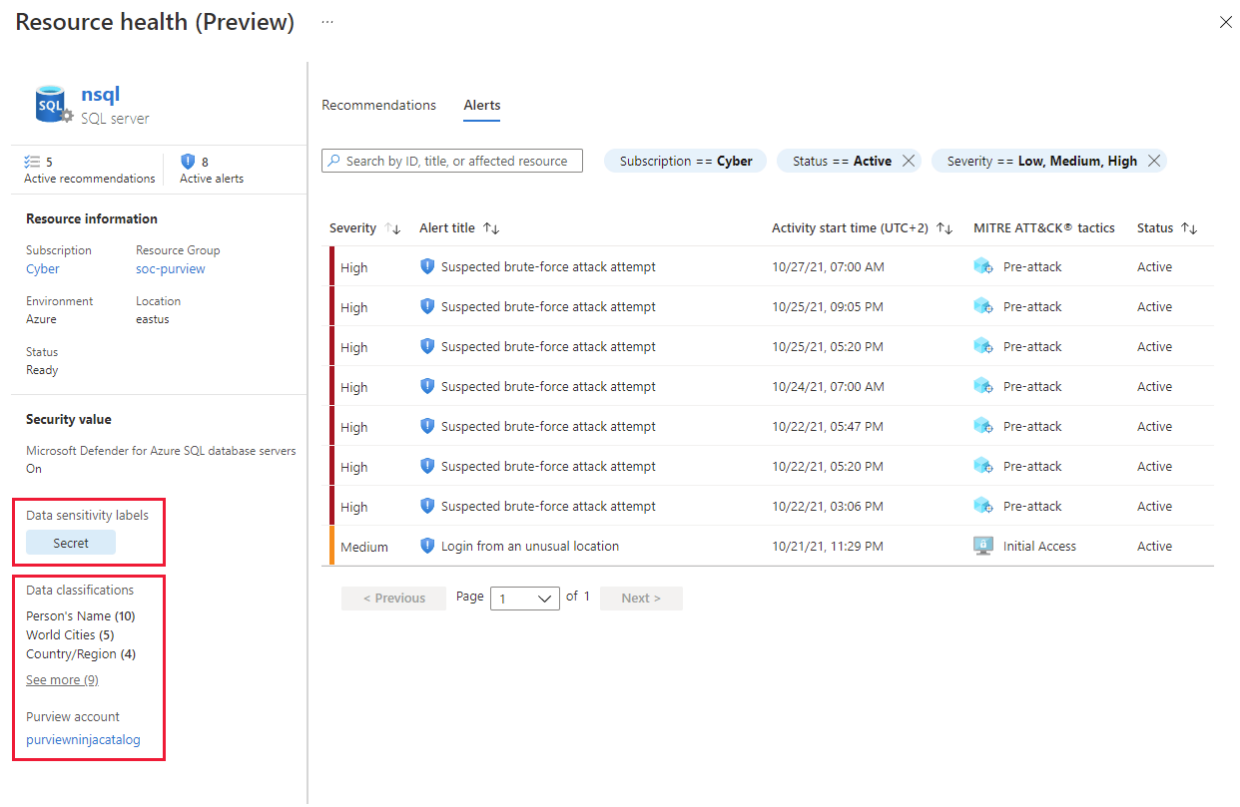

Overview tile

The dedicated **Information protection** tile in Defender for Cloud's overview dashboard shows Azure Purview's coverage. It also shows the resource types with the most sensitive data discovered.

A graph shows the number of recommendations and alerts by classified resource types. The tile also includes a link to Azure Purview to scan additional resources. Select the tile to see classified resources in Defender for Cloud's asset inventory page.

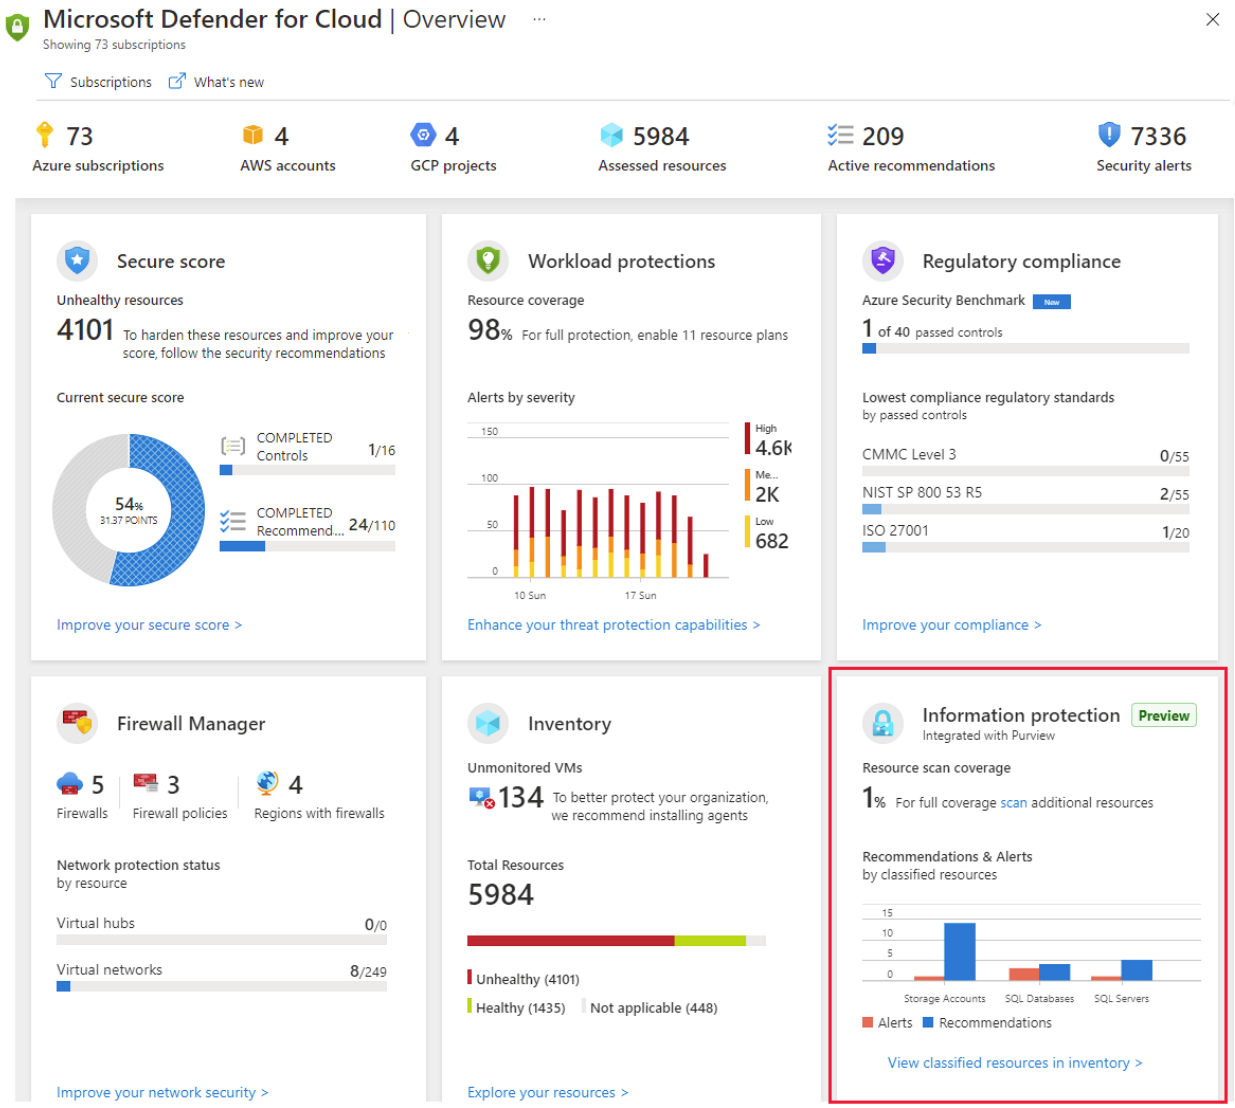

Next steps

For related information, see:

- What is Azure Purview?
- Azure Purview's supported data sources and file types and supported data stores
- Azure Purview deployment best practices
- How to label to your data in Azure Purview

Use asset inventory to manage your resources' security posture

2/15/2022 • 7 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The asset inventory page of Microsoft Defender for Cloud provides a single page for viewing the security posture of the resources you've connected to Microsoft Defender for Cloud.

Defender for Cloud periodically analyzes the security state of resources connected to your subscriptions to identify potential security vulnerabilities. It then provides you with recommendations on how to remediate those vulnerabilities.

When any resource has outstanding recommendations, they'll appear in the inventory.

Use this view and its filters to address such questions as:

- Which of my subscriptions with enhanced security features enabled have outstanding recommendations?
- Which of my machines with the tag 'Production' are missing the Log Analytics agent?
- How many of my machines tagged with a specific tag have outstanding recommendations?
- Which machines in a specific resource group have a known vulnerability (using a CVE number)?

The asset management possibilities for this tool are substantial and continue to grow.

> **TIP**
> The security recommendations on the asset inventory page are the same as those on the **Recommendations** page, but here they're shown according to the affected resource. For information about how to resolve recommendations, see Implementing security recommendations in Microsoft Defender for Cloud.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free* * Some features of the inventory page, such as the software inventory require paid solutions to be in-place
Required roles and permissions:	All users
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

ASPECT	DETAILS

What are the key features of asset inventory?

The inventory page provides the following tools:

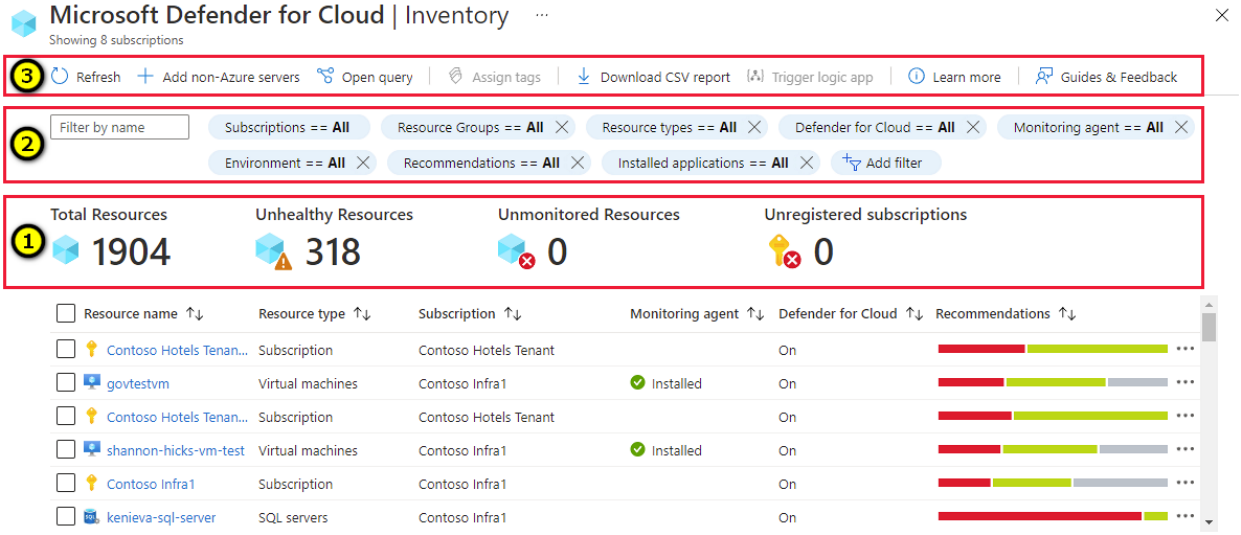

1 - Summaries

Before you define any filters, a prominent strip of values at the top of the inventory view shows:

- **Total resources**: The total number of resources connected to Defender for Cloud.
- **Unhealthy resources**: Resources with active security recommendations. Learn more about security recommendations.
- **Unmonitored resources**: Resources with agent monitoring issues - they have the Log Analytics agent deployed, but the agent isn't sending data or has other health issues.
- **Unregistered subscriptions**: Any subscription in the selected scope that haven't yet been connected to Microsoft Defender for Cloud.

2 - Filters

The multiple filters at the top of the page provide a way to quickly refine the list of resources according to the question you're trying to answer. For example, if you wanted to answer the question *Which of my machines with the tag 'Production' are missing the Log Analytics agent?* you could combine the **Agent monitoring** filter with the **Tags** filter.

As soon as you've applied filters, the summary values are updated to relate to the query results.

3 - Export and asset management tools

Export options - Inventory includes an option to export the results of your selected filter options to a CSV file. You can also export the query itself to Azure Resource Graph Explorer to further refine, save, or modify the Kusto Query Language (KQL) query.

> **TIP**
> The KQL documentation provides a database with some sample data together with some simple queries to get the "feel" for the language. Learn more in this KQL tutorial.

Asset management options - When you've found the resources that match your queries, inventory provides shortcuts for operations such as:

- Assign tags to the filtered resources - select the checkboxes alongside the resources you want to tag.
- Onboard new servers to Defender for Cloud - use the **Add non-Azure servers** toolbar button.
- Automate workloads with Azure Logic Apps - use the **Trigger Logic App** button to run a logic app on one or more resources. Your logic apps have to be prepared in advance, and accept the relevant trigger type (HTTP request). Learn more about logic apps.

How does asset inventory work?

Asset inventory utilizes Azure Resource Graph (ARG), an Azure service that provides the ability to query Defender for Cloud's security posture data across multiple subscriptions.

ARG is designed to provide efficient resource exploration with the ability to query at scale.

Using the Kusto Query Language (KQL), asset inventory can quickly produce deep insights by cross-referencing Defender for Cloud data with other resource properties.

How to use asset inventory

1. From Defender for Cloud's sidebar, select **Inventory**.
2. Use the **Filter by name** box to display a specific resource, or use the filters as described below.
3. Select the relevant options in the filters to create the specific query you want to perform.

   By default, the resources are sorted by the number of active security recommendations.

   > **IMPORTANT**
   >
   > The options in each filter are specific to the resources in the currently selected subscriptions **and** your selections in the other filters.
   >
   > For example, if you've selected only one subscription, and the subscription has no resources with outstanding security recommendations to remediate (0 unhealthy resources), the **Recommendations** filter will have no options.

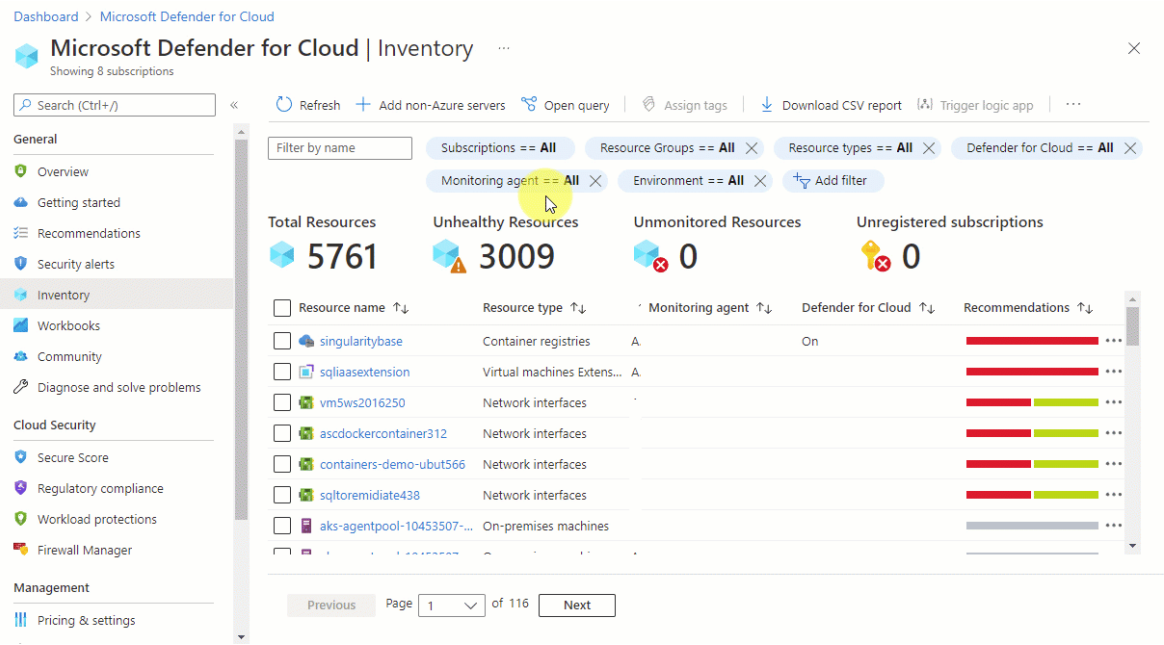

4. To use the **Security findings contain** filter, enter free text from the ID, security check, or CVE name of a vulnerability finding to filter to the affected resources:

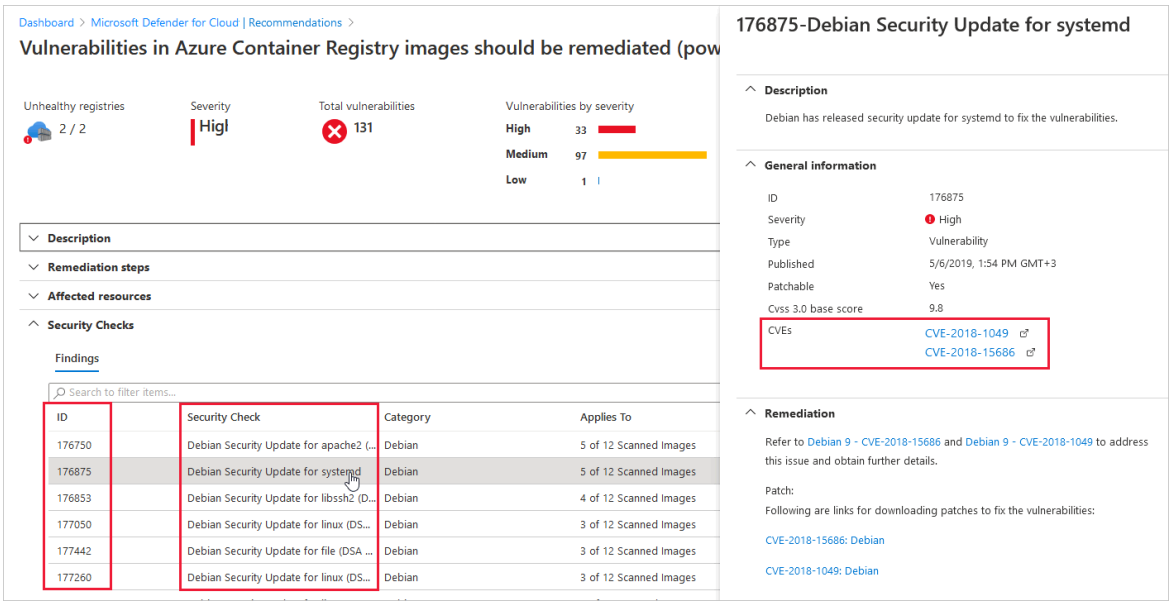

> **TIP**
>
> The **Security findings contain** and **Tags** filters only accept a single value. To filter by more than one, use **Add filters**.

5. To use the **Defender for Cloud** filter, select one or more options (Off, On, or Partial):

   - **Off** - Resources that aren't protected by a Microsoft Defender plan. You can right-click on any of these and upgrade them:

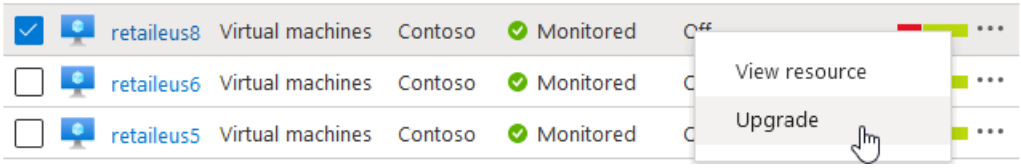

   - **On** - Resources that are protected by a Microsoft Defender plan
   - **Partial** - This applies to **subscriptions** that have some but not all of the Microsoft Defender plans disabled. For example, the following subscription has seven Microsoft Defender plans disabled.

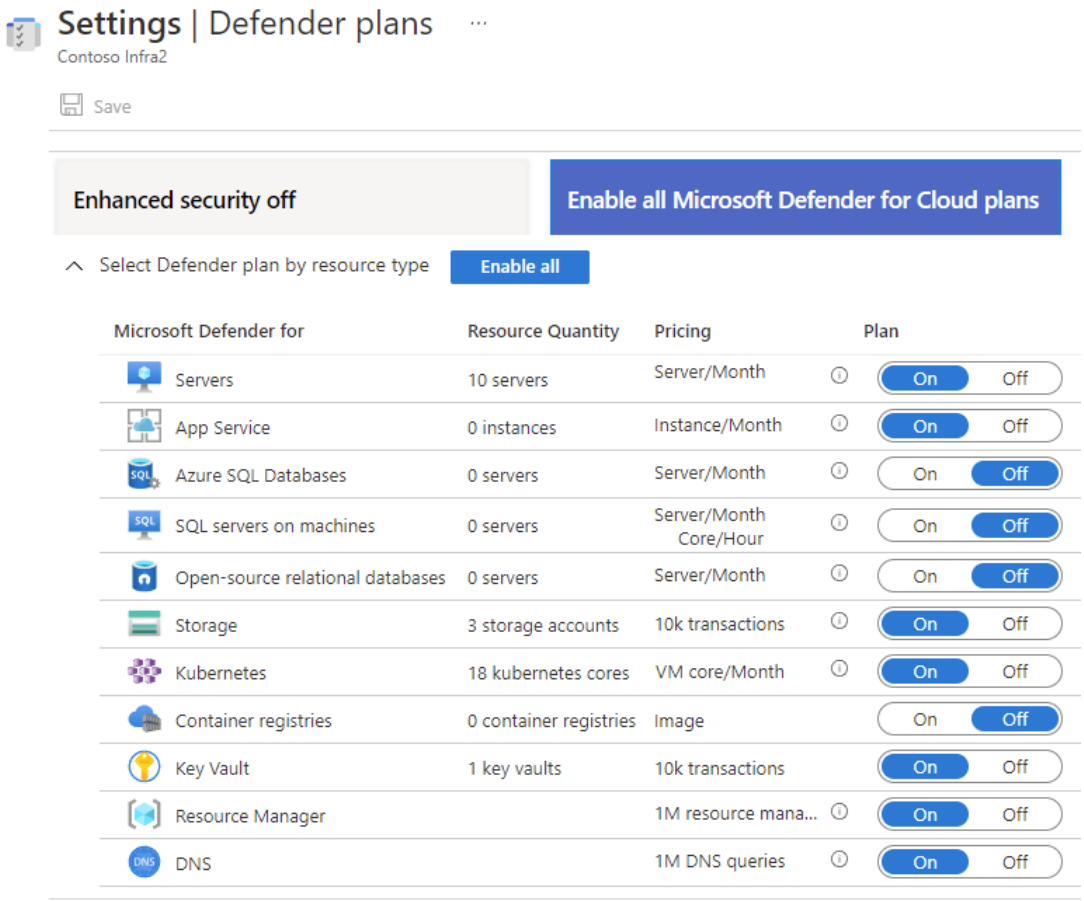

6. To further examine the results of your query, select the resources that interest you.

7. To view the current selected filter options as a query in Resource Graph Explorer, select **Open query**.

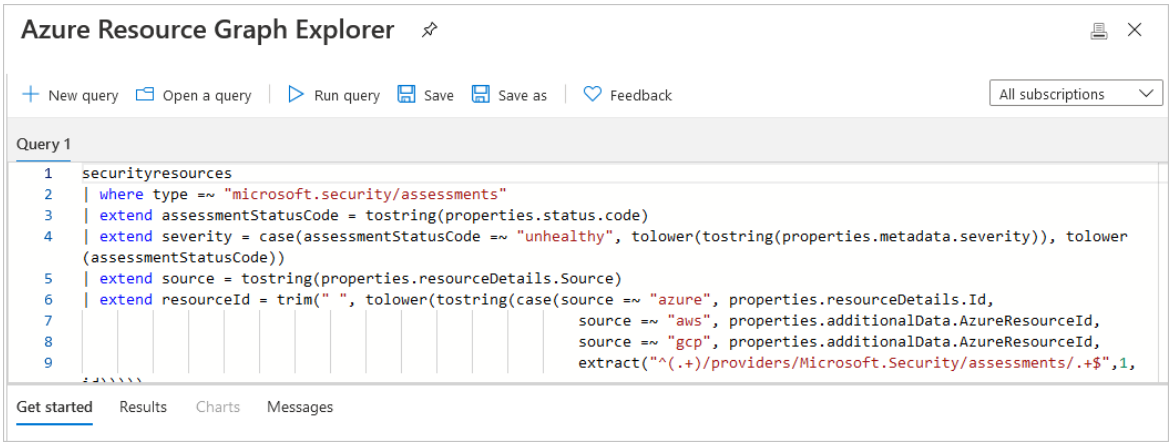

8. If you've defined some filters and left the page open, Defender for Cloud won't update the results automatically. Any changes to resources won't impact the displayed results unless you manually reload the page or select **Refresh**.

Access a software inventory

If you've enabled the integration with Microsoft Defender for Endpoint and enabled Microsoft Defender for servers, you'll have access to the software inventory.

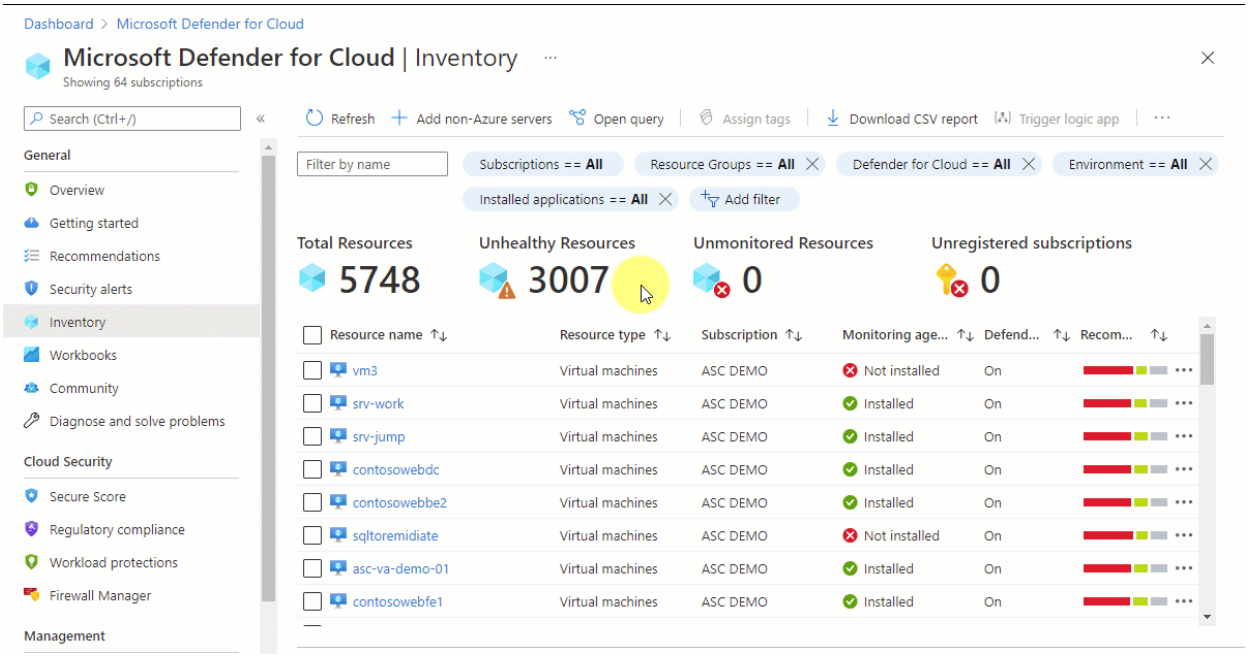

NOTE

The "Blank" option shows machines without Microsoft Defender for Endpoint (or without Microsoft Defender for servers).

As well as the filters in the asset inventory page, you can explore the software inventory data from Azure Resource Graph Explorer.

Examples of using Azure Resource Graph Explorer to access and explore software inventory data:

1. Open **Azure Resource Graph Explorer**.

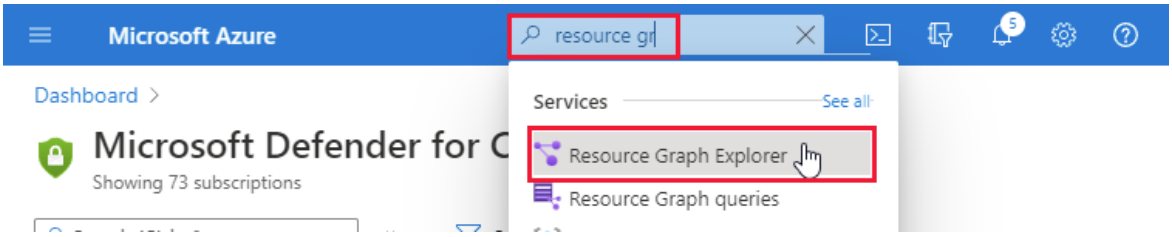

2. Select the following subscription scope: securityresources/softwareinventories
3. Enter any of the following queries (or customize them or write your own!) and select **Run query**.
   - To generate a basic list of installed software:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| project id, Vendor=properties.vendor, Software=properties.softwareName,
Version=properties.version
```

   - To filter by version numbers:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| project id, Vendor=properties.vendor, Software=properties.softwareName,
Version=tostring(properties.    version)
| where Software=="windows_server_2019" and parse_version(Version)
<=parse_version("10.0.17763.1999")
```

- To find machines with a combination of software products:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| extend vmId = properties.azureVmId
| where properties.softwareName == "apache_http_server" or properties.softwareName == "mysql"
| summarize count() by tostring(vmId)
| where count_ > 1
```

- Combination of a software product with another security recommendation:

  (In this example – machines having MySQL installed and exposed management ports)

```
securityresources
| where type == "microsoft.security/softwareinventories"
| extend vmId = tolower(properties.azureVmId)
| where properties.softwareName == "mysql"
| join (
securityresources
| where type == "microsoft.security/assessments"
| where properties.displayName == "Management ports should be closed on your virtual machines"
and properties.status.code == "Unhealthy"
| extend vmId = tolower(properties.resourceDetails.Id)
) on vmId
```

FAQ - Inventory

Why aren't all of my subscriptions, machines, storage accounts, etc. shown?

The inventory view lists your Defender for Cloud connected resources from a Cloud Security Posture Management (CSPM) perspective. The filters don't return every resource in your environment; only the ones with outstanding (or 'active') recommendations.

For example, the following screenshot shows a user with access to 8 subscriptions but only 7 currently have recommendations. So when they filter by **Resource type = Subscriptions**, only those 7 subscriptions with active recommendations appear in the inventory:

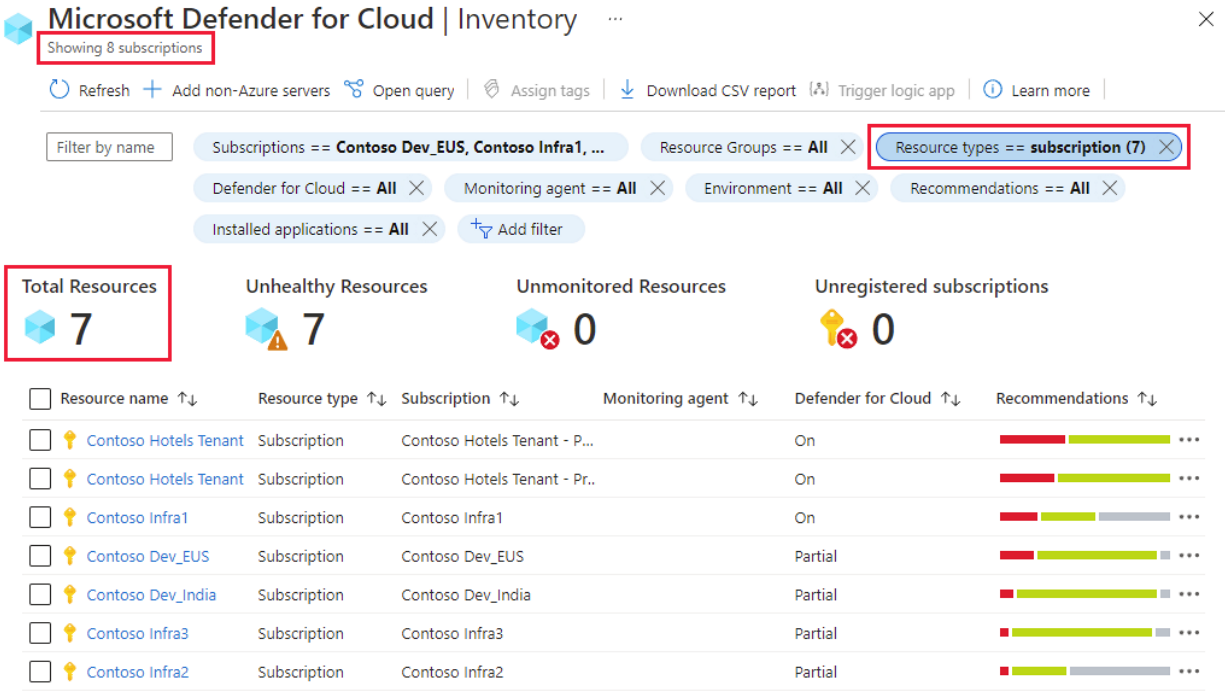

Why do some of my resources show blank values in the Defender for Cloud or monitoring agent columns?

Not all Defender for Cloud monitored resources have agents. For example, Azure Storage accounts or PaaS resources such as disks, Logic Apps, Data Lake Analysis, and Event Hub don't need agents to be monitored by Defender for Cloud.

When pricing or agent monitoring isn't relevant for a resource, nothing will be shown in those columns of inventory.

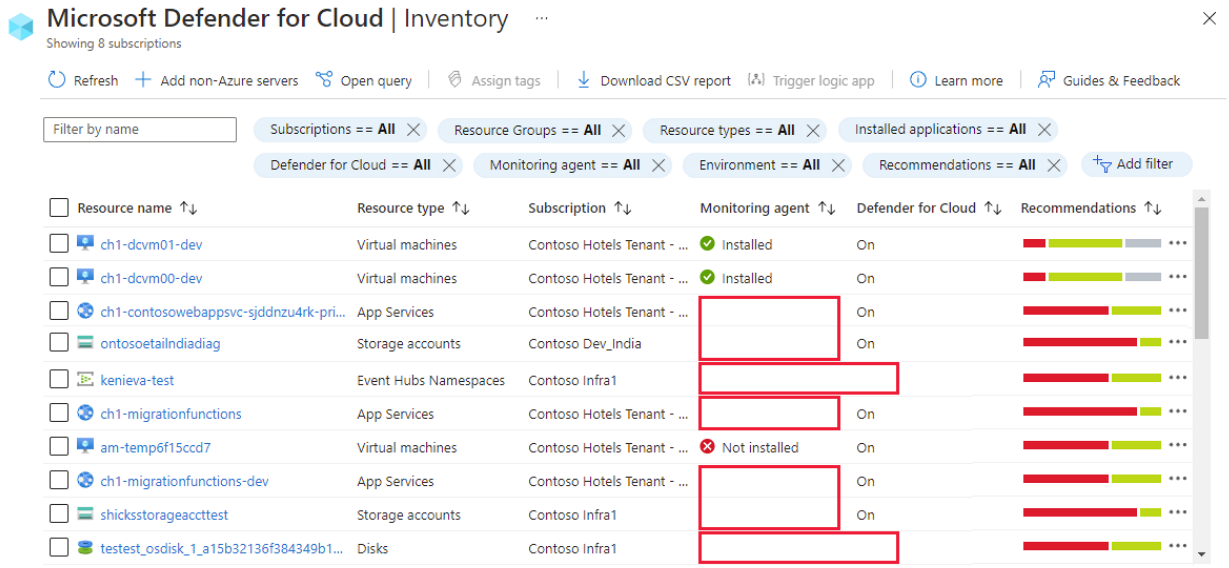

Next steps

This article described the asset inventory page of Microsoft Defender for Cloud.

For more information on related tools, see the following pages:

- Azure Resource Graph (ARG)
- Kusto Query Language (KQL)

Create rich, interactive reports of Defender for Cloud data

2/15/2022 • 7 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Monitor Workbooks provide a flexible canvas for data analysis and the creation of rich visual reports within the Azure portal. They allow you to tap into multiple data sources from across Azure, and combine them into unified interactive experiences.

Workbooks provide a rich set of capabilities for visualizing your Azure data. For detailed examples of each visualization type, see the visualizations examples and documentation.

Within Microsoft Defender for Cloud, you can access the built-in workbooks to track your organization's security posture. You can also build custom workbooks to view a wide range of data from Defender for Cloud or other supported data sources.

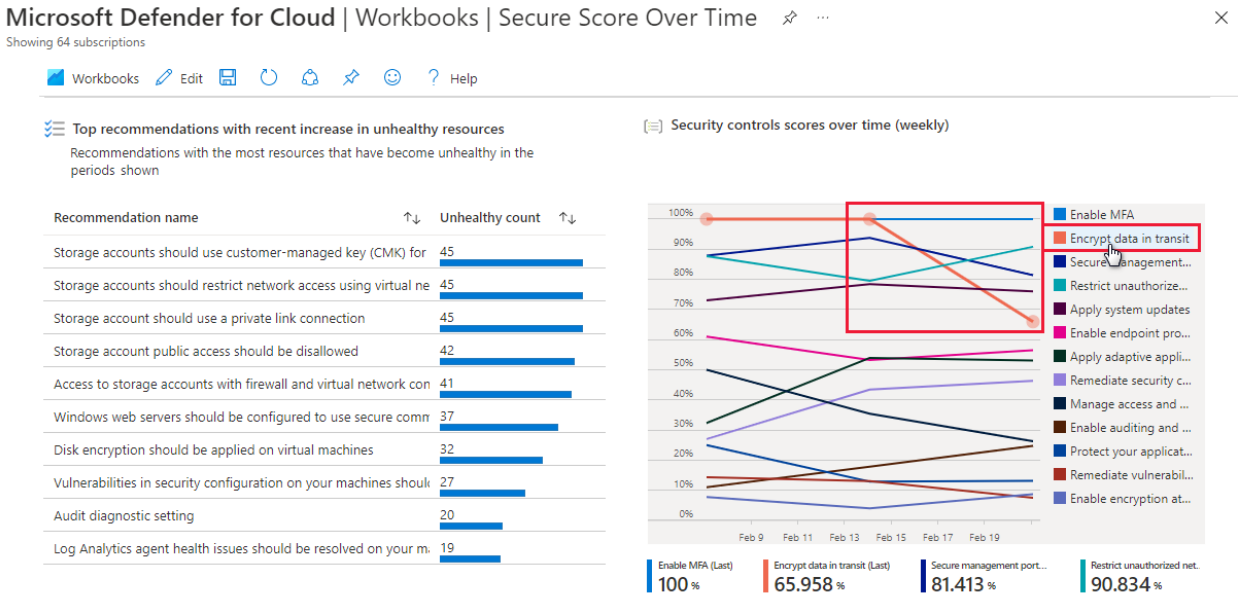

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free
Required roles and permissions:	To save workbooks, you must have at least Workbook Contributor permissions on the target resource group

ASPECT	DETAILS
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

Workbooks gallery in Microsoft Defender for Cloud

With the integrated Azure Workbooks functionality, Microsoft Defender for Cloud makes it straightforward to build your own custom, interactive workbooks. Defender for Cloud also includes a gallery with the following workbooks ready for your customization:

- 'Secure Score Over Time' workbook - Track your subscriptions' scores and changes to recommendations for your resources
- 'System Updates' workbook - View missing system updates by resources, OS, severity, and more
- 'Vulnerability Assessment Findings' workbook - View the findings of vulnerability scans of your Azure resources
- 'Compliance Over Time' workbook - View the status of a subscription's compliance with the regulatory or industry standards you've selected
- 'Active Alerts' workbook - view active alerts by severity, type, tag, MITRE ATT&CK tactics, and location.

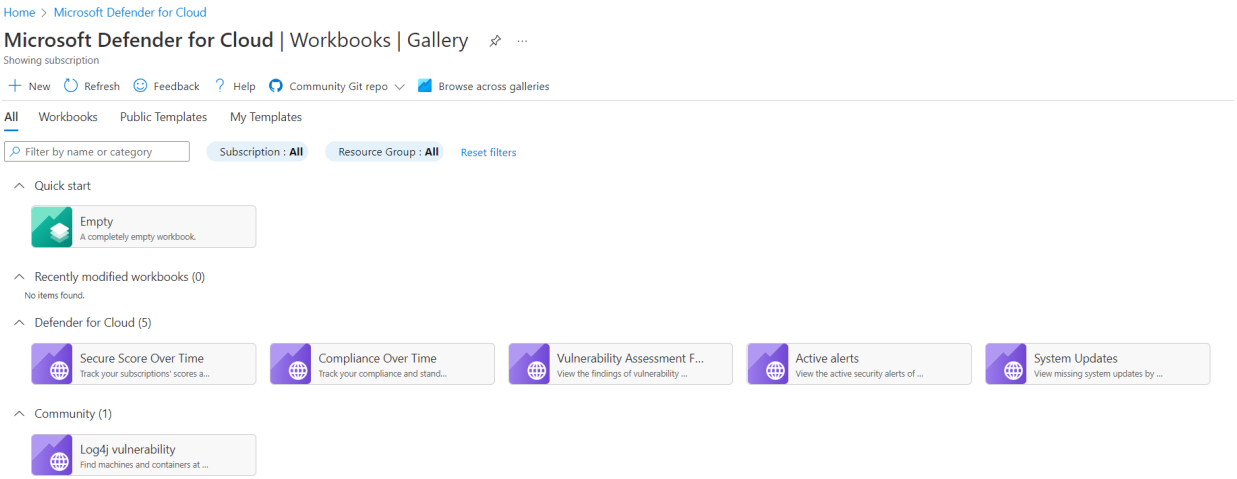

Choose one of the supplied workbooks or create your own.

TIP

Use the **Edit** button to customize any of the supplied workbooks to your satisfaction. When you're done editing, select **Save** and your changes will be saved to a new workbook.

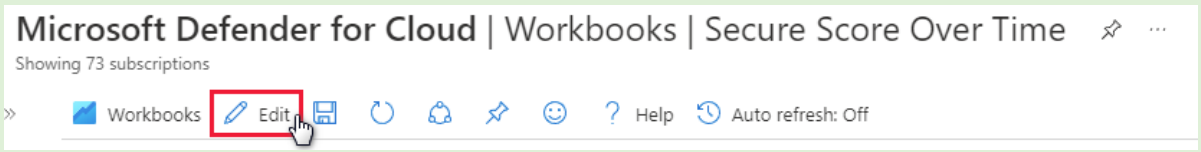

Use the 'Secure Score Over Time' workbook

This workbook uses secure score data from your Log Analytics workspace. That data needs to be exported from the continuous export tool as described in Configure continuous export from the Defender for Cloud pages in Azure portal.

When you set up the continuous export, set the export frequency to both **streaming updates** and **snapshots**.

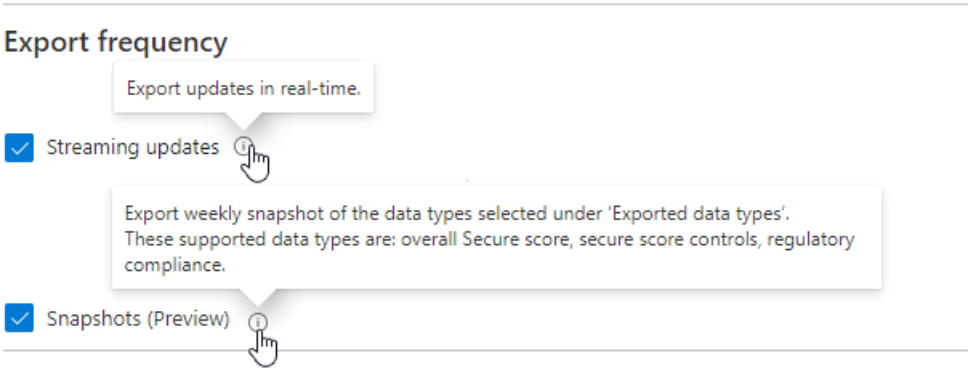

NOTE

Snapshots get exported weekly, so you'll need to wait at least one week for the first snapshot to be exported before you can view data in this workbook.

TIP

To configure continuous export across your organization, use the supplied Azure Policy 'DeployIfNotExist' policies described in Configure continuous export at scale.

The secure score over time workbook has five graphs for the subscriptions reporting to the selected workspaces:

GRAPH	EXAMPLE
Score trends for the last week and month Use this section to monitor the current score and general trends of the scores for your subscriptions.	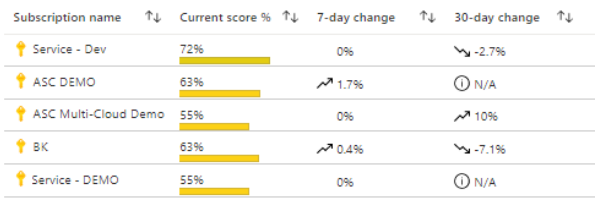
Aggregated score for all selected subscriptions Hover your mouse over any point in the trend line to see the aggregated score at any date in the selected time range.	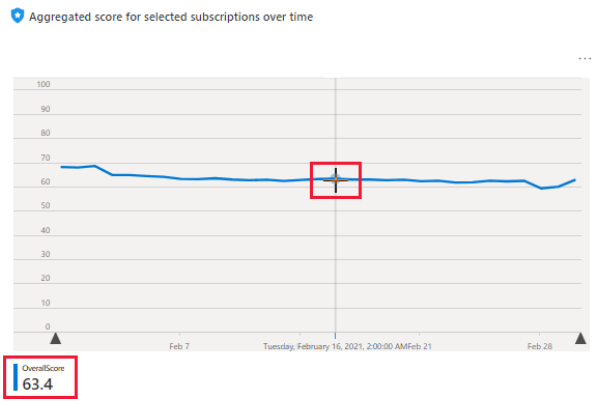

GRAPH	EXAMPLE
Recommendations with the most unhealthy resources This table helps you triage the recommendations that have had the most resources changed to unhealthy over the selected period.	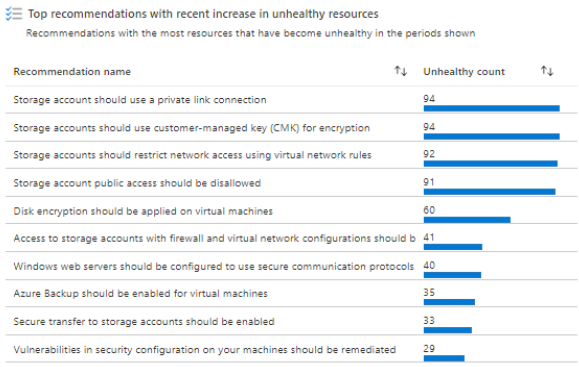
Scores for specific security controls Defender for Cloud's security controls are logical groupings of recommendations. This chart shows you, at a glance, the weekly scores for all of your controls.	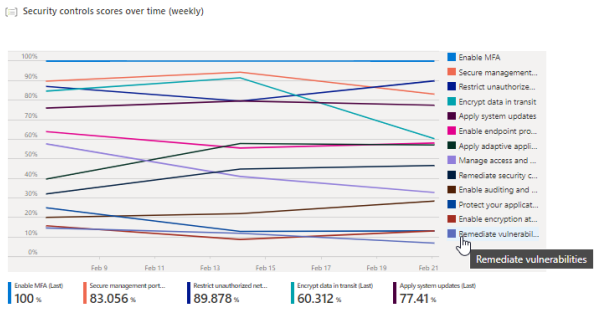
Resources changes Recommendations with the most resources that have changed state (healthy, unhealthy, or not applicable) during the selected period are listed here. Select any recommendation from the list to open a new table listing the specific resources.	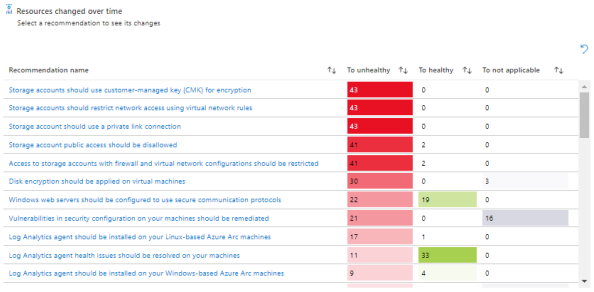

Use the 'System Updates' workbook

This workbook is based on the security recommendation "System updates should be installed on your machines".

The workbook helps you identify machines with outstanding updates.

You can view the situation for the selected subscriptions according to:

- The list of resources with outstanding updates
- The list of updates missing from your resources

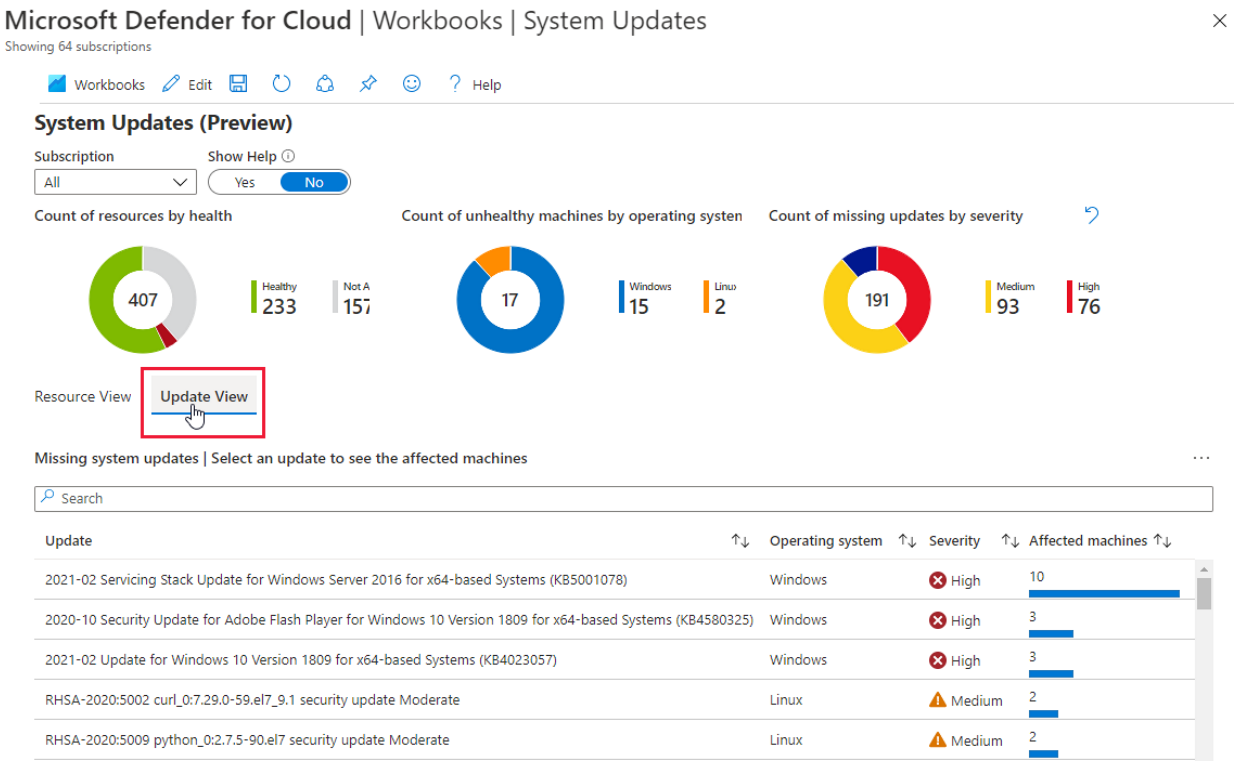

Use the 'Vulnerability Assessment Findings' workbook

Defender for Cloud includes vulnerability scanners for your machines, containers in container registries, and SQL servers.

Learn more about using these scanners:

- Find vulnerabilities with Microsoft threat and vulnerability management
- Find vulnerabilities with the integrated Qualys scanner
- Scan your registry images for vulnerabilities
- Scan your SQL resources for vulnerabilities

Findings for each resource type are reported in separate recommendations:

- Vulnerabilities in your virtual machines should be remediated (includes findings from Microsoft threat and vulnerability management, the integrated Qualys scanner, and any configured BYOL VA solutions)
- Container registry images should have vulnerability findings resolved
- SQL databases should have vulnerability findings resolved
- SQL servers on machines should have vulnerability findings resolved

This workbook gathers these findings and organizes them by severity, resource type, and category.

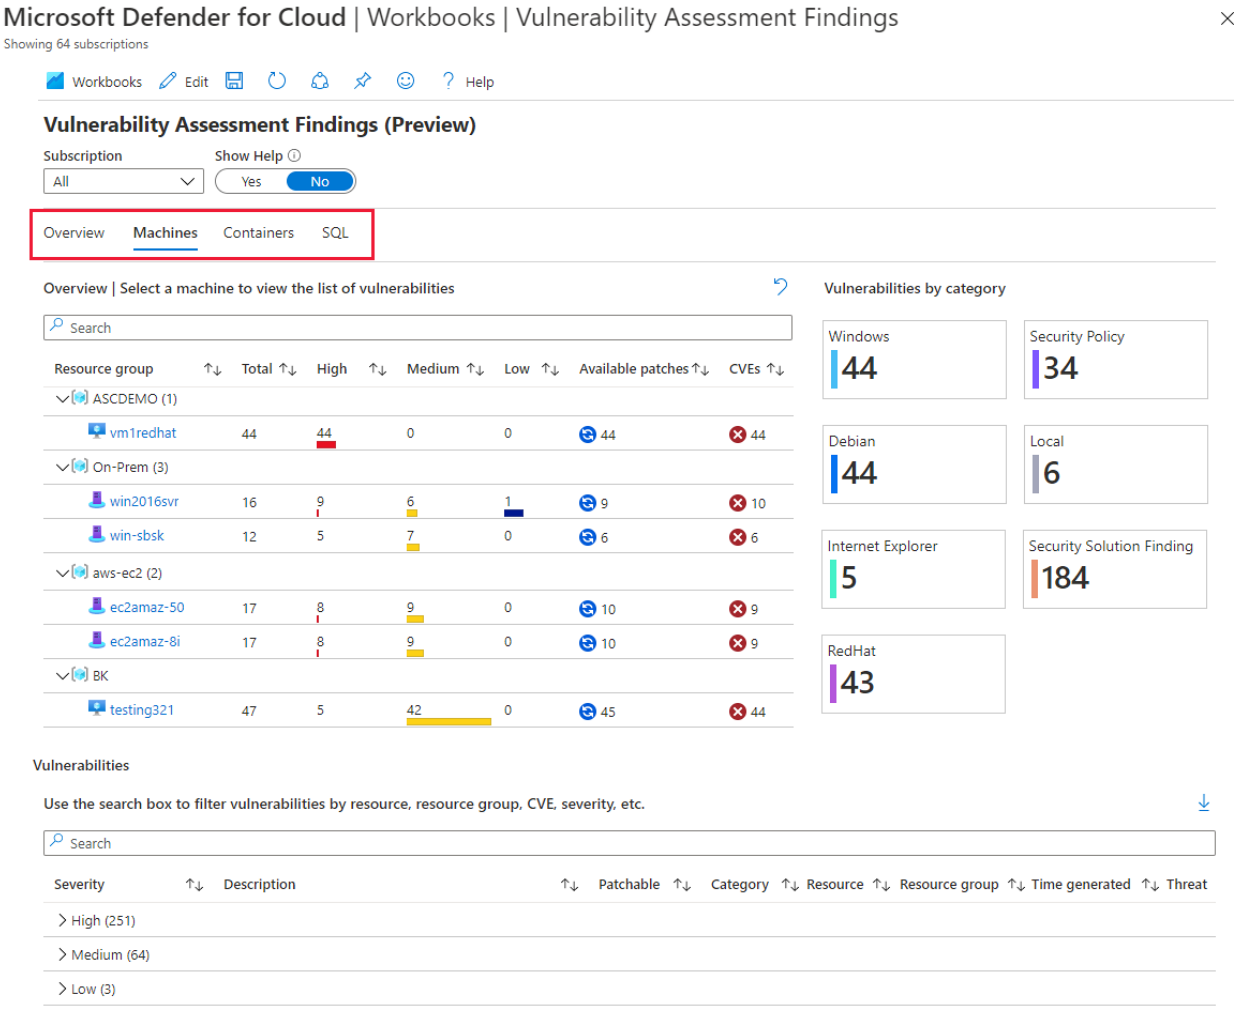

Use the 'Compliance Over Time' workbook

Microsoft Defender for Cloud continually compares the configuration of your resources with requirements in industry standards, regulations, and benchmarks. Built-in standards include NIST SP 800-53, SWIFT CSP CSCF v2020, Canada Federal PBMM, HIPAA HITRUST, and more. You can select the specific standards relevant to your organization using the regulatory compliance dashboard. Learn more in Customize the set of standards in your regulatory compliance dashboard.

This workbook tracks your compliance status over time with the various standards you've added to your dashboard.

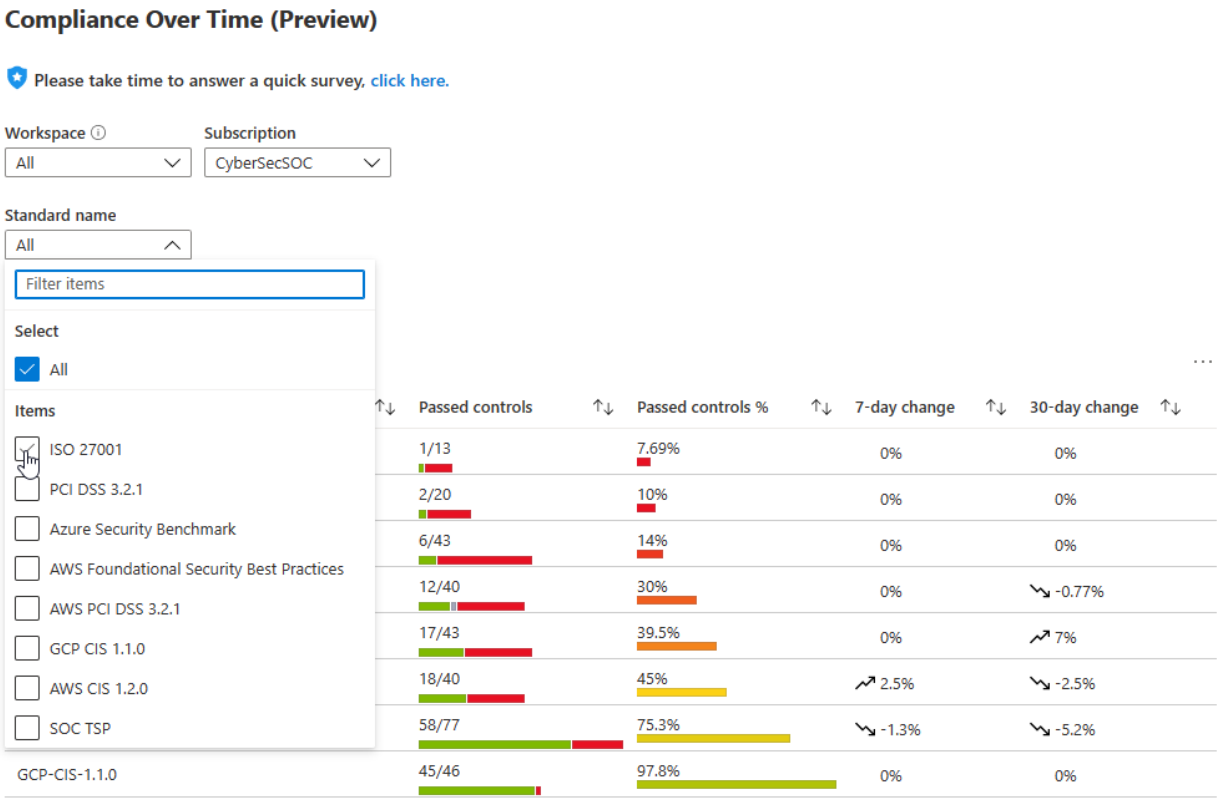

When you select a standard from the overview area of the report, the lower pane reveals a more detailed breakdown:

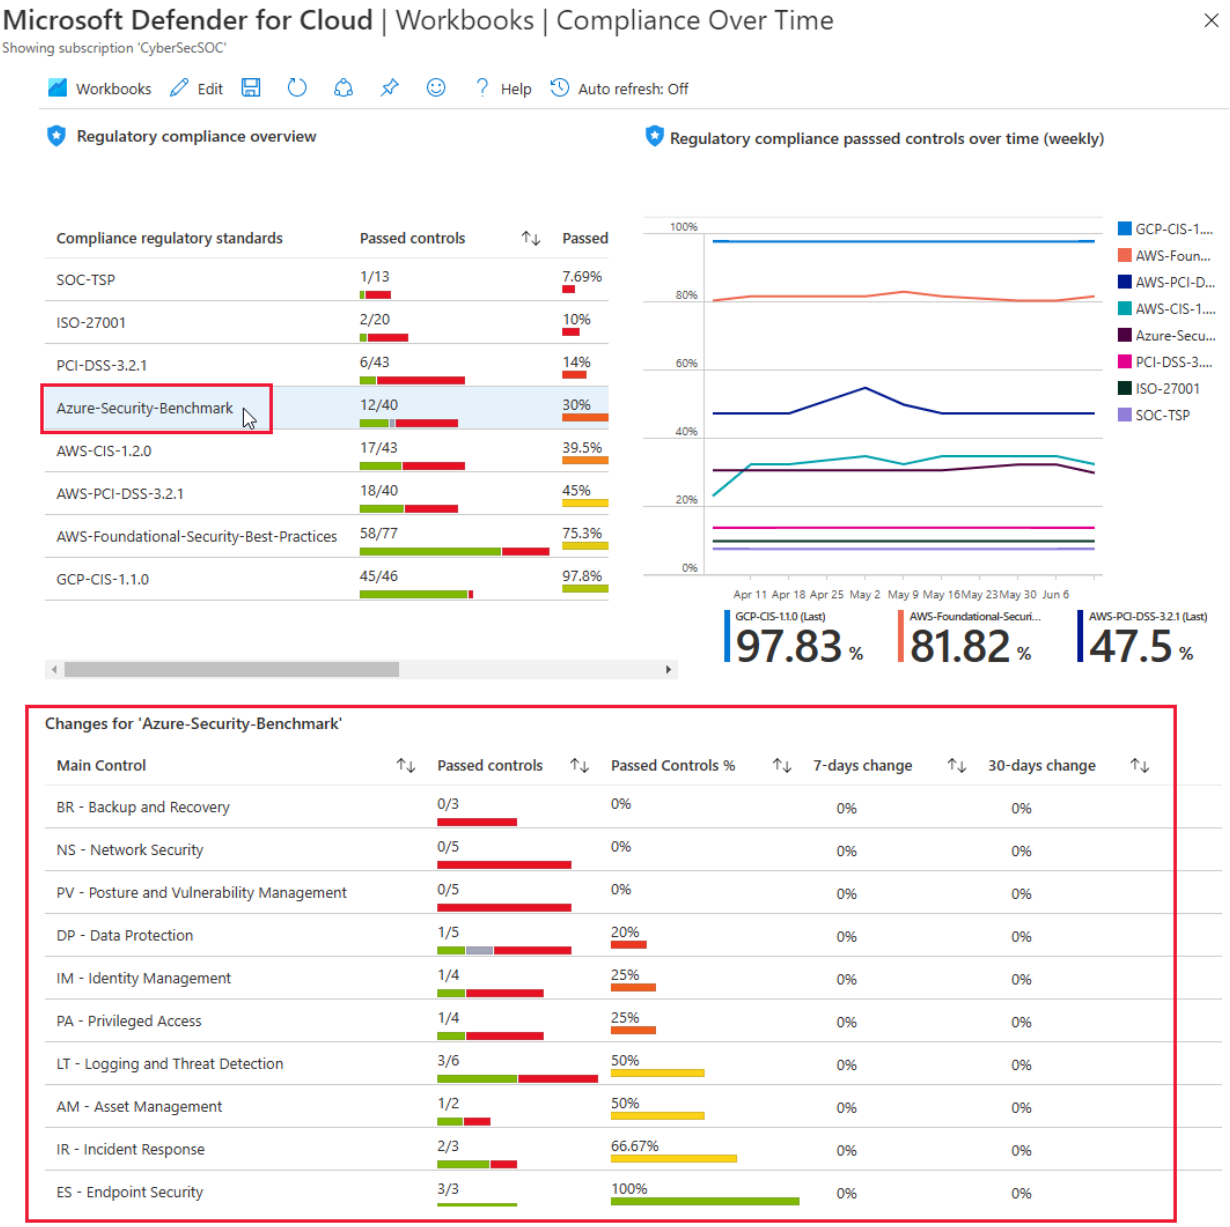

Changes for 'Azure-Security-Benchmark'

Main Control	Passed controls	Passed Controls %	7-days change	30-days change
BR - Backup and Recovery	0/3	0%	0%	0%
NS - Network Security	0/5	0%	0%	0%
PV - Posture and Vulnerability Management	0/5	0%	0%	0%
DP - Data Protection	1/5	20%	0%	0%
IM - Identity Management	1/4	25%	0%	0%
PA - Privileged Access	1/4	25%	0%	0%
LT - Logging and Threat Detection	3/6	50%	0%	0%
AM - Asset Management	1/2	50%	0%	0%
IR - Incident Response	2/3	66.67%	0%	0%
ES - Endpoint Security	3/3	100%	0%	0%

You can keep drilling down - right down to the recommendation level - to view the resources that have passed or failed each control.

TIP

For each panel of the report, you can export the data to Excel with the "Export to Excel" option.

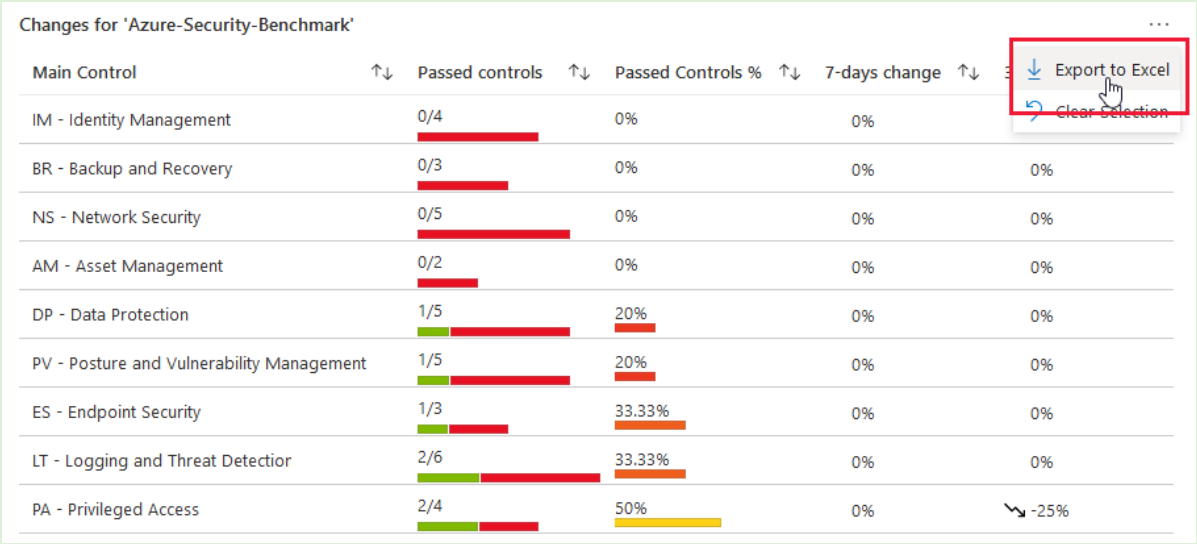

Main Control	Passed controls	Passed Controls %	7-days change	
IM - Identity Management	0/4	0%	0%	
BR - Backup and Recovery	0/3	0%	0%	0%
NS - Network Security	0/5	0%	0%	0%
AM - Asset Management	0/2	0%	0%	0%
DP - Data Protection	1/5	20%	0%	0%
PV - Posture and Vulnerability Management	1/5	20%	0%	0%
ES - Endpoint Security	1/3	33.33%	0%	0%
LT - Logging and Threat Detectior	2/6	33.33%	0%	0%
PA - Privileged Access	2/4	50%	0%	-25%

Use the 'Active Alerts' workbook

This workbook displays the active security alerts for your subscriptions on one dashboard. Security alerts are the notifications that Defender for Cloud generates when it detects threats on your resources. Defender for Cloud prioritizes, and lists the alerts, along with information needed for quick investigation and remediation.

This workbook benefits you by letting you understand the active threats on your environment, and allows you to prioritize between the active alerts.

NOTE

Most workbooks use Azure Resource Graph (ARG) to query their data. For example, to display the Map View, Log Analytics workspace is used to query the data. Continuous export should be enabled, and export the security alerts to the Log Analytics workspace.

You can view the active alerts by severity, resource group, or tag.

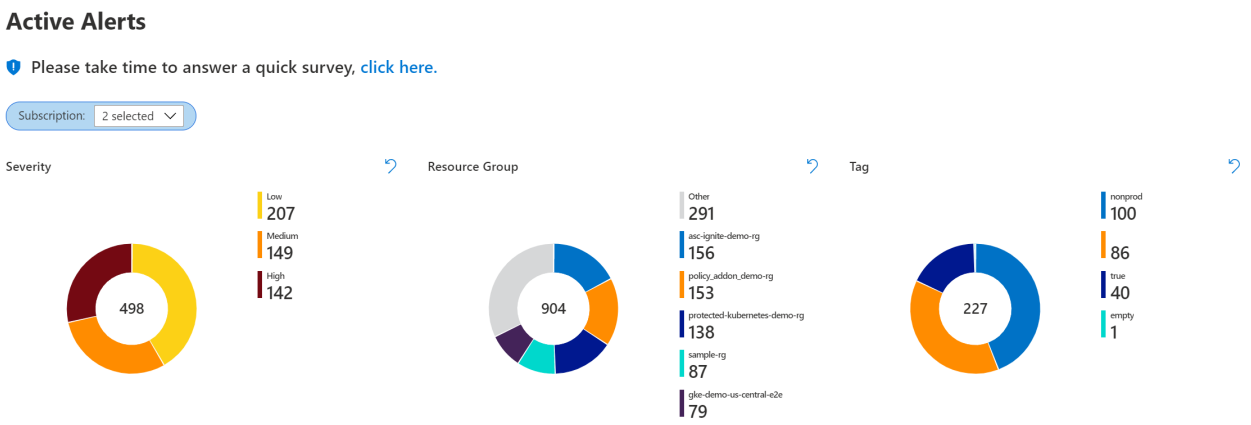

You can also view your subscription's top alerts by attacked resources, alert types, and new alerts.

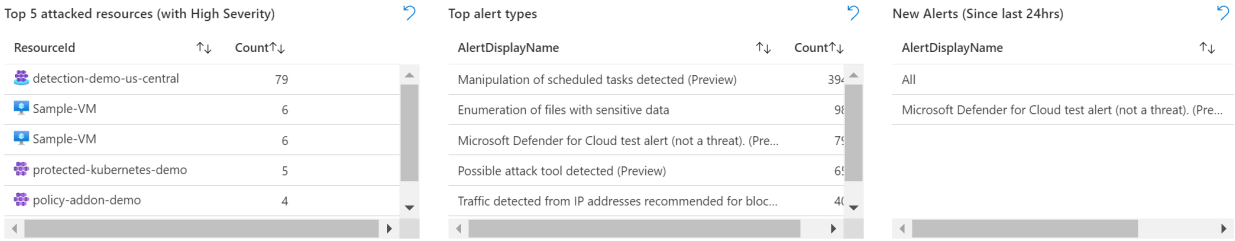

You can get more details on any of these alerts by selecting it.

The MITRE ATT&CK tactics displays by the order of the kill-chain, and the number of alerts the subscription has at each stage.

MITRE ATT&CK tactics

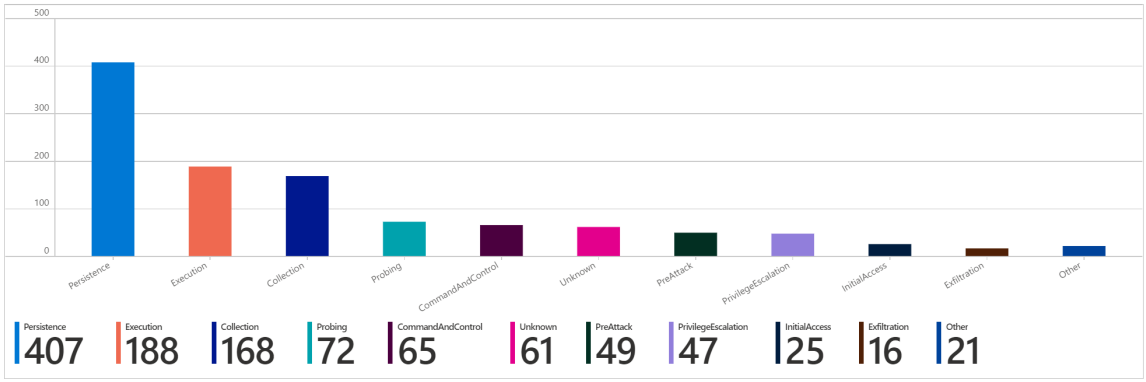

You can see all of the active alerts in a table with the ability to filter by columns. By selecting an alert, the alert view button appears.

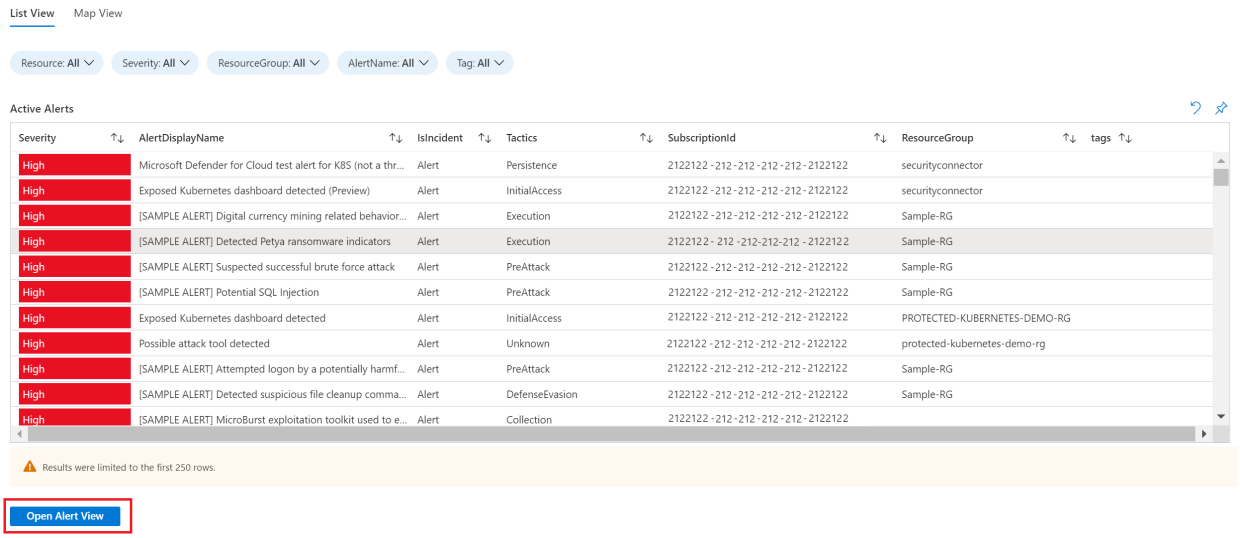

By selecting the Open Alert View button, you can see all the details of that specific alert.

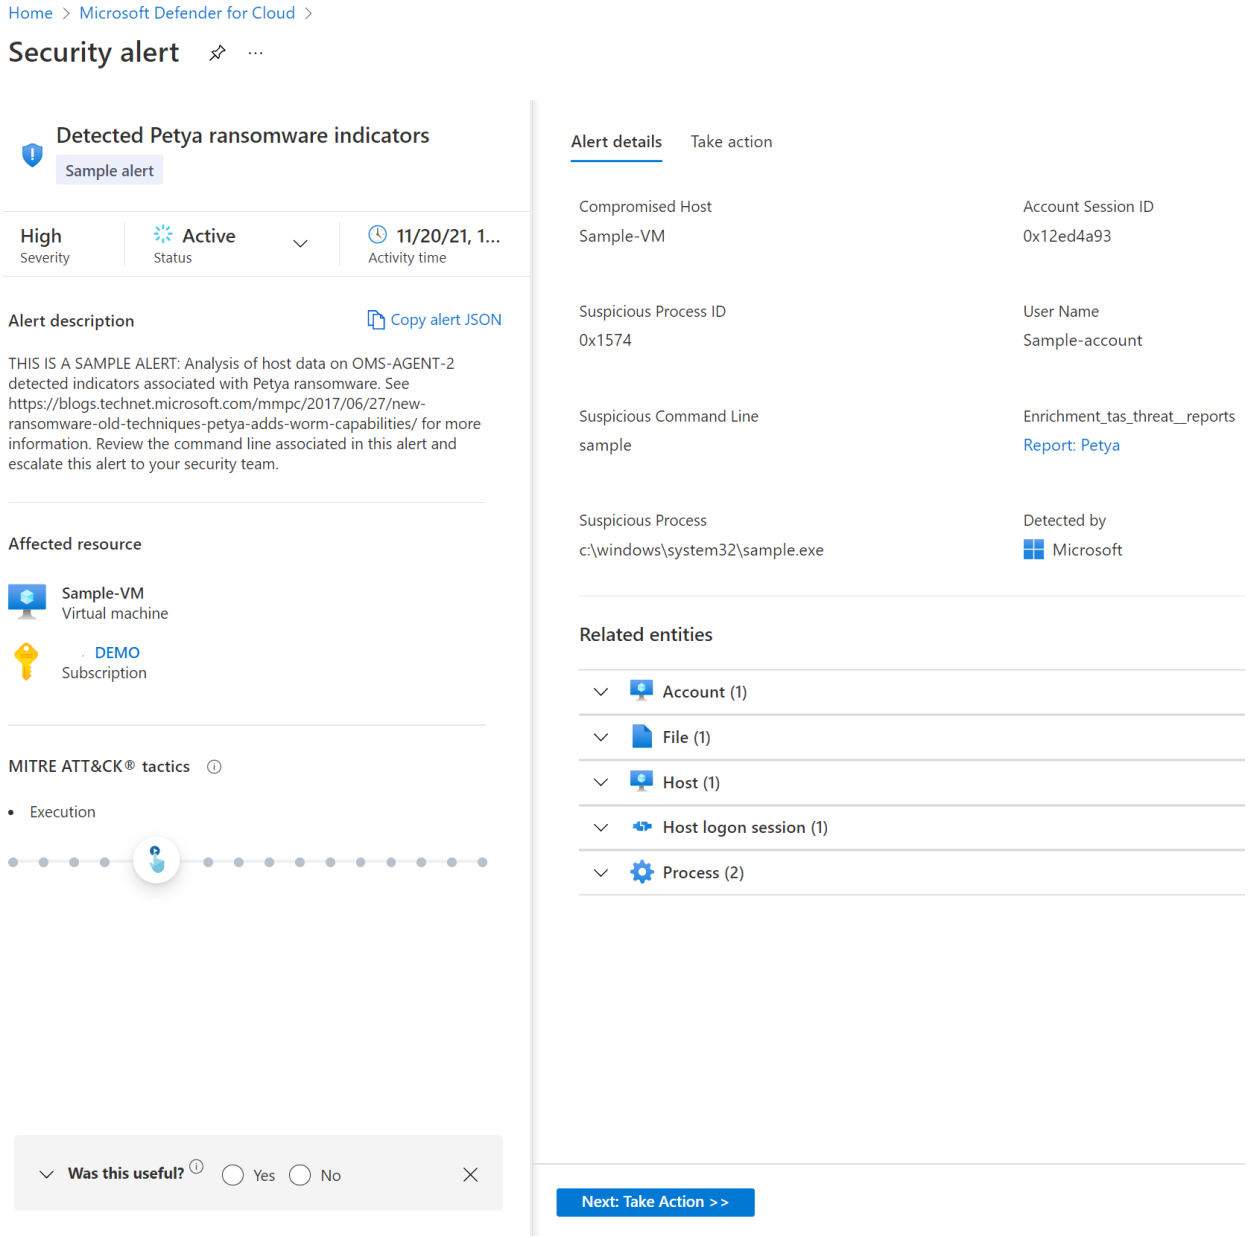

By selecting Map View, you can also see all alerts based on their location.

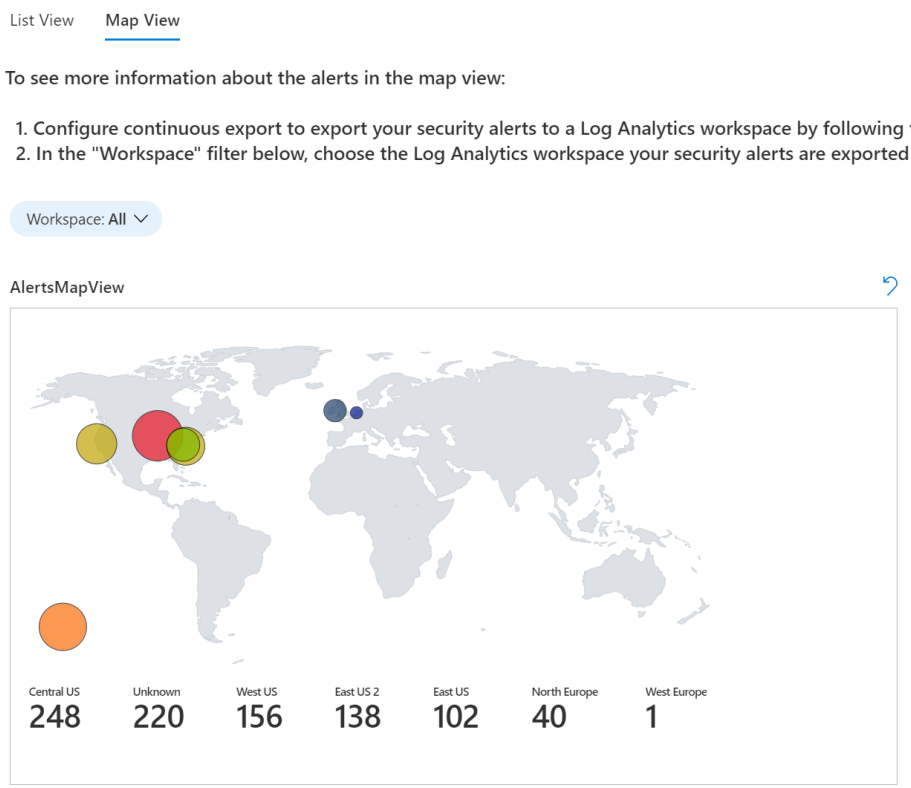

By selecting a location on the map you will be able to view all of the alerts for that location.

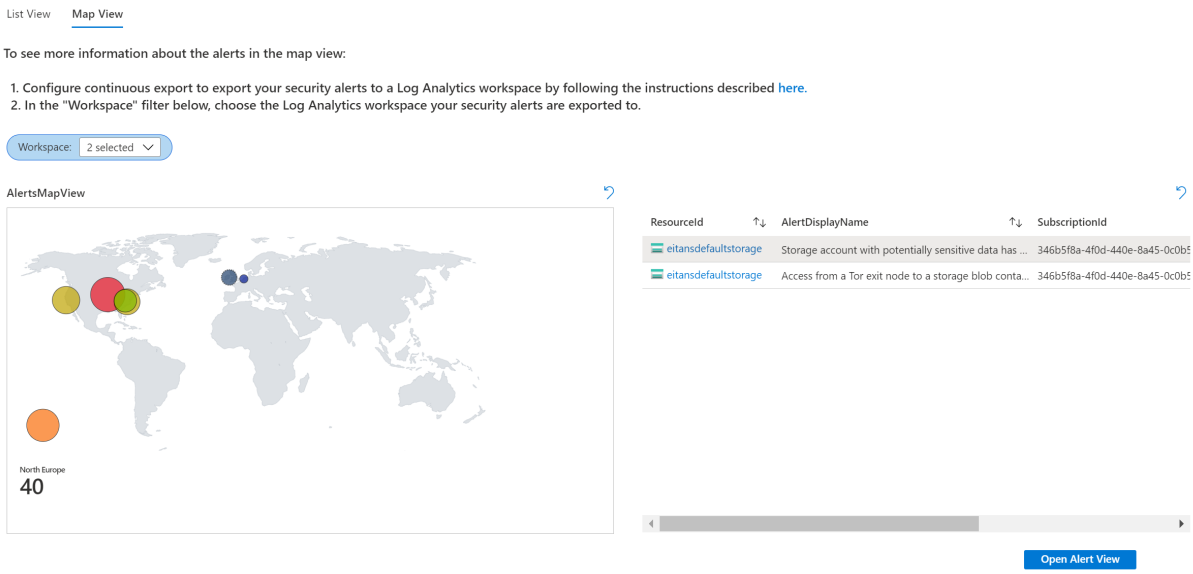

You can see the details for that alert with the Open Alert View button.

Import workbooks from other workbook galleries

If you've built workbooks in other Azure services and want to move them into your Microsoft Defender for Cloud workbooks gallery:

1. Open the target workbook.
2. From the toolbar, select **Edit**.

3. From the toolbar, select </> to enter the Advanced Editor.

4. Copy the workbook's Gallery Template JSON.
5. Open the workbooks gallery in Defender for Cloud and from the menu bar select **New**.
6. Select the **</>** to enter the Advanced Editor.
7. Paste in the entire Gallery Template JSON.
8. Select **Apply**.
9. From the toolbar, select **Save As**.

10. Enter the required details for saving the workbook:
    a. A name for the workbook
    b. The desired region
    c. Subscription, resource group, and sharing as appropriate.

You'll find your saved workbook in the **Recently modified workbooks** category.

Next steps

This article described Defender for Cloud's integrated Azure Monitor Workbooks page with built-in reports and the option to build your own custom, interactive reports.

- Learn more about Azure Monitor Workbooks
- The built-in workbooks pull their data from Defender for Cloud's recommendations. Learn about the many security recommendations in Security recommendations - a reference guide

Review your security recommendations

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This topic explains how to view and understand the recommendations in Microsoft Defender for Cloud to help you protect your Azure resources.

Monitor recommendations

Defender for Cloud analyzes the security state of your resources to identify potential vulnerabilities.

1. From Defender for Cloud's menu, open the **Recommendations** page to see the recommendations applicable to your environment. Recommendations are grouped into security controls.

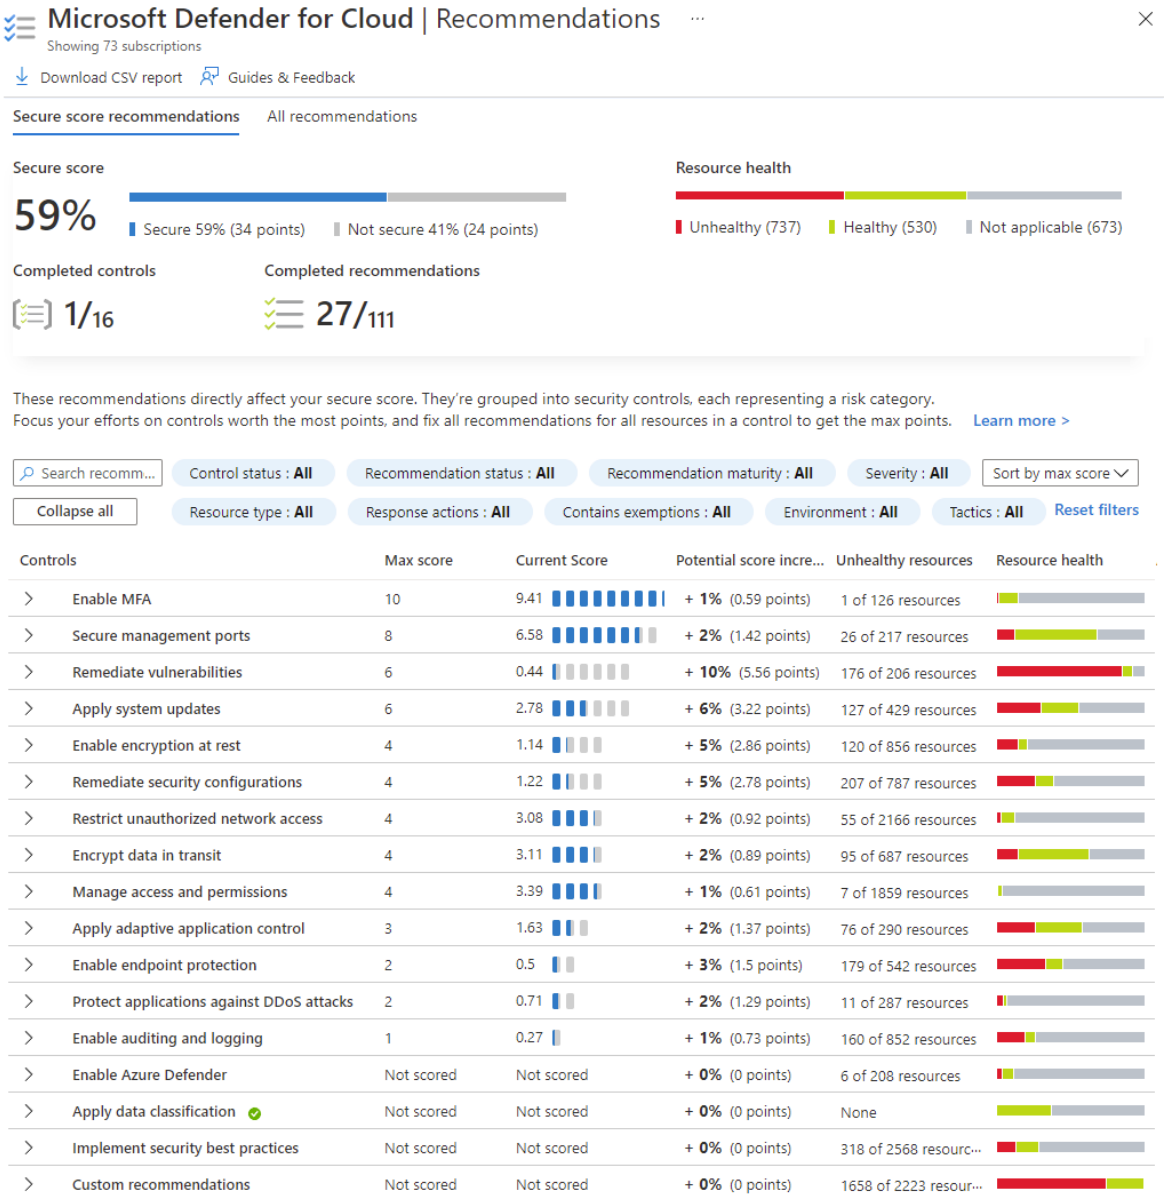

2. To find recommendations specific to the resource type, severity, environment, or other criteria that are important to you, use the optional filters above the list of recommendations.

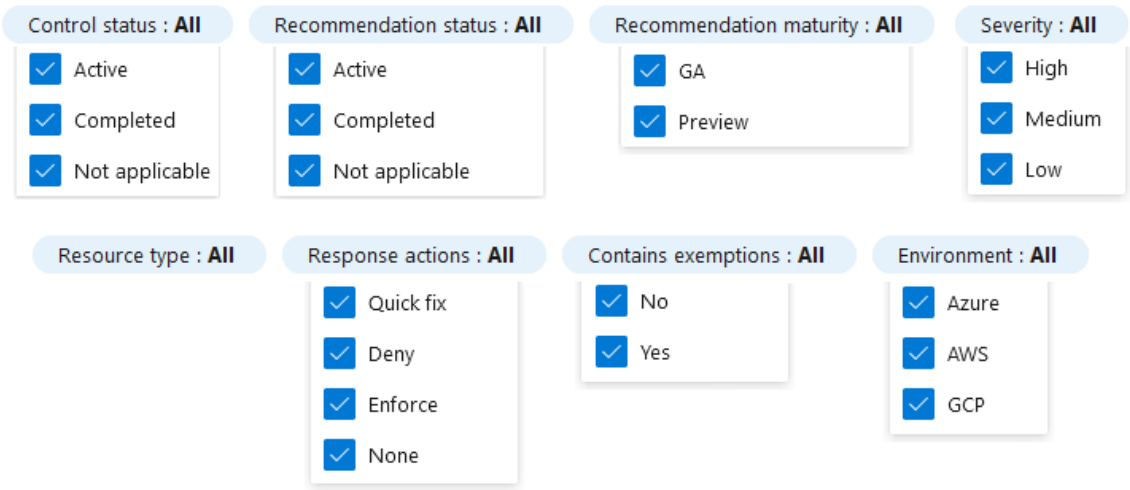

3. Expand a control and select a specific recommendation to view the recommendation details page.

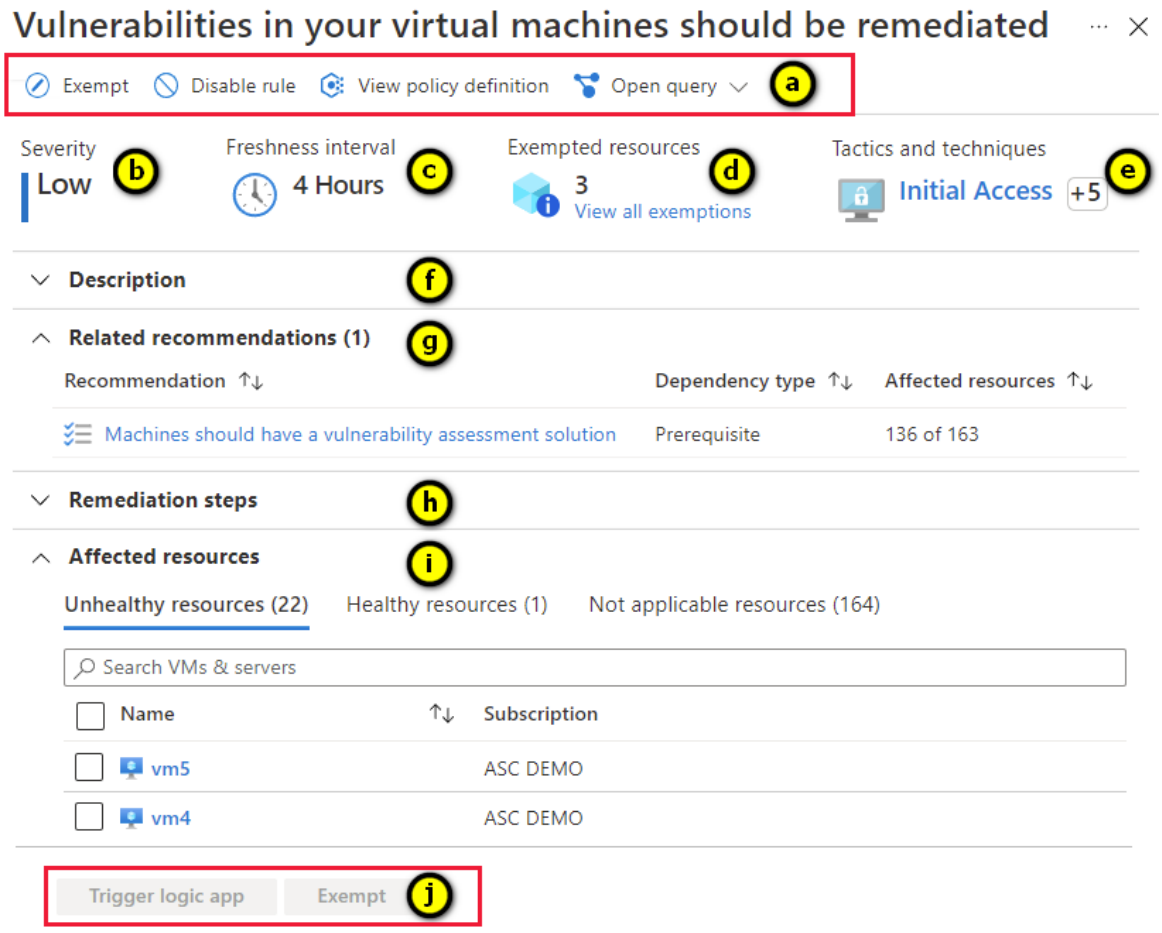

The page includes:

a. For supported recommendations, the top toolbar shows any or all of the following buttons:

- **Enforce** and **Deny** (see Prevent misconfigurations with Enforce/Deny recommendations).
- **View policy definition** to go directly to the Azure Policy entry for the underlying policy.
- **Open query** - All recommendations have the option to view the detailed information about the affected resources using Azure Resource Graph Explorer.

b. **Severity indicator**.

c. **Freshness interval** (where relevant).

d. **Count of exempted resources** if exemptions exist for a recommendation, this shows the number of resources that have been exempted with a link to view the specific resources.

e. **Mapping to MITRE ATT&CK ® tactics and techniques** if a recommendation has defined tactics and techniques, select the icon for links to the relevant pages on MITRE's site.

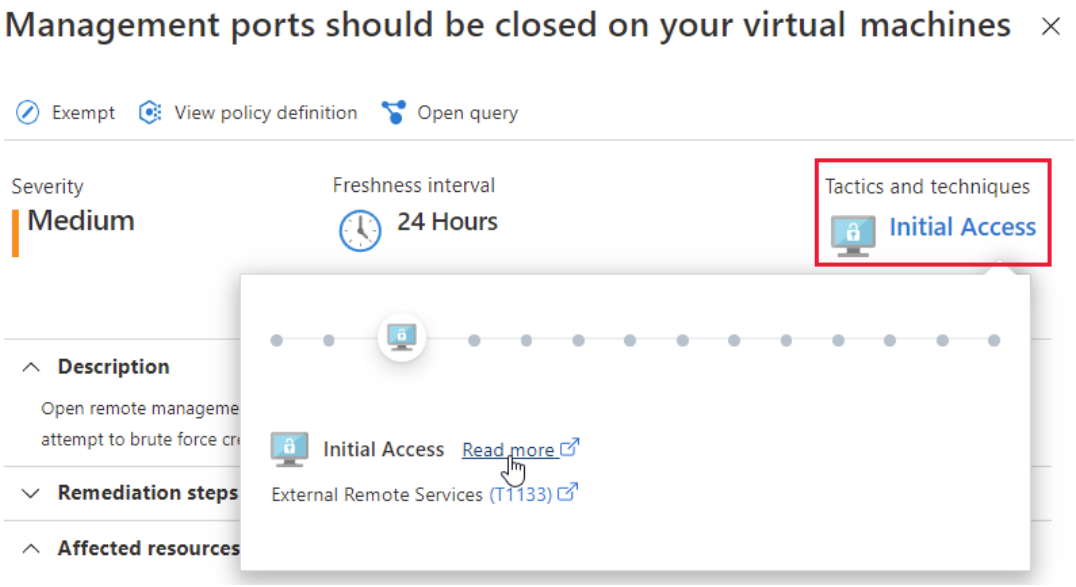

f. **Description** - A short description of the security issue.

g. When relevant, the details page also includes a table of **related recommendations**:

The relationship types are:

- **Prerequisite** - A recommendation that must be completed before the selected recommendation
- **Alternative** - A different recommendation which provides another way of achieving the goals of the selected recommendation
- **Dependent** - A recommendation for which the selected recommendation is a prerequisite

For each related recommendation, the number of unhealthy resources is shown in the "Affected resources" column.

> **TIP**
> If a related recommendation is grayed out, its dependency isn't yet completed and so isn't available.

h. **Remediation steps** - A description of the manual steps required to remediate the security issue on the affected resources. For recommendations with the **Fix** option**, you can select **View remediation logic** before applying the suggested fix to your resources.

i. **Affected resources** - Your resources are grouped into tabs:

- **Healthy resources** – Relevant resources which either aren't impacted or on which you've already remediated the issue.
- **Unhealthy resources** – Resources which are still impacted by the identified issue.
- **Not applicable resources** – Resources for which the recommendation can't give a definitive answer. The not applicable tab also includes reasons for each resource.

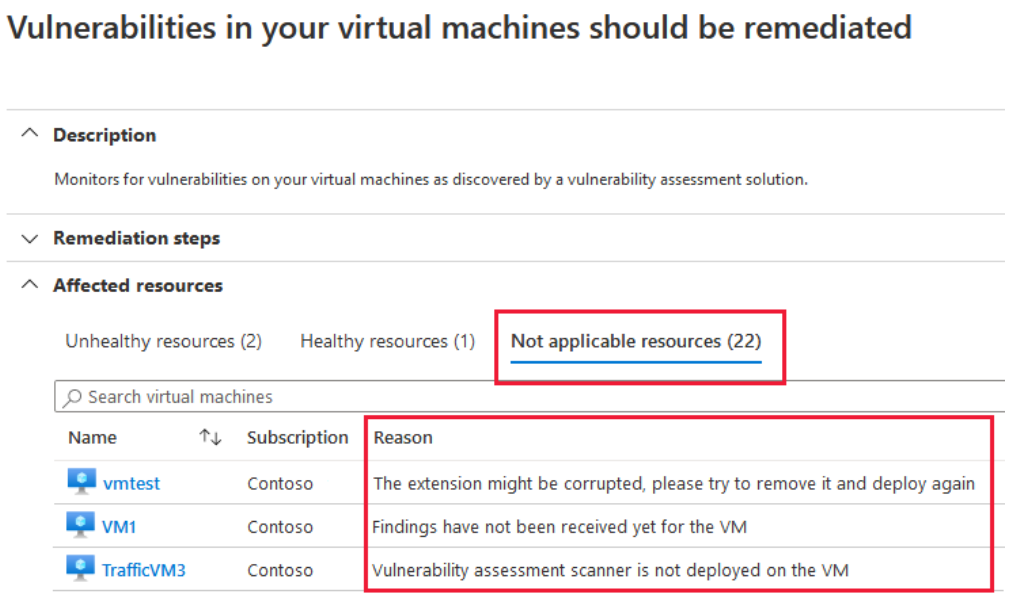

j. Action buttons to remediate the recommendation or trigger a logic app.

Review recommendation data in Azure Resource Graph Explorer (ARG)

The toolbar on the recommendation details page includes an **Open query** button to explore the details in Azure Resource Graph (ARG), an Azure service that provides the ability to query - across multiple subscriptions - Defender for Cloud's security posture data.

ARG is designed to provide efficient resource exploration with the ability to query at scale across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal.

Using the Kusto Query Language (KQL), you can cross-reference Defender for Cloud data with other resource properties.

For example, this recommendation details page shows fifteen affected resources:

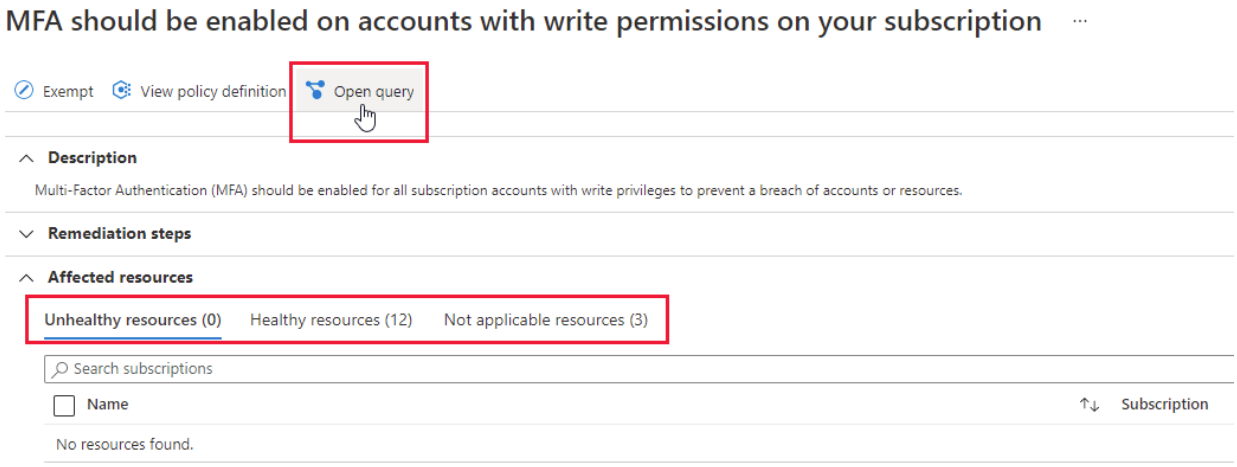

When you open the underlying query, and run it, Azure Resource Graph Explorer returns the same fifteen resources and their health status for this recommendation:

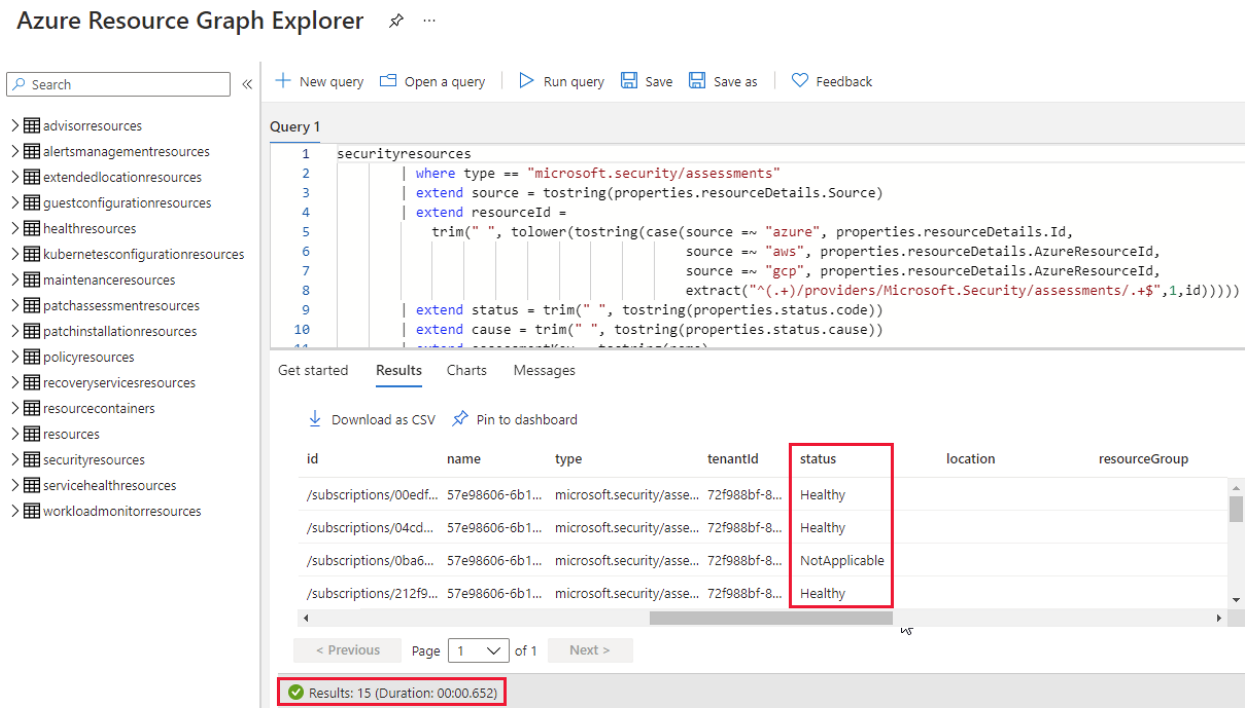

Preview recommendations

Recommendations flagged as **Preview** aren't included in the calculations of your secure score.

They should still be remediated wherever possible, so that when the preview period ends they'll contribute towards your score.

An example of a preview recommendation:

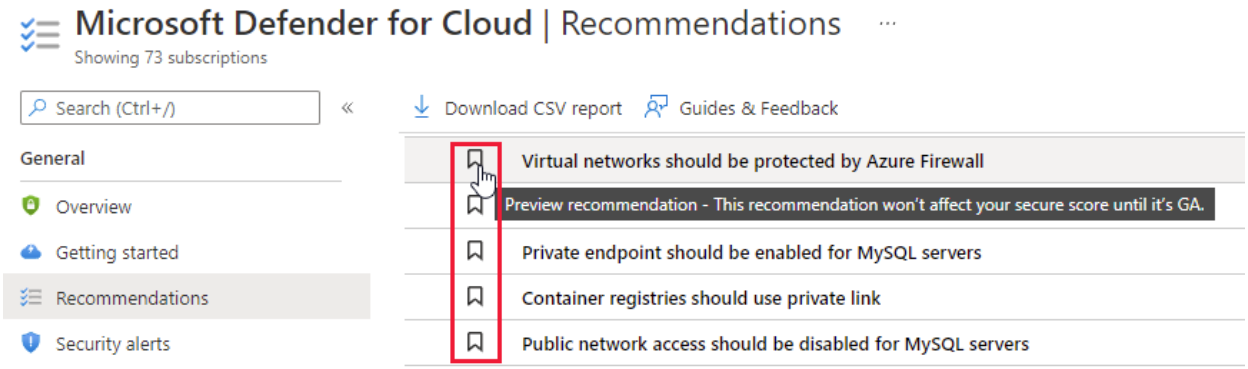

Next steps

In this document, you were introduced to security recommendations in Defender for Cloud. For related information:

- Remediate recommendations--Learn how to configure security policies for your Azure subscriptions and resource groups.
- Prevent misconfigurations with Enforce/Deny recommendations.
- Automate responses to Defender for Cloud triggers--Automate responses to recommendations
- Exempt a resource from a recommendation
- Security recommendations - a reference guide

Implement security recommendations in Microsoft Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Recommendations give you suggestions on how to better secure your resources. You implement a recommendation by following the remediation steps provided in the recommendation.

Remediation steps

After reviewing all the recommendations, decide which one to remediate first. We recommend that you prioritize the security controls with the highest potential to increase your secure score.

1. From the list, select a recommendation.
2. Follow the instructions in the **Remediation steps** section. Each recommendation has its own set of instructions. The following screenshot shows remediation steps for configuring applications to only allow traffic over HTTPS.

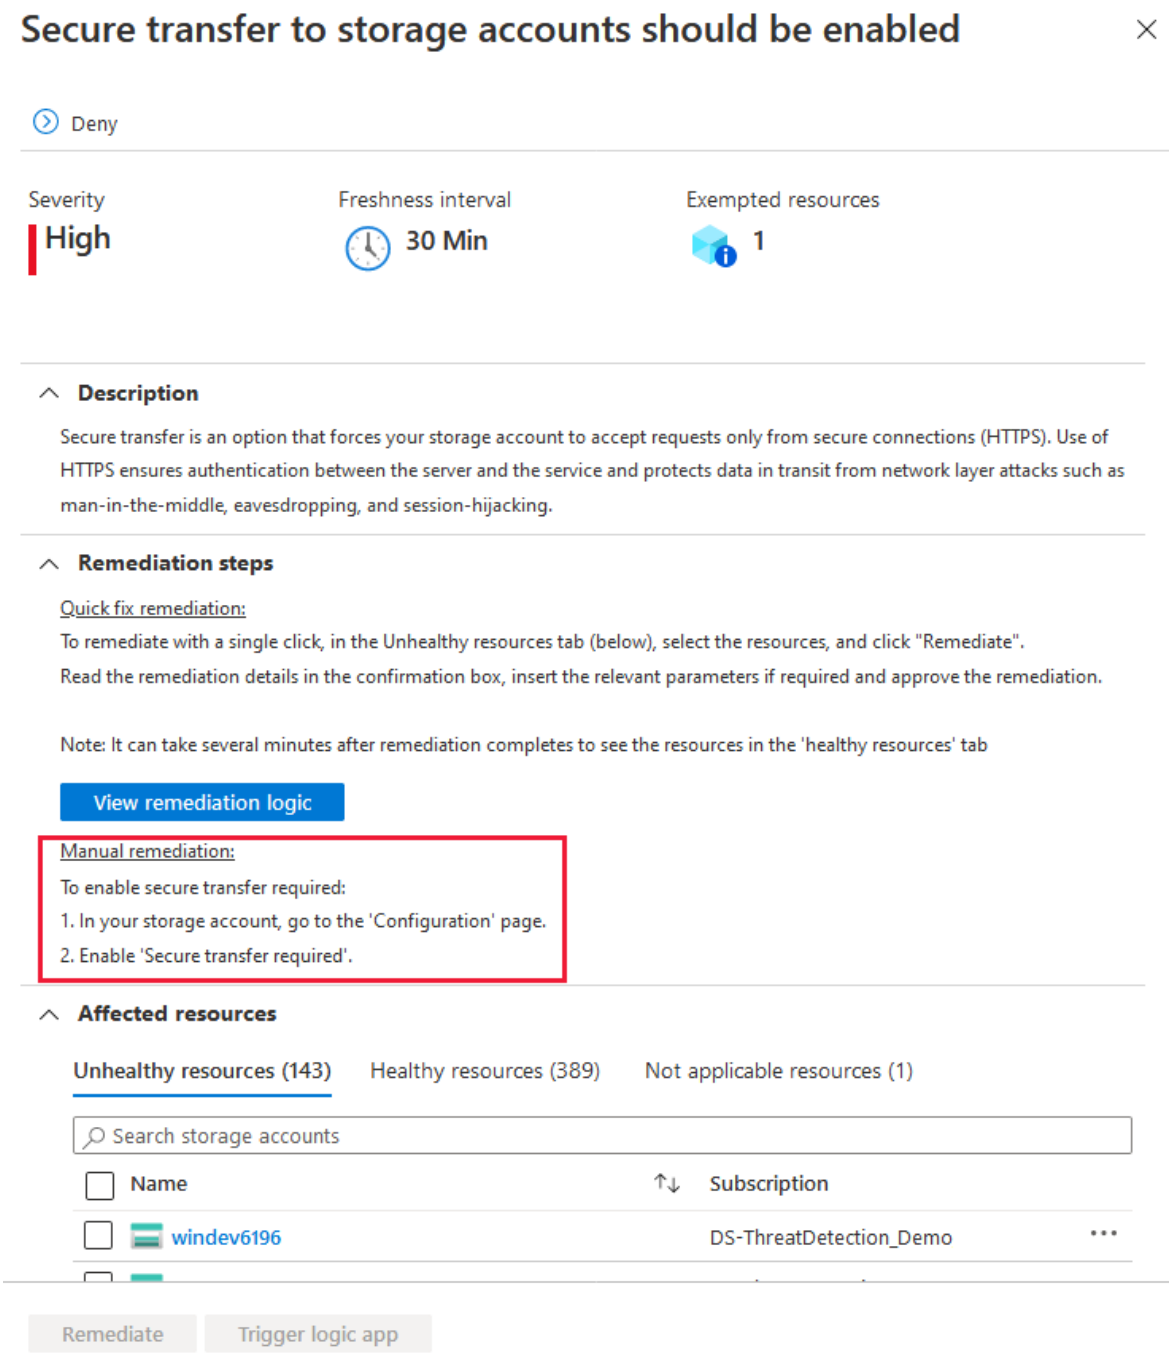

3. Once completed, a notification appears informing you whether the issue is resolved.

Fix button

To simplify remediation and improve your environment's security (and increase your secure score), many recommendations include a **Fix** option.

Fix helps you quickly remediate a recommendation on multiple resources.

TIP

The **Fix** feature is only available for specific recommendations. To find recommendations that have an available fix, use the **Response actions** filter for the list of recommendations:

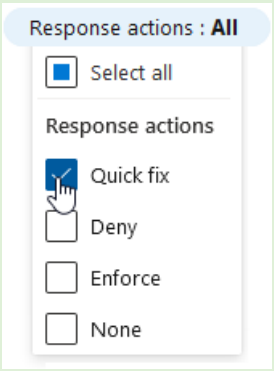

To implement a **Fix**:

1. From the list of recommendations that have the **Fix** action icon, , select a recommendation.

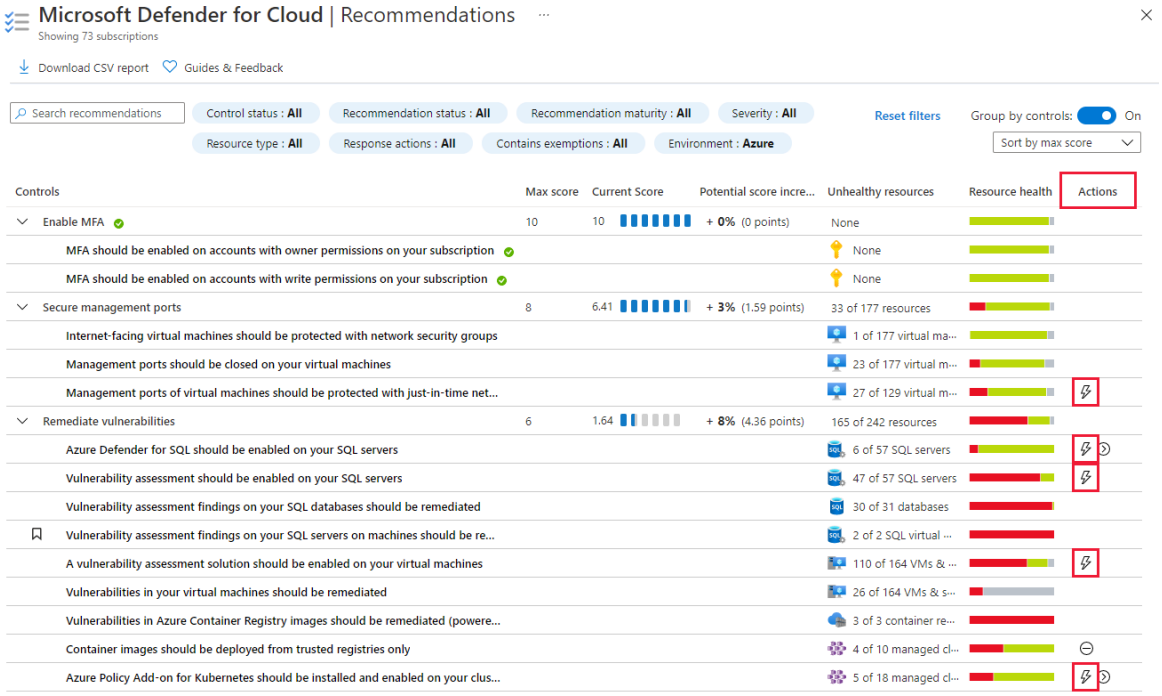

2. From the **Unhealthy resources** tab, select the resources that you want to implement the recommendation on, and select **Remediate**.

> **NOTE**
>
> Some of the listed resources might be disabled, because you don't have the appropriate permissions to modify them.

3. In the confirmation box, read the remediation details and implications.

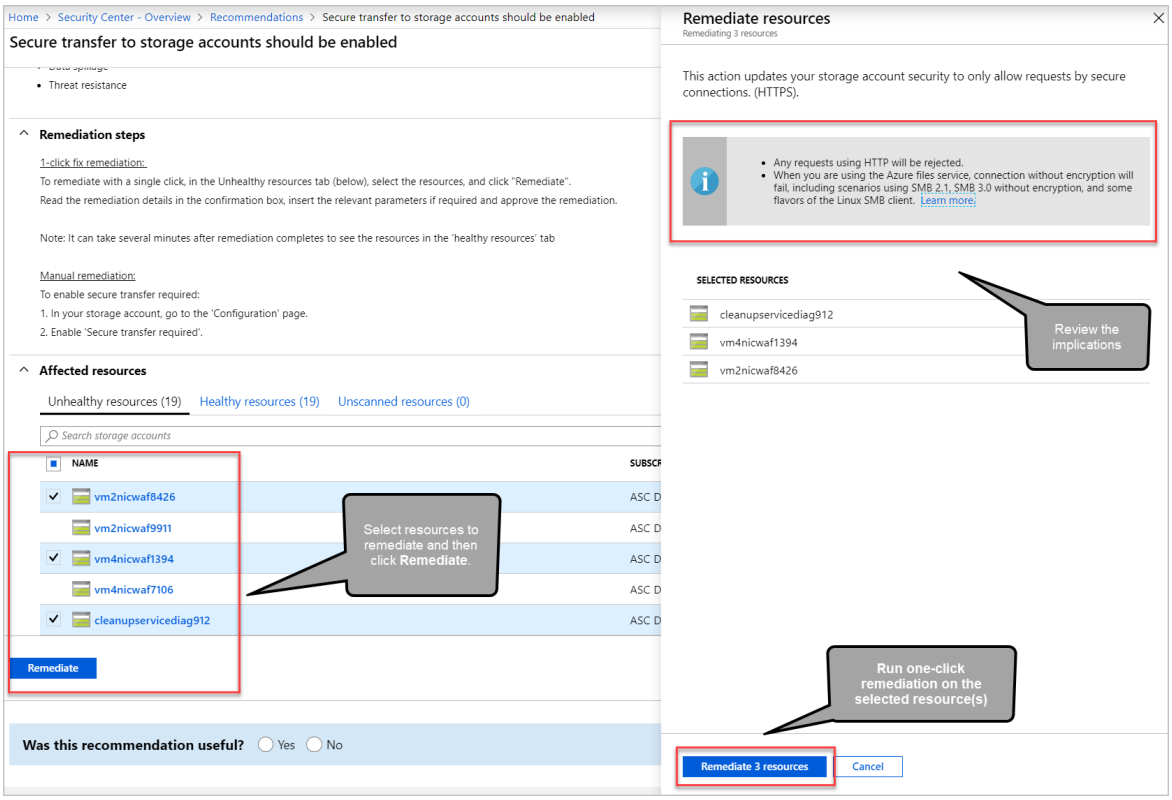

> **NOTE**
>
> The implications are listed in the grey box in the **Remediate resources** window that opens after clicking **Remediate**. They list what changes happen when proceeding with the **Fix**.

4. Insert the relevant parameters if necessary, and approve the remediation.

> **NOTE**
>
> It can take several minutes after remediation completes to see the resources in the **Healthy resources** tab. To view the remediation actions, check the activity log.

5. Once completed, a notification appears informing you if the remediation succeeded.

Fix actions logged to the activity log

The remediation operation uses a template deployment or REST API `PATCH` request to apply the configuration on the resource. These operations are logged in Azure activity log.

Next steps

In this document, you were shown how to remediate recommendations in Defender for Cloud. To learn how recommendations are defined and selected for your environment, see the following page:

- What are security policies, initiatives, and recommendations?

Prevent misconfigurations with Enforce/Deny recommendations

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Security misconfigurations are a major cause of security incidents. Defender for Cloud can help *prevent* misconfigurations of new resources with regard to specific recommendations.

This feature can help keep your workloads secure and stabilize your secure score.

Enforcing a secure configuration, based on a specific recommendation, is offered in two modes:

- Using the **Deny** effect of Azure Policy, you can stop unhealthy resources from being created
- Using the **Enforce** option, you can take advantage of Azure Policy's **DeployIfNotExist** effect and automatically remediate non-compliant resources upon creation

This can be found at the top of the resource details page for selected security recommendations (see Recommendations with deny/enforce options).

Prevent resource creation

1. Open the recommendation that your new resources must satisfy, and select the **Deny** button at the top of the page.

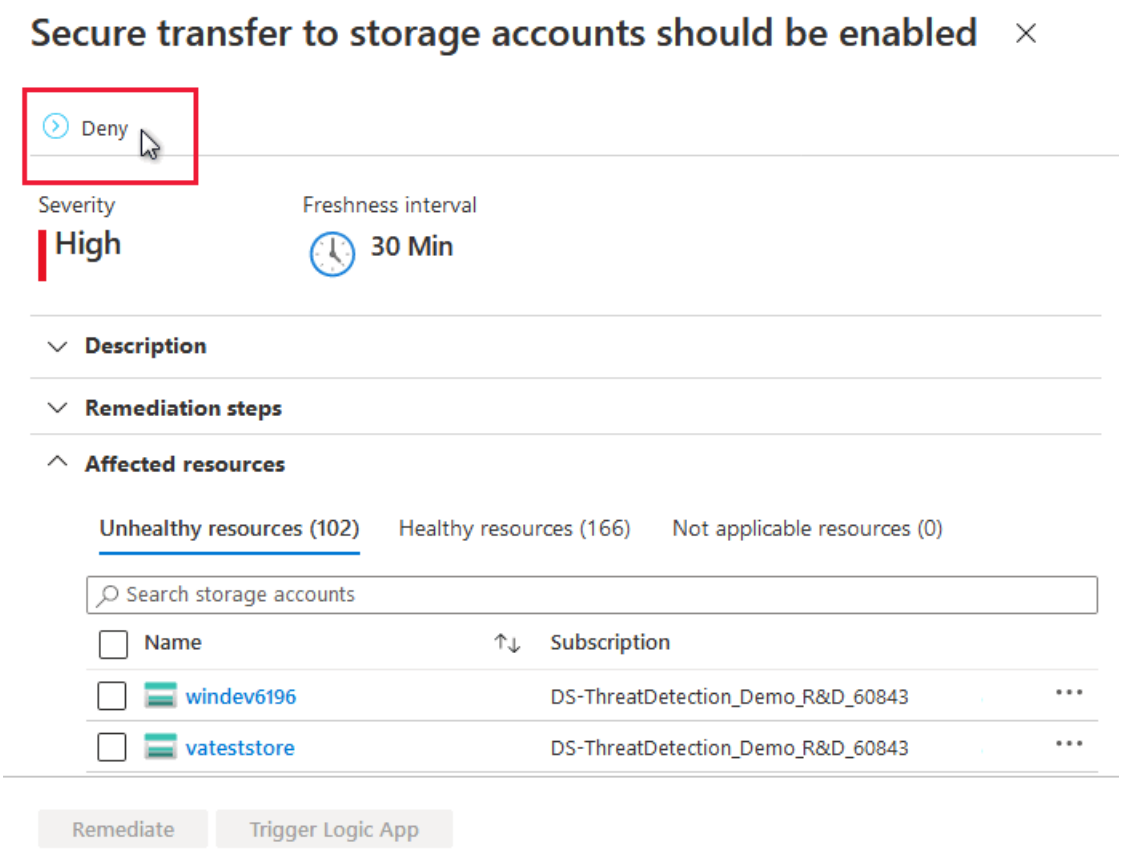

The configuration pane opens listing the scope options.

2. Set the scope by selecting the relevant subscription or management group.

> **TIP**
> You can use the three dots at the end of the row to change a single subscription, or use the checkboxes to select multiple subscriptions or groups then select **Change to Deny**.

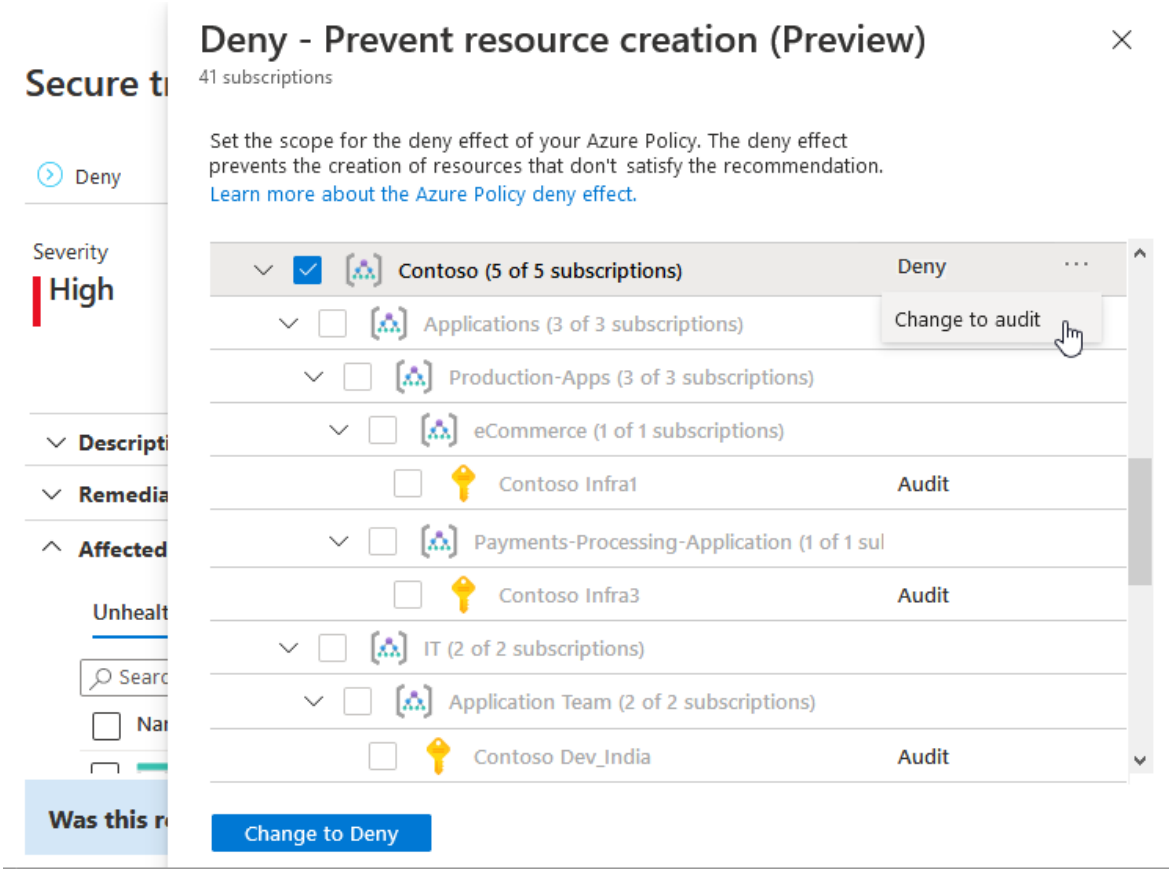

Enforce a secure configuration

1. Open the recommendation that you'll deploy a template deployment for if new resources don't satisfy it, and select the **Enforce** button at the top of the page.

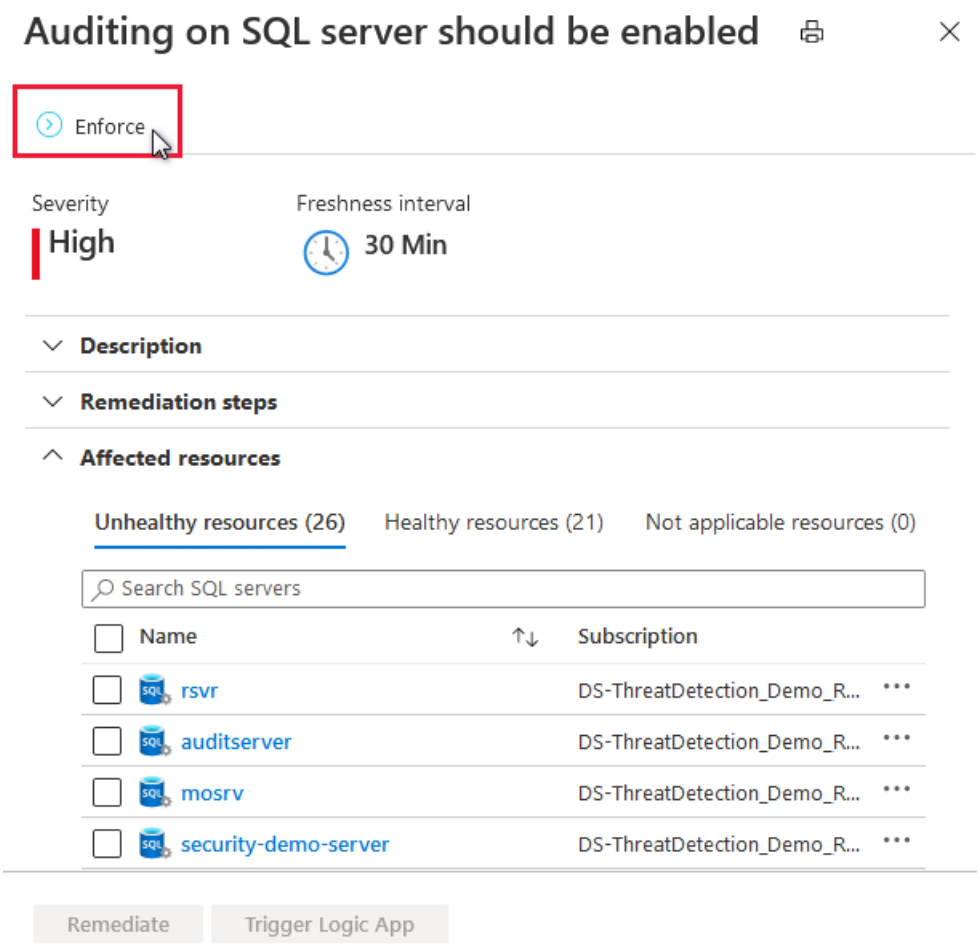

The configuration pane opens with all of the policy configuration options.

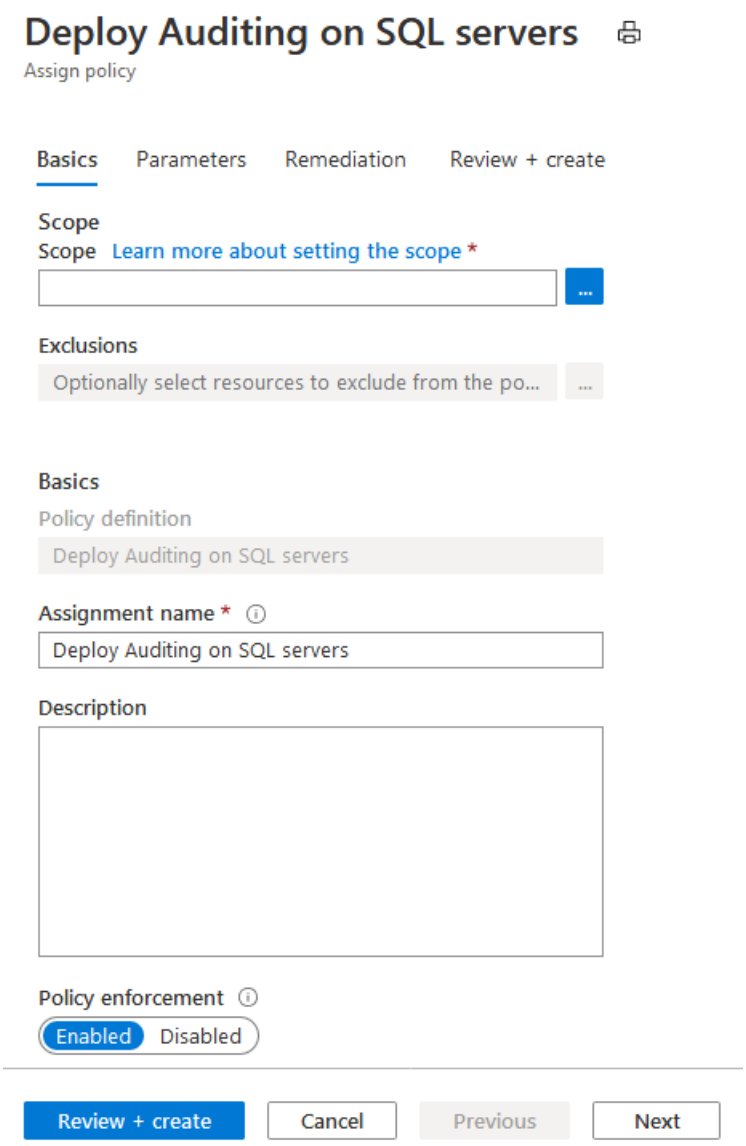

2. Set the scope, assignment name, and other relevant options.
3. Select **Review + create**.

Recommendations with deny/enforce options

These recommendations can be used with the **deny** option:

- [Enable if required] Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest
- [Enable if required] Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK)
- [Enable if required] Cognitive Services accounts should enable data encryption with a customer-managed key (CMK)
- [Enable if required] Container registries should be encrypted with a customer-managed key (CMK)
- Access to storage accounts with firewall and virtual network configurations should be restricted
- Automation account variables should be encrypted
- Azure Cache for Redis should reside within a virtual network
- Azure Spring Cloud should use network injection
- Container CPU and memory limits should be enforced
- Container images should be deployed from trusted registries only

- Container with privilege escalation should be avoided
- Containers sharing sensitive host namespaces should be avoided
- Containers should listen on allowed ports only
- Immutable (read-only) root filesystem should be enforced for containers
- Key Vault keys should have an expiration date
- Key Vault secrets should have an expiration date
- Key vaults should have purge protection enabled
- Key vaults should have soft delete enabled
- Least privileged Linux capabilities should be enforced for containers
- Overriding or disabling of containers AppArmor profile should be restricted
- Privileged containers should be avoided
- Redis Cache should allow access only via SSL
- Running containers as root user should be avoided
- Secure transfer to storage accounts should be enabled
- Service Fabric clusters should have the ClusterProtectionLevel property set to EncryptAndSign
- Service Fabric clusters should only use Azure Active Directory for client authentication
- Services should listen on allowed ports only
- Storage account public access should be disallowed
- Storage accounts should be migrated to new Azure Resource Manager resources
- Storage accounts should restrict network access using virtual network rules
- Usage of host networking and ports should be restricted
- Usage of pod HostPath volume mounts should be restricted to a known list to restrict node access from compromised containers
- Validity period of certificates stored in Azure Key Vault should not exceed 12 months
- Virtual machines should be migrated to new Azure Resource Manager resources
- Web Application Firewall (WAF) should be enabled for Application Gateway
- Web Application Firewall (WAF) should be enabled for Azure Front Door Service service

These recommendations can be used with the **enforce** option:

- Auditing on SQL server should be enabled
- Azure Arc-enabled Kubernetes clusters should have Microsoft Defender for Cloud's extension installed
- Azure Backup should be enabled for virtual machines
- Microsoft Defender for App Service should be enabled
- Microsoft Defender for container registries should be enabled
- Microsoft Defender for DNS should be enabled
- Microsoft Defender for Key Vault should be enabled
- Microsoft Defender for Kubernetes should be enabled
- Microsoft Defender for Resource Manager should be enabled
- Microsoft Defender for servers should be enabled
- Microsoft Defender for Azure SQL Database servers should be enabled
- Microsoft Defender for SQL servers on machines should be enabled
- Microsoft Defender for SQL should be enabled for unprotected Azure SQL servers
- Microsoft Defender for Storage should be enabled
- Azure Policy Add-on for Kubernetes should be installed and enabled on your clusters
- Diagnostic logs in Azure Stream Analytics should be enabled
- Diagnostic logs in Batch accounts should be enabled
- Diagnostic logs in Data Lake Analytics should be enabled

- Diagnostic logs in Event Hub should be enabled
- Diagnostic logs in Key Vault should be enabled
- Diagnostic logs in Logic Apps should be enabled
- Diagnostic logs in Search services should be enabled
- Diagnostic logs in Service Bus should be enabled

Automate responses to Microsoft Defender for Cloud triggers

2/15/2022 • 6 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Every security program includes multiple workflows for incident response. These processes might include notifying relevant stakeholders, launching a change management process, and applying specific remediation steps. Security experts recommend that you automate as many steps of those procedures as you can. Automation reduces overhead. It can also improve your security by ensuring the process steps are done quickly, consistently, and according to your predefined requirements.

This article describes the workflow automation feature of Microsoft Defender for Cloud. This feature can trigger Logic Apps on security alerts, recommendations, and changes to regulatory compliance. For example, you might want Defender for Cloud to email a specific user when an alert occurs. You'll also learn how to create Logic Apps using Azure Logic Apps.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free
Required roles and permissions:	**Security admin role** or **Owner** on the resource group Must also have write permissions for the target resource To work with Azure Logic Apps workflows, you must also have the following Logic Apps roles/permissions: - Logic App Operator permissions are required or Logic App read/trigger access (this role can't create or edit logic apps; only *run* existing ones) - Logic App Contributor permissions are required for Logic App creation and modification If you want to use Logic App connectors, you may need additional credentials to sign in to their respective services (for example, your Outlook/Teams/Slack instances)
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

Create a logic app and define when it should automatically run

1. From Defender for Cloud's sidebar, select **Workflow automation**.

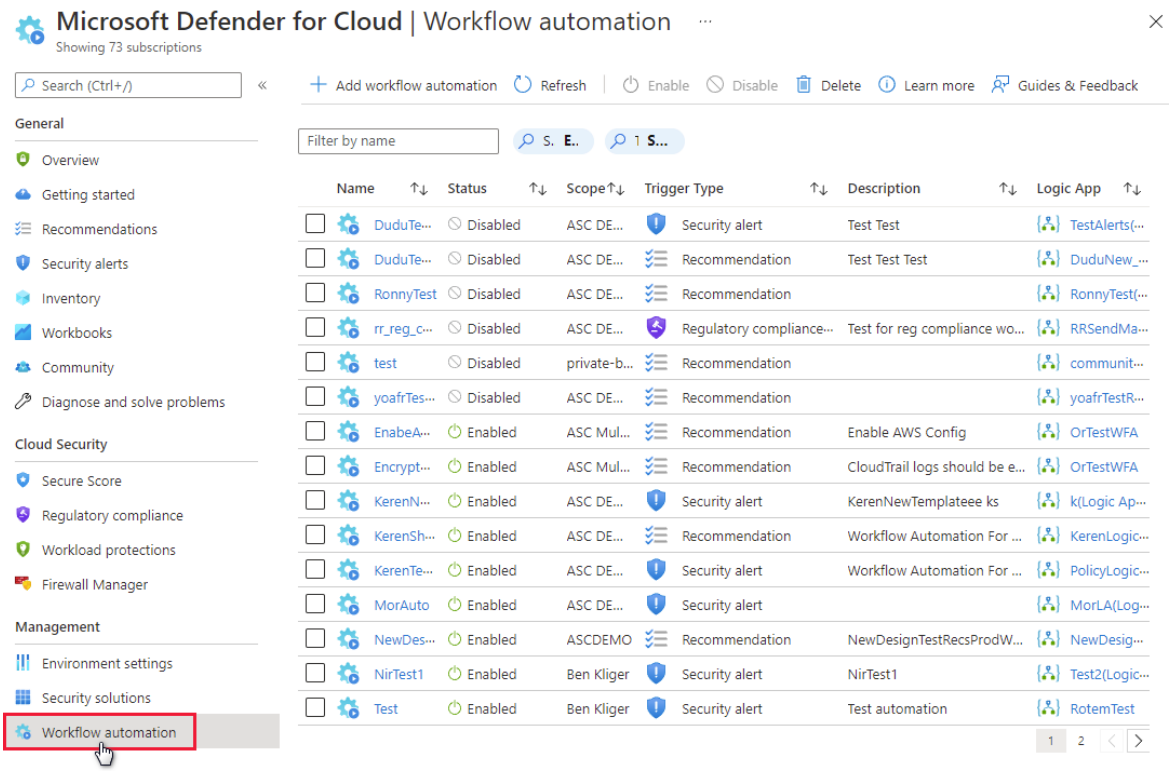

From this page you can create new automation rules, as well as enable, disable, or delete existing ones.

2. To define a new workflow, click **Add workflow automation**. The options pane for your new automation opens.

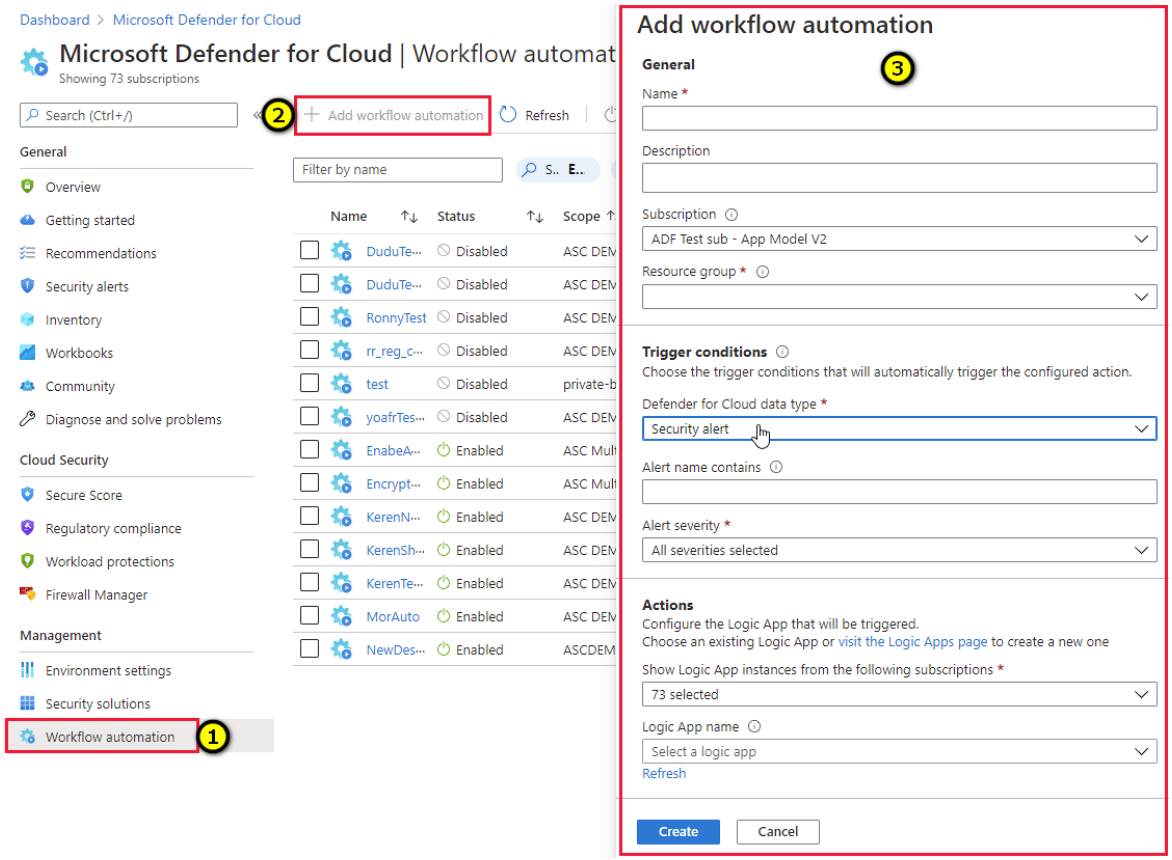

Here you can enter:

a. A name and description for the automation.

b. The triggers that will initiate this automatic workflow. For example, you might want your Logic App to run when a security alert that contains "SQL" is generated.

> **NOTE**
>
> If your trigger is a recommendation that has "sub-recommendations", for example **Vulnerability assessment findings on your SQL databases should be remediated**, the logic app will not trigger for every new security finding; only when the status of the parent recommendation changes.

c. The Logic App that will run when your trigger conditions are met.

3. From the Actions section, select **visit the Logic Apps page** to begin the Logic App creation process.

   You'll be taken to Azure Logic Apps.

4. Select **Add**.

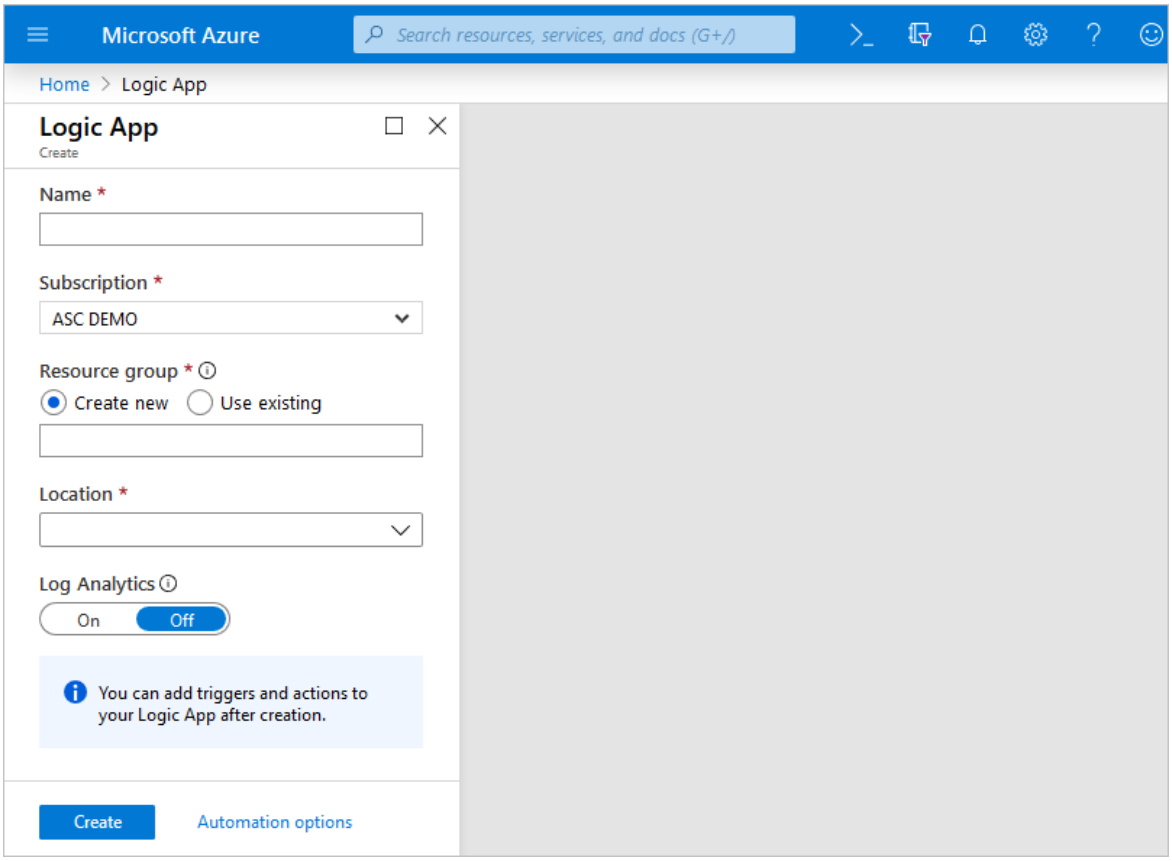

5. Enter a name, resource group, and location, and select **Review and create** > **Create**.

   The message **Deployment is in progress** appears. Wait for the deployment complete notification to appear and select **Go to resource** from the notification.

6. In your new logic app, you can choose from built-in, predefined templates from the security category. Or you can define a custom flow of events to occur when this process is triggered.

> **TIP**
>
> Sometimes in a logic app, parameters are included in the connector as part of a string and not in their own field. For an example of how to extract parameters, see step #14 of Working with logic app parameters while building Microsoft Defender for Cloud workflow automations.

   The logic app designer supports these Defender for Cloud triggers:

- **When a Microsoft Defender for Cloud Recommendation is created or triggered** - If your logic app relies on a recommendation that gets deprecated or replaced, your automation will stop

working and you'll need to update the trigger. To track changes to recommendations, use the release notes.

- **When a Defender for Cloud Alert is created or triggered** - You can customize the trigger so that it relates only to alerts with the severity levels that interest you.
- **When a Defender for Cloud regulatory compliance assessment is created or triggered** - Trigger automations based on updates to regulatory compliance assessments.

> **NOTE**
>
> If you are using the legacy trigger "When a response to a Microsoft Defender for Cloud alert is triggered", your logic apps will not be launched by the Workflow Automation feature. Instead, use either of the triggers mentioned above.

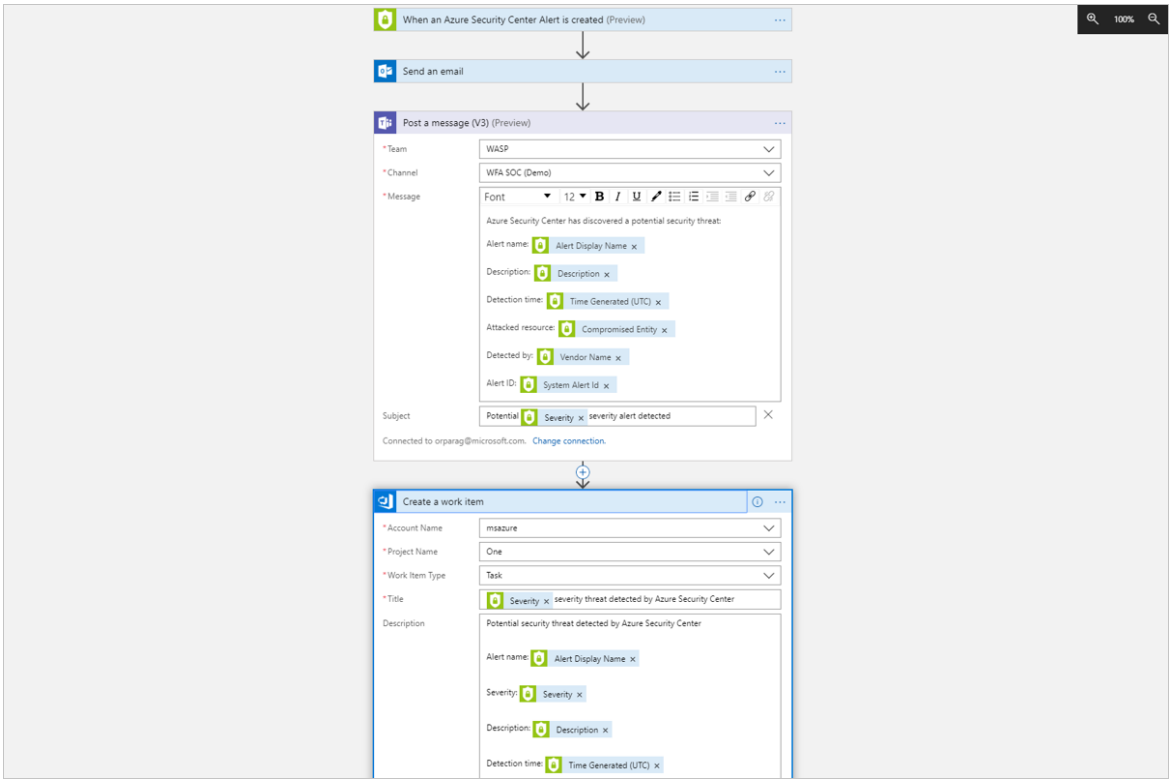

7. After you've defined your logic app, return to the workflow automation definition pane ("Add workflow automation"). Click **Refresh** to ensure your new Logic App is available for selection.

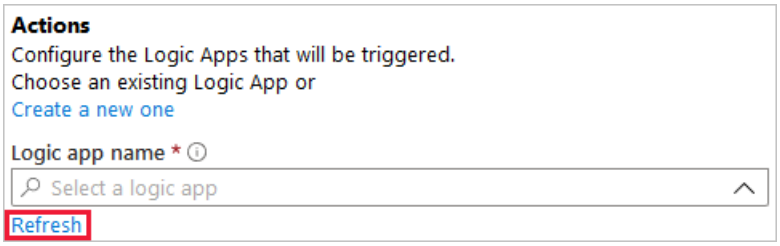

8. Select your logic app and save the automation. Note that the Logic App dropdown only shows Logic Apps with supporting Defender for Cloud connectors mentioned above.

Manually trigger a Logic App

You can also run Logic Apps manually when viewing any security alert or recommendation.

To manually run a Logic App, open an alert or a recommendation and click **Trigger Logic App**:

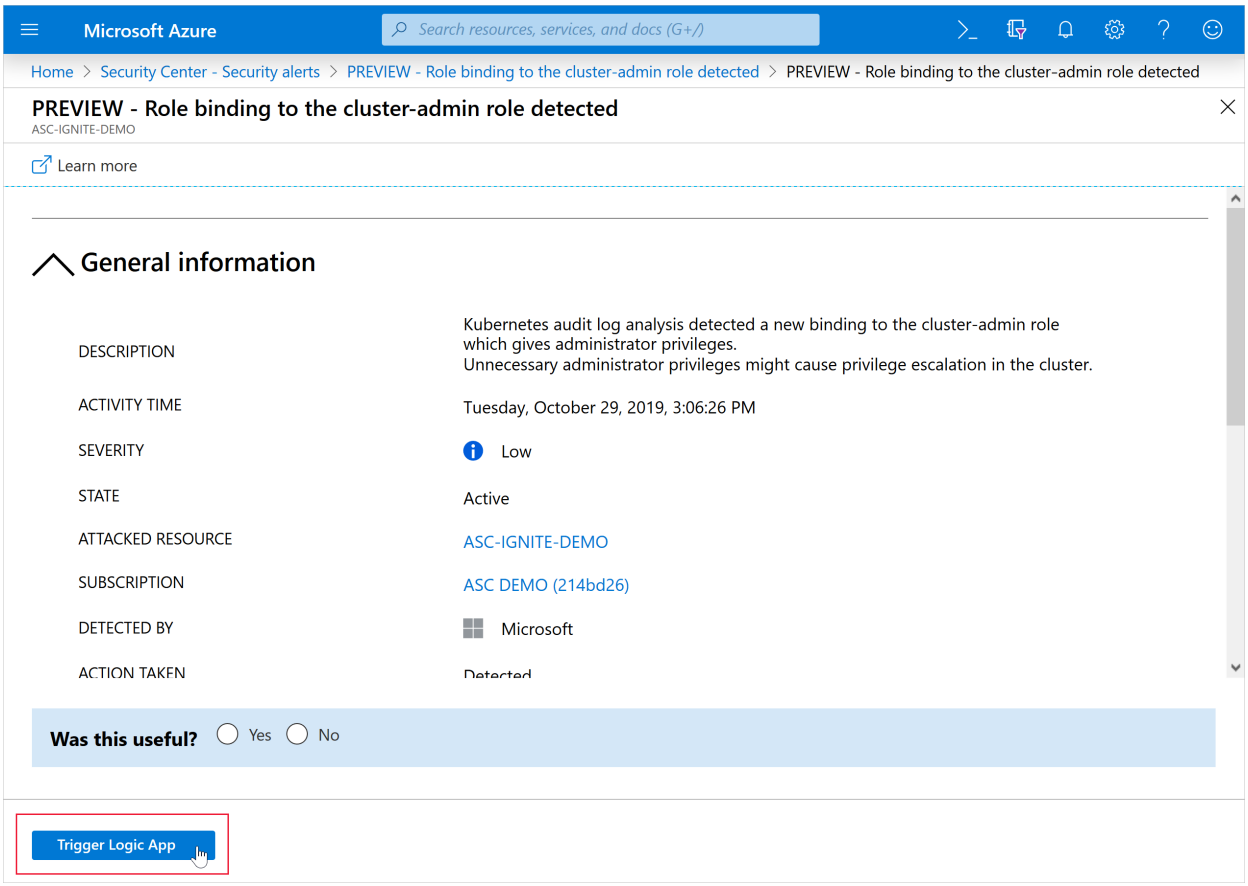

Configure workflow automation at scale using the supplied policies

Automating your organization's monitoring and incident response processes can greatly improve the time it takes to investigate and mitigate security incidents.

To deploy your automation configurations across your organization, use the supplied Azure Policy 'DeployIfNotExist' policies described below to create and configure workflow automation procedures.

Get started with workflow automation templates.

To implement these policies:

1. From the table below, select the policy you want to apply:

GOAL	POLICY	POLICY ID
Workflow automation for security alerts	Deploy Workflow Automation for Microsoft Defender for Cloud alerts	f1525828-9a90-4fcf-be48-268cdd02361e
Workflow automation for security recommendations	Deploy Workflow Automation for Microsoft Defender for Cloud recommendations	73d6ab6c-2475-4850-afd6-43795f3492ef
Workflow automation for regulatory compliance changes	Deploy Workflow Automation for Microsoft Defender for Cloud regulatory compliance	509122b9-ddd9-47ba-a5f1-d0dac20be63c

TIP

You can also find these by searching Azure Policy:

1. Open Azure Policy.

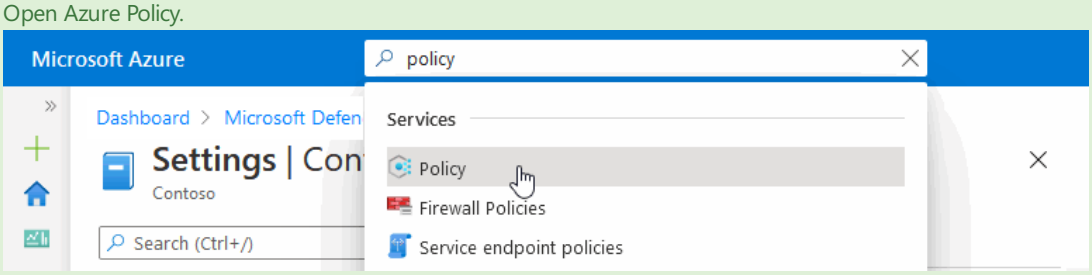

2. From the Azure Policy menu, select **Definitions** and search for them by name.

2. From the relevant Azure Policy page, select **Assign**.

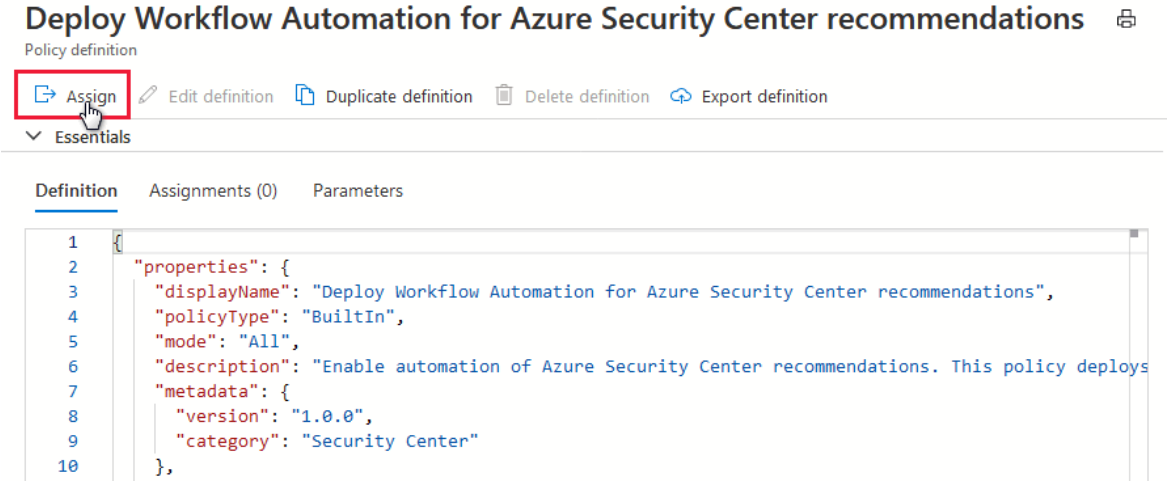

3. Open each tab and set the parameters as desired:

   a. In the **Basics** tab, set the scope for the policy. To use centralized management, assign the policy to the Management Group containing the subscriptions that will use the workflow automation configuration.

   b. In the **Parameters** tab, set the resource group and data type details.

TIP

Each parameter has a tooltip explaining the options available to you.

Azure Policy's parameters tab (1) provides access to similar configuration options as Defender for Cloud's workflow automation page (2).

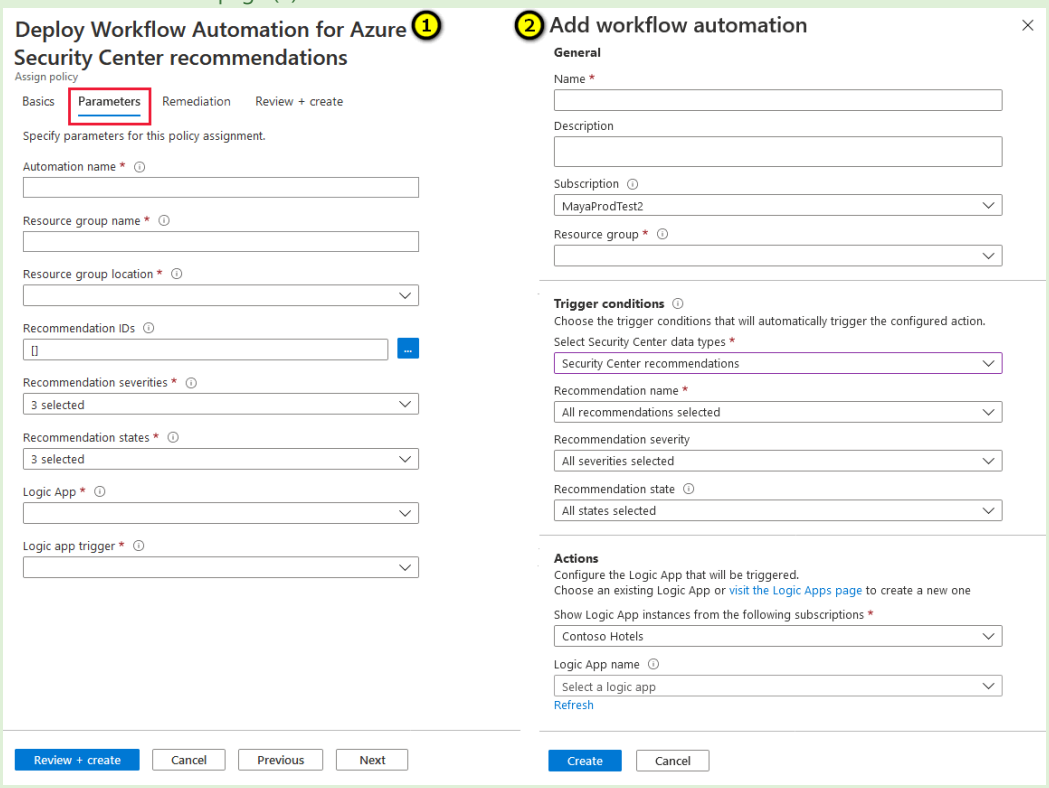

c. Optionally, to apply this assignment to existing subscriptions, open the **Remediation** tab and select the option to create a remediation task.

4. Review the summary page and select **Create**.

Data types schemas

To view the raw event schemas of the security alerts or recommendations events passed to the Logic App instance, visit the Workflow automation data types schemas. This can be useful in cases where you are not using Defender for Cloud's built-in Logic App connectors mentioned above, but instead are using Logic App's generic HTTP connector - you could use the event JSON schema to manually parse it as you see fit.

FAQ - Workflow automation

Does workflow automation support any business continuity or disaster recovery (BCDR) scenarios?

When preparing your environment for BCDR scenarios, where the target resource is experiencing an outage or other disaster, it's the organization's responsibility to prevent data loss by establishing backups according to the guidelines from Azure Event Hubs, Log Analytics workspace, and Logic App.

For every active automation, we recommend you create an identical (disabled) automation and store it in a different location. When there's an outage, you can enable these backup automations and maintain normal operations.

Learn more about Business continuity and disaster recovery for Azure Logic Apps.

Next steps

In this article, you learned about creating Logic Apps, automating their execution in Defender for Cloud, and running them manually.

For related material, see:

- The Microsoft Learn module on how to use workflow automation to automate a security response
- Security recommendations in Microsoft Defender for Cloud
- Security alerts in Microsoft Defender for Cloud
- About Azure Logic Apps
- Connectors for Azure Logic Apps
- Workflow automation data types schemas

Manage security policies

2/15/2022 • 5 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how security policies are configured, and how to view them in Microsoft Defender for Cloud.

To understand the relationships between initiatives, policies, and recommendations, see What are security policies, initiatives, and recommendations?

Who can edit security policies?

Defender for Cloud uses Azure role-based access control (Azure RBAC), which provides built-in roles you can assign to Azure users, groups, and services. When users open Defender for Cloud, they see only information related to the resources they can access. Which means users are assigned the role of *owner*, *contributor*, or *reader* to the resource's subscription. There are also two specific Defender for Cloud roles:

- **Security reader**: Has rights to view Defender for Cloud items such as recommendations, alerts, policy, and health. Can't make changes.
- **Security admin**: Has the same view rights as *security reader*. Can also update the security policy and dismiss alerts.

You can edit security policies through the Azure Policy portal, via REST API or using Windows PowerShell.

Manage your security policies

To view your security policies in Defender for Cloud:

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the relevant subscription or management group whose policies you want to view.
3. Open the **Security policy** page.
4. The security policy page for that subscription or management group appears. It shows the available and assigned policies.

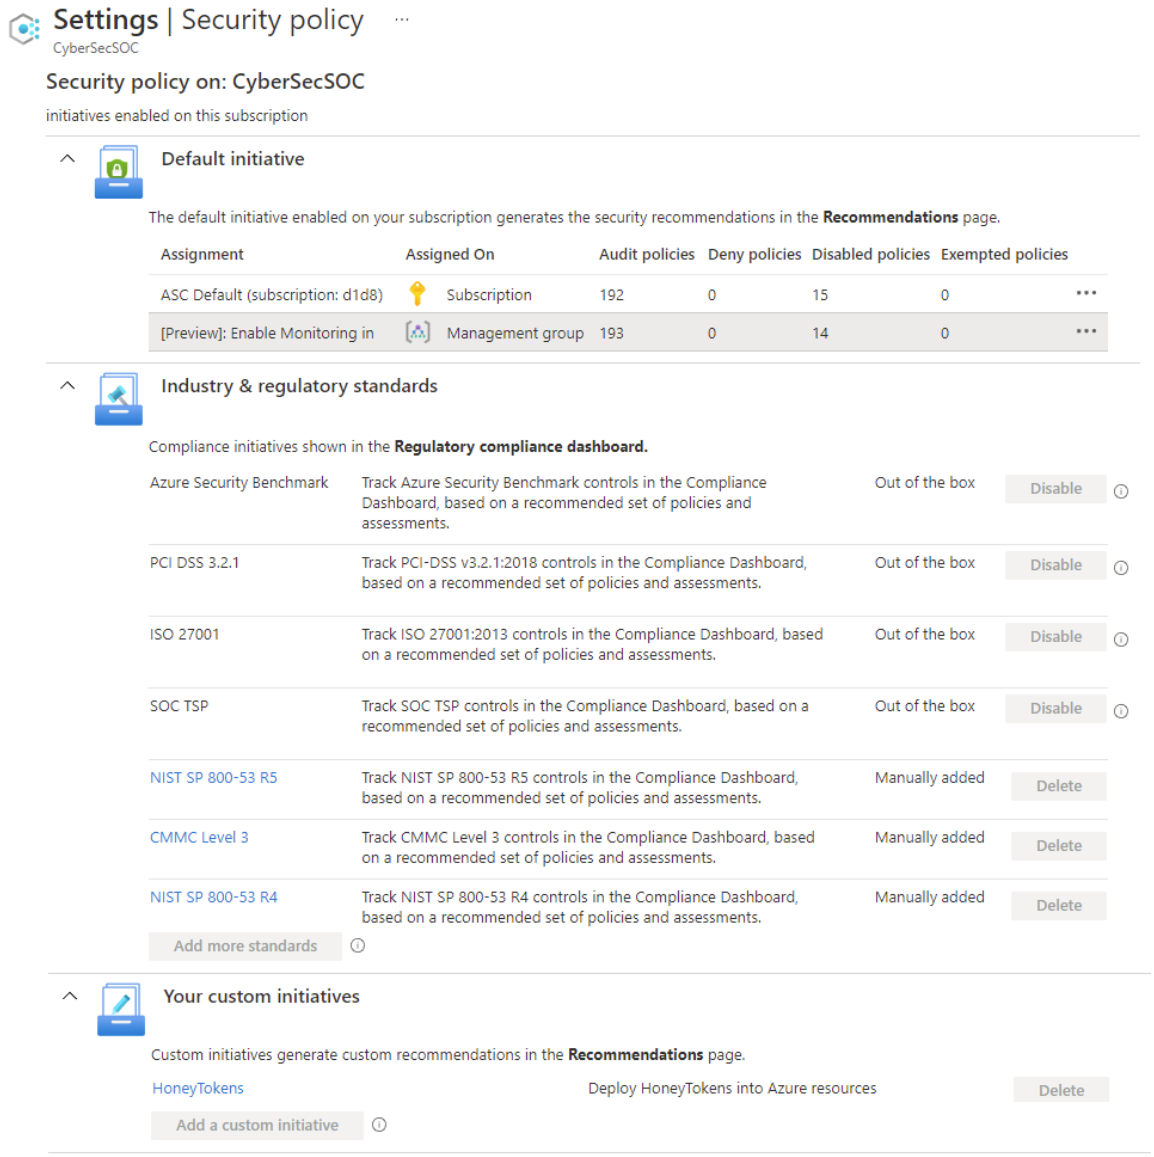

> **NOTE**
>
> If there is a label "MG Inherited" alongside your default initiative, it means that the initiative has been assigned to a management group and inherited by the subscription you're viewing.

5. Choose from the available options on this page:

   a. To work with industry standards, select **Add more standards**. For more information, see Customize the set of standards in your regulatory compliance dashboard.

   b. To assign and manage custom initiatives, select **Add custom initiatives**. For more information, see Using custom security initiatives and policies.

   c. To view and edit the default initiative, select it and proceed as described below.

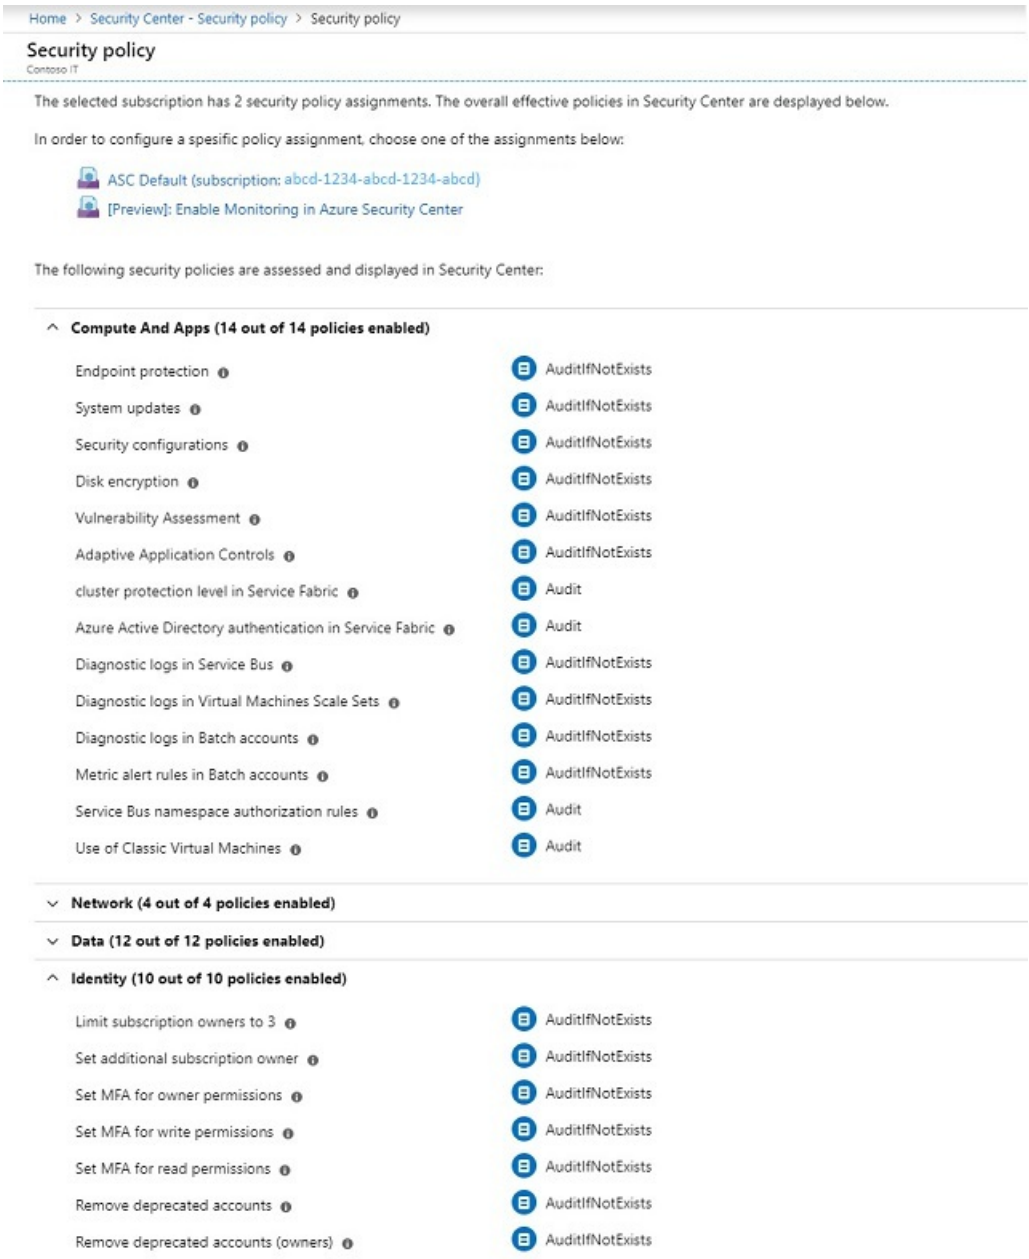

This **Security policy** screen reflects the action taken by the policies assigned on the subscription or management group you selected.

- Use the links at the top to open a policy **assignment** that applies on the subscription or management group. These links let you access the assignment and edit or disable the policy. For example, if you see that a particular policy assignment is effectively denying endpoint protection, use the link to edit or disable the policy.
- In the list of policies, you can see the effective application of the policy on your subscription or management group. The settings of each policy that apply to the scope are taken into consideration and the cumulative outcome of actions taken by the policy is shown. For example, if in one assignment of the policy is disabled, but in another it's set to AuditIfNotExist, then the cumulative effect applies AuditIfNotExist. The more active effect always takes precedence.
- The policies' effect can be: Append, Audit, AuditIfNotExists, Deny, DeployIfNotExists, Disabled. For more information on how effects are applied, see Understand Policy effects.

> **NOTE**
> When you view assigned policies, you can see multiple assignments and you can see how each assignment is configured on its own.

Disable security policies and disable recommendations

When your security initiative triggers a recommendation that's irrelevant for your environment, you can prevent that recommendation from appearing again. To disable a recommendation, disable the specific policy that generates the recommendation.

The recommendation you want to disable will still appear if it's required for a regulatory standard you've applied with Defender for Cloud's regulatory compliance tools. Even if you've disabled a policy in the built-in initiative, a policy in the regulatory standard's initiative will still trigger the recommendation if it's necessary for compliance. You can't disable policies from regulatory standard initiatives.

For more information about recommendations, see Managing security recommendations.

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the subscription or management group for which you want to disable the recommendation (and policy).

   > **NOTE**
   > Remember that a management group applies its policies to its subscriptions. Therefore, if you disable a subscription's policy, and the subscription belongs to a management group that still uses the same policy, then you will continue to receive the policy recommendations. The policy will still be applied from the management level and the recommendations will still be generated.

3. Open the **Security policy** page.
4. From the **Default initiative** or **Your custom initiatives** sections, select the relevant initiative containing the policy you want to disable.
5. Open the **Parameters** section and search for the policy that invokes the recommendation that you want to disable.
6. From the dropdown list, change the value for the corresponding policy to **Disabled**.

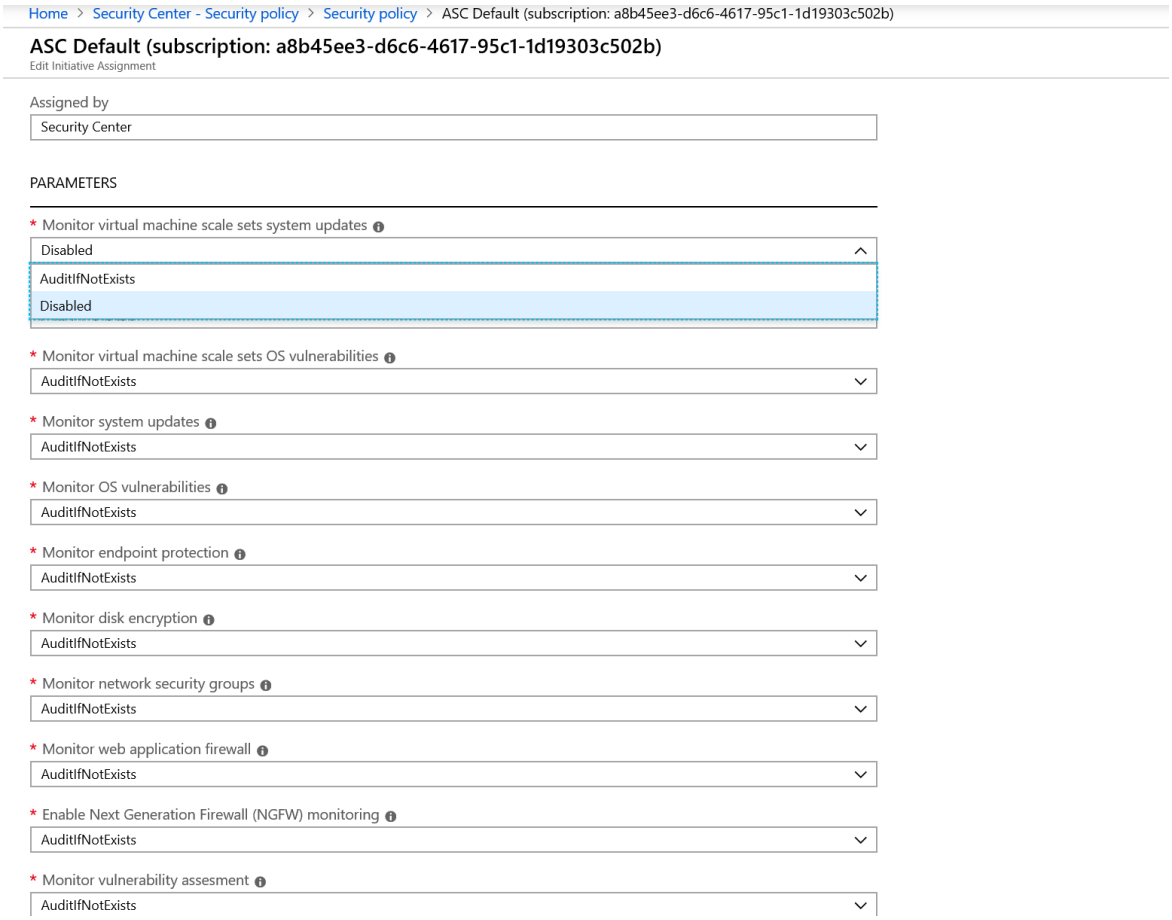

7. Select **Save**.

> **NOTE**
>
> The change might take up to 12 hours to take effect.

Enable a security policy

Some policies in your initiatives might be disabled by default. For example, in the Azure Security Benchmark initiative, some policies are provided for you to enable only if they meet a specific regulatory or compliance requirement for your organization. Such policies include recommendations to encrypt data at rest with customer-managed keys, such as "Container registries should be encrypted with a customer-managed key (CMK)".

To enable a disabled policy and ensure it's assessed for your resources:

1. From Defender for Cloud's menu, open the **Environment settings** page. Here, you can see the management groups, subscriptions, and the initiatives applied to each.
2. Select the subscription or management group for which you want to enable the recommendation (and policy).
3. Open the **Security policy** page.
4. From the **Default initiative**, **Industry & regulatory standards**, or **Your custom initiatives** sections, select the relevant initiative with the policy you want to enable.
5. Open the **Parameters** section and search for the policy that invokes the recommendation that you want to disable.
6. From the dropdown list, change the value for the corresponding policy to **AuditIfNotExists** or **Enforce**.

7. Select **Save**.

> **NOTE**
>
> The change might take up to 12 hours to take effect.

Next steps

This page explained security policies. For related information, see the following pages:

- Learn how to set policies using PowerShell
- Learn how to edit a security policy in Azure Policy
- Learn how to set a policy across subscriptions or on Management groups using Azure Policy
- Learn how to enable Defender for Cloud on all subscriptions in a management group

Exempting resources and recommendations from your secure score

2/15/2022 • 9 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

A core priority of every security team is to ensure analysts can focus on the tasks and incidents that matter to the organization. Defender for Cloud has many features for customizing the experience and making sure your secure score reflects your organization's security priorities. The **exempt** option is one such feature.

When you investigate your security recommendations in Microsoft Defender for Cloud, one of the first pieces of information you review is the list of affected resources.

Occasionally, a resource will be listed that you feel shouldn't be included. Or a recommendation will show in a scope where you feel it doesn't belong. The resource might have been remediated by a process not tracked by Defender for Cloud. The recommendation might be inappropriate for a specific subscription. Or perhaps your organization has simply decided to accept the risks related to the specific resource or recommendation.

In such cases, you can create an exemption for a recommendation to:

- **Exempt a resource** to ensure it isn't listed with the unhealthy resources in the future, and doesn't impact your secure score. The resource will be listed as not applicable and the reason will be shown as "exempted" with the specific justification you select.
- **Exempt a subscription or management group** to ensure that the recommendation doesn't impact your secure score and won't be shown for the subscription or management group in the future. This relates to existing resources and any you create in the future. The recommendation will be marked with the specific justification you select for the scope that you selected.

Availability

ASPECT	DETAILS
Release state:	Preview The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	This is a premium Azure Policy capability that's offered at no additional cost for customers with Microsoft Defender for Cloud's enhanced security features enabled. For other users, charges might apply in the future.

ASPECT	DETAILS
Required roles and permissions:	**Owner** or **Resource Policy Contributor** to create an exemption To create a rule, you need permissions to edit policies in Azure Policy. Learn more in Azure RBAC permissions in Azure Policy.
Limitations:	Exemptions can be created only for recommendations included in Defender for Cloud's default initiative, Azure Security Benchmark, or any of the supplied regulatory standard initiatives. Recommendations that are generated from custom initiatives cannot be exempted. Learn more about the relationships between policies, initiatives, and recommendations.
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

Define an exemption

To fine-tune the security recommendations that Defender for Cloud makes for your subscriptions, management group, or resources, you can create an exemption rule to:

- Mark a specific **recommendation** or as "mitigated" or "risk accepted". You can create recommendation exemptions for a subscription, multiple subscriptions, or an entire management group.
- Mark **one or more resources** as "mitigated" or "risk accepted" for a specific recommendation.

> **NOTE**
> Exemptions can be created only for recommendations included in Defender for Cloud's default initiative, Azure Security Benchmark or any of the supplied regulatory standard initiatives. Recommendations that are generated from any custom initiatives assigned to your subscriptions cannot be exempted. Learn more about the relationships between policies, initiatives, and recommendations.

> **TIP**
> You can also create exemptions using the API. For an example JSON, and an explanation of the relevant structures see Azure Policy exemption structure.

To create an exemption rule:

1. Open the recommendations details page for the specific recommendation.
2. From the toolbar at the top of the page, select **Exempt**.

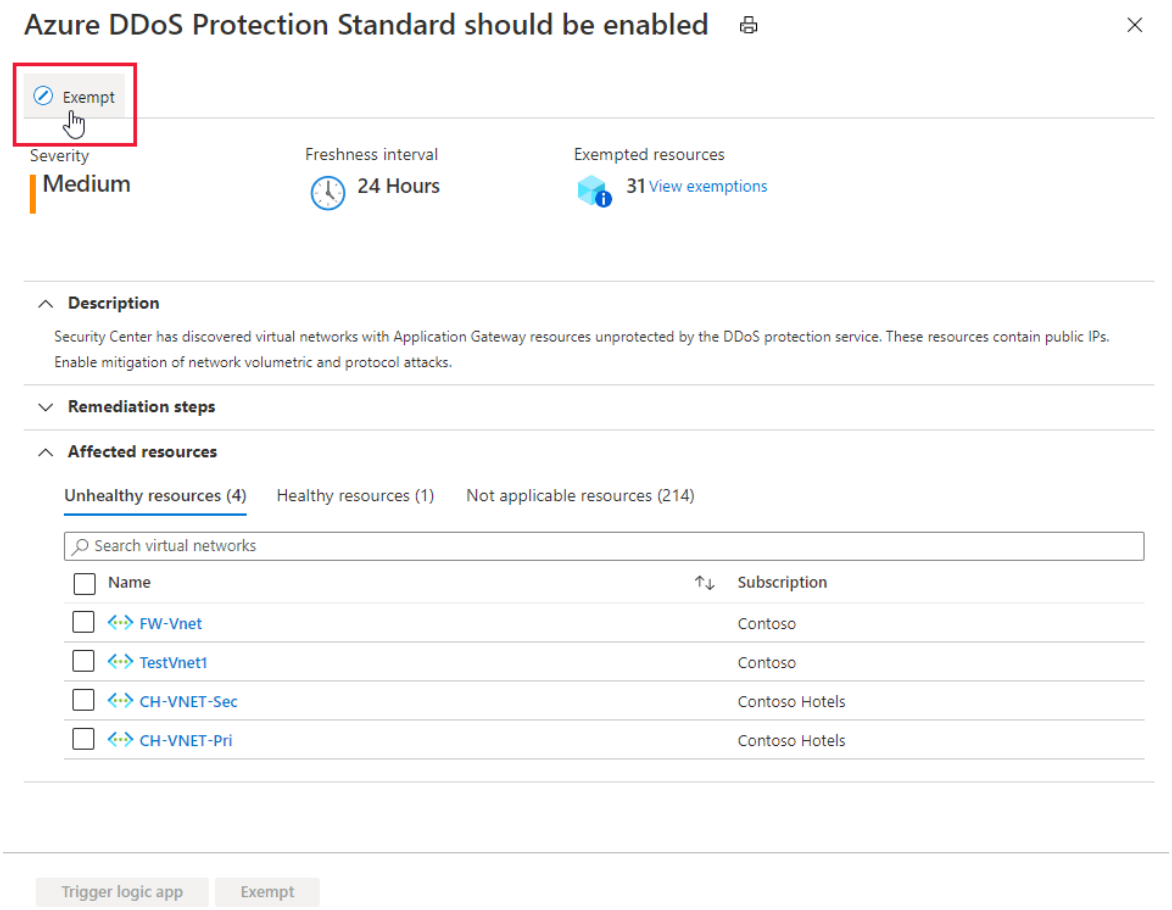

3. In the **Exempt** pane:

   a. Select the scope for this exemption rule:

   - If you select a management group, the recommendation will be exempted from all subscriptions within that group
   - If you're creating this rule to exempt one or more resources from the recommendation, choose "Selected resources"" and select the relevant ones from the list

   b. Enter a name for this exemption rule.

   c. Optionally, set an expiration date.

   d. Select the category for the exemption:

   - **Resolved through 3rd party (mitigated)** – if you're using a third-party service that Defender for Cloud hasn't identified.

     > **NOTE**
     > When you exempt a recommendation as mitigated, you aren't given points towards your secure score. But because points aren't *removed* for the unhealthy resources, the result is that your score will increase.

   - **Risk accepted (waiver)** – if you've decided to accept the risk of not mitigating this recommendation

   e. Enter a description.

   f. Select **Create**.

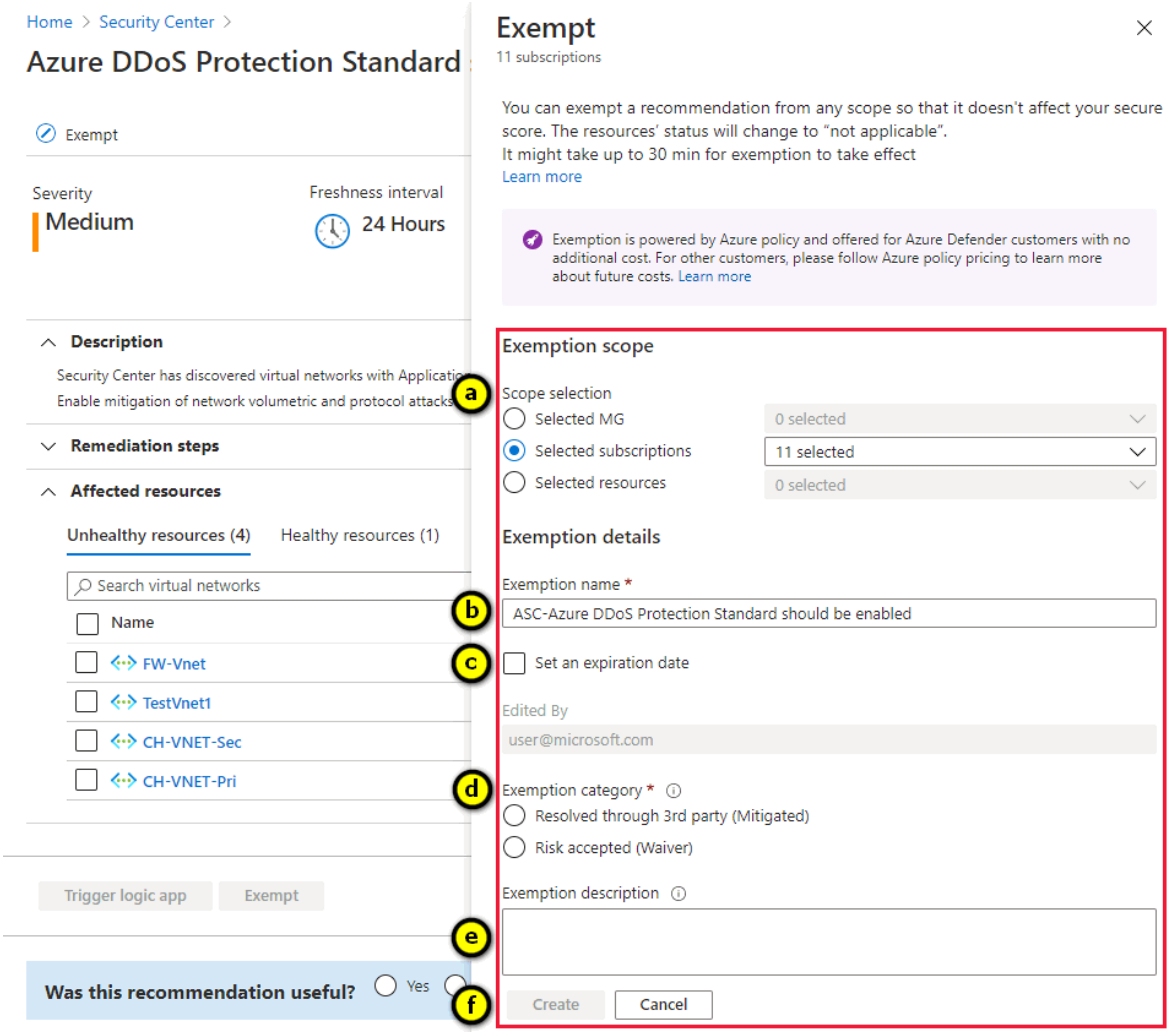

When the exemption takes effect (it might take up to 30 minutes):

- The recommendation or resources won't impact your secure score.
- If you've exempted specific resources, they'll be listed in the **Not applicable** tab of the recommendation details page.
- If you've exempted a recommendation, it will be hidden by default on Defender for Cloud's recommendations page. This is because the default options of the **Recommendation status** filter on that page are to exclude **Not applicable** recommendations. The same is true if you exempt all recommendations in a security control.

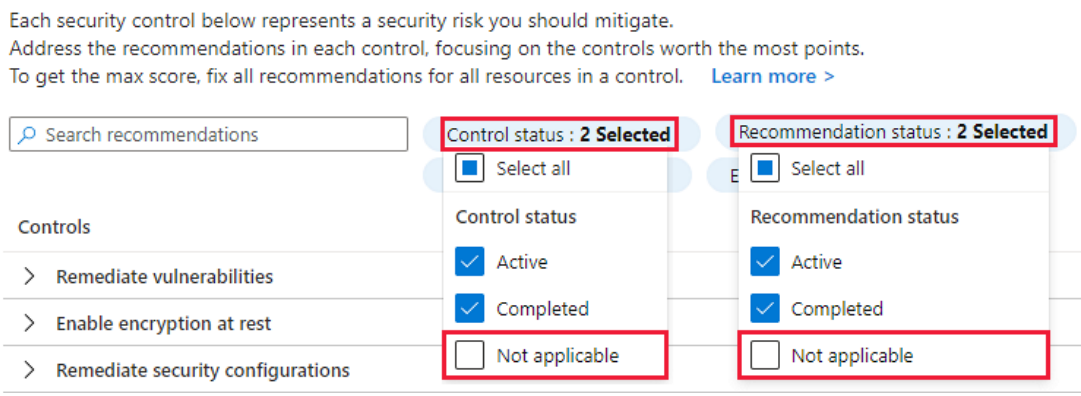

- The information strip at the top of the recommendation details page updates the number of exempted resources:

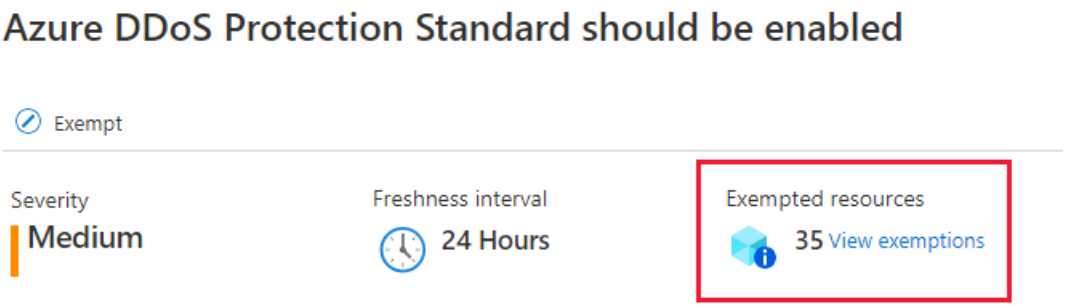

4. To review your exempted resources, open the **Not applicable** tab:

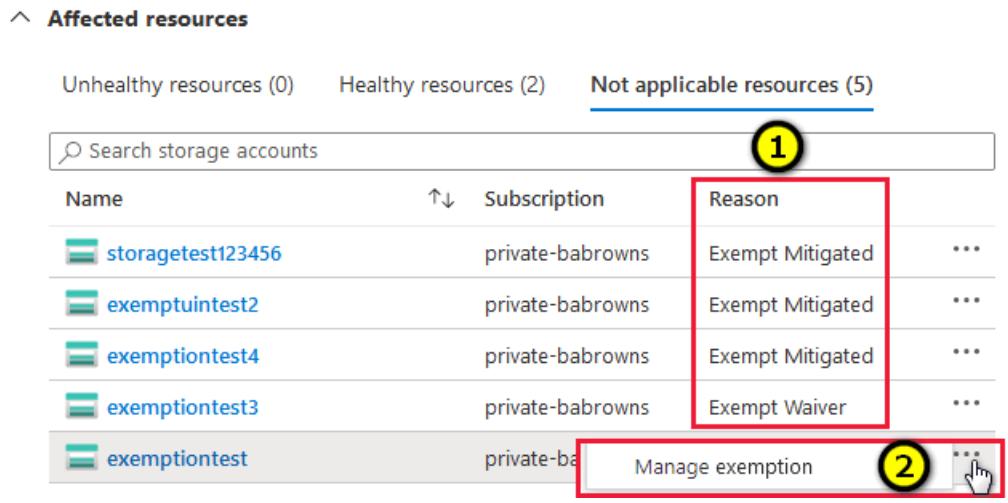

   The reason for each exemption is included in the table (1).

   To modify or delete an exemption, select the ellipsis menu ("...") as shown (2).

5. To review all of the exemption rules on your subscription, select **View exemptions** from the information strip:

> **IMPORTANT**
> To see the specific exemptions relevant to one recommendation, filter the list according to the relevant scope and recommendation name.

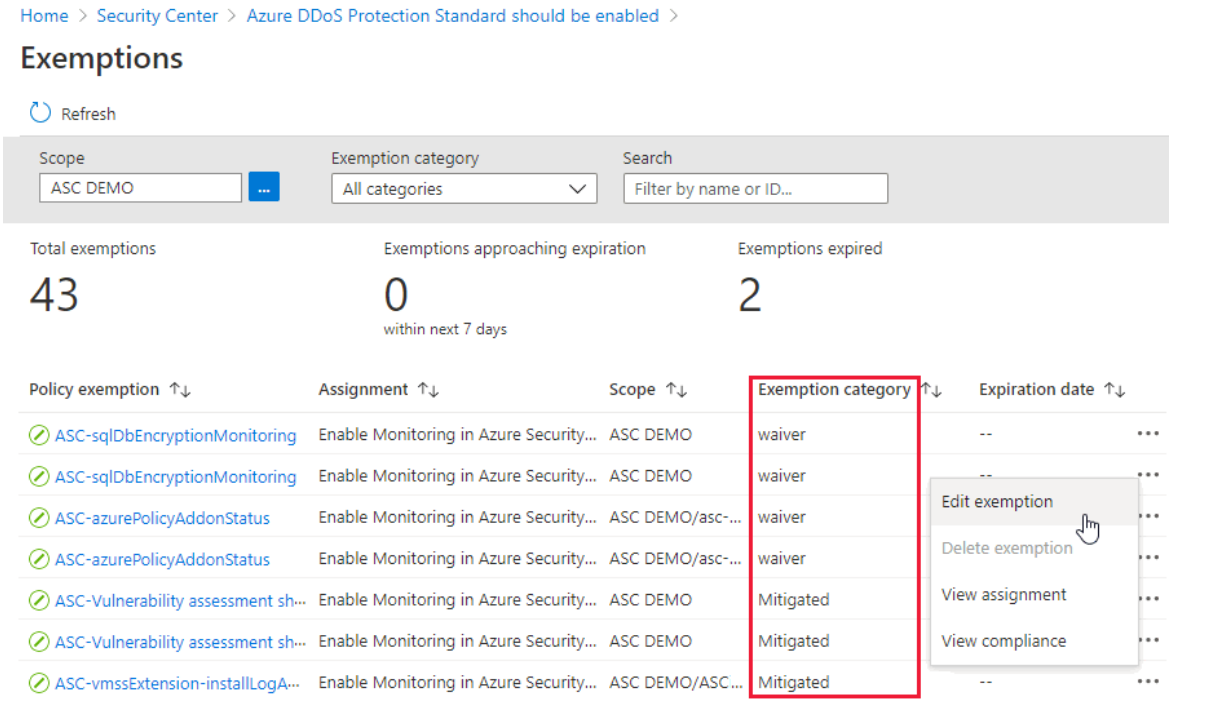

TIP

Alternatively, use Azure Resource Graph to find recommendations with exemptions.

Monitor exemptions created in your subscriptions

As explained earlier on this page, exemption rules are a powerful tool providing granular control over the recommendations affecting resources in your subscriptions and management groups.

To keep track of how your users are exercising this capability, we've created an Azure Resource Manager (ARM) template that deploys a Logic App Playbook and all necessary API connections to notify you when an exemption has been created.

- To learn more about the playbook, see the tech community blog post How to keep track of Resource Exemptions in Microsoft Defender for Cloud
- You'll find the ARM template in the Microsoft Defender for Cloud GitHub repository
- To deploy all the necessary components, use this automated process

Use the inventory to find resources that have exemptions applied

The asset inventory page of Microsoft Defender for Cloud provides a single page for viewing the security posture of the resources you've connected to Defender for Cloud. Learn more in Explore and manage your resources with asset inventory.

The inventory page includes many filters to let you narrow the list of resources to the ones of most interest for any given scenario. One such filter is the **Contains exemptions**. Use this filter to find all resources that have been exempted from one or more recommendation.

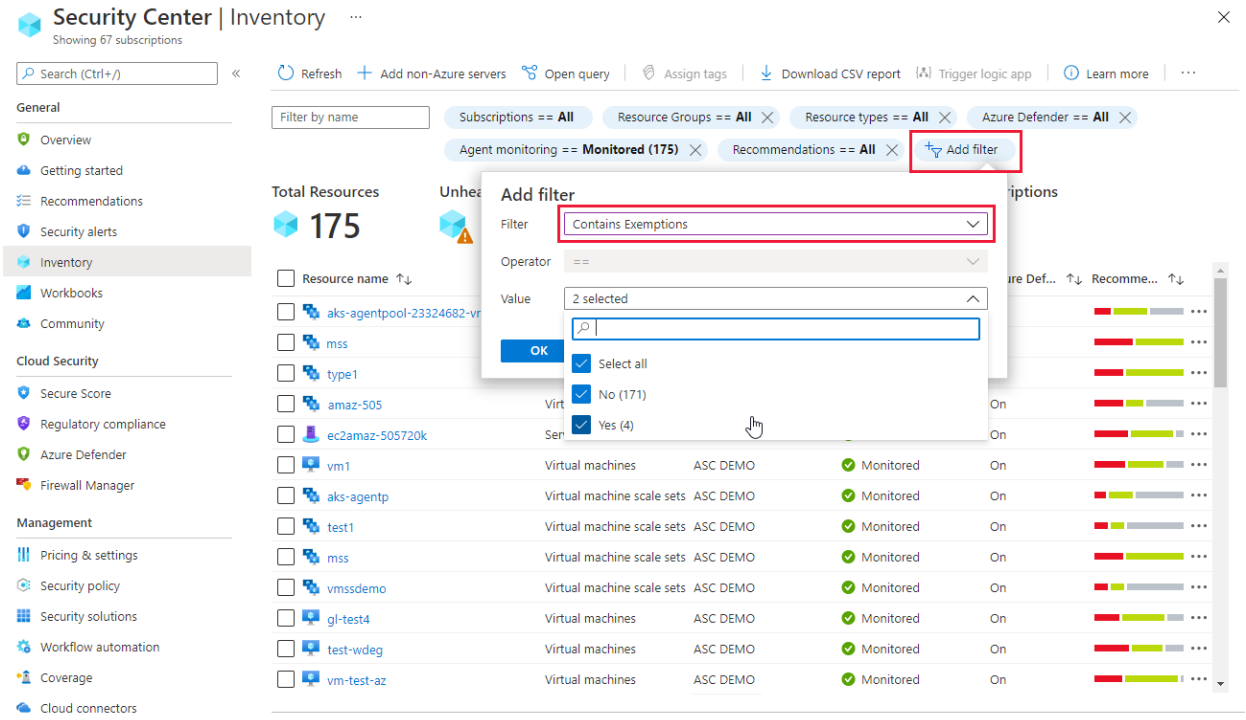

Find recommendations with exemptions using Azure Resource Graph

Azure Resource Graph (ARG) provides instant access to resource information across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal.

To view all recommendations that have exemption rules:

1. Open **Azure Resource Graph Explorer**.

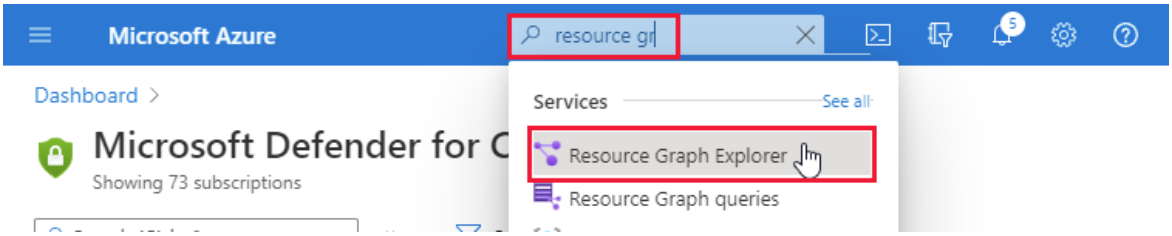

2. Enter the following query and select **Run query**.

```
securityresources
| where type == "microsoft.security/assessments"
// Get recommendations in useful format
| project
 ['TenantID'] = tenantId,
 ['SubscriptionID'] = subscriptionId,
 ['AssessmentID'] = name,
 ['DisplayName'] = properties.displayName,
 ['ResourceType'] = tolower(split(properties.resourceDetails.Id,"/").[7]),
 ['ResourceName'] = tolower(split(properties.resourceDetails.Id,"/").[8]),
 ['ResourceGroup'] = resourceGroup,
 ['ContainsNestedRecom'] = tostring(properties.additionalData.subAssessmentsLink),
 ['StatusCode'] = properties.status.code,
 ['StatusDescription'] = properties.status.description,
 ['PolicyDefID'] = properties.metadata.policyDefinitionId,
 ['Description'] = properties.metadata.description,
 ['RecomType'] = properties.metadata.assessmentType,
 ['Remediation'] = properties.metadata.remediationDescription,
 ['Severity'] = properties.metadata.severity,
 ['Link'] = properties.links.azurePortal
 | where StatusDescription contains "Exempt"
```

Learn more in the following pages:

- Learn more about Azure Resource Graph.
- How to create queries with Azure Resource Graph Explorer
- Kusto Query Language (KQL)

FAQ - Exemption rules

- What happens when one recommendation is in multiple policy initiatives?
- Are there any recommendations that don't support exemption?

What happens when one recommendation is in multiple policy initiatives?

Sometimes, a security recommendation appears in more than one policy initiative. If you've got multiple instances of the same recommendation assigned to the same subscription, and you create an exemption for the recommendation, it will affect all of the initiatives that you have permission to edit.

For example, the recommendation **** is part of the default policy initiative assigned to all Azure subscriptions by Microsoft Defender for Cloud. It's also in XXXXX.

If you try to create an exemption for this recommendation, you'll see one of the two following messages:

- If you **have** the necessary permissions to edit both initiatives, you'll see:

  This recommendation is included in several policy initiatives: [initiative names separated by comma]. Exemptions will be created on all of them.

- If you **don't have** sufficient permissions on both initiatives, you'll see this message instead:

  You have limited permissions to apply the exemption on all the policy initiatives, the exemptions will be created only on the initiatives with sufficient permissions.

Are there any recommendations that don't support exemption?

These generally available recommendations don't support exemption:

- All advanced threat protection types should be enabled in SQL managed instance advanced data security settings
- All advanced threat protection types should be enabled in SQL server advanced data security settings

- Container CPU and memory limits should be enforced
- Container images should be deployed from trusted registries only
- Container with privilege escalation should be avoided
- Containers sharing sensitive host namespaces should be avoided
- Containers should listen on allowed ports only
- Default IP Filter Policy should be Deny
- Immutable (read-only) root filesystem should be enforced for containers
- IoT Devices - Open Ports On Device
- IoT Devices - Permissive firewall policy in one of the chains was found
- IoT Devices - Permissive firewall rule in the input chain was found
- IoT Devices - Permissive firewall rule in the output chain was found
- IP Filter rule large IP range
- Least privileged Linux capabilities should be enforced for containers
- Machines should be configured securely
- Overriding or disabling of containers AppArmor profile should be restricted
- Privileged containers should be avoided
- Running containers as root user should be avoided
- Services should listen on allowed ports only
- SQL servers should have an Azure Active Directory administrator provisioned
- Usage of host networking and ports should be restricted
- Usage of pod HostPath volume mounts should be restricted to a known list to restrict node access from compromised containers

Next steps

In this article, you learned how to exempt a resource from a recommendation so that it doesn't impact your secure score. For more information about secure score, see:

- Secure score in Microsoft Defender for Cloud

Create custom security initiatives and policies

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

To help secure your systems and environment, Microsoft Defender for Cloud generates security recommendations. These recommendations are based on industry best practices, which are incorporated into the generic, default security policy supplied to all customers. They can also come from Defender for Cloud's knowledge of industry and regulatory standards.

With this feature, you can add your own *custom* initiatives. You'll then receive recommendations if your environment doesn't follow the policies you create. Any custom initiatives you create will appear alongside the built-in initiatives in the regulatory compliance dashboard, as described in the tutorial Improve your regulatory compliance.

As discussed in the Azure Policy documentation, when you specify a location for your custom initiative, it must be a management group or a subscription.

> **TIP**
> For an overview of the key concepts on this page, see What are security policies, initiatives, and recommendations?.

To add a custom initiative to your subscription

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription or management group to which you would like to add a custom initiative.

   > **NOTE**
   > For your custom initiatives to be evaluated and displayed in Defender for Cloud, you must add them at the subscription level (or higher). We recommend that you select the widest scope available.

3. Open the **Security policy** page, and in the **Your custom initiatives** area, select **Add a custom initiative**.

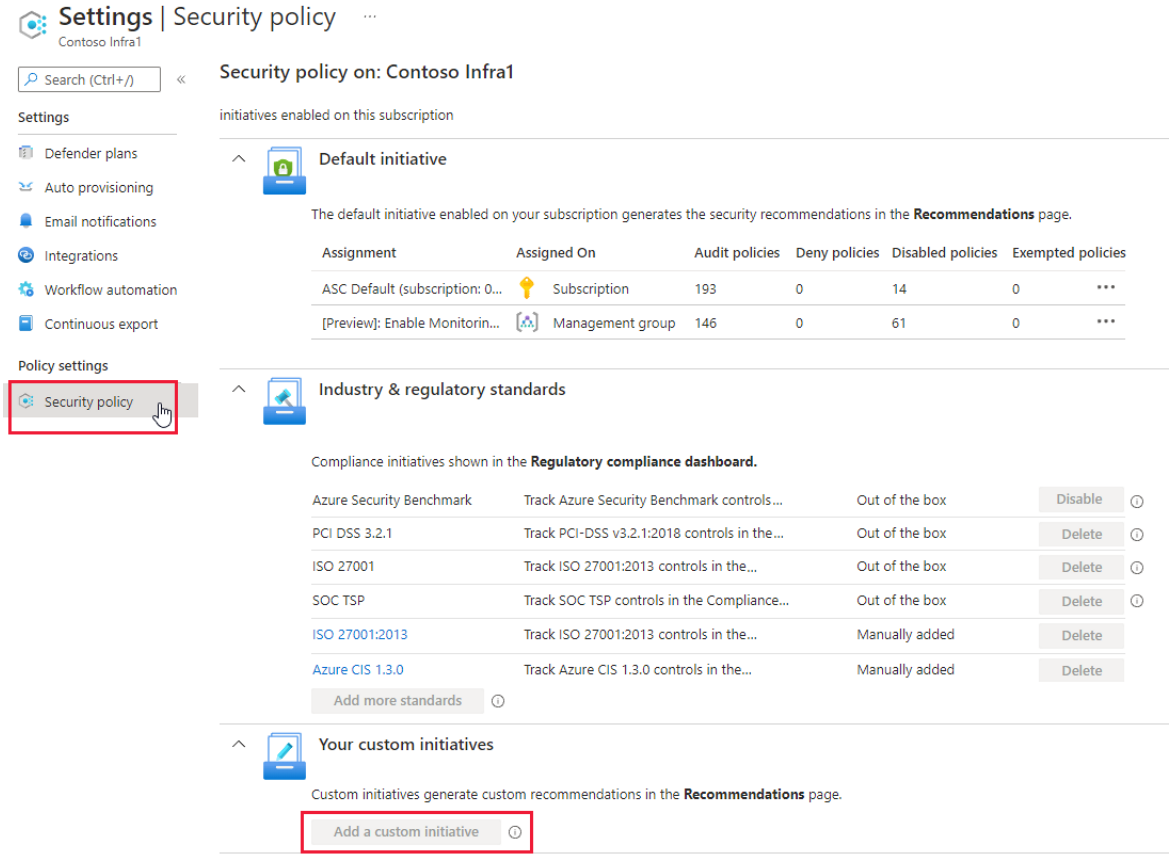

4. In the Add custom initiatives page, review the list of custom policies already created in your organization.

   - If you see one you want to assign to your subscription, select **Add**.
   - If there isn't an initiative in the list that meets your needs, create a new custom initiative:

     a. Select **Create new**.
     b. Enter the definition's location and name.
     c. Select the policies to include and select **Add**.
     d. Enter any desired parameters.
     e. Select **Save**.
     f. In the Add custom initiatives page, click refresh. Your new initiative will be available.
     g. Select **Add** and assign it to your subscription.

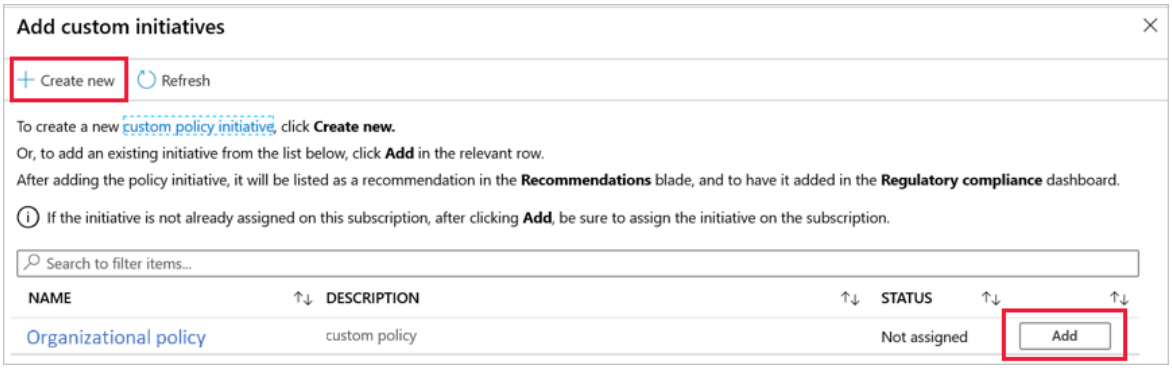

> **NOTE**
>
> Creating new initiatives requires subscription owner credentials. For more information about Azure roles, see Permissions in Microsoft Defender for Cloud.

Your new initiative takes effect and you can see the impact in the following two ways:

- From the Defender for Cloud menu, select **Regulatory compliance**. The compliance dashboard opens to show your new custom initiative alongside the built-in initiatives.
- You'll begin to receive recommendations if your environment doesn't follow the policies you've defined.

5. To see the resulting recommendations for your policy, click **Recommendations** from the sidebar to open the recommendations page. The recommendations will appear with a "Custom" label and be available within approximately one hour.

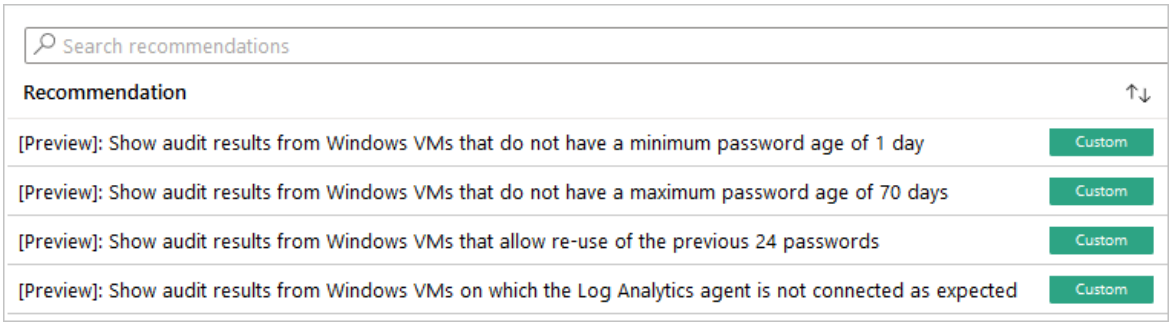

Configure a security policy in Azure Policy using the REST API

As part of the native integration with Azure Policy, Microsoft Defender for Cloud enables you to take advantage Azure Policy's REST API to create policy assignments. The following instructions walk you through creation of policy assignments, as well as customization of existing assignments.

Important concepts in Azure Policy:

- A **policy definition** is a rule
- An **initiative** is a collection of policy definitions (rules)
- An **assignment** is an application of an initiative or a policy to a specific scope (management group, subscription, etc.)

Defender for Cloud has a built-in initiative, Azure Security Benchmark, that includes all of its security policies. To assess Defender for Cloud's policies on your Azure resources, you should create an assignment on the management group, or subscription you want to assess.

The built-in initiative has all of Defender for Cloud's policies enabled by default. You can choose to disable certain policies from the built-in initiative. For example, to apply all of Defender for Cloud's policies except **web application firewall**, change the value of the policy's effect parameter to **Disabled**.

API examples

In the following examples, replace these variables:

- **{scope}** enter the name of the management group or subscription to which you're applying the policy
- **{policyAssignmentName}** enter the name of the relevant policy assignment
- **{name}** enter your name, or the name of the administrator who approved the policy change

This example shows you how to assign the built-in Defender for Cloud initiative on a subscription or management group

```
   PUT

https://management.azure.com/{scope}/providers/Microsoft.Authorization/policyAssignments/{policyAssignmentNa
me}?api-version=2018-05-01

   Request Body (JSON)

   {

     "properties":{

   "displayName":"Enable Monitoring in Microsoft Defender for Cloud",

   "metadata":{

   "assignedBy":"{Name}"

   },

   "policyDefinitionId":"/providers/Microsoft.Authorization/policySetDefinitions/1f3afdf9-d0c9-4c3d-847f-
89da613e70a8",

   "parameters":{},

   }

   }
```

This example shows you how to assign the built-in Defender for Cloud initiative on a subscription, with the following policies disabled:

- System updates (“systemUpdatesMonitoringEffect”)
- Security configurations ("systemConfigurationsMonitoringEffect")
- Endpoint protection ("endpointProtectionMonitoringEffect")

```
   PUT
https://management.azure.com/{scope}/providers/Microsoft.Authorization/policyAssignments/{policyAssignmentNa
me}?api-version=2018-05-01

   Request Body (JSON)

   {

     "properties":{

   "displayName":"Enable Monitoring in Microsoft Defender for Cloud",

   "metadata":{

   "assignedBy":"{Name}"

   },

   "policyDefinitionId":"/providers/Microsoft.Authorization/policySetDefinitions/1f3afdf9-d0c9-4c3d-847f-
89da613e70a8",

   "parameters":{

   "systemUpdatesMonitoringEffect":{"value":"Disabled"},

   "systemConfigurationsMonitoringEffect":{"value":"Disabled"},

   "endpointProtectionMonitoringEffect":{"value":"Disabled"},

   },

    }

   }
```

This example shows you how to remove an assignment:

```
   DELETE

https://management.azure.com/{scope}/providers/Microsoft.Authorization/policyAssignments/{policyAssignmentNa
me}?api-version=2018-05-01
```

Enhance your custom recommendations with detailed information

The built-in recommendations supplied with Microsoft Defender for Cloud include details such as severity levels and remediation instructions. If you want to add this type of information to your custom recommendations so that it appears in the Azure portal or wherever you access your recommendations, you'll need to use the REST API.

The two types of information you can add are:

- **RemediationDescription** – String
- **Severity** – Enum [Low, Medium, High]

The metadata should be added to the policy definition for a policy that is part of the custom initiative. It should be in the 'securityCenter' property, as shown:

```
 "metadata": {
 "securityCenter": {
  "RemediationDescription": "Custom description goes here",
  "Severity": "High"
    },
```

Below is an example of a custom policy including the metadata/securityCenter property:

```
{
"properties": {
 "displayName": "Security - ERvNet - AuditRGLock",
 "policyType": "Custom",
 "mode": "All",
 "description": "Audit required resource groups lock",
 "metadata": {
  "securityCenter": {
   "RemediationDescription": "Resource Group locks can be set via Azure Portal -> Resource Group -> Locks",
   "Severity": "High"
  }
 },
 "parameters": {
  "expressRouteLockLevel": {
   "type": "String",
   "metadata": {
    "displayName": "Lock level",
    "description": "Required lock level for ExpressRoute resource groups."
   },
   "allowedValues": [
    "CanNotDelete",
    "ReadOnly"
   ]
  }
 },
 "policyRule": {
  "if": {
   "field": "type",
   "equals": "Microsoft.Resources/subscriptions/resourceGroups"
  },
  "then": {
   "effect": "auditIfNotExists",
   "details": {
    "type": "Microsoft.Authorization/locks",
    "existenceCondition": {
     "field": "Microsoft.Authorization/locks/level",
     "equals": "[parameters('expressRouteLockLevel')]"
    }
   }
  }
 }
}
}
```

For another example of using the securityCenter property, see this section of the REST API documentation.

Next steps

In this article, you learned how to create custom security policies.

For other related material, see the following articles:

- The overview of security policies
- A list of the built-in security policies

Introduction to Microsoft Defender for servers

2/15/2022 • 7 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for servers is one of the enhanced security features of Microsoft Defender for Cloud. Use it to add threat detection and advanced defenses to your Windows and Linux machines whether they're running in Azure, on-premises, or in a multi-cloud environment.

To protect machines in hybrid and multi-cloud environments, Defender for Cloud uses Azure Arc. Connect your hybrid and multi-cloud machines as explained in the relevant quickstart:

- Connect your non-Azure machines to Microsoft Defender for Cloud
- Connect your AWS accounts to Microsoft Defender for Cloud

> **TIP**
> For details of which Defender for servers features are relevant for machines running on other cloud environments, see Supported features for virtual machines and servers.

What are the benefits of Microsoft Defender for servers?

The threat detection and protection capabilities provided with Microsoft Defender for servers include:

- **Integrated license for Microsoft Defender for Endpoint** - Microsoft Defender for servers includes Microsoft Defender for Endpoint. Together, they provide comprehensive endpoint detection and response (EDR) capabilities. For more information, see Protect your endpoints.

  When Defender for Endpoint detects a threat, it triggers an alert. The alert is shown in Defender for Cloud. From Defender for Cloud, you can also pivot to the Defender for Endpoint console, and perform a detailed investigation to uncover the scope of the attack. Learn more about Microsoft Defender for Endpoint.

  > **IMPORTANT**
  > Defender for Cloud's integration with Microsoft Defender for Endpoint is enabled by default. So when you enable Microsoft Defender for servers, you give consent for Defender for Cloud to access the Microsoft Defender for Endpoint data related to vulnerabilities, installed software, and alerts for your endpoints.
  >
  > Learn more in Protect your endpoints with Defender for Cloud's integrated EDR solution: Microsoft Defender for Endpoint.

- **Vulnerability assessment tools for machines** - Microsoft Defender for servers includes a choice of vulnerability discovery and management tools for your machines. From Defender for Cloud's settings pages, you can select which of these tools to deploy to your machines and the discovered vulnerabilities will be shown in a security recommendation.

- **Microsoft threat and vulnerability management** - Discover vulnerabilities and misconfigurations in real time with Microsoft Defender for Endpoint, and without the need of additional agents or periodic scans. Threat and vulnerability management prioritizes vulnerabilities based on the threat landscape, detections in your organization, sensitive information on vulnerable devices, and business context. Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management
- **Vulnerability scanner powered by Qualys** - Qualys' scanner is one of the leading tools for real-time identification of vulnerabilities in your Azure and hybrid virtual machines. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Defender for Cloud. Learn more in Defender for Cloud's integrated Qualys scanner for Azure and hybrid machines.

- **Just-in-time (JIT) virtual machine (VM) access** - Threat actors actively hunt accessible machines with open management ports, like RDP or SSH. All of your virtual machines are potential targets for an attack. When a VM is successfully compromised, it's used as the entry point to attack further resources within your environment.

  When you enable Microsoft Defender for servers, you can use just-in-time VM access to lock down the inbound traffic to your VMs, reducing exposure to attacks while providing easy access to connect to VMs when needed. For more information, see Understanding JIT VM access.

- **File integrity monitoring (FIM)** - File integrity monitoring (FIM), also known as change monitoring, examines files and registries of operating system, application software, and others for changes that might indicate an attack. A comparison method is used to determine if the current state of the file is different from the last scan of the file. You can use this comparison to determine if valid or suspicious modifications have been made to your files.

  When you enable Microsoft Defender for servers, you can use FIM to validate the integrity of Windows files, your Windows registries, and Linux files. For more information, see File integrity monitoring in Microsoft Defender for Cloud.

- **Adaptive application controls (AAC)** - Adaptive application controls are an intelligent and automated solution for defining allowlists of known-safe applications for your machines.

  When you've enabled and configured adaptive application controls, you'll get security alerts if any application runs other than the ones you've defined as safe. For more information, see Use adaptive application controls to reduce your machines' attack surfaces.

- **Adaptive network hardening (ANH)** - Applying network security groups (NSG) to filter traffic to and from resources, improves your network security posture. However, there can still be some cases in which the actual traffic flowing through the NSG is a subset of the NSG rules defined. In these cases, further improving the security posture can be achieved by hardening the NSG rules, based on the actual traffic patterns.

  Adaptive Network Hardening provides recommendations to further harden the NSG rules. It uses a machine learning algorithm that factors in actual traffic, known trusted configuration, threat intelligence, and other indicators of compromise, and then provides recommendations to allow traffic only from specific IP/port tuples. For more information, see Improve your network security posture with adaptive network hardening.

- **Docker host hardening** - Microsoft Defender for Cloud identifies unmanaged containers hosted on IaaS Linux VMs, or other Linux machines running Docker containers. Defender for Cloud continuously assesses the configurations of these containers. It then compares them with the Center for Internet Security (CIS) Docker Benchmark. Defender for Cloud includes the entire ruleset of the CIS Docker Benchmark and alerts you if your containers don't satisfy any of the controls. For more information, see

Harden your Docker hosts.

- **Fileless attack detection** - Fileless attacks inject malicious payloads into memory to avoid detection by disk-based scanning techniques. The attacker's payload then persists within the memory of compromised processes and performs a wide range of malicious activities.

  With fileless attack detection, automated memory forensic techniques identify fileless attack toolkits, techniques, and behaviors. This solution periodically scans your machine at runtime, and extracts insights directly from the memory of processes. Specific insights include the identification of:

  - Well-known toolkits and crypto mining software
  - Shellcode, which is a small piece of code typically used as the payload in the exploitation of a software vulnerability.
  - Injected malicious executable in process memory

  Fileless attack detection generates detailed security alerts that include descriptions with process metadata such as network activity. These details accelerate alert triage, correlation, and downstream response time. This approach complements event-based EDR solutions, and provides increased detection coverage.

  For details of the fileless attack detection alerts, see the Reference table of alerts.

- **Linux auditd alerts and Log Analytics agent integration (Linux only)** - The auditd system consists of a kernel-level subsystem, which is responsible for monitoring system calls. It filters them by a specified rule set, and writes messages for them to a socket. Defender for Cloud integrates functionalities from the auditd package within the Log Analytics agent. This integration enables collection of auditd events in all supported Linux distributions, without any prerequisites.

  Log Analytics agent for Linux collects auditd records and enriches and aggregates them into events. Defender for Cloud continuously adds new analytics that use Linux signals to detect malicious behaviors on cloud and on-premises Linux machines. Similar to Windows capabilities, these analytics span across suspicious processes, dubious sign-in attempts, kernel module loading, and other activities. These activities can indicate a machine is either under attack or has been breached.

  For a list of the Linux alerts, see the Reference table of alerts.

How does Defender for servers collect data?

For Windows, Microsoft Defender for Cloud integrates with Azure services to monitor and protect your Windows-based machines. Defender for Cloud presents the alerts and remediation suggestions from all of these services in an easy-to-use format.

For Linux, Defender for Cloud collects audit records from Linux machines by using auditd, one of the most common Linux auditing frameworks.

For hybrid and multi-cloud scenarios, Defender for Cloud integrates with Azure Arc to ensure these non-Azure machines are seen as Azure resources.

Simulating alerts

You can simulate alerts by downloading one of the following playbooks:

- For Windows: Microsoft Defender for Cloud Playbook: Security Alerts
- For Linux: Microsoft Defender for Cloud Playbook: Linux Detections.

Next steps

In this article, you learned about Microsoft Defender for servers.

Enable enhanced protections

For related material, see the following page:

- Whether an alert is generated by Defender for Cloud, or received by Defender for Cloud from a different security product, you can export it. To export your alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Exporting alerts to a SIEM.

Apply Azure security baselines to machines

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

To reduce a machine's attack surface and avoid known risks, it's important to configure the operating system (OS) as securely as possible.

The Azure Security Benchmark has guidance for OS hardening which has led to security baseline documents for Windows and Linux.

Use the security recommendations described in this article to assess the machines in your environment and:

- Identify gaps in the security configurations
- Learn how to remediate those gaps

Availability

ASPECT	DETAILS
Release state:	Preview. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	Free
Prerequisites:	Machines must (1) be members of a workgroup, (2) have the Guest Configuration extension, (3) have a system-assigned managed-identity, and (4) be running a supported OS: • Windows Server 2012, 2012r2, 2016 or 2019 • Ubuntu 14.04, 16.04, 17.04, 18.04 or 20.04 • Debian 7, 8, 9, or 10 • CentOS 7 or 8 • Red Hat Enterprise Linux (RHEL) 7 or 8 • Oracle Linux 7 or 8 • SUSE Linux Enterprise Server 12
Required roles and permissions:	To install the Guest Configuration extension and its prerequisites, **write** permission is required on the relevant machines. To **view** the recommendations and explore the OS baseline data, **read** permission is required at the subscription level.
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

What are the hardening recommendations?

Microsoft Defender for Cloud includes two recommendations that check whether the configuration of Windows and Linux machines in your environment meet the Azure security baseline configurations:

- For **Windows** machines, Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration) compares the configuration with the Windows security baseline.
- For **Linux** machines, Vulnerabilities in security configuration on your Linux machines should be remediated (powered by Guest Configuration) compares the configuration with the Linux security baseline.

These recommendations use the guest configuration feature of Azure Policy to compare the OS configuration of a machine with the baseline defined in the Azure Security Benchmark.

Compare machines in your subscriptions with the OS security baselines

To compare machines with the OS security baselines:

1. From Defender for Cloud's portal pages, open the **Recommendations** page.
2. Select the relevant recommendation:
   - For **Windows** machines, Vulnerabilities in security configuration on your Windows machines should be remediated (powered by Guest Configuration)
   - For **Linux** machines, Vulnerabilities in security configuration on your Linux machines should be remediated (powered by Guest Configuration)

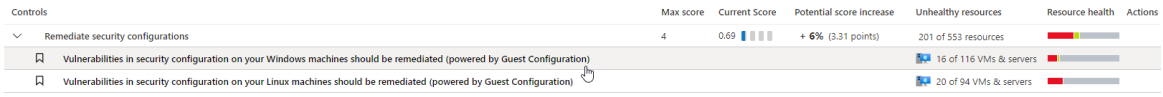

3. On the recommendation details page you can see:
   a. The affected resources.
   b. The specific security checks that failed.

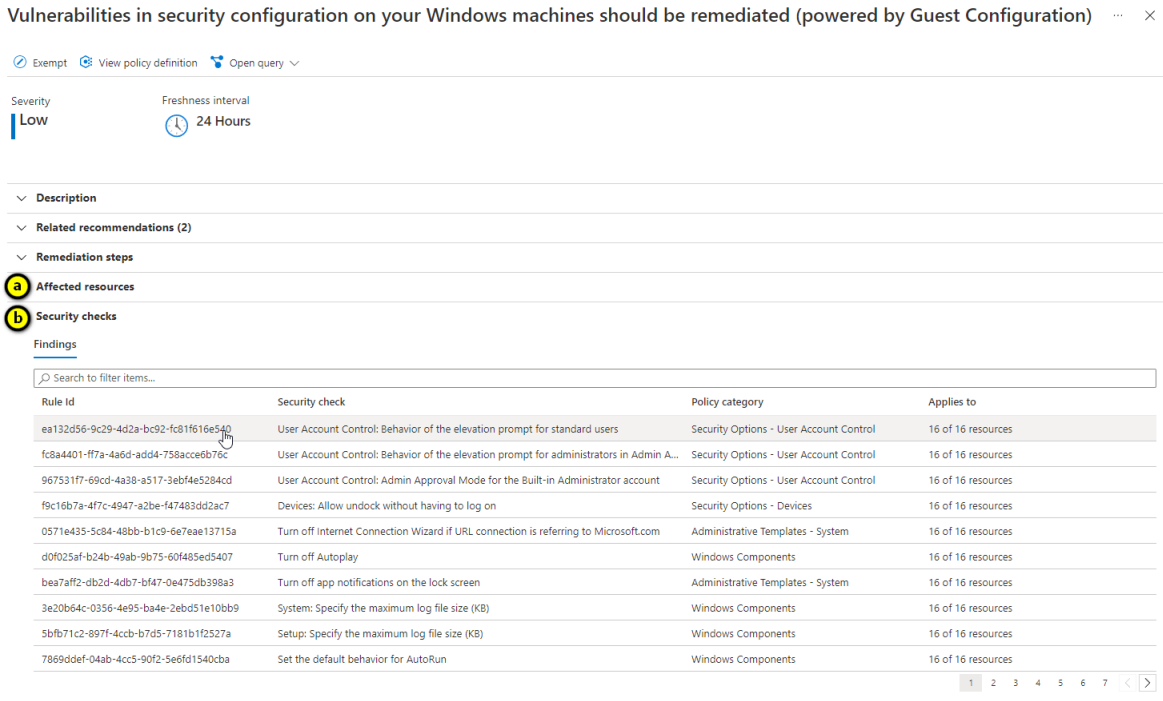

4. To learn more about a specific finding, select it.

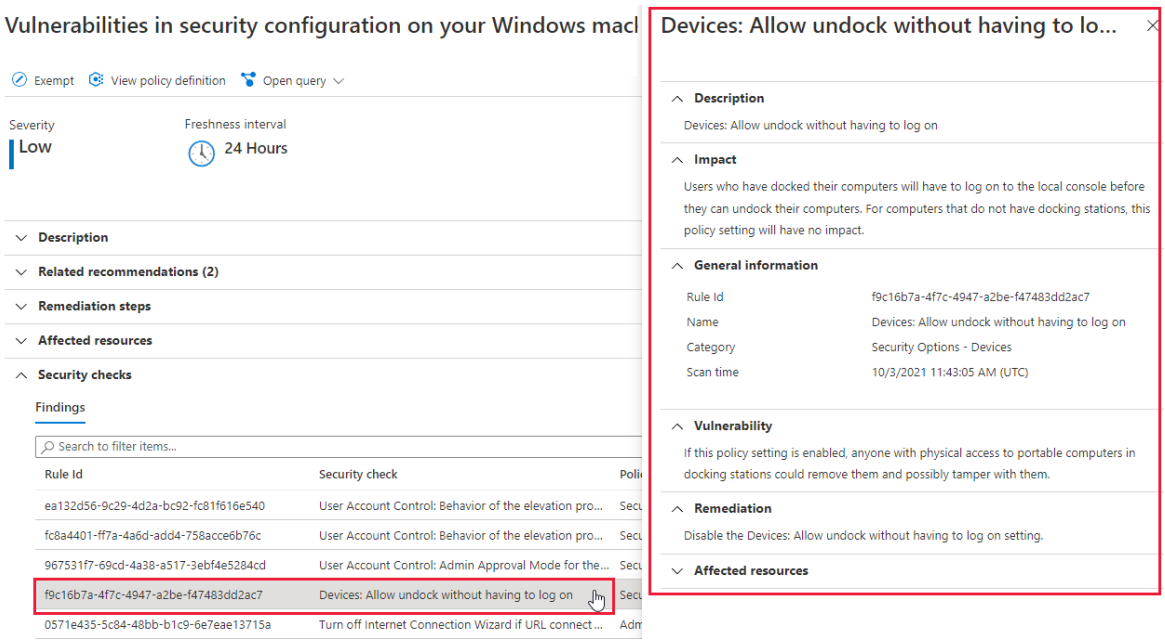

5. Other investigation possibilities:

   - To view the list of machines that have been assessed, open **Affected resources**.
   - To view the list of findings for one machine, select a machine from the **Unhealthy resources** tab. A page will open listing only the findings for that machine.

FAQ - Hardening an OS according to the security baseline

- How do I deploy the prerequisites for the security configuration recommendations?
- Why is a machine shown as not applicable?

How do I deploy the prerequisites for the security configuration recommendations?

To deploy the Guest Configuration extension with its prerequisites:

- For selected machines, follow the security recommendation **Guest Configuration extension should be installed on your machines** from the **Implement security best practices** security control.
- At scale, assign the policy initiative **Deploy prerequisites to enable Guest Configuration policies on virtual machines**.

Why is a machine shown as not applicable?

The list of resources in the **Not applicable** tab includes a **Reason** column. Some of the common reasons include:

REASON	DETAILS
No scan data available on the machine	There aren't any compliance results for this machine in Azure Resource Graph. All compliance results are written to Azure Resource Graph by the Guest Configuration extension. You can check the data in Azure Resource Graph using the sample queries in Azure Policy Guest Configuration - sample ARG queries.
Guest Configuration extension is not installed on the machine	The machine is missing the Guest Configuration extension, which is a prerequisite for assessing the compliance with the Azure security baseline.

REASON	DETAILS
System managed identity is not configured on the machine	A system-assigned, managed identity must be deployed on the machine.
The recommendation is disabled in policy	The policy definition that assesses the OS baseline is disabled on the scope that includes the relevant machine.

Next steps

In this document, you learned how to use Defender for Cloud's guest configuration recommendations to compare the hardening of your OS with the Azure security baseline.

To learn more about these configuration settings, see:

- Windows security baseline
- Linux security baseline
- Azure Security Benchmark

Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft's threat and vulnerability management is a built-in module in Microsoft Defender for Endpoint that can:

- Discover vulnerabilities and misconfigurations in near real time
- Prioritize vulnerabilities based on the threat landscape and detections in your organization

If you've enabled the integration with Microsoft Defender for Endpoint, you'll automatically get the threat and vulnerability management findings without the need for additional agents.

As it's a built-in module for Microsoft Defender for Endpoint, **threat and vulnerability management** doesn't require periodic scans.

For a quick overview of threat and vulnerability management, watch this video:

> **TIP**
> As well as alerting you to vulnerabilities, threat and vulnerability management provides additional functionality for Defender for Cloud's asset inventory tool. Learn more in Software inventory.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Machine types:	✅ Azure virtual machines ✅ Azure Arc-enabled machines Supported machines
Pricing:	Requires Microsoft Defender for servers
Prerequisites:	Enable the integration with Microsoft Defender for Endpoint
Required roles and permissions:	Owner (resource group level) can deploy the scanner Security Reader can view findings
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

Onboarding your machines to threat and vulnerability management

The integration with Microsoft Defender for Cloud doesn't involve any changes at the endpoint level: it takes place in the background between the two platforms.

- **To manually onboard one or more machines** to threat and vulnerability management, use the security recommendation "Machines should have a vulnerability assessment solution":

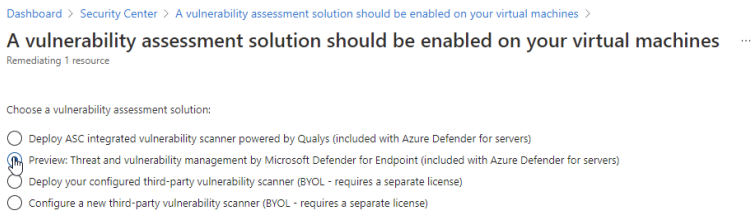

- **To automatically surface the vulnerabilities**, on existing and new machines, without the need to manually remediate the recommendation mentioned above, see Automatically configure vulnerability assessment for your machines.
- **To onboard via the REST API**, run PUT/DELETE using this URL:

```
https://management.azure.com/subscriptions/.../resourceGroups/.../providers/Microsoft.Compute/virtualMachines/.../providers/Microsoft.Security/serverV
api-version=2015-06-01-preview
```

The findings for **all** vulnerability assessment tools are provided in a Defender for Cloud recommendation

Vulnerabilities in your virtual machines should be remediated. Learn about how to View and remediate findings from vulnerability assessment solutions on your VMs

Next steps

Remediate the findings from your vulnerability assessment solution

Defender for Cloud also offers vulnerability analysis for your:

- SQL databases - see Explore vulnerability assessment reports in the vulnerability assessment dashboard
- Azure Container Registry images - see Use Microsoft Defender for container registries to scan your images for vulnerabilities

Defender for Cloud's integrated Qualys vulnerability scanner for Azure and hybrid machines

2/15/2022 • 10 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

A core component of every cyber risk and security program is the identification and analysis of vulnerabilities.

Defender for Cloud regularly checks your connected machines to ensure they're running vulnerability assessment tools.

When a machine is found that doesn't have vulnerability assessment solution deployed, Defender for Cloud generates the following security recommendation:

Machines should have a vulnerability assessment solution

Use this recommendation to deploy the vulnerability assessment solution to your Azure virtual machines and your Azure Arc-enabled hybrid machines.

Deploy the vulnerability assessment solution that best meets your needs and budget:

- **Microsoft Defender for Endpoint's threat and vulnerability management tools** - Discover vulnerabilities and misconfigurations in real time with sensors, and without the need of agents or periodic scans. It prioritizes vulnerabilities based on the threat landscape, detections in your organization, sensitive information on vulnerable devices, and business context. Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management.
- **Integrated vulnerability assessment solution (powered by Qualys)** - Defender for Cloud includes vulnerability scanning for your machines at no extra cost. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Defender for Cloud. This page provides details of this scanner and instructions for how to deploy it.

  TIP

  The integrated vulnerability assessment solution supports both Azure virtual machines and hybrid machines. To deploy the vulnerability assessment scanner to your on-premises and multi-cloud machines, connect them to Azure first with Azure Arc as described in Connect your non-Azure machines to Defender for Cloud.

  Defender for Cloud's integrated vulnerability assessment solution works seamlessly with Azure Arc. When you've deployed Azure Arc, your machines will appear in Defender for Cloud and no Log Analytics agent is required.

- **Bring your own license (BYOL) solutions** - Defender for Cloud supports the integration of tools from other vendors, but you'll need to handle the licensing costs, deployment, and configuration. By deploying your tool with Defender for Cloud, you'll get information about which Azure virtual machines are missing the tool. You'll also be able to view findings within Defender for Cloud. If you'd prefer to use your organization's private Qualys or Rapid7 license instead of the Qualys license included with Defender for Cloud, see How to deploy a BYOL solution.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Machine types (hybrid scenarios):	✅ Azure virtual machines ✅ Azure Arc-enabled machines
Pricing:	Requires Microsoft Defender for servers
Required roles and permissions:	Owner (resource group level) can deploy the scanner Security Reader can view findings
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet) ✅ Connected AWS accounts

Overview of the integrated vulnerability scanner

The vulnerability scanner included with Microsoft Defender for Cloud is powered by Qualys. Qualys' scanner is one of the leading tools for real-time identification of vulnerabilities. It's only available with Microsoft Defender for servers. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Defender for Cloud.

How the integrated vulnerability scanner works

The vulnerability scanner extension works as follows:

1. **Deploy** - Microsoft Defender for Cloud monitors your machines and provides recommendations to deploy the Qualys extension on your selected machine/s.
2. **Gather information** - The extension collects artifacts and sends them for analysis in the Qualys cloud service in the defined region.
3. **Analyze** - Qualys' cloud service conducts the vulnerability assessment and sends its findings to Defender for Cloud.

   > **IMPORTANT**
   > To ensure the privacy, confidentiality, and security of our customers, we don't share customer details with Qualys. Learn more about the privacy standards built into Azure.

4. **Report** - The findings are available in Defender for Cloud.

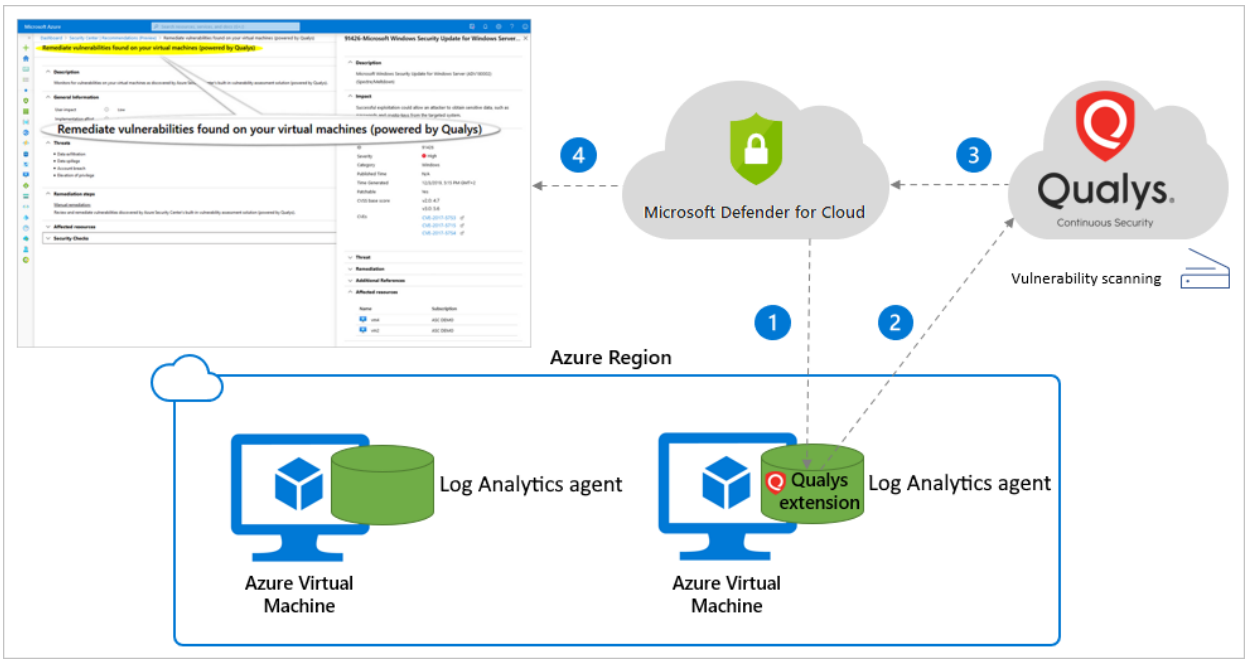

Deploy the integrated scanner to your Azure and hybrid machines

1. From the Azure portal, open **Defender for Cloud**.
2. From Defender for Cloud's menu, open the **Recommendations** page.
3. Select the recommendation **Machines should have a vulnerability assessment solution**.

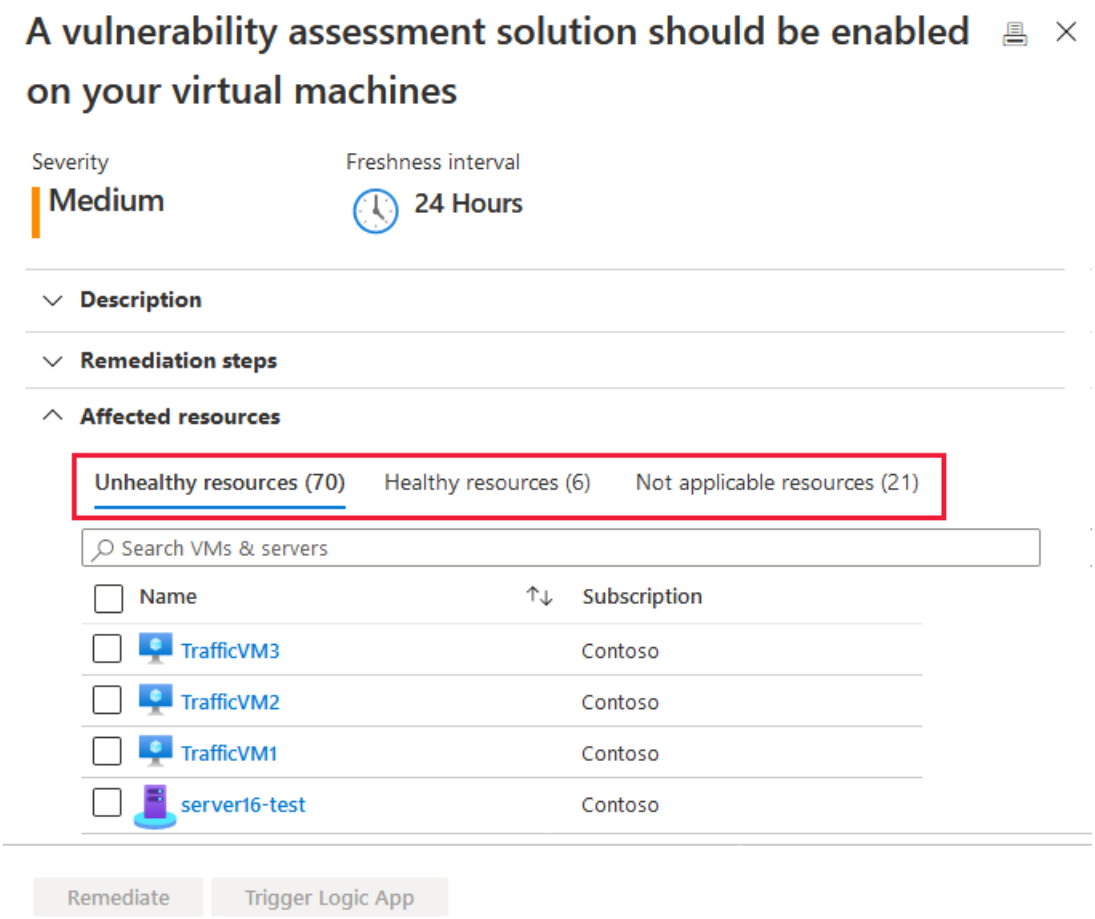

TIP

The machine "server16-test" above, is an Azure Arc-enabled machine. To deploy the vulnerability assessment scanner to your on-premises and multi-cloud machines, see Connect your non-Azure machines to Defender for Cloud.

Defender for Cloud works seamlessly with Azure Arc. When you've deployed Azure Arc, your machines will appear in Defender for Cloud and no Log Analytics agent is required.

Your machines will appear in one or more of the following groups:

- **Healthy resources** – Defender for Cloud has detected a vulnerability assessment solution running on these machines.
- **Unhealthy resources** – A vulnerability scanner extension can be deployed to these machines.
- **Not applicable resources** – these machines can't have a vulnerability scanner extension deployed. Your machine might be in this tab because it's an image in an AKS cluster, it's part of a virtual machine scale set, or it's not running one of the supported operating systems for the integrated vulnerability scanner:

VENDOR	OS	SUPPORTED VERSIONS
Microsoft	Windows	All
Amazon	Amazon Linux	2015.09-2018.03
Amazon	Amazon Linux 2	2017.03-2.0.2021
Red Hat	Enterprise Linux	5.4+, 6, 7-7.9, 8-8.3
Red Hat	CentOS	5.4+, 6, 7, 7.1-7.8, 8-8.4
Red Hat	Fedora	22-33
SUSE	Linux Enterprise Server (SLES)	11, 12, 15
SUSE	openSUSE	12, 13, 15.0-15.2
SUSE	Leap	42.1
Oracle	Enterprise Linux	5.11, 6, 7-7.9, 8-8.4
Debian	Debian	7.x-10.x
Ubuntu	Ubuntu	12.04 LTS, 14.04 LTS, 15.x, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS

4. From the list of unhealthy machines, select the ones to receive a vulnerability assessment solution and select **Remediate**.

IMPORTANT

Depending on your configuration, this list might appear differently.

- If you haven't got a third-party vulnerability scanner configured, you won't be offered the opportunity to deploy it.
- If your selected machines aren't protected by Microsoft Defender for servers, the Defender for Cloud integrated vulnerability scanner option won't be available.

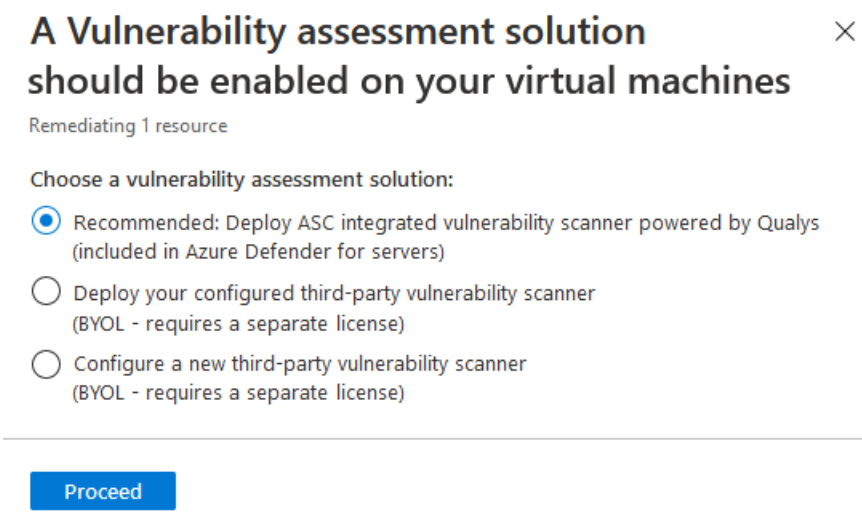

5. Choose the recommended option, **Deploy integrated vulnerability scanner**, and **Proceed**.
6. You'll be asked for one further confirmation. Select **Remediate**.

   The scanner extension will be installed on all of the selected machines within a few minutes.

   Scanning begins automatically as soon as the extension is successfully deployed. Scans will then run every 12 hours. This interval isn't configurable.

IMPORTANT

If the deployment fails on one or more machines, ensure the target machines can communicate with Qualys' cloud service by adding the following IPs to your allow lists (via port 443 - the default for HTTPS):

- `https://qagpublic.qg3.apps.qualys.com` - Qualys' US data center
- `https://qagpublic.qg2.apps.qualys.eu` - Qualys' European data center

If your machine is in a European Azure region, its artifacts will be processed in Qualys' European data center. Artifacts for virtual machines located elsewhere are sent to the US data center.

Automate at-scale deployments

NOTE

All of the tools described in this section are available from Defender for Cloud's GitHub community repository. There, you can find scripts, automations, and other useful resources to use throughout your Defender for Cloud deployment.

Some of these tools only affect new machines connected after you enable at scale deployment. Others also deploy to existing machines. You can combine multiple approaches.

Some of the ways you can automate deployment at scale of the integrated scanner:

- **Azure Resource Manager** – This method is available from **view recommendation logic** in the Azure portal. The remediation script includes the relevant ARM template you can use for your automation:

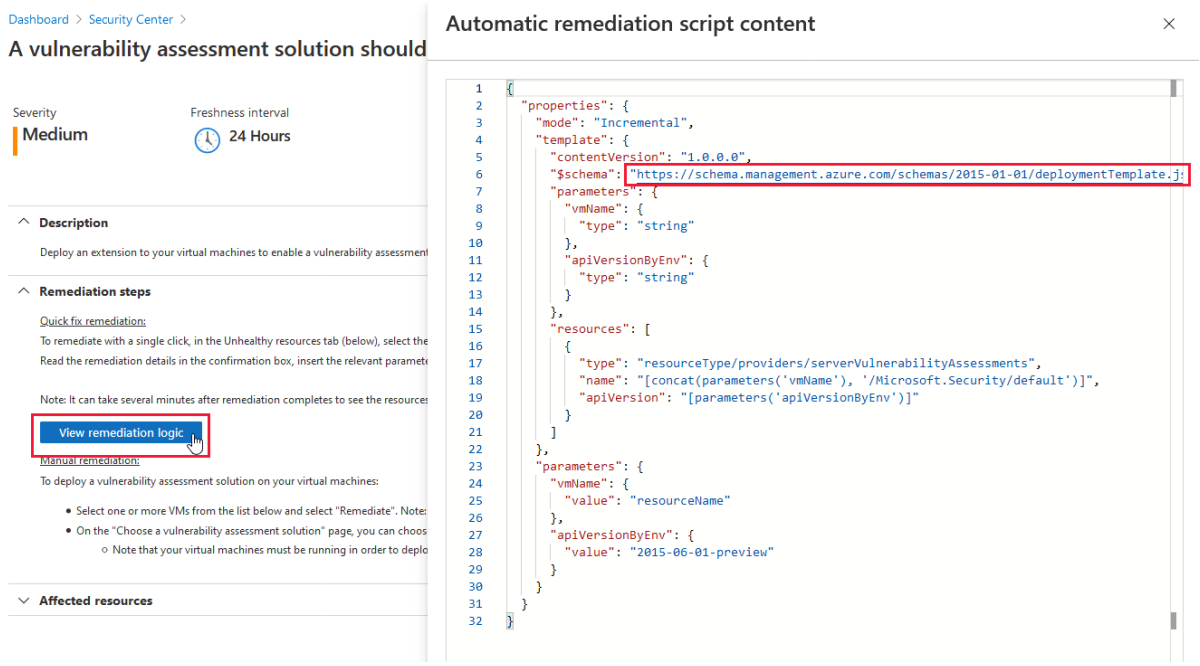

- **DeployIfNotExists policy** – A custom policy for ensuring all newly created machines receive the scanner. Select **Deploy to Azure** and set the relevant parameters. You can assign this policy at the level of resource groups, subscriptions, or management groups.
- **PowerShell Script** – Use the `Update qualys-remediate-unhealthy-vms.ps1` script to deploy the extension for all unhealthy virtual machines. To install on new resources, automate the script with Azure Automation. The script finds all unhealthy machines discovered by the recommendation and executes an Azure Resource Manager call.
- **Azure Logic Apps** – Build a logic app based on the sample app. Use Defender for Cloud's workflow automation tools to trigger your logic app to deploy the scanner whenever the **Machines should have a vulnerability assessment solution** recommendation is generated for a resource.
- **REST API** – To deploy the integrated vulnerability assessment solution using the Defender for Cloud REST API, make a PUT request for the following URL and add the relevant resource ID:

```
https://management.azure.com/<resourceId>/providers/Microsoft.Security/serverVulnerabilityAssessments/default?api-Version=2015-06-01-preview
```

Trigger an on-demand scan

You can trigger an on-demand scan from the machine itself, using locally or remotely executed scripts or Group Policy Object (GPO). Alternatively, you can integrate it into your software distribution tools at the end of a patch deployment job.

The following commands trigger an on-demand scan:

- **Windows machines**:

```
REG ADD HKLM\SOFTWARE\Qualys\QualysAgent\ScanOnDemand\Vulnerability /v "ScanOnDemand" /t REG_DWORD /d "1" /f
```

- **Linux machines**: `sudo /usr/local/qualys/cloud-agent/bin/cloudagentctl.sh action=demand type=vm`

FAQ - Integrated vulnerability scanner (powered by Qualys)

Are there any additional charges for the Qualys license?

No. The built-in scanner is free to all Microsoft Defender for servers users. The recommendation deploys the scanner with its licensing and configuration information. No additional licenses are required.

What prerequisites and permissions are required to install the Qualys extension?

You'll need write permissions for any machine on which you want to deploy the extension.

The Microsoft Defender for Cloud vulnerability assessment extension (powered by Qualys), like other extensions, runs on top of the Azure Virtual Machine agent. So it runs as Local Host on Windows, and Root on Linux.

During setup, Defender for Cloud checks to ensure that the machine can communicate with the following two Qualys data centers (via port 443 - the default for HTTPS):

- `https://qagpublic.qg3.apps.qualys.com` - Qualys' US data center
- `https://qagpublic.qg2.apps.qualys.eu` - Qualys' European data center

The extension doesn't currently accept any proxy configuration details.

Can I remove the Defender for Cloud Qualys extension?

If you want to remove the extension from a machine, you can do it manually or with any of your programmatic tools.

You'll need the following details:

- On Linux, the extension is called "LinuxAgent.AzureSecurityCenter" and the publisher name is "Qualys"
- On Windows, the extension is called "WindowsAgent.AzureSecurityCenter" and the provider name is "Qualys"

How does the extension get updated?

Like the Microsoft Defender for Cloud agent itself and all other Azure extensions, minor updates of the Qualys scanner might automatically happen in the background. All agents and extensions are tested extensively before being automatically deployed.

Why does my machine show as "not applicable" in the recommendation?

The recommendation details page groups your machines into the following lists: **healthy**, **unhealthy**, and **not applicable**.

If you have machines in the **not applicable** resources group, it means Defender for Cloud can't deploy the vulnerability scanner extension on those machines.

Your machine might be in this tab because:

- It's not protected by Defender for Cloud - As explained above, the vulnerability scanner included with Microsoft Defender for Cloud is only available for machines protected by Microsoft Defender for servers.
- It's an image in an AKS cluster or part of a virtual machine scale set - This extension doesn't support VMs that are PaaS resources.
- It's not running one of the supported operating systems:

VENDOR	OS	SUPPORTED VERSIONS
Microsoft	Windows	All
Amazon	Amazon Linux	2015.09-2018.03
Amazon	Amazon Linux 2	2017.03-2.0.2021
Red Hat	Enterprise Linux	5.4+, 6, 7-7.9, 8-8.3
Red Hat	CentOS	5.4+, 6, 7, 7.1-7.8, 8-8.4

VENDOR	OS	SUPPORTED VERSIONS
Red Hat	Fedora	22-33
SUSE	Linux Enterprise Server (SLES)	11, 12, 15
SUSE	openSUSE	12, 13, 15.0-15.2
SUSE	Leap	42.1
Oracle	Enterprise Linux	5.11, 6, 7-7.9, 8-8.4
Debian	Debian	7.x-10.x
Ubuntu	Ubuntu	12.04 LTS, 14.04 LTS, 15.x, 16.04 LTS, 18.04 LTS, 19.10, 20.04 LTS

What is scanned by the built-in vulnerability scanner?

The scanner runs on your machine to look for vulnerabilities of the machine itself. From the machine, it can't scan your network.

Does the scanner integrate with my existing Qualys console?

The Defender for Cloud extension is a separate tool from your existing Qualys scanner. Licensing restrictions mean that it can only be used within Microsoft Defender for Cloud.

How quickly will the scanner identify newly disclosed critical vulnerabilities?

Within 48 hrs of the disclosure of a critical vulnerability, Qualys incorporates the information into their processing and can identify affected machines.

Next steps

Remediate the findings from your vulnerability assessment solution

Defender for Cloud also offers vulnerability analysis for your:

- SQL databases - see Explore vulnerability assessment reports in the vulnerability assessment dashboard
- Azure Container Registry images - see Use Microsoft Defender for container registries to scan your images for vulnerabilities

Deploy a bring your own license (BYOL) vulnerability assessment solution

2/15/2022 • 5 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

If you've enabled **Microsoft Defender for servers**, you're able to use Microsoft Defender for Cloud's built-in vulnerability assessment tool as described in Integrated Qualys vulnerability scanner for virtual machines. This tool is integrated into Defender for Cloud and doesn't require any external licenses - everything's handled seamlessly inside Defender for Cloud. In addition, the integrated scanner supports Azure Arc-enabled machines.

Alternatively, you might want to deploy your own privately licensed vulnerability assessment solution from Qualys or Rapid7. You can install one of these partner solutions on multiple VMs belonging to the same subscription (but not to Azure Arc-enabled machines).

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Machine types:	✔ Azure virtual machines ✘ Azure Arc-enabled machines
Pricing:	Free
Required roles and permissions:	**Resource owner** can deploy the scanner **Security reader** can view findings
Clouds:	✔ Commercial clouds ✘ National (Azure Government, Azure China 21Vianet)

Deploy a BYOL solution from the Azure portal

The BYOL options refer to supported third-party vulnerability assessment solutions. Currently both Qualys and Rapid7 are supported providers.

Supported solutions report vulnerability data to the partner's management platform. In turn, that platform provides vulnerability and health monitoring data back to Defender for Cloud. You can identify vulnerable VMs on the workload protection dashboard and switch to the partner management console directly from Defender for Cloud for reports and more information.

1. From the Azure portal, open **Defender for Cloud**.

2. From Defender for Cloud's menu, open the **Recommendations** page.
3. Select the recommendation **Machines should have a vulnerability assessment solution**.

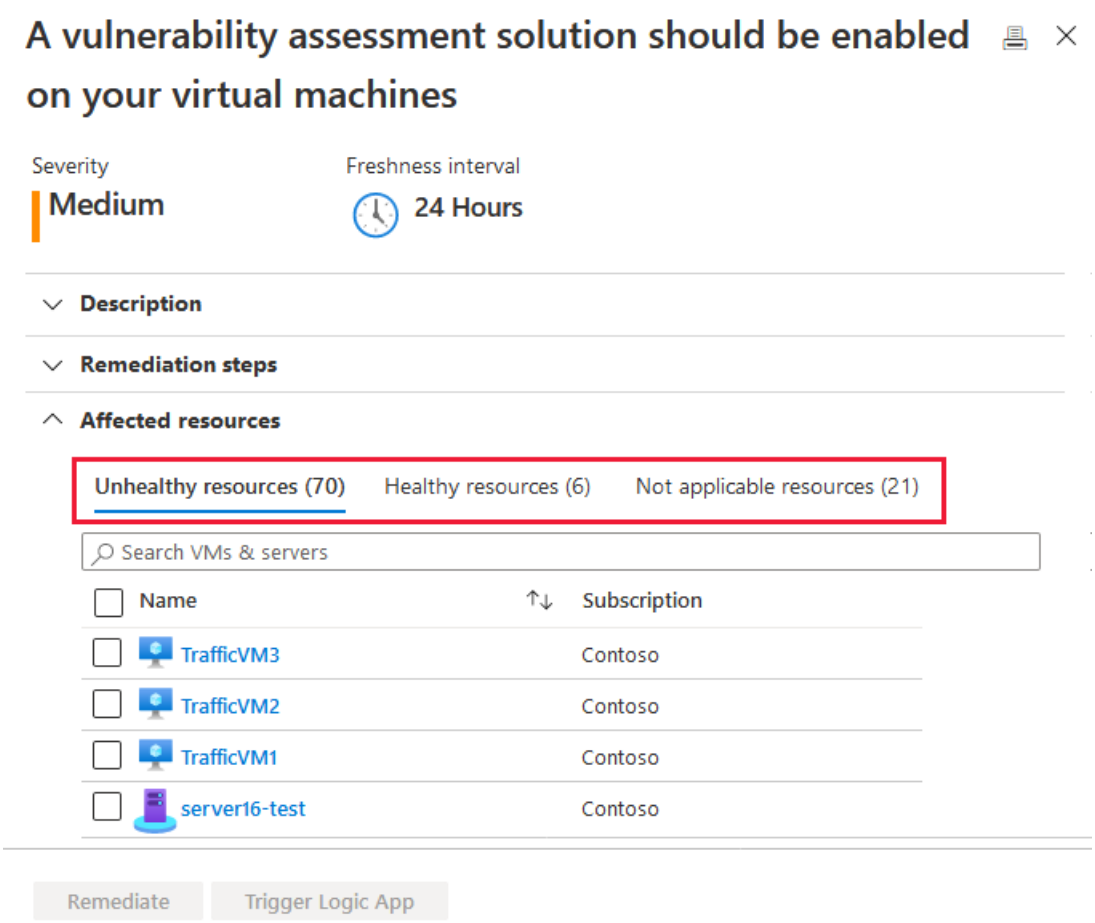

Your VMs will appear in one or more of the following groups:

- **Healthy resources** – Defender for Cloud has detected a vulnerability assessment solution running on these VMs.
- **Unhealthy resources** – A vulnerability scanner extension can be deployed to these VMs.
- **Not applicable resources** – these VMs can't have a vulnerability scanner extension deployed.

4. From the list of unhealthy machines, select the ones to receive a vulnerability assessment solution and select **Remediate**.

> **IMPORTANT**
>
> Depending on your configuration, you might only see a subset of this list.
>
> - If you haven't got a third-party vulnerability scanner configured, you won't be offered the opportunity to deploy it.
> - If your selected VMs aren't protected by Microsoft Defender for servers, the Defender for Cloud integrated vulnerability scanner option will be unavailable.

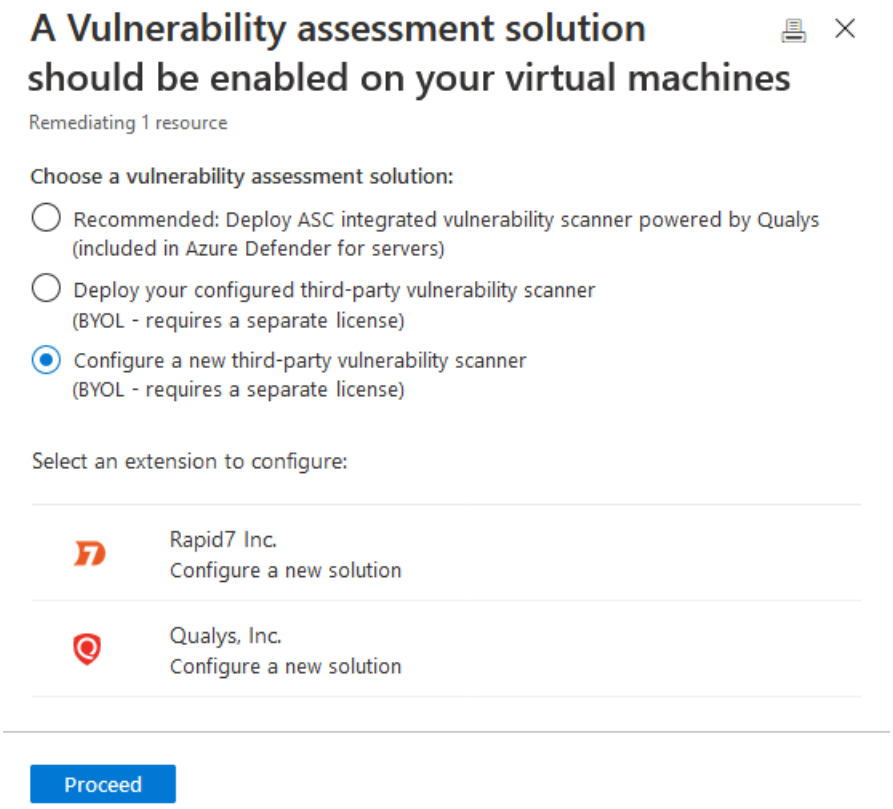

5. If you're setting up a new BYOL configuration, select **Configure a new third-party vulnerability scanner**, select the relevant extension, select **Proceed**, and enter the details from the provider as follows:

    a. For **Resource group**, select **Use existing**. If you later delete this resource group, the BYOL solution won't be available.
    b. For **Location**, select where the solution is geographically located.
    c. For Qualys, enter the license provided by Qualys into the **License code** field.
    d. For Rapid7, upload the **Rapid7 Configuration File**.
    e. In the **Public key** box, enter the public key information provided by the partner.
    f. To automatically install this vulnerability assessment agent on all discovered VMs in the subscription of this solution, select **Auto deploy**.
    g. Select **OK**.
6. If you've already set up your BYOL solution, select **Deploy your configured third-party vulnerability scanner**, select the relevant extension, and select **Proceed**.

After the vulnerability assessment solution is installed on the target machines, Defender for Cloud runs a scan to detect and identify vulnerabilities in the system and application. It might take a couple of hours for the first scan to complete. After that, it runs hourly.

Deploy a BYOL solution using PowerShell and the REST API

To programmatically deploy your own privately licensed vulnerability assessment solution from Qualys or Rapid7, use the supplied script PowerShell > Vulnerability Solution. This script uses the REST API to create a new security solution in Defender for Cloud. You'll need a license and a key provided by your service provider (Qualys or Rapid7).

> **IMPORTANT**
> Only one solution can be created per license. Attempting to create another solution using the same name/license/key will fail.

Prerequisites

Required PowerShell modules:

- Install-module Az
- Install-module Az.security

Run the script

To run the script, you'll need the relevant information for the parameters below.

PARAMETER	REQUIRED	NOTES
SubscriptionId	✔	The subscriptionID of the Azure Subscription that contains the resources you want to analyze.
ResourceGroupName	✔	Name of the resource group. Use any existing resource group including the default ("DefaultResourceGroup-xxx"). Since the solution isn't an Azure resource, it won't be listed under the resource group, but it's still attached to it. If you later delete the resource group, the BYOL solution will be unavailable.
vaSolutionName	✔	The name of the new solution.
vaType	✔	Qualys or Rapid7.
licenseCode	✔	Vendor provided license string.
publicKey	✔	Vendor provided public key.
autoUpdate	-	Enable (true) or disable (false) auto deploy for this VA solution. When enabled, every new VM on the subscription will automatically attempt to link to the solution. (Default: False)

Syntax:

```
.\New-ASCVASolution.ps1 -subscriptionId <Subscription Id> -resourceGroupName <RG Name>
-vaSolutionName <New solution name> -vaType <Qualys / Rapid7> -autoUpdate <true/false>
-licenseCode <License code from vendor> -publicKey <Public Key received from vendor>
```

Example (this example doesn't include valid license details):

```
.\New-ASCVASolution.ps1 -subscriptionId 'f4cx1b69-dtgb-4ch6-6y6f-ea2e95373d3b' -resourceGroupName
'DefaultResourceGroup-WEU' -vaSolutionName 'QualysVa001' -vaType 'Qualys' -autoUpdate 'false' `
-licenseCode
'eyJjaWQiOiJkZDg5OTYzXe4iMTMzLWM4NTAtODM5FD2mZWM1N2Q3ZGU5MjgiLCJgbTYuOiIyMmM5NDg3MS1lNTVkLTQ1OGItYjhlMC03OTR
hMmM3YWM1ZGQiLCJwd3NVcmwiOiJodHRwczovL3FhZ3B1YmxpYy1wMDEuaW50LnF1YWx5cy5jb20vQ2xvdSKJY6VudC8iLCJwd3NQb3J0Ijo
iNDQzIn0=' `
-publicKey
'MIGfMA0GCSqGSIb3DQEBAQUAA4GNADCBiQKBgQCOiOLXjOywMfLZIBGPZLwSocf1Q64GASLK9OHFEmanBl1nkJhZDrZ4YD5lM98fThYbAx1
Rde2iYV1ze/wDlX4cIvFAyXuN7HbdkeIlBl6vWXEBZpUU17bOdJOUGolzEzNBhtxi/elEZLghq9Chmah82me/okGMIhJJsCiTtglVQIDAQAB
'
```

FAQ - BYOL vulnerability scanner

- If I deploy a Qualys agent, what communications settings are required?
- Why do I have to specify a resource group when configuring a BYOL solution?

If I deploy a Qualys agent, what communications settings are required?

The Qualys Cloud Agent is designed to communicate with Qualys's SOC at regular intervals for updates, and to perform the various operations required for product functionality. To allow the agent to communicate seamlessly with the SOC, configure your network security to allow inbound and outbound traffic to the Qualys SOC CIDR and URLs.

There are multiple Qualys platforms across various geographic locations. The SOC CIDR and URLs will differ depending on the host platform of your Qualys subscription. To identify your Qualys host platform, use this page https://www.qualys.com/platform-identification/.

Why do I have to specify a resource group when configuring a BYOL solution?

When you set up your solution, you must choose a resource group to attach it to. The solution isn't an Azure resource, so it won't be included in the list of the resource group's resources. Nevertheless, it's attached to that resource group. If you later delete the resource group, the BYOL solution will be unavailable.

Next steps

Remediate the findings from your vulnerability assessment solution

Defender for Cloud also offers vulnerability analysis for your:

- SQL databases - see Explore vulnerability assessment reports in the vulnerability assessment dashboard
- Azure Container Registry images - see Use Microsoft Defender for container registries to scan your images for vulnerabilities

Automatically configure vulnerability assessment for your machines

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Defender for Cloud collects data from your machines using agents and extensions. Those agents and extensions *can* be installed manually (see Manual installation of the Log Analytics agent). However, **auto provisioning** reduces management overhead by installing all required agents and extensions on existing - and new - machines to ensure faster security coverage for all supported resources. Learn more in Configure auto provisioning for agents and extensions from Microsoft Defender for Cloud.

To assess your machines for vulnerabilities, you can use one of the following solutions:

- Microsoft's threat and vulnerability management module of Microsoft Defender for Endpoint (included with Microsoft Defender for servers)
- An integrated Qualys agent (included with Microsoft Defender for servers)
- A Qualys or Rapid7 scanner which you have licensed separately and configured within Defender for Cloud (this is called the Bring Your Own License, or BYOL, scenario)

> **NOTE**
> To automatically configure a BYOL solution, see Integrate security solutions in Microsoft Defender for Cloud.

Automatically enable a vulnerability assessment solution

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.
3. Open the **Auto provisioning** page.
4. Set the status of auto provisioning for the vulnerability assessment for machines to **On** and select the relevant solution.

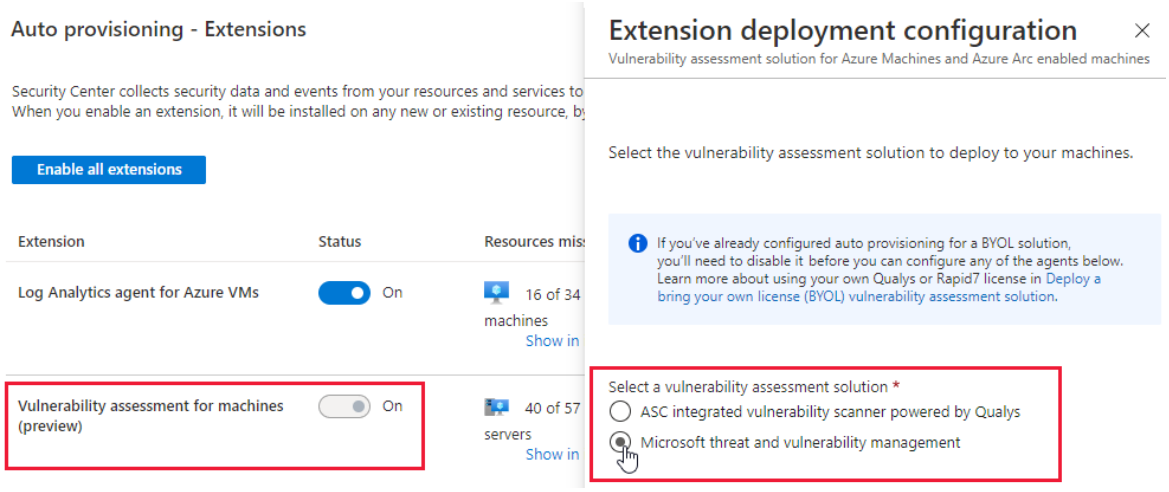

TIP

Defender for Cloud enables the following policy: (Preview) Configure machines to receive a vulnerability assessment provider.

5. Select **Apply** and **Save**.

6. To view the findings for **all** supported vulnerability assessment solutions, see the **Machines should have vulnerability findings resolved** recommendation.

   Learn more in View and remediate findings from vulnerability assessment solutions on your machines.

Next steps

Remediate the discovered vulnerabilities

Defender for Cloud also offers vulnerability assessment for your:

- SQL databases - see Explore vulnerability assessment reports in the vulnerability assessment dashboard
- Azure Container Registry images - see Use Microsoft Defender for container registries to scan your images for vulnerabilities

View and remediate findings from vulnerability assessment solutions on your VMs

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

When your vulnerability assessment tool reports vulnerabilities to Defender for Cloud, Defender for Cloud presents the findings and related information as recommendations. In addition, the findings include related information such as remediation steps, relevant CVEs, CVSS scores, and more. You can view the identified vulnerabilities for one or more subscriptions, or for a specific VM.

View findings from the scans of your virtual machines

To view vulnerability assessment findings (from all of your configured scanners) and remediate identified vulnerabilities:

1. From Defender for Cloud's menu, open the **Recommendations** page.
2. Select the recommendation **Machines should have vulnerability findings resolved**.

   Defender for Cloud shows you all the findings for all VMs in the currently selected subscriptions. The findings are ordered by severity.

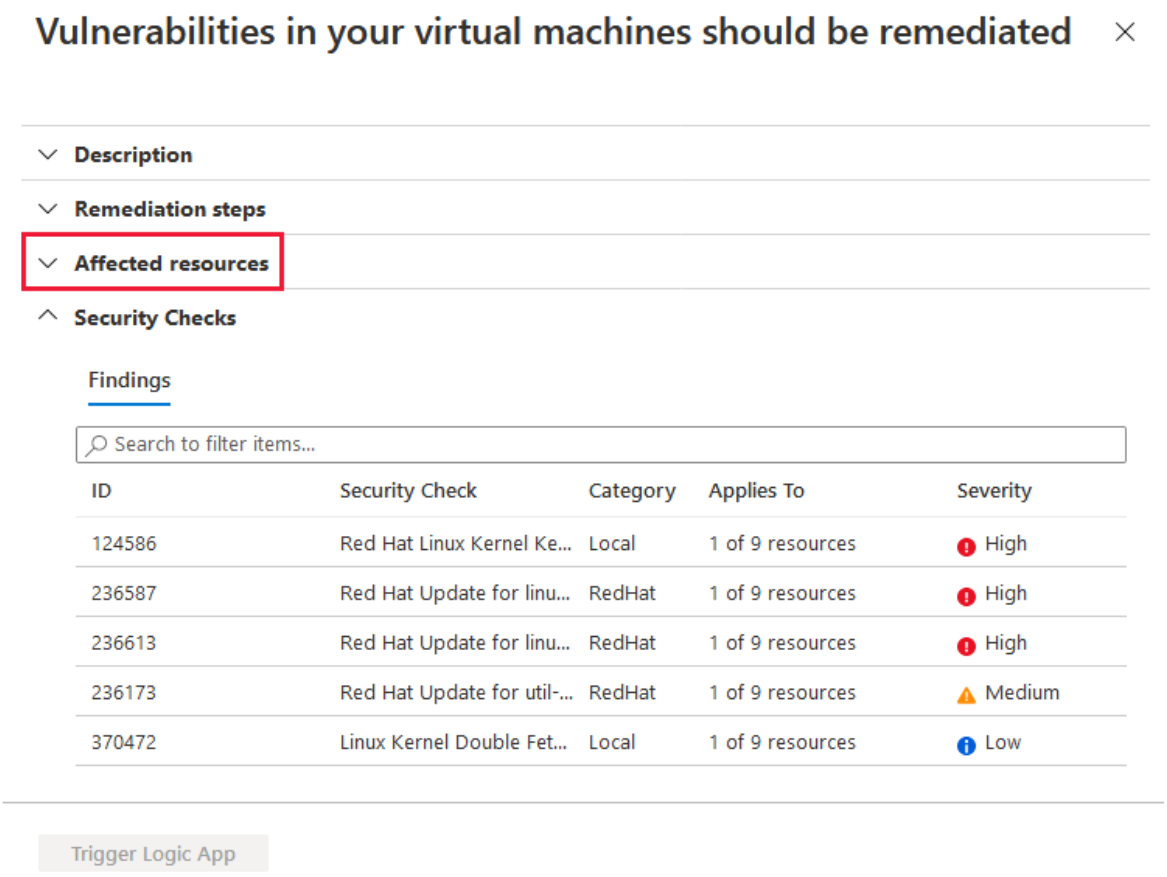

3. To filter the findings by a specific VM, open the "Affected resources" section and click the VM that interests you. Or you can select a VM from the resource health view, and view all relevant recommendations for that resource.

   Defender for Cloud shows the findings for that VM, ordered by severity.

4. To learn more about a specific vulnerability, select it.

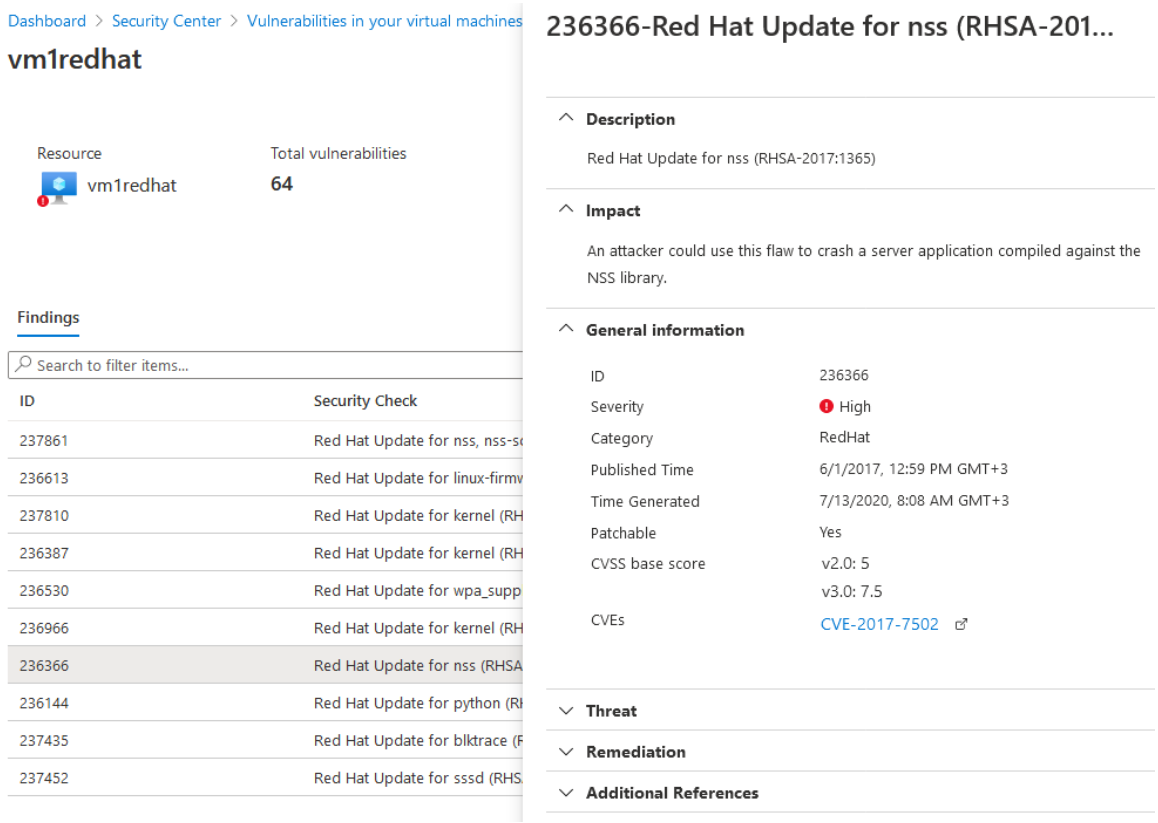

The details pane that appears contains extensive information about the vulnerability, including:

- Links to all relevant CVEs (where available)
- Remediation steps
- Any additional reference pages

5. To remediate a finding, follow the remediation steps from this details pane.

Disable specific findings

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

When a finding matches the criteria you've defined in your disable rules, it won't appear in the list of findings. Typical scenarios include:

- Disable findings with severity below medium
- Disable findings that are non-patchable
- Disable findings with CVSS score below 6.5
- Disable findings with specific text in the security check or category (for example, "RedHat", "CentOS Security Update for sudo")

> **IMPORTANT**
> To create a rule, you need permissions to edit a policy in Azure Policy. Learn more in Azure RBAC permissions in Azure Policy.

To create a rule:

1. From the recommendations detail page for Machines should have vulnerability findings resolved, select **Disable rule**.
2. Select the relevant scope.
3. Define your criteria. You can use any of the following criteria:
   - Finding ID
   - Category
   - Security check
   - CVSS scores (v2, v3)
   - Severity
   - Patchable status
4. Select **Apply rule**.

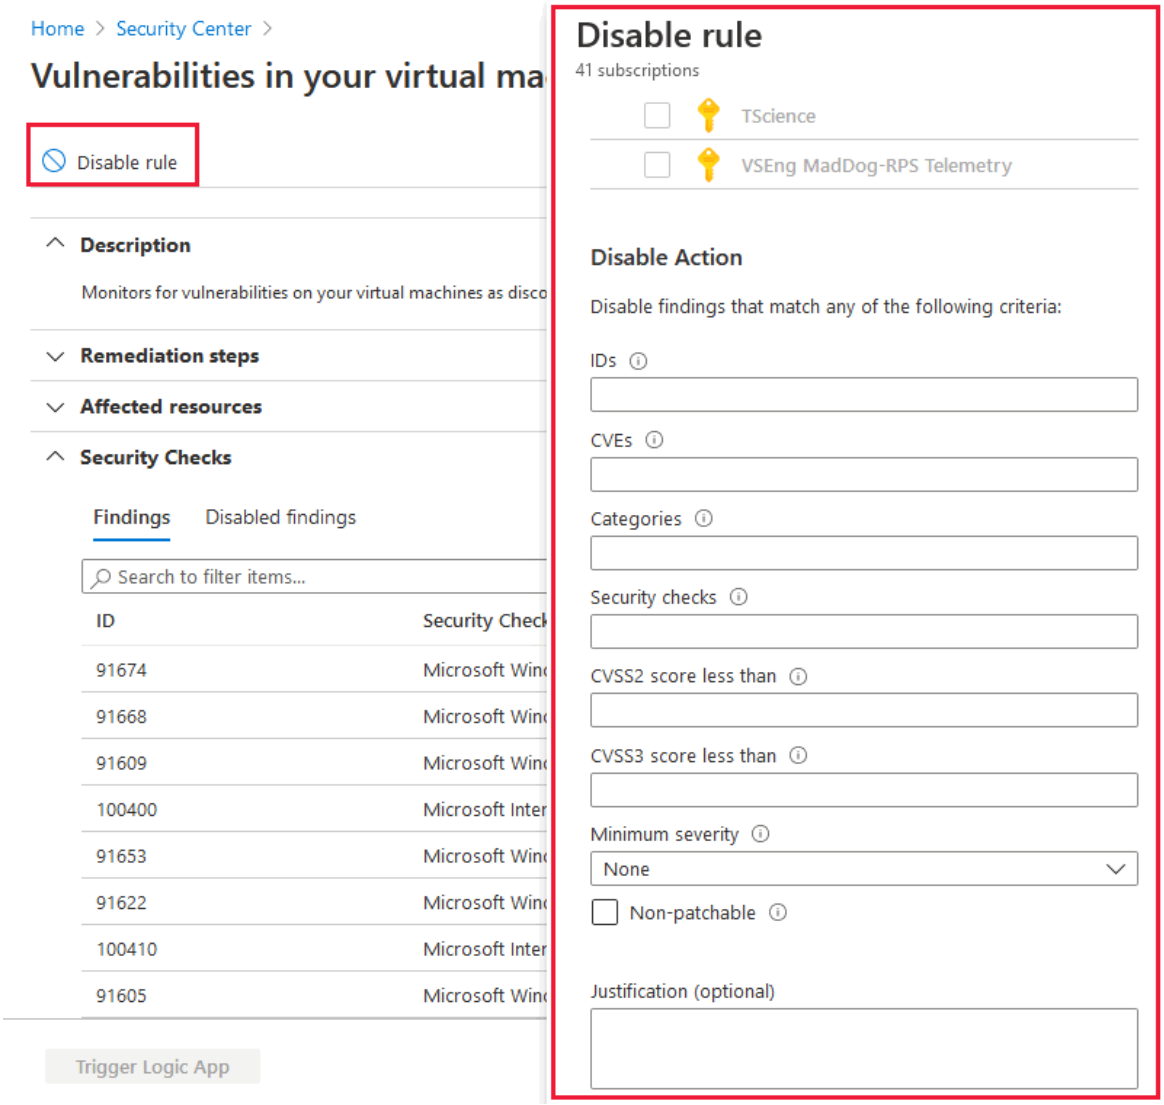

IMPORTANT

Changes might take up to 24hrs to take effect.

5. To view, override, or delete a rule:

   a. Select **Disable rule**.

   b. From the scope list, subscriptions with active rules show as **Rule applied**.

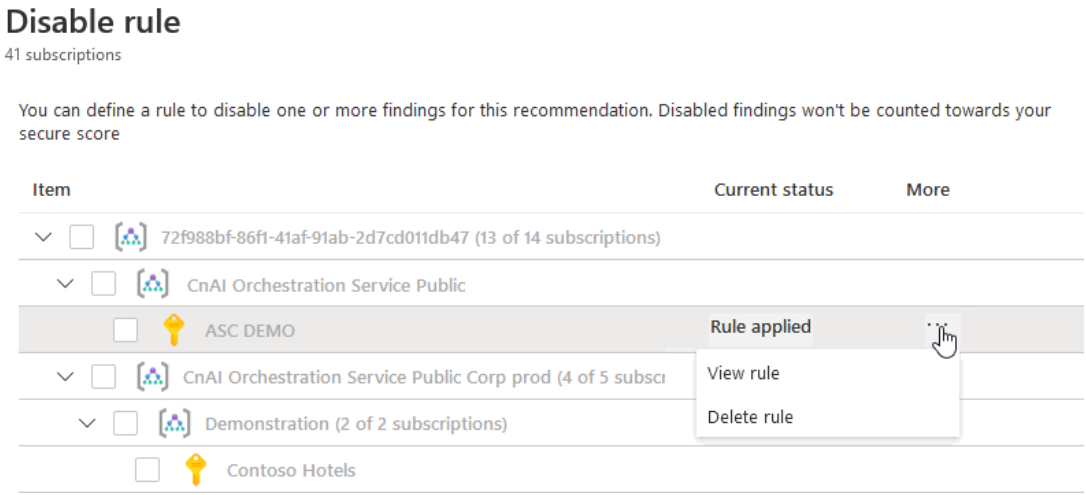

   c. To view or delete the rule, select the ellipsis menu ("...").

Export the results

To export vulnerability assessment results, you'll need to use Azure Resource Graph (ARG). This tool provides instant access to resource information across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal.

For full instructions and a sample ARG query, see the following Tech Community post: Exporting vulnerability assessment results in Microsoft Defender for Cloud.

Next steps

This article described the Microsoft Defender for Cloud vulnerability assessment extension (powered by Qualys) for scanning your VMs. For related material, see the following articles:

- Learn about the different elements of a recommendation
- Learn how to remediate recommendations

View and remediate findings from vulnerability assessment solutions on your VMs

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

When your vulnerability assessment tool reports vulnerabilities to Defender for Cloud, Defender for Cloud presents the findings and related information as recommendations. In addition, the findings include related information such as remediation steps, relevant CVEs, CVSS scores, and more. You can view the identified vulnerabilities for one or more subscriptions, or for a specific VM.

View findings from the scans of your virtual machines

To view vulnerability assessment findings (from all of your configured scanners) and remediate identified vulnerabilities:

1. From Defender for Cloud's menu, open the **Recommendations** page.
2. Select the recommendation **Machines should have vulnerability findings resolved**.

   Defender for Cloud shows you all the findings for all VMs in the currently selected subscriptions. The findings are ordered by severity.

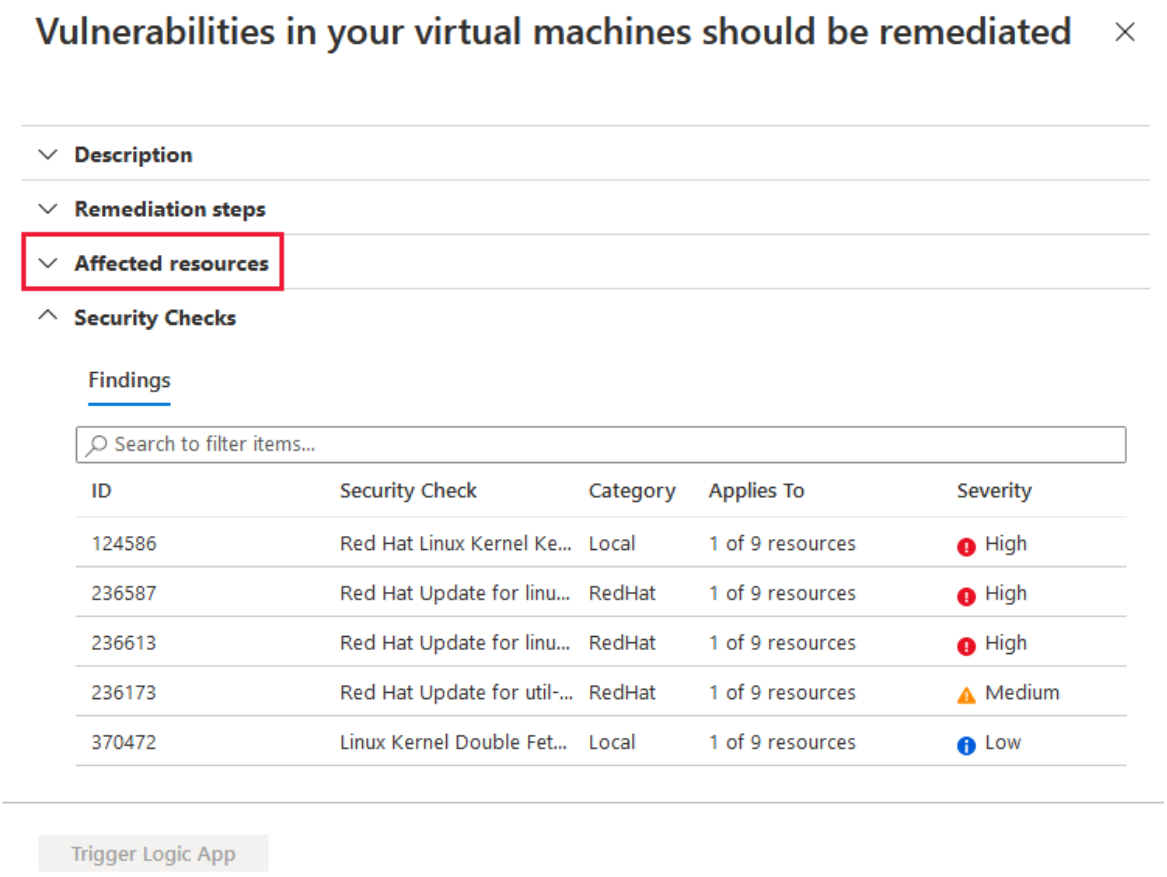

3. To filter the findings by a specific VM, open the "Affected resources" section and click the VM that interests you. Or you can select a VM from the resource health view, and view all relevant recommendations for that resource.

   Defender for Cloud shows the findings for that VM, ordered by severity.

4. To learn more about a specific vulnerability, select it.

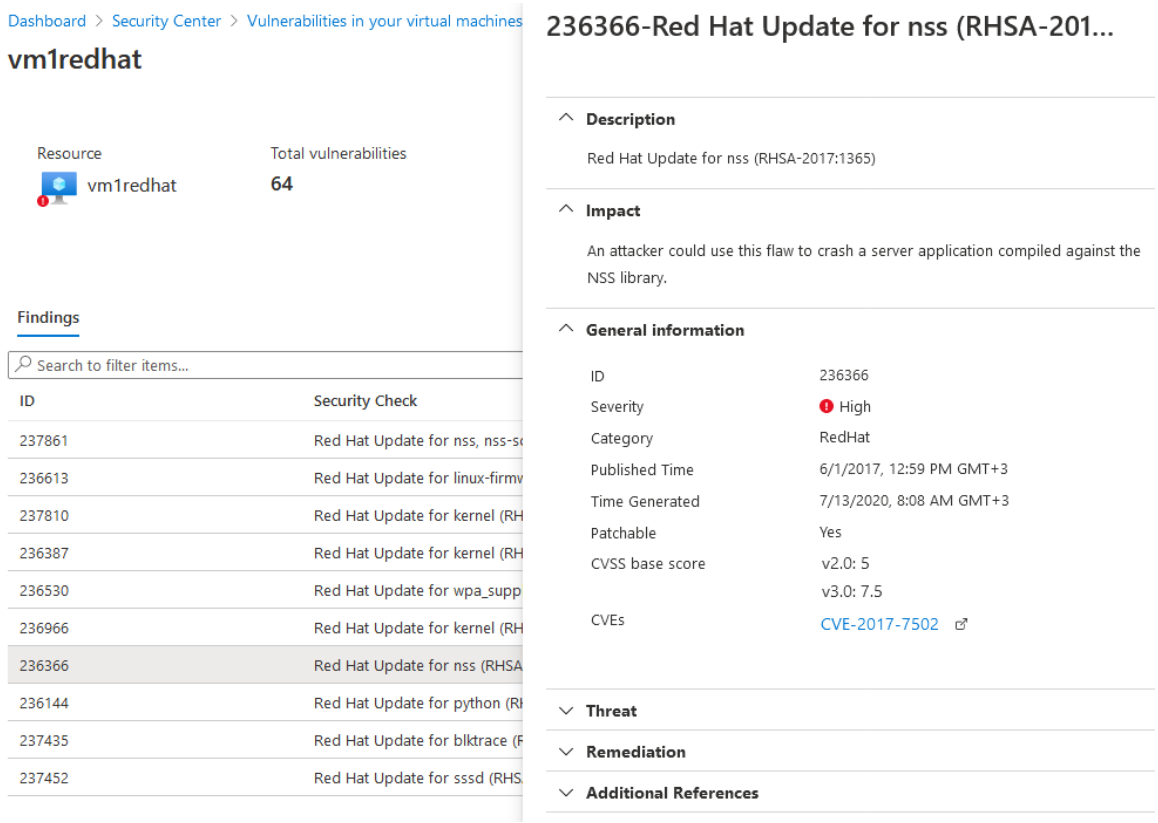

The details pane that appears contains extensive information about the vulnerability, including:

- Links to all relevant CVEs (where available)
- Remediation steps
- Any additional reference pages

5. To remediate a finding, follow the remediation steps from this details pane.

Disable specific findings

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

When a finding matches the criteria you've defined in your disable rules, it won't appear in the list of findings. Typical scenarios include:

- Disable findings with severity below medium
- Disable findings that are non-patchable
- Disable findings with CVSS score below 6.5
- Disable findings with specific text in the security check or category (for example, "RedHat", "CentOS Security Update for sudo")

> **IMPORTANT**
> To create a rule, you need permissions to edit a policy in Azure Policy. Learn more in Azure RBAC permissions in Azure Policy.

To create a rule:

1. From the recommendations detail page for Machines should have vulnerability findings resolved, select **Disable rule**.
2. Select the relevant scope.
3. Define your criteria. You can use any of the following criteria:
   - Finding ID
   - Category
   - Security check
   - CVSS scores (v2, v3)
   - Severity
   - Patchable status
4. Select **Apply rule**.

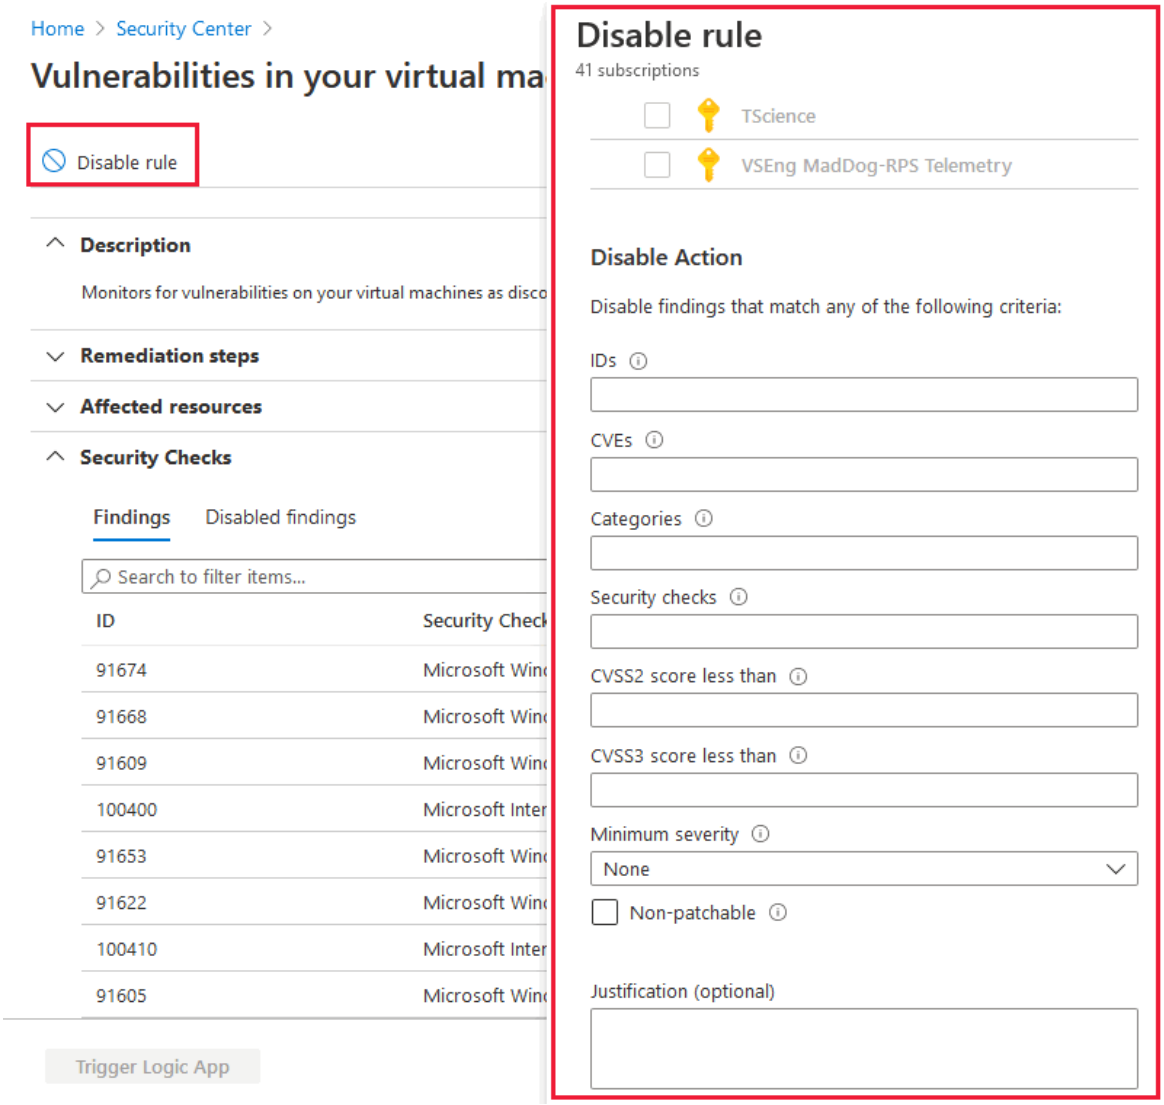

IMPORTANT

Changes might take up to 24hrs to take effect.

5. To view, override, or delete a rule:

   a. Select **Disable rule**.

   b. From the scope list, subscriptions with active rules show as **Rule applied**.

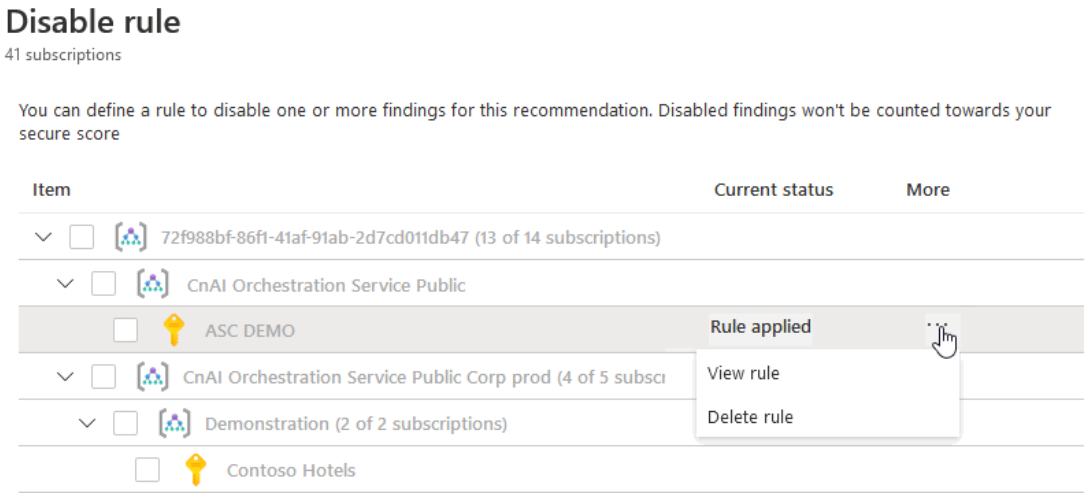

   c. To view or delete the rule, select the ellipsis menu ("...").

Export the results

To export vulnerability assessment results, you'll need to use Azure Resource Graph (ARG). This tool provides instant access to resource information across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal.

For full instructions and a sample ARG query, see the following Tech Community post: Exporting vulnerability assessment results in Microsoft Defender for Cloud.

Next steps

This article described the Microsoft Defender for Cloud vulnerability assessment extension (powered by Qualys) for scanning your VMs. For related material, see the following articles:

- Learn about the different elements of a recommendation
- Learn how to remediate recommendations

Understanding just-in-time (JIT) VM access

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains the principles behind Microsoft Defender for Cloud's just-in-time (JIT) VM access feature and the logic behind the recommendation.

To learn how to apply JIT to your VMs using the Azure portal (either Defender for Cloud or Azure Virtual Machines) or programmatically, see How to secure your management ports with JIT.

The risk of open management ports on a virtual machine

Threat actors actively hunt accessible machines with open management ports, like RDP or SSH. All of your virtual machines are potential targets for an attack. When a VM is successfully compromised, it's used as the entry point to attack further resources within your environment.

Why JIT VM access is the solution

As with all cybersecurity prevention techniques, your goal should be to reduce the attack surface. In this case, that means having fewer open ports, especially management ports.

Your legitimate users also use these ports, so it's not practical to keep them closed.

To solve this dilemma, Microsoft Defender for Cloud offers JIT. With JIT, you can lock down the inbound traffic to your VMs, reducing exposure to attacks while providing easy access to connect to VMs when needed.

How JIT operates with network security groups and Azure Firewall

When you enable just-in-time VM access, you can select the ports on the VM to which inbound traffic will be blocked. Defender for Cloud ensures "deny all inbound traffic" rules exist for your selected ports in the network security group (NSG) and Azure Firewall rules. These rules restrict access to your Azure VMs' management ports and defend them from attack.

If other rules already exist for the selected ports, then those existing rules take priority over the new "deny all inbound traffic" rules. If there are no existing rules on the selected ports, then the new rules take top priority in the NSG and Azure Firewall.

When a user requests access to a VM, Defender for Cloud checks that the user has Azure role-based access control (Azure RBAC) permissions for that VM. If the request is approved, Defender for Cloud configures the NSGs and Azure Firewall to allow inbound traffic to the selected ports from the relevant IP address (or range), for the amount of time that was specified. After the time has expired, Defender for Cloud restores the NSGs to their previous states. Connections that are already established are not interrupted.

NOTE

JIT does not support VMs protected by Azure Firewalls controlled by Azure Firewall Manager. The Azure Firewall must be configured with Rules (Classic) and cannot use Firewall policies.

How Defender for Cloud identifies which VMs should have JIT applied

The diagram below shows the logic that Defender for Cloud applies when deciding how to categorize your supported VMs:

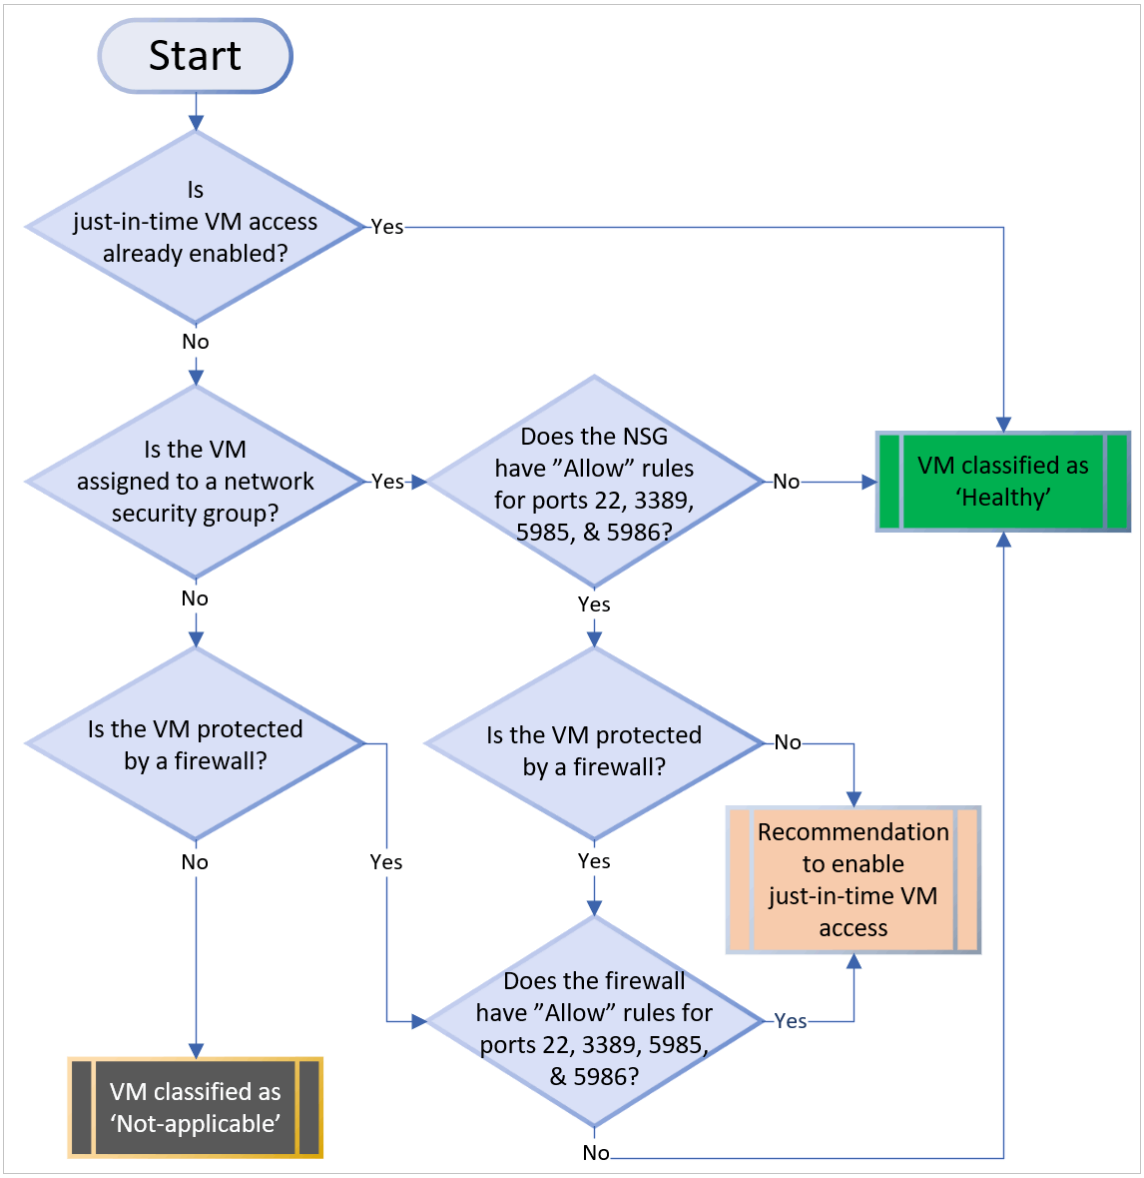

When Defender for Cloud finds a machine that can benefit from JIT, it adds that machine to the recommendation's **Unhealthy resources** tab.

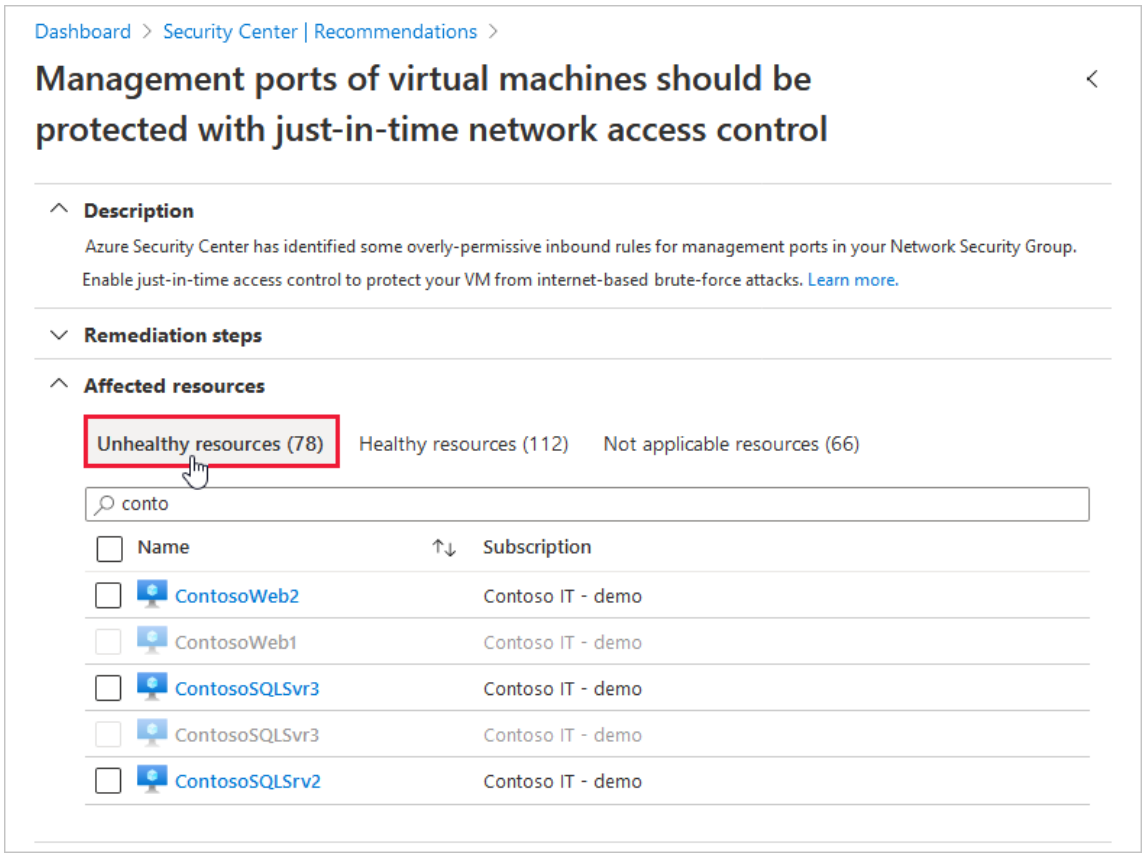

FAQ - Just-in-time virtual machine access

What permissions are needed to configure and use JIT?

JIT requires Microsoft Defender for servers to be enabled on the subscription.

Reader and **SecurityReader** roles can both view the JIT status and parameters.

If you want to create custom roles that can work with JIT, you'll need the details from the table below.

> **TIP**
> To create a least-privileged role for users that need to request JIT access to a VM, and perform no other JIT operations, use the Set-JitLeastPrivilegedRole script from the Defender for Cloud GitHub community pages.

TO ENABLE A USER TO:	PERMISSIONS TO SET
Configure or edit a JIT policy for a VM	*Assign these actions to the role:* • On the scope of a subscription or resource group that is associated with the VM: `Microsoft.Security/locations/jitNetworkAccessPolicies/write` • On the scope of a subscription or resource group of VM: `Microsoft.Compute/virtualMachines/write`

TO ENABLE A USER TO:	PERMISSIONS TO SET
Request JIT access to a VM	*Assign these actions to the user:* • On the scope of a subscription or resource group that is associated with the VM: `Microsoft.Security/locations/jitNetworkAccessPolicies/initiate/a` • On the scope of a subscription or resource group that is associated with the VM: `Microsoft.Security/locations/jitNetworkAccessPolicies/*/read` • On the scope of a subscription or resource group or VM: `Microsoft.Compute/virtualMachines/read` • On the scope of a subscription or resource group or VM: `Microsoft.Network/networkInterfaces/*/read`
Read JIT policies	*Assign these actions to the user:* • `Microsoft.Security/locations/jitNetworkAccessPolicies/read` • `Microsoft.Security/locations/jitNetworkAccessPolicies/initiate/a` • `Microsoft.Security/policies/read` • `Microsoft.Security/pricings/read` • `Microsoft.Compute/virtualMachines/read` • `Microsoft.Network/*/read`

Next steps

This page explained *why* just-in-time (JIT) virtual machine (VM) access should be used. To learn about *how* to enable JIT and request access to your JIT-enabled VMs, see the following:

How to secure your management ports with JIT

Secure your management ports with just-in-time access

2/15/2022 • 10 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Lock down inbound traffic to your Azure Virtual Machines with Microsoft Defender for Cloud's just-in-time (JIT) virtual machine (VM) access feature. This reduces exposure to attacks while providing easy access when you need to connect to a VM.

For a full explanation about how JIT works and the underlying logic, see Just-in-time explained.

For a full explanation of the privilege requirements, see What permissions are needed to configure and use JIT?.

This page teaches you how to include JIT in your security program. You'll learn how to:

- **Enable JIT on your VMs** - You can enable JIT with your own custom options for one or more VMs using Defender for Cloud, PowerShell, or the REST API. Alternatively, you can enable JIT with default, hard-coded parameters, from Azure virtual machines. When enabled, JIT locks down inbound traffic to your Azure VMs by creating a rule in your network security group.
- **Request access to a VM that has JIT enabled** - The goal of JIT is to ensure that even though your inbound traffic is locked down, Defender for Cloud still provides easy access to connect to VMs when needed. You can request access to a JIT-enabled VM from Defender for Cloud, Azure virtual machines, PowerShell, or the REST API.
- **Audit the activity** - To ensure your VMs are secured appropriately, review the accesses to your JIT-enabled VMs as part of your regular security checks.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Supported VMs:	✔ VMs deployed through Azure Resource Manager. ✘ VMs deployed with classic deployment models. Learn more about these deployment models. ✘ VMs protected by Azure Firewalls[1] controlled by Azure Firewall Manager.

ASPECT	DETAILS
Required roles and permissions:	**Reader** and **SecurityReader** roles can both view the JIT status and parameters. To create custom roles that can work with JIT, see What permissions are needed to configure and use JIT?. To create a least-privileged role for users that need to request JIT access to a VM, and perform no other JIT operations, use the Set-JitLeastPrivilegedRole script from the Defender for Cloud GitHub community pages.
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) ❌ Connected AWS accounts

[1] For any VM protected by Azure Firewall, JIT will only fully protect the machine if it's in the same VNET as the firewall. VMs using VNET peering will not be fully protected.

Enable JIT VM access

You can enable JIT VM access with your own custom options for one or more VMs using Defender for Cloud or programmatically.

Alternatively, you can enable JIT with default, hard-coded parameters, from Azure Virtual machines.

Each of these options is explained in a separate tab below.

- Microsoft Defender for Cloud
- Azure virtual machines
- PowerShell
- REST API

Enable JIT on your VMs from Microsoft Defender for Cloud

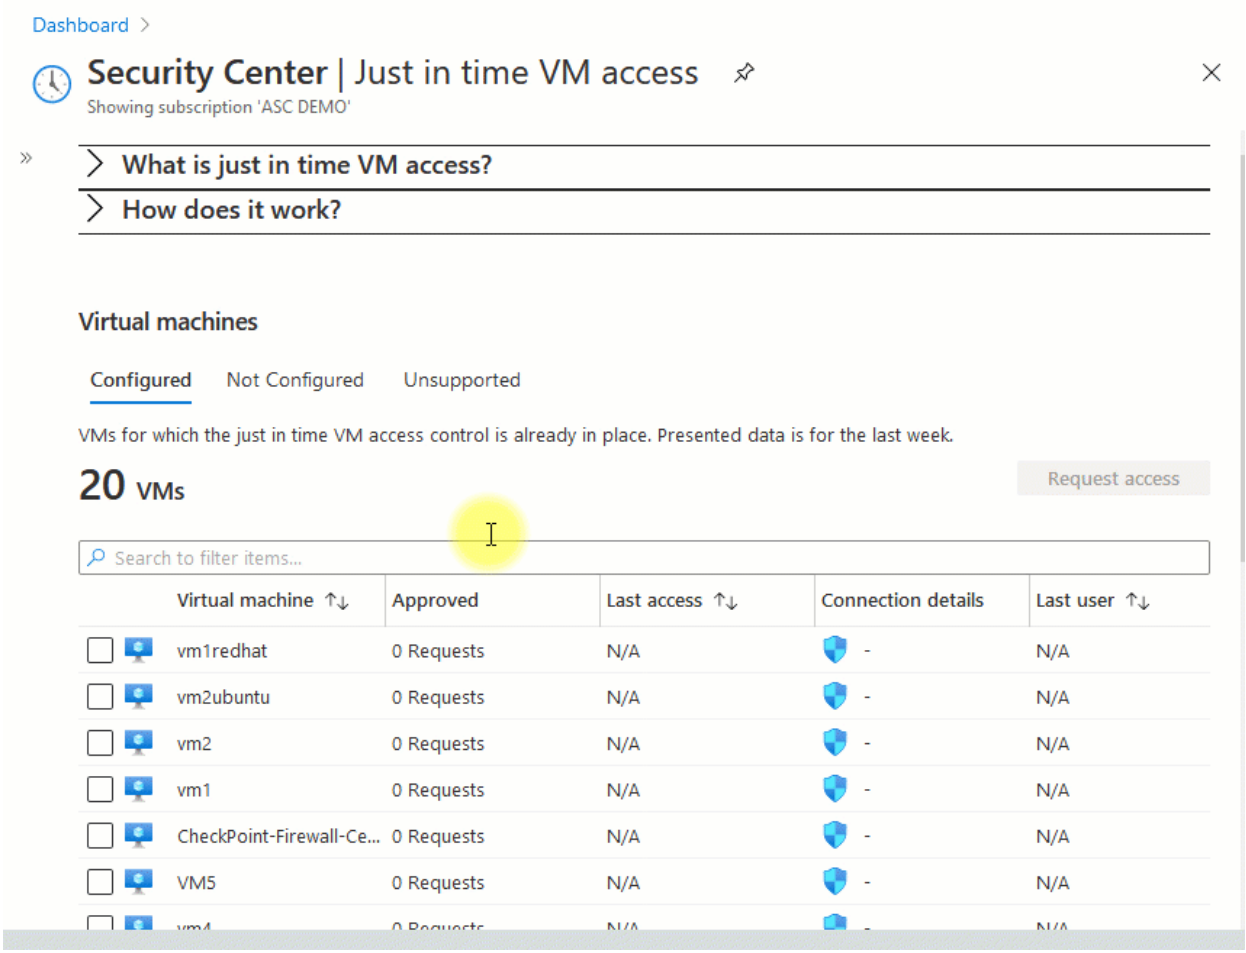

From Defender for Cloud, you can enable and configure the JIT VM access.

1. Open the **Workload protections dashboard** and from the advanced protection area, select **Just-in-time VM access**.

   The **Just-in-time VM access** page opens with your VMs grouped into the following tabs:

   - **Configured** - VMs that have been already been configured to support just-in-time VM access. For each VM, the configured tab shows:
     - the number of approved JIT requests in the last seven days
     - the last access date and time
     - the connection details configured
     - the last user
   - **Not configured** - VMs without JIT enabled, but that can support JIT. We recommend that you enable JIT for these VMs.
   - **Unsupported** - VMs without JIT enabled and which don't support the feature. Your VM might be in this tab for the following reasons:
     - Missing network security group (NSG) or Azure Firewall - JIT requires an NSG to be configured or a Firewall configuration (or both)
     - Classic VM - JIT supports VMs that are deployed through Azure Resource Manager, not 'classic deployment'. Learn more about classic vs Azure Resource Manager deployment models.
     - Other - Your VM might be in this tab if the JIT solution is disabled in the security policy of the subscription or the resource group.
2. From the **Not configured** tab, mark the VMs to protect with JIT and select **Enable JIT on VMs**.

   The JIT VM access page opens listing the ports that Defender for Cloud recommends protecting:

   - 22 - SSH

- 3389 - RDP
- 5985 - WinRM
- 5986 - WinRM

To accept the default settings, select **Save**.

3. To customize the JIT options:

   - Add custom ports with the **Add** button.
   - Modify one of the default ports, by selecting it from the list.

   For each port (custom and default) the **Add port configuration** pane offers the following options:

   - **Protocol**- The protocol that is allowed on this port when a request is approved
   - **Allowed source IPs**- The IP ranges that are allowed on this port when a request is approved
   - **Maximum request time**- The maximum time window during which a specific port can be opened

   a. Set the port security to your needs.

   b. Select **OK**.

4. Select **Save**.

Edit the JIT configuration on a JIT-enabled VM using Defender for Cloud

You can modify a VM's just-in-time configuration by adding and configuring a new port to protect for that VM, or by changing any other setting related to an already protected port.

To edit the existing JIT rules for a VM:

1. Open the **Workload protections dashboard** and from the advanced protection area, select **Just-in-time VM access**.

2. From the **Configured** tab, right-click on the VM to which you want to add a port, and select edit.

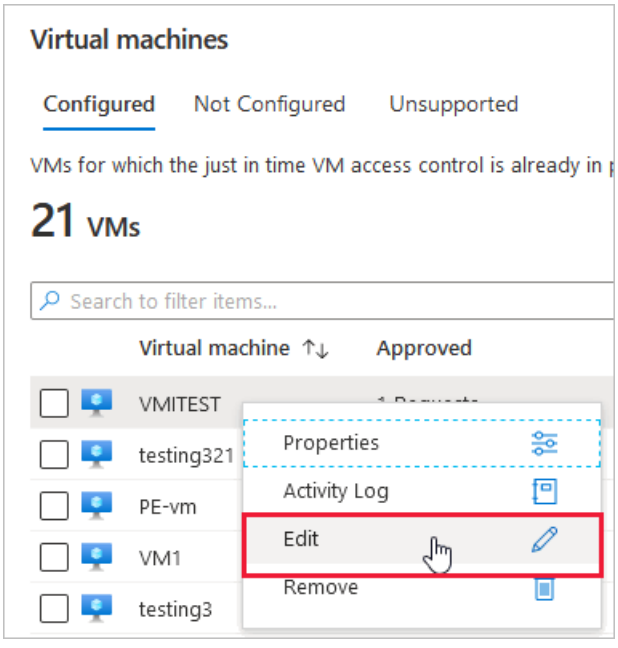

3. Under **JIT VM access configuration**, you can either edit the existing settings of an already protected port or add a new custom port.

4. When you've finished editing the ports, select **Save**.

Request access to a JIT-enabled VM

You can request access to a JIT-enabled VM from the Azure portal (in Defender for Cloud or Azure Virtual

machines) or programmatically.

Each of these options is explained in a separate tab below.

- Microsoft Defender for Cloud
- Azure virtual machines
- PowerShell
- REST API

Request access to a JIT-enabled VM from Microsoft Defender for Cloud

When a VM has a JIT enabled, you have to request access to connect to it. You can request access in any of the supported ways, regardless of how you enabled JIT.

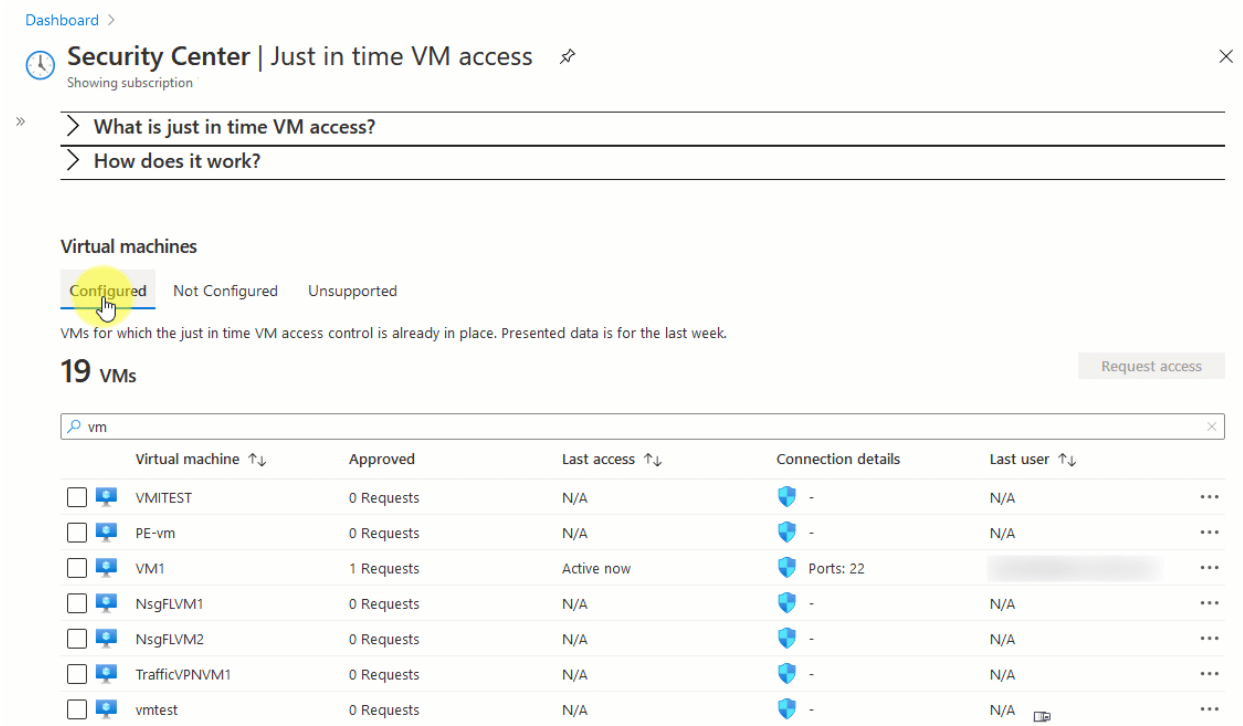

1. From the **Just-in-time VM access** page, select the **Configured** tab.
2. Mark the VMs you want to access.
   - The icon in the **Connection Details** column indicates whether JIT is enabled on the network security group or firewall. If it's enabled on both, only the firewall icon appears.
   - The **Connection Details** column provides the information required to connect the VM, and its open ports.
3. Select **Request access**. The **Request access** window opens.
4. Under **Request access**, for each VM, configure the ports that you want to open and the source IP addresses that the port is opened on and the time window for which the port will be open. It will only be possible to request access to the configured ports. Each port has a maximum allowed time derived from the JIT configuration you've created.
5. Select **Open ports**.

> **NOTE**
>
> If a user who is requesting access is behind a proxy, the option **My IP** may not work. You may need to define the full IP address range of the organization.

Audit JIT access activity in Defender for Cloud

You can gain insights into VM activities using log search. To view the logs:

1. From **Just-in-time VM access**, select the **Configured** tab.
2. For the VM that you want to audit, open the ellipsis menu at the end of the row.
3. Select **Activity Log** from the menu.

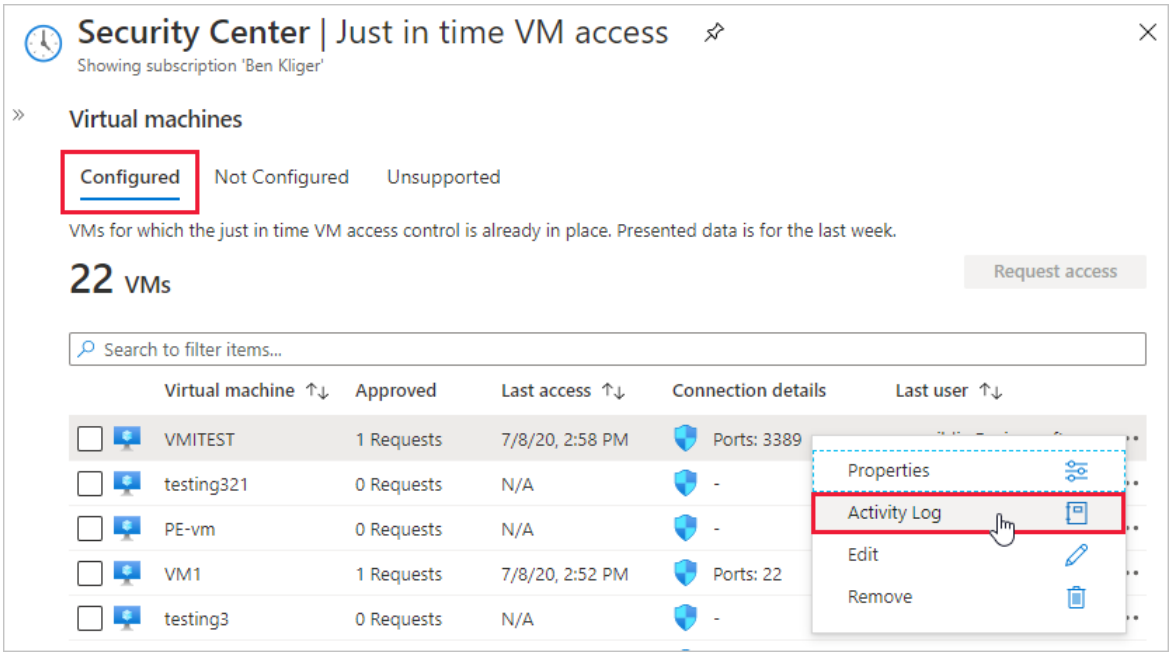

   The activity log provides a filtered view of previous operations for that VM along with time, date, and subscription.

4. To download the log information, select **Download as CSV**.

Next steps

In this article, you learned *how* to configure and use just-in-time VM access. To learn *why* JIT should be used, read the concept article explaining the threats it defends against:

JIT explained

Protect your endpoints with Defender for Cloud's integrated EDR solution: Microsoft Defender for Endpoint

2/15/2022 • 12 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Endpoint is a holistic, cloud-delivered, endpoint security solution. Its main features are:

- Risk-based vulnerability management and assessment
- Attack surface reduction
- Behavioral based and cloud-powered protection
- Endpoint detection and response (EDR)
- Automatic investigation and remediation
- Managed hunting services

> **TIP**
>
> Originally launched as **Windows Defender ATP**, in 2019, this EDR product was renamed **Microsoft Defender ATP**.
>
> At Ignite 2020, we launched the Microsoft Defender for Cloud XDR suite, and this EDR component was renamed **Microsoft Defender for Endpoint**.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers
Supported environments:	✅ Azure Arc-enabled machines running Windows/Linux ✅ Azure VMs running Linux (supported versions) ✅ Azure VMs running Windows Server 2022, 2019, 2016, 2012 R2, 2008 R2 SP1, Azure Virtual Desktop (formerly Windows Virtual Desktop), Windows 10 Enterprise multi-session (formerly Enterprise for Virtual Desktops) ❌ Azure VMs running Windows 11 or Windows 10 (except if running Azure Virtual Desktop or Windows 10 Enterprise multi-session)

ASPECT	DETAILS
Required roles and permissions:	* To enable/disable the integration: **Security admin** or **Owner** * To view Defender for Endpoint alerts in Defender for Cloud: **Security reader**, **Reader**, **Resource Group Contributor**, **Resource Group Owner**, **Security admin**, **Subscription owner**, or **Subscription Contributor**
Clouds:	✅ Commercial clouds ✅ Azure Government ❌ Azure China 21Vianet ✅ Connected AWS accounts

Benefits of integrating Microsoft Defender for Endpoint with Defender for Cloud

Microsoft Defender for Endpoint protects your Windows and Linux machines whether they're hosted in Azure, hybrid clouds (on-premises), or AWS. Protections include:

- **Advanced post-breach detection sensors**. Defender for Endpoint's sensors collect a vast array of behavioral signals from your machines.
- **Vulnerability assessment from the Microsoft threat and vulnerability management solution**. With Microsoft Defender for Endpoint enabled, Defender for Cloud can show vulnerabilities discovered by the threat and vulnerability management module and also offer this module as a supported vulnerability assessment solution. Learn more in Investigate weaknesses with Microsoft Defender for Endpoint's threat and vulnerability management.

  This module also brings the software inventory features described in Access a software inventory and can be automatically enabled for supported machines with the auto deploy settings.
- **Analytics-based, cloud-powered, post-breach detection**. Defender for Endpoint quickly adapts to changing threats. It uses advanced analytics and big data. It's amplified by the power of the Intelligent Security Graph with signals across Windows, Azure, and Office to detect unknown threats. It provides actionable alerts and enables you to respond quickly.
- **Threat intelligence**. Defender for Endpoint generates alerts when it identifies attacker tools, techniques, and procedures. It uses data generated by Microsoft threat hunters and security teams, augmented by intelligence provided by partners.

By integrating Defender for Endpoint with Defender for Cloud, you'll benefit from the following extra capabilities:

- **Automated onboarding**. Defender for Cloud automatically enables the Defender for Endpoint sensor on all supported machines connected to Defender for Cloud.
- **Single pane of glass**. The Defender for Cloud portal pages display Defender for Endpoint alerts. To investigate further, use Microsoft Defender for Endpoint's own portal pages where you'll see additional information such as the alert process tree and the incident graph. You can also see a detailed machine timeline that shows every behavior for a historical period of up to six months.

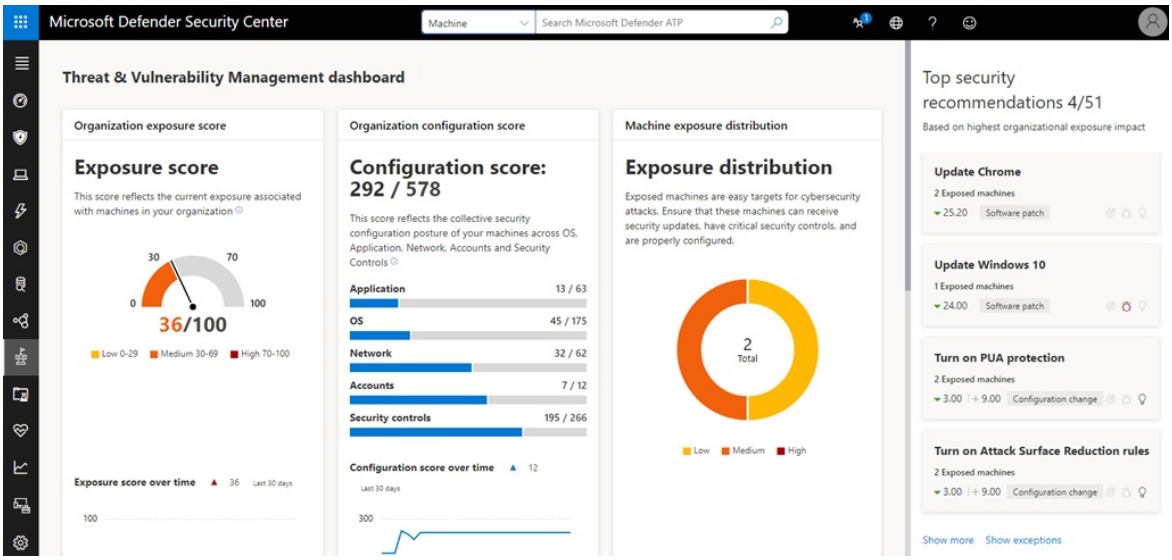

What are the requirements for the Microsoft Defender for Endpoint tenant?

When you use Defender for Cloud to monitor your machines, a Defender for Endpoint tenant is automatically created.

- **Location:** Data collected by Defender for Endpoint is stored in the geo-location of the tenant as identified during provisioning. Customer data - in pseudonymized form - may also be stored in the central storage and processing systems in the United States. After you've configured the location, you can't change it. If you have your own license for Microsoft Defender for Endpoint and need to move your data to another location, contact Microsoft support to reset the tenant.
- **Moving subscriptions:** If you've moved your Azure subscription between Azure tenants, some manual preparatory steps are required before Defender for Cloud will deploy Defender for Endpoint. For full details, contact Microsoft support.

Enable the Microsoft Defender for Endpoint integration

Prerequisites

Confirm that your machine meets the necessary requirements for Defender for Endpoint:

1. Ensure the machine is connected to Azure and the internet as required:
   - **Azure virtual machines (Windows or Linux)** - Configure the network settings described in configure device proxy and internet connectivity settings: Windows or Linux.
   - **On-premises machines** - Connect your target machines to Azure Arc as explained in Connect hybrid machines with Azure Arc-enabled servers.
2. Enable **Microsoft Defender for servers**. See Quickstart: Enable Defender for Cloud's enhanced security features.

   > **IMPORTANT**
   > Defender for Cloud's integration with Microsoft Defender for Endpoint is enabled by default. So when you enable enhanced security features, you give consent for Microsoft Defender for servers to access the Microsoft Defender for Endpoint data related to vulnerabilities, installed software, and alerts for your endpoints.

3. If you've moved your subscription between Azure tenants, some manual preparatory steps are also required. For full details, contact Microsoft support.

Enable the integration

- Windows
- Linux

1. From Defender for Cloud's menu, select **Environment settings** and select the subscription with the Windows machines that you want to receive Defender for Endpoint.
2. Select **Integrations**.
3. Select **Allow Microsoft Defender for Endpoint to access my data**, and select **Save**.

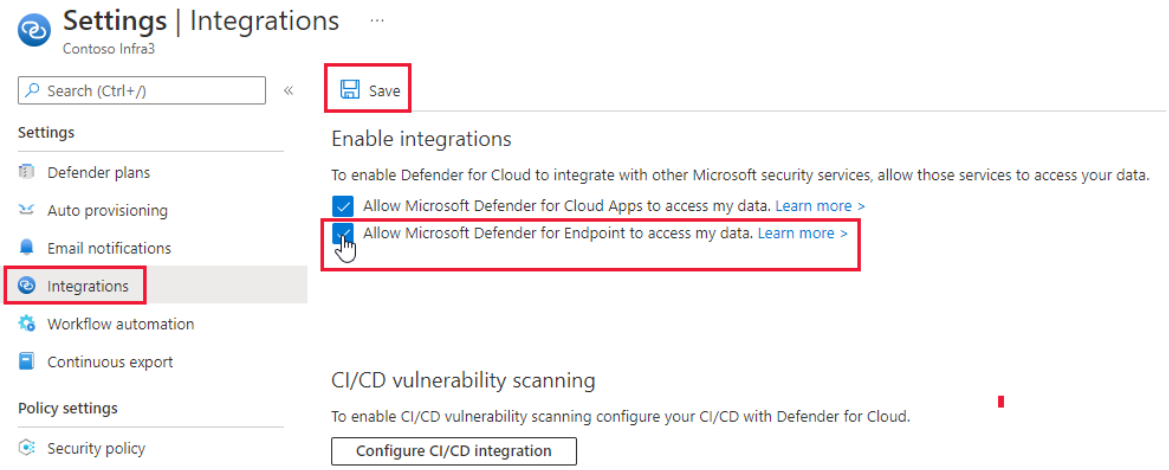

Microsoft Defender for Cloud will automatically onboard your machines to Microsoft Defender for Endpoint. Onboarding might take up to 12 hours. For new machines created after the integration has been enabled, onboarding takes up to an hour.

Access the Microsoft Defender for Endpoint portal

1. Ensure the user account has the necessary permissions. Learn more in Assign user access to Microsoft Defender Security Center.
2. Check whether you have a proxy or firewall that is blocking anonymous traffic. The Defender for Endpoint sensor connects from the system context, so anonymous traffic must be permitted. To ensure unhindered access to the Defender for Endpoint portal, follow the instructions in Enable access to service URLs in the proxy server.
3. Open the Defender for Endpoint Security Center portal. Learn more about the portal's features and icons, in Defender for Endpoint Security Center portal overview.

Send a test alert

To generate a benign test alert from Defender for Endpoint, select the tab for the relevant operating system of your endpoint:

- Windows
- Linux

For endpoints running Windows:

1. Create a folder 'C:\test-MDATP-test'.
2. Use Remote Desktop to access your machine.

3. Open a command-line window.
4. At the prompt, copy and run the following command. The command prompt window will close automatically.

```
powershell.exe -NoExit -ExecutionPolicy Bypass -WindowStyle Hidden (New-Object
System.Net.WebClient).DownloadFile('http://127.0.0.1/1.exe', 'C:\\test-MDATP-test\\invoice.exe');
Start-Process 'C:\\test-MDATP-test\\invoice.exe'
```

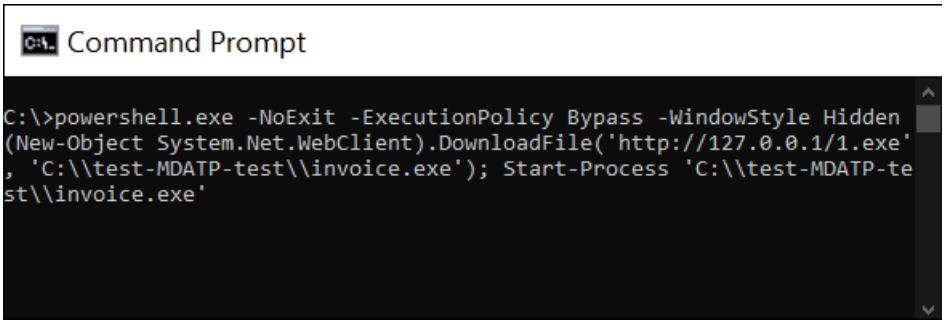

If the command is successful, you'll see a new alert on the workload protection dashboard and the Microsoft Defender for Endpoint portal. This alert might take a few minutes to appear.

5. To review the alert in Defender for Cloud, go to **Security alerts** > **Suspicious PowerShell CommandLine**.
6. From the investigation window, select the link to go to the Microsoft Defender for Endpoint portal.

> **TIP**
> The alert is triggered with **Informational** severity.

Remove Defender for Endpoint from a machine

To remove the Defender for Endpoint solution from your machines:

1. Disable the integration:
   a. From Defender for Cloud's menu, select **Environment settings** and select the subscription with the relevant machines.
   b. Open **Integrations** and clear the checkbox for **Allow Microsoft Defender for Endpoint to access my data**.
   c. Select **Save**.
2. Remove the MDE.Windows/MDE.Linux extension from the machine.
3. Follow the steps in Offboard devices from the Microsoft Defender for Endpoint service from the Defender for Endpoint documentation.

FAQ - Microsoft Defender for Cloud integration with Microsoft Defender for Endpoint

- What's this "MDE.Windows" / "MDE.Linux" extension running on my machine?
- What are the licensing requirements for Microsoft Defender for Endpoint?
- If I already have a license for Microsoft Defender for Endpoint, can I get a discount for Microsoft Defender for servers?
- How do I switch from a third-party EDR tool?

What's this "MDE.Windows" / "MDE.Linux" extension running on my machine?

In the past, Microsoft Defender for Endpoint was provisioned by the Log Analytics agent. When we expanded support to include Windows Server 2019 and Linux, we also added an extension to perform the automatic onboarding.

Defender for Cloud automatically deploys the extension to machines running:

- Windows Server 2019 & 2022.
- Windows 10 Virtual Desktop (WVD).
- Other versions of Windows Server if Defender for Cloud doesn't recognize the OS version (for example, when a custom VM image is used). In this case, Microsoft Defender for Endpoint is still provisioned by the Log Analytics agent.
- Linux.

> **IMPORTANT**
> If you delete the MDE.Windows/MDE.Linux extension, it will not remove Microsoft Defender for Endpoint. to 'offboard', see Offboard Windows servers..

I've enabled the solution by the "MDE.Windows" / "MDE.Linux" extension isn't showing on my machine

If you've enabled the integration, but still don't see the extension running on your machines, check the following:

1. If 12 hours hasn't passed since you enabled the solution, you'll need to wait until the end of this period to be sure there's an issue to investigate.
2. After 12 hours have passed, if you still don't see the extension running on your machines, check that you've met Prerequisites for the integration.
3. Ensure you've enabled the Microsoft Defender for servers plan for the subscriptions related to the machines you're investigating.
4. If you've moved your Azure subscription between Azure tenants, some manual preparatory steps are required before Defender for Cloud will deploy Defender for Endpoint. For full details, contact Microsoft support.

What are the licensing requirements for Microsoft Defender for Endpoint?

Defender for Endpoint is included at no extra cost with **Microsoft Defender for servers**. Alternatively, it can be purchased separately for 50 machines or more.

If I already have a license for Microsoft Defender for Endpoint, can I get a discount for Microsoft Defender for servers?

If you've already got a license for **Microsoft Defender for Endpoint for Servers** , you won't have to pay for that part of your Microsoft Defender for servers license. Learn more about this license.

To request your discount, contact Defender for Cloud's support team. You'll need to provide the relevant workspace ID, region, and number of Microsoft Defender for Endpoint for servers licenses applied for machines in the given workspace.

The discount will be effective starting from the approval date, and won't take place retroactively.

How do I switch from a third-party EDR tool?

Full instructions for switching from a non-Microsoft endpoint solution are available in the Microsoft Defender for Endpoint documentation: Migration overview.

Next steps

- Platforms and features supported by Microsoft Defender for Cloud

- Learn how recommendations help you protect your Azure resources

Use adaptive application controls to reduce your machines' attack surfaces

2/15/2022 • 9 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Learn about the benefits of Microsoft Defender for Cloud's adaptive application controls and how you can enhance your security with this data-driven, intelligent feature.

What are adaptive application controls?

Adaptive application controls are an intelligent and automated solution for defining allow lists of known-safe applications for your machines.

Often, organizations have collections of machines that routinely run the same processes. Microsoft Defender for Cloud uses machine learning to analyze the applications running on your machines and create a list of the known-safe software. Allow lists are based on your specific Azure workloads, and you can further customize the recommendations using the instructions below.

When you've enabled and configured adaptive application controls, you'll get security alerts if any application runs other than the ones you've defined as safe.

What are the benefits of adaptive application controls?

By defining lists of known-safe applications, and generating alerts when anything else is executed, you can achieve multiple oversight and compliance goals:

- Identify potential malware, even any that might be missed by antimalware solutions
- Improve compliance with local security policies that dictate the use of only licensed software
- Identify outdated or unsupported versions of applications
- Identify software that's banned by your organization but is nevertheless running on your machines
- Increase oversight of apps that access sensitive data

No enforcement options are currently available. Adaptive application controls are intended to provide security alerts if any application runs other than the ones you've defined as safe.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers

ASPECT	DETAILS
Supported machines:	✅ Azure and non-Azure machines running Windows and Linux ✅ Azure Arc machines
Required roles and permissions:	**Security Reader** and **Reader** roles can both view groups and the lists of known-safe applications **Contributor** and **Security Admin** roles can both edit groups and the lists of known-safe applications
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) ✅ Connected AWS accounts

Enable application controls on a group of machines

If Microsoft Defender for Cloud has identified groups of machines in your subscriptions that consistently run a similar set of applications, you'll be prompted with the following recommendation: **Adaptive application controls for defining safe applications should be enabled on your machines**.

Select the recommendation, or open the adaptive application controls page to view the list of suggested known-safe applications and groups of machines.

1. Open the Workload protections dashboard and from the advanced protection area, select **Adaptive application controls**.

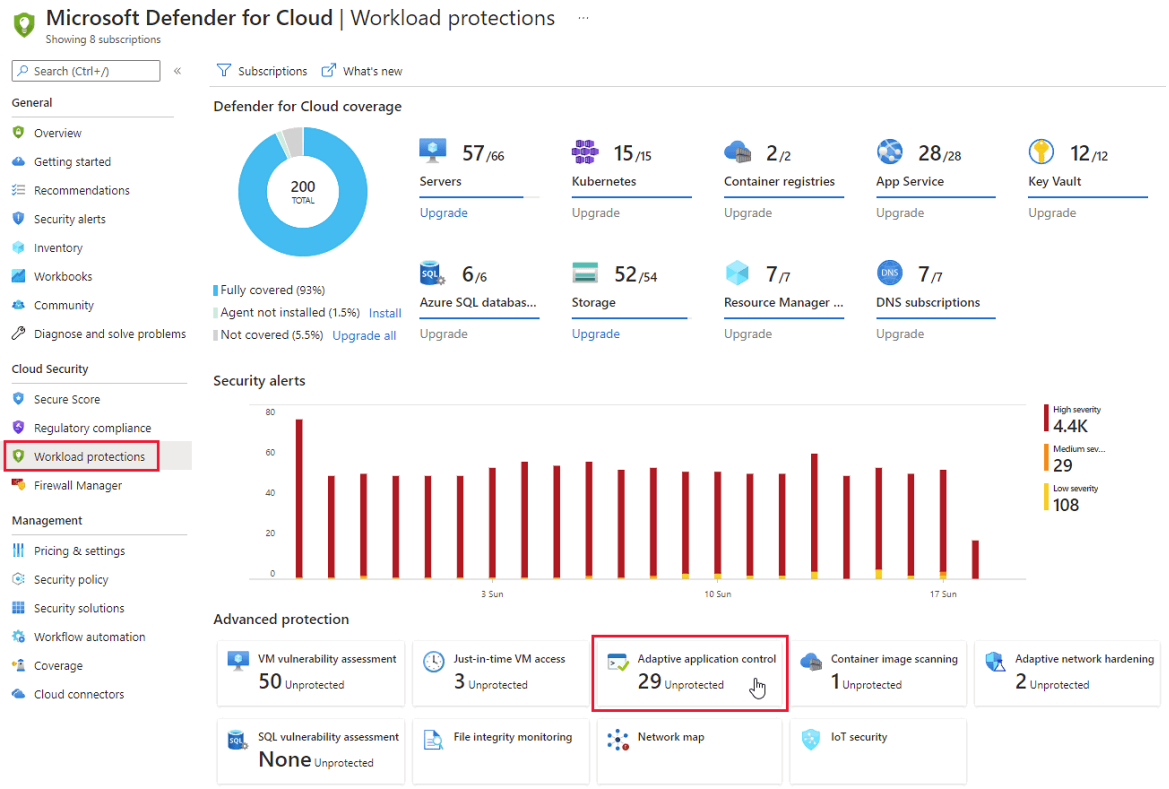

The **Adaptive application controls** page opens with your VMs grouped into the following tabs:

- **Configured** - Groups of machines that already have a defined allow list of applications. For each group, the configured tab shows:
    - the number of machines in the group

- recent alerts

- **Recommended** - Groups of machines that consistently run the same applications, and don't have an allow list configured. We recommend that you enable adaptive application controls for these groups.

  > **TIP**
  > If you see a group name with the prefix "REVIEWGROUP", it contains machines with a partially consistent list of applications. Microsoft Defender for Cloud can't see a pattern but recommends reviewing this group to see whether *you* can manually define some adaptive application controls rules as described in Editing a group's adaptive application controls rule.
  >
  > You can also move machines from this group to other groups as described in Move a machine from one group to another.

- **No recommendation** - Machines without a defined allow list of applications, and which don't support the feature. Your machine might be in this tab for the following reasons:

  - It's missing a Log Analytics agent
  - The Log Analytics agent isn't sending events
  - It's a Windows machine with a pre-existing AppLocker policy enabled by either a GPO or a local security policy

  > **TIP**
  > Defender for Cloud needs at least two weeks of data to define the unique recommendations per group of machines. Machines that have recently been created, or which belong to subscriptions that were only recently protected by Microsoft Defender for servers, will appear under the **No recommendation** tab.

2. Open the **Recommended** tab. The groups of machines with recommended allow lists appears.

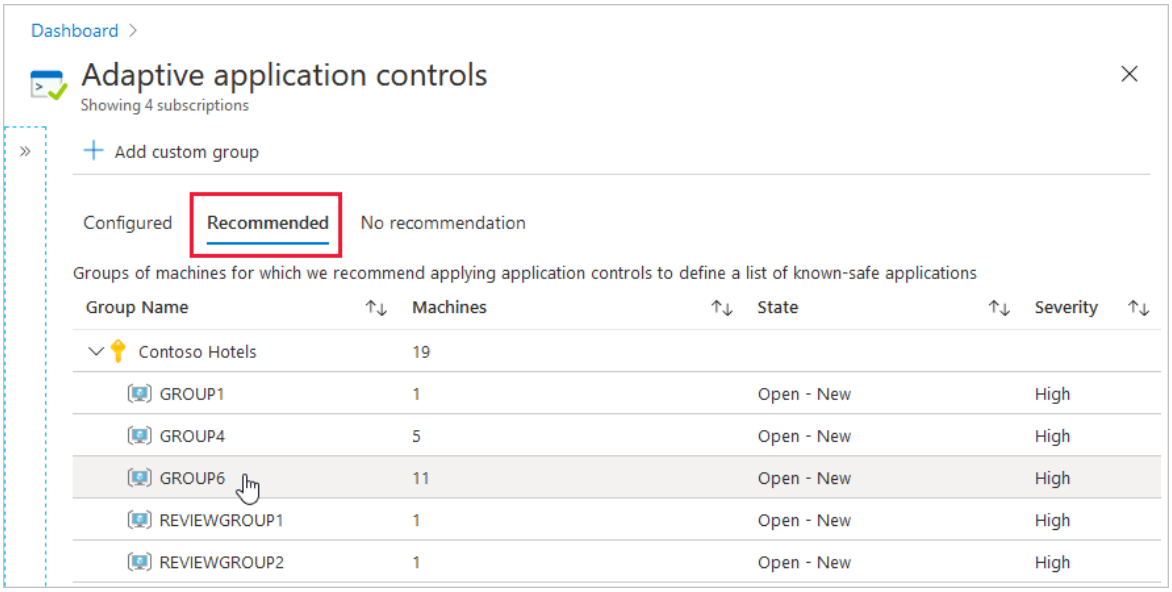

3. Select a group.

4. To configure your new rule, review the various sections of this **Configure application control rules** page and the contents, which will be unique to this specific group of machines:

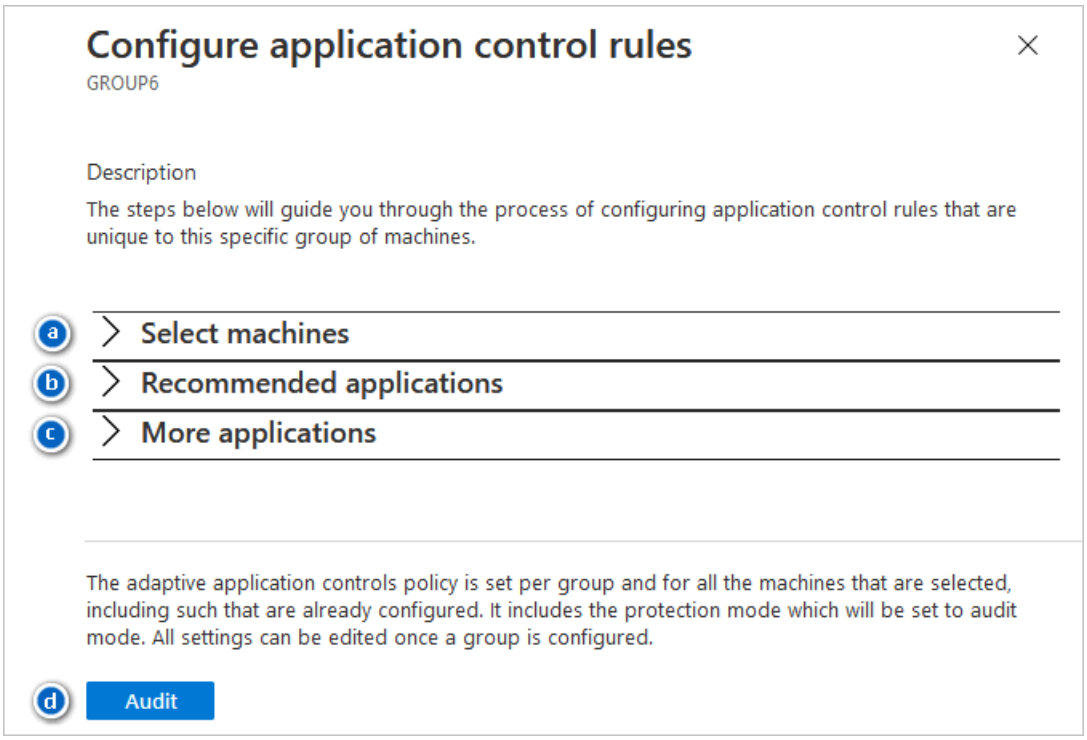

a. **Select machines** - By default, all machines in the identified group are selected. Unselect any to removed them from this rule.

b. **Recommended applications** - Review this list of applications that are common to the machines within this group, and recommended to be allowed to run.

c. **More applications** - Review this list of applications that are either seen less frequently on the machines within this group, or are known to be exploitable. A warning icon indicates that a specific application could be used by an attacker to bypass an application allow list. We recommend that you carefully review these applications.

> **TIP**
>
> Both application lists include the option to restrict a specific application to certain users. Adopt the principle of least privilege whenever possible.
>
> Applications are defined by their publishers, if an application doesn't have publisher information (it's unsigned), a path rule is created for the full path of the specific application.

d. To apply the rule, select **Audit**.

Edit a group's adaptive application controls rule

You might decide to edit the allow list for a group of machines because of known changes in your organization.

To edit the rules for a group of machines:

1. Open the **Workload protections dashboard** and from the advanced protection area, select **Adaptive application controls**.
2. From the **Configured** tab, select the group with the rule you want to edit.
3. Review the various sections of the **Configure application control rules** page as described in Enable adaptive application controls on a group of machines.
4. Optionally, add one or more custom rules:

a. Select **Add rule**.

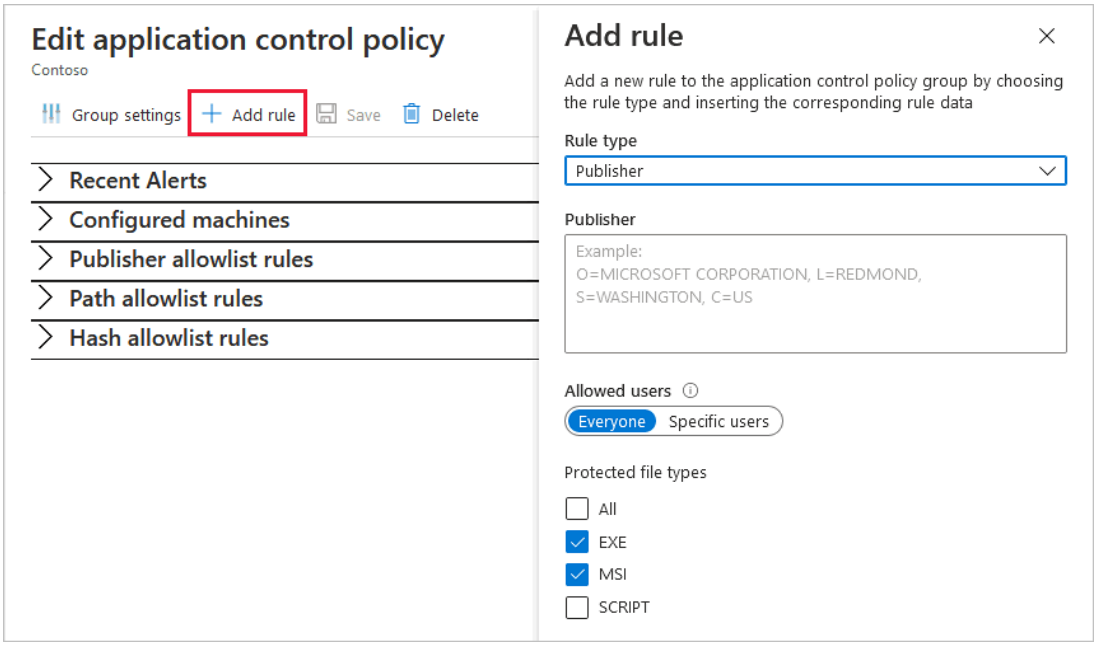

b. If you're defining a known safe path, change the **Rule type** to 'Path' and enter a single path. You can include wildcards in the path.

> **TIP**
>
> Some scenarios for which wildcards in a path might be useful:
>
> - Using a wildcard at the end of a path to allow all executables within this folder and sub-folders.
> - Using a wildcard in the middle of a path to enable a known executable name with a changing folder name (for example, personal user folders containing a known executable, automatically generated folder names, etc).

c. Define the allowed users and protected file types.

d. When you've finished defining the rule, select **Add**.

5. To apply the changes, select **Save**.

Review and edit a group's settings

1. To view the details and settings of your group, select **Group settings**

   This pane shows the name of the group (which can be modified), the OS type, the location, and other relevant details.

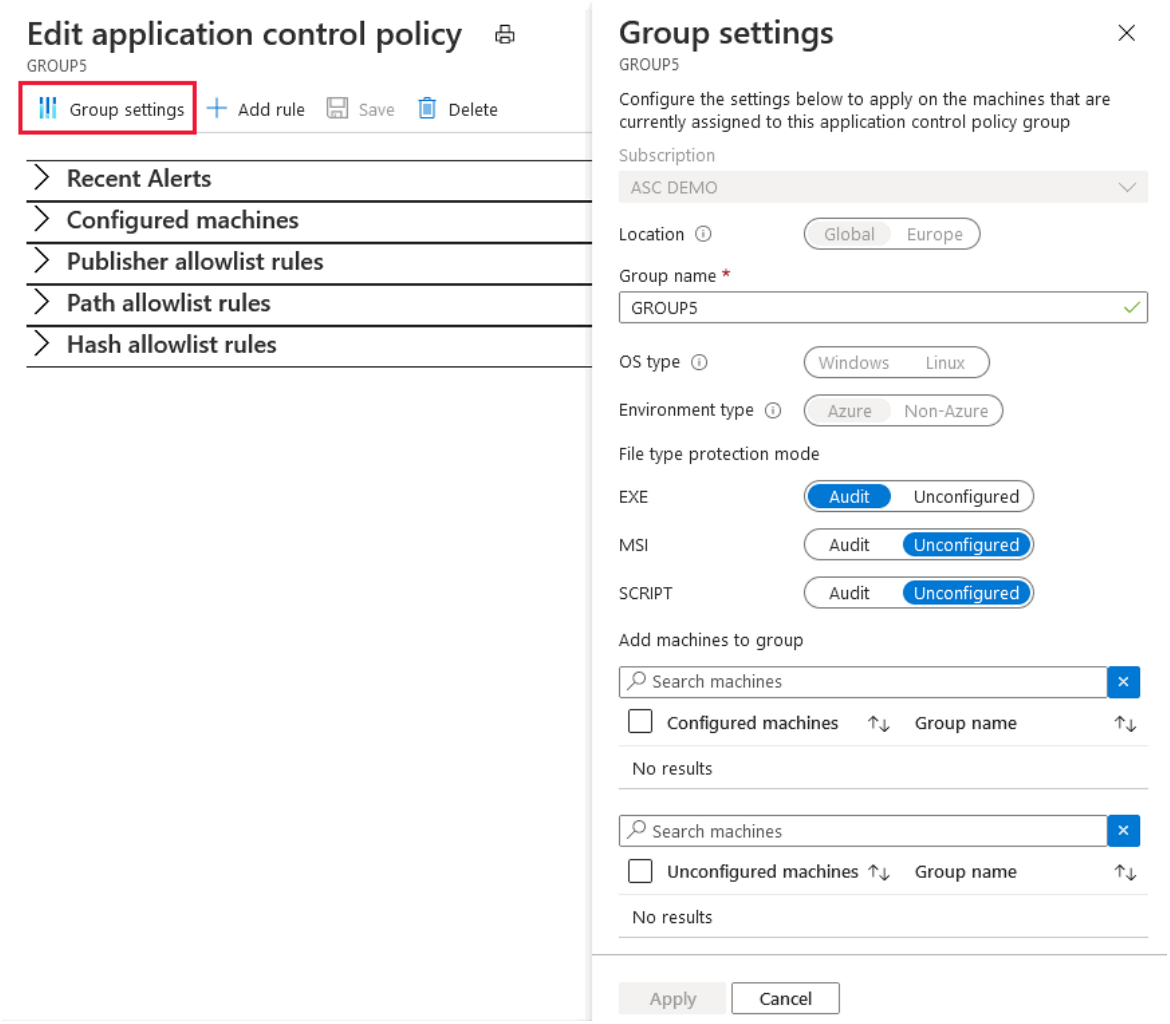

2. Optionally, modify the group's name or file type protection modes.
3. Select **Apply** and **Save**.

Respond to the "Allowlist rules in your adaptive application control policy should be updated" recommendation

You'll see this recommendation when Defender for Cloud's machine learning identifies potentially legitimate behavior that hasn't previously been allowed. The recommendation suggests new rules for your existing definitions to reduce the number of false positive alerts.

To remediate the issues:

1. From the recommendations page, select the **Allowlist rules in your adaptive application control policy should be updated** recommendation to see groups with newly identified, potentially legitimate behavior.
2. Select the group with the rule you want to edit.
3. Review the various sections of the **Configure application control rules** page as described in Enable adaptive application controls on a group of machines.
4. To apply the changes, select **Audit**.

Audit alerts and violations

1. Open the **Workload protections dashboard** and from the advanced protection area, select **Adaptive application controls**.

2. To see groups with machines that have recent alerts, review the groups listed in the **Configured** tab.
3. To investigate further, select a group.

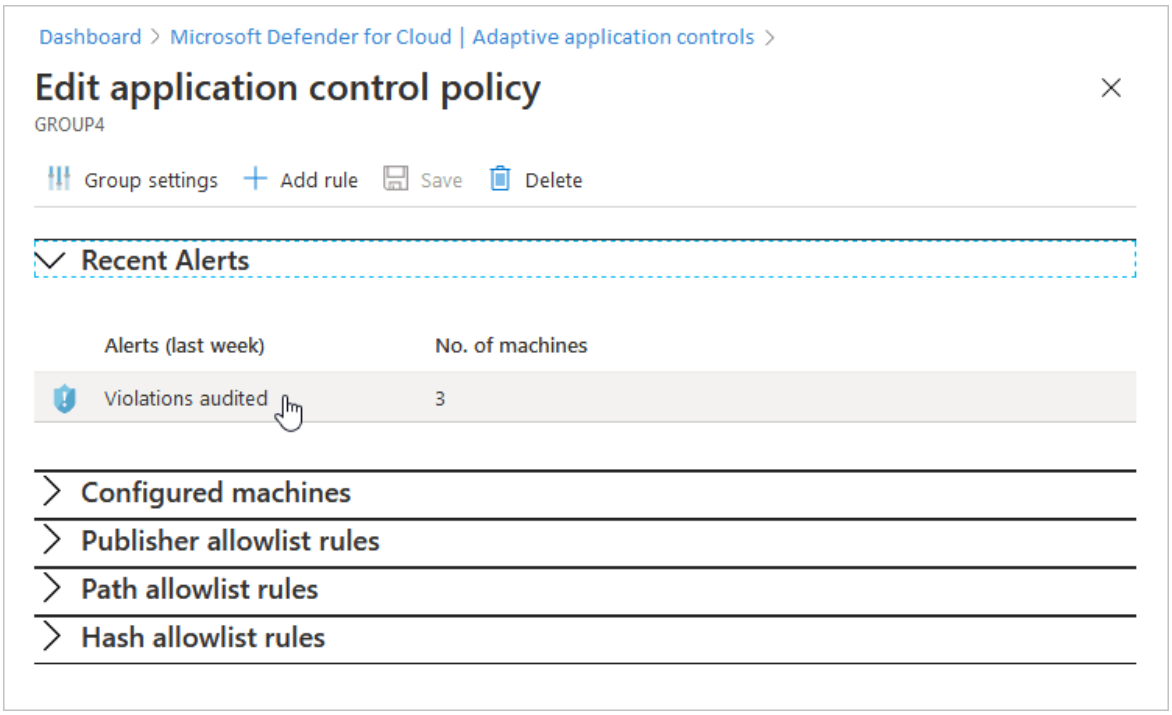

4. For further details, and the list of affected machines, select an alert.

   The alerts page shows the more details of the alerts and provides a **Take action** link with recommendations of how to mitigate the threat.

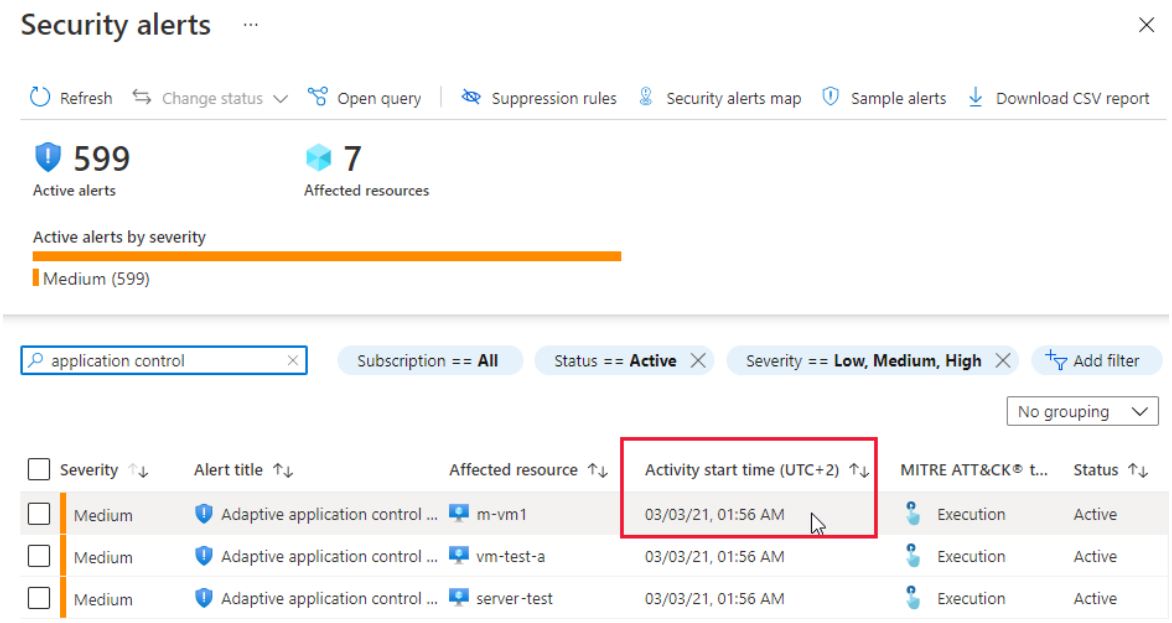

> **NOTE**
>
> Adaptive application controls calculates events once every twelve hours. The "activity start time" shown in the alerts page is the time that adaptive application controls created the alert, **not** the time that the suspicious process was active.

Move a machine from one group to another

When you move a machine from one group to another, the application control policy applied to it changes to the

settings of the group that you moved it to. You can also move a machine from a configured group to a non-configured group, doing so removes any application control rules that were applied to the machine.

1. Open the **Workload protections dashboard** and from the advanced protection area, select **Adaptive application controls**.
2. From the **Adaptive application controls** page, from the **Configured** tab, select the group containing the machine to be moved.
3. Open the list of **Configured machines**.
4. Open the machine's menu from three dots at the end of the row, and select **Move**. The **Move machine to a different group** pane opens.
5. Select the destination group, and select **Move machine**.
6. To save your changes, select **Save**.

Manage application controls via the REST API

To manage your adaptive application controls programmatically, use our REST API.

The relevant API documentation is available in the Adaptive Application Controls section of Defender for Cloud's API docs.

Some of the functions that are available from the REST API:

- **List** retrieves all your group recommendations and provides a JSON with an object for each group.
- **Get** retrieves the JSON with the full recommendation data (that is, list of machines, publisher/path rules, and so on).
- **Put** configures your rule (use the JSON you retrieved with **Get** as the body for this request).

> **IMPORTANT**
>
> The **Put** function expects fewer parameters than the JSON returned by the Get command contains.
>
> Remove the following properties before using the JSON in the Put request: recommendationStatus, configurationStatus, issues, location, and sourceSystem.

FAQ - Adaptive application controls

- Are there any options to enforce the application controls?
- Why do I see a Qualys app in my recommended applications?

Are there any options to enforce the application controls?

No enforcement options are currently available. Adaptive application controls are intended to provide **security alerts** if any application runs other than the ones you've defined as safe. They have a range of benefits (What are the benefits of adaptive application controls?) and are extremely customizable as shown on this page.

Why do I see a Qualys app in my recommended applications?

Microsoft Defender for servers includes vulnerability scanning for your machines at no extra cost. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Defender for Cloud. For details of this scanner and instructions for how to deploy it, see Defender for Cloud's integrated Qualys vulnerability assessment solution.

To ensure no alerts are generated when Defender for Cloud deploys the scanner, the adaptive application

controls recommended allow list includes the scanner for all machines.

Next steps

On this page, you learned how to use adaptive application control in Microsoft Defender for Cloud to define allow lists of applications running on your Azure and non-Azure machines. To learn more about some other cloud workload protection features, see:

- Understanding just-in-time (JIT) VM access
- Securing your Azure Kubernetes clusters

Use asset inventory to manage your resources' security posture

2/15/2022 • 7 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The asset inventory page of Microsoft Defender for Cloud provides a single page for viewing the security posture of the resources you've connected to Microsoft Defender for Cloud.

Defender for Cloud periodically analyzes the security state of resources connected to your subscriptions to identify potential security vulnerabilities. It then provides you with recommendations on how to remediate those vulnerabilities.

When any resource has outstanding recommendations, they'll appear in the inventory.

Use this view and its filters to address such questions as:

- Which of my subscriptions with enhanced security features enabled have outstanding recommendations?
- Which of my machines with the tag 'Production' are missing the Log Analytics agent?
- How many of my machines tagged with a specific tag have outstanding recommendations?
- Which machines in a specific resource group have a known vulnerability (using a CVE number)?

The asset management possibilities for this tool are substantial and continue to grow.

> **TIP**
> The security recommendations on the asset inventory page are the same as those on the **Recommendations** page, but here they're shown according to the affected resource. For information about how to resolve recommendations, see Implementing security recommendations in Microsoft Defender for Cloud.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free* * Some features of the inventory page, such as the software inventory require paid solutions to be in-place
Required roles and permissions:	All users
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet)

ASPECT	DETAILS

What are the key features of asset inventory?

The inventory page provides the following tools:

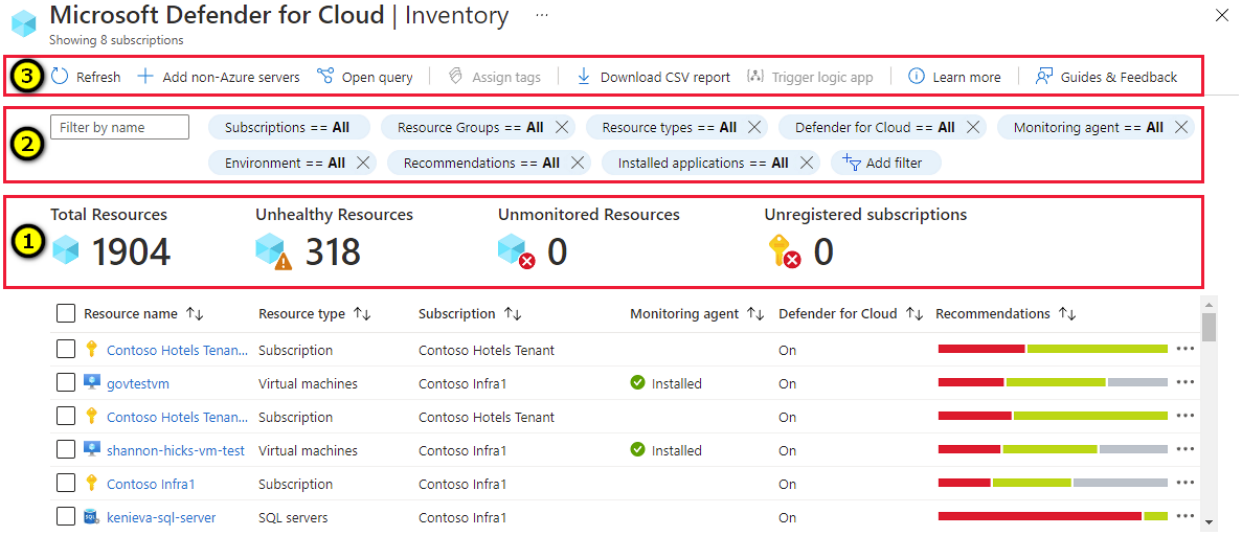

1 - Summaries

Before you define any filters, a prominent strip of values at the top of the inventory view shows:

- **Total resources**: The total number of resources connected to Defender for Cloud.
- **Unhealthy resources**: Resources with active security recommendations. Learn more about security recommendations.
- **Unmonitored resources**: Resources with agent monitoring issues - they have the Log Analytics agent deployed, but the agent isn't sending data or has other health issues.
- **Unregistered subscriptions**: Any subscription in the selected scope that haven't yet been connected to Microsoft Defender for Cloud.

2 - Filters

The multiple filters at the top of the page provide a way to quickly refine the list of resources according to the question you're trying to answer. For example, if you wanted to answer the question *Which of my machines with the tag 'Production' are missing the Log Analytics agent?* you could combine the **Agent monitoring** filter with the **Tags** filter.

As soon as you've applied filters, the summary values are updated to relate to the query results.

3 - Export and asset management tools

Export options - Inventory includes an option to export the results of your selected filter options to a CSV file. You can also export the query itself to Azure Resource Graph Explorer to further refine, save, or modify the Kusto Query Language (KQL) query.

> **TIP**
> The KQL documentation provides a database with some sample data together with some simple queries to get the "feel" for the language. Learn more in this KQL tutorial.

Asset management options - When you've found the resources that match your queries, inventory provides shortcuts for operations such as:

- Assign tags to the filtered resources - select the checkboxes alongside the resources you want to tag.
- Onboard new servers to Defender for Cloud - use the **Add non-Azure servers** toolbar button.
- Automate workloads with Azure Logic Apps - use the **Trigger Logic App** button to run a logic app on one or more resources. Your logic apps have to be prepared in advance, and accept the relevant trigger type (HTTP request). Learn more about logic apps.

How does asset inventory work?

Asset inventory utilizes Azure Resource Graph (ARG), an Azure service that provides the ability to query Defender for Cloud's security posture data across multiple subscriptions.

ARG is designed to provide efficient resource exploration with the ability to query at scale.

Using the Kusto Query Language (KQL), asset inventory can quickly produce deep insights by cross-referencing Defender for Cloud data with other resource properties.

How to use asset inventory

1. From Defender for Cloud's sidebar, select **Inventory**.
2. Use the **Filter by name** box to display a specific resource, or use the filters as described below.
3. Select the relevant options in the filters to create the specific query you want to perform.

   By default, the resources are sorted by the number of active security recommendations.

   > **IMPORTANT**
   >
   > The options in each filter are specific to the resources in the currently selected subscriptions **and** your selections in the other filters.
   >
   > For example, if you've selected only one subscription, and the subscription has no resources with outstanding security recommendations to remediate (0 unhealthy resources), the **Recommendations** filter will have no options.

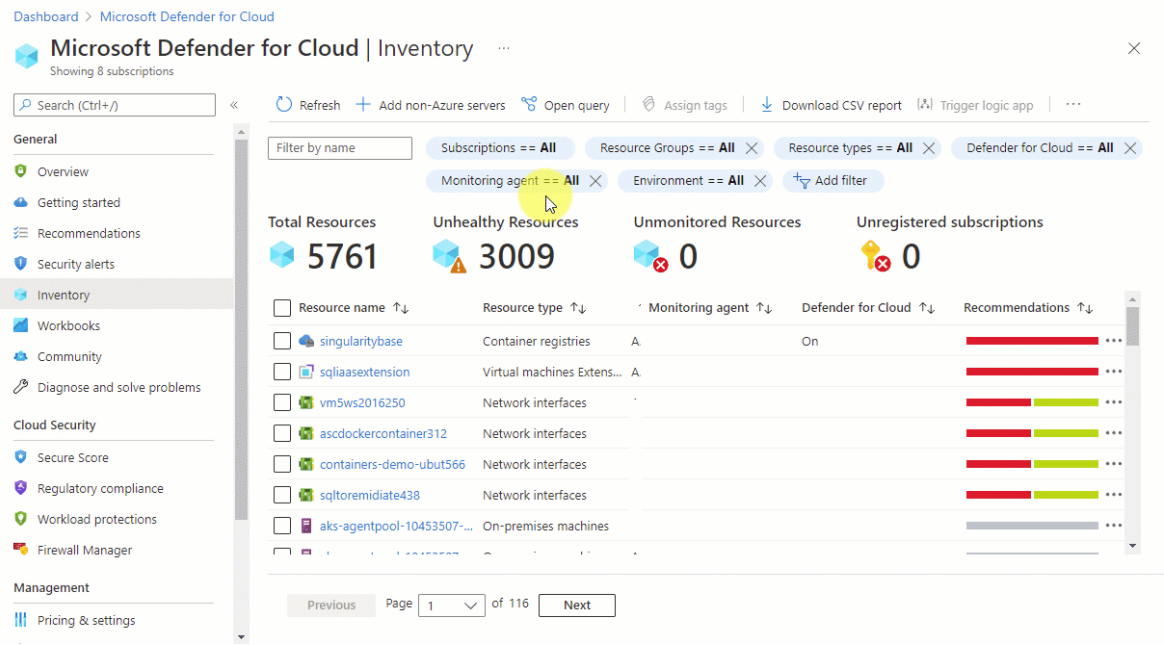

4. To use the **Security findings contain** filter, enter free text from the ID, security check, or CVE name of a vulnerability finding to filter to the affected resources:

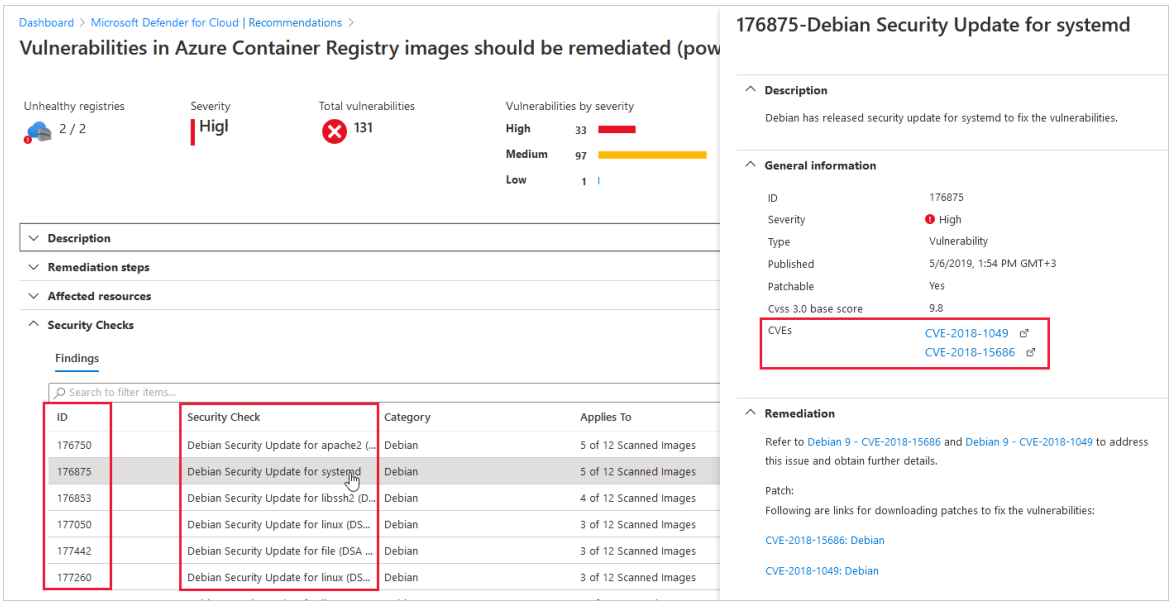

> **TIP**
>
> The **Security findings contain** and **Tags** filters only accept a single value. To filter by more than one, use **Add filters**.

5. To use the **Defender for Cloud** filter, select one or more options (Off, On, or Partial):

   - **Off** - Resources that aren't protected by a Microsoft Defender plan. You can right-click on any of these and upgrade them:

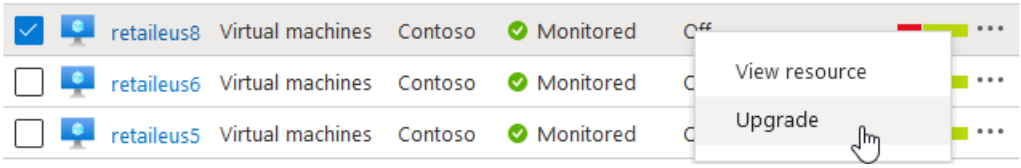

   - **On** - Resources that are protected by a Microsoft Defender plan
   - **Partial** - This applies to **subscriptions** that have some but not all of the Microsoft Defender plans disabled. For example, the following subscription has seven Microsoft Defender plans disabled.

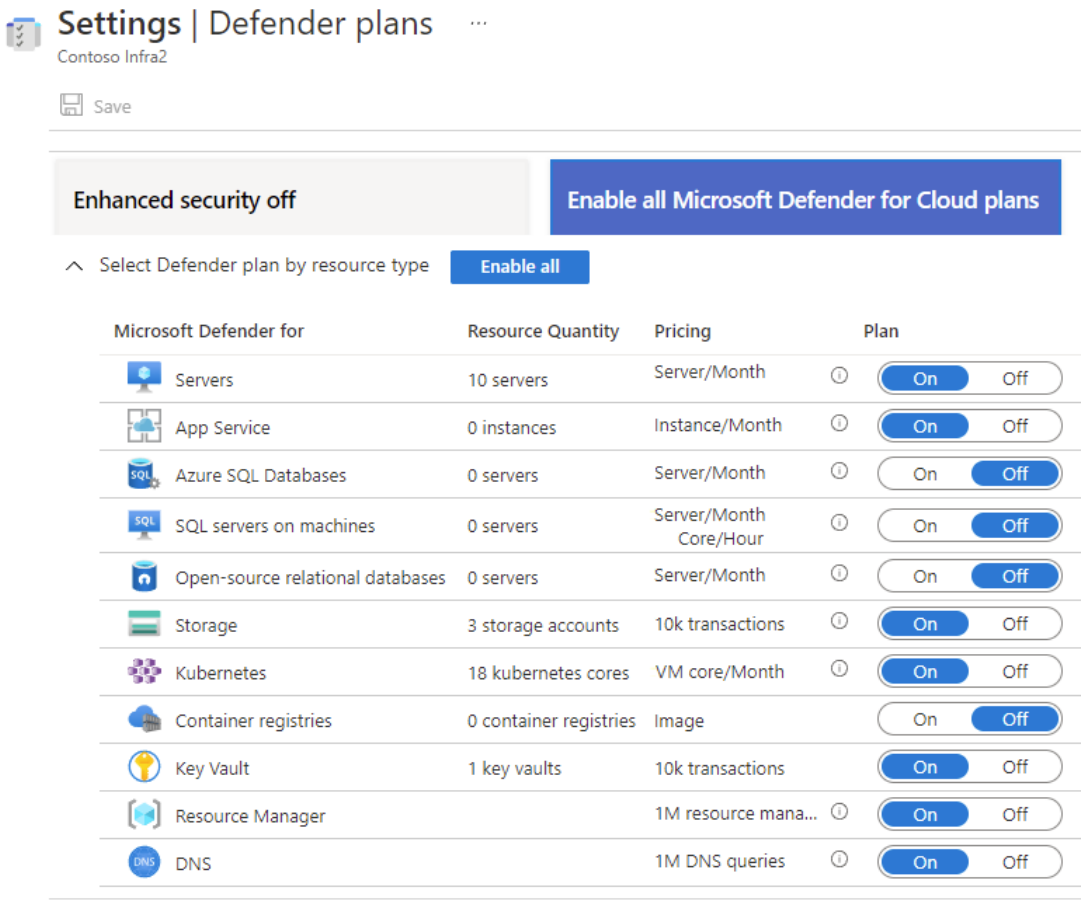

6. To further examine the results of your query, select the resources that interest you.

7. To view the current selected filter options as a query in Resource Graph Explorer, select **Open query**.

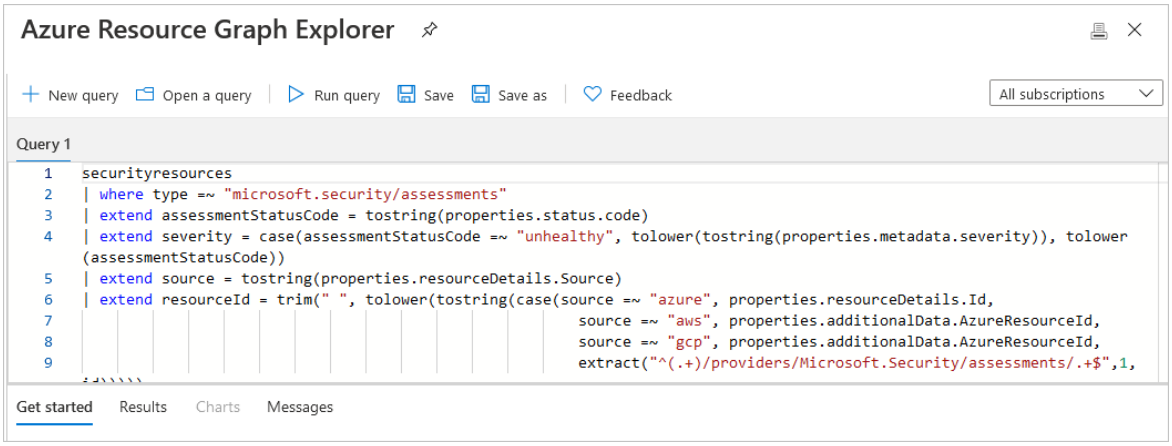

8. If you've defined some filters and left the page open, Defender for Cloud won't update the results automatically. Any changes to resources won't impact the displayed results unless you manually reload the page or select **Refresh**.

Access a software inventory

If you've enabled the integration with Microsoft Defender for Endpoint and enabled Microsoft Defender for servers, you'll have access to the software inventory.

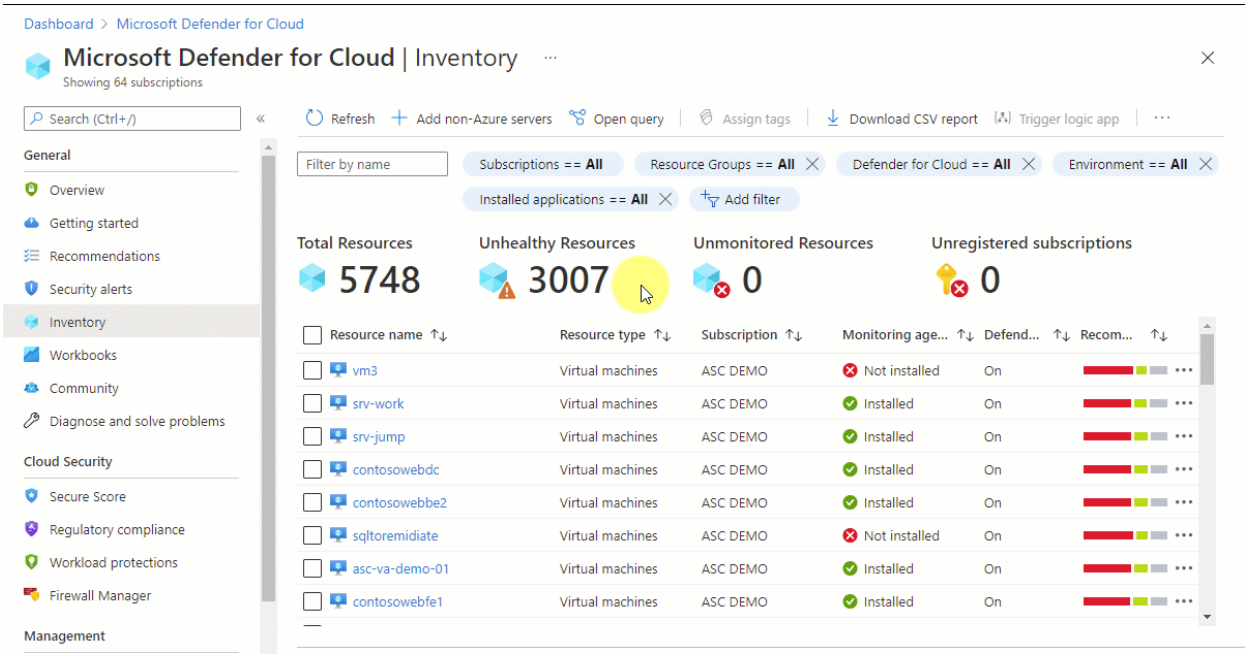

NOTE

The "Blank" option shows machines without Microsoft Defender for Endpoint (or without Microsoft Defender for servers).

As well as the filters in the asset inventory page, you can explore the software inventory data from Azure Resource Graph Explorer.

Examples of using Azure Resource Graph Explorer to access and explore software inventory data:

1. Open **Azure Resource Graph Explorer**.

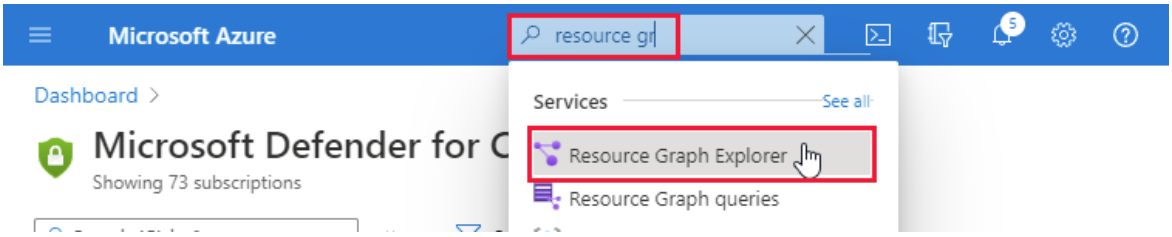

2. Select the following subscription scope: securityresources/softwareinventories
3. Enter any of the following queries (or customize them or write your own!) and select **Run query**.
   - To generate a basic list of installed software:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| project id, Vendor=properties.vendor, Software=properties.softwareName,
Version=properties.version
```

   - To filter by version numbers:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| project id, Vendor=properties.vendor, Software=properties.softwareName,
Version=tostring(properties.    version)
| where Software=="windows_server_2019" and parse_version(Version)
<=parse_version("10.0.17763.1999")
```

- To find machines with a combination of software products:

```
securityresources
| where type == "microsoft.security/softwareinventories"
| extend vmId = properties.azureVmId
| where properties.softwareName == "apache_http_server" or properties.softwareName == "mysql"
| summarize count() by tostring(vmId)
| where count_ > 1
```

- Combination of a software product with another security recommendation:

  (In this example – machines having MySQL installed and exposed management ports)

```
securityresources
| where type == "microsoft.security/softwareinventories"
| extend vmId = tolower(properties.azureVmId)
| where properties.softwareName == "mysql"
| join (
securityresources
| where type == "microsoft.security/assessments"
| where properties.displayName == "Management ports should be closed on your virtual machines"
and properties.status.code == "Unhealthy"
| extend vmId = tolower(properties.resourceDetails.Id)
) on vmId
```

FAQ - Inventory

Why aren't all of my subscriptions, machines, storage accounts, etc. shown?

The inventory view lists your Defender for Cloud connected resources from a Cloud Security Posture Management (CSPM) perspective. The filters don't return every resource in your environment; only the ones with outstanding (or 'active') recommendations.

For example, the following screenshot shows a user with access to 8 subscriptions but only 7 currently have recommendations. So when they filter by **Resource type = Subscriptions**, only those 7 subscriptions with active recommendations appear in the inventory:

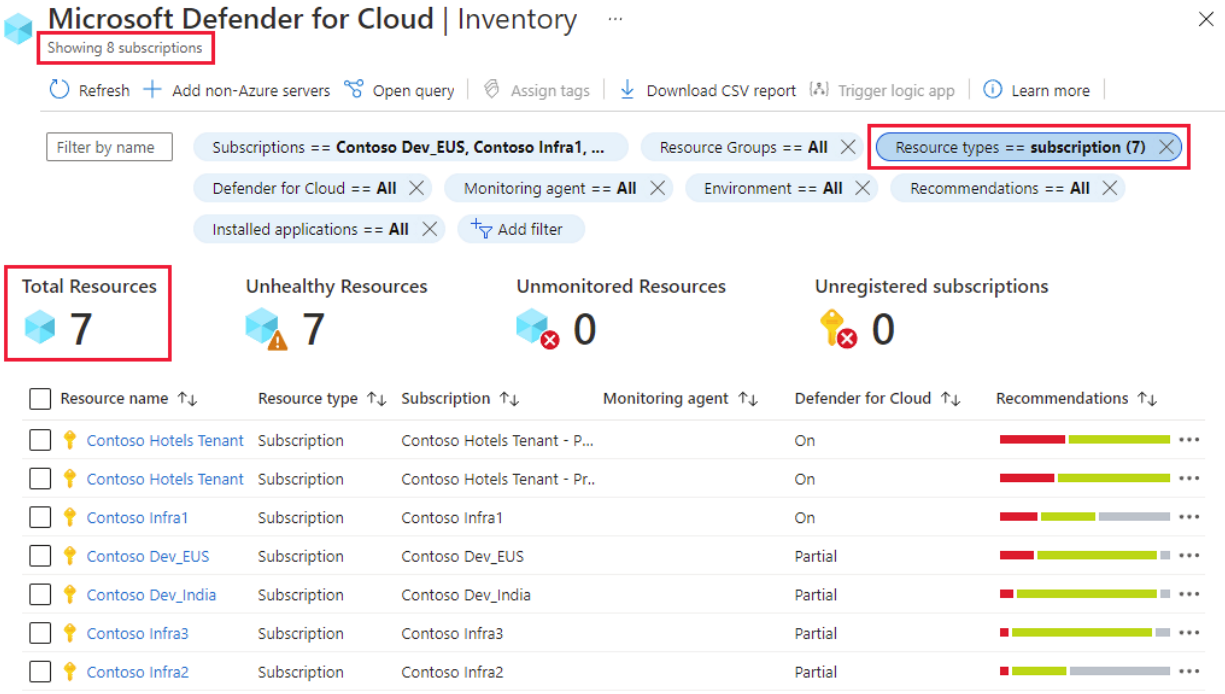

Why do some of my resources show blank values in the Defender for Cloud or monitoring agent columns?

Not all Defender for Cloud monitored resources have agents. For example, Azure Storage accounts or PaaS resources such as disks, Logic Apps, Data Lake Analysis, and Event Hub don't need agents to be monitored by Defender for Cloud.

When pricing or agent monitoring isn't relevant for a resource, nothing will be shown in those columns of inventory.

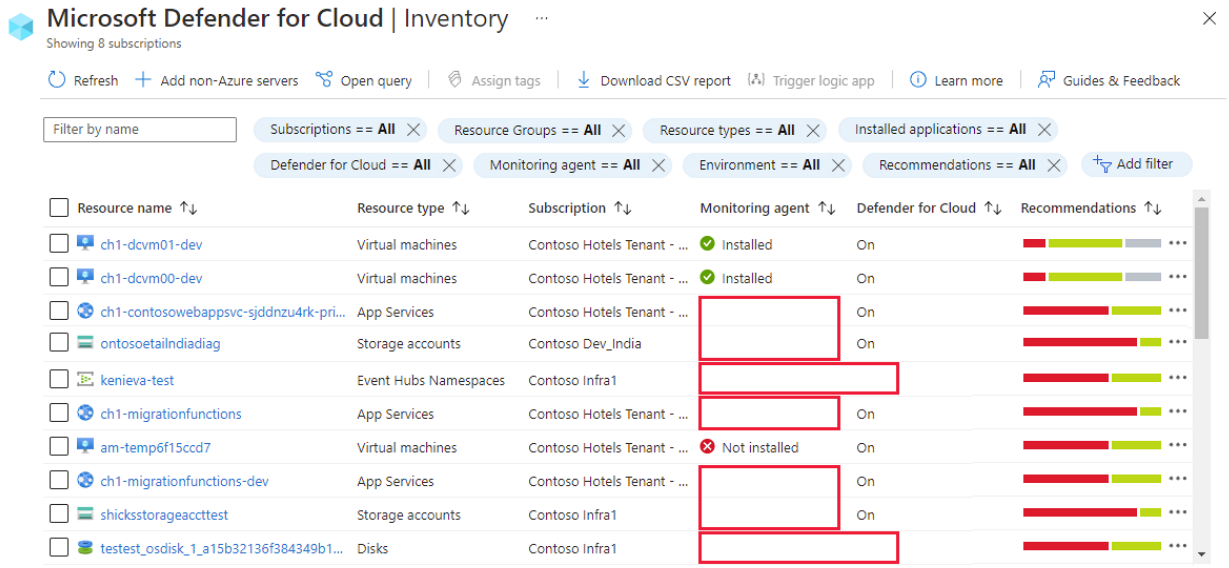

Next steps

This article described the asset inventory page of Microsoft Defender for Cloud.

For more information on related tools, see the following pages:

- Azure Resource Graph (ARG)
- Kusto Query Language (KQL)

File integrity monitoring in Microsoft Defender for Cloud

2/15/2022 • 7 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Learn how to configure file integrity monitoring (FIM) in Microsoft Defender for Cloud using this walkthrough.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers. Using the Log Analytics agent, FIM uploads data to the Log Analytics workspace. Data charges apply, based on the amount of data you upload. See Log Analytics pricing to learn more.
Required roles and permissions:	**Workspace owner** can enable/disable FIM (for more information, see Azure Roles for Log Analytics). **Reader** can view results.
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) Supported only in regions where Azure Automation's change tracking solution is available. ✅ Azure Arc enabled devices. See Supported regions for linked Log Analytics workspace. Learn more about change tracking. ✅ Connected AWS accounts

What is FIM in Defender for Cloud?

File integrity monitoring (FIM), also known as change monitoring, examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack.

Defender for Cloud recommends entities to monitor with FIM, and you can also define your own FIM policies or entities to monitor. FIM informs you about suspicious activity such as:

- File and registry key creation or removal
- File modifications (changes in file size, access control lists, and hash of the content)
- Registry modifications (changes in size, access control lists, type, and the content)

In this tutorial you'll learn how to:

- Review the list of suggested entities to monitor with FIM
- Define your own, custom FIM rules
- Audit changes to your monitored entities
- Use wildcards to simplify tracking across directories

How does FIM work?

The Log Analytics agent uploads data to the Log Analytics workspace. By comparing the current state of these items with the state during the previous scan, FIM notifies you if suspicious modifications have been made.

FIM uses the Azure Change Tracking solution to track and identify changes in your environment. When file integrity monitoring is enabled, you have a **Change Tracking** resource of type **Solution**. For data collection frequency details, see Change Tracking data collection details.

> **NOTE**
> If you remove the **Change Tracking** resource, you will also disable the file integrity monitoring feature in Defender for Cloud.

Which files should I monitor?

When choosing which files to monitor, consider the files that are critical for your system and applications. Monitor files that you don't expect to change without planning. If you choose files that are frequently changed by applications or operating system (such as log files and text files) it'll create a lot of noise, making it difficult to identify an attack.

Defender for Cloud provides the following list of recommended items to monitor based on known attack patterns.

LINUX FILES	WINDOWS FILES	WINDOWS REGISTRY KEYS (HKLM = HKEY_LOCAL_MACHINE)
/bin/login	C:\autoexec.bat	HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType 0\CryptSIPDllRemoveSignedDataMsg{C689AAB8-8E78-11D0-8C47-00C04FC295EE}
/bin/passwd	C:\boot.ini	HKLM\SOFTWARE\Microsoft\Cryptography\OID\EncodingType 0\CryptSIPDllRemoveSignedDataMsg{603BCC1F-4B59-4E08-B724-D2C6297EF351}
/etc/*.conf	C:\config.sys	HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\IniFileMapping\SYSTEM.ini\boot
/usr/bin	C:\Windows\system.ini	HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Windows
/usr/sbin	C:\Windows\win.ini	HKLM\SOFTWARE\Microsoft\Windows NT\CurrentVersion\Winlogon

LINUX FILES	WINDOWS FILES	WINDOWS REGISTRY KEYS (HKLM = HKEY_LOCAL_MACHINE)
/bin	C:\Windows\regedit.exe	HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders
/sbin	C:\Windows\System32\userinit.exe	HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders
/boot	C:\Windows\explorer.exe	HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\Run
/usr/local/bin	C:\Program Files\Microsoft Security Client\msseces.exe	HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnce
/usr/local/sbin		HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\RunOnceEx
/opt/bin		HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\RunServices
/opt/sbin		HKLM\SOFTWARE\Microsoft\Windows\CurrentVersion\RunServicesOnce
/etc/crontab		HKLM\SOFTWARE\WOW6432Node\Microsoft\Cryptography\OID\EncodingType 0\CryptSIPDllRemoveSignedDataMsg{C689AAB8-8E78-11D0-8C47-00C04FC295EE}
/etc/init.d		HKLM\SOFTWARE\WOW6432Node\Microsoft\Cryptography\OID\EncodingType 0\CryptSIPDllRemoveSignedDataMsg{603BCC1F-4B59-4E08-B724-D2C6297EF351}
/etc/cron.hourly		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows NT\CurrentVersion\IniFileMapping\system.ini\boot
/etc/cron.daily		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows NT\CurrentVersion\Windows
/etc/cron.weekly		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows NT\CurrentVersion\Winlogon
/etc/cron.monthly		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\Explorer\Shell Folders

LINUX FILES	WINDOWS FILES	WINDOWS REGISTRY KEYS (HKLM = HKEY_LOCAL_MACHINE)
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\Explorer\User Shell Folders
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\Run
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\RunOnce
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\RunOnceEx
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\RunServices
		HKLM\SOFTWARE\WOW6432Node\Microsoft\Windows\CurrentVersion\RunServicesOnce
		HKLM\SYSTEM\CurrentControlSet\Control\hivelist
		HKLM\SYSTEM\CurrentControlSet\Control\Session Manager\KnownDLLs
		HKLM\SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\DomainProfile
		HKLM\SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\PublicProfile
		HKLM\SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile

Enable file integrity monitoring

FIM is only available from Defender for Cloud's pages in the Azure portal. There is currently no REST API for working with FIM.

1. From the **Workload protections** dashboard's **Advanced protection** area, select **File integrity monitoring**.

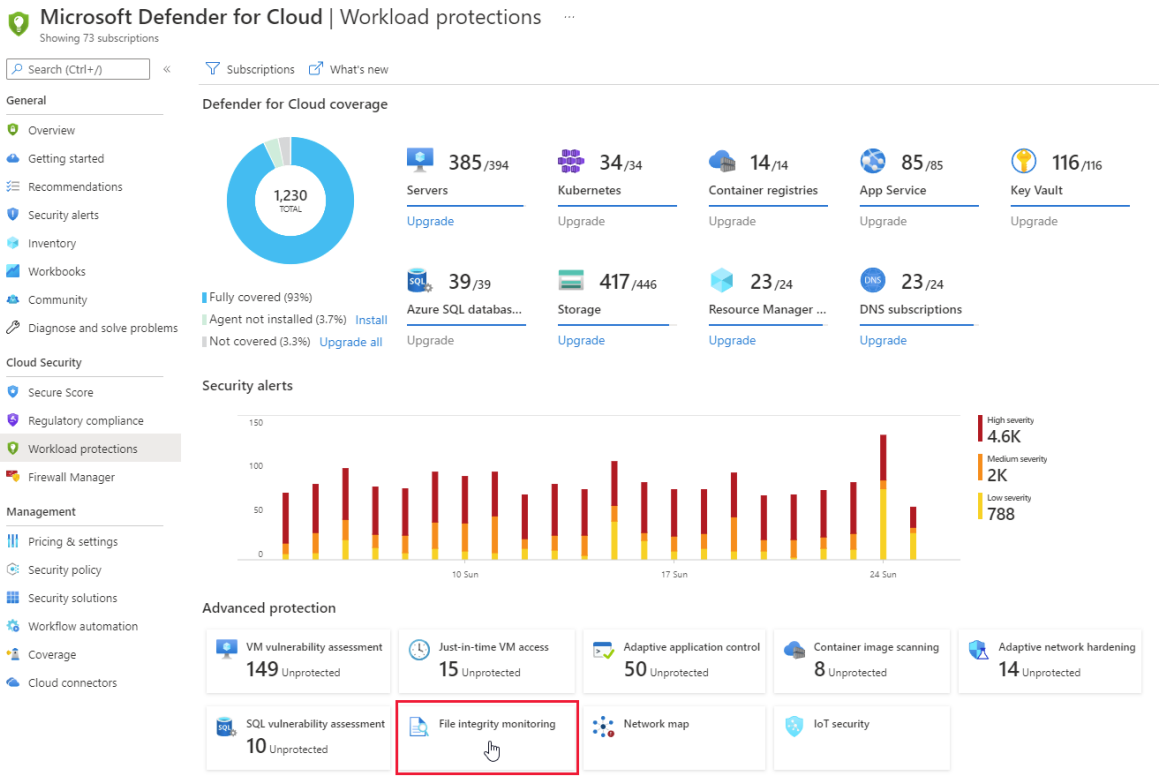

The **File integrity monitoring** configuration page opens.

The following information is provided for each workspace:

- Total number of changes that occurred in the last week (you may see a dash "-" if FIM is not enabled on the workspace)
- Total number of computers and VMs reporting to the workspace
- Geographic location of the workspace
- Azure subscription that the workspace is under

2. Use this page to:

- Access and view the status and settings of each workspace
- UPGRADE PLAN Upgrade the workspace to use enhanced security features. This icon Indicates that the workspace or subscription isn't protected with Microsoft Defender for servers. To use the FIM features, your subscription must be protected with this plan. For more information, see Microsoft Defender for Cloud's enhanced security features.
- ENABLE Enable FIM on all machines under the workspace and configure the FIM options. This icon indicates that FIM is not enabled for the workspace.

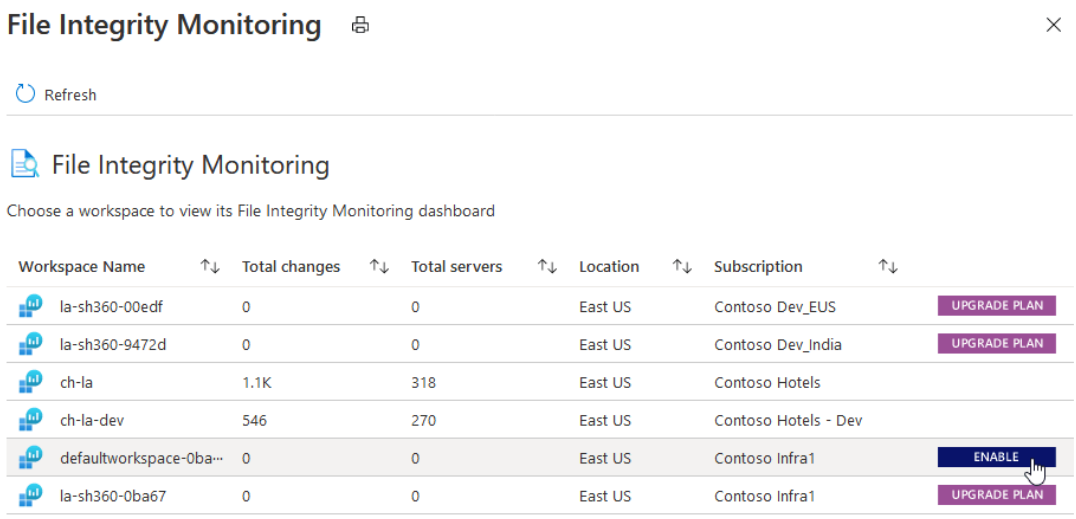

> **TIP**
> If there's no enable or upgrade button, and the space is blank, it means that FIM is already enabled on the workspace.

3. Select **ENABLE**. The details of the workspace including the number of Windows and Linux machines under the workspace is shown.

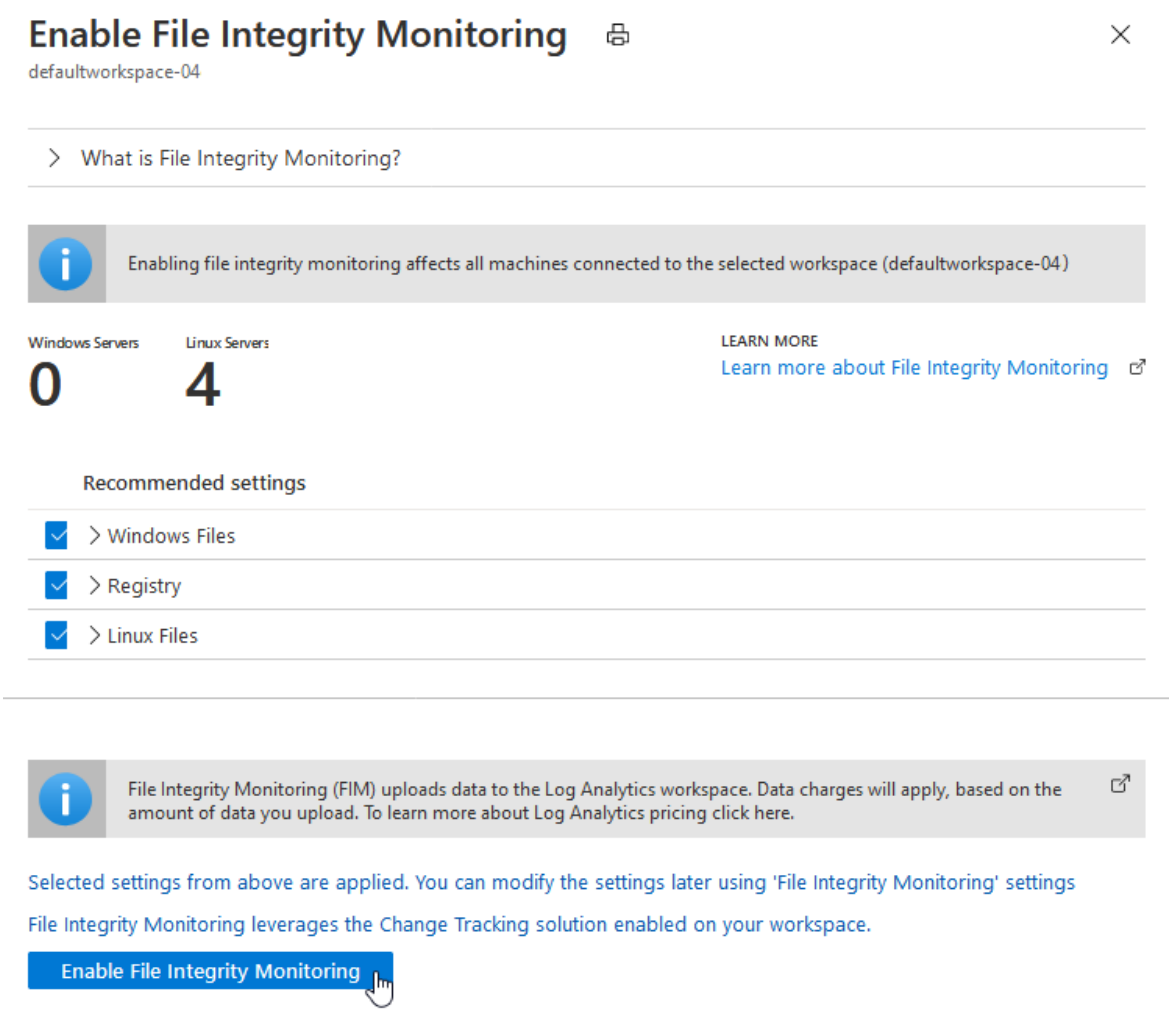

   The recommended settings for Windows and Linux are also listed. Expand **Windows files**, **Registry**, and **Linux files** to see the full list of recommended items.

4. Clear the checkboxes for any recommended entities you do not want to be monitored by FIM.

5. Select **Apply file integrity monitoring** to enable FIM.

> **NOTE**
> You can change the settings at any time. See Edit monitored entities below to learn more.

Audit monitored workspaces

The **File integrity monitoring** dashboard displays for workspaces where FIM is enabled. The FIM dashboard opens after you enable FIM on a workspace or when you select a workspace in the **file integrity monitoring** window that already has FIM enabled.

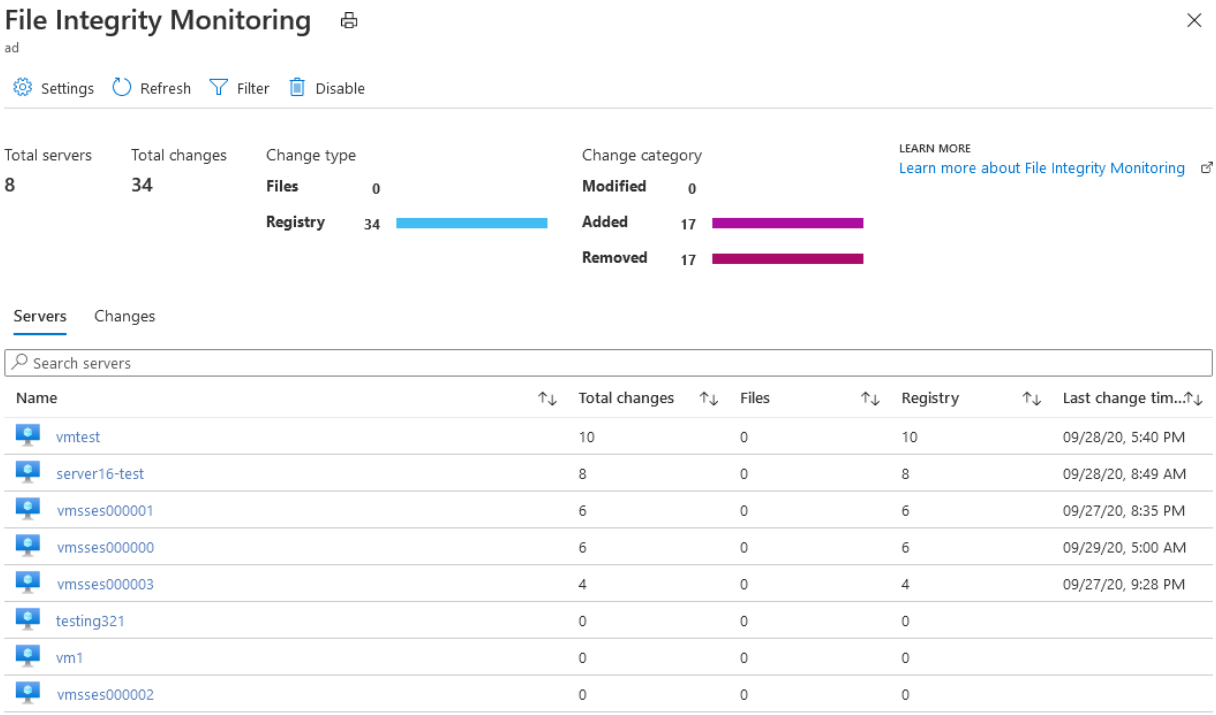

The FIM dashboard for a workspace displays the following details:

- Total number of machines connected to the workspace
- Total number of changes that occurred during the selected time period
- A breakdown of change type (files, registry)
- A breakdown of change category (modified, added, removed)

Select **Filter** at the top of the dashboard to change the time period for which changes are shown.

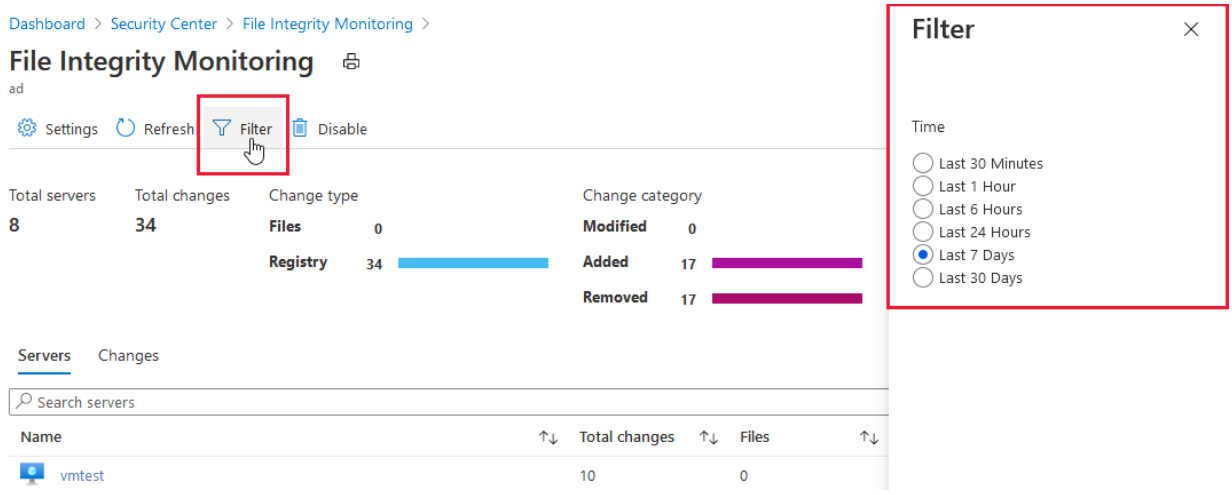

The **Servers** tab lists the machines reporting to this workspace. For each machine, the dashboard lists:

- Total changes that occurred during the selected period of time
- A breakdown of total changes as file changes or registry changes

When you select a machine, the query appears along with the results that identify the changes made during the selected time period for the machine. You can expand a change for more information.

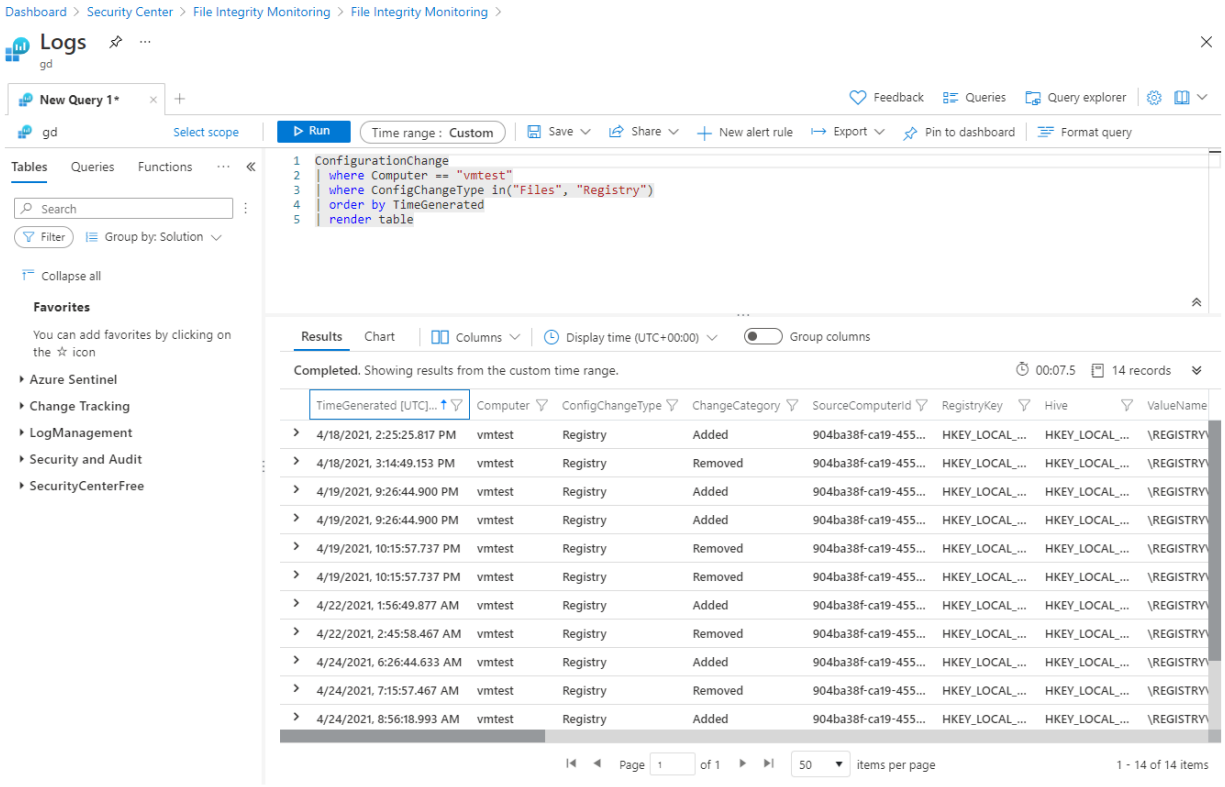

The **Changes** tab (shown below) lists all changes for the workspace during the selected time period. For each entity that was changed, the dashboard lists the:

- Machine that the change occurred on
- Type of change (registry or file)
- Category of change (modified, added, removed)
- Date and time of change

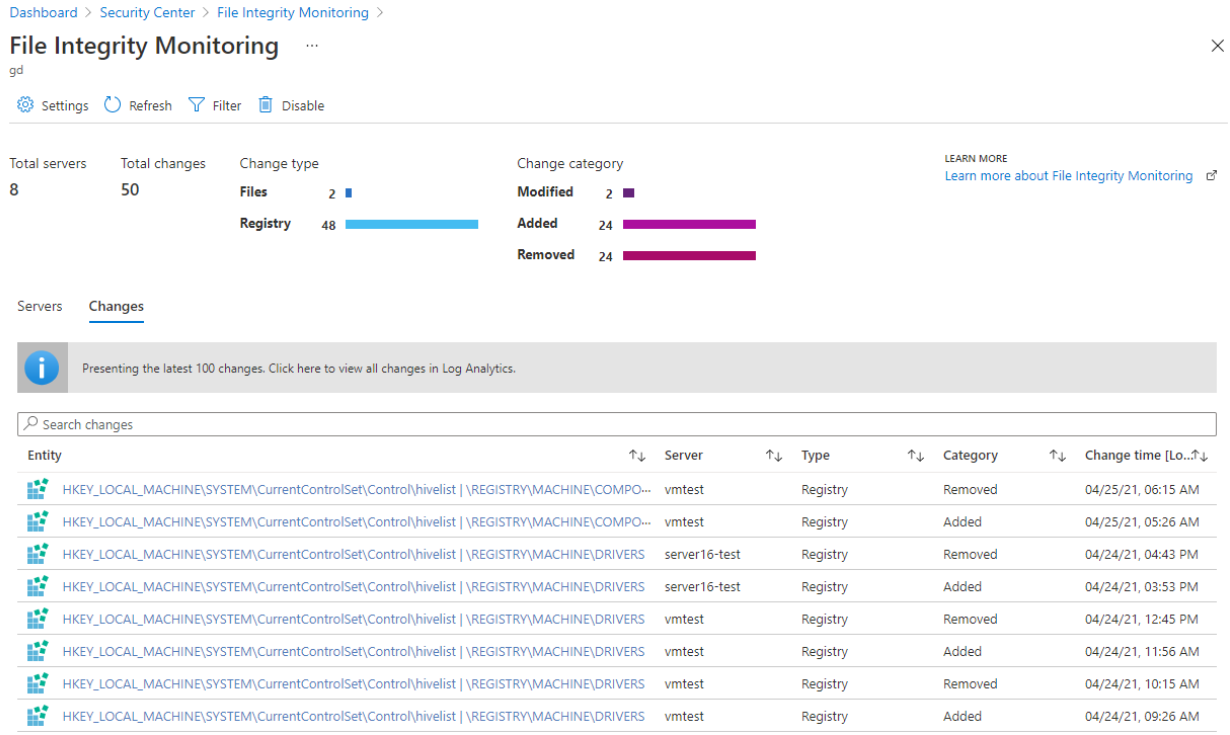

Change details opens when you enter a change in the search field or select an entity listed under the **Changes** tab.

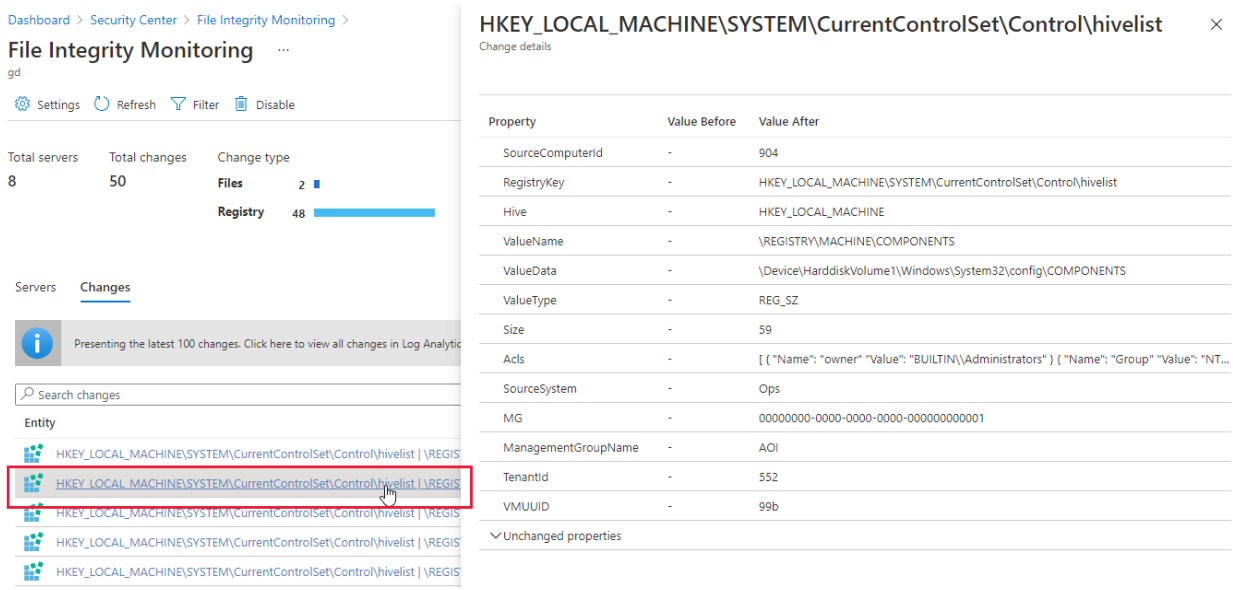

Edit monitored entities

1. From the **File integrity monitoring dashboard** for a workspace, select **Settings** from the toolbar.

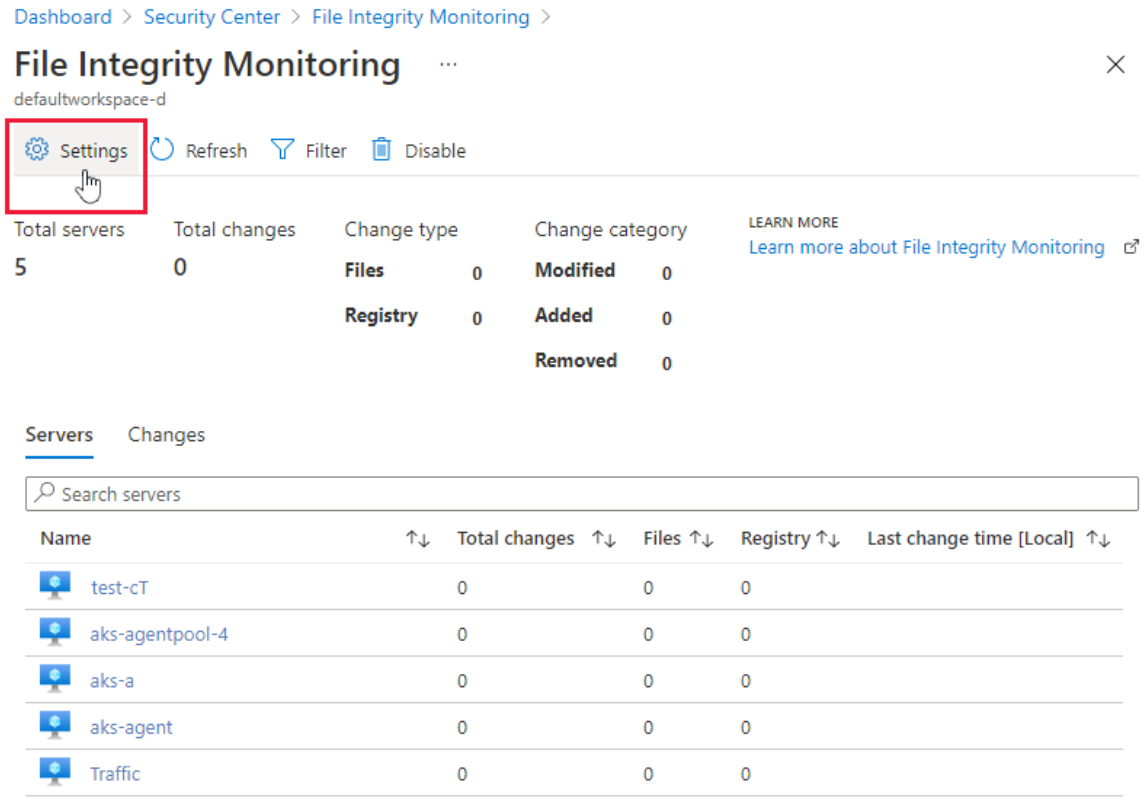

Workspace Configuration opens with tabs for each type of element that can be monitored:

- Windows registry
- Windows files
- Linux Files
- File content
- Windows services

Each tab lists the entities that you can edit in that category. For each entity listed, Defender for Cloud identifies whether FIM is enabled (true) or not enabled (false). Edit the entity to enable or disable FIM.

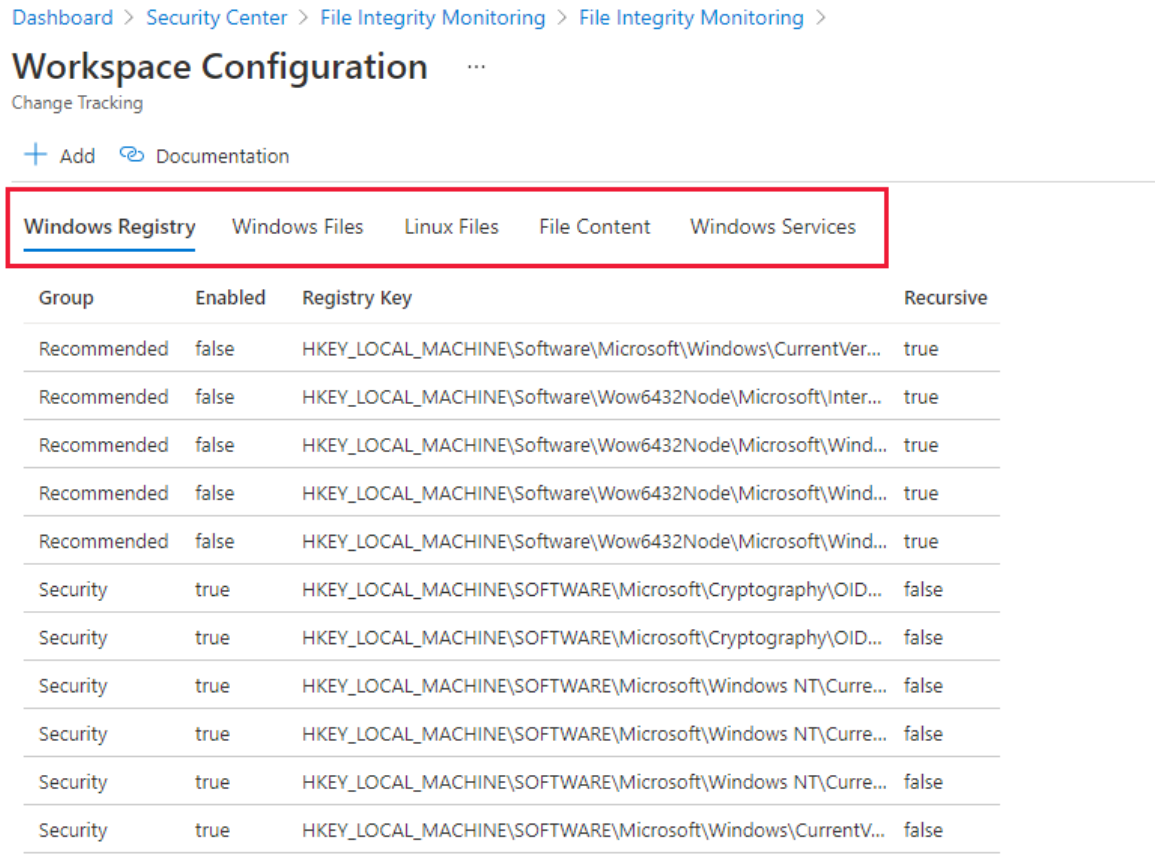

2. Select an entry from one of the tabs and edit any of the available fields in the **Edit for Change Tracking** pane. Options include:

   - Enable (True) or disable (False) file integrity monitoring
   - Provide or change the entity name
   - Provide or change the value or path
   - Delete the entity

3. Discard or save your changes.

Add a new entity to monitor

1. From the **File integrity monitoring dashboard** for a workspace, select **Settings** from the toolbar.

   The **Workspace Configuration** opens.

2. One the **Workspace Configuration**:

   a. Select the tab for the type of entity that you want to add: Windows registry, Windows files, Linux Files, file content, or Windows services.

   b. Select **Add**.

      In this example, we selected **Linux Files**.

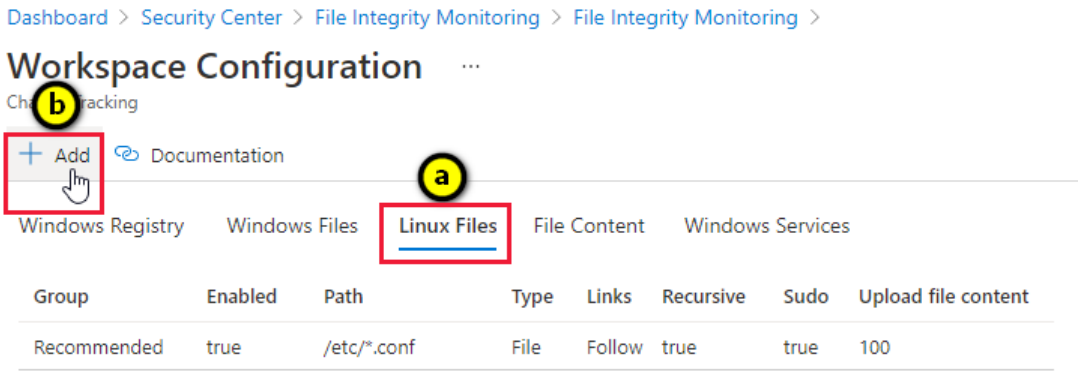

3. Select **Add**. **Add for Change Tracking** opens.

4. Enter the necessary information and select **Save**.

Folder and path monitoring using wildcards

Use wildcards to simplify tracking across directories. The following rules apply when you configure folder monitoring using wildcards:

- Wildcards are required for tracking multiple files.
- Wildcards can only be used in the last segment of a path, such as C:\folder\file or /etc/*.conf
- If an environment variable includes a path that is not valid, validation will succeed but the path will fail when inventory runs.
- When setting the path, avoid general paths such as c:*.* which will result in too many folders being traversed.

Disable FIM

You can disable FIM. FIM uses the Azure Change Tracking solution to track and identify changes in your environment. By disabling FIM, you remove the Change Tracking solution from selected workspace.

To disable FIM:

1. From the **File integrity monitoring dashboard** for a workspace, select **Disable**.

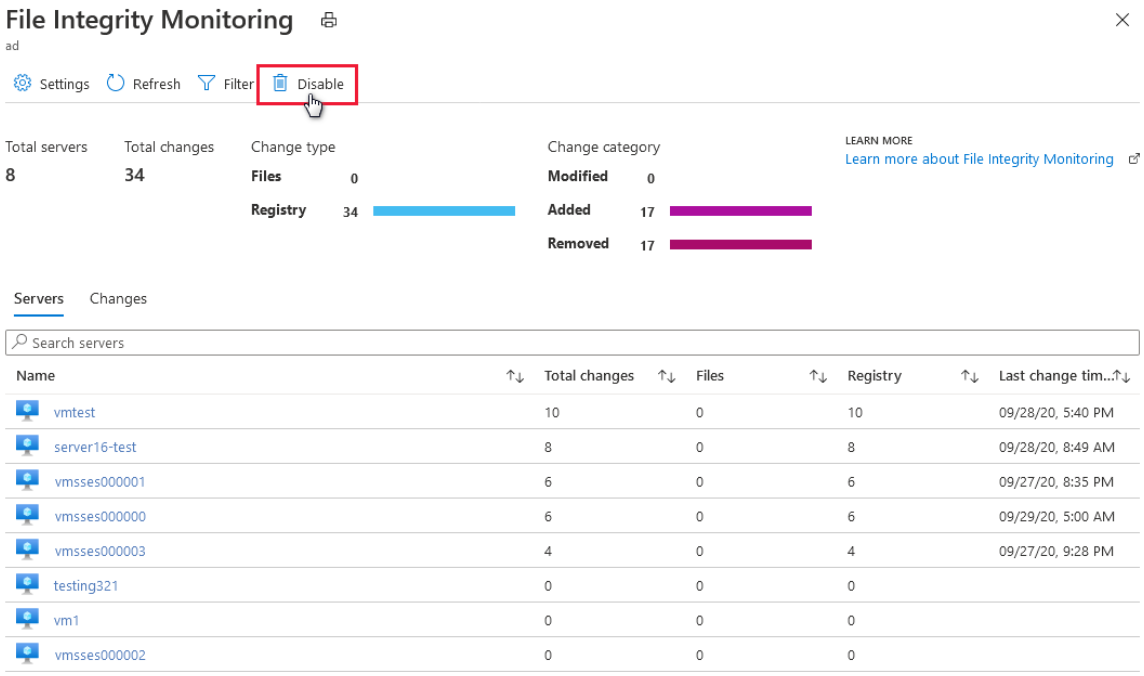

2. Select **Remove**.

Next steps

In this article, you learned to use file integrity monitoring (FIM) in Defender for Cloud. To learn more about Defender for Cloud, see the following pages:

- Setting security policies -- Learn how to configure security policies for your Azure subscriptions and resource groups.
- Managing security recommendations -- Learn how recommendations help you protect your Azure resources.
- Azure Security blog--Get the latest Azure security news and information.

Compare baselines using File Integrity Monitoring (FIM)

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

File Integrity Monitoring (FIM) informs you when changes occur to sensitive areas in your resources, so you can investigate and address unauthorized activity. FIM monitors Windows files, Windows registries, and Linux files.

This topic explains how to enable FIM on the files and registries. For more information about FIM, see File Integrity Monitoring in Microsoft Defender for Cloud.

Why use FIM?

Operating system, applications, and associated configurations control the behavior and security state of your resources. Therefore, attackers target the files that control your resources, in order to overtake a resource's operating system and/or execute activities without being detected.

In fact, many regulatory compliance standards such as PCI-DSS & ISO 17799 require implementing FIM controls.

Enable built-in recursive registry checks

The FIM registry hive defaults provide a convenient way to monitor recursive changes within common security areas. For example, an adversary may configure a script to execute in LOCAL_SYSTEM context by configuring an execution at startup or shutdown. To monitor changes of this type, enable the built-in check.

Recommended	false	HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Group Policy\Scripts\Shutdown
Recommended	false	HKEY_LOCAL_MACHINE\Software\Microsoft\Windows\CurrentVersion\Group Policy\Scripts\Startup
Recommended	false	HKEY_LOCAL_MACHINE\Software\Wow6432Node\Microsoft\Internet Explorer\Extensions

NOTE

Recursive checks apply only to recommended security hives and not to custom registry paths.

Add a custom registry check

FIM baselines start by identifying characteristics of a known-good state for the operating system and supporting application. For this example, we will focus on the password policy configurations for Windows Server 2008 and higher.

POLICY NAME	REGISTRY SETTING
Domain controller: Refuse machine account password changes	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\RefusePasswordChange
Domain member: Digitally encrypt or sign secure channel data (always)	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\RequireSignOrSeal
Domain member: Digitally encrypt secure channel data (when possible)	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\SealSecureChannel
Domain member: Digitally sign secure channel data (when possible)	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\SignSecureChannel
Domain member: Disable machine account password changes	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\DisablePasswordChange
Domain member: Maximum machine account password age	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\MaximumPasswordAge
Domain member: Require strong (Windows 2000 or later) session key	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\RequireStrongKey
Network security: Restrict NTLM: NTLM authentication in this domain	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\RestrictNTLMInDomain
Network security: Restrict NTLM: Add server exceptions in this domain	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\DCAllowedNTLMServers
Network security: Restrict NTLM: Audit NTLM authentication in this domain	MACHINE\System\CurrentControlSet\Services \Netlogon\Parameters\AuditNTLMInDomain

> **NOTE**
>
> To learn more about registry settings supported by various operating system versions, refer to the Group Policy Settings reference spreadsheet.

To configure FIM to monitor registry baselines:

1. In the **Add Windows Registry for Change Tracking** window, in the **Windows Registry Key** text box, enter the following registry key:

```
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Netlogon\Parameters
```

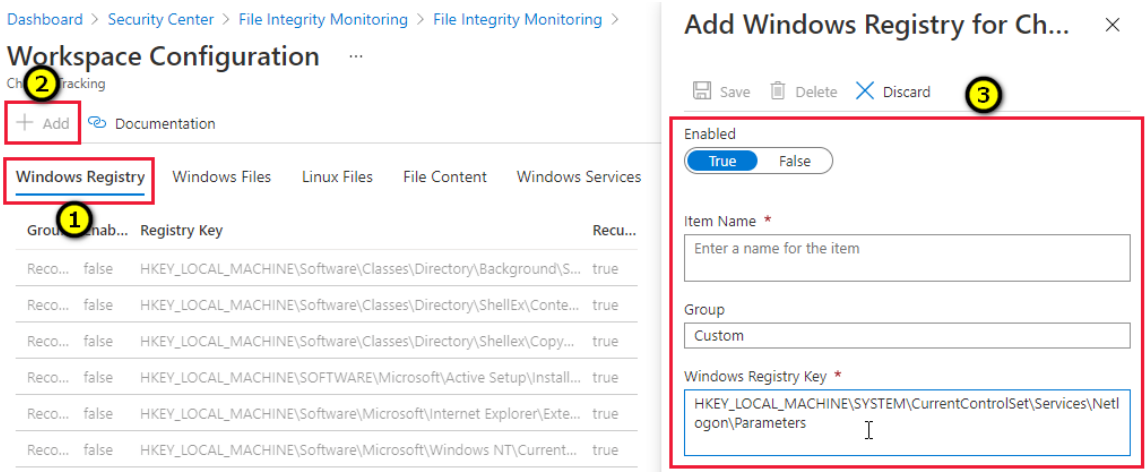

Track changes to Windows files

1. In the **Add Windows File for Change Tracking** window, in the **Enter path** text box, enter the folder which contains the files that you want to track. In the example in the following figure, **Contoso Web App** resides in the D:\ drive within the **ContosWebApp** folder structure.

2. Create a custom Windows file entry by providing a name of the setting class, enabling recursion, and specifying the top folder with a wildcard (*) suffix.

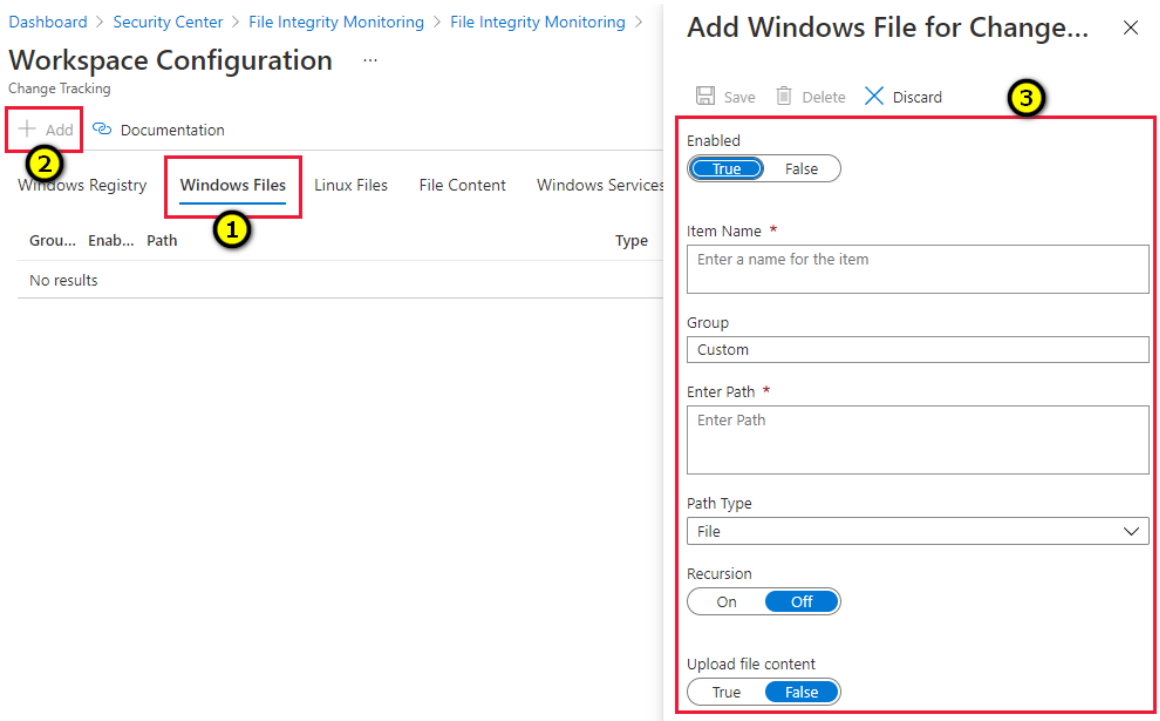

Retrieve change data

File Integrity Monitoring data resides within the Azure Log Analytics / ConfigurationChange table set.

1. Set a time range to retrieve a summary of changes by resource. In the following example, we are retrieving all changes in the last fourteen days in the categories of registry and files:

```
ConfigurationChange
| where TimeGenerated > ago(14d)
| where ConfigChangeType in ('Registry', 'Files')
| summarize count() by Computer, ConfigChangeType
```

2. To view details of the registry changes:

   a. Remove **Files** from the **where** clause,

   b. Remove the summarization line and replace it with an ordering clause:

```
ConfigurationChange
| where TimeGenerated > ago(14d)
| where ConfigChangeType in ('Registry')
| order by Computer, RegistryKey
```

Reports can be exported to CSV for archival and/or channeled to a Power BI report.

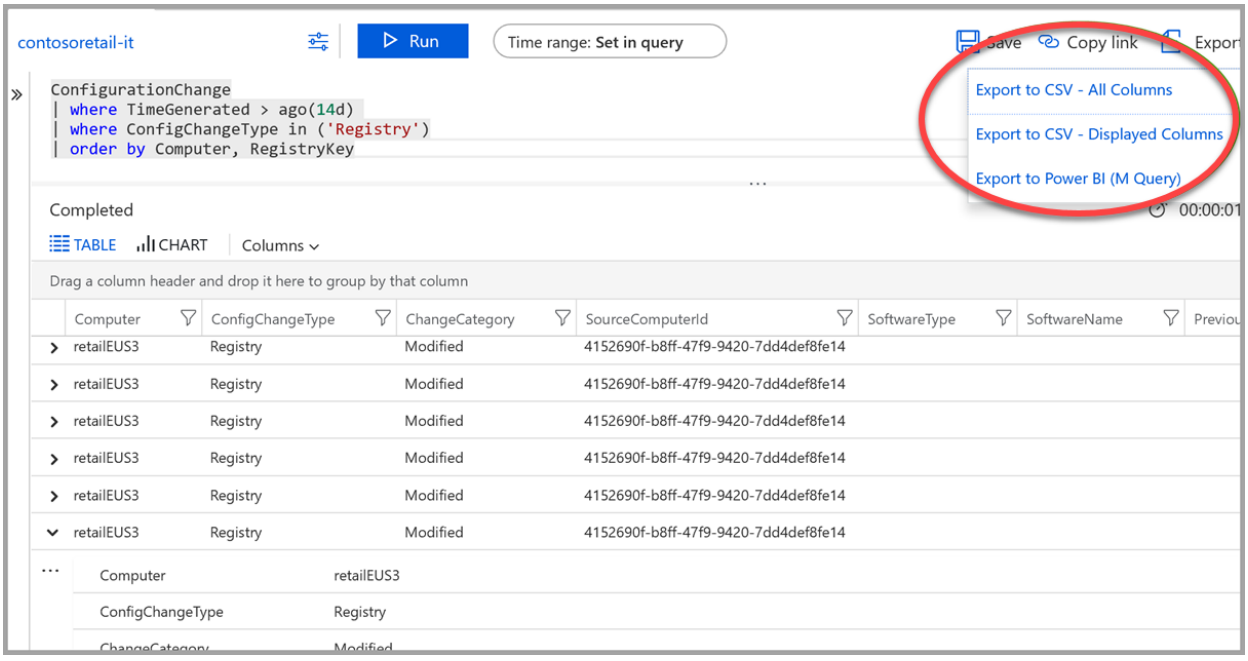

Improve your network security posture with adaptive network hardening

2/15/2022 • 6 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Adaptive network hardening is an agentless feature of Microsoft Defender for Cloud - nothing needs to be installed on your machines to benefit from this network hardening tool.

This page explains how to configure and manage adaptive network hardening in Defender for Cloud.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers
Required roles and permissions:	Write permissions on the machine's NSGs
Clouds:	✅ Commercial clouds ❌ National (Azure Government, Azure China 21Vianet) ❌ Connected AWS accounts

What is adaptive network hardening?

Applying network security groups (NSG) to filter traffic to and from resources, improves your network security posture. However, there can still be some cases in which the actual traffic flowing through the NSG is a subset of the NSG rules defined. In these cases, further improving the security posture can be achieved by hardening the NSG rules, based on the actual traffic patterns.

Adaptive network hardening provides recommendations to further harden the NSG rules. It uses a machine learning algorithm that factors in actual traffic, known trusted configuration, threat intelligence, and other indicators of compromise, and then provides recommendations to allow traffic only from specific IP/port tuples.

For example, let's say the existing NSG rule is to allow traffic from 140.20.30.10/24 on port 22. Based on traffic analysis, adaptive network hardening might recommend narrowing the range to allow traffic from 140.23.30.10/29, and deny all other traffic to that port. For the full list of supported ports, see the FAQ entry Which ports are supported?.

View hardening alerts and recommended rules

1. From Defender for Cloud's menu, open the **Workload protections** dashboard.
2. Select the adaptive network hardening tile (1), or the insights panel item related to adaptive network hardening (2).

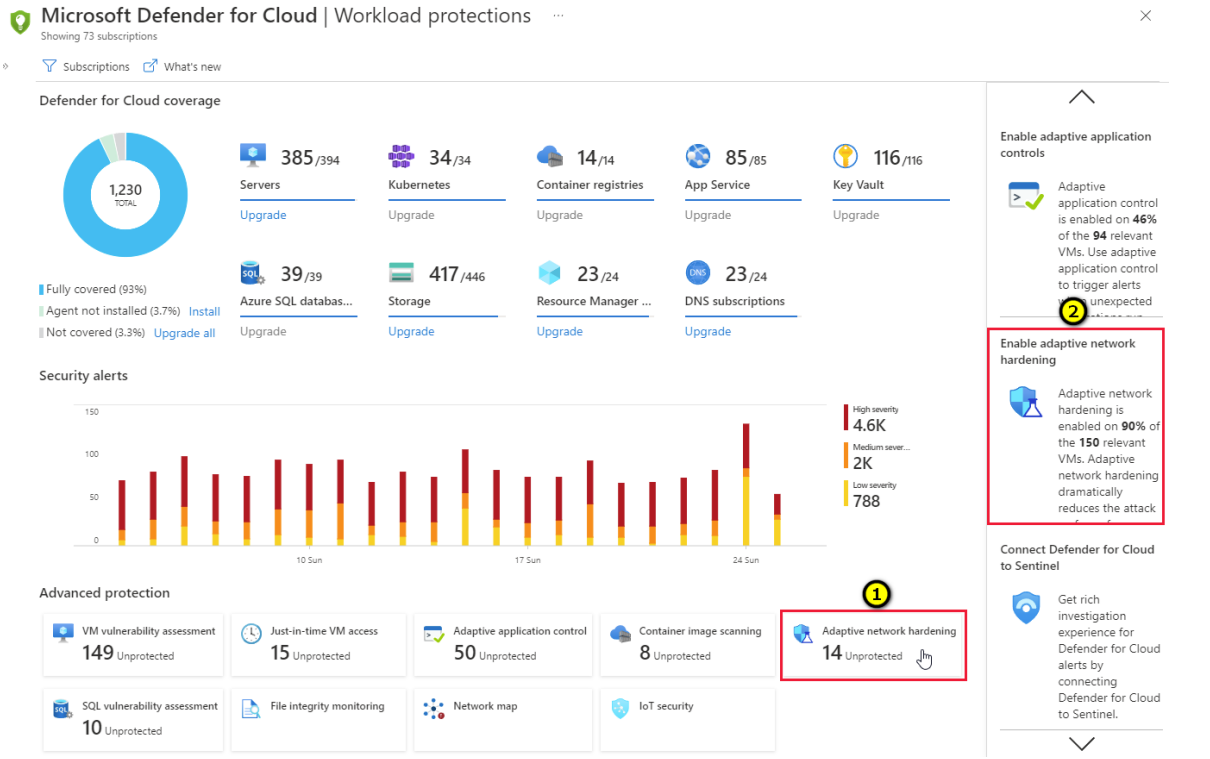

> **TIP**
>
> The insights panel shows the percentage of your VMs currently defended with adaptive network hardening.

3. The details page for the **Adaptive Network Hardening recommendations should be applied on internet facing virtual machines** recommendation opens with your network VMs grouped into three tabs:

   - **Unhealthy resources**: VMs that currently have recommendations and alerts that were triggered by running the adaptive network hardening algorithm.
   - **Healthy resources**: VMs without alerts and recommendations.
   - **Unscanned resources**: VMs that the adaptive network hardening algorithm cannot be run on because of one of the following reasons:
     - **VMs are Classic VMs**: Only Azure Resource Manager VMs are supported.
     - **Not enough data is available**: In order to generate accurate traffic hardening recommendations, Defender for Cloud requires at least 30 days of traffic data.
     - **VM is not protected by Microsoft Defender for servers**: Only VMs protected with Microsoft Defender for servers are eligible for this feature.

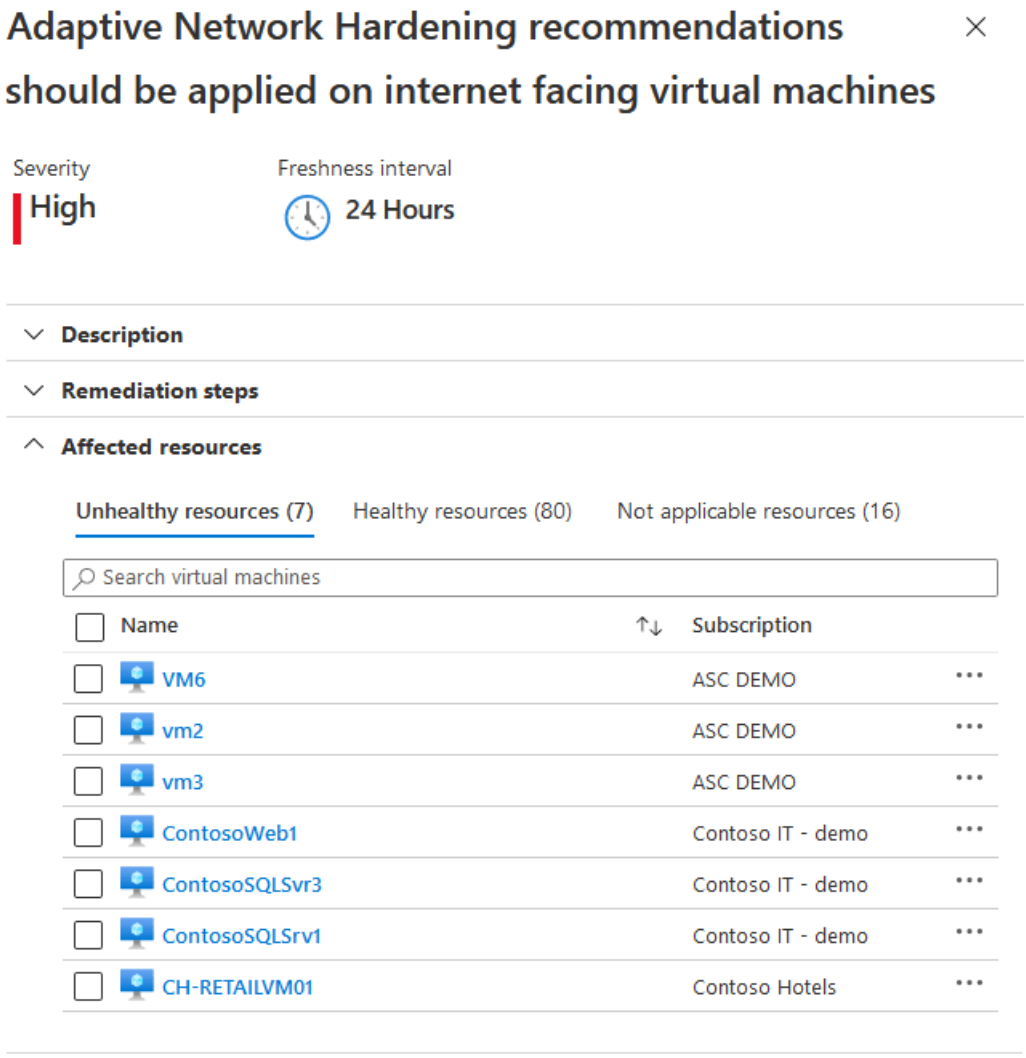

4. From the **Unhealthy resources** tab, select a VM to view its alerts and the recommended hardening rules to apply.

   - The **Rules** tab lists the rules that adaptive network hardening recommends you add
   - The **Alerts** tab lists the alerts that were generated due to traffic, flowing to the resource, which is not within the IP range allowed in the recommended rules.

5. Optionally, edit the rules:

   - Modify a rule
   - Delete a rule
   - Add a rule

6. Select the rules that you want to apply on the NSG, and select **Enforce**.

> **TIP**
> If the allowed source IP ranges shows as 'None', it means that recommended rule is a *deny* rule, otherwise, it is an *allow* rule.

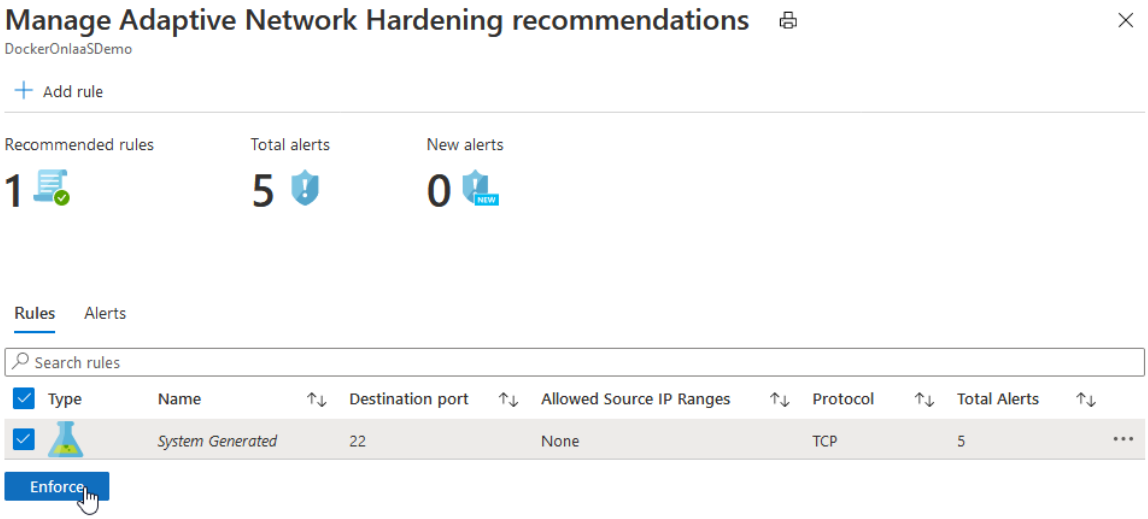

> **NOTE**
>
> The enforced rules are added to the NSG(s) protecting the VM. (A VM could be protected by an NSG that is associated to its NIC, or the subnet in which the VM resides, or both)

Modify a rule

You may want to modify the parameters of a rule that has been recommended. For example, you may want to change the recommended IP ranges.

Some important guidelines for modifying an adaptive network hardening rule:

- You cannot change **allow** rules to become **deny** rules.
- You can modify the parameters of **allow** rules only.

  Creating and modifying "deny" rules is done directly on the NSG. For more information, see Create, change, or delete a network security group.
- A **Deny all traffic** rule is the only type of "deny" rule that would be listed here, and it cannot be modified. You can, however, delete it (see Delete a rule). To learn about this type of rule, see the FAQ entry When should I use a "Deny all traffic" rule?.

To modify an adaptive network hardening rule:

1. To modify some of the parameters of a rule, in the **Rules** tab, select on the three dots (...) at the end of the rule's row, and select **Edit**.

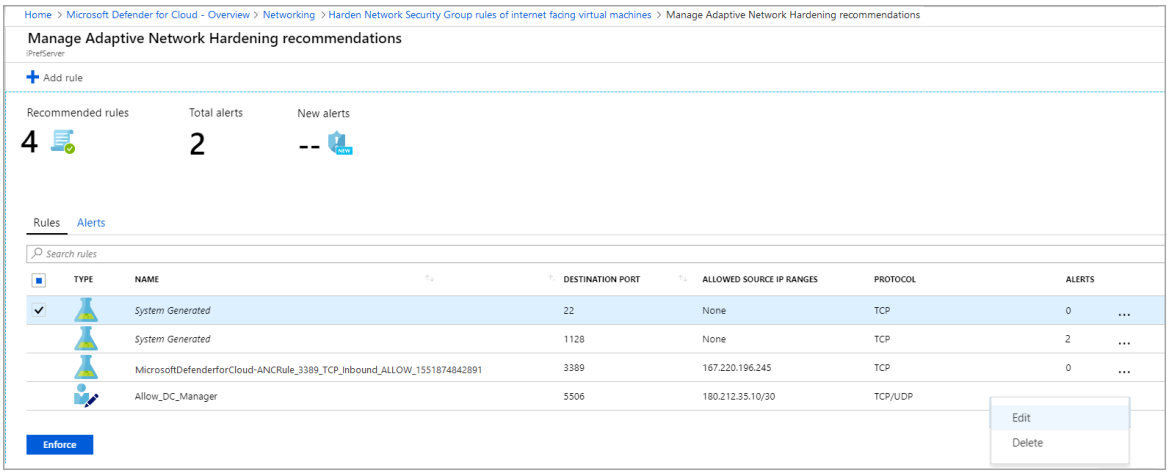

2. In the **Edit rule** window, update the details that you want to change, and select **Save**.

> **NOTE**
>
> After selecting **Save**, you have successfully changed the rule. *However, you have not applied it to the NSG.* To apply it, you must select the rule in the list, and select **Enforce** (as explained in the next step).

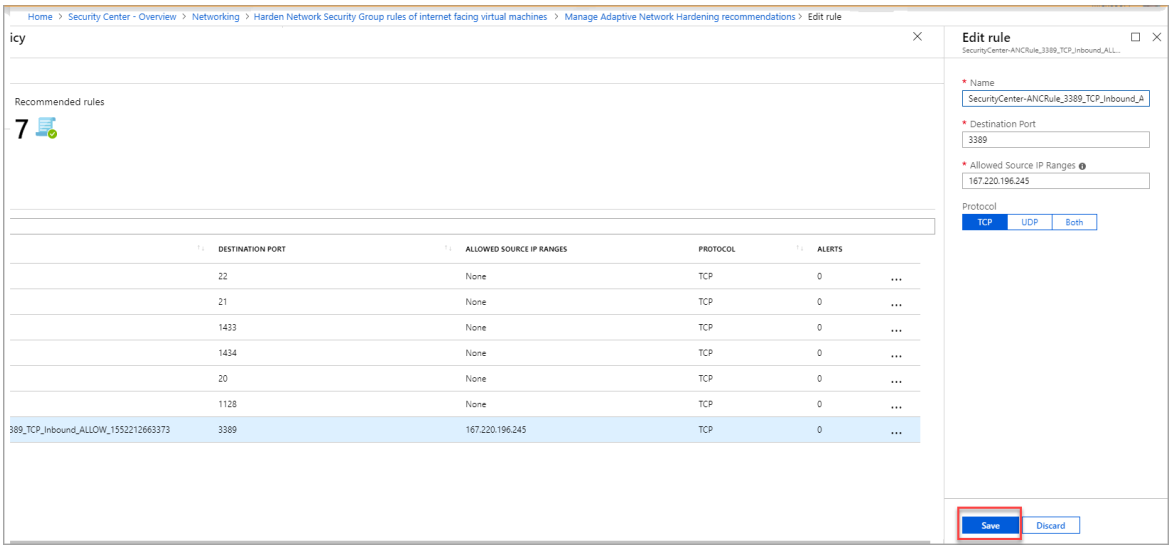

3. To apply the updated rule, from the list, select the updated rule and select **Enforce**.

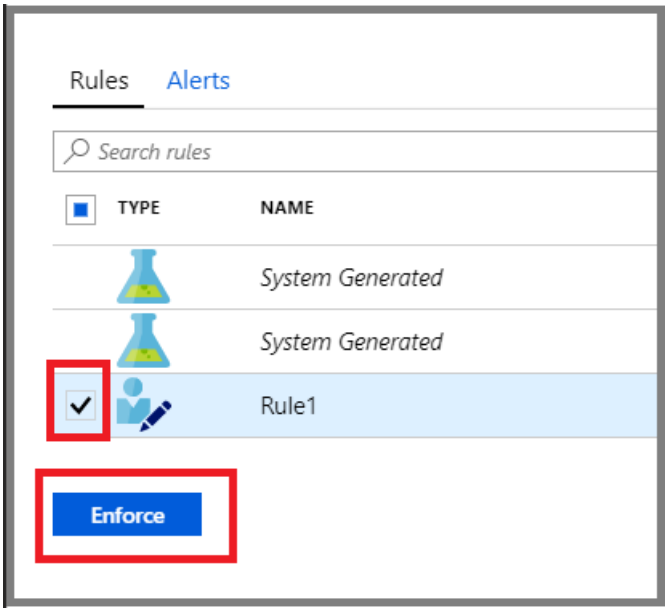

Add a new rule

You can add an "allow" rule that was not recommended by Defender for Cloud.

> **NOTE**
>
> Only "allow" rules can be added here. If you want to add "deny" rules, you can do so directly on the NSG. For more information, see Create, change, or delete a network security group.

To add an adaptive network hardening rule:

1. From the top toolbar, select **Add rule**.

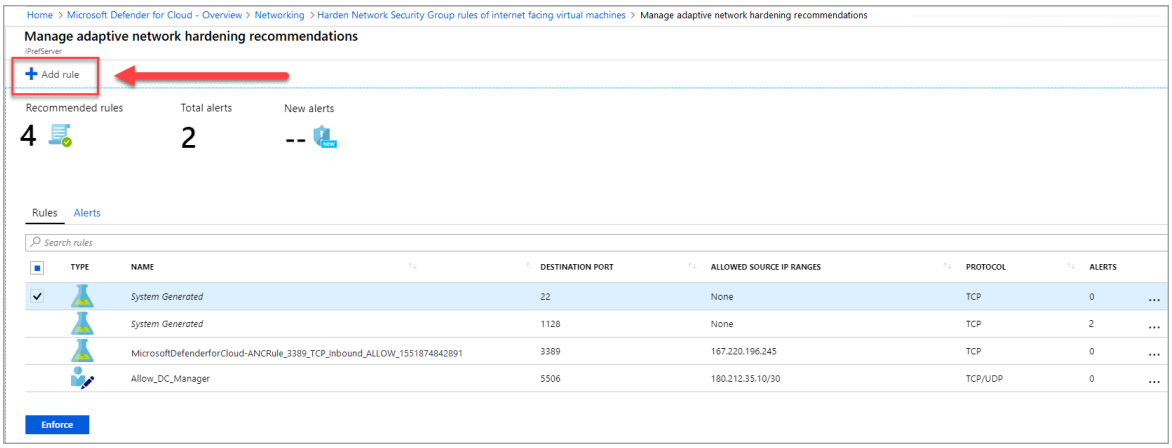

2. In the **New rule** window, enter the details and select **Add**.

> **NOTE**
>
> After selecting **Add**, you have successfully added the rule, and it is listed with the other recommended rules. However, you have not *applied* it on the NSG. To activate it, you must select the rule in the list, and select **Enforce** (as explained in the next step).

3. To apply the new rule, from the list, select the new rule and select **Enforce**.

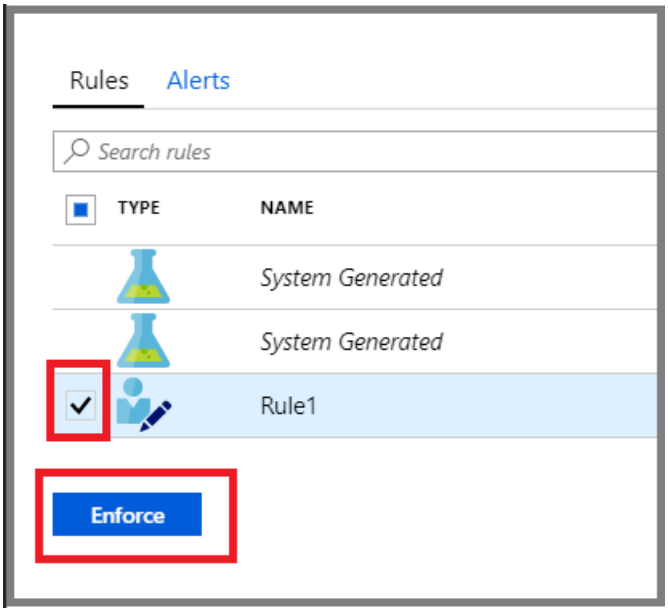

Delete a rule

When necessary, you can delete a recommended rule for the current session. For example, you may determine that applying a suggested rule could block legitimate traffic.

To delete an adaptive network hardening rule for your current session:

- In the **Rules** tab, select the three dots (...) at the end of the rule's row, and select **Delete**.

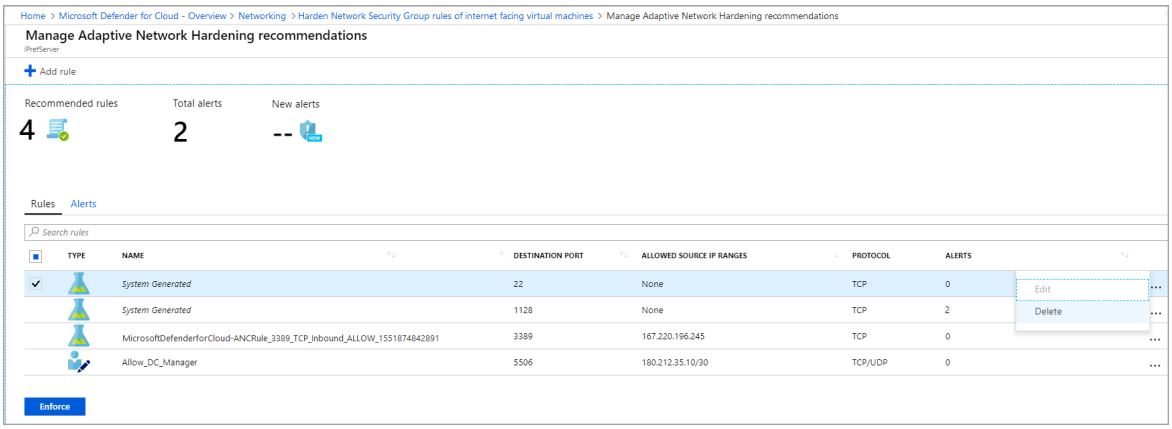

FAQ - Adaptive network hardening

- Which ports are supported?
- Are there any prerequisites or VM extensions required for adaptive network hardening?

Which ports are supported?

Adaptive network hardening recommendations are only supported on the following specific ports (for both UDP and TCP):

13, 17, 19, 22, 23, 53, 69, 81, 111, 119, 123, 135, 137, 138, 139, 161, 162, 389, 445, 512, 514, 593, 636, 873, 1433, 1434, 1900, 2049, 2301, 2323, 2381, 3268, 3306, 3389, 4333, 5353, 5432, 5555, 5800, 5900, 5900, 5985, 5986, 6379, 6379, 7000, 7001, 7199, 8081, 8089, 8545, 9042, 9160, 9300, 11211, 16379, 26379, 27017, 37215

Are there any prerequisites or VM extensions required for adaptive network hardening?

Adaptive network hardening is an agentless feature of Microsoft Defender for Cloud - nothing needs to be installed on your machines to benefit from this network hardening tool.

When should I use a "Deny all traffic" rule?

A **Deny all traffic** rule is recommended when, as a result of running the algorithm, Defender for Cloud does not identify traffic that should be allowed, based on the existing NSG configuration. Therefore, the recommended rule is to deny all traffic to the specified port. The name of this type of rule is displayed as "*System Generated*". After enforcing this rule, its actual name in the NSG will be a string comprised of the protocol, traffic direction, "DENY", and a random number.

Harden your Docker hosts

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud identifies unmanaged containers hosted on IaaS Linux VMs, or other Linux machines running Docker containers. Defender for Cloud continuously assesses the configurations of these containers. It then compares them with the Center for Internet Security (CIS) Docker Benchmark.

Defender for Cloud includes the entire ruleset of the CIS Docker Benchmark and alerts you if your containers don't satisfy any of the controls. When it finds misconfigurations, Defender for Cloud generates security recommendations. Use Defender for Cloud's **recommendations page** to view recommendations and remediate issues.

When vulnerabilities are found, they're grouped inside a single recommendation.

> **NOTE**
> These CIS benchmark checks will not run on AKS-managed instances or Databricks-managed VMs.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Requires Microsoft Defender for servers
Required roles and permissions:	**Reader** on the workspace to which the host connects
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) ✅ Connected AWS accounts

Identify and remediate security vulnerabilities in your Docker configuration

1. From Defender for Cloud's menu, open the **Recommendations** page.
2. Filter to the recommendation **Vulnerabilities in container security configurations should be remediated** and select the recommendation.

   The recommendation page shows the affected resources (Docker hosts).

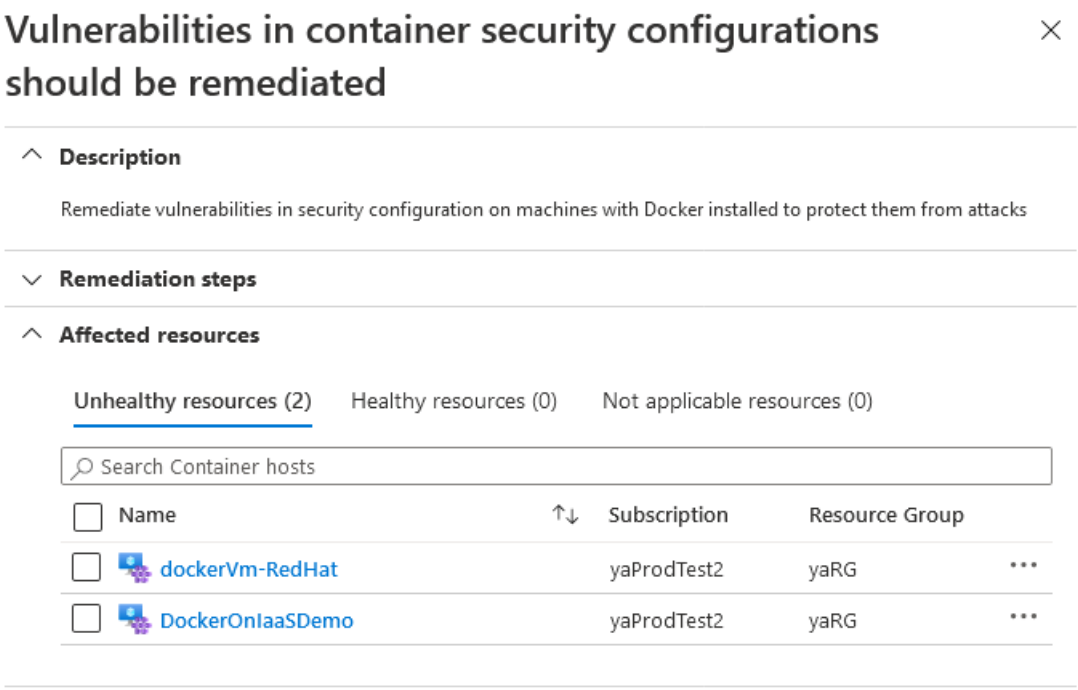

NOTE

Machines that aren't running Docker will be shown in the **Not applicable resources** tab. They'll appear in Azure Policy as Compliant.

3. To view and remediate the CIS controls that a specific host failed, select the host you want to investigate.

TIP

If you started at the asset inventory page and reached this recommendation from there, select the **Take action** button on the recommendation page.

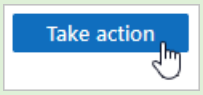

Log Analytics opens with a custom operation ready to run. The default custom query includes a list of all failed rules that were assessed, along with guidelines to help you resolve the issues.

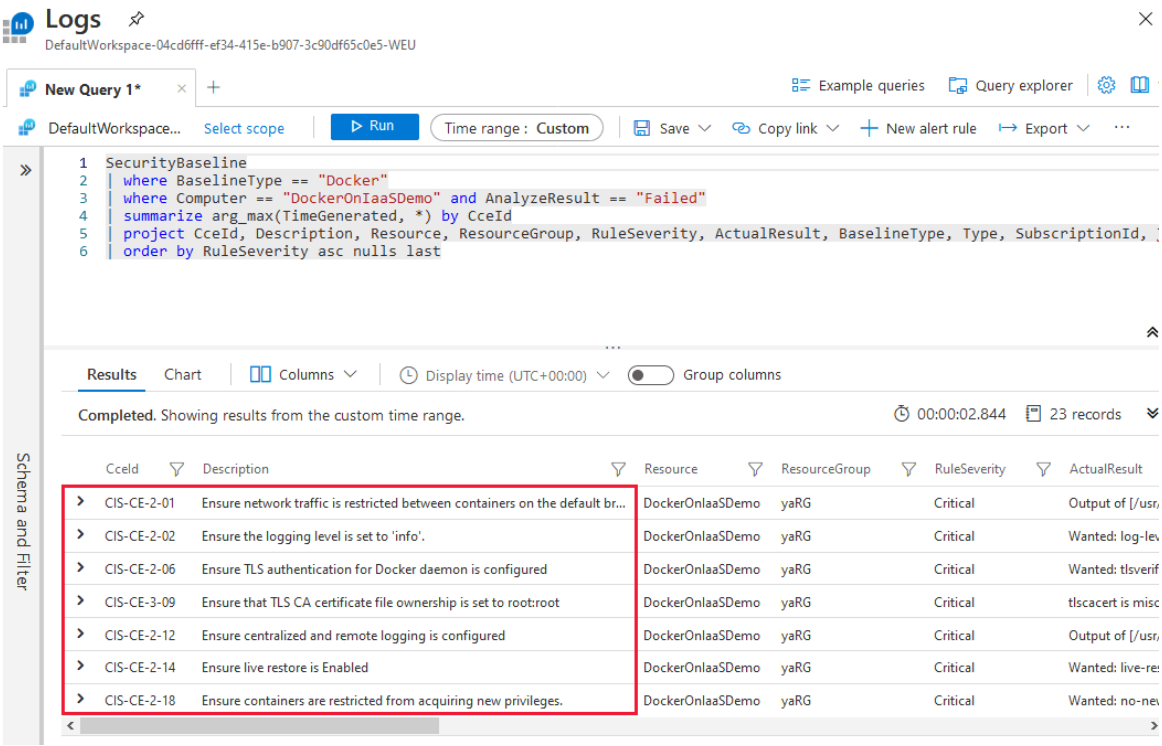

4. Tweak the query parameters if necessary.

5. When you're sure the command is appropriate and ready for your host, select **Run**.

Next steps

Docker hardening is just one aspect of Defender for Cloud's container security features.

Learn more Container security in Defender for Cloud.

Introduction to Microsoft Defender for SQL

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for SQL includes two Microsoft Defender plans that extend Microsoft Defender for Cloud's data security package to secure your databases and their data wherever they're located. Microsoft Defender for SQL includes functionalities for discovering and mitigating potential database vulnerabilities, and detecting anomalous activities that could indicate a threat to your databases.

Availability

ASPECT	DETAILS
Release state:	**Microsoft Defender for Azure SQL database servers** - Generally available (GA) **Microsoft Defender for SQL servers on machines** - Generally available (GA)
Pricing:	The two plans that form **Microsoft Defender for SQL** are billed as shown on the pricing page
Protected SQL versions:	SQL on Azure virtual machines SQL Server on Azure Arc-enabled servers On-premises SQL servers on Windows machines without Azure Arc Azure SQL single databases and elastic pools Azure SQL Managed Instance Azure Synapse Analytics (formerly SQL DW) dedicated SQL pool
Clouds:	✅ Commercial clouds ✅ Azure Government ✅ Azure China 21Vianet (**Partial**: Subset of alerts and vulnerability assessment for SQL servers. Behavioral threat protections aren't available.)

What does Microsoft Defender for SQL protect?

Microsoft Defender for SQL comprises two separate Microsoft Defender plans:

- **Microsoft Defender for Azure SQL database servers** protects:
  - Azure SQL Database
  - Azure SQL Managed Instance
  - Dedicated SQL pool in Azure Synapse

- **Microsoft Defender for SQL servers on machines** extends the protections for your Azure-native SQL Servers to fully support hybrid environments and protect SQL servers (all supported version) hosted in Azure, other cloud environments, and even on-premises machines:
  - SQL Server on Virtual Machines
  - On-premises SQL servers:
    - Azure Arc-enabled SQL Server (preview)
    - SQL Server running on Windows machines without Azure Arc

When you enable either of these plans, all supported resources that exist within the subscription are protected. Future resources created on the same subscription will also be protected.

What are the benefits of Microsoft Defender for SQL?

These two plans include functionality for identifying and mitigating potential database vulnerabilities and detecting anomalous activities that could indicate threats to your databases.

A vulnerability assessment service discovers, tracks, and helps you remediate potential database vulnerabilities. Assessment scans provide an overview of your SQL machines' security state, and details of any security findings.

- Learn more about vulnerability assessment for Azure SQL Database.
- Learn more about vulnerability assessment for Azure SQL servers on machines.

An advanced threat protection service continuously monitors your SQL servers for threats such as SQL injection, brute-force attacks, and privilege abuse. This service provides action-oriented security alerts in Microsoft Defender for Cloud with details of the suspicious activity, guidance on how to mitigate to the threats, and options for continuing your investigations with Microsoft Sentinel. Learn more about advanced threat protection.

> **TIP**
> View the list of security alerts for SQL servers in the alerts reference page.

Is there a performance impact from deploying Microsoft Defender for SQL on machines?

The focus of **Microsoft Defender for SQL on machines** is obviously security. But we also care about your business and so we've prioritized performance to ensure the minimal impact on your SQL servers.

The service has a split architecture to balance data uploading and speed with performance:

- some of our detectors run on the machine for real-time speed advantages
- others run in the cloud to spare the machine from heavy computational loads

Lab tests of our solution, comparing it against benchmark loads, showed CPU usage averaging 3% for peak slices. An analysis of the telemetry for our current users shows a negligible impact on CPU and memory usage.

Of course, performance always varies between environments, machines, and loads. The statements and numbers above are provided as a general guideline, not a guarantee for any individual deployment.

What kind of alerts does Microsoft Defender for SQL provide?

Threat intelligence enriched security alerts are triggered when there's:

- **Potential SQL injection attacks** - including vulnerabilities detected when applications generate a faulty SQL statement in the database
- **Anomalous database access and query patterns** - for example, an abnormally high number of failed sign-in attempts with different credentials (a brute force attempt)
- **Suspicious database activity** - for example, a legitimate user accessing an SQL Server from a breached computer which communicated with a crypto-mining C&C server

Alerts include details of the incident that triggered them, as well as recommendations on how to investigate and remediate threats.

Next steps

In this article, you learned about Microsoft Defender for SQL. To use the services that have been described:

- Use Microsoft Defender for SQL servers on machines to scan your SQL servers for vulnerabilities
- For a presentation of Microsoft Defender for SQL, see how Microsoft Defender for SQL can protect SQL servers anywhere

Enable Microsoft Defender for SQL servers on machines

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This Microsoft Defender plan detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases.

You'll see alerts when there are suspicious database activities, potential vulnerabilities, or SQL injection attacks, and anomalous database access and query patterns.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for SQL servers on machines** is billed as shown on the pricing page
Protected SQL versions:	SQL Server (versions currently supported by Microsoft)
Clouds:	✅ Commercial clouds ✅ Azure Government ❌ Azure China 21Vianet

Set up Microsoft Defender for SQL servers on machines

To enable this plan:

Step 1. Install the agent extension

Step 2. Provision the Log Analytics agent on your SQL server's host:

Step 3. Enable the optional plan in Defender for Cloud's environment settings page:

Step 1. Install the agent extension

- **SQL Server on Azure VM** - Register your SQL Server VM with the SQL IaaS Agent extension as explained in Register SQL Server VM with SQL IaaS Agent Extension.
- **SQL Server on Azure Arc-enabled servers** - Install the Azure Arc agent by following the installation methods described in the Azure Arc documentation.

Step 2. Provision the Log Analytics agent on your SQL server's host:

- **SQL Server on Azure VM** - If your SQL machine is hosted on an Azure VM, you can enable auto provisioning of the Log Analytics agent
  . Alternatively, you can follow the manual procedure for Onboard your Azure Stack Hub VMs.

- **SQL Server on Azure Arc-enabled servers** - If your SQL Server is managed by Azure Arc enabled servers, you can deploy the Log Analytics agent using the Defender for Cloud recommendation "Log Analytics agent should be installed on your Windows-based Azure Arc machines (Preview)".

- **SQL Server on-prem** - If your SQL Server is hosted on an on-premises Windows machine without Azure Arc, you have two options for connecting it to Azure:

  - **Deploy Azure Arc** - You can connect any Windows machine to Defender for Cloud. However, Azure Arc provides deeper integration across *all* of your Azure environment. If you set up Azure Arc, you'll see the **SQL Server – Azure Arc** page in the portal and your security alerts will appear on a dedicated **Security** tab on that page. So the first and recommended option is to set up Azure Arc on the host and follow the instructions for **SQL Server on Azure Arc**, above.

  - **Connect the Windows machine without Azure Arc** - If you choose to connect a SQL Server running on a Windows machine without using Azure Arc, follow the instructions in Connect Windows machines to Azure Monitor.

Step 3. Enable the optional plan in Defender for Cloud's environment settings page:

1. From Defender for Cloud's menu, open the **Environment settings** page.

   - If you're using **Microsoft Defender for Cloud's default workspace** (named "defaultworkspace-[your subscription ID]-[region]"), select the relevant **subscription**.

   - If you're using **a non-default workspace**, select the relevant **workspace** (enter the workspace's name in the filter if necessary).

2. Set the option for **Microsoft Defender for SQL servers on machines** plan to **on**.

   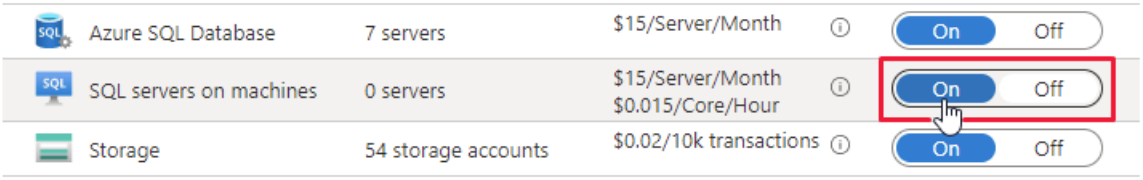

   The plan will be enabled on all SQL servers connected to the selected workspace. The protection will be fully active after the first restart of the SQL Server instance.

   > **TIP**
   > To create a new workspace, follow the instructions in Create a Log Analytics workspace.

3. Optionally, configure email notification for security alerts.

   You can set a list of recipients to receive an email notification when Defender for Cloud alerts are generated. The email contains a direct link to the alert in Microsoft Defender for Cloud with all the relevant details. For more information, see Set up email notifications for security alerts.

Microsoft Defender for SQL alerts

Alerts are generated by unusual and potentially harmful attempts to access or exploit SQL machines. These events can trigger alerts shown in the alerts reference page.

Explore and investigate security alerts

Microsoft Defender for SQL alerts are available in Defender for Cloud's alerts page, the machine's security page, the workload protections dashboard, or through the direct link in the alert emails.

1. To view alerts, select **Security alerts** from Defender for Cloud's menu and select an alert.
2. Alerts are designed to be self-contained, with detailed remediation steps and investigation information in each one. You can investigate further by using other Microsoft Defender for Cloud and Microsoft Sentinel capabilities for a broader view:
   - Enable SQL Server's auditing feature for further investigations. If you're a Microsoft Sentinel user, you can upload the SQL auditing logs from the Windows Security Log events to Sentinel and enjoy a rich investigation experience. Learn more about SQL Server Auditing.
   - To improve your security posture, use Defender for Cloud's recommendations for the host machine indicated in each alert. This will reduce the risks of future attacks.

   Learn more about managing and responding to alerts.

FAQ - Microsoft Defender for SQL servers on machines

If I enable this Microsoft Defender plan on my subscription, are all SQL servers on the subscription protected?

No. To defend a SQL Server deployment on an Azure virtual machine, or a SQL Server running on an Azure Arc-enabled machine, Defender for Cloud requires the following:

- a Log Analytics agent on the machine
- the relevant Log Analytics workspace to have the Microsoft Defender for SQL solution enabled

The subscription *status*, shown in the SQL server page in the Azure portal, reflects the default workspace status and applies to all connected machines. Only the SQL servers on hosts with a Log Analytics agent reporting to that workspace are protected by Defender for Cloud.

Next steps

For related material, see the following article:

- Security alerts for SQL Database and Azure Synapse Analytics
- Set up email notifications for security alerts
- Learn more about Microsoft Sentinel

Introduction to Microsoft Defender for open-source relational databases

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This plan brings threat protections for the following open-source relational databases:

- Azure Database for PostgreSQL
- Azure Database for MySQL
- Azure Database for MariaDB

Defender for Cloud detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. The plan makes it simple to address potential threats to databases without the need to be a security expert or manage advanced security monitoring systems.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for open-source relational databases** is billed as shown on the pricing page
Supported environments:	✔ PaaS ✘ Azure Arc-enabled machines
Protected versions of PostgreSQL:	Single Server - General Purpose and Memory Optimized. Learn more in PostgreSQL pricing tiers.
Protected versions of MySQL:	Single Server - General Purpose and Memory Optimized. Learn more in MySQL pricing tiers.
Protected versions of MariaDB:	General Purpose and Memory Optimized. Learn more in MariaDB pricing tiers.
Clouds:	✔ Commercial clouds ✔ Azure Government ✘ Azure China 21Vianet

What are the benefits of Microsoft Defender for open-source relational databases?

Defender for Cloud provides security alerts on anomalous activities so that you can detect potential threats and respond to them as they occur.

When you enable this plan, Defender for Cloud will provide alerts when it detects anomalous database access and query patterns as well as suspicious database activities.

These alerts appear in Defender for Cloud's security alerts page and include:

- details of the suspicious activity that triggered them
- the associated MITRE ATT&CK tactic
- recommended actions for how to investigate and mitigate the threat
- options for continuing your investigations with Microsoft Sentinel

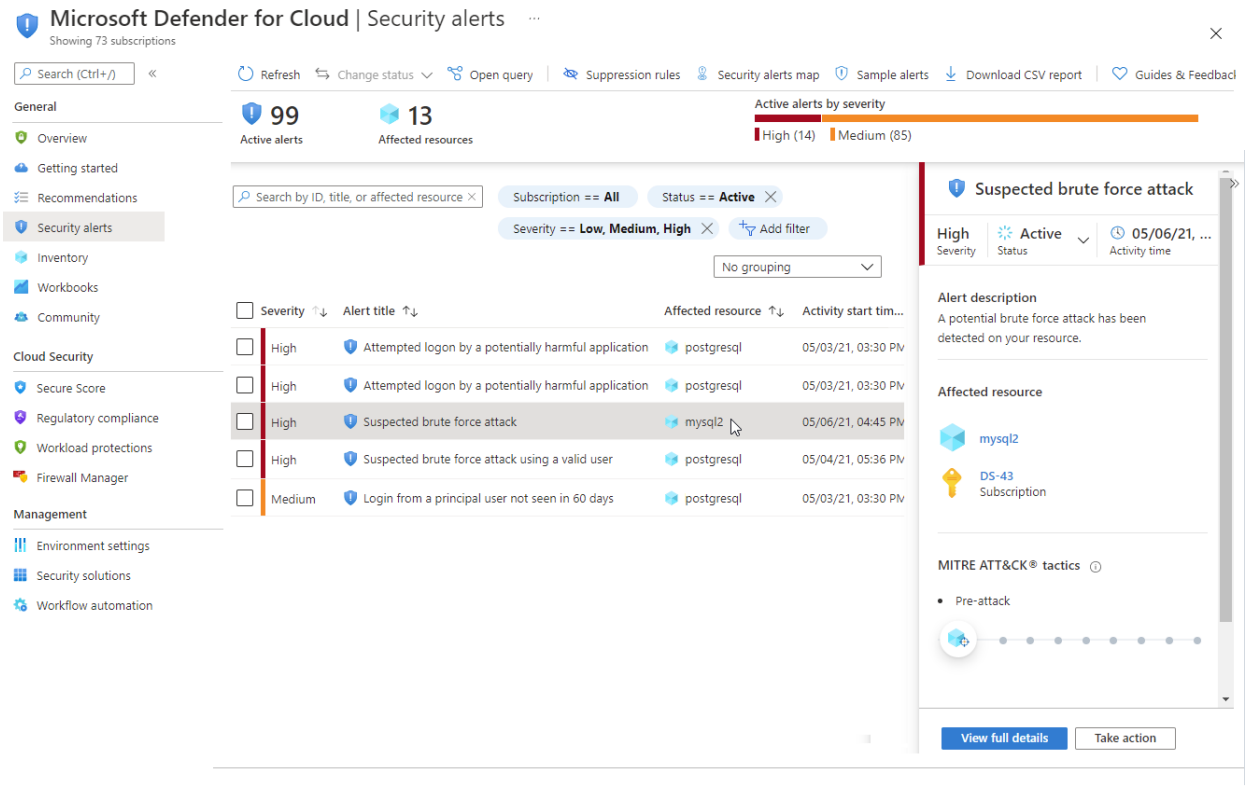

What kind of alerts does Microsoft Defender for open-source relational databases provide?

Threat intelligence enriched security alerts are triggered when there are:

- **Anomalous database access and query patterns** - For example, an abnormally high number of failed sign-in attempts with different credentials (a brute force attempt)
- **Suspicious database activities** - For example, a legitimate user accessing an SQL Server from a breached computer which communicated with a crypto-mining C&C server
- **Brute-force attacks** – With the ability to separate simple brute force from brute force on a valid user or a successful brute force

> **TIP**
> View the full list of security alerts for database servers in the alerts reference page.

Next steps

In this article, you learned about Microsoft Defender for open-source relational databases.

Enable enhanced protections

Enable Microsoft Defender for open-source relational databases and respond to alerts

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases for the following services:

- Azure Database for PostgreSQL
- Azure Database for MySQL
- Azure Database for MariaDB

To get alerts from the Microsoft Defender plan you'll first need to enable it as shown below.

Learn more about this Microsoft Defender plan in Introduction to Microsoft Defender for open-source relational databases.

Enable enhanced security

1. From the Azure portal, open the configuration page of the database server you want to protect.
2. From the security menu on the left, select **Microsoft Defender for Cloud**.
3. If enhanced security isn't enabled, you'll see a button as shown in the following screenshot. Select **Enable Microsoft Defender for [Database type]** (for example, "Microsoft Defender for MySQL") and select **Save**.

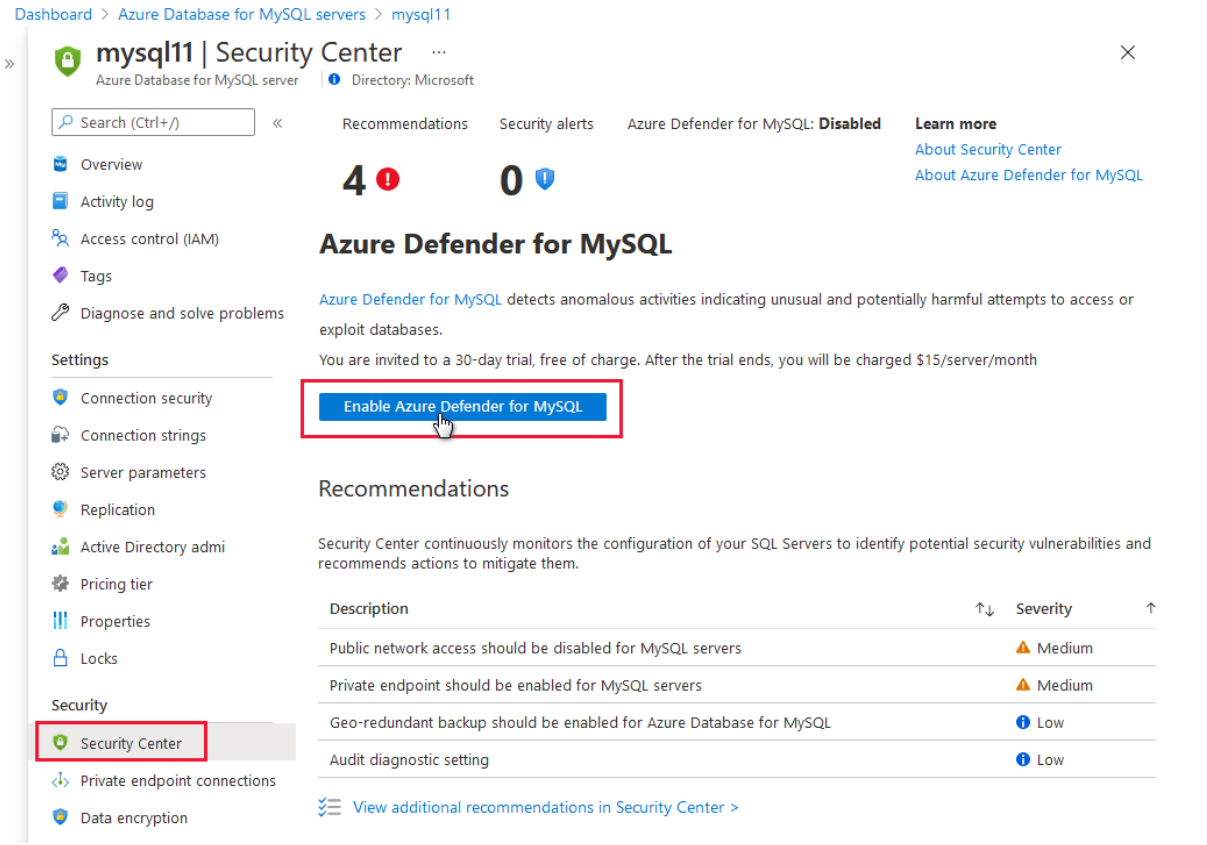

> **TIP**
> This page in the portal will be the same regardless of the database type (PostgreSQL, MySQL, or MariaDB).

Respond to security alerts

When Microsoft Defender for Cloud is enabled on your database, it detects anomalous activities and generates alerts. These alerts are available from multiple locations, including:

- In the Azure portal:
  - **Microsoft Defender for Cloud's security alerts page** - Shows alerts for all resources protected by Defender for Cloud in the subscriptions you've got permissions to view.
  - The resource's **Microsoft Defender for Cloud** page - Shows alerts and recommendations for one specific resource, as shown above in Enable enhanced security.
- In the inbox of whoever in your organization has been designated to receive email alerts.

> **TIP**
> A live tile on Microsoft Defender for Cloud's overview dashboard tracks the status of active threats to all your resources including databases. Select the tile to launch the Defender for Cloud alerts page and get an overview of active threats detected on your databases.
>
> For detailed steps and the recommended method to respond to security alerts, see Respond to a security alert.

Respond to email notifications of security alerts

Defender for Cloud sends email notifications when it detects anomalous database activities. The email includes details of the suspicious security event such as the nature of the anomalous activities, database name, server name, application name, and event time. The email also provides information on possible causes and recommended actions to investigate and mitigate any potential threats to the database.

1. From the email, select the **View the full alert** link to launch the Azure portal and show the alerts page, which provides an overview of active threats detected on the database.

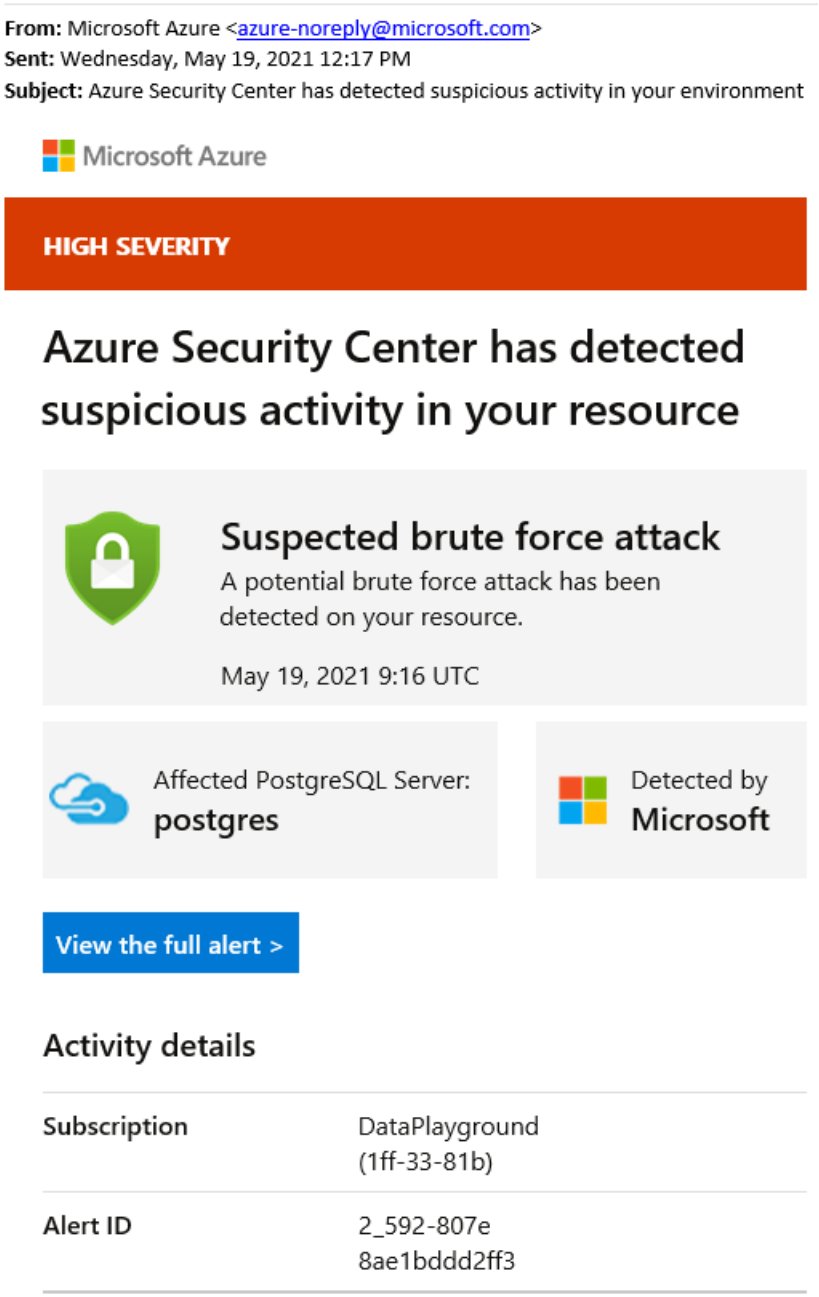

View active threats at the subscription level from within the Defender for Cloud portal pages:

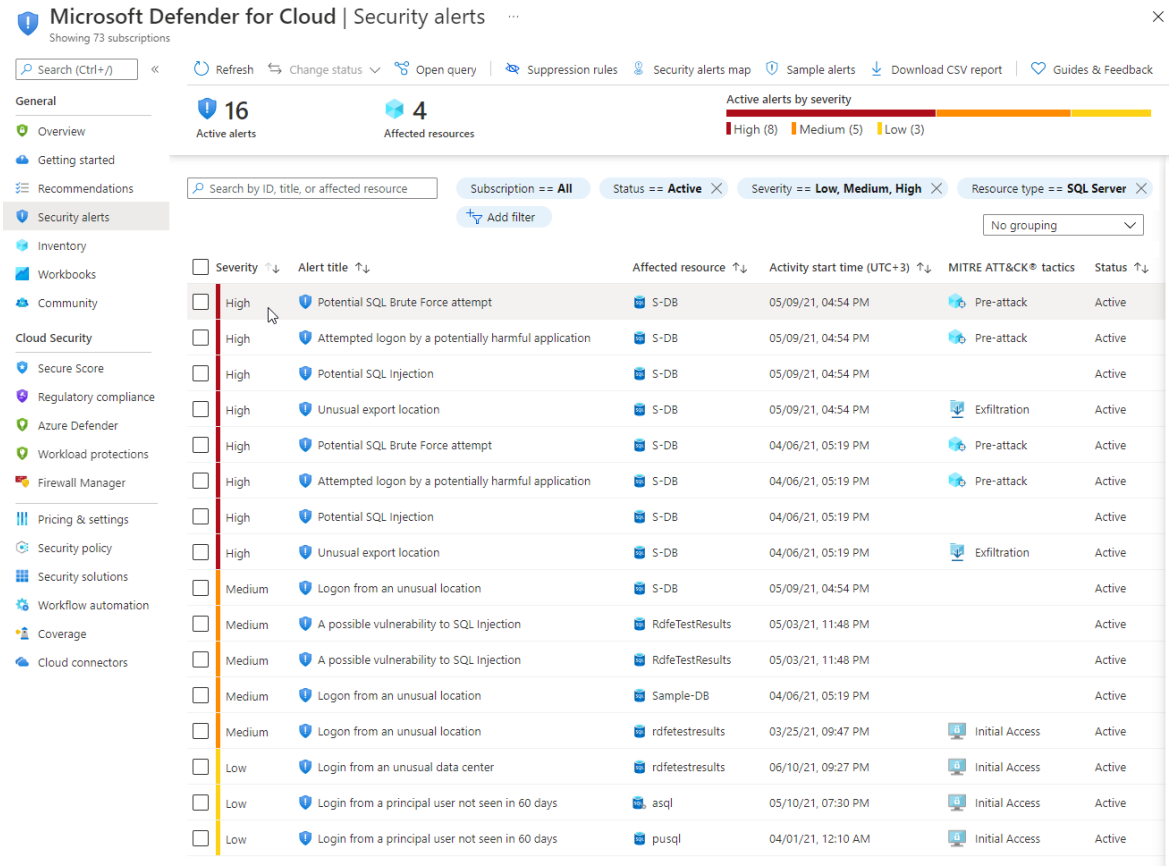

2. For additional details and recommended actions for investigating the current threat and remediating future threats, select a specific alert.

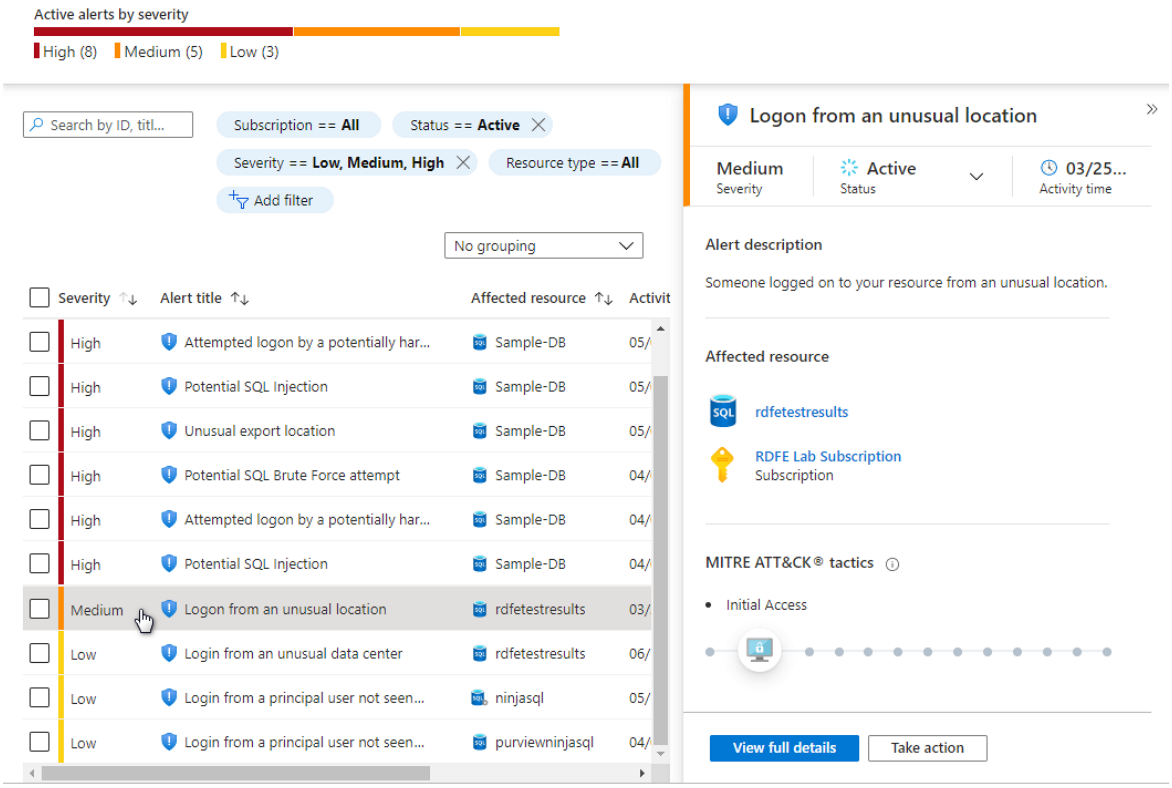

TIP

For a detailed tutorial on how to handle your alerts, see Tutorial: Triage, investigate, and respond to security alerts.

Next steps

- Automate responses to Defender for Cloud triggers
- Stream alerts to a SIEM, SOAR, or ITSM solution
- Suppress alerts from Defender for Cloud

Scan your SQL servers for vulnerabilities

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for SQL servers on machines extends the protections for your Azure-native SQL Servers to fully support hybrid environments and protect SQL servers (all supported version) hosted in Azure, other cloud environments, and even on-premises machines:

- SQL Server on Virtual Machines
- On-premises SQL servers:
  - SQL Server on Azure Arc-enabled servers
  - SQL Server running on Windows machines without Azure Arc

The integrated vulnerability assessment scanner discovers, tracks, and helps you remediate potential database vulnerabilities. Assessment scans findings provide an overview of your SQL machines' security state, and details of any security findings.

> **NOTE**
> The scan is lightweight, safe, only takes a few seconds per database to run and is entirely read-only. It does not make any changes to your database.

Explore vulnerability assessment reports

The vulnerability assessment service scans your databases every 12 hours.

The vulnerability assessment dashboard provides an overview of your assessment results across all your databases, along with a summary of healthy and unhealthy databases, and an overall summary of failing checks according to risk distribution.

You can view the vulnerability assessment results directly from Defender for Cloud.

1. From Defender for Cloud's sidebar, open the **Recommendations** page.
2. Select the recommendation SQL servers on machines should have vulnerability findings resolved. For more information, see the Defender for Cloud recommendations reference page.

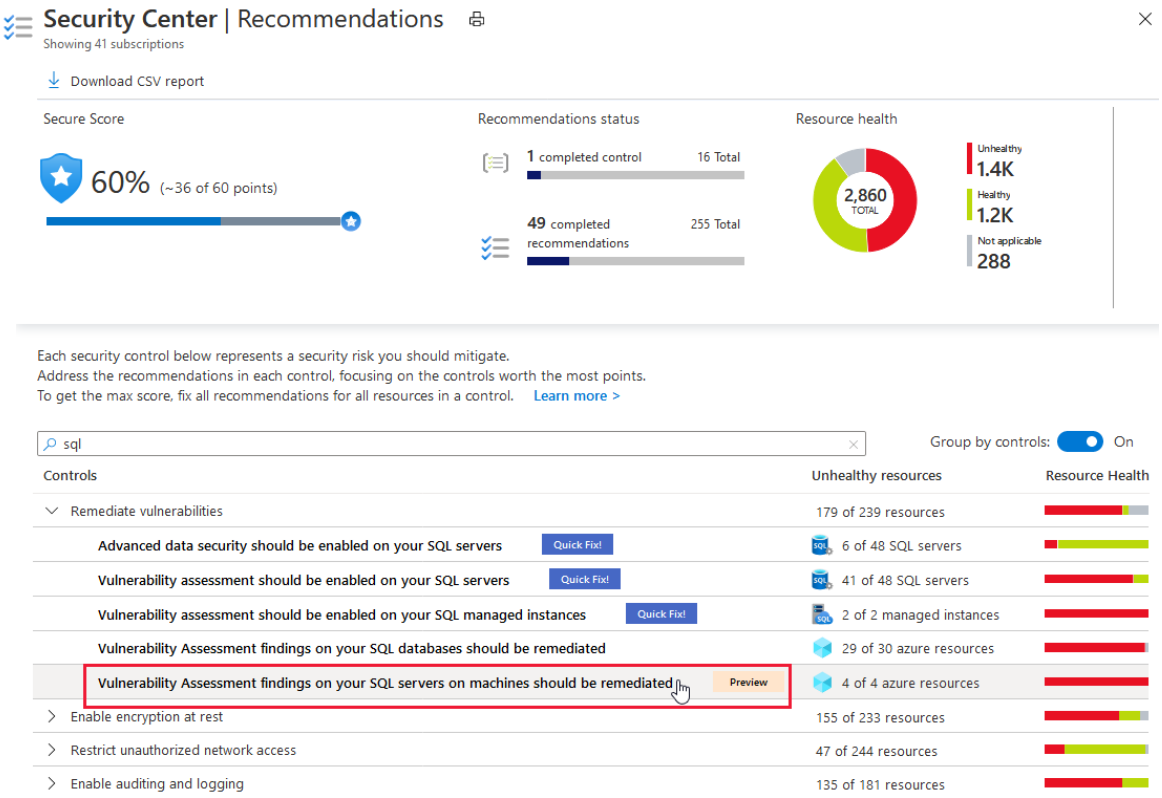

The detailed view for this recommendation appears.

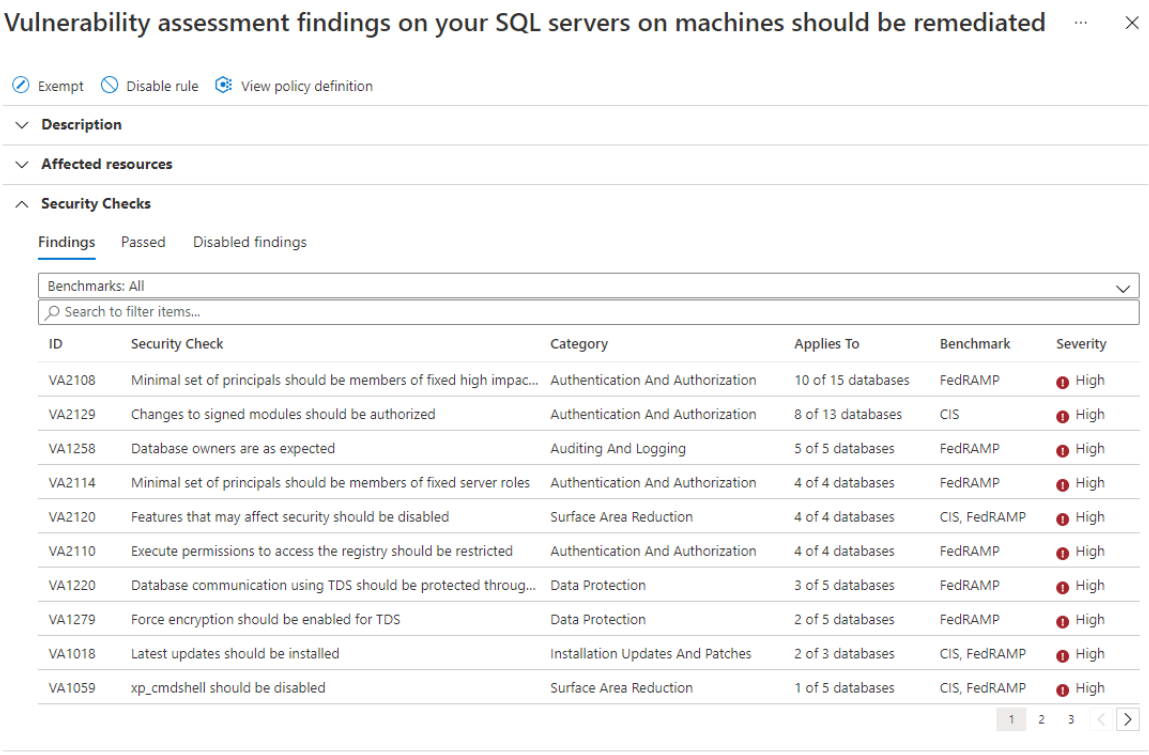

3. For more details, drill down:

   - For an overview of scanned resources (databases) and the list of security checks that were tested, open the **Affected resources** and select the server of interest.
   - For an overview of the vulnerabilities grouped by a specific SQL database, select the database of interest.

   In each view, the security checks are sorted by **Severity**. Select a specific security check to see a details pane with a **Description**, how to **Remediate** it, and other related information such as **Impact** or **Benchmark**.

Set a baseline

As you review your assessment results, you can mark results as being an acceptable baseline in your environment. The baseline is essentially a customization of how the results are reported. Results that match the baseline are considered as passing in subsequent scans. After you've established your baseline security state, the vulnerability assessment scanner only reports on deviations from the baseline. In this way, you can focus your attention on the relevant issues.

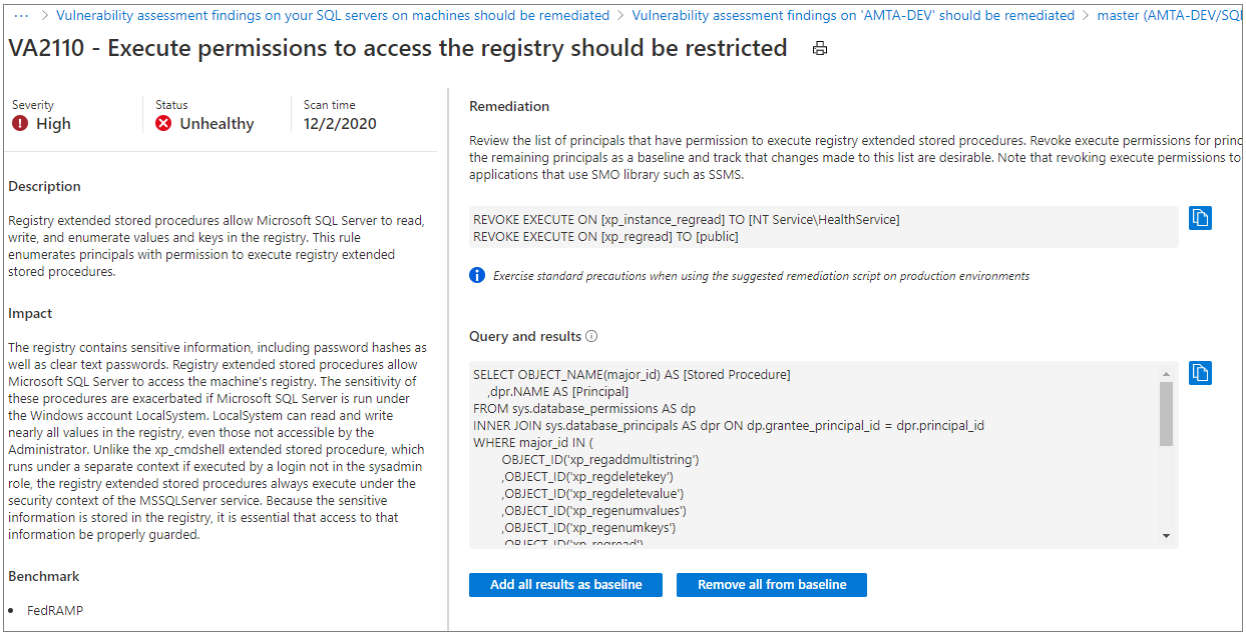

Export results

Use the Continuous export feature of Microsoft Defender for Cloud to export vulnerability assessment findings to Azure Event Hub or to Log Analytics workspace.

View vulnerabilities in graphical, interactive reports

Defender for Cloud's integrated Azure Monitor Workbooks gallery includes an interactive report of all findings from the vulnerability scanners for machines, containers in container registries, and SQL servers.

Findings for each of these scanners are reported in separate recommendations:

- Machines should have vulnerability findings resolved
- Container registry images should have vulnerability findings resolved (powered by Qualys)
- SQL databases should have vulnerability findings resolved
- SQL servers on machines should have vulnerability findings resolved

The 'Vulnerability Assessment Findings' report gathers all of these findings and organizes them by severity, resource type, and category. You can find the report in the workbooks gallery available from Defender for Cloud's sidebar.

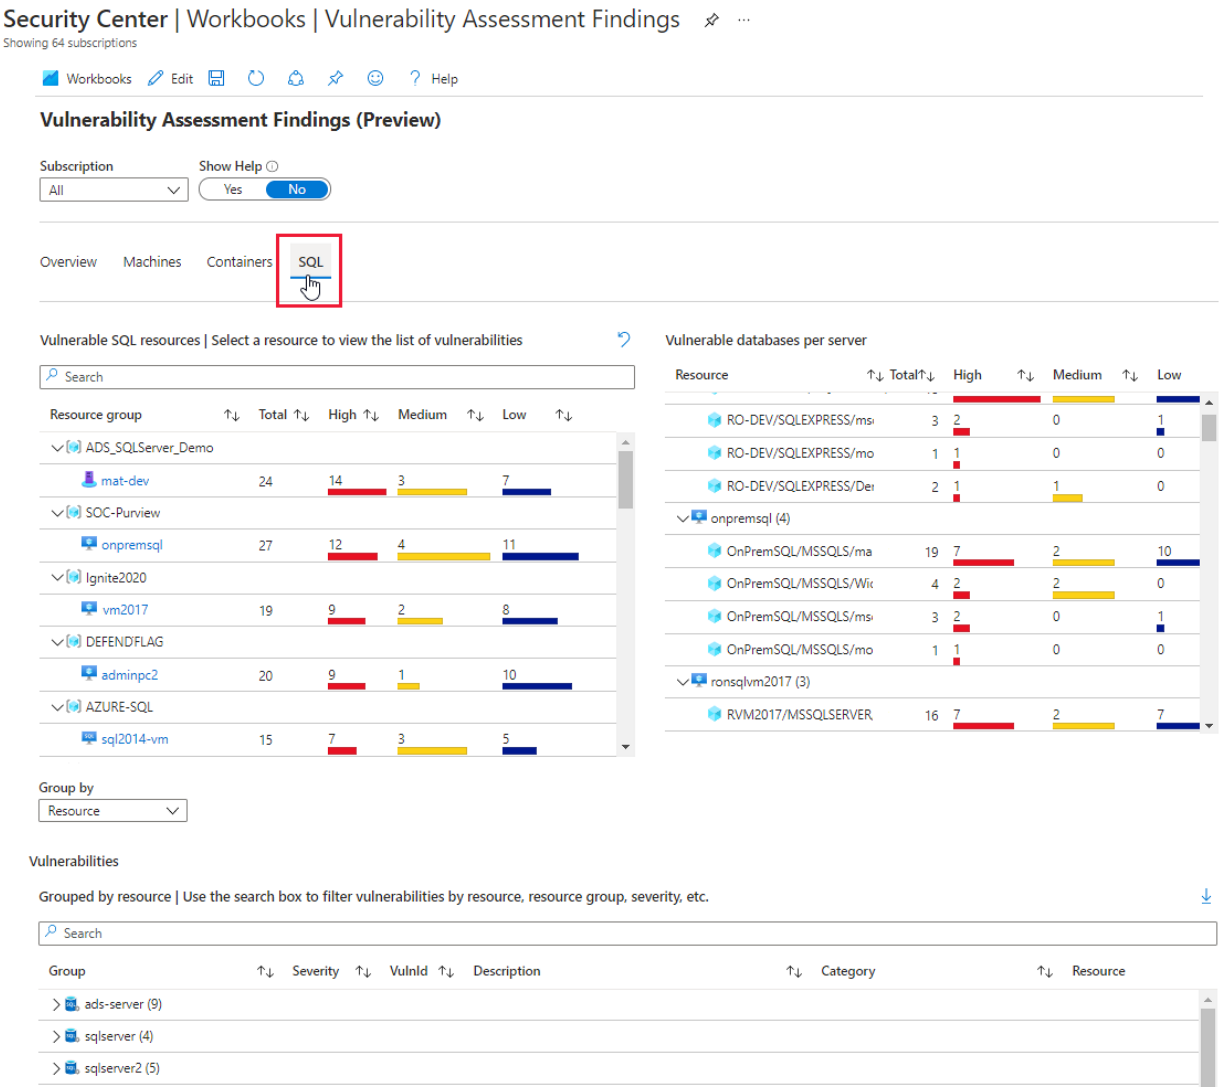

Disable specific findings

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

When a finding matches the criteria you've defined in your disable rules, it won't appear in the list of findings. Typical scenarios include:

- Disable findings with severity below medium
- Disable findings that are non-patchable
- Disable findings from benchmarks that aren't of interest for a defined scope

> **IMPORTANT**
>
> To disable specific findings, you need permissions to edit a policy in Azure Policy. Learn more in Azure RBAC permissions in Azure Policy.

To create a rule:

1. From the recommendations detail page for SQL servers on machines should have vulnerability findings resolved, select **Disable rule**.
2. Select the relevant scope.
3. Define your criteria. You can use any of the following criteria:

- Finding ID
- Severity
- Benchmarks

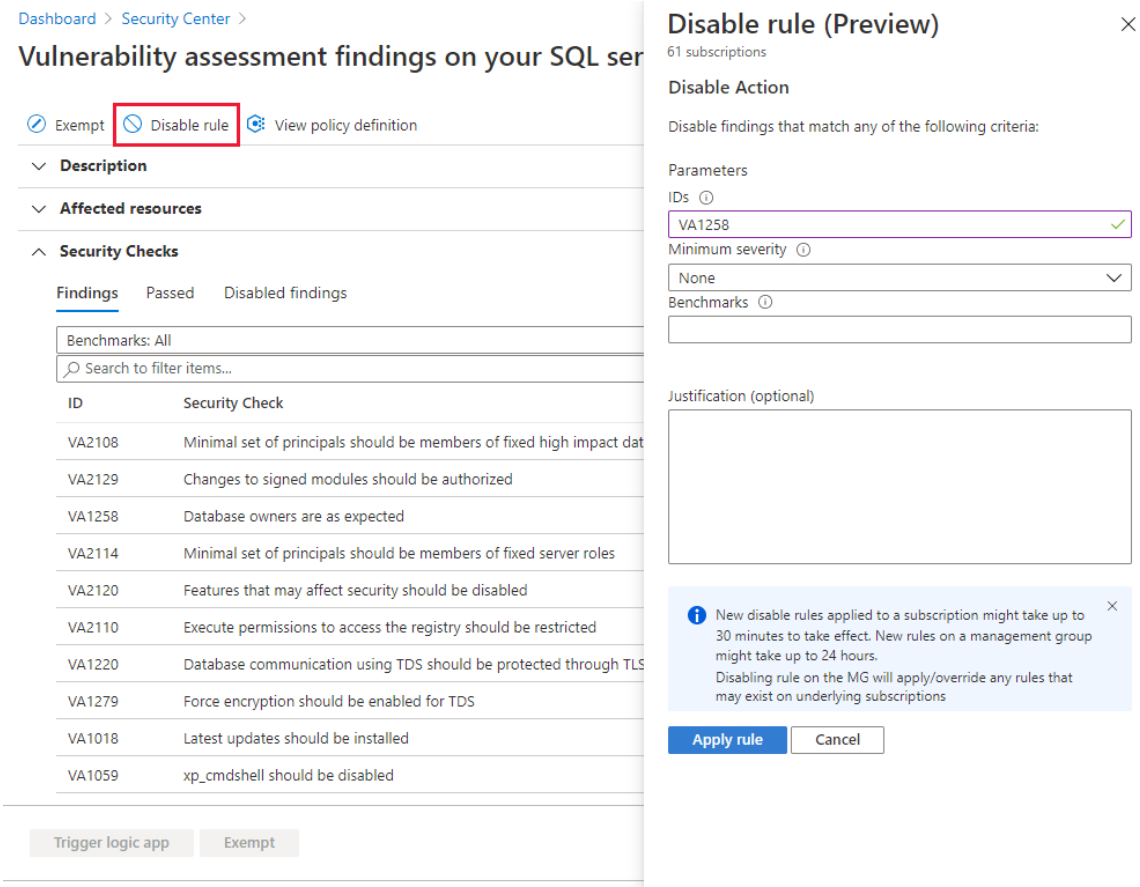

4. Select **Apply rule**. Changes might take up to 24 hrs to take effect.
5. To view, override, or delete a rule:
   a. Select **Disable rule**.
   b. From the scope list, subscriptions with active rules show as **Rule applied**.

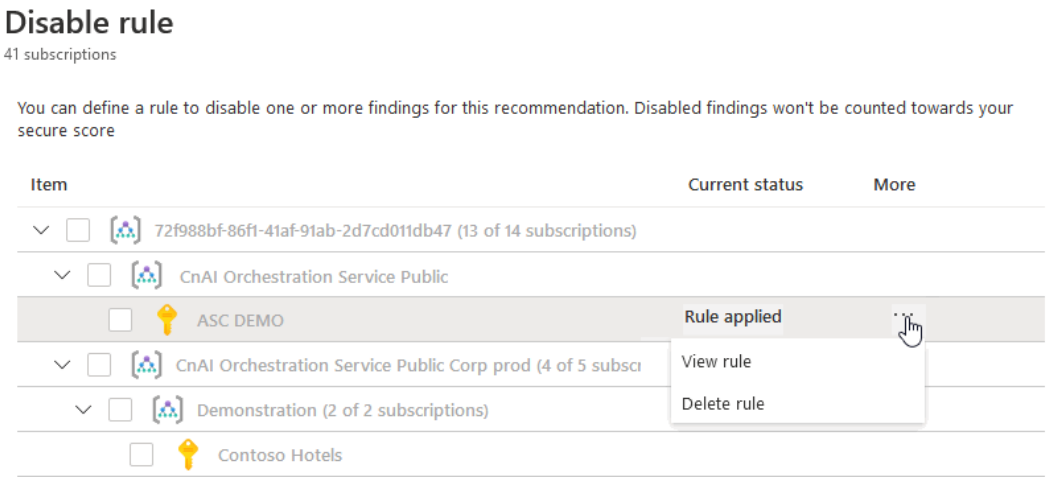

   c. To view or delete the rule, select the ellipsis menu ("...").

Manage vulnerability assessments programmatically

Using Azure PowerShell

You can use Azure PowerShell cmdlets to programmatically manage your vulnerability assessments. The

supported cmdlets are:

CMDLET NAME AS A LINK	DESCRIPTION
Add-AzSecuritySqlVulnerabilityAssessmentBaseline	Add SQL Vulnerability Assessment baseline.
Get-AzSecuritySqlVulnerabilityAssessmentBaseline	Get SQL Vulnerability Assessment baseline.
Get-AzSecuritySqlVulnerabilityAssessmentScanResult	Gets SQL Vulnerability Assessment scan results.
Get-AzSecuritySqlVulnerabilityAssessmentScanRecord	Gets SQL Vulnerability Assessment scan records.
Remove-AzSecuritySqlVulnerabilityAssessmentBaseline	Removes SQL Vulnerability Assessment baseline.
Set-AzSecuritySqlVulnerabilityAssessmentBaseline	Sets new SQL Vulnerability Assessment baseline on a specific database discards old baseline if any exists.

Data residency

SQL Vulnerability Assessment queries the SQL server using publicly available queries under Defender for Cloud recommendations for SQL Vulnerability Assessment, and stores the query results. SQL Vulnerability Assessment data is stored in the location of the Log Analytics workspace that the machine is connected to. For example, if the user connects a SQL Virtual Machine to a Log Analytics workspace in West Europe, the results will be stored in West Europe. This data will be collected only if the SQL Vulnerability Assessment solution is enabled on the Log Analytics workspace.

Metadata information about the connected machine is also collected. Specifically:

- Operating system name, type, and version
- Computer fully qualified domain name (FQDN)
- Connected Machine agent version
- UUID (BIOS ID)
- SQL server name and underlying database names

You can specify the region where your SQL Vulnerability Assessment data will be stored by choosing the Log Analytics workspace location. Microsoft may replicate to other regions for data resiliency, but Microsoft does not replicate data outside the geography.

Next steps

Learn more about Defender for Cloud's protections for SQL resources in Introduction to Microsoft Defender for SQL.

SQL information protection policy in Microsoft Defender for Cloud

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

SQL information protection's data discovery and classification mechanism provides advanced capabilities for discovering, classifying, labeling, and reporting the sensitive data in your databases. It's built into Azure SQL Database, Azure SQL Managed Instance, and Azure Synapse Analytics.

The classification mechanism is based on the following two elements:

- **Labels** – The main classification attributes, used to define the *sensitivity level of the data* stored in the column.
- **Information Types** – Provides additional granularity into the *type of data* stored in the column.

The information protection policy options within Defender for Cloud provide a predefined set of labels and information types which serve as the defaults for the classification engine. You can customize the policy, according to your organization's needs, as described below.

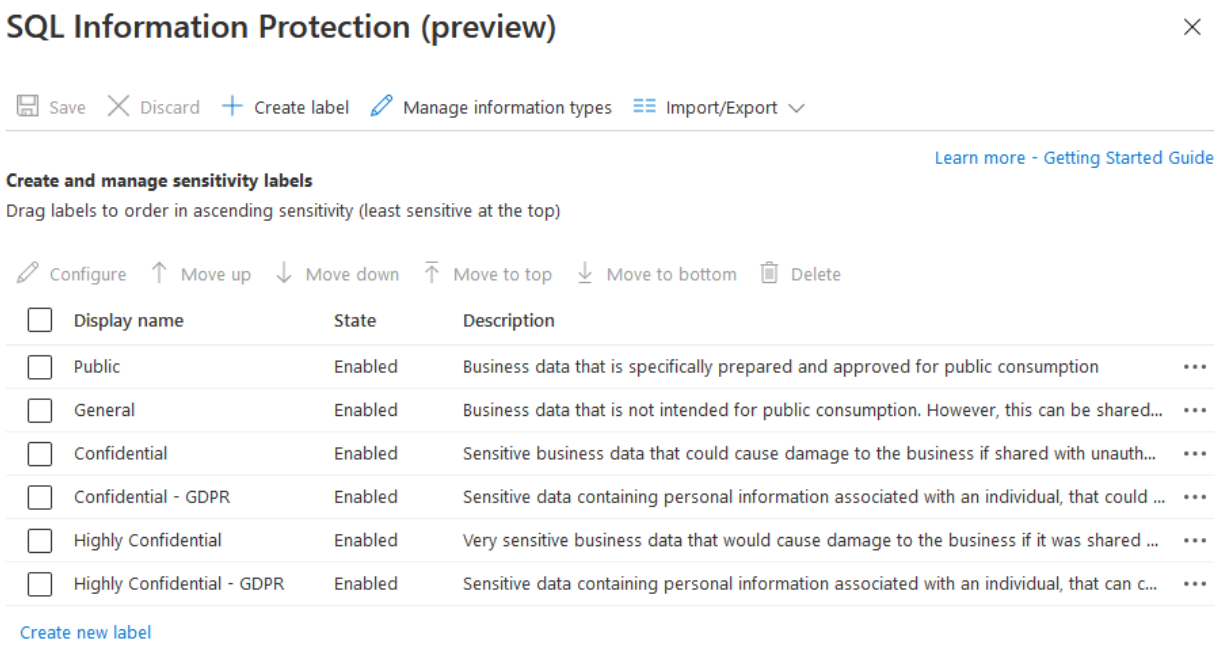

How do I access the SQL information protection policy?

There are three ways to access the information protection policy:

- **(Recommended)** From the **Environment settings** page of Defender for Cloud
- From the security recommendation "Sensitive data in your SQL databases should be classified"
- From the Azure SQL DB data discovery page

Each of these is shown in the relevant tab below.

- From Defender for Cloud's settings
- From Defender for Cloud's recommendation
- From Azure SQL

Access the policy from Defender for Cloud's environment settings page

From Defender for Cloud's **Environment settings** page, select **SQL information protection**.

> **NOTE**
> This option only appears for users with tenant-level permissions. Grant tenant-wide permissions to yourself.

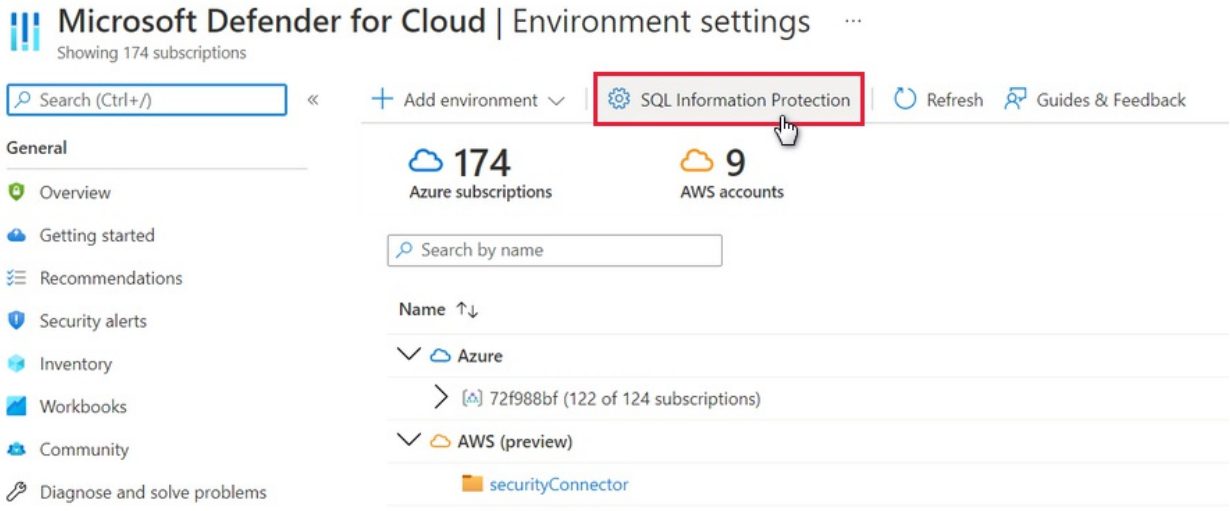

Customize your information types

To manage and customize information types:

1. Select **Manage information types**.

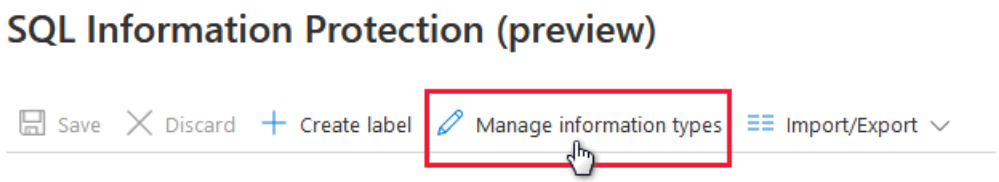

2. To add a new type, select **Create information type**. You can configure a name, description, and search pattern strings for the information type. Search pattern strings can optionally use keywords with wildcard characters (using the character '%'), which the automated discovery engine uses to identify sensitive data in your databases, based on the columns' metadata.

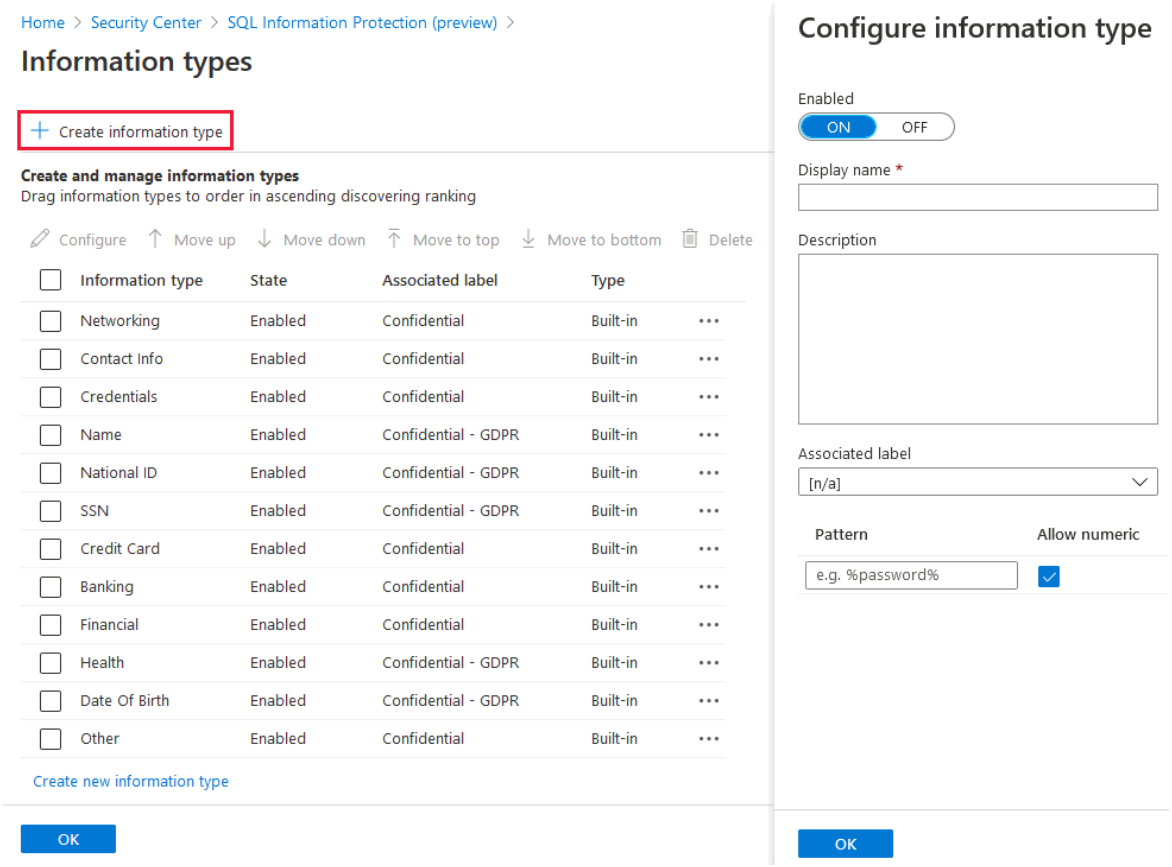

3. You can also modify the built-in types by adding additional search pattern strings, disabling some of the existing strings, or by changing the description.

   > **TIP**
   > You can't delete built-in types or change their names.

4. **Information types** are listed in order of ascending discovery ranking, meaning that the types higher in the list will attempt to match first. To change the ranking between information types, drag the types to the right spot in the table, or use the **Move up** and **Move down** buttons to change the order.

5. Select **OK** when you are done.

6. After you completed managing your information types, be sure to associate the relevant types with the relevant labels, by clicking **Configure** for a particular label, and adding or deleting information types as appropriate.

7. To apply your changes, select **Save** in the main **Labels** page.

Exporting and importing a policy

You can download a JSON file with your defined labels and information types, edit the file in the editor of your choice, and then import the updated file.

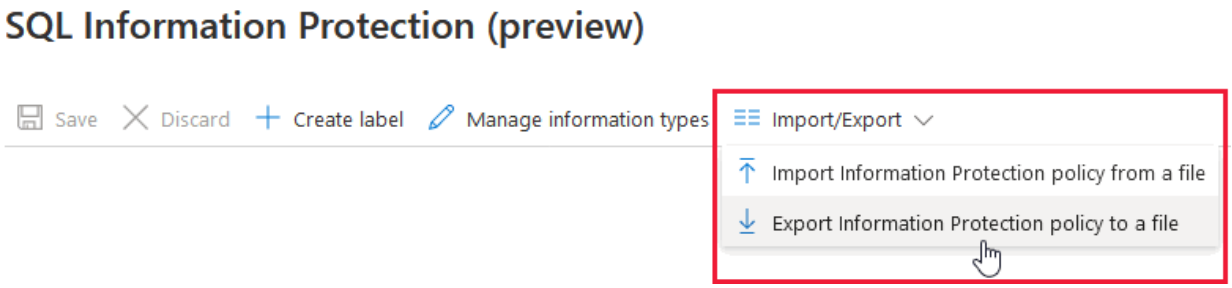

> **NOTE**
> You'll need tenant level permissions to import a policy file.

Permissions

To customize the information protection policy for your Azure tenant, you'll need the following actions on the tenant's root management group:

- Microsoft.Security/informationProtectionPolicies/read
- Microsoft.Security/informationProtectionPolicies/write

Learn more in Grant and request tenant-wide visibility.

Manage SQL information protection using Azure PowerShell

- Get-AzSqlInformationProtectionPolicy: Retrieves the effective tenant SQL information protection policy.
- Set-AzSqlInformationProtectionPolicy: Sets the effective tenant SQL information protection policy.

Next steps

In this article, you learned about defining an information protection policy in Microsoft Defender for Cloud. To learn more about using SQL Information Protection to classify and protect sensitive data in your SQL databases, see Azure SQL Database Data Discovery and Classification.

For more information on security policies and data security in Defender for Cloud, see the following articles:

- Setting security policies in Microsoft Defender for Cloud: Learn how to configure security policies for your Azure subscriptions and resource groups
- Microsoft Defender for Cloud data security: Learn how Defender for Cloud manages and safeguards data

Overview of Microsoft Defender for Containers

2/15/2022 • 12 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Containers is the cloud-native solution for securing your containers.

On this page, you'll learn how you can use Defender for Containers to improve, monitor, and maintain the security of your clusters, containers, and their applications.

Availability

ASPECT	DETAILS
Release state:	General availability (GA) Where indicated, specific features are in preview. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	**Microsoft Defender for Containers** is billed as shown on the pricing page
Registries and images:	**Supported** • Linux images in Azure Container Registry (ACR) registries accessible from the public internet with shell access • Private registries with access granted to Trusted Services • ACR registries protected with Azure Private Link **Unsupported** • Windows images • Super-minimalist images such as Docker scratch images • "Distroless" images that only contain an application and its runtime dependencies without a package manager, shell, or OS • Images with Open Container Initiative (OCI) Image Format Specification

ASPECT	DETAILS
Kubernetes distributions and configurations:	**Supported** • Any Cloud Native Computing Foundation (CNCF) certified Kubernetes clusters **Unsupported** • Any taints applied to your nodes *might* disrupt the configuration of Defender for Containers •The AKS Defender profile doesn't support AKS clusters that don't have RBAC role enabled. **Tested on** • Azure Kubernetes Service • Amazon Elastic Kubernetes Service (EKS) • Azure Kubernetes Service on Azure Stack HCI • Kubernetes • AKS Engine • Azure Red Hat OpenShift • Red Hat OpenShift (version 4.6 or newer) • VMware Tanzu Kubernetes Grid • Rancher Kubernetes Engine
Required roles and permissions:	• To auto provision the required components, Contributor, Log Analytics Contributor, or Azure Kubernetes Service Contributor Role • **Security admin** can dismiss alerts • **Security reader** can view vulnerability assessment findings See also Azure Container Registry roles and permissions
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) (Except for preview features) ✅ Connected AWS accounts (Preview)

What are the benefits of Microsoft Defender for Containers?

Defender for Containers helps with the core aspects of container security:

- **Environment hardening** - Defender for Containers protects your Kubernetes clusters whether they're running on Azure Kubernetes Service, Kubernetes on-prem / IaaS, or Amazon EKS. By continuously assessing clusters, Defender for Containers provides visibility into misconfigurations and guidelines to help mitigate identified threats. Learn more in Environment hardening through security recommendations.
- **Vulnerability assessment** - Vulnerability assessment and management tools for images **stored** in ACR registries and **running** in Azure Kubernetes Service. Learn more in Vulnerability assessment.
- **Run-time threat protection for nodes and clusters** - Threat protection for clusters and Linux nodes generates security alerts for suspicious activities. Learn more in Run-time protection for Kubernetes nodes, clusters, and hosts.

Architecture overview

The architecture of the various elements involved in the full range of protections provided by Defender for Containers varies depending on where your Kubernetes clusters are hosted.

Defender for Containers protects your clusters whether they're running in:

- **Azure Kubernetes Service (AKS)** - Microsoft's managed service for developing, deploying, and managing containerized applications.
- **Amazon Elastic Kubernetes Service (EKS) in a connected Amazon Web Services (AWS) account** - Amazon's managed service for running Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes.
- **An unmanaged Kubernetes distribution** (using Azure Arc-enabled Kubernetes) - Cloud Native Computing Foundation (CNCF) certified Kubernetes clusters hosted on-premises or on IaaS.

> **NOTE**
> Defender for Containers' support for Arc-enabled Kubernetes clusters (and therefore AWS EKS too) is a preview feature.

For high-level diagrams of each scenario, see the relevant tabs below.

In the diagrams you'll see that the items received and analyzed by Defender for Cloud include:

- Audit logs and security events from the API server
- Cluster configuration information from the control plane
- Workload configuration from Azure Policy
- Security signals and events from the node level

- AKS cluster
- Azure Arc-enabled Kubernetes
- AWS EKS

Architecture diagram of Defender for Cloud and AKS clusters

When Defender for Cloud protects a cluster hosted in Azure Kubernetes Service, the collection of audit log data is agentless and frictionless.

The **Defender profile (preview)** deployed to each node provides the runtime protections and collects signals from nodes using eBPF technology.

The **Azure Policy add-on for Kubernetes** collects cluster and workload configuration for admission control policies as explained in Protect your Kubernetes workloads.

> **NOTE**
> Defender for Containers' **Defender profile** is a preview feature.

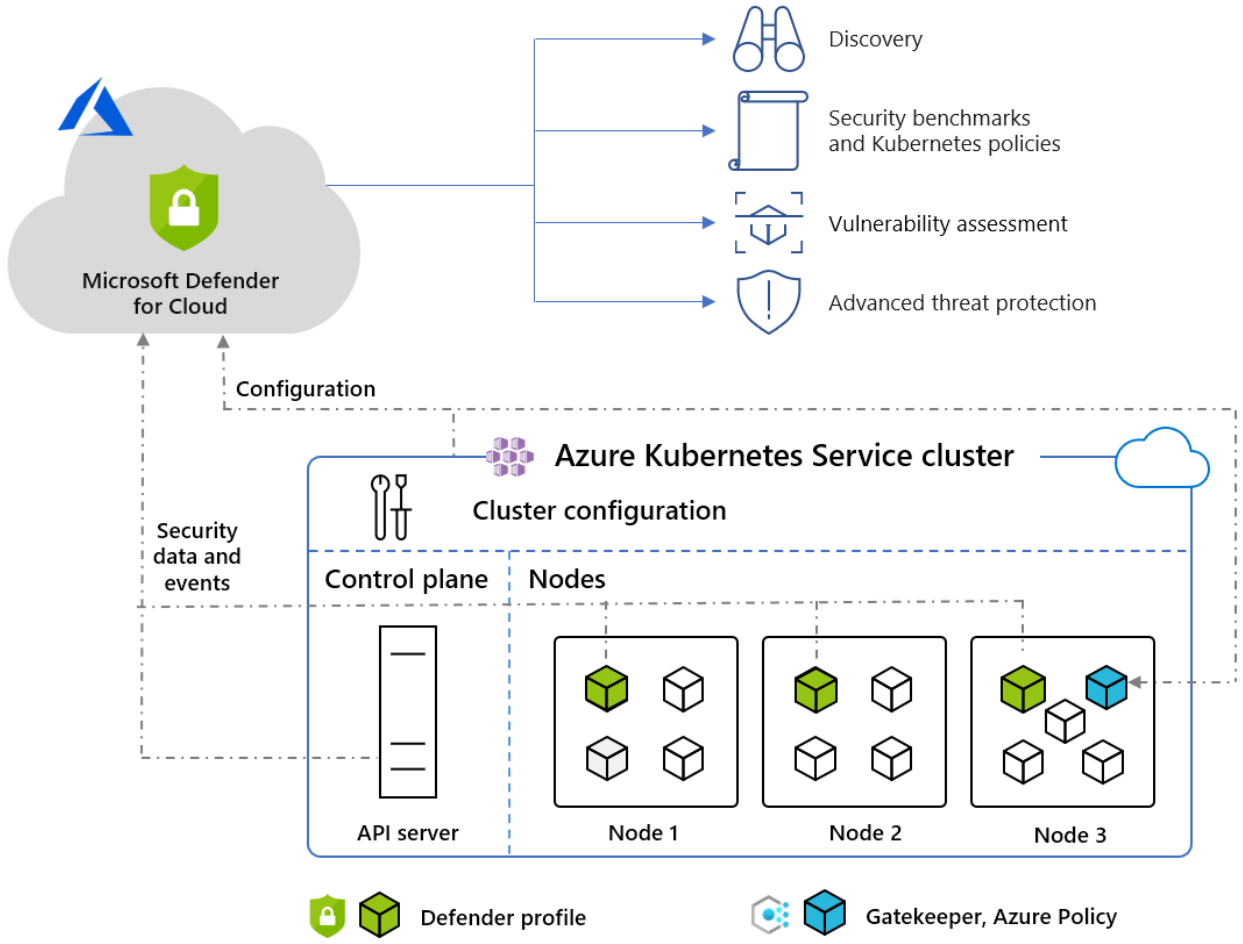

Defender profile component details

POD NAME	NAMESPACE	KIND	SHORT DESCRIPTION	CAPABILITIES	RESOURCE LIMITS	EGRESS REQUIRED
azuredefender-collector-ds-*	kube-system	DeamonSet	A set of containers that focus on collecting inventory and security events from the Kubernetes environment.	SYS_ADMIN, SYS_RESOURCE, SYS_PTRACE	memory: 64Mi cpu: 60m	No
azuredefender-collector-misc-*	kube-system	Deployment	A set of containers that focus on collecting inventory and security events from the Kubernetes environment that aren't bounded to a specific node.	N/A	memory: 64Mi cpu: 60m	No

POD NAME	NAMESPACE	KIND	SHORT DESCRIPTION	CAPABILITIES	RESOURCE LIMITS	EGRESS REQUIRED
azuredefender-publisher-ds-*	kube-system	DeamonSet	Publish the collected data to Microsoft Defender for Containers' backend service where the data will be processed for and analyzed.	N/A	memory: 64Mi cpu: 60m	Https 443 Learn more about the outbound access prerequisites

* resource limits are not configurable

Environment hardening through security recommendations

Continuous monitoring of your Kubernetes clusters - wherever they're hosted

Defender for Cloud continuously assesses the configurations of your clusters and compares them with the initiatives applied to your subscriptions. When it finds misconfigurations, Defender for Cloud generates security recommendations. Use Defender for Cloud's **recommendations page** to view recommendations and remediate issues. For details of the relevant Defender for Cloud recommendations that might appear for this feature, see the compute section of the recommendations reference table.

For Kubernetes clusters on EKS, you'll need to connect your AWS account to Microsoft Defender for Cloud via the environment settings page as described in Connect your AWS accounts to Microsoft Defender for Cloud. Then ensure you've enabled the CSPM plan.

When reviewing the outstanding recommendations for your container-related resources, whether in asset inventory or the recommendations page, you can use the resource filter:

Workload protection best-practices using Kubernetes admission control

For a bundle of recommendations to protect the workloads of your Kubernetes containers, install the **Azure Policy for Kubernetes**. You can also auto deploy this component as explained in enable auto provisioning of agents and extensions. By default, auto provisioning is enabled when you enable Defender for Containers.

With the add-on on your AKS cluster, every request to the Kubernetes API server will be monitored against the predefined set of best practices before being persisted to the cluster. You can then configure to **enforce** the best practices and mandate them for future workloads.

For example, you can mandate that privileged containers shouldn't be created, and any future requests to do so will be blocked.

Learn more in Protect your Kubernetes workloads.

Vulnerability assessment

Scanning images in ACR registries

Defender for Containers includes an integrated vulnerability scanner for scanning images in Azure Container Registry registries.

There are four triggers for an image scan:

- **On push** - Whenever an image is pushed to your registry, Defender for container registries automatically scans that image. To trigger the scan of an image, push it to your repository.

- **Recently pulled** - Since new vulnerabilities are discovered every day, **Microsoft Defender for Containers** also scans, on a weekly basis, any image that has been pulled within the last 30 days. There's no extra charge for these rescans; as mentioned above, you're billed once per image.
- **On import** - Azure Container Registry has import tools to bring images to your registry from Docker Hub, Microsoft Container Registry, or another Azure container registry. **Microsoft Defender for container Containers** scans any supported images you import. Learn more in Import container images to a container registry.
- **Continuous scan**- This trigger has two modes:
    - A Continuous scan based on an image pull. This scan is performed every 7 days after an image was pulled, and only for 30 days after the image was pulled. This mode doesn't require the security profile, or extension.
    - (Preview) Continuous scan for running images. This scan is performed every 7 days for as long as the image runs. This mode runs instead of the above mode when the Defender profile, or extension is running on the cluster.

This scan typically completes within 2 minutes, but it might take up to 40 minutes. For every vulnerability identified, Defender for Cloud provides actionable recommendations, along with a severity classification, and guidance for how to remediate the issue.

Defender for Cloud filters, and classifies findings from the scanner. When an image is healthy, Defender for Cloud marks it as such. Defender for Cloud generates security recommendations only for images that have issues to be resolved. By only notifying when there are problems, Defender for Cloud reduces the potential for unwanted informational alerts.

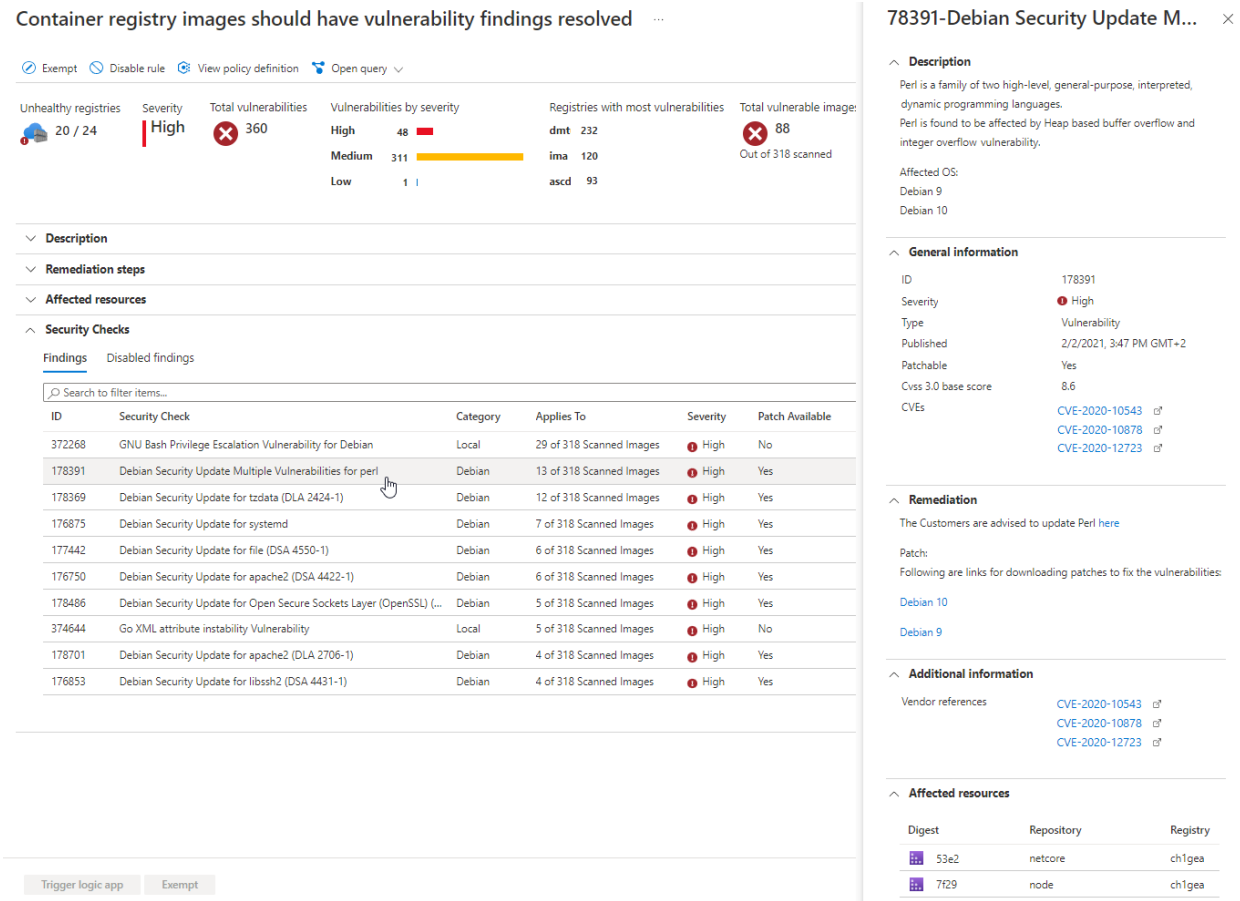

View vulnerabilities for running images

Defender for Containers expands on the registry scanning features of the Defender for container registries plan by introducing the **preview feature** of run-time visibility of vulnerabilities powered by the Defender profile, or extension.

The new recommendation, **Running container images should have vulnerability findings resolved**, only shows vulnerabilities for running images, and relies on the Defender security profile, or extension to discover which images are currently running. This recommendation groups running images that have vulnerabilities, and provides details about the issues discovered, and how to remediate them. The Defender profile, or extension is used to gain visibility into vulnerable containers that are active.

This recommendation shows running images, and their vulnerabilities based on ACR image image. Images that are deployed from a non ACR registry, will not be scanned, and will appear under the Not applicable tab.

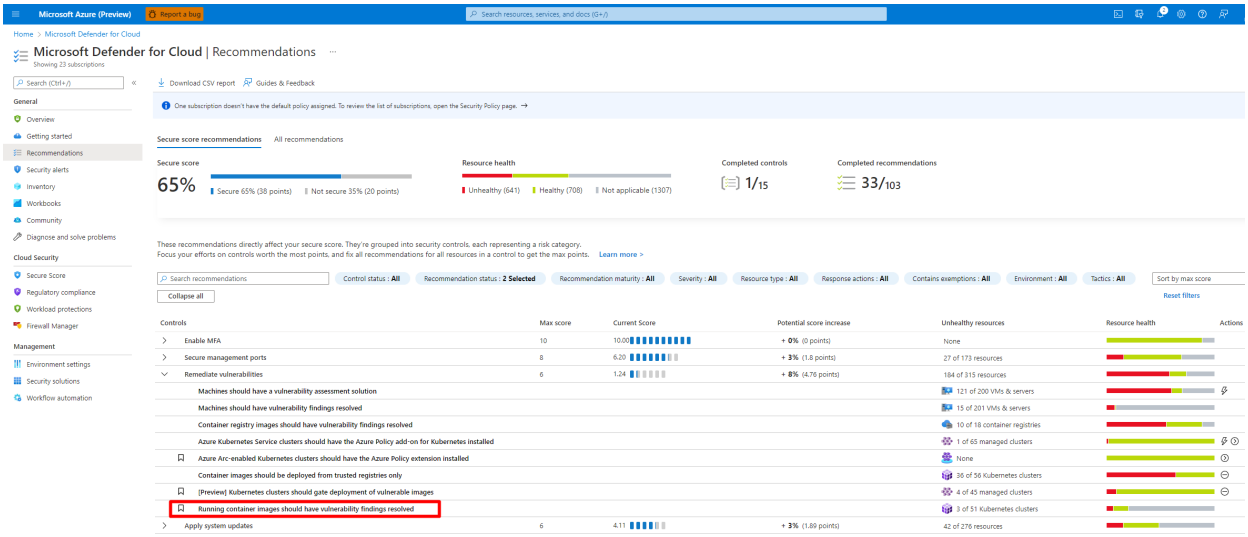

Run-time protection for Kubernetes nodes and clusters

Defender for Cloud provides real-time threat protection for your containerized environments and generates alerts for suspicious activities. You can use this information to quickly remediate security issues and improve the security of your containers.

Threat protection at the cluster level is provided by the Defender profile and analysis of the Kubernetes audit logs. Examples of events at this level include exposed Kubernetes dashboards, creation of high-privileged roles, and the creation of sensitive mounts.

In addition, our threat detection goes beyond the Kubernetes management layer. Defender for Containers includes **host-level threat detection** with over 60 Kubernetes-aware analytics, AI, and anomaly detections based on your runtime workload. Our global team of security researchers constantly monitor the threat landscape. They add container-specific alerts and vulnerabilities as they're discovered. Together, this solution monitors the growing attack surface of multi-cloud Kubernetes deployments and tracks the MITRE ATT&CK® matrix for Containers, a framework that was developed by the Center for Threat-Informed Defense in close partnership with Microsoft and others.

The full list of available alerts can be found in the Reference table of alerts.

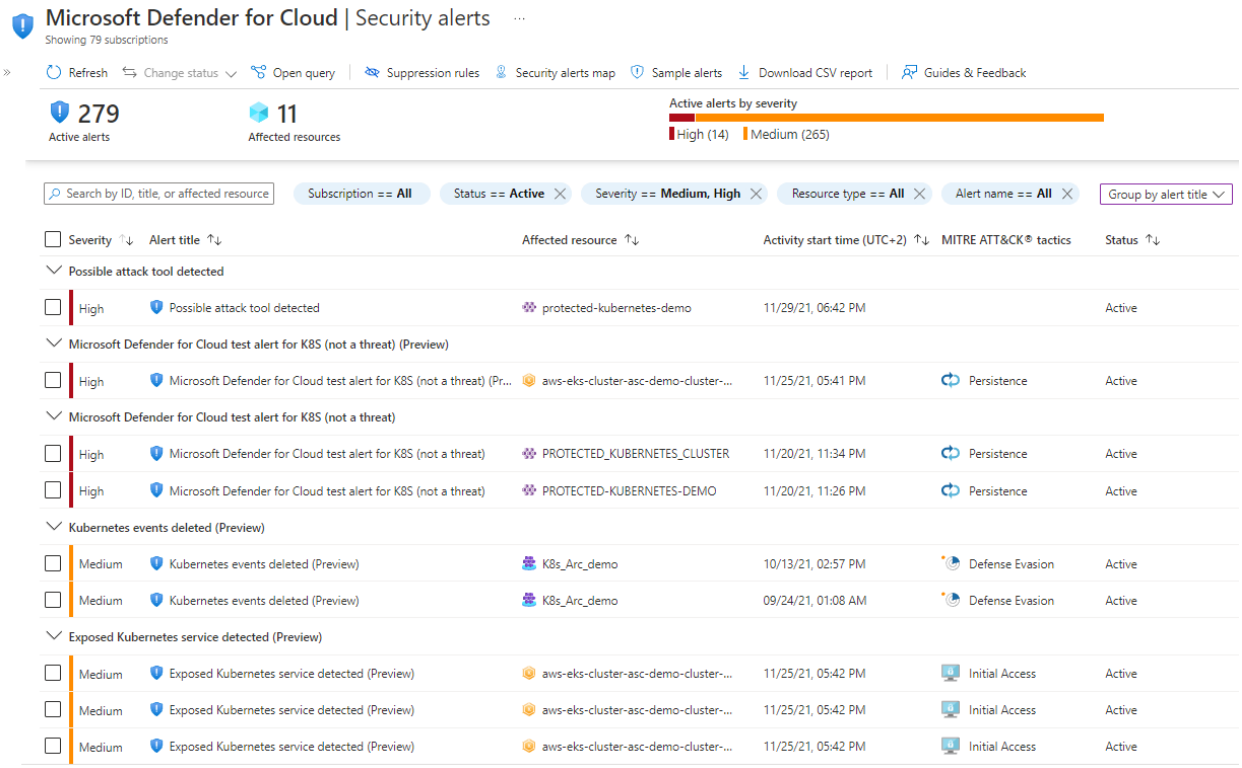

FAQ - Defender for Containers

- What happens to subscriptions with Microsoft Defender for Kubernetes or Microsoft Defender for container registries enabled?
- Is Defender for Containers a mandatory upgrade?
- Does the new plan reflect a price increase?
- What are the options to enable the new plan at scale?

What happens to subscriptions with Microsoft Defender for Kubernetes or Microsoft Defender for container registries enabled?

Subscriptions that already have one of these plans enabled can continue to benefit from it.

If you haven't enabled them yet, or create a new subscription, these plans can no longer be enabled.

Is Defender for Containers a mandatory upgrade?

No. Subscriptions that have either Microsoft Defender for Kubernetes or Microsoft Defender for container registries enabled don't need to be upgraded to the new Microsoft Defender for Containers plan. However, they won't benefit from the new and improved capabilities and they'll have an upgrade icon shown alongside them in the Azure portal.

Does the new plan reflect a price increase?

No. There's no direct price increase. The new comprehensive Container security plan combines Kubernetes protection and container registry image scanning, and removes the previous dependency on the (paid) Defender for Servers plan.

What are the options to enable the new plan at scale?

We've rolled out a new policy in Azure Policy, **Configure Microsoft Defender for Containers to be enabled**, to make it easier to enable the new plan at scale.

Next steps

In this overview, you learned about the core elements of container security in Microsoft Defender for Cloud. To enable the plan, see:

Enable Defender for Containers

Enable Microsoft Defender for Containers

2/15/2022 • 15 minutes to read • Edit Online

Microsoft Defender for Containers is the cloud-native solution for securing your containers.

Defender for Containers protects your clusters whether they're running in:

- **Azure Kubernetes Service (AKS)** - Microsoft's managed service for developing, deploying, and managing containerized applications.
- **Amazon Elastic Kubernetes Service (EKS) in a connected Amazon Web Services (AWS) account** - Amazon's managed service for running Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes.
- **An unmanaged Kubernetes distribution** (using Azure Arc-enabled Kubernetes) - Cloud Native Computing Foundation (CNCF) certified Kubernetes clusters hosted on-premises or on IaaS.

Learn about this plan in Overview of Microsoft Defender for Containers.

> **NOTE**
> Defender for Containers' support for Arc-enabled Kubernetes clusters (and therefore AWS EKS too) is a preview feature.
>
> The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

Prerequisites

Validate the following endpoints are configured for outbound access so that the Defender profile can connect to Microsoft Defender for Cloud to send security data and events:

See the required FQDN/application rules for Microsoft Defender for Containers.

By default, AKS clusters have unrestricted outbound (egress) internet access.

Prerequisites

Validate the following endpoints are configured for outbound access so that the Defender extension can connect to Microsoft Defender for Cloud to send security data and events:

For Azure public cloud deployments:

DOMAIN	PORT
*.ods.opinsights.azure.com	443
*.oms.opinsights.azure.com	443
login.microsoftonline.com	443

Enable the plan

1. From Defender for Cloud's menu, open the Environment settings page and select the relevant subscription.

2. In the Defender plans page, enable **Defender for Containers**

> **TIP**
>
> If the subscription already has Defender for Kubernetes and/or Defender for container registries enabled, an update notice is shown. Otherwise, the only option will be **Defender for Containers**.
>
> 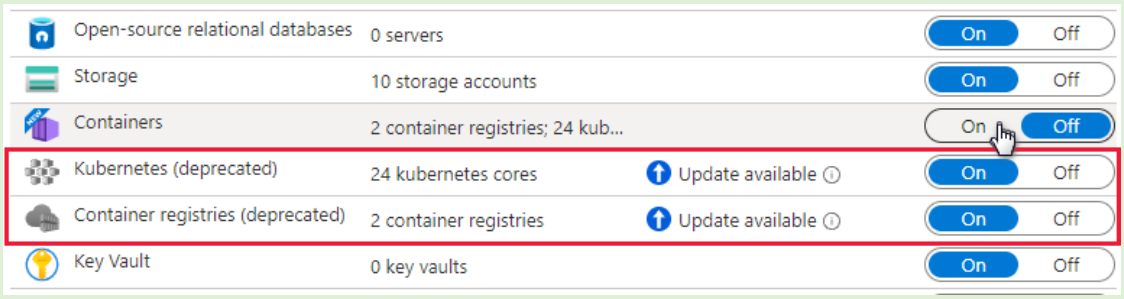
>

3. By default, the plan is configured to automatically defend any supported Kubernetes cluster that is attached to this subscription. To optionally modify the configuration, select *configure** from the configuration column.

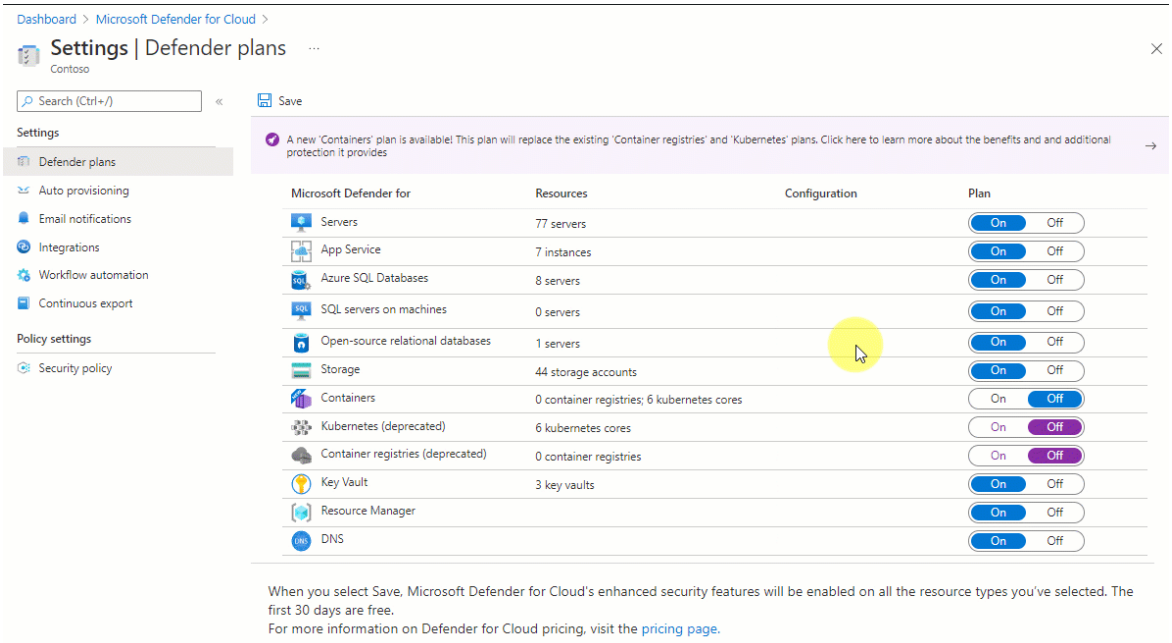

You can also modify this configuration from the Auto provisioning page on the **Microsoft Defender for Containers components (preview)** row:

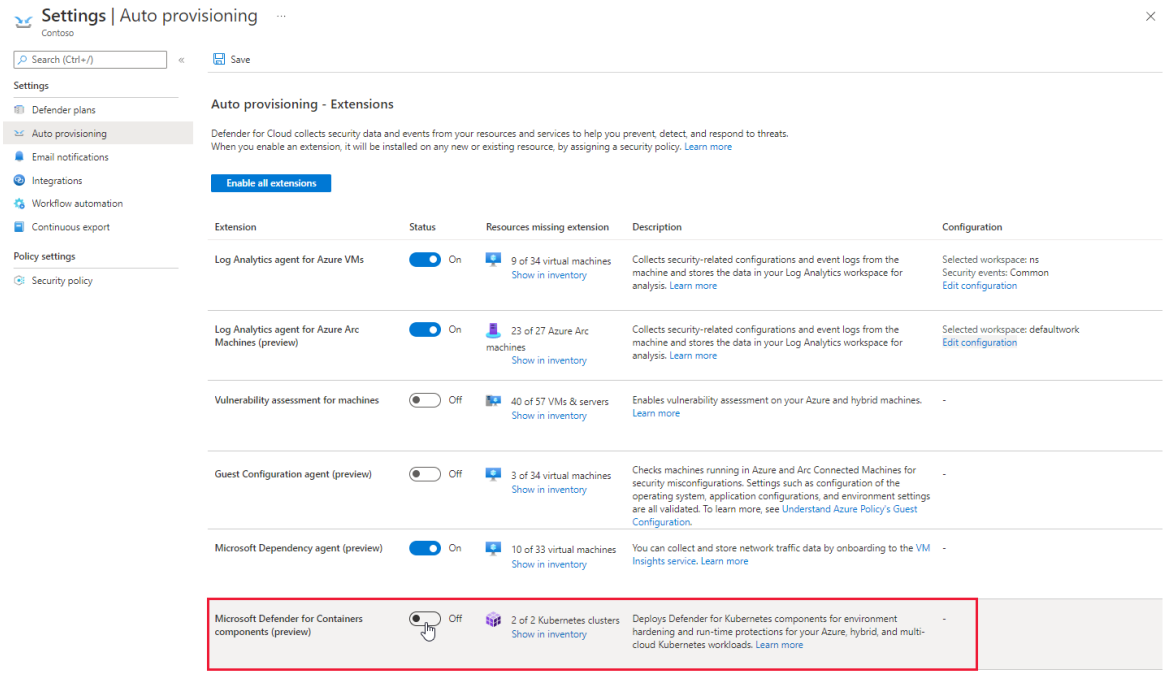

> **NOTE**
>
> If you choose to **disable the plan** at any time after enabling if through the portal as shown above, you'll need to manually disable auto provisioning of the Defender for Containers components. This will not remove the components from machines on which they've already been deployed.

4. If you disable the auto provisioning of any component, you can easily deploy the component to one or more clusters using the appropriate recommendation:

   - Policy Add-on for Kubernetes - Azure Kubernetes Service clusters should have the Azure Policy Add-on for Kubernetes installed
   - Azure Kubernetes Service profile - Azure Kubernetes Service clusters should have Defender profile enabled
   - Azure Arc-enabled Kubernetes extension - Azure Arc-enabled Kubernetes clusters should have the Defender extension installed

Deploy the Defender profile

You can enable the Defender for Containers plan and deploy all of the relevant components from the Azure portal, the REST API, or with a Resource Manager template. For detailed steps, select the relevant tab.

The Defender security profile is a preview feature. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

- **Azure portal**
- **REST API**
- **Resource Manager**

Use the fix button from the Defender for Cloud recommendation

A streamlined, frictionless, process lets you use the Azure portal pages to enable the Defender for Cloud plan and setup auto provisioning of all the necessary components for defending your Kubernetes clusters at scale.

A dedicated Defender for Cloud recommendation provides:

- **Visibility** about which of your clusters has the Defender profile deployed
- **Fix** button to deploy it to those clusters without the extension

1. From Microsoft Defender for Cloud's recommendations page, open the **Enable enhanced security** security control.
2. Use the filter to find the recommendation named **Azure Kubernetes Service clusters should have Defender profile enabled**.

   > **TIP**
   > Notice the **Fix** icon in the actions column

3. Select the clusters to see the details of the healthy and unhealthy resources - clusters with and without the profile.
4. From the unhealthy resources list, select a cluster and select **Remediate** to open the pane with the remediation confirmation.
5. Select **Fix *[x]* resources**.

Enable the plan

1. From Defender for Cloud's menu, open the Environment settings page and select the relevant subscription.
2. In the Defender plans page, enable **Defender for Containers**

   > **TIP**
   > If the subscription already has Defender for Kubernetes and/or Defender for container registries enabled, an update notice is shown. Otherwise, the only option will be **Defender for Containers**.

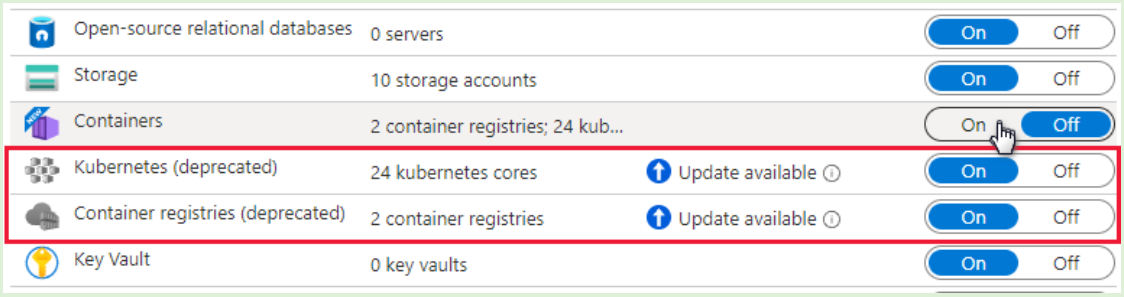

3. By default, the plan is configured to automatically defend any supported Kubernetes cluster that is attached to this subscription. To optionally modify the configuration, select *configure** from the configuration column.

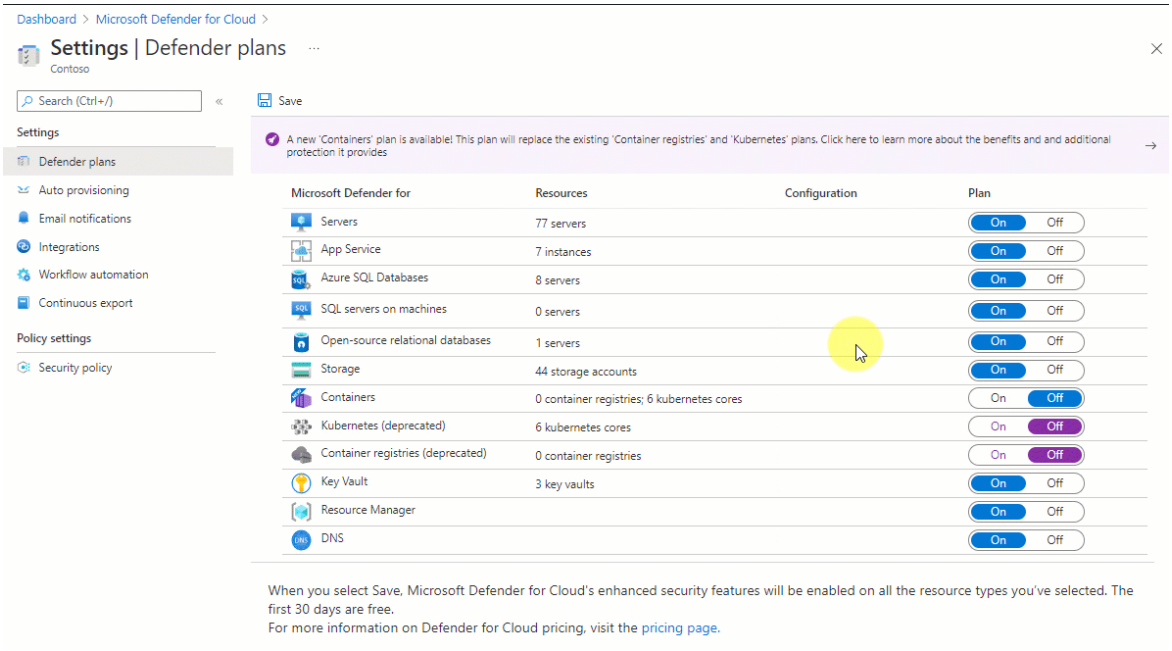

You can also modify this configuration from the Auto provisioning page on the **Microsoft Defender for Containers components (preview)** row:

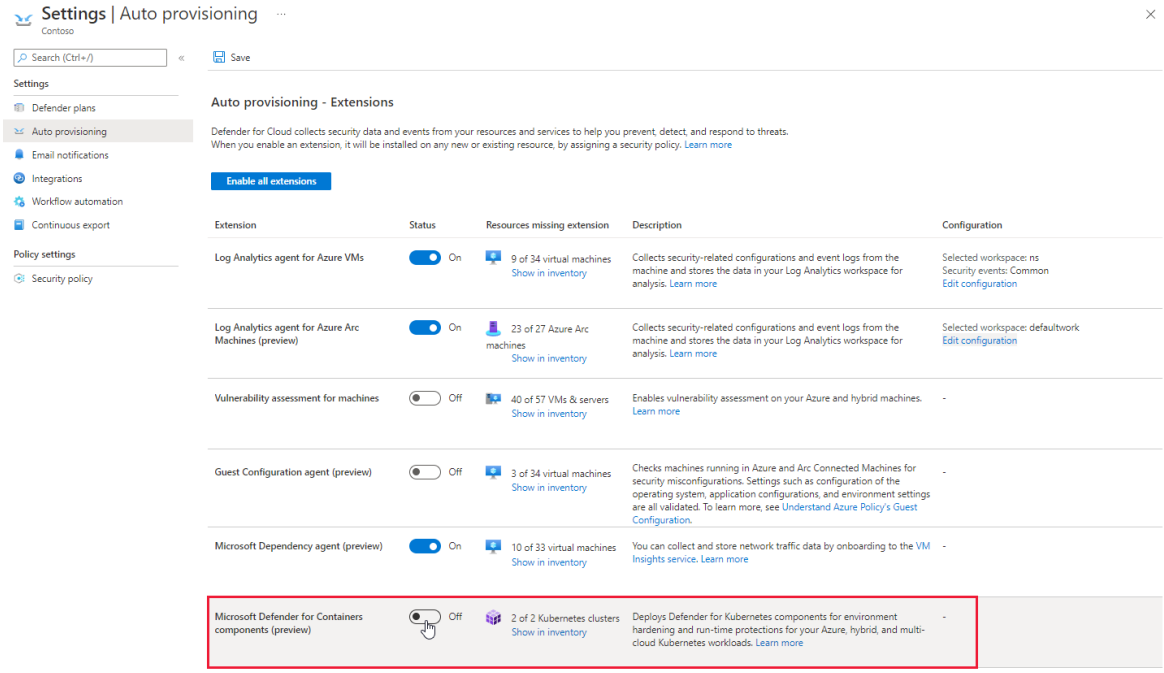

> **NOTE**
>
> If you choose to **disable the plan** at any time after enabling if through the portal as shown above, you'll need to manually disable auto provisioning of the Defender for Containers components. This will not remove the components from machines on which they've already been deployed.

4. If you disable the auto provisioning of any component, you can easily deploy the component to one or more clusters using the appropriate recommendation:

   - Policy Add-on for Kubernetes - Azure Kubernetes Service clusters should have the Azure Policy Add-on for Kubernetes installed
   - Azure Kubernetes Service profile - Azure Kubernetes Service clusters should have Defender profile enabled
   - Azure Arc-enabled Kubernetes extension - Azure Arc-enabled Kubernetes clusters should have the

Defender extension installed

Additional Prerequisites

Before deploying the extension, ensure you:

- Connect the Kubernetes cluster to Azure Arc
- Complete the pre-requisites listed under the generic cluster extensions documentation.

Deploy the Defender extension

You can deploy the Defender extension using a range of methods. For detailed steps, select the relevant tab.

- **Azure portal**
- **Azure CLI**
- **Resource Manager**
- **REST API**

Use the fix button from the Defender for Cloud recommendation

A dedicated Defender for Cloud recommendation provides:

- **Visibility** about which of your clusters has the Defender for Kubernetes extension deployed
- **Fix** button to deploy it to those clusters without the extension

1. From Microsoft Defender for Cloud's recommendations page, open the **Enable enhanced security** security control.

2. Use the filter to find the recommendation named **Azure Arc-enabled Kubernetes clusters should have Defender for Cloud's extension installed**.

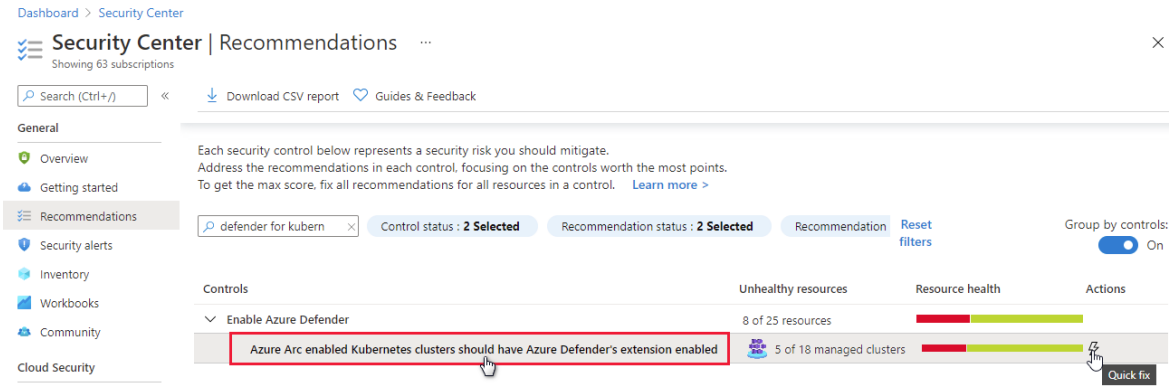

> **TIP**
>
> Notice the **Fix** icon in the actions column

3. Select the extension to see the details of the healthy and unhealthy resources - clusters with and without the extension.

4. From the unhealthy resources list, select a cluster and select **Remediate** to open the pane with the remediation options.

5. Select the relevant Log Analytics workspace and select **Remediate x resource**.

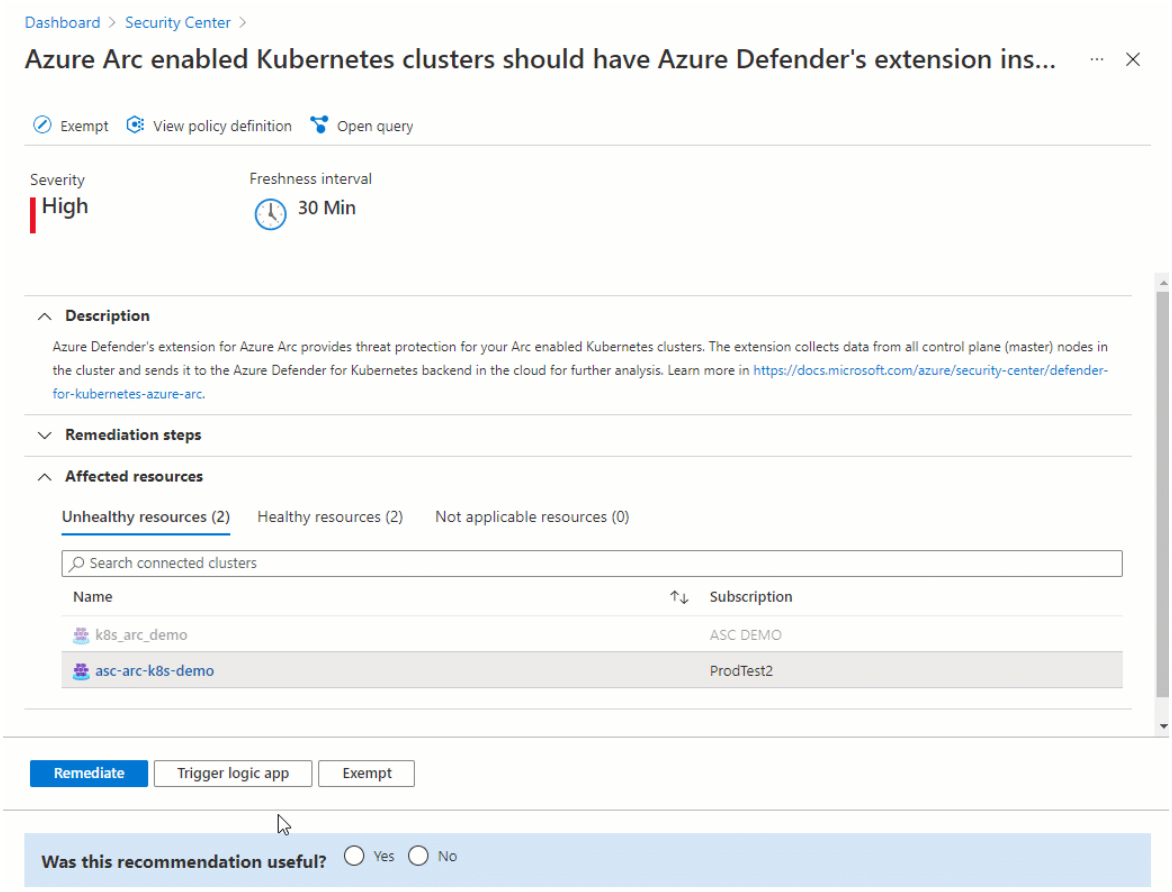

Verify the deployment

To verify that your cluster has the Defender extension installed on it, follow the steps in one of the tabs below:

- Azure portal - Defender for Cloud
- Azure portal - Azure Arc
- Azure CLI
- REST API

Use Defender for Cloud recommendation to verify the status of your extension

1. From Microsoft Defender for Cloud's recommendations page, open the **Enable Microsoft Defender for Cloud** security control.

2. Select the recommendation named **Azure Arc-enabled Kubernetes clusters should have Microsoft Defender for Cloud's extension installed**.

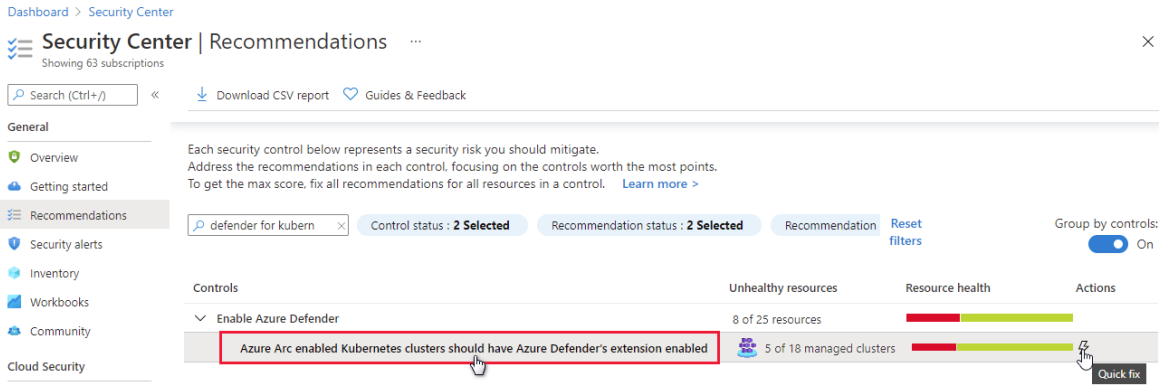

3. Check that the cluster on which you deployed the extension is listed as **Healthy**.

Protect Amazon Elastic Kubernetes Service clusters

> **IMPORTANT**
>
> If you haven't already connected an AWS account, do so now using the instructions in Connect your AWS accounts to Microsoft Defender for Cloud and skip to step 3 below.

To protect your EKS clusters, enable the Containers plan on the relevant account connector:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the AWS connector.

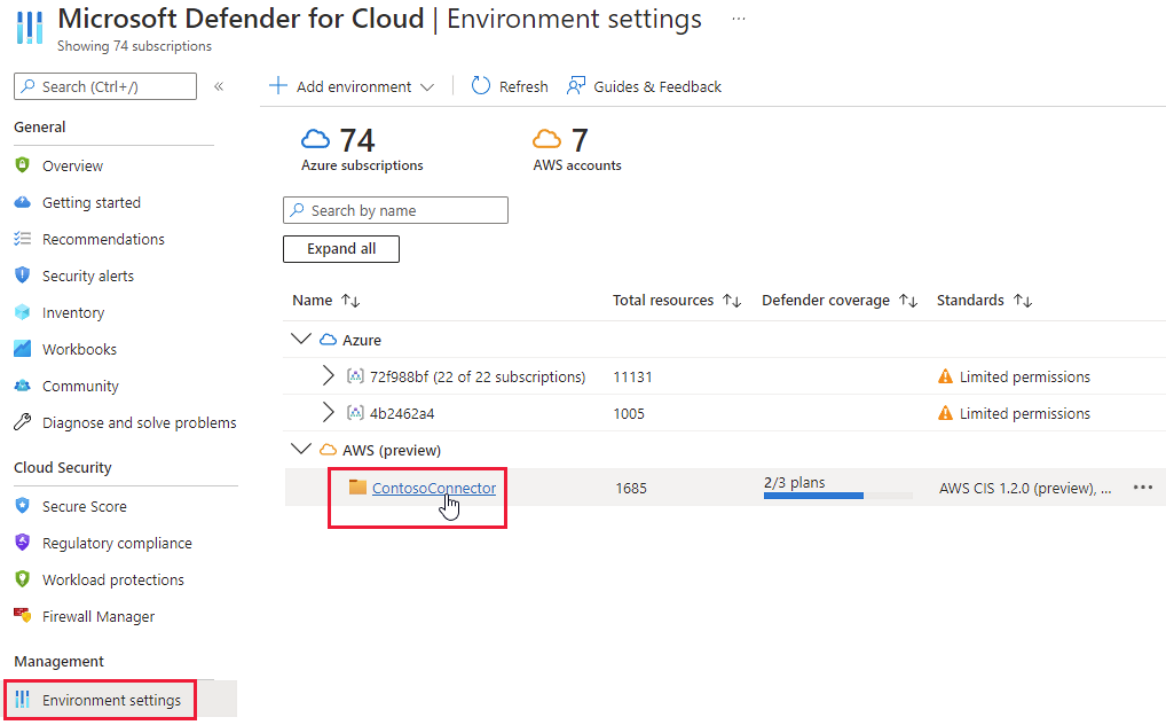

3. Set the toggle for the **Containers** plan to **On**.

4. Optionally, to change the retention period for your audit logs, select **Configure**, enter the required timeframe, and select **Save**.

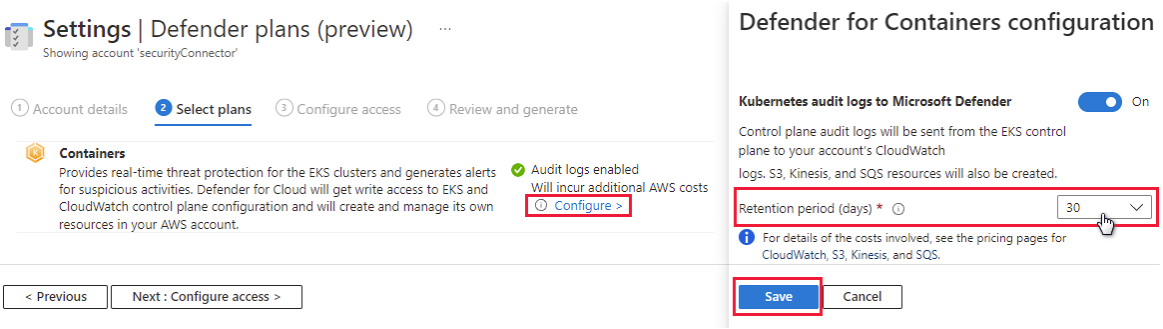

5. Continue through the remaining pages of the connector wizard.
6. Azure Arc-enabled Kubernetes and the Defender extension should be installed and running on your EKS clusters. A dedicated Defender for Cloud recommendation deploys the extension (and Arc if necessary):

a. From Defender for Cloud's **Recommendations** page, search for **EKS clusters should have Azure Defender's extension for Azure Arc installed**.

b. Select an unhealthy cluster.

> **IMPORTANT**
>
> You must select the clusters one at a time.
>
> Don't select the clusters by their hyperlinked names: select anywhere else in the relevant row.

c. Select **Fix**.

d. Defender for Cloud generates a script in the language of your choice: select Bash (for Linux) or PowerShell (for Windows).

e. Select **Download remediation logic**.

f. Run the generated script on your cluster.

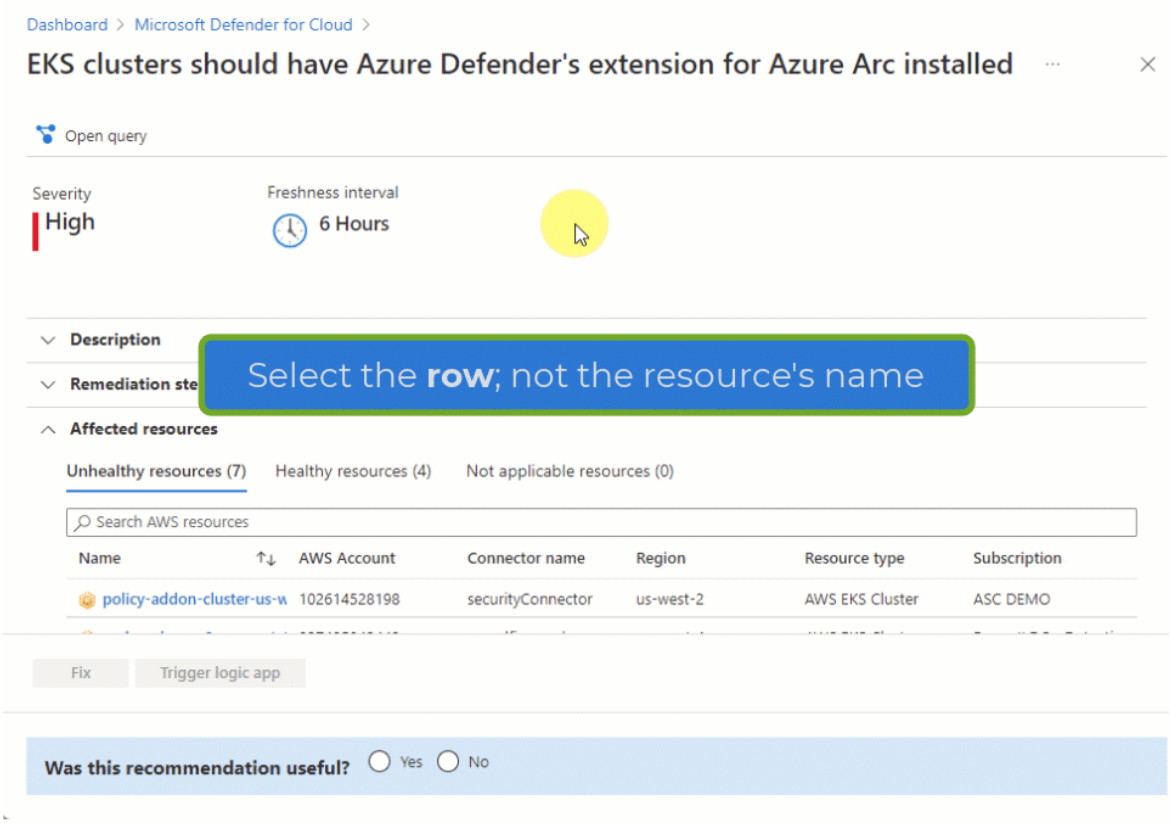

View recommendations and alerts for your EKS clusters

> **TIP**
>
> You can simulate container alerts by following the instructions in this blog post.

To view the alerts and recommendations for your EKS clusters, use the filters on the alerts, recommendations, and inventory pages to filter by resource type **AWS EKS cluster**.

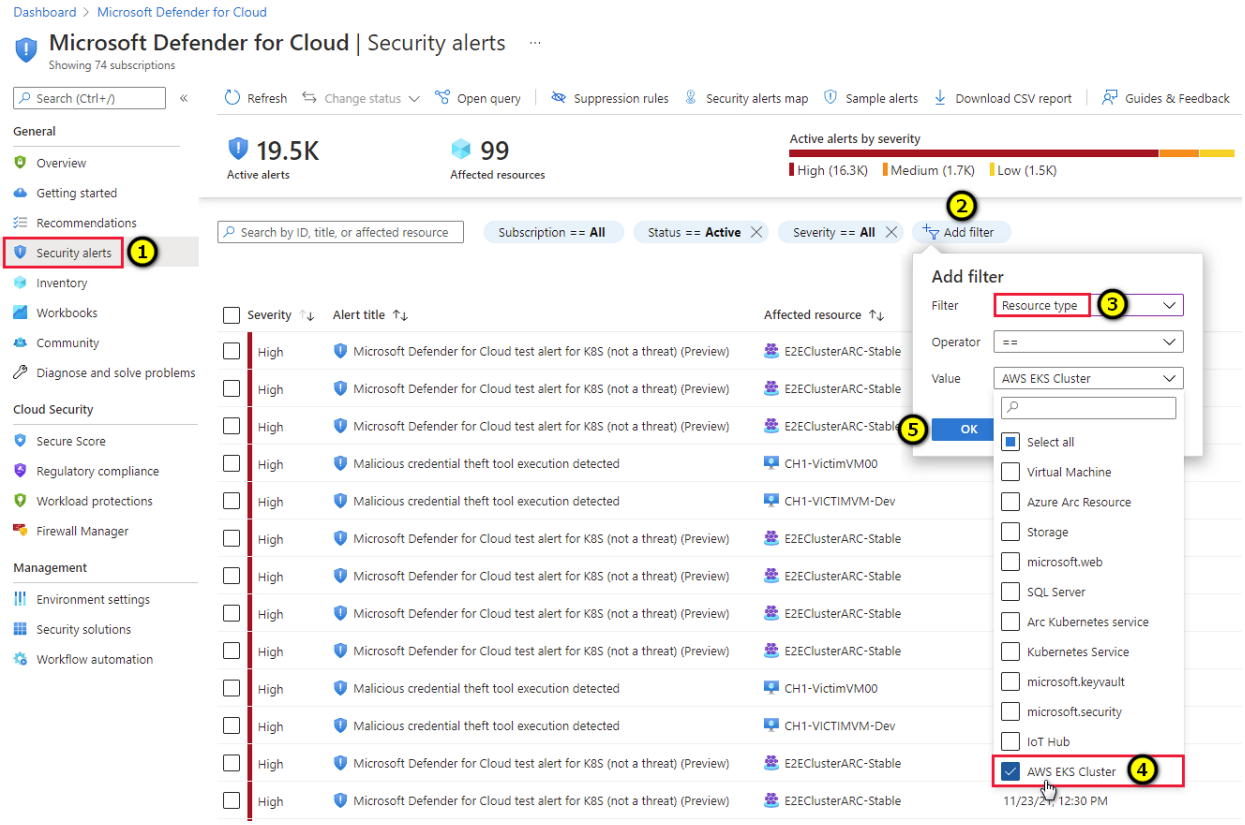

Simulate security alerts from Microsoft Defender for Containers

A full list of supported alerts is available in the reference table of all Defender for Cloud security alerts.

1. To simulate a security alert, run the following command from the cluster:

   ```
   kubectl get pods --namespace=asc-alerttest-662jfi039n
   ```

   The expected response is "No resource found".

   Within 30 minutes, Defender for Cloud will detect this activity and trigger a security alert.

2. In the Azure portal, open Microsoft Defender for Cloud's security alerts page and look for the alert on the relevant resource:

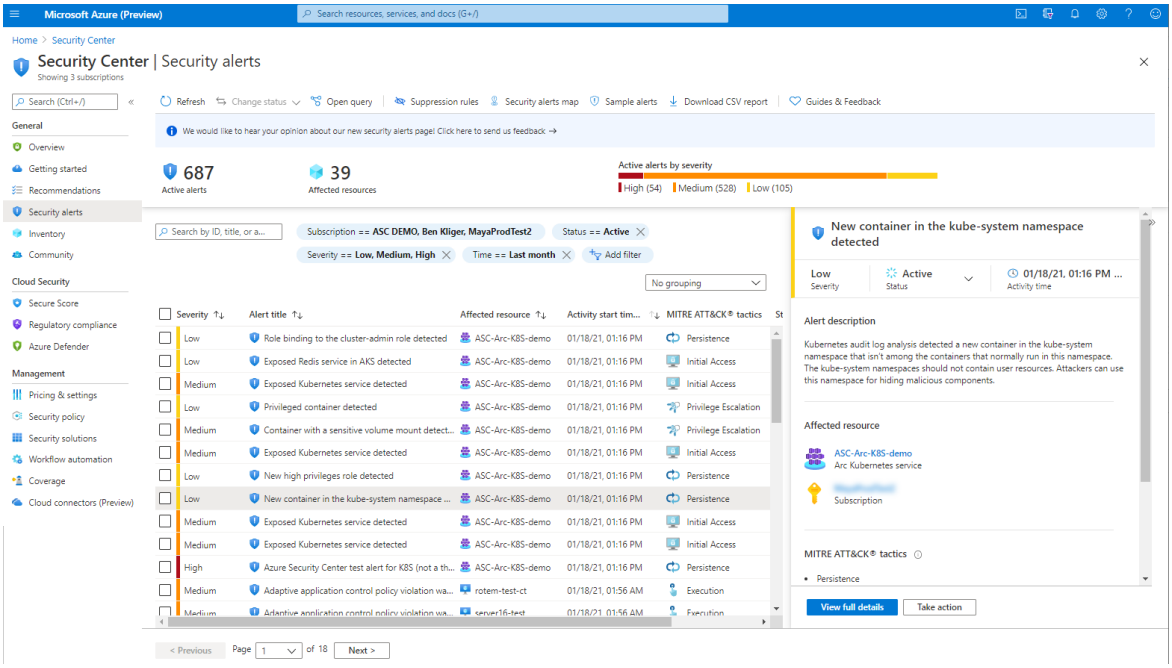

Remove the Defender extension

To remove this - or any - Defender for Cloud extension, it's not enough to turn off auto provisioning:

- **Enabling** auto provisioning, potentially impacts *existing* and *future* machines.
- **Disabling** auto provisioning for an extension, only affects the *future* machines - nothing is uninstalled by disabling auto provisioning.

Nevertheless, to ensure the Defender for Containers components aren't automatically provisioned to your resources from now on, disable auto provisioning of the extensions as explained in Configure auto provisioning for agents and extensions from Microsoft Defender for Cloud.

You can remove the extension using Azure portal, Azure CLI, or REST API as explained in the tabs below.

- Azure portal - Arc
- Azure CLI
- REST API

Use Azure portal to remove the extension

1. From the Azure portal, open Azure Arc.
2. From the infrastructure list, select **Kubernetes clusters** and then select the specific cluster.
3. Open the extensions page. The extensions on the cluster are listed.
4. Select the cluster and select **Uninstall**.

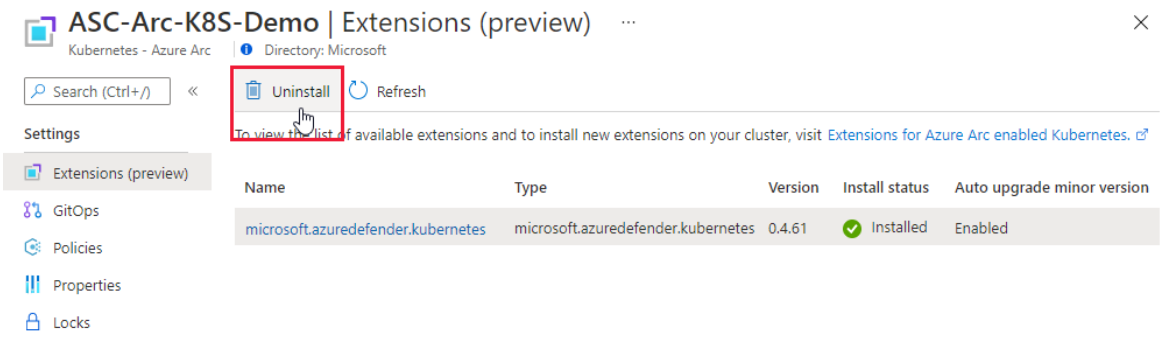

Remove the Defender profile

To remove this - or any - Defender for Cloud extension, it's not enough to turn off auto provisioning:

- **Enabling** auto provisioning, potentially impacts *existing* and *future* machines.
- **Disabling** auto provisioning for an extension, only affects the *future* machines - nothing is uninstalled by disabling auto provisioning.

Nevertheless, to ensure the Defender for Containers components aren't automatically provisioned to your resources from now on, disable auto provisioning of the extensions as explained in Configure auto provisioning for agents and extensions from Microsoft Defender for Cloud.

You can remove the profile using the REST API or a Resource Manager template as explained in the tabs below.

- REST API
- Resource Manager

Use REST API to remove the Defender profile from AKS

To remove the profile using the REST API, run the following PUT command:

```
https://management.azure.com/subscriptions/{{SubscriptionId}}/resourcegroups/{{ResourceGroup}}/providers/Microsoft.ContainerService/managedClusters/{{ClusterName}}?api-version={{ApiVersion}}
```

NAME	DESCRIPTION	MANDATORY
SubscriptionId	Cluster's subscription ID	Yes
ResourceGroup	Cluster's resource group	Yes
ClusterName	Cluster's name	Yes
ApiVersion	API version, must be >= 2021-07-01	Yes

Request body:

```
{
  "location": "{{Location}}",
  "properties": {
    "securityProfile": {
            "azureDefender": {
                "enabled": false
            }
        }
    }
}
```

Request body parameters:

NAME	DESCRIPTION	MANDATORY
location	Cluster's location	Yes
properties.securityProfile.azureDefender.enabled	Determines whether to enable or disable Microsoft Defender for Containers on the cluster	Yes

Introduction to Microsoft Defender for Kubernetes (deprecated)

2/15/2022 • 6 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud provides environment hardening, workload protection, and run-time protections as outlined in Container security in Defender for Cloud.

Defender for Kubernetes protects your Kubernetes clusters whether they're running in:

- **Azure Kubernetes Service (AKS)** - Microsoft's managed service for developing, deploying, and managing containerized applications.
- **Amazon Elastic Kubernetes Service (EKS) in a connected Amazon Web Services (AWS) account** (preview) - Amazon's managed service for running Kubernetes on AWS without needing to install, operate, and maintain your own Kubernetes control plane or nodes.
- **An unmanaged Kubernetes distribution** - Cloud Native Computing Foundation (CNCF) certified Kubernetes clusters on premises or on IaaS. Learn more in Defend Azure Arc-enabled Kubernetes clusters running in on-premises and multi-cloud environments.

Host-level threat detection for your Linux AKS nodes is available if you enable Microsoft Defender for servers and its Log Analytics agent. However, if your cluster is deployed on an Azure Kubernetes Service virtual machine scale set, the Log Analytics agent is not currently supported.

Availability

> **IMPORTANT**
> Microsoft Defender for Kubernetes has been replaced with **Microsoft Defender for Containers**. If you've already enabled Defender for Kubernetes on a subscription, you can continue to use it. However, you won't get Defender for Containers' improvements and new features.
>
> This plan is no longer available for subscriptions where it isn't already enabled.
>
> To upgrade to Microsoft Defender for Containers, open the Defender plans page in the portal and enable the new plan:

Learn more about this change in the release note.

ASPECT	DETAILS
Release state:	General availability (GA) Protections for EKS clusters are preview. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	**Microsoft Defender for Kubernetes** is billed as shown on the pricing page. **Containers plan** for EKS clusters in connected AWS accounts is free while it's in preview.
Required roles and permissions:	**Security admin** can dismiss alerts. **Security reader** can view findings.
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet) ✅ Connected AWS accounts (Preview)

What are the benefits of Microsoft Defender for Kubernetes?

Our global team of security researchers constantly monitor the threat landscape. As container-specific alerts and vulnerabilities are discovered, these researchers add them to our threat intelligence feeds and Defender for Cloud alerts you to any that are relevant for your environment.

In addition, Microsoft Defender for Kubernetes provides **cluster-level threat protection** by monitoring your clusters' logs. This means that security alerts are only triggered for actions and deployments that occur *after* you've enabled Defender for Kubernetes on your subscription.

TIP

For EKS-based clusters, we monitor the control plane audit logs. These are enabled in the containers plan configuration:

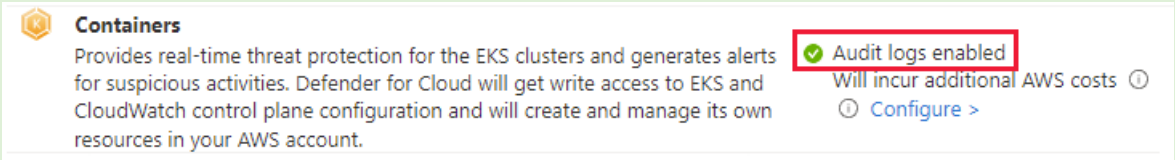

Examples of security events that Microsoft Defender for Kubernetes monitors include:

- Exposed Kubernetes dashboards
- Creation of high privileged roles
- Creation of sensitive mounts.

For a full list of the cluster level alerts, see the reference table of alerts.

Protect Azure Kubernetes Service (AKS) clusters

To protect your AKS clusters, enable the Defender plan on the relevant subscription:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.

3. In the **Defender plans** page, set the status of Microsoft Defender for Kubernetes to **On**.

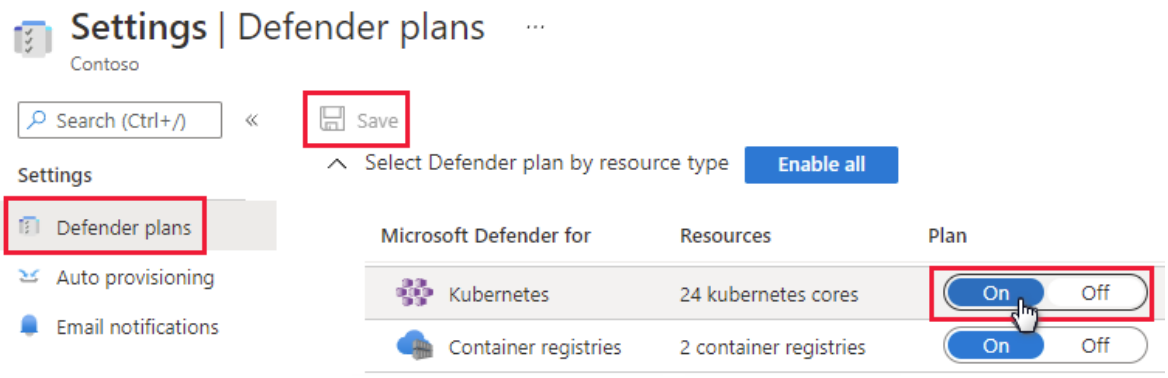

4. Select **Save**.

Protect Amazon Elastic Kubernetes Service clusters

> **IMPORTANT**
> If you haven't already connected an AWS account, do so now using the instructions in Connect your AWS accounts to Microsoft Defender for Cloud and skip to step 3 below.

To protect your EKS clusters, enable the Containers plan on the relevant account connector:

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the AWS connector.

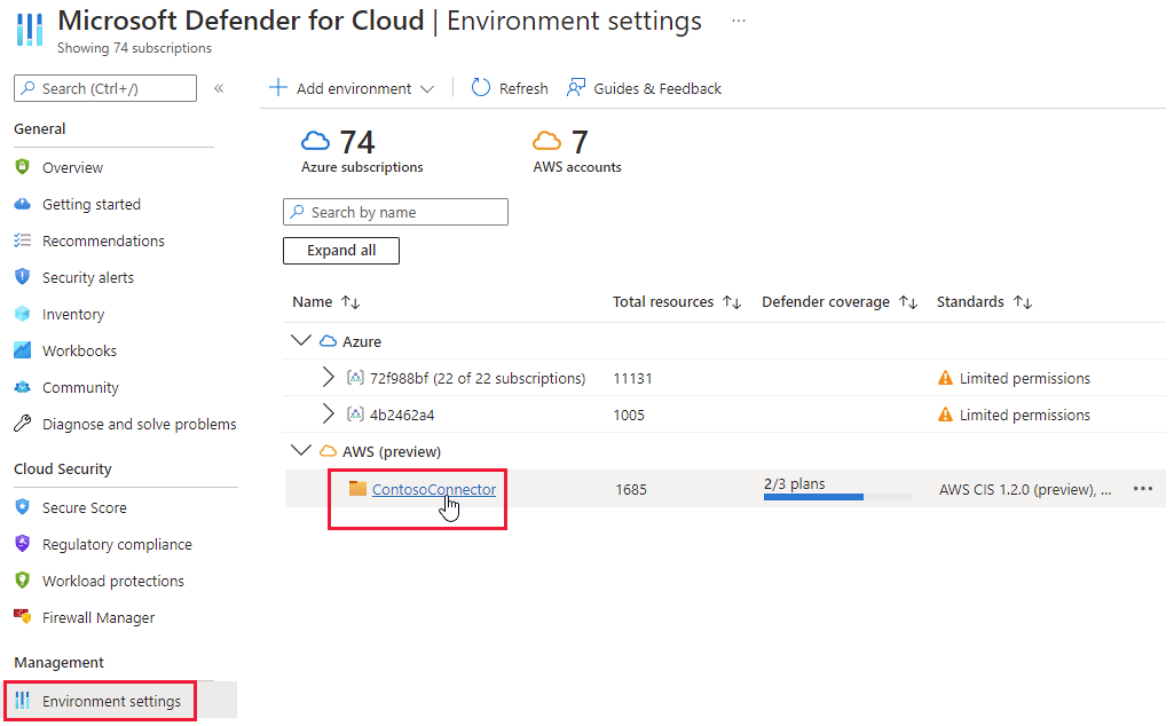

3. Set the toggle for the **Containers** plan to **On**.

4. Optionally, to change the retention period for your audit logs, select **Configure**, enter the desired timeframe, and select **Save**.

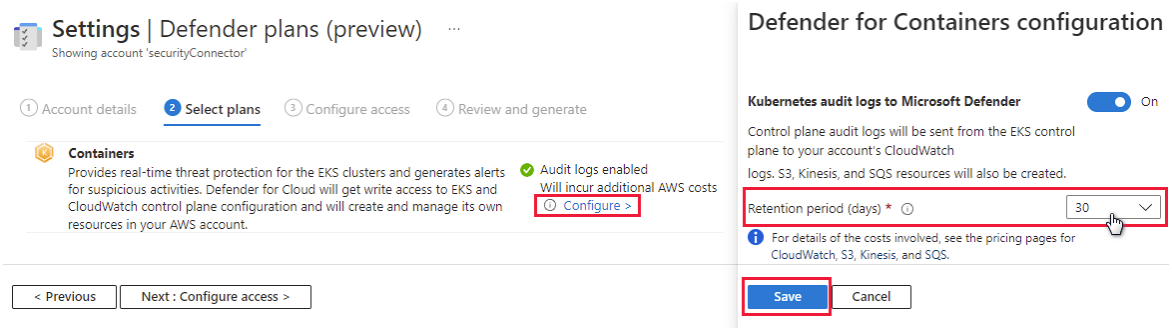

5. Continue through the remaining pages of the connector wizard.

6. Azure Arc-enabled Kubernetes and the Defender extension should be installed and running on your EKS clusters. A dedicated Defender for Cloud recommendation deploys the extension (and Arc if necessary):

   a. From Defender for Cloud's **Recommendations** page, search for **EKS clusters should have Azure Defender's extension for Azure Arc installed**.

   b. Select an unhealthy cluster.

   > **IMPORTANT**
   >
   > You must select the clusters one at a time.
   >
   > Don't select the clusters by their hyperlinked names: select anywhere else in the relevant row.

   c. Select **Fix**.

   d. Defender for Cloud generates a script in the language of your choice: select Bash (for Linux) or PowerShell (for Windows).

   e. Select **Download remediation logic**.

   f. Run the generated script on your cluster.

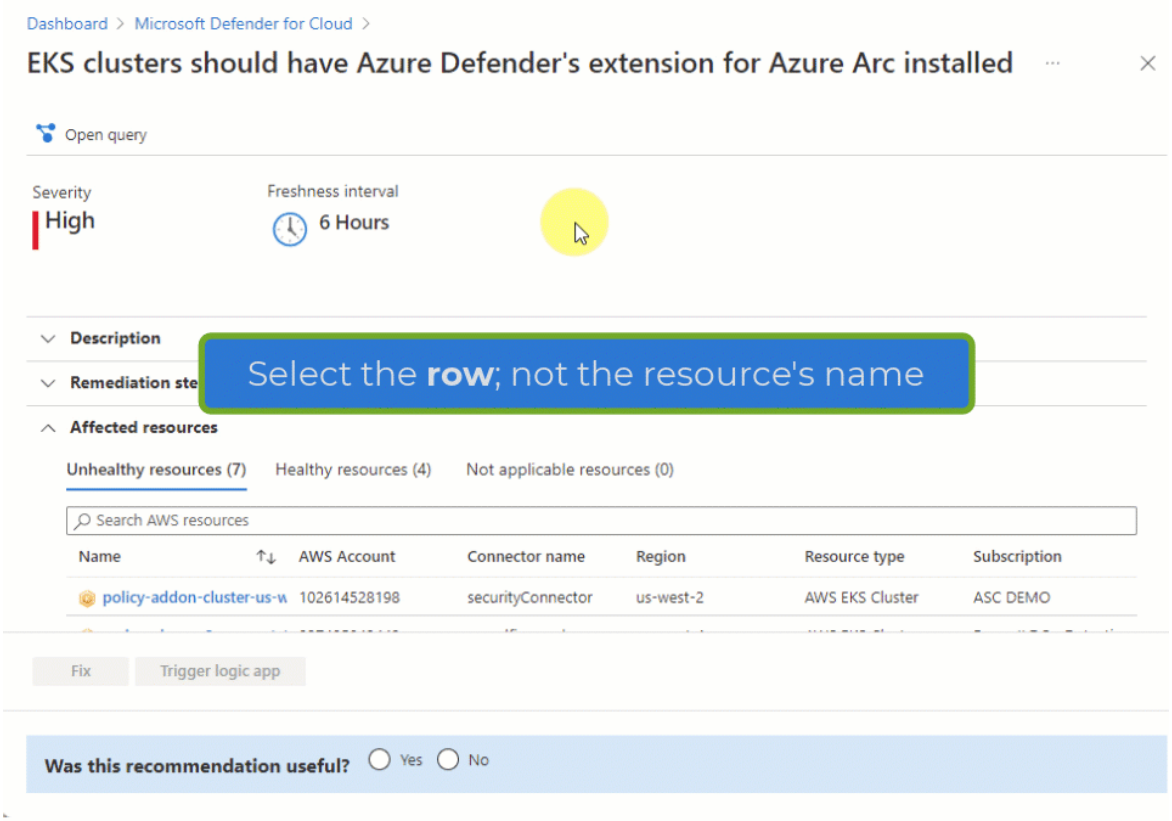

View recommendations and alerts for your EKS clusters

> **TIP**
> You can simulate container alerts by following the instructions in this blog post.

To view the alerts and recommendations for your EKS clusters, use the filters on the alerts, recommendations, and inventory pages to filter by resource type **AWS EKS cluster**.

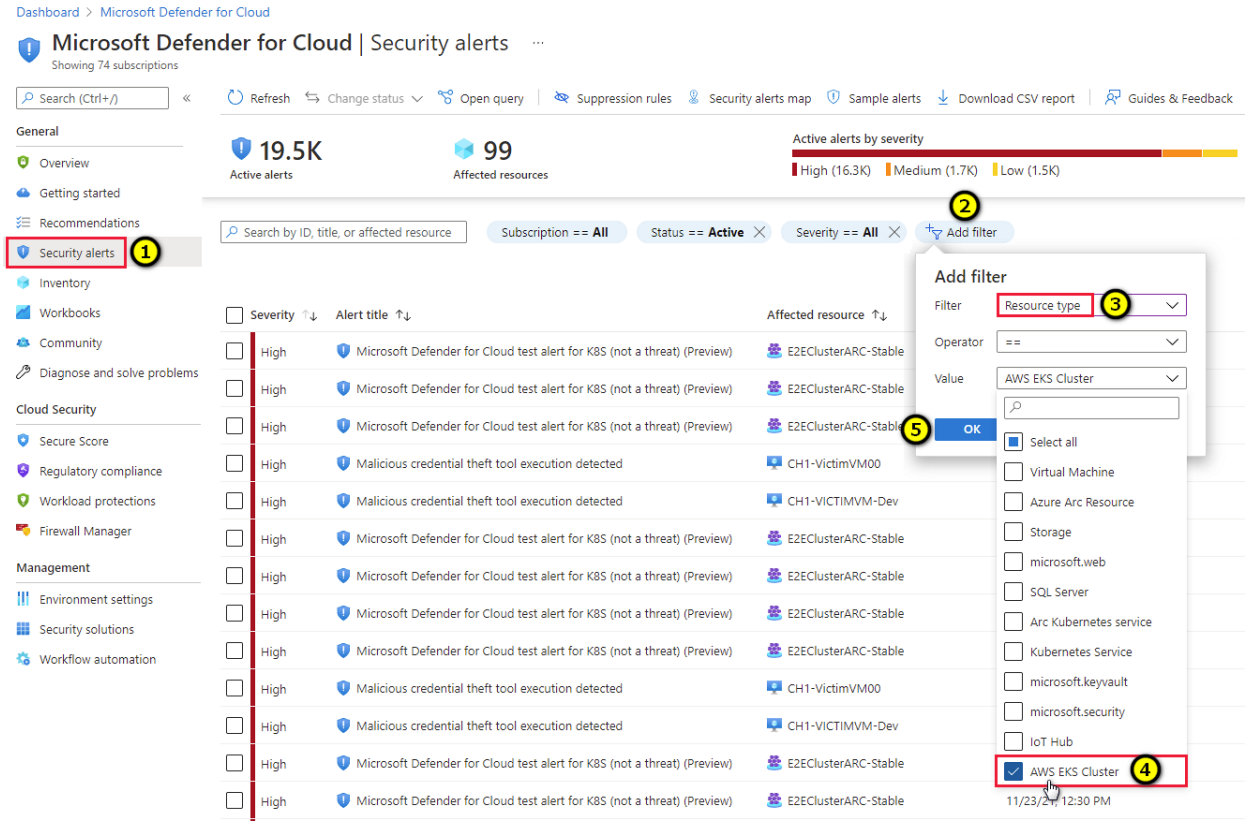

FAQ - Microsoft Defender for Kubernetes

- Can I still get cluster protections without the Log Analytics agent?
- Does AKS allow me to install custom VM extensions on my AKS nodes?
- If my cluster is already running an Azure Monitor for containers agent, do I need the Log Analytics agent too?
- Does Microsoft Defender for Kubernetes support AKS with virtual machine scale set nodes?

Can I still get cluster protections without the Log Analytics agent?

Microsoft Defender for Kubernetes provides protections at the cluster level. If you also deploy the Log Analytics agent of **Microsoft Defender for servers**, you'll get the threat protection for your nodes that's provided with that plan. Learn more in Introduction to Microsoft Defender for servers.

We recommend deploying both, for the most complete protection possible.

If you choose not to install the agent on your hosts, you'll only receive a subset of the threat protection benefits and security alerts. You'll still receive alerts related to network analysis and communications with malicious servers.

Does AKS allow me to install custom VM extensions on my AKS nodes?

For Defender for Cloud to monitor your AKS nodes, they must be running the Log Analytics agent.

AKS is a managed service and since the Log Analytics agent is a Microsoft-managed extension, it is also supported on AKS clusters. However, if your cluster is deployed on an Azure Kubernetes Service virtual machine

scale set, the Log Analytics agent isn't currently supported.

If my cluster is already running an Azure Monitor for containers agent, do I need the Log Analytics agent too?

For Defender for Cloud to monitor your nodes, they must be running the Log Analytics agent.

If your clusters are already running the Azure Monitor for containers agent, you can install the Log Analytics agent too and the two agents can work alongside one another without any problems.

Learn more about the Azure Monitor for containers agent.

Does Microsoft Defender for Kubernetes support AKS with virtual machine scale set nodes?

If your cluster is deployed on an Azure Kubernetes Service virtual machine scale set, the Log Analytics agent is not currently supported.

Next steps

In this article, you learned about Kubernetes protection in Defender for Cloud, including Microsoft Defender for Kubernetes.

Enable enhanced protections

For related material, see the following articles:

- Stream alerts to a SIEM, SOAR, or IT Service Management solution
- Reference table of alerts

Introduction to Microsoft Defender for container registries (deprecated)

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Container Registry (ACR) is a managed, private Docker registry service that stores and manages your container images for Azure deployments in a central registry. It's based on the open-source Docker Registry 2.0.

To protect the Azure Resource Manager based registries in your subscription, enable **Microsoft Defender for container registries** at the subscription level. Defender for Cloud will then scan all images when they're pushed to the registry, imported into the registry, or pulled within the last 30 days. You'll be charged for every image that gets scanned – once per image.

Availability

> **IMPORTANT**
> Microsoft Defender for container registries has been replaced with **Microsoft Defender for Containers**. If you've already enabled Defender for container registries on a subscription, you can continue to use it. However, you won't get Defender for Containers' improvements and new features.
>
> This plan is no longer available for subscriptions where it isn't already enabled.
>
> To upgrade to Microsoft Defender for Containers, open the Defender plans page in the portal and enable the new plan:

> Learn more about this change in the release note.

ASPECT	DETAILS
Release state:	Generally available (GA)
Pricing:	**Microsoft Defender for container registries** is billed as shown on the pricing page
Supported registries and images:	Linux images in ACR registries accessible from the public internet with shell access ACR registries protected with Azure Private Link

ASPECT	DETAILS
Unsupported registries and images:	Windows images 'Private' registries (unless access is granted to Trusted Services) Super-minimalist images such as Docker scratch images, or "Distroless" images that only contain an application and its runtime dependencies without a package manager, shell, or OS Images with Open Container Initiative (OCI) Image Format Specification
Required roles and permissions:	**Security reader** and Azure Container Registry roles and permissions
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

What are the benefits of Microsoft Defender for container registries?

Defender for Cloud identifies Azure Resource Manager based ACR registries in your subscription and seamlessly provides Azure-native vulnerability assessment and management for your registry's images.

Microsoft Defender for container registries includes a vulnerability scanner to scan the images in your Azure Resource Manager-based Azure Container Registry registries and provide deeper visibility into your images' vulnerabilities. The integrated scanner is powered by Qualys, the industry-leading vulnerability scanning vendor.

When issues are found – by Qualys or Defender for Cloud – you'll get notified in the workload protection dashboard. For every vulnerability, Defender for Cloud provides actionable recommendations, along with a severity classification, and guidance for how to remediate the issue. For details of Defender for Cloud's recommendations for containers, see the reference list of recommendations.

Defender for Cloud filters and classifies findings from the scanner. When an image is healthy, Defender for Cloud marks it as such. Defender for Cloud generates security recommendations only for images that have issues to be resolved. Defender for Cloud provides details of each reported vulnerability and a severity classification. Additionally, it gives guidance for how to remediate the specific vulnerabilities found on each image.

By only notifying when there are problems, Defender for Cloud reduces the potential for unwanted informational alerts.

When are images scanned?

There are three triggers for an image scan:

- **On push** - Whenever an image is pushed to your registry, Defender for container registries automatically scans that image. To trigger the scan of an image, push it to your repository.
- **Recently pulled** - Since new vulnerabilities are discovered every day, **Microsoft Defender for container registries** also scans, on a weekly basis, any image that has been pulled within the last 30 days. There's no additional charge for these rescans; as mentioned above, you're billed once per image.
- **On import** - Azure Container Registry has import tools to bring images to your registry from Docker Hub, Microsoft Container Registry, or another Azure container registry. **Microsoft Defender for**

container registries scans any supported images you import. Learn more in Import container images to a container registry.

The scan completes typically within 2 minutes, but it might take up to 40 minutes. Findings are made available as security recommendations such as this one:

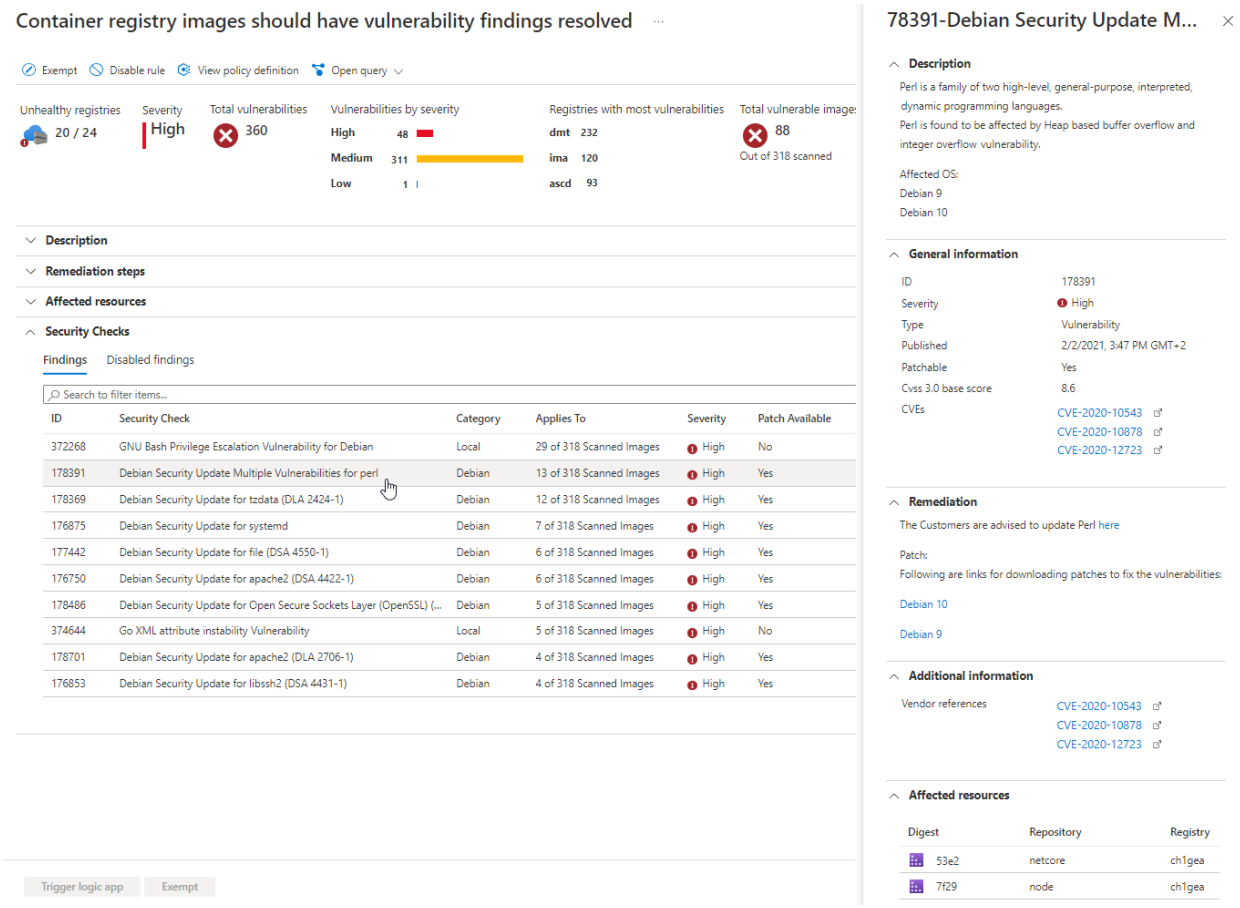

How does Defender for Cloud work with Azure Container Registry

Below is a high-level diagram of the components and benefits of protecting your registries with Defender for Cloud.

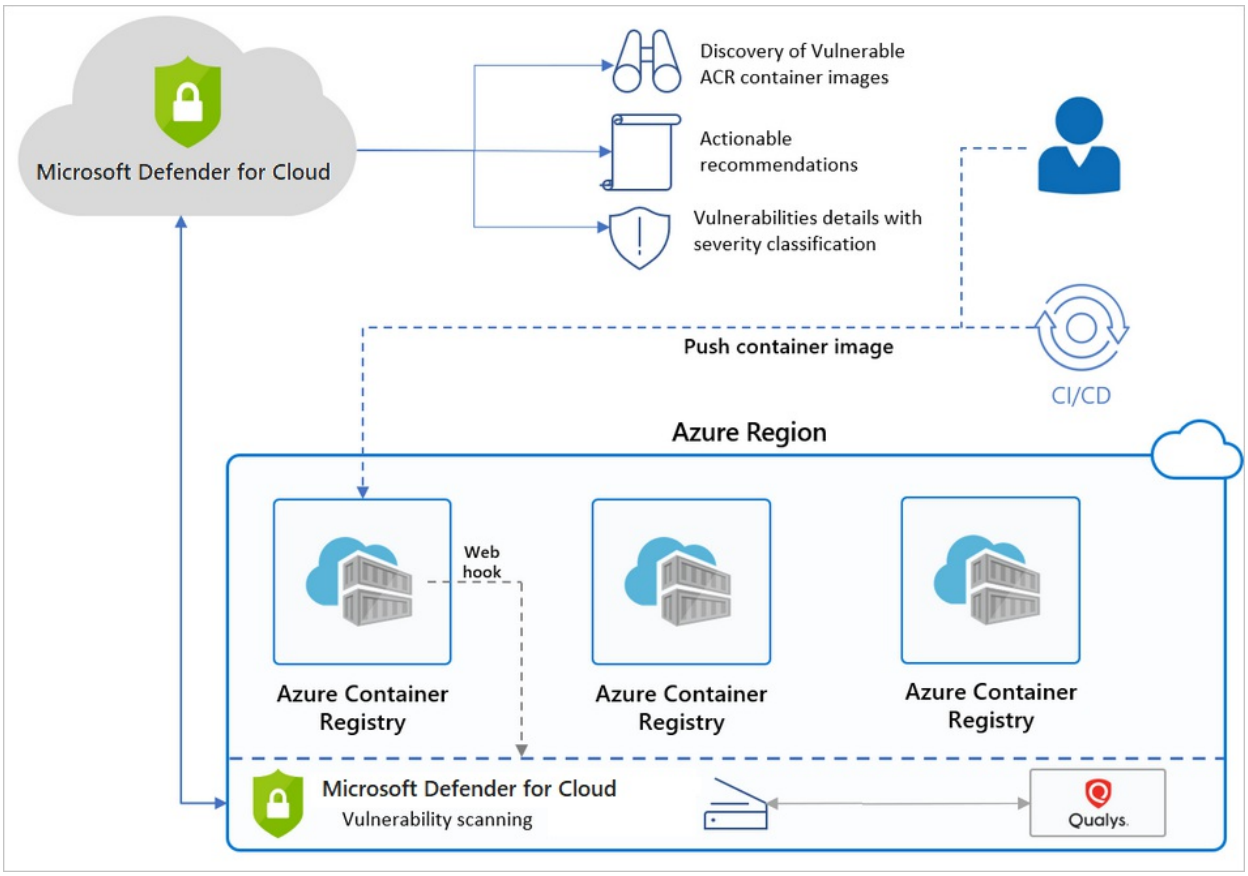

FAQ - Azure Container Registry image scanning

How does Defender for Cloud scan an image?

Defender for Cloud pulls the image from the registry and runs it in an isolated sandbox with the Qualys scanner. The scanner extracts a list of known vulnerabilities.

Defender for Cloud filters and classifies findings from the scanner. When an image is healthy, Defender for Cloud marks it as such. Defender for Cloud generates security recommendations only for images that have issues to be resolved. By only notifying you when there are problems, Defender for Cloud reduces the potential for unwanted informational alerts.

Can I get the scan results via REST API?

Yes. The results are under Sub-Assessments Rest API. Also, you can use Azure Resource Graph (ARG), the Kusto-like API for all of your resources: a query can fetch a specific scan.

What registry types are scanned? What types are billed?

For a list of the types of container registries supported by Microsoft Defender for container registries, see Availability.

If you connect unsupported registries to your Azure subscription, Defender for Cloud won't scan them and won't bill you for them.

Can I customize the findings from the vulnerability scanner?

Yes. If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

Learn about creating rules to disable findings from the integrated vulnerability assessment tool.

Why is Defender for Cloud alerting me to vulnerabilities about an image that isn't in my registry?

Defender for Cloud provides vulnerability assessments for every image pushed or pulled in a registry. Some images may reuse tags from an image that was already scanned. For example, you may reassign the tag "Latest"

every time you add an image to a digest. In such cases, the 'old' image does still exist in the registry and may still be pulled by its digest. If the image has security findings and is pulled, it'll expose security vulnerabilities.

Next steps

Scan your images for vulnerabilities

Use Microsoft Defender for container registries to scan your images for vulnerabilities

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how to use the built-in vulnerability scanner to scan the container images stored in your Azure Resource Manager-based Azure Container Registry.

When **Microsoft Defender for container registries** is enabled, any image you push to your registry will be scanned immediately. In addition, any image pulled within the last 30 days is also scanned.

When the scanner reports vulnerabilities to Defender for Cloud, Defender for Cloud presents the findings and related information as recommendations. In addition, the findings include related information such as remediation steps, relevant CVEs, CVSS scores, and more. You can view the identified vulnerabilities for one or more subscriptions, or for a specific registry.

> **TIP**
> You can also scan container images for vulnerabilities as the images are built in your CI/CD GitHub workflows. Learn more in Identify vulnerable container images in your CI/CD workflows.

Identify vulnerabilities in images in Azure container registries

To enable vulnerability scans of images stored in your Azure Resource Manager-based Azure Container Registry:

1. Enable **Microsoft Defender for container registries** for your subscription. Defender for Cloud is now ready to scan images in your registries.

   > **NOTE**
   > This feature is charged per image.

2. Image scans are triggered on every push or import, and if the image has been pulled within the last 30 days.

   When the scan completes (typically after approximately 2 minutes, but can be up to 15 minutes), findings are available as Defender for Cloud recommendations.

3. View and remediate findings as explained below.

Identify vulnerabilities in images in other container registries

1. Use the ACR tools to bring images to your registry from Docker Hub or Microsoft Container Registry. When the import completes, the imported images are scanned by the built-in vulnerability assessment

solution.

Learn more in Import container images to a container registry

When the scan completes (typically after approximately 2 minutes, but can be up to 15 minutes), findings are available as Defender for Cloud recommendations.

2. View and remediate findings as explained below.

View and remediate findings

1. To view the findings, open the **Recommendations** page. If issues were found, you'll see the recommendation Container registry images should have vulnerability findings resolved (powered by Qualys).

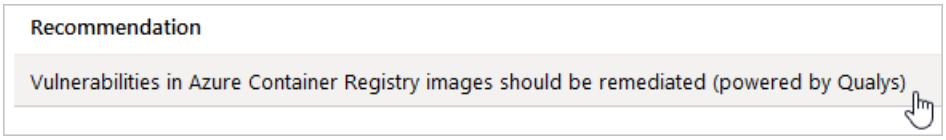

2. Select the recommendation.

   The recommendation details page opens with additional information. This information includes the list of registries with vulnerable images ("Affected resources") and the remediation steps.

3. Select a specific registry to see the repositories within it that have vulnerable repositories.

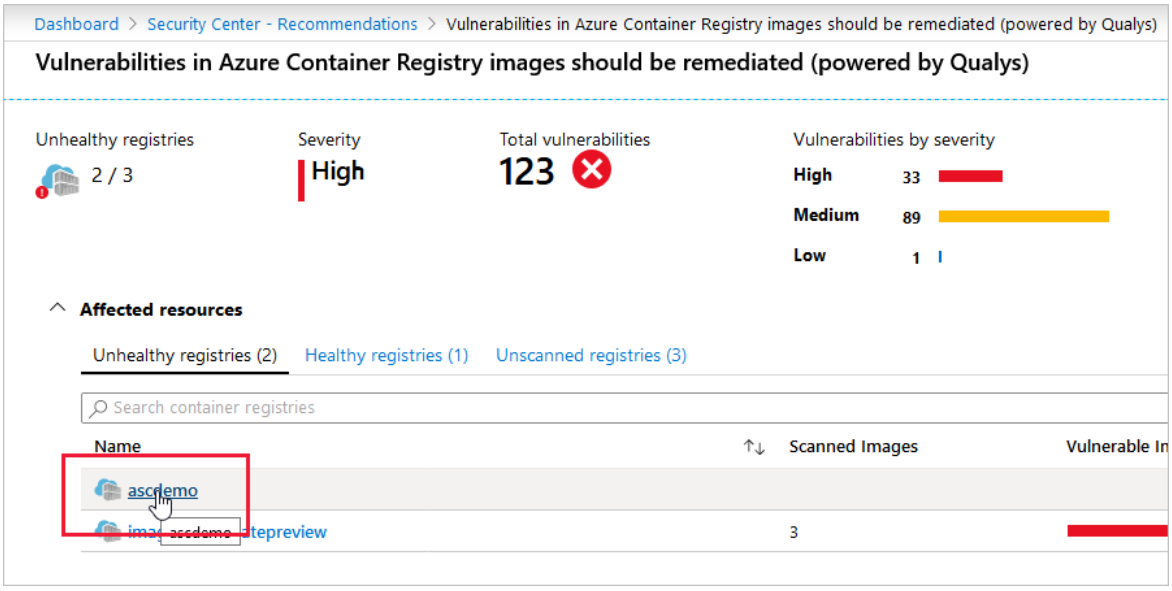

   The registry details page opens with the list of affected repositories.

4. Select a specific repository to see the repositories within it that have vulnerable images.

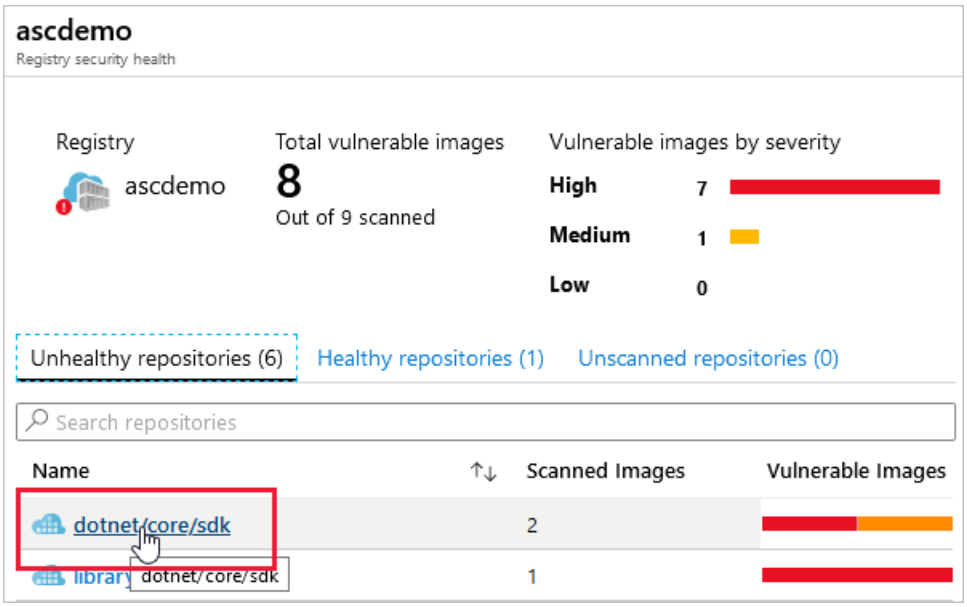

The repository details page opens. It lists the vulnerable images together with an assessment of the severity of the findings.

5. Select a specific image to see the vulnerabilities.

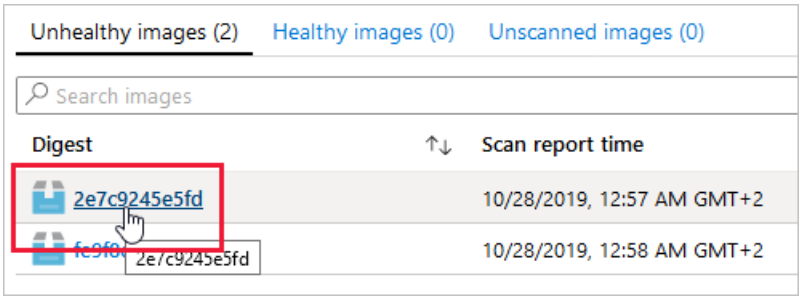

The list of findings for the selected image opens.

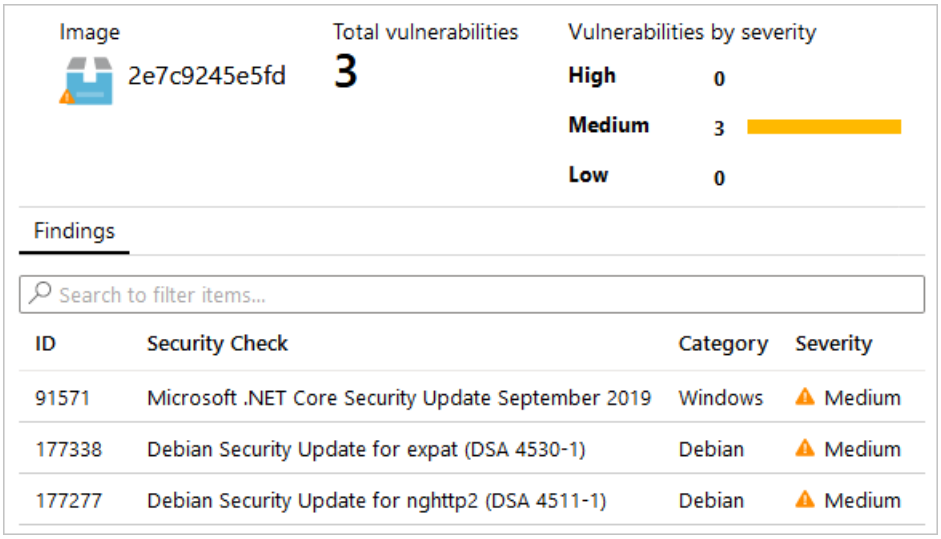

6. To learn more about a finding, select the finding.

   The findings details pane opens.

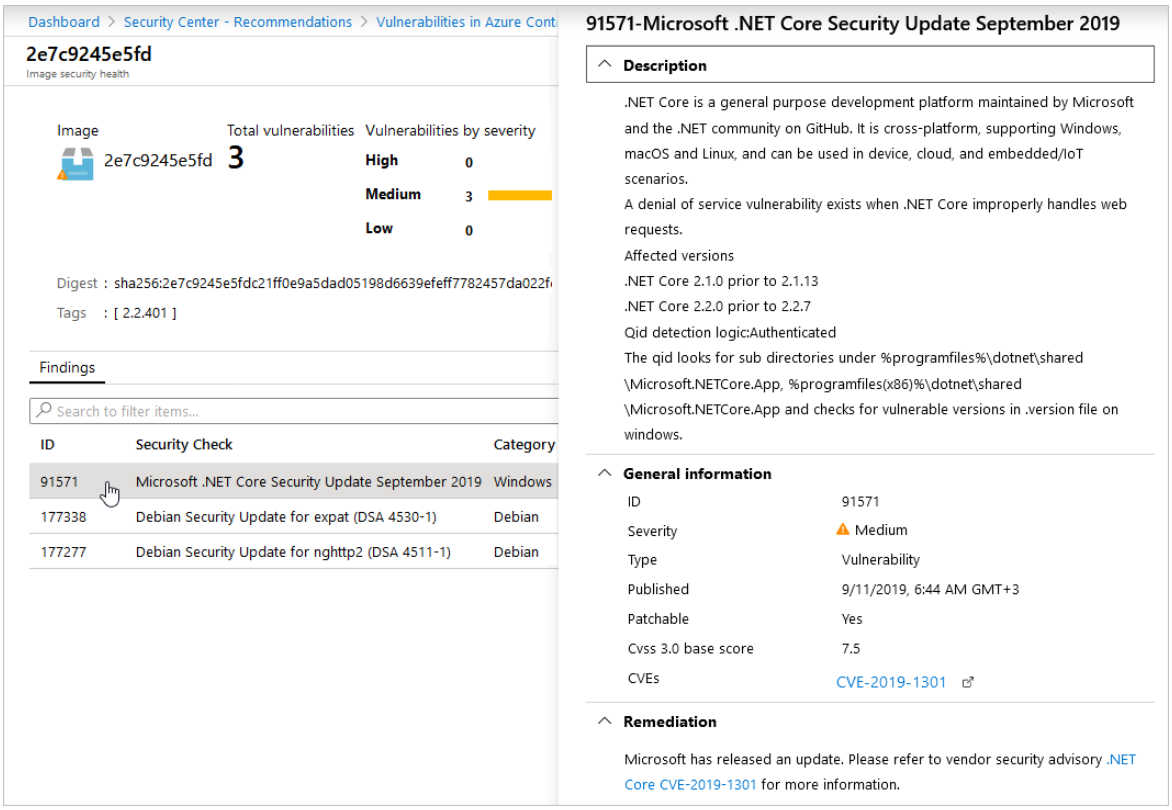

This pane includes a detailed description of the issue and links to external resources to help mitigate the threats.

7. Follow the steps in the remediation section of this pane.
8. When you have taken the steps required to remediate the security issue, replace the image in your registry:

   a. Push the updated image. This will trigger a scan.

   b. Check the recommendations page for the recommendation Container registry images should have vulnerability findings resolved (powered by Qualys).

      If the recommendation still appears and the image you've handled still appears in the list of vulnerable images, check the remediation steps again.

   c. When you are sure the updated image has been pushed, scanned, and is no longer appearing in the recommendation, delete the "old" vulnerable image from your registry.

Disable specific findings

> **NOTE**
>
> The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

When a finding matches the criteria you've defined in your disable rules, it won't appear in the list of findings. Typical scenarios include:

- Disable findings with severity below medium
- Disable findings that are non-patchable

- Disable findings with CVSS score below 6.5
- Disable findings with specific text in the security check or category (for example, "RedHat", "CentOS Security Update for sudo")

> **IMPORTANT**
>
> To create a rule, you need permissions to edit a policy in Azure Policy.
>
> Learn more in Azure RBAC permissions in Azure Policy.

You can use any of the following criteria:

- Finding ID
- Category
- Security check
- CVSS v3 scores
- Severity
- Patchable status

To create a rule:

1. From the recommendations detail page for Container registry images should have vulnerability findings resolved (powered by Qualys), select **Disable rule**.
2. Select the relevant scope.
3. Define your criteria.
4. Select **Apply rule**.

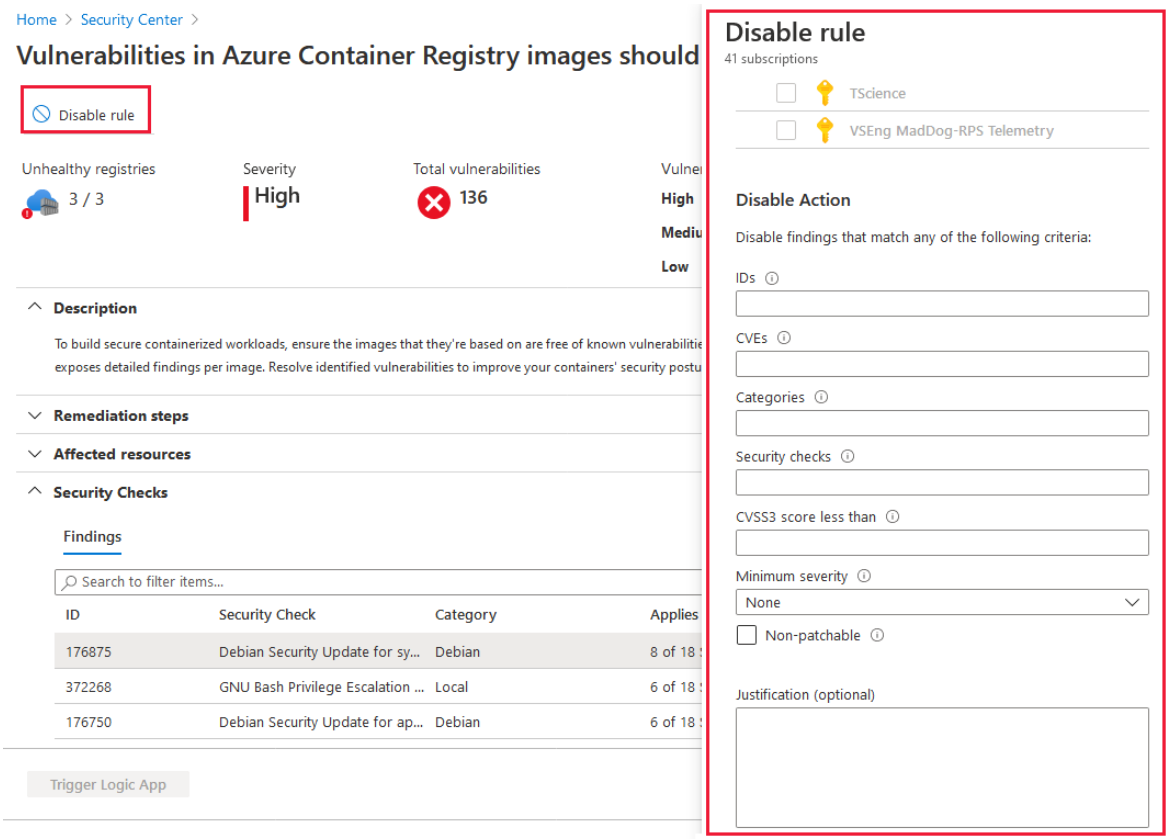

5. To view, override, or delete a rule:
   a. Select **Disable rule**.

b. From the scope list, subscriptions with active rules show as **Rule applied**.

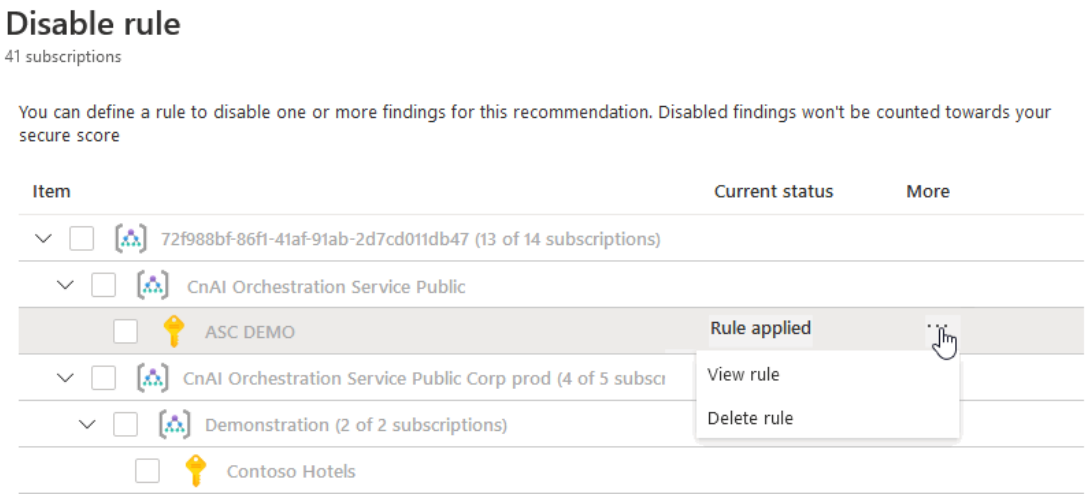

c. To view or delete the rule, select the ellipsis menu ("...").

Next steps

Learn more about the advanced protection plans of Microsoft Defender for Cloud.

Identify vulnerable container images in your CI/CD workflows

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how to scan your Azure Container Registry-based container images with the integrated vulnerability scanner when they're built as part of your GitHub workflows.

To set up the scanner, you'll need to enable **Microsoft Defender for container registries** and the CI/CD integration. When your CI/CD workflows push images to your registries, you can view registry scan results and a summary of CI/CD scan results.

The findings of the CI/CD scans are an enrichment to the existing registry scan findings by Qualys. Defender for Cloud's CI/CD scanning is powered by Aqua Trivy.

You'll get traceability information such as the GitHub workflow and the GitHub run URL, to help identify the workflows that are resulting in vulnerable images.

TIP

The vulnerabilities identified in a scan of your registry might differ from the findings of your CI/CD scans. One reason for these differences is that the registry scanning is continuous, whereas the CI/CD scanning happens immediately before the workflow pushes the image into the registry.

Availability

ASPECT	DETAILS
Release state:	**This CI/CD integration is in preview.** We recommend that you experiment with it on non-production workflows only. The Azure Preview Supplemental Terms include additional legal terms that apply to Azure features that are in beta, preview, or otherwise not yet released into general availability.
Pricing:	**Microsoft Defender for container registries** is billed as shown on the pricing page
Clouds:	✔ Commercial clouds ✘ National (Azure Government, Azure China 21Vianet)

Prerequisites

To scan your images as they're pushed by CI/CD workflows into your registries, you must have **Microsoft Defender for container registries** enabled on the subscription.

Set up vulnerability scanning of your CI/CD workflows

To enable vulnerability scans of images in your GitHub workflows:

Step 1. Enable the CI/CD integration in Defender for Cloud

Step 2. Add the necessary lines to your GitHub workflow

Step 1. Enable the CI/CD integration in Defender for Cloud

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the relevant subscription.
3. From the sidebar of the settings page for that subscription, select **Integrations**.
4. In the pane that appears, select an Application Insights account to push the CI/CD scan results from your workflow.
5. Copy the authentication token and connection string into your GitHub workflow.

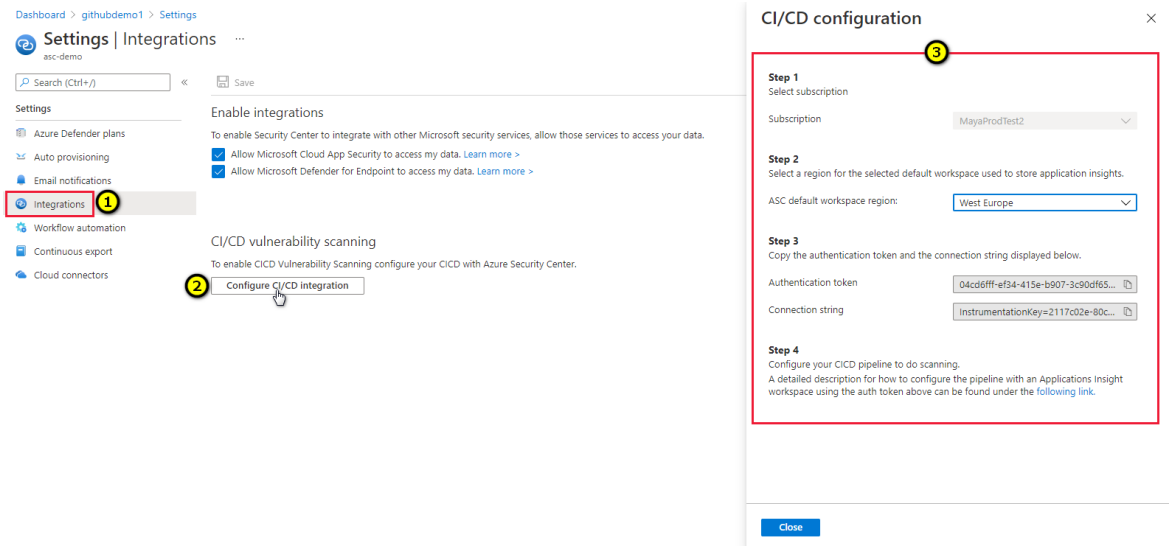

> **IMPORTANT**
>
> The authentication token and connection string are used to correlate the ingested security telemetry with resources in the subscription. If you use invalid values for these parameters, it'll lead to dropped telemetry.

Step 2. Add the necessary lines to your GitHub workflow and perform a scan

1. From your GitHub workflow, enable CI/CD scanning as follows:

> **TIP**
>
> We recommend creating two secrets in your repository to reference in your YAML file as shown below. The secrets can be named according to your own naming conventions. In this example, the secrets are referenced as **AZ_APPINSIGHTS_CONNECTION_STRING** and **AZ_SUBSCRIPTION_TOKEN**.

> **IMPORTANT**
>
> The push to the registry must happen prior to the results being published.

```
- run: |
  echo "github.sha=$GITHUB_SHA"
  docker build -t githubdemo1.azurecr.io/k8sdemo:${{ github.sha }}

- uses: Azure/container-scan@v0
  name: Scan image for vulnerabilities
  id: container-scan
  continue-on-error: true
  with:
    image-name: githubdemo1.azurecr.io/k8sdemo:${{ github.sha }}

- name: Push Docker image - githubdemo1.azurecr.io/k8sdemo:${{ github.sha }}
  run: |
  docker push githubdemo1.azurecr.io/k8sdemo:${{ github.sha }}

- name: Post logs to appinsights
  uses: Azure/publish-security-assessments@v0
  with:
    scan-results-path: ${{ steps.container-scan.outputs.scan-report-path }}
    connection-string: ${{ secrets.AZ_APPINSIGHTS_CONNECTION_STRING }}
    subscription-token: ${{ secrets.AZ_SUBSCRIPTION_TOKEN }}
```

2. Run the workflow that will push the image to the selected container registry. Once the image is pushed into the registry, a scan of the registry runs and you can view the CI/CD scan results along with the registry scan results within Microsoft Defender for Cloud.

3. View CI/CD scan results.

View CI/CD scan results

1. To view the findings, open the **Recommendations** page. If issues were found, you'll see the recommendation Container registry images should have vulnerability findings resolved (powered by Qualys).

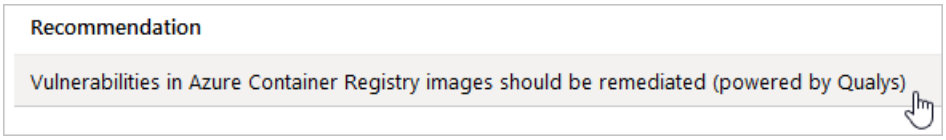

2. Select the recommendation.

   The recommendation details page opens with additional information. This information includes the list of registries with vulnerable images ("Affected resources") and the remediation steps.

3. Open the **affected resources** list and select an unhealthy registry to see the repositories within it that have vulnerable images.

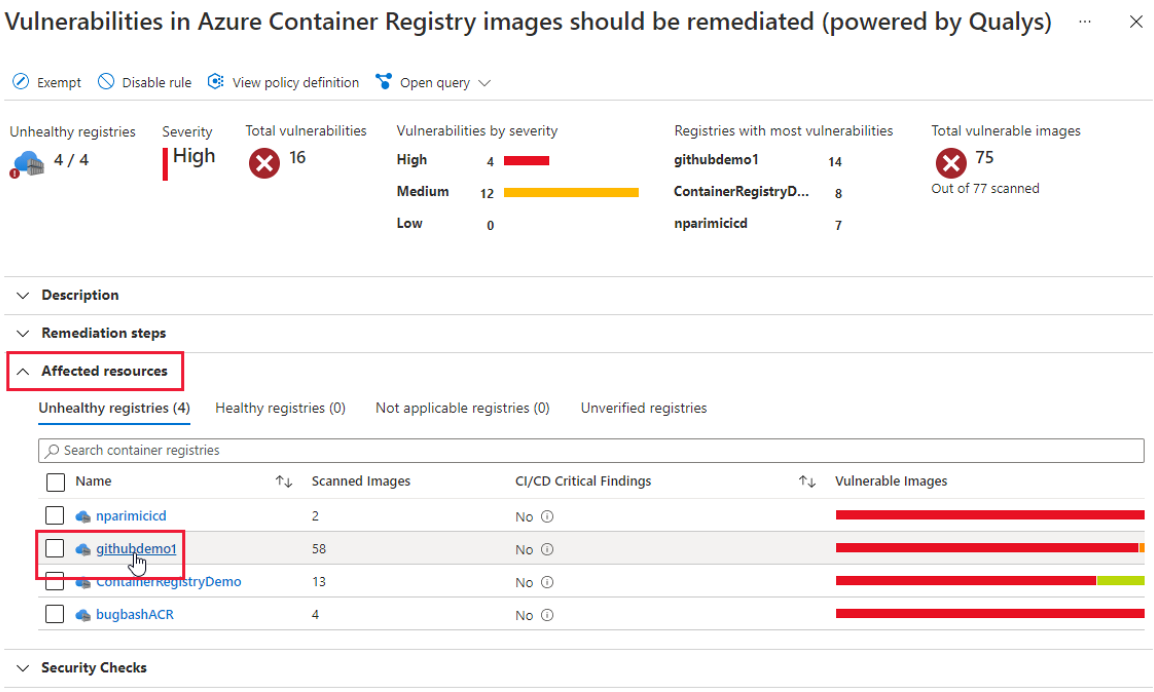

The registry details page opens with the list of affected repositories.

4. Select a specific repository to see the repositories within it that have vulnerable images.

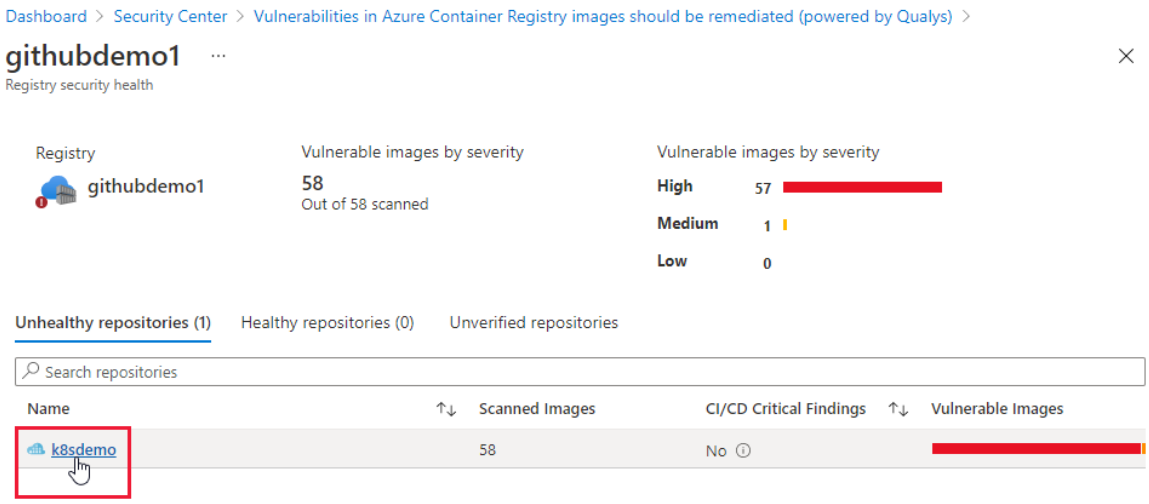

The repository details page opens. It lists the vulnerable images together with an assessment of the severity of the findings.

5. Select a specific image to see the vulnerabilities.

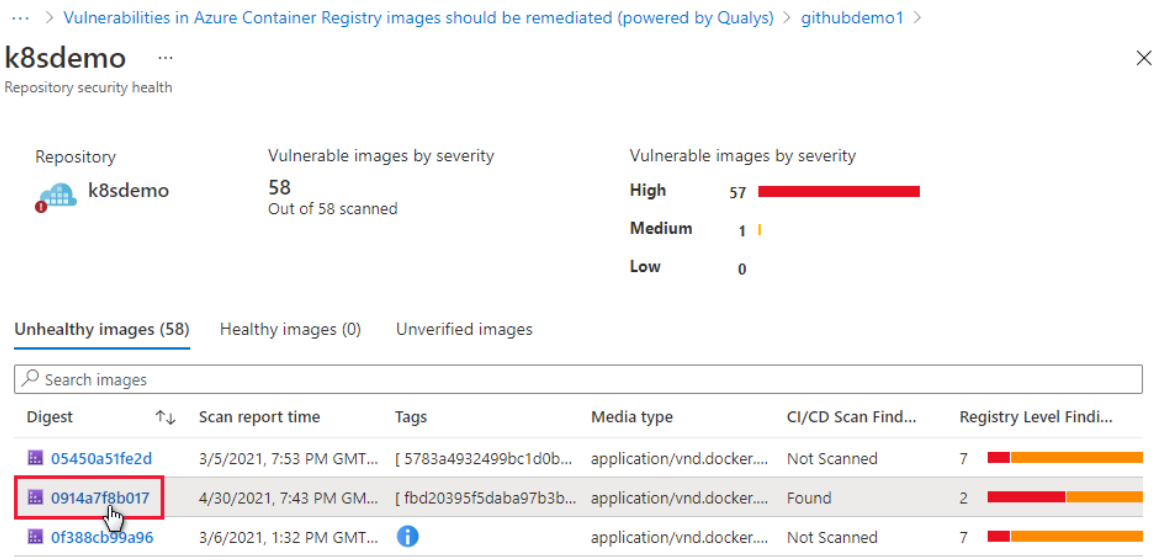

The list of findings for the selected image opens.

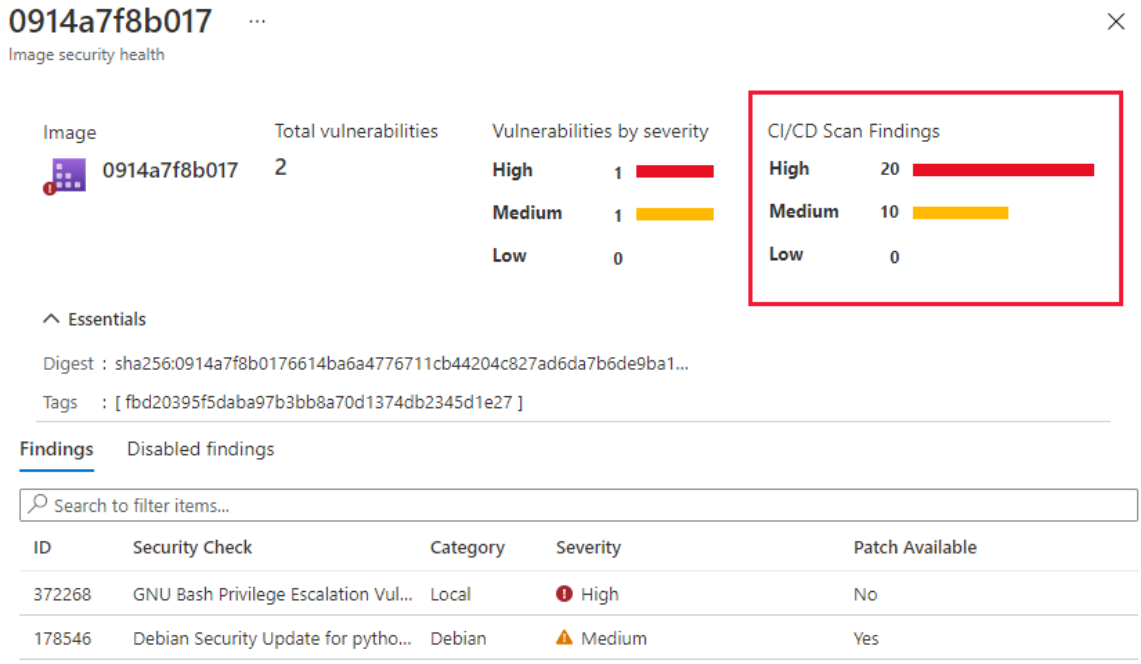

6. To learn more about which GitHub workflow is pushing these vulnerable images, select the information bubble:

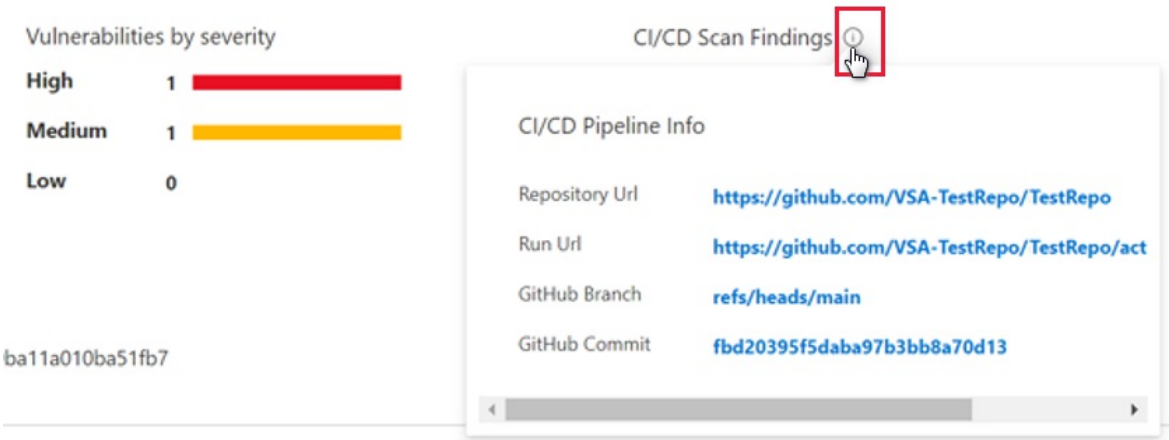

Next steps

Learn more about the advanced protection plans of Microsoft Defender for Cloud.

Protect your Kubernetes workloads

2/15/2022 • 6 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page describes how to use Microsoft Defender for Cloud's set of security recommendations dedicated to Kubernetes workload protection.

> **TIP**
> For a list of the security recommendations that might appear for Kubernetes clusters and nodes, see the Container recommendations of the recommendations reference table.

Availability

ASPECT	DETAILS
Release state:	AKS - General availability (GA) Arc enabled Kubernetes - Preview
Pricing:	Free for AKS workloads For Azure Arc-enabled Kubernetes, it's billed according to the Microsoft Defender for Containers plan
Required roles and permissions:	**Owner** or **Security admin** to edit an assignment **Reader** to view the recommendations
Environment requirements:	Kubernetes v1.14 (or newer) is required No PodSecurityPolicy resource (old PSP model) on the clusters Windows nodes are not supported
Azure Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)
Non-Azure Clouds, and On-prem:	supported via Arc enabled Kubernetes.

Set up your workload protection

Microsoft Defender for Cloud includes a bundle of recommendations that are available once you've installed the **Azure Policy add-on for Kubernetes or extensions**.

Prerequisites

- Add the Required FQDN/application rules for Azure policy.
- (For non AKS clusters) Connect an existing Kubernetes cluster to Azure Arc.

Enable Kubernetes workload protection

When you enable Microsoft Defender for Containers, Azure Kubernetes Service clusters, and Azure Arc enabled Kubernetes clusters (Preview) protection are both enabled by default. You can configure your Kubernetes workload protections, when you enable Microsoft Defender for Containers.

To enable Azure Kubernetes Service clusters and Azure Arc enabled Kubernetes clusters (Preview):

1. Sign in to the Azure portal.
2. Navigate to **Microsoft Defender for Cloud** > **Environment settings**.
3. Select the relevant subscription.
4. On the Defender plans page, ensure that Containers is toggled to **On**.
5. Select **Configure**.

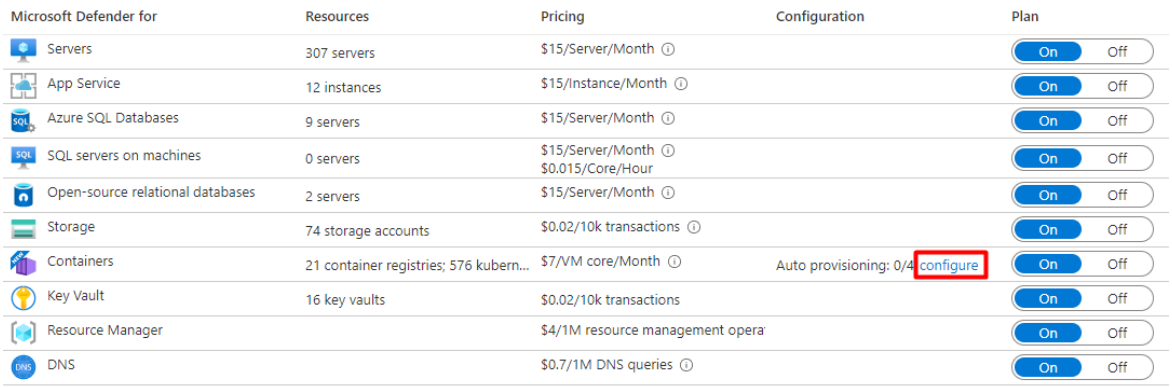

6. On the Advanced configuration page, toggle each relevant component to **On**.

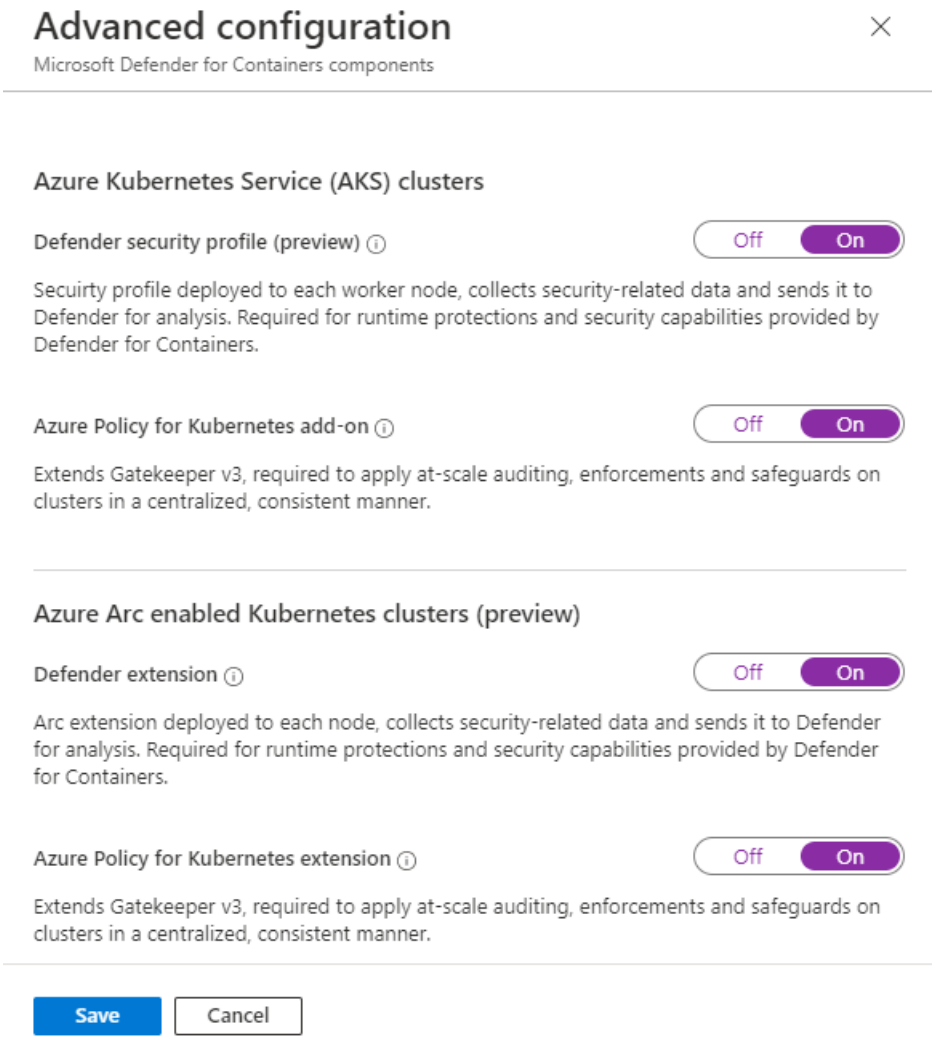

7. Select **Save**.

Configure Defender for Containers components

If you disabled any of the default protections when you enabled Microsoft Defender for Containers, you can change the configurations and reenable them via auto provisioning.

To configure the Defender for Containers components:

1. Sign in to the Azure portal.
2. Navigate to **Microsoft Defender for Cloud** > **Environment settings**.
3. Select the relevant subscription.
4. From the left side tool bar, select **Auto provisioning**.
5. Ensure that Microsoft Defenders for Containers components (preview) is toggled to On.

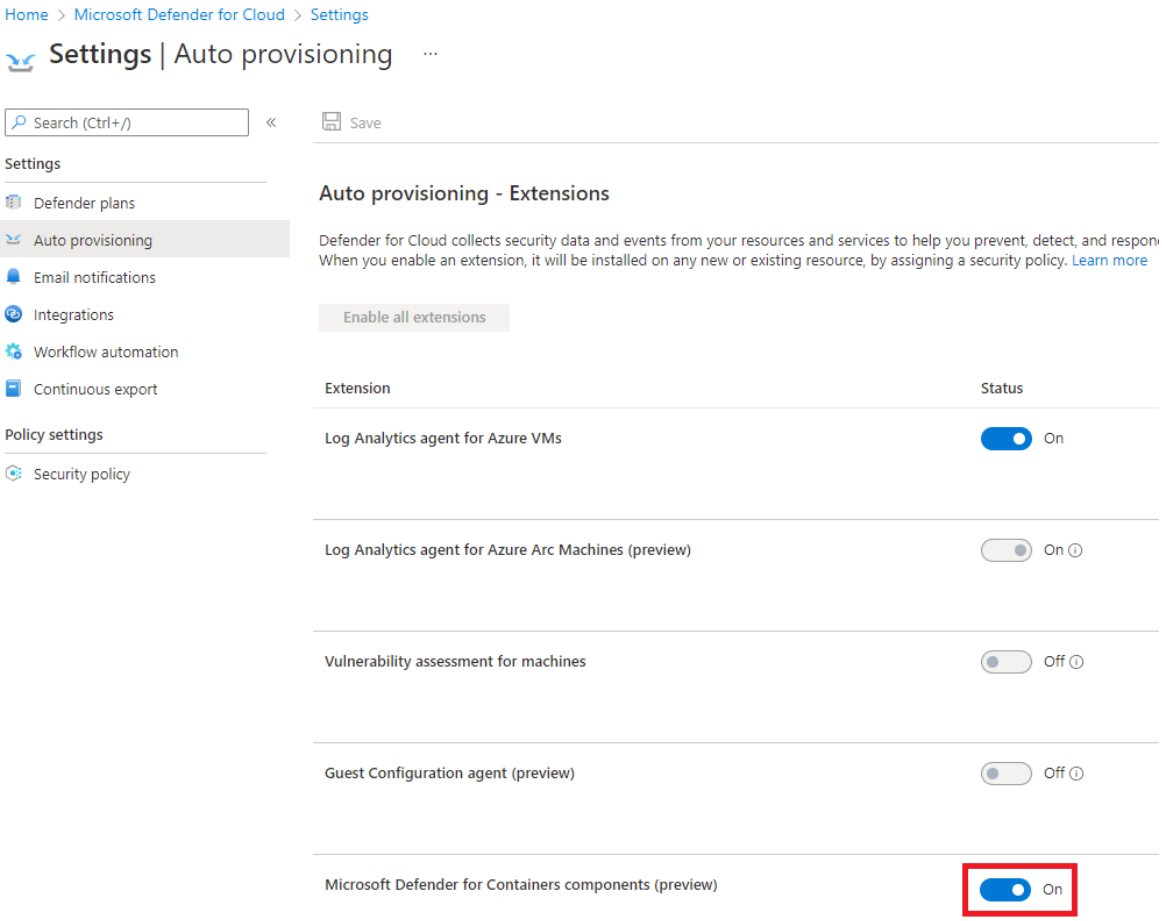

6. Select **Edit configuration**.

7. On the Advanced configuration page, toggle each relevant component to **On**.

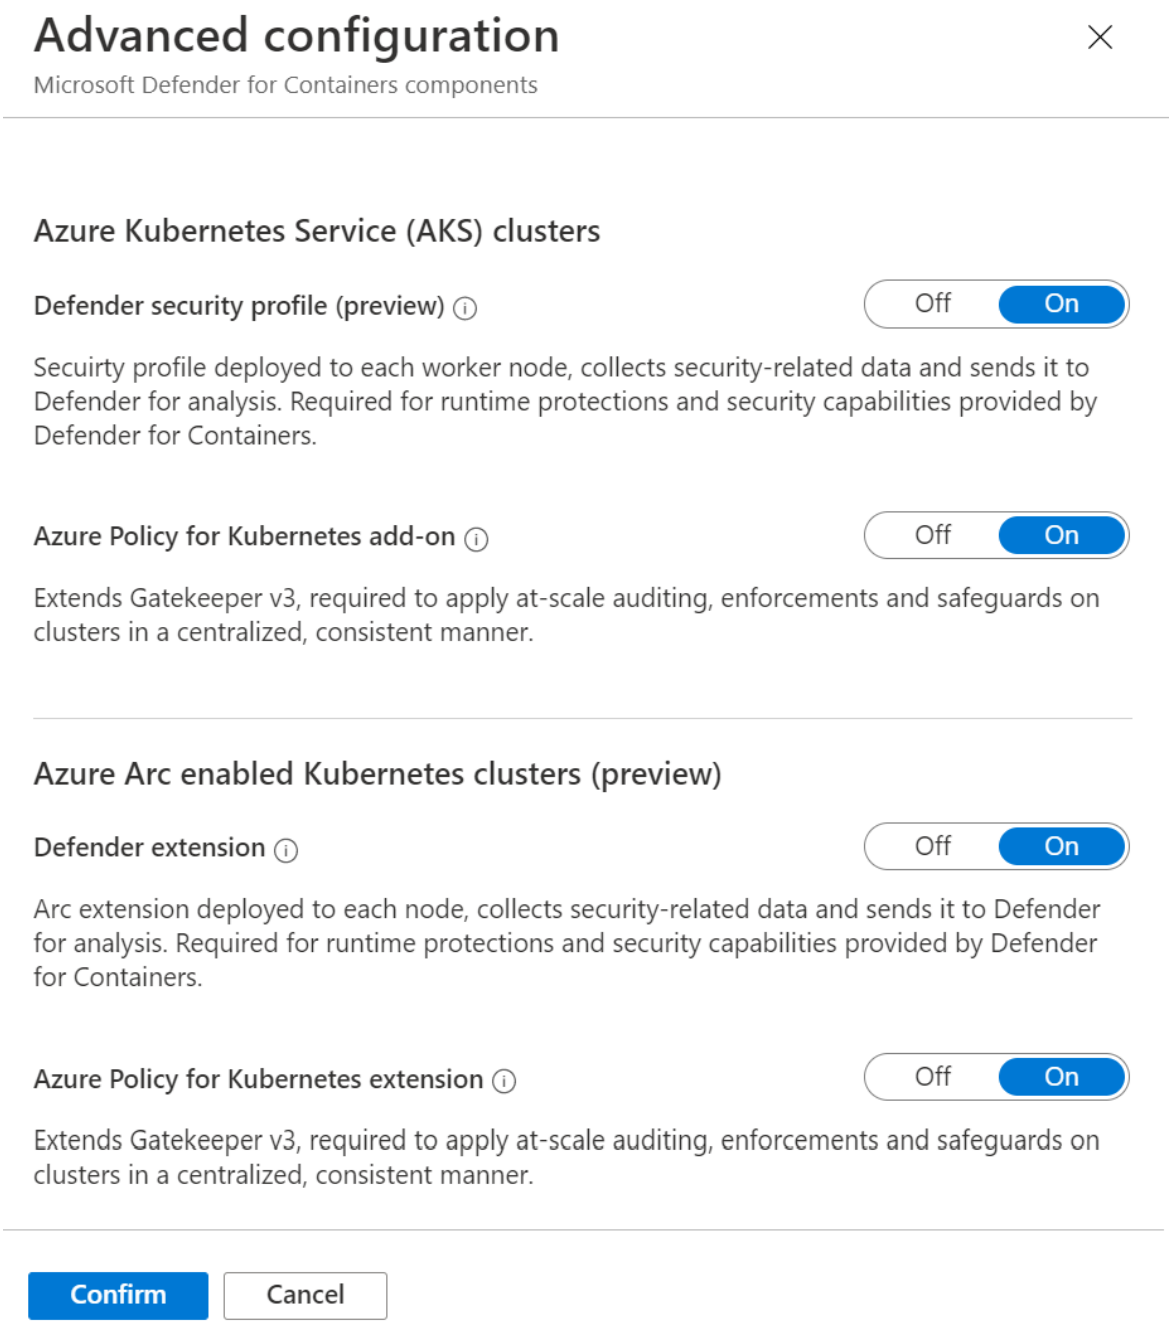

8. Select **Confirm**.

Deploy the add-on to specified clusters

You can manually configure the Kubernetes workload add-on, or extension protection through the Recommendations page. This can be accomplished by remediating the `Azure Policy add-on for Kubernetes should be installed and enabled on your clusters` recommendation.

To Deploy the add-on to specified clusters:

1. From the recommendations page, search for the recommendation `Azure Policy add-on for Kubernetes should be installed and enabled on your clusters`.

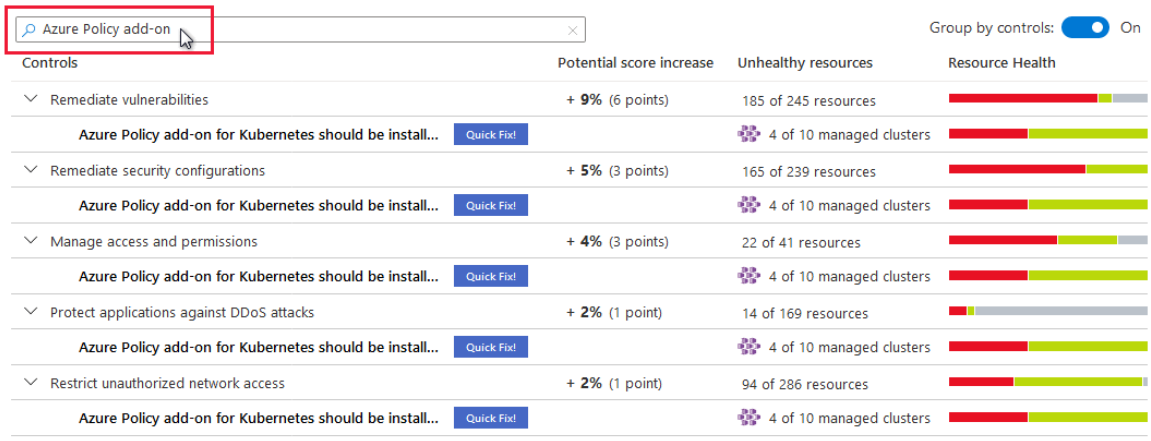

TIP

The recommendation is included in five different security controls and it doesn't matter which one you select in the next step.

2. From any of the security controls, select the recommendation to see the resources on which you can install the add-on.

3. Select the relevant cluster, and **Remediate**.

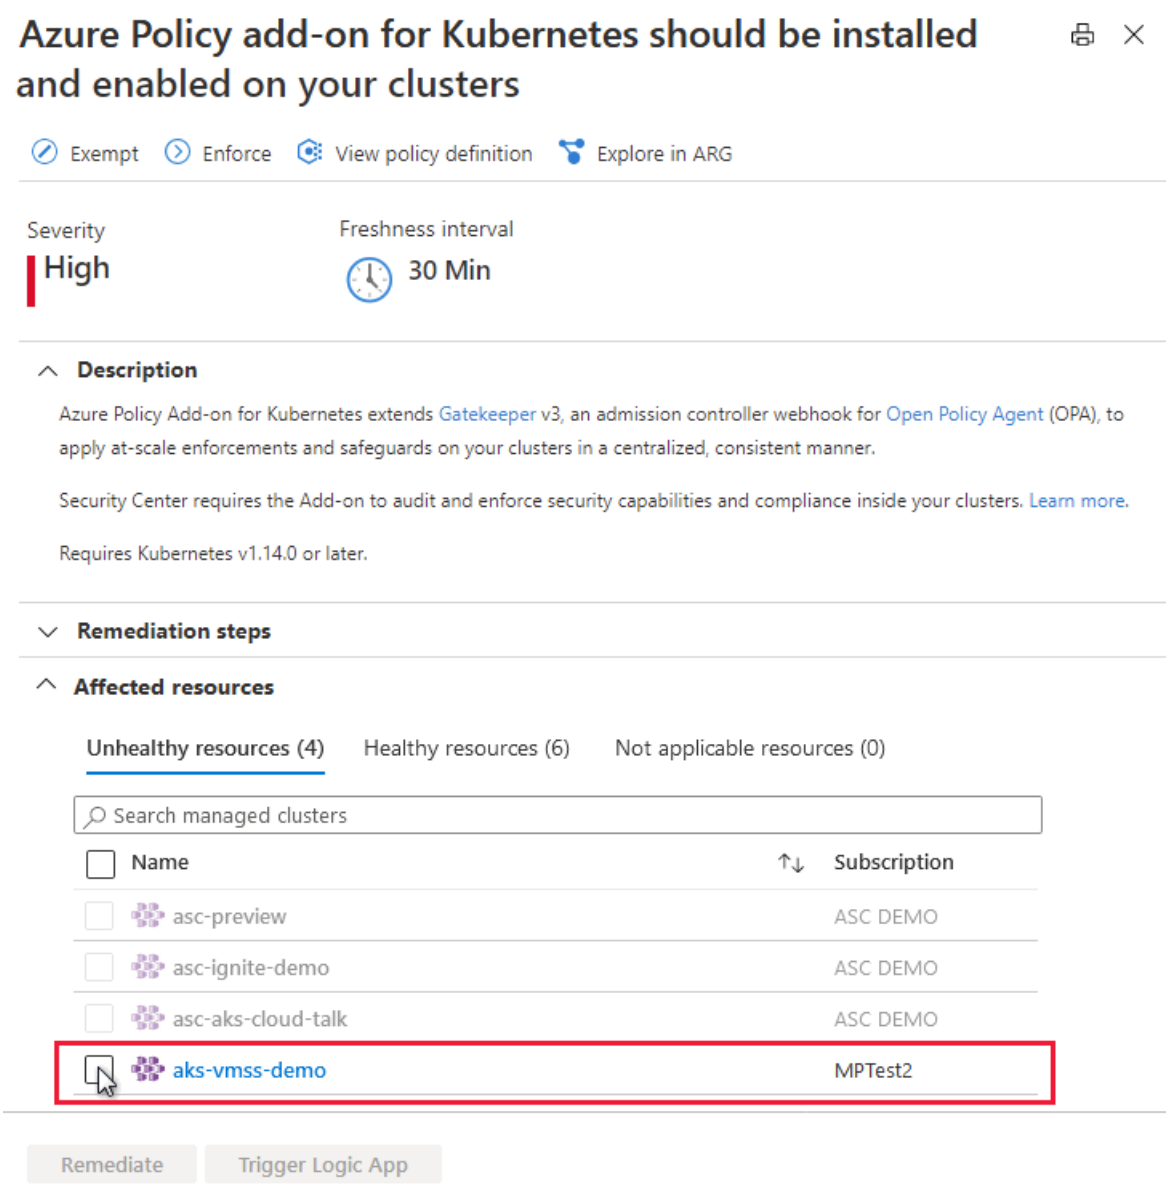

View and configure the bundle of recommendations

1. Approximately 30 minutes after the add-on installation completes, Defender for Cloud shows the clusters' health status for the following recommendations, each in the relevant security control as shown:

> **NOTE**
>
> If you're installing the add-on for the first time, these recommendations will appear as new additions in the list of recommendations.

> **TIP**
>
> Some recommendations have parameters that must be customized via Azure Policy to use them effectively. For example, to benefit from the recommendation **Container images should bedeployedonlyfrom trusted registries**, you'll have to define your trusted registries.
>
> If you don't enter the necessary parameters for the recommendations that require configuration, your workloads will be shown as unhealthy.

RECOMMENDATION NAME	SECURITY CONTROL	CONFIGURATION REQUIRED
Container CPU and memory limits should be enforced	Protect applications against DDoSattack	**Yes**
Container images should bedeployedonlyfrom trusted registries	Remediatevulnerabilities	**Yes**
Containers should listen on allowed ports only	Restrict unauthorized networkaccess	**Yes**
Least privilegedLinux capabilitiesshould be enforced for containers	Manage access andpermissions	**Yes**
Overriding or disabling of containers AppArmor profile should be restricted	Remediate security configurations	**Yes**
Services should listen on allowed ports only	Restrict unauthorized networkaccess	**Yes**
Usage of host networking and ports should be restricted	Restrict unauthorized network access	**Yes**
Usage of pod HostPath volume mounts should be restricted to a known list	Manage access and permissions	**Yes**
Container with privilege escalation should be avoided	Manage access andpermissions	No
Containers sharing sensitive host namespaces should be avoided	Manage access and permissions	No

RECOMMENDATION NAME	SECURITY CONTROL	CONFIGURATION REQUIRED
Immutable (read-only) root filesystem should be enforced for containers	Manage access andpermissions	No
Kubernetes clusters should be accessible only over HTTPS	Encrypt data in transit	No
Kubernetes clusters should disable automounting API credentials	Manage access andpermissions	No
Kubernetes clusters should not use the default namespace	Implement security best practices	No
Privileged containers should be avoided	Manage access andpermissions	No
Running containers as root user should be avoided	Manage access andpermissions	No

For recommendations with parameters that need to be customized, you will need to set the parameters:

To set the parameters:

1. Sign in to the Azure portal.
2. Navigate to **Microsoft Defender for Cloud** > **Environment settings**.
3. Select the relevant subscription.
4. From Defender for Cloud's menu, select **Security policy**.
5. Select the relevant assignment. The default assignment is `ASC default`.
6. Open the **Parameters** tab and modify the values as required.

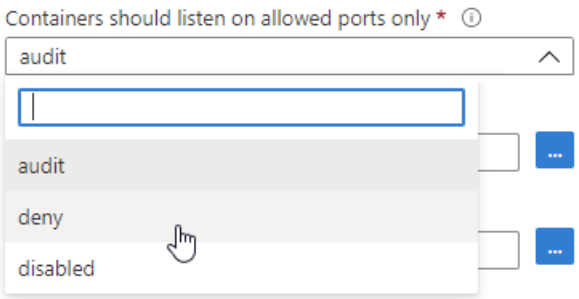

7. Select **Review + save**.
8. Select **Save**.

To enforce any of the recommendations:

1. Open the recommendation details page and select **Deny**:

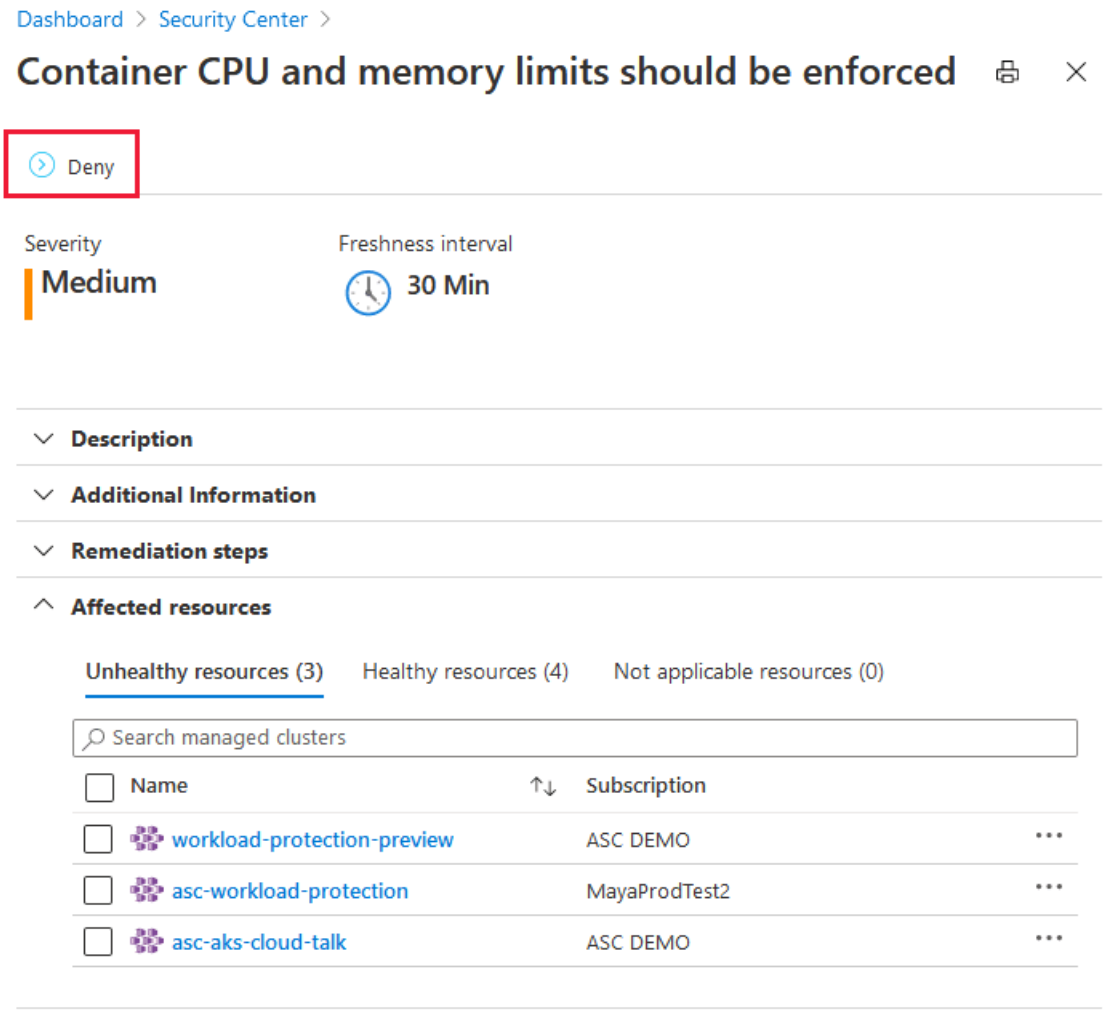

This will open the pane where you set the scope.

2. When you've set the scope, select **Change to deny**.

To see which recommendations apply to your clusters:

1. Open Defender for Cloud's asset inventory page and use the resource type filter to **Kubernetes services**.

2. Select a cluster to investigate and review the available recommendations available for it.

When viewing a recommendation from the workload protection set, you'll see the number of affected pods ("Kubernetes components") listed alongside the cluster. For a list of the specific pods, select the cluster and then select **Take action**.

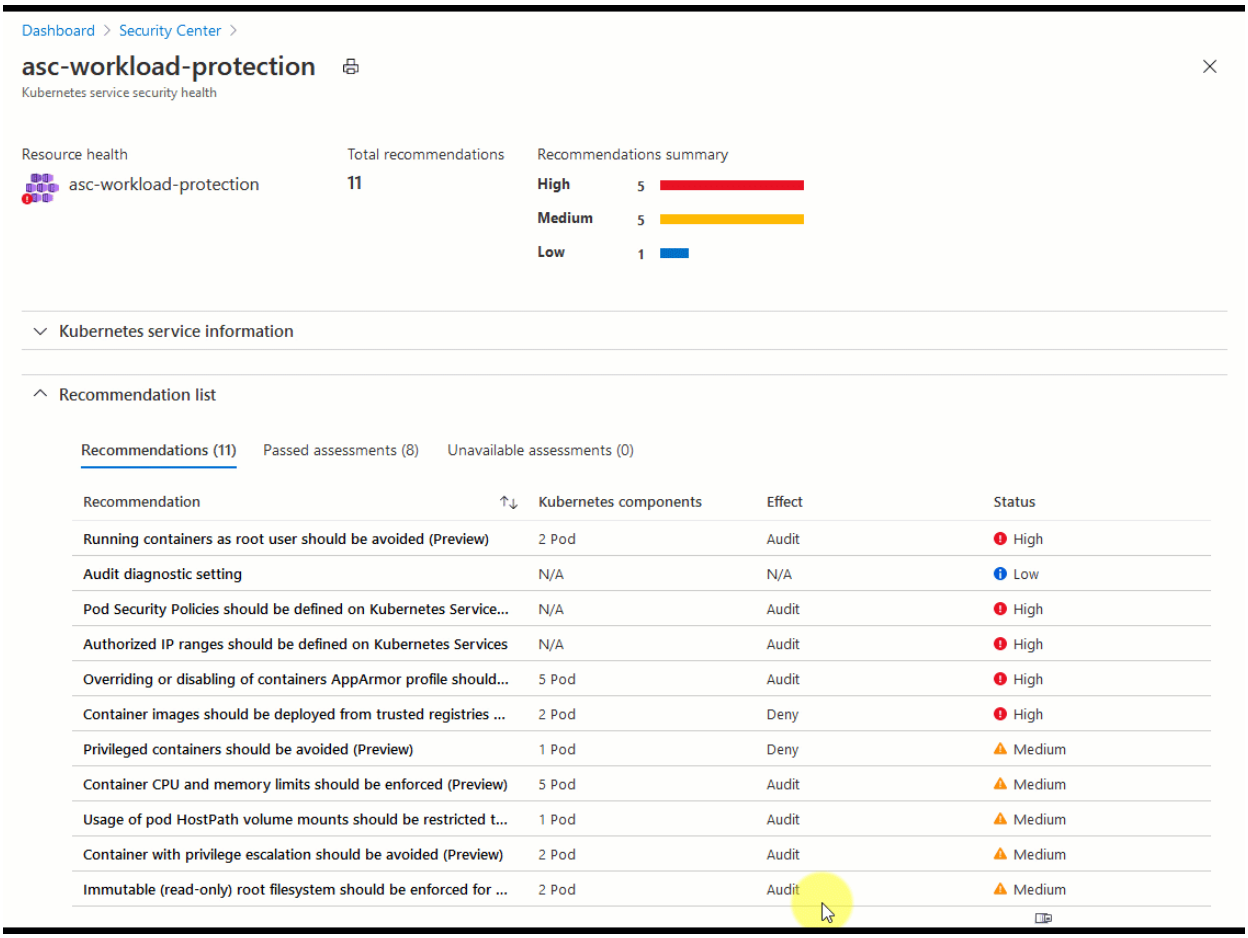

To test the enforcement, use the two Kubernetes deployments below:

- One is for a healthy deployment, compliant with the bundle of workload protection recommendations.
- The other is for an unhealthy deployment, non-compliant with *any* of the recommendations.

Deploy the example .yaml files as-is, or use them as a reference to remediate your own workload (step VIII).

Healthy deployment example .yaml file

```yaml
apiVersion: apps/v1
kind: Deployment
metadata:
  name: redis-healthy-deployment
  labels:
    app: redis
spec:
  replicas: 3
  selector:
    matchLabels:
      app: redis
  template:
    metadata:
      labels:
        app: redis
      annotations:
        container.apparmor.security.beta.kubernetes.io/redis: runtime/default
    spec:
      containers:
      - name: redis
        image: <customer-registry>.azurecr.io/redis:latest
        ports:
        - containerPort: 80
        resources:
          limits:
            cpu: 100m
            memory: 250Mi
        securityContext:
          privileged: false
          readOnlyRootFilesystem: true
          allowPrivilegeEscalation: false
          runAsNonRoot: true
          runAsUser: 1000
---
apiVersion: v1
kind: Service
metadata:
  name: redis-healthy-service
spec:
  type: LoadBalancer
  selector:
    app: redis
  ports:
  - port: 80
    targetPort: 80
```

Unhealthy deployment example .yaml file

```
apiVersion: apps/v1
kind: Deployment
metadata:
  name: redis-unhealthy-deployment
  labels:
    app: redis
spec:
  replicas: 3
  selector:
    matchLabels:
      app: redis
  template:
    metadata:
      labels:
        app: redis
    spec:
      hostNetwork: true
      hostPID: true
      hostIPC: true
      containers:
      - name: redis
        image: redis:latest
        ports:
        - containerPort: 9001
          hostPort: 9001
        securityContext:
          privileged: true
          readOnlyRootFilesystem: false
          allowPrivilegeEscalation: true
          runAsUser: 0
          capabilities:
            add:
              - NET_ADMIN
        volumeMounts:
        - mountPath: /test-pd
          name: test-volume
          readOnly: true
      volumes:
      - name: test-volume
        hostPath:
          # directory location on host
          path: /tmp
---
apiVersion: v1
kind: Service
metadata:
  name: redis-unhealthy-service
spec:
  type: LoadBalancer
  selector:
    app: redis
  ports:
  - port: 6001
    targetPort: 9001
```

Next steps

In this article, you learned how to configure Kubernetes workload protection.

For other related material, see the following pages:

- Defender for Cloud recommendations for compute
- Alerts for AKS cluster level

Protect your web apps and APIs

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Prerequisites

Defender for Cloud is natively integrated with App Service, eliminating the need for deployment and onboarding - the integration is transparent.

To protect your Azure App Service plan with Microsoft Defender for App Service, you'll need:

- A supported App Service plan associated with dedicated machines. Supported plans are listed in Availability.
- Defender for Cloud's enhanced protections enabled on your subscription as described in Quickstart: Enable enhanced security features.

> **TIP**
> You can optionally enable individual Microsoft Defender plans, like Microsoft Defender for App Service.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Microsoft Defender for App Service is billed as shown on the pricing page Billing is according to total compute instances in all plans
Supported App Service plans:	All App Service plans are supported except Azure Functions on the consumption plan.
Clouds:	✔ Commercial clouds ✘ National (Azure Government, Azure China 21Vianet)

What are the benefits of Microsoft Defender for App Service?

Azure App Service is a fully managed platform for building and hosting your web apps and APIs. Since the platform is fully managed, you don't have to worry about the infrastructure. It provides management, monitoring, and operational insights to meet enterprise-grade performance, security, and compliance requirements. For more information, see Azure App Service.

Microsoft Defender for App Service uses the scale of the cloud to identify attacks targeting applications running over App Service. Attackers probe web applications to find and exploit weaknesses. Before being routed to specific environments, requests to applications running in Azure go through several gateways, where they're inspected and logged. This data is then used to identify exploits and attackers, and to learn new patterns that will be used later.

When you enable Microsoft Defender for App Service, you immediately benefit from the following services offered by this Defender plan:

- **Secure** - Defender for App Service assesses the resources covered by your App Service plan and generates security recommendations based on its findings. Use the detailed instructions in these recommendations to harden your App Service resources.
- **Detect** - Defender for App Service detects a multitude of threats to your App Service resources by monitoring:
  - the VM instance in which your App Service is running, and its management interface
  - the requests and responses sent to and from your App Service apps
  - the underlying sandboxes and VMs
  - App Service internal logs - available thanks to the visibility that Azure has as a cloud provider

As a cloud-native solution, Defender for App Service can identify attack methodologies applying to multiple targets. For example, from a single host it would be difficult to identify a distributed attack from a small subset of IPs, crawling to similar endpoints on multiple hosts.

The log data and the infrastructure together can tell the story: from a new attack circulating in the wild to compromises in customer machines. Therefore, even if Microsoft Defender for App Service is deployed after a web app has been exploited, it might be able to detect ongoing attacks.

What threats can Defender for App Service detect?

Threats by MITRE ATT&CK tactics

Defender for Cloud monitors for many threats to your App Service resources. The alerts cover almost the complete list of MITRE ATT&CK tactics from pre-attack to command and control.

- **Pre-attack threats** - Defender for Cloud can detect the execution of multiple types of vulnerability scanners that attackers frequently use to probe applications for weaknesses.
- **Initial access threats** - Microsoft Threat Intelligence powers these alerts that include triggering an alert when a known malicious IP address connects to your Azure App Service FTP interface.
- **Execution threats** - Defender for Cloud can detect attempts to run high privilege commands, Linux commands on a Windows App Service, fileless attack behavior, digital currency mining tools, and many other suspicious and malicious code execution activities.

Dangling DNS detection

Defender for App Service also identifies any DNS entries remaining in your DNS registrar when an App Service website is decommissioned - these are known as dangling DNS entries. When you remove a website and don't remove its custom domain from your DNS registrar, the DNS entry is pointing at a non-existent resource and your subdomain is vulnerable to a takeover. Defender for Cloud doesn't scan your DNS registrar for *existing* dangling DNS entries; it alerts you when an App Service website is decommissioned and its custom domain (DNS entry) isn't deleted.

Subdomain takeovers are a common, high-severity threat for organizations. When a threat actor detects a dangling DNS entry, they create their own site at the destination address. The traffic intended for the organization's domain is then directed to the threat actor's site, and they can use that traffic for a wide range of

malicious activity.

Dangling DNS protection is available whether your domains are managed with Azure DNS or an external domain registrar and applies to App Service on both Windows and Linux.

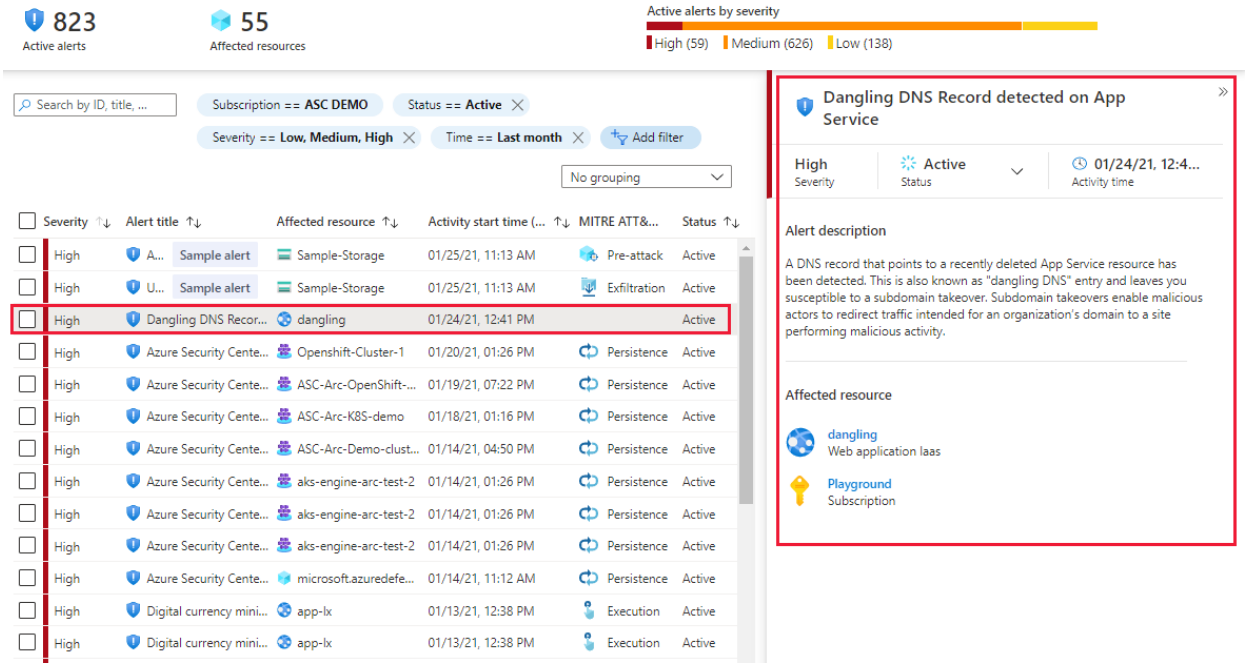

Learn more about dangling DNS and the threat of subdomain takeover, in Prevent dangling DNS entries and avoid subdomain takeover.

For a full list of the App Service alerts, see the Reference table of alerts.

> **NOTE**
>
> Defender for Cloud might not trigger dangling DNS alerts if your custom domain doesn't point directly to an App Service resource, or if Defender for Cloud hasn't monitored traffic to your website since the dangling DNS protection was enabled (because there won't be logs to help identify the custom domain).

Next steps

In this article, you learned about Microsoft Defender for App Service.

Enable enhanced protections

For related material, see the following articles:

- To export your alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Stream alerts to a SIEM, SOAR, or IT Service Management solution.
- For a list of the Microsoft Defender for App Service alerts, see the Reference table of alerts.
- For more information on App Service plans, see App Service plans.

Introduction to Microsoft Defender for Storage

2/15/2022 • 7 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Storage is an Azure-native layer of security intelligence that detects unusual and potentially harmful attempts to access or exploit your storage accounts. It uses advanced threat detection capabilities and Microsoft Threat Intelligence data to provide contextual security alerts. Those alerts also include steps to mitigate the detected threats and prevent future attacks.

You can enable **Microsoft Defender for Storage** at either the subscription level (recommended) or the resource level.

Defender for Storage continually analyzes the telemetry stream generated by the Azure Blob Storage and Azure Files services. When potentially malicious activities are detected, security alerts are generated. These alerts are displayed in Microsoft Defender for Cloud together with the details of the suspicious activity along with the relevant investigation steps, remediation actions, and security recommendations.

Analyzed telemetry of Azure Blob Storage includes operation types such as **Get Blob**, **Put Blob**, **Get Container ACL**, **List Blobs**, and **Get Blob Properties**. Examples of analyzed Azure Files operation types include **Get File**, **Create File**, **List Files**, **Get File Properties**, and **Put Range**.

Defender for Storage doesn't access the Storage account data and has no impact on its performance.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for Storage** is billed as shown on the pricing page
Protected storage types:	Blob Storage (Standard/Premium StorageV2, Block Blobs) Azure Files (over REST API and SMB) Azure Data Lake Storage Gen2 (Standard/Premium accounts with hierarchical namespaces enabled)
Clouds:	✅ Commercial clouds ✅ Azure Government ❌ Azure China 21Vianet ❌ Connected AWS accounts

What are the benefits of Microsoft Defender for Storage?

Defender for Storage provides:

- **Azure-native security** - With 1-click enablement, Defender for Storage protects data stored in Azure Blob, Azure Files, and Data Lakes. As an Azure-native service, Defender for Storage provides centralized security across all data assets that are managed by Azure and is integrated with other Azure security services such as Microsoft Sentinel.
- **Rich detection suite** - Powered by Microsoft Threat Intelligence, the detections in Defender for Storage cover the top storage threats such as unauthenticated access, compromised credentials, social engineering attacks, data exfiltration, privilege abuse, and malicious content.
- **Response at scale** - Defender for Cloud's automation tools make it easier to prevent and respond to identified threats. Learn more in Automate responses to Defender for Cloud triggers.

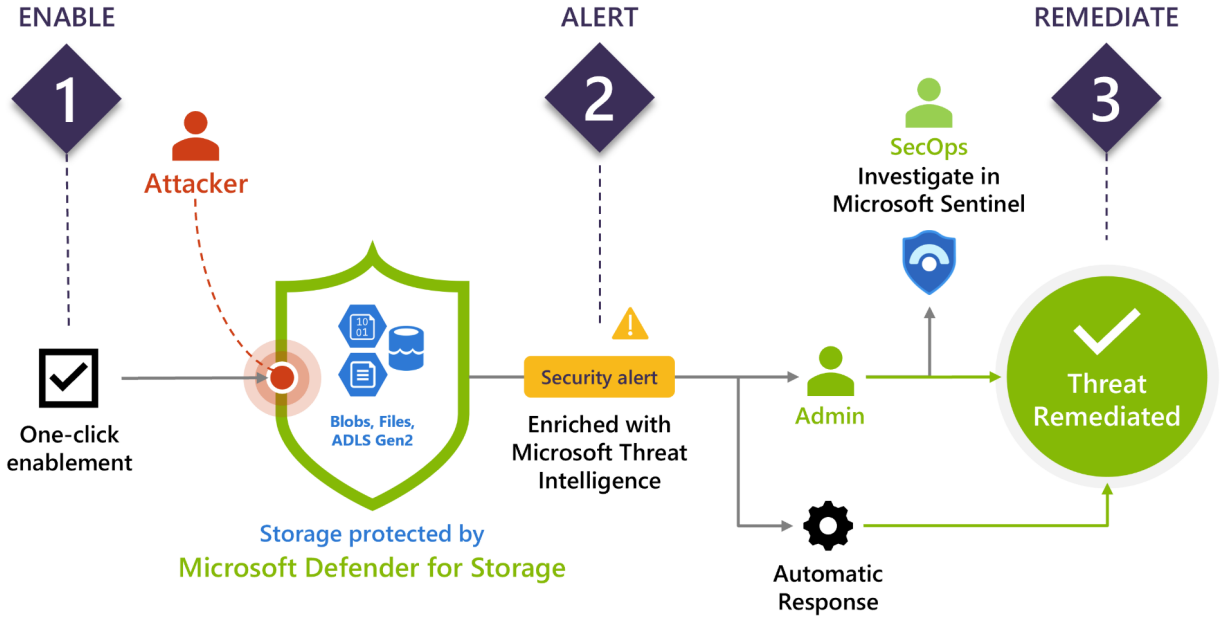

Security threats in cloud-based storage services

Microsoft security researchers have analyzed the attack surface of storage services. Storage accounts can be subject to data corruption, exposure of sensitive content, malicious content distribution, data exfiltration, unauthorized access, and more.

The potential security risks are described in the threat matrix for cloud-based storage services and are based on the MITRE ATT&CK® framework, a knowledge base for the tactics and techniques employed in cyber attacks.

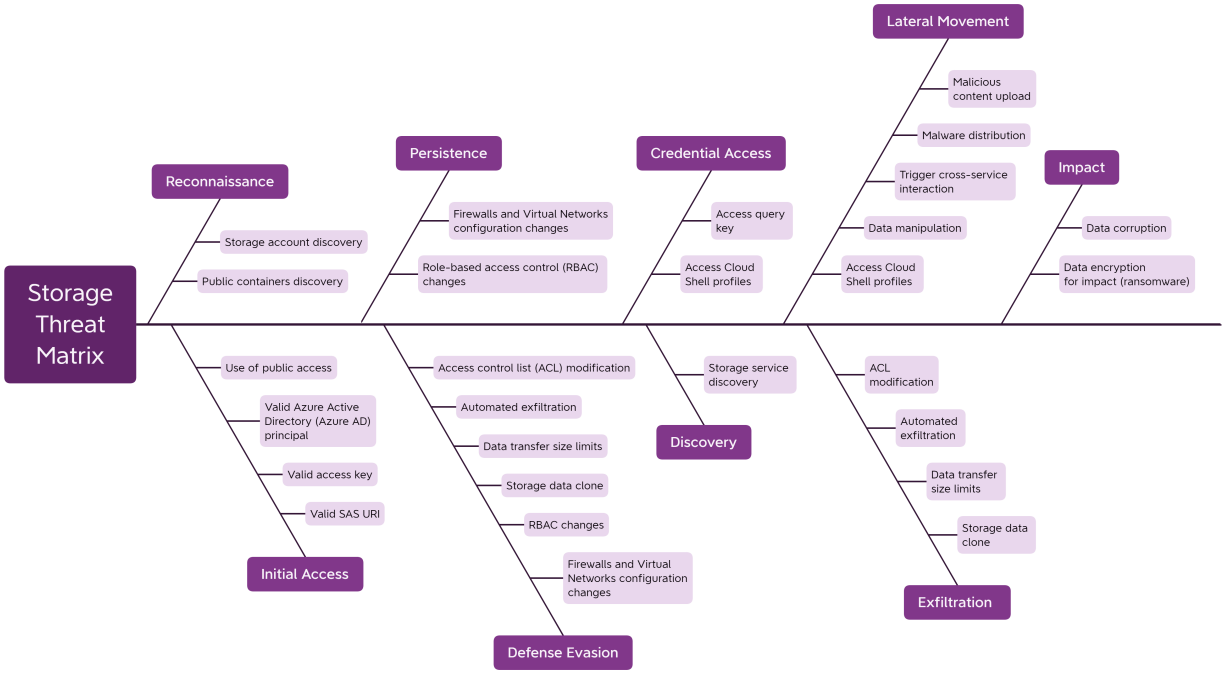

What kind of alerts does Microsoft Defender for Storage provide?

Security alerts are triggered for the following scenarios (typically from 1-2 hours after the event):

TYPE OF THREAT	DESCRIPTION
Unusual access to an account	For example, access from a TOR exit node, suspicious IP addresses, unusual applications, unusual locations, and anonymous access without authentication.
Unusual behavior in an account	Behavior that deviates from a learned baseline, such as a change of access permissions in an account, unusual access inspection, unusual data exploration, unusual deletion of blobs/files, or unusual data extraction.
Hash reputation based Malware detection	Detection of known malware based on full blob/file hash. This can help detect ransomware, viruses, spyware, and other malware uploaded to an account, prevent it from entering the organization, and spreading to more users and resources. See also Limitations of hash reputation analysis.
Unusual file uploads	Unusual cloud service packages and executable files that have been uploaded to an account.
Public visibility	Potential break-in attempts by scanning containers and pulling potentially sensitive data from publicly accessible containers.
Phishing campaigns	When content that's hosted on Azure Storage is identified as part of a phishing attack that's impacting Microsoft 365 users.

Alerts include details of the incident that triggered them, and recommendations on how to investigate and remediate threats. Alerts can be exported to Microsoft Sentinel or any other third-party SIEM or any other external tool. Learn more in Stream alerts to a SIEM, SOAR, or IT Service Management solution.

> **TIP**
> For a comprehensive list of all Defender for Storage alerts, see the alerts reference page. This is useful for workload owners who want to know what threats can be detected and help SOC teams gain familiarity with detections before investigating them. Learn more about what's in a Defender for Cloud security alert, and how to manage your alerts in Manage and respond to security alerts in Microsoft Defender for Cloud.

Limitations of hash reputation analysis

- **Hash reputation isn't deep file inspection** - Microsoft Defender for Storage uses hash reputation analysis supported by Microsoft Threat Intelligence to determine whether an uploaded file is suspicious. The threat protection tools don't scan the uploaded files; rather they analyze the telemetry generated from the Blobs Storage and Files services. Defender for Storage then compares the hashes of newly uploaded files with hashes of known viruses, trojans, spyware, and ransomware.
- **Hash reputation analysis isn't supported for all files protocols and operation types** - Some, but not all, of the telemetry logs contain the hash value of the related blob or file. In some cases, the telemetry doesn't contain a hash value. As a result, some operations can't be monitored for known malware uploads. Examples of such unsupported use cases include SMB file-shares and when a blob is created using Put Block and Put Block List.

> **TIP**
> When a file is suspected to contain malware, Defender for Cloud displays an alert and can optionally email the storage owner for approval to delete the suspicious file. To set up this automatic removal of files that hash reputation analysis indicates contain malware, deploy a workflow automation to trigger on alerts that contain "Potential malware uploaded to a storage account".

Enable Defender for Storage

When you enable this Defender plan on a subscription, all existing Azure Storage accounts will be protected and any storage resources added to that subscription in the future will also be automatically protected.

You can enable Defender for Storage in any of several ways, described in Set up Microsoft Defender for Cloud in the Azure Storage documentation.

Trigger a test alert for Microsoft Defender for Storage

To test the security alerts from Microsoft Defender for Storage in your environment, generate the alert "Access from a Tor exit node to a storage account" with the following steps:

1. Open a storage account with Microsoft Defender for Storage enabled.
2. From the sidebar, select "Containers" and open an existing container or create a new one.

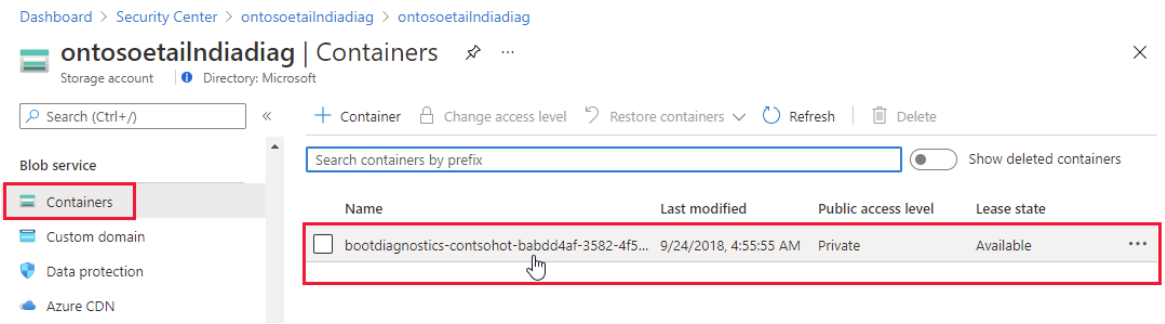

3. Upload a file to that container.

Caution

Don't upload a file containing sensitive data.

4. Use the context menu on the uploaded file to select "Generate SAS".

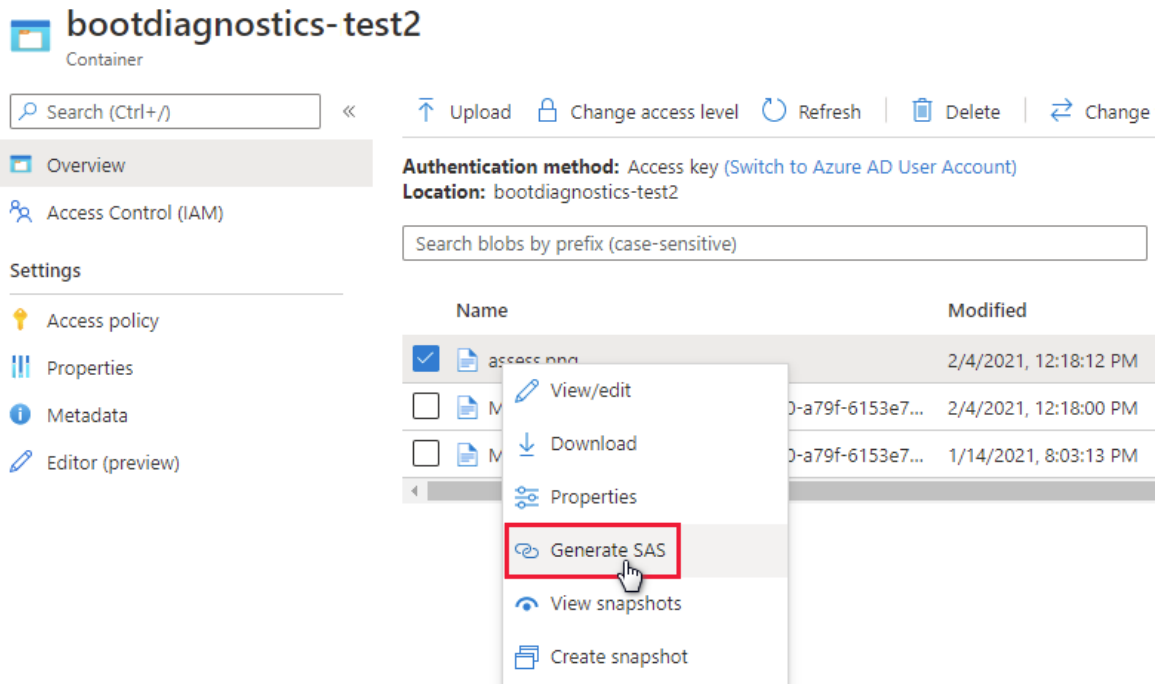

5. Leave the default options and select **Generate SAS token and URL**.

6. Copy the generated SAS URL.

7. On your local machine, open the Tor browser.

> **TIP**
>
> You can download Tor from the Tor Project site https://www.torproject.org/download/.

8. In the Tor browser, navigate to the SAS URL.

9. Download the file you uploaded in step 3.

   Within two hours you'll get the following security alert from Defender for Cloud:

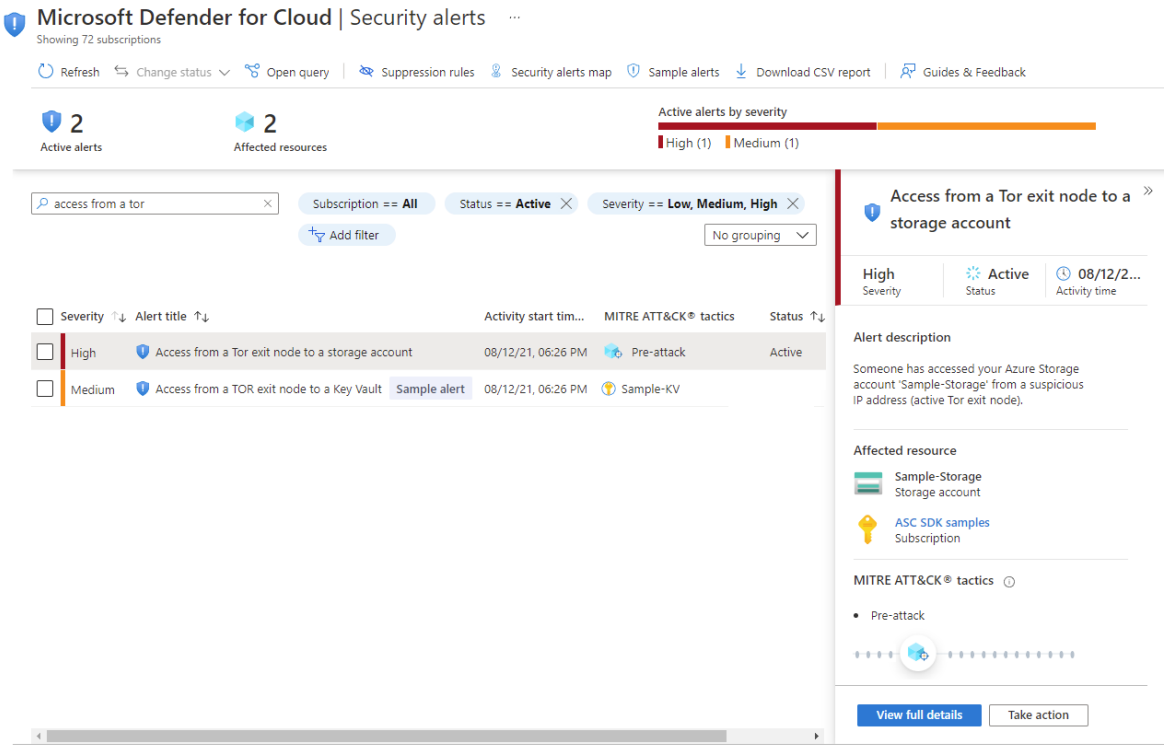

FAQ - Microsoft Defender for Storage

- How do I estimate charges at the account level?
- Can I exclude a specific Azure Storage account from a protected subscription?
- How do I configure automatic responses for security alerts?

How do I estimate charges at the account level?

To optimize costs, you might want to exclude specific Storage accounts associated with high traffic from Defender for Storage protections. To get an estimate of Defender for Storage costs, use the Price Estimation Dashboard.

Can I exclude a specific Azure Storage account from a protected subscription?

To exclude a specific Storage account when Defender for Storage is enabled on a subscription, follow the instructions in Exclude a storage account from Microsoft Defender for Storage protections.

How do I configure automatic responses for security alerts?

Use workflow automation to trigger automatic responses to Defender for Cloud security alerts.

For example, you can set up automation to open tasks or tickets for specific personnel or teams in an external task management system.

> **TIP**
>
> Explore the automations available from the Defender for Cloud community pages: ServiceNow automation, Jira automation, Azure DevOps automation, Slack automation or build your own.

Use automation for automatic response - to define your own or use ready-made automation from the community (such as removing malicious files upon detection). For more solutions, visit the Microsoft community on GitHub.

Next steps

In this article, you learned about Microsoft Defender for Storage.

Enable Defender for Storage

- The full list of Microsoft Defender for Storage alerts
- Stream alerts to a SIEM, SOAR, or IT Service Management solution
- Save Storage telemetry for investigation

Exclude a storage account from Microsoft Defender for Storage protections

2/15/2022 • 4 minutes to read • Edit Online

Caution

Excluding resources from advanced threat protection is not recommended and leaves your cloud workload exposed.

When you enable Microsoft Defender for Storage on a subscription, all existing Azure Storage accounts will be protected and any storage resources added to that subscription in the future will also be automatically protected.

If you need to exempt a specific Azure Storage account from this Defender plan, use the instructions on this page.

> **TIP**
> We recommend enabling Microsoft Defender for Resource Manager for any accounts with unprotected Azure Storage resources. Defender for Resource Manager automatically monitors your organization's resource management operations, whether they're performed through the Azure portal, Azure REST APIs, Azure CLI, or other Azure programmatic clients.

Exclude a specific storage account

To exclude specific storage accounts from Microsoft Defender for Storage when the plan is enabled on a subscription:

- PowerShell
- Azure CLI
- Azure portal

Use PowerShell to exclude an Azure Storage account

1. If you don't have the Azure Az PowerShell module installed, install it using the instructions from the Azure PowerShell documentation.
2. Using an authenticated account, connect to Azure with the `Connect-AzAccount` cmdlet, as explained in Sign in with Azure PowerShell.
3. Define the AzDefenderPlanAutoEnable tag on the storage account with the `Update-AzTag` cmdlet (replace the ResourceId with the resource ID of the relevant storage account):

   ```
   Update-AzTag -ResourceId <resourceID> -Tag @{"AzDefenderPlanAutoEnable" = "off"} -Operation Merge
   ```

   If you skip this stage, your untagged resources will continue receiving daily updates from the subscription level enablement policy. That policy will enable Defender for Storage again on the account.

   > **TIP**
   > Learn more about tags in Use tags to organize your Azure resources and management hierarchy.

4. Disable Microsoft Defender for Storage for the desired account on the relevant subscription with the `Disable-AzSecurityAdvancedThreatProtection` cmdlet (using the same resource ID):

```
Disable-AzSecurityAdvancedThreatProtection -ResourceId <resourceId>
```

Learn more about this cmdlet.

Exclude an Azure Databricks Storage account

Exclude an active Databricks workspace

Microsoft Defender for Storage can exclude specific active Databricks workspace storage accounts, when the plan is already enabled on a subscription.

To exclude an active Databricks workspace:

1. Sign in to the Azure portal.
2. Navigate to **Azure Databricks** > `Your Databricks workspace` > **Tags**.
3. In the Name field, enter `AzDefenderPlanAutoEnable`.
4. In the Value field, enter `off`.
5. Select **Apply**.

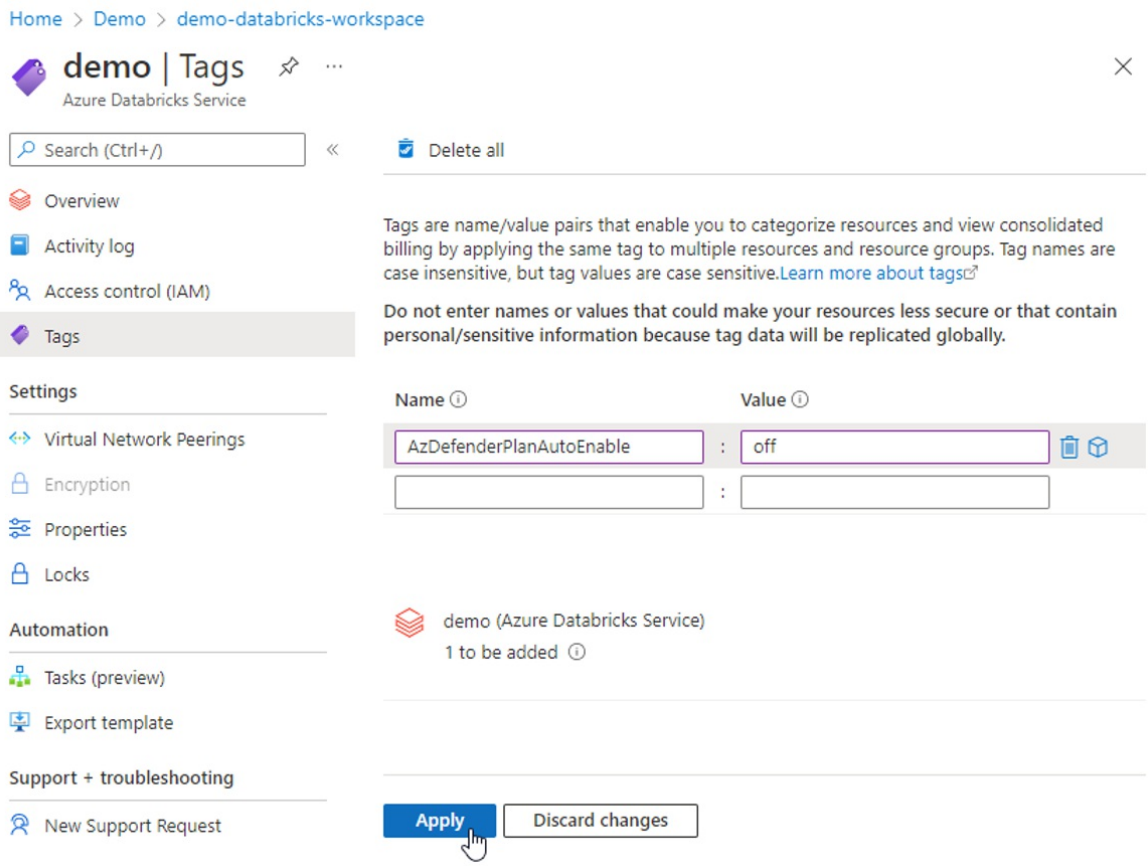

6. Navigate to **Microsoft Defender for Cloud** > **Environment settings** > `Your subscription`.
7. Toggle the Defender for Storage plan to **Off**.

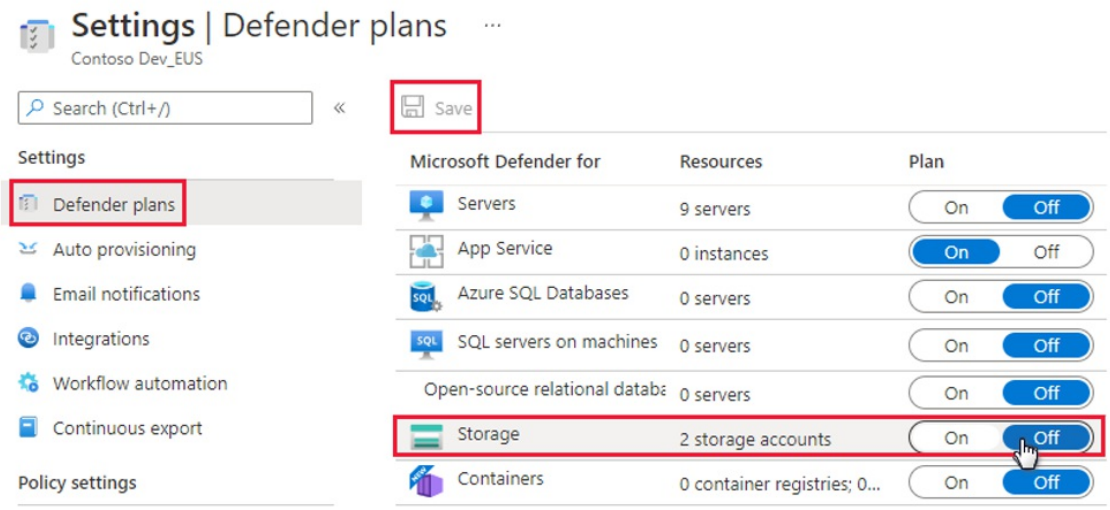

8. Select **Save**.
9. Toggle the Defender for Storage plan to **On**.
10. Select **Save**.

The tags will be inherited by the Storage account of the Databricks workspace and prevent Defender for Storage from turning on.

> **NOTE**
> Tags can't be added directly to the Databricks Storage account, or its Managed Resource Group.

Prevent auto-enabling on a new Databricks workspace storage account

When you create a new Databricks workspace, you have the ability to add a tag that will prevent your Microsoft Defender for Storage account from enabling automatically.

To prevent auto-enabling on a new Databricks workspace storage account:

1. Follow these steps to create a new Azure Databricks workspace.
2. In the Tags tab, enter a tag named `AzDefenderPlanAutoEnable`.
3. Enter the value `off`.

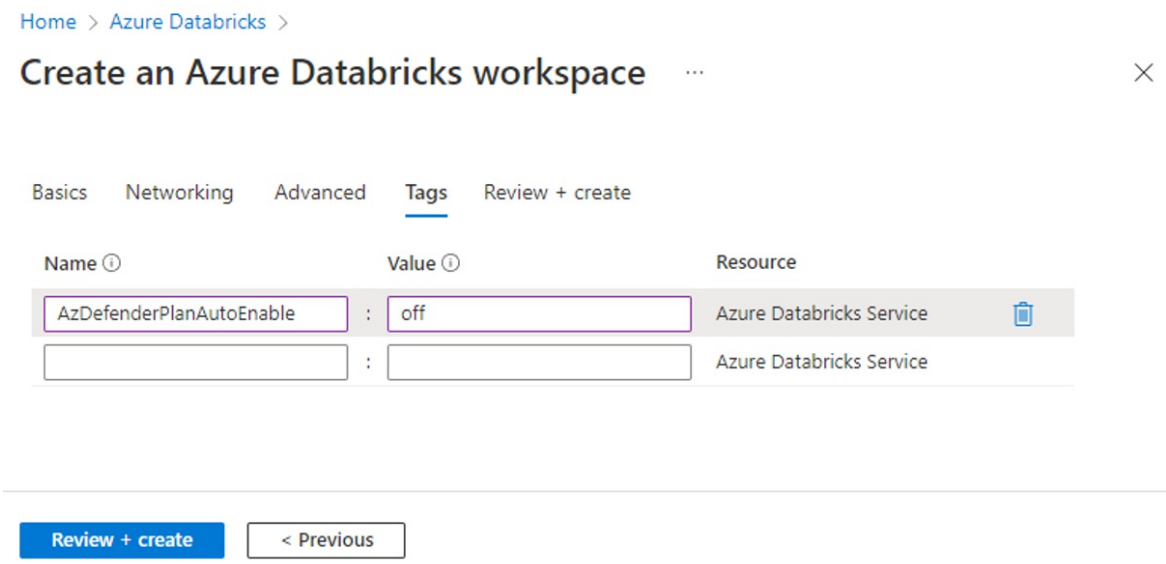

4. Continue following the instructions to create your new Azure Databricks workspace.

The Microsoft Defender for Storage account will inherit the tag of the Databricks workspace, which will prevent Defender for Storage from turning on automatically.

Next steps

- Explore the Microsoft Defender for Storage – Price Estimation Dashboard

Introduction to Microsoft Defender for Key Vault

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Key Vault is a cloud service that safeguards encryption keys and secrets like certificates, connection strings, and passwords.

Enable **Microsoft Defender for Key Vault** for Azure-native, advanced threat protection for Azure Key Vault, providing an additional layer of security intelligence.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for Key Vault** is billed as shown on the pricing page
Clouds:	✅ Commercial clouds ❎ National (Azure Government, Azure China 21Vianet)

What are the benefits of Microsoft Defender for Key Vault?

Microsoft Defender for Key Vault detects unusual and potentially harmful attempts to access or exploit Key Vault accounts. This layer of protection helps you address threats even if you're not a security expert, and without the need to manage third-party security monitoring systems.

When anomalous activities occur, Defender for Key Vault shows alerts and optionally sends them via email to relevant members of your organization. These alerts include the details of the suspicious activity and recommendations on how to investigate and remediate threats.

Microsoft Defender for Key Vault alerts

When you get an alert from Microsoft Defender for Key Vault, we recommend you investigate and respond to the alert as described in Respond to Microsoft Defender for Key Vault. Microsoft Defender for Key Vault protects applications and credentials, so even if you're familiar with the application or user that triggered the alert, it's important to check the situation surrounding every alert.

The alerts appear in Key Vault's **Security** page, the Workload protections, and Defender for Cloud's alerts page.

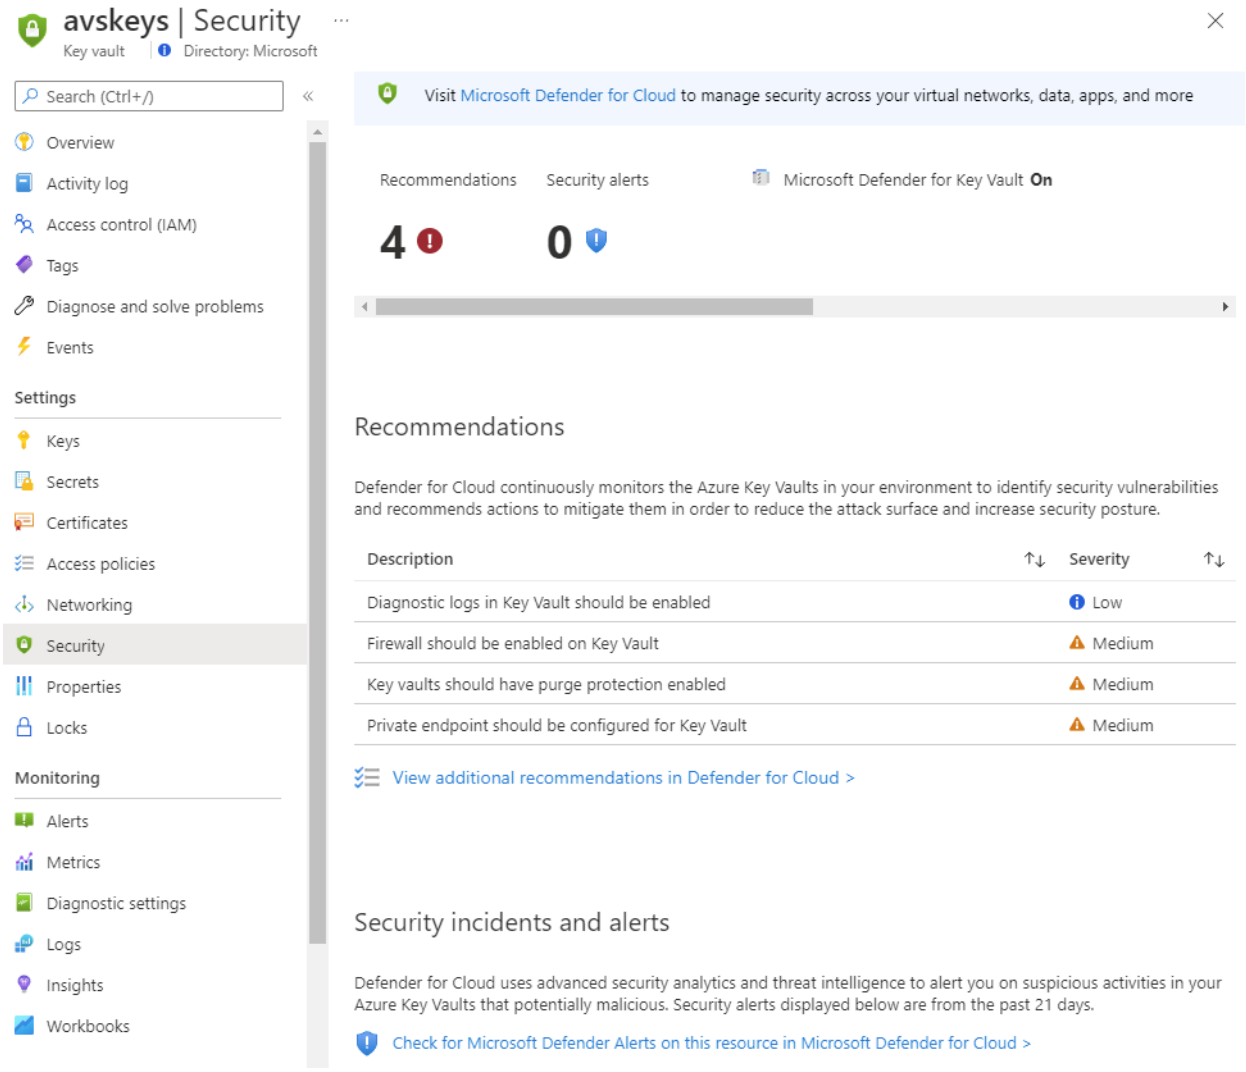

TIP

You can simulate Microsoft Defender for Key Vault alerts by following the instructions in Validating Azure Key Vault threat detection in Microsoft Defender for Cloud.

Respond to Microsoft Defender for Key Vault alerts

When you receive an alert from Microsoft Defender for Key Vault, we recommend you investigate and respond to the alert as described below. Microsoft Defender for Key Vault protects applications and credentials, so even if you're familiar with the application or user that triggered the alert, it's important to verify the situation surrounding every alert.

Alerts from Microsoft Defender for Key Vault includes these elements:

- Object ID
- User Principal Name or IP address of the suspicious resource

Depending on the *type* of access that occurred, some fields might not be available. For example, if your key vault was accessed by an application, you won't see an associated User Principal Name. If the traffic originated from outside of Azure, you won't see an Object ID.

> **TIP**
> Azure virtual machines are assigned Microsoft IPs. This means that an alert might contain a Microsoft IP even though it relates to activity performed from outside of Microsoft. So even if an alert has a Microsoft IP, you should still investigate as described on this page.

Step 1. Identify the source

1. Verify whether the traffic originated from within your Azure tenant. If the key vault firewall is enabled, it's likely that you've provided access to the user or application that triggered this alert.
2. If you can't verify the source of the traffic, continue to Step 2. Respond accordingly.
3. If you can identify the source of the traffic in your tenant, contact the user or owner of the application.

Caution

Microsoft Defender for Key Vault is designed to help identify suspicious activity caused by stolen credentials. **Don't** dismiss the alert simply because you recognize the user or application. Contact the owner of the application or the user and verify the activity was legitimate. You can create a suppression rule to eliminate noise if necessary. Learn more in Suppress security alerts.

Step 2. Respond accordingly

If you don't recognize the user or application, or if you think the access shouldn't have been authorized:

- If the traffic came from an unrecognized IP Address:
  1. Enable the Azure Key Vault firewall as described in Configure Azure Key Vault firewalls and virtual networks.
  2. Configure the firewall with trusted resources and virtual networks.
- If the source of the alert was an unauthorized application or suspicious user:
  1. Open the key vault's access policy settings.
  2. Remove the corresponding security principal, or restrict the operations the security principal can perform.
- If the source of the alert has an Azure Active Directory role in your tenant:
  1. Contact your administrator.
  2. Determine whether there's a need to reduce or revoke Azure Active Directory permissions.

Step 3. Measure the impact

When the event has been mitigated, investigate the secrets in your key vault that were affected:

1. Open the **Security** page on your Azure key vault and view the triggered alert.
2. Select the specific alert that was triggered and review the list of the secrets that were accessed and the timestamp.
3. Optionally, if you have key vault diagnostic logs enabled, review the previous operations for the corresponding caller IP, user principal, or object ID.

Step 4. Take action

When you've compiled your list of the secrets, keys, and certificates that were accessed by the suspicious user or application, you should rotate those objects immediately.

1. Affected secrets should be disabled or deleted from your key vault.
2. If the credentials were used for a specific application:
   a. Contact the administrator of the application and ask them to audit their environment for any uses of the compromised credentials since they were compromised.
   b. If the compromised credentials were used, the application owner should identify the information that

was accessed and mitigate the impact.

Next steps

In this article, you learned about Microsoft Defender for Key Vault.

For related material, see the following articles:

- Key Vault security alerts--The Key Vault section of the reference table for all Microsoft Defender for Cloud alerts
- Continuously export Defender for Cloud data
- Suppress security alerts

Introduction to Microsoft Defender for Key Vault

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Key Vault is a cloud service that safeguards encryption keys and secrets like certificates, connection strings, and passwords.

Enable **Microsoft Defender for Key Vault** for Azure-native, advanced threat protection for Azure Key Vault, providing an additional layer of security intelligence.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for Key Vault** is billed as shown on the pricing page
Clouds:	✅ Commercial clouds ❎ National (Azure Government, Azure China 21Vianet)

What are the benefits of Microsoft Defender for Key Vault?

Microsoft Defender for Key Vault detects unusual and potentially harmful attempts to access or exploit Key Vault accounts. This layer of protection helps you address threats even if you're not a security expert, and without the need to manage third-party security monitoring systems.

When anomalous activities occur, Defender for Key Vault shows alerts and optionally sends them via email to relevant members of your organization. These alerts include the details of the suspicious activity and recommendations on how to investigate and remediate threats.

Microsoft Defender for Key Vault alerts

When you get an alert from Microsoft Defender for Key Vault, we recommend you investigate and respond to the alert as described in Respond to Microsoft Defender for Key Vault. Microsoft Defender for Key Vault protects applications and credentials, so even if you're familiar with the application or user that triggered the alert, it's important to check the situation surrounding every alert.

The alerts appear in Key Vault's **Security** page, the Workload protections, and Defender for Cloud's alerts page.

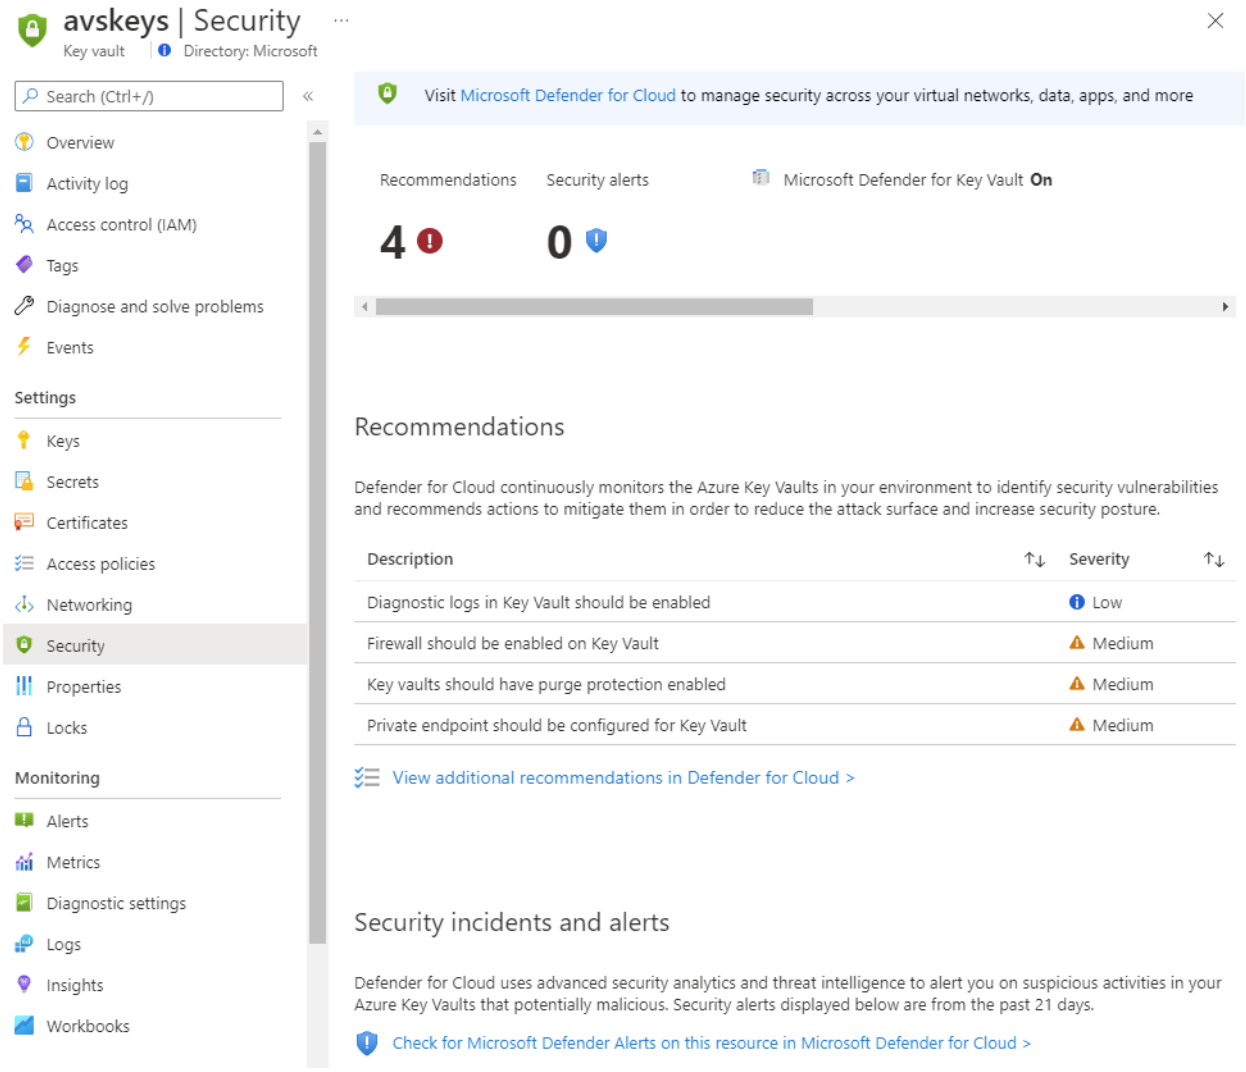

TIP

You can simulate Microsoft Defender for Key Vault alerts by following the instructions in Validating Azure Key Vault threat detection in Microsoft Defender for Cloud.

Respond to Microsoft Defender for Key Vault alerts

When you receive an alert from Microsoft Defender for Key Vault, we recommend you investigate and respond to the alert as described below. Microsoft Defender for Key Vault protects applications and credentials, so even if you're familiar with the application or user that triggered the alert, it's important to verify the situation surrounding every alert.

Alerts from Microsoft Defender for Key Vault includes these elements:

- Object ID
- User Principal Name or IP address of the suspicious resource

Depending on the *type* of access that occurred, some fields might not be available. For example, if your key vault was accessed by an application, you won't see an associated User Principal Name. If the traffic originated from outside of Azure, you won't see an Object ID.

> **TIP**
> Azure virtual machines are assigned Microsoft IPs. This means that an alert might contain a Microsoft IP even though it relates to activity performed from outside of Microsoft. So even if an alert has a Microsoft IP, you should still investigate as described on this page.

Step 1. Identify the source

1. Verify whether the traffic originated from within your Azure tenant. If the key vault firewall is enabled, it's likely that you've provided access to the user or application that triggered this alert.
2. If you can't verify the source of the traffic, continue to Step 2. Respond accordingly.
3. If you can identify the source of the traffic in your tenant, contact the user or owner of the application.

Caution

Microsoft Defender for Key Vault is designed to help identify suspicious activity caused by stolen credentials. **Don't** dismiss the alert simply because you recognize the user or application. Contact the owner of the application or the user and verify the activity was legitimate. You can create a suppression rule to eliminate noise if necessary. Learn more in Suppress security alerts.

Step 2. Respond accordingly

If you don't recognize the user or application, or if you think the access shouldn't have been authorized:

- If the traffic came from an unrecognized IP Address:
  1. Enable the Azure Key Vault firewall as described in Configure Azure Key Vault firewalls and virtual networks.
  2. Configure the firewall with trusted resources and virtual networks.
- If the source of the alert was an unauthorized application or suspicious user:
  1. Open the key vault's access policy settings.
  2. Remove the corresponding security principal, or restrict the operations the security principal can perform.
- If the source of the alert has an Azure Active Directory role in your tenant:
  1. Contact your administrator.
  2. Determine whether there's a need to reduce or revoke Azure Active Directory permissions.

Step 3. Measure the impact

When the event has been mitigated, investigate the secrets in your key vault that were affected:

1. Open the **Security** page on your Azure key vault and view the triggered alert.
2. Select the specific alert that was triggered and review the list of the secrets that were accessed and the timestamp.
3. Optionally, if you have key vault diagnostic logs enabled, review the previous operations for the corresponding caller IP, user principal, or object ID.

Step 4. Take action

When you've compiled your list of the secrets, keys, and certificates that were accessed by the suspicious user or application, you should rotate those objects immediately.

1. Affected secrets should be disabled or deleted from your key vault.
2. If the credentials were used for a specific application:
   a. Contact the administrator of the application and ask them to audit their environment for any uses of the compromised credentials since they were compromised.
   b. If the compromised credentials were used, the application owner should identify the information that

was accessed and mitigate the impact.

Next steps

In this article, you learned about Microsoft Defender for Key Vault.

For related material, see the following articles:

- Key Vault security alerts--The Key Vault section of the reference table for all Microsoft Defender for Cloud alerts
- Continuously export Defender for Cloud data
- Suppress security alerts

Introduction to Microsoft Defender for Resource Manager

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Azure Resource Manager is the deployment and management service for Azure. It provides a management layer that enables you to create, update, and delete resources in your Azure account. You use management features, like access control, locks, and tags, to secure and organize your resources after deployment.

The cloud management layer is a crucial service connected to all your cloud resources. Because of this, it is also a potential target for attackers. Consequently, we recommend security operations teams monitor the resource management layer closely.

Microsoft Defender for Resource Manager automatically monitors the resource management operations in your organization, whether they're performed through the Azure portal, Azure REST APIs, Azure CLI, or other Azure programmatic clients. Defender for Cloud runs advanced security analytics to detect threats and alerts you about suspicious activity.

> **NOTE**
> Some of these analytics are powered by Microsoft Defender for Cloud Apps (formerly known as Microsoft Cloud App Security). To benefit from these analytics, you must activate a Defender for Cloud Apps license. If you have a Defender for Cloud Apps license, then these alerts are enabled by default. To disable the alerts:
>
> 1. From Defender for Cloud's menu, open **Environment settings**.
> 2. Select the subscription you want to change.
> 3. Select **Integrations**.
> 4. Clear **Allow Microsoft Defender for Cloud Apps to access my data**, and select **Save**.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for Resource Manager** is billed as shown on the pricing page
Clouds:	✅ Commercial clouds ✅ Azure Government ✅ Azure China 21Vianet

What are the benefits of Microsoft Defender for Resource Manager?

Microsoft Defender for Resource Manager protects against issues including:

- **Suspicious resource management operations**, such as operations from malicious IP addresses, disabling antimalware and suspicious scripts running in VM extensions
- **Use of exploitation toolkits** like Microburst or PowerZure
- **Lateral movement** from the Azure management layer to the Azure resources data plane

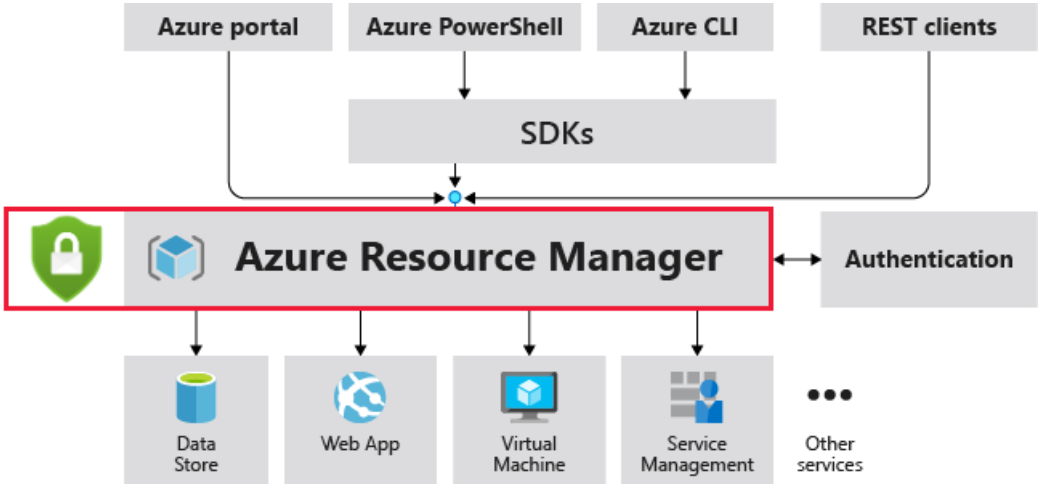

A full list of the alerts provided by Microsoft Defender for Resource Manager is on the alerts reference page.

How to investigate alerts from Microsoft Defender for Resource Manager

Security alerts from Microsoft Defender for Resource Manager are based on threats detected by monitoring Azure Resource Manager operations. Defender for Cloud uses internal log sources of Azure Resource Manager as well as Azure Activity log, a platform log in Azure that provides insight into subscription-level events.

Learn more about Azure Activity log.

To investigate security alerts from Microsoft Defender for Resource Manager:

1. Open Azure Activity log.

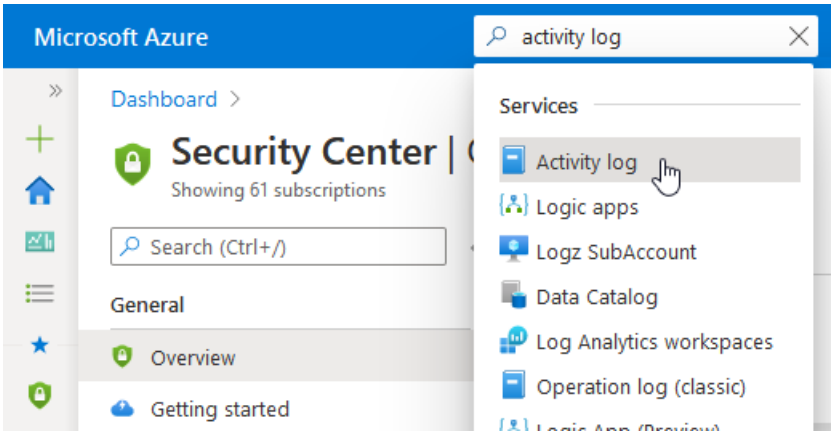

2. Filter the events to:
   - The subscription mentioned in the alert
   - The timeframe of the detected activity
   - The related user account (if relevant)
3. Look for suspicious activities.

TIP

For a better, richer investigation experience, stream your Azure activity logs to Microsoft Sentinel as described in Connect data from Azure Activity log.

Next steps

In this article, you learned about Microsoft Defender for Resource Manager.

Enable enhanced protections

For related material, see the following article:

- Security alerts might be generated or received by Defender for Cloud from different security products. To export all of these alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Exporting alerts to a SIEM solution.

Respond to Microsoft Defender for Resource Manager alerts

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

When you receive an alert from Microsoft Defender for Resource Manager, we recommend you investigate and respond to the alert as described below. Microsoft Defender for Resource Manager protects all connected resources, so even if you're familiar with the application or user that triggered the alert, it's important to verify the situation surrounding every alert.

Step 1. Contact

1. Contact the resource owner to determine whether the behavior was expected or intentional.
2. If the activity is expected, dismiss the alert.
3. If the activity is unexpected, treat the related user accounts, subscriptions, and virtual machines as compromised and mitigate as described in the following step.

Step 2. Immediate mitigation

1. Remediate compromised user accounts:
   - If they're unfamiliar, delete them as they may have been created by a threat actor
   - If they're familiar, change their authentication credentials
   - Use Azure Activity Logs to review all activities performed by the user and identify any that are suspicious
2. Remediate compromised subscriptions:
   - Remove any unfamiliar Runbooks from the compromised automation account
   - Review IAM permissions for the subscription and remove permissions for any unfamiliar user account
   - Review all Azure resources in the subscription and delete any that are unfamiliar
   - Review and investigate any security alerts for the subscription in Microsoft Defender for Cloud
   - Use Azure Activity Logs to review all activities performed in the subscription and identify any that are suspicious
3. Remediate the compromised virtual machines
   - Change the passwords for all users
   - Run a full antimalware scan on the machine
   - Reimage the machines from a malware-free source

Next steps

This page explained the process of responding to an alert from Microsoft Defender for Resource Manager. For related information see the following pages:

- Introduction to Microsoft Defender for Resource Manager
- Suppress security alerts
- Continuously export Defender for Cloud data

Introduction to Microsoft Defender for DNS

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for DNS provides an additional layer of protection for resources that use Azure DNS's Azure-provided name resolution capability.

From within Azure DNS, Defender for DNS monitors the queries from these resources and detects suspicious activities without the need for any additional agents on your resources.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for DNS** is billed as shown on the pricing page
Clouds:	✅ Commercial clouds ✅ Azure China 21Vianet ✅ Azure Government

What are the benefits of Microsoft Defender for DNS?

Microsoft Defender for DNS detects suspicious and anomalous activities such as:

- **Data exfiltration** from your Azure resources using DNS tunneling
- **Malware** communicating with command and control servers
- **DNS attacks** - communication with malicious DNS resolvers
- **Communication with domains used for malicious activities** such as phishing and crypto mining

A full list of the alerts provided by Microsoft Defender for DNS is on the alerts reference page.

Dependencies

Microsoft Defender for DNS doesn't use any agents.

To protect your DNS layer, enable Microsoft Defender for DNS for each of your subscriptions as described in Enable enhanced protections.

Respond to Microsoft Defender for DNS alerts

When you receive an alert from Microsoft Defender for DNS, we recommend you investigate and respond to

the alert as described below. Microsoft Defender for DNS protects all connected resources, so even if you're familiar with the application or user that triggered the alert, it's important to verify the situation surrounding every alert.

Step 1. Contact

1. Contact the resource owner to determine whether the behavior was expected or intentional.
2. If the activity is expected, dismiss the alert.
3. If the activity is unexpected, treat the resource as potentially compromised and mitigate as described in the next step.

Step 2. Immediate mitigation

1. Isolate the resource from the network to prevent lateral movement.
2. Run a full antimalware scan on the resource, following any resulting remediation advice.
3. Review installed and running software on the resource, removing any unknown or unwanted packages.
4. Revert the machine to a known good state, reinstalling the operating system if required, and restore software from a verified malware-free source.
5. Resolve any Microsoft Defender for Cloud recommendations for the machine, remediating highlighted security issues to prevent future breaches.

Next steps

In this article, you learned about Microsoft Defender for DNS.

Enable enhanced protections

For related material, see the following article:

- Security alerts might be generated by Defender for Cloud or received from other security products. To export all of these alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Exporting alerts to a SIEM.

Introduction to Microsoft Defender for DNS

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for DNS provides an additional layer of protection for resources that use Azure DNS's Azure-provided name resolution capability.

From within Azure DNS, Defender for DNS monitors the queries from these resources and detects suspicious activities without the need for any additional agents on your resources.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	**Microsoft Defender for DNS** is billed as shown on the pricing page
Clouds:	✅ Commercial clouds ✅ Azure China 21Vianet ✅ Azure Government

What are the benefits of Microsoft Defender for DNS?

Microsoft Defender for DNS detects suspicious and anomalous activities such as:

- **Data exfiltration** from your Azure resources using DNS tunneling
- **Malware** communicating with command and control servers
- **DNS attacks** - communication with malicious DNS resolvers
- **Communication with domains used for malicious activities** such as phishing and crypto mining

A full list of the alerts provided by Microsoft Defender for DNS is on the alerts reference page.

Dependencies

Microsoft Defender for DNS doesn't use any agents.

To protect your DNS layer, enable Microsoft Defender for DNS for each of your subscriptions as described in Enable enhanced protections.

Respond to Microsoft Defender for DNS alerts

When you receive an alert from Microsoft Defender for DNS, we recommend you investigate and respond to

the alert as described below. Microsoft Defender for DNS protects all connected resources, so even if you're familiar with the application or user that triggered the alert, it's important to verify the situation surrounding every alert.

Step 1. Contact

1. Contact the resource owner to determine whether the behavior was expected or intentional.
2. If the activity is expected, dismiss the alert.
3. If the activity is unexpected, treat the resource as potentially compromised and mitigate as described in the next step.

Step 2. Immediate mitigation

1. Isolate the resource from the network to prevent lateral movement.
2. Run a full antimalware scan on the resource, following any resulting remediation advice.
3. Review installed and running software on the resource, removing any unknown or unwanted packages.
4. Revert the machine to a known good state, reinstalling the operating system if required, and restore software from a verified malware-free source.
5. Resolve any Microsoft Defender for Cloud recommendations for the machine, remediating highlighted security issues to prevent future breaches.

Next steps

In this article, you learned about Microsoft Defender for DNS.

Enable enhanced protections

For related material, see the following article:

- Security alerts might be generated by Defender for Cloud or received from other security products. To export all of these alerts to Microsoft Sentinel, any third-party SIEM, or any other external tool, follow the instructions in Exporting alerts to a SIEM.

Protect Windows Admin Center resources with Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Windows Admin Center is a management tool for your Windows servers. It's a single location for system administrators to access the majority of the most commonly used admin tools. From within Windows Admin Center, you can directly onboard your on-premises servers into Microsoft Defender for Cloud. You can then view a summary of your security recommendations and alerts directly in the Windows Admin Center experience.

> **NOTE**
> Your Azure subscription and the associated Log Analytics workspace both need to have Microsoft Defender for Cloud's enhanced security features enabled in order to enable the Windows Admin Center integration. Enhanced security features are free for the first 30 days if you haven't previously used it on the subscription and workspace. For pricing details in your local currency or region, see the pricing page.

When you've successfully onboarded a server from Windows Admin Center to Microsoft Defender for Cloud, you can:

- View security alerts and recommendations inside the Defender for Cloud extension in Windows Admin Center
- View the security posture and retrieve additional detailed information of your Windows Admin Center managed servers in Defender for Cloud within the Azure portal (or via an API)

By combining these two tools, Defender for Cloud becomes your single pane of glass to view all your security information, whatever the resource: protecting your Windows Admin Center managed on-premises servers, your VMs, and any additional PaaS workloads.

Onboard Windows Admin Center managed servers into Defender for Cloud

1. From Windows Admin Center, select one of your servers, and in the **Tools** pane, select the Microsoft Defender for Cloud extension:

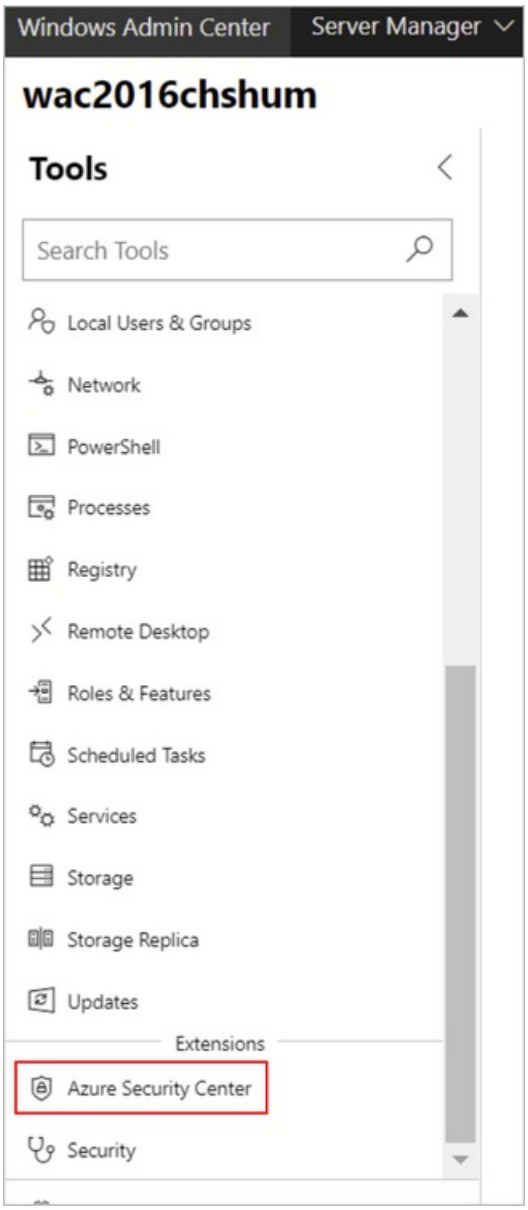

NOTE

If the server is already onboarded to Defender for Cloud, the set-up window will not appear.

2. Click **Sign in to Azure and set up**.

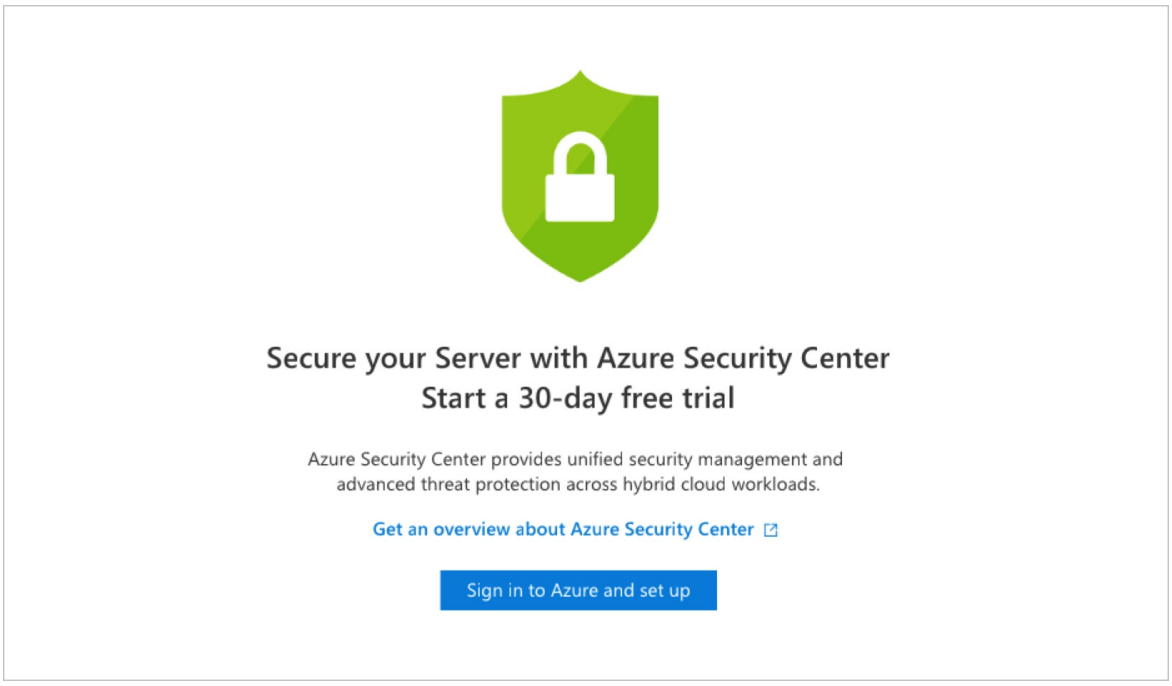

3. Follow the instructions to connect your server to Defender for Cloud. After you've entered the necessary details and confirmed, Defender for Cloud makes the necessary configuration changes to ensure that all of the following are true:

   - An Azure Gateway is registered.
   - The server has a workspace to report to and an associated subscription.
   - Defender for Cloud's Log Analytics solution is enabled on the workspace. This solution provides Microsoft Defender for Cloud's features for *all* servers and virtual machines reporting to this workspace.
   - Microsoft Defender for servers is enabled on the subscription.
   - The Log Analytics agent is installed on the server and configured to report to the selected workspace. If the server already reports to another workspace, it's configured to report to the newly selected workspace as well.

> **NOTE**
> It may take some time after onboarding for recommendations to appear. In fact, depending on on your server activity you may not receive *any* alerts. To generate test alerts to test your alerts are working correctly, follow the instructions in the alert validation procedure.

View security recommendations and alerts in Windows Admin Center

Once onboarded, you can view your alerts and recommendations directly in the Microsoft Defender for Cloud area of Windows Admin Center. Click a recommendation or an alert to view them in the Azure portal. There, you'll get additional information and learn how to remediate issues.

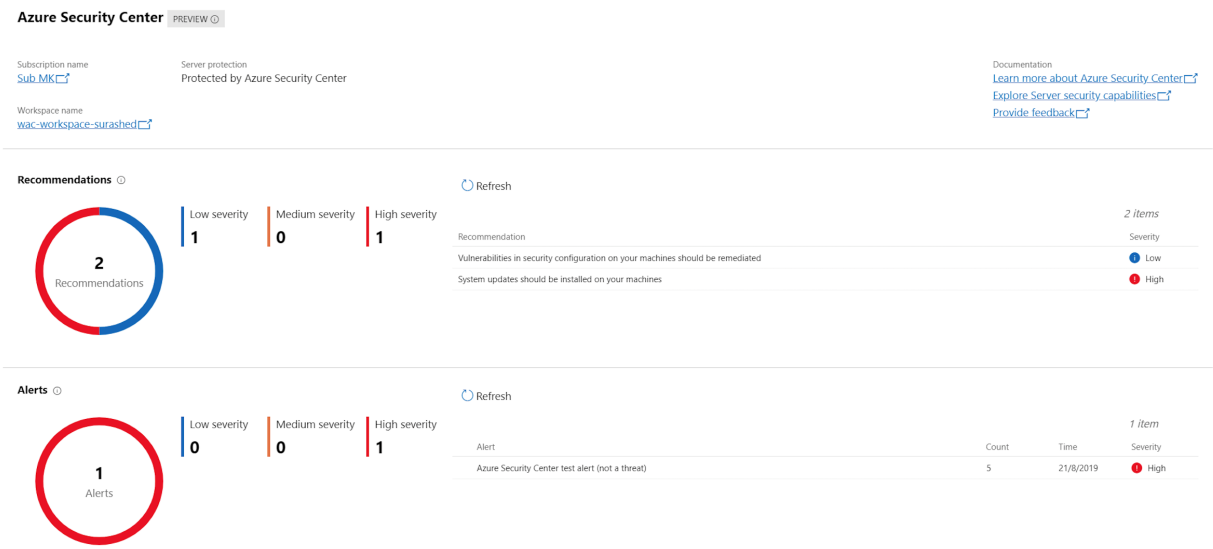

View security recommendations and alerts for Windows Admin Center managed servers in Defender for Cloud

From Microsoft Defender for Cloud:

- To view security recommendations for all your Windows Admin Center servers, open asset inventory and filter to the machine type that you want to investigate. select the **VMs and Computers** tab.
- To view security alerts for all your Windows Admin Center servers, open **Security alerts**. Click **Filter** and ensure **only** "Non-Azure" is selected:

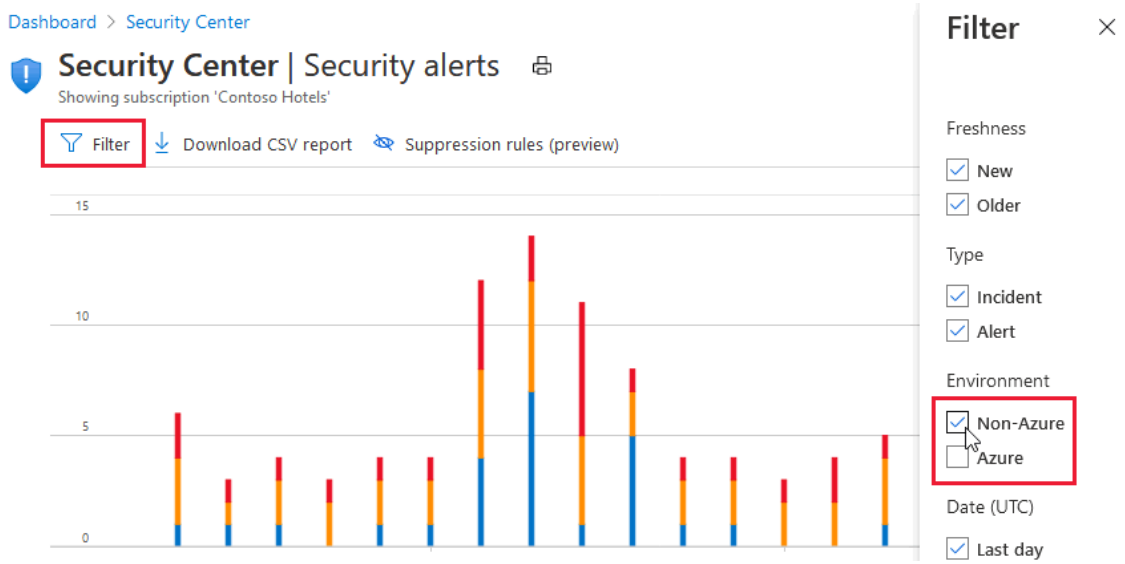

Integrate security solutions in Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This document helps you to manage security solutions already connected to Microsoft Defender for Cloud and add new ones.

Integrated Azure security solutions

Defender for Cloud makes it easy to enable integrated security solutions in Azure. Benefits include:

- **Simplified deployment**: Defender for Cloud offers streamlined provisioning of integrated partner solutions. For solutions like antimalware and vulnerability assessment, Defender for Cloud can provision the agent on your virtual machines. For firewall appliances, Defender for Cloud can take care of much of the network configuration required.
- **Integrated detections**: Security events from partner solutions are automatically collected, aggregated, and displayed as part of Defender for Cloud alerts and incidents. These events also are fused with detections from other sources to provide advanced threat-detection capabilities.
- **Unified health monitoring and management**: Customers can use integrated health events to monitor all partner solutions at a glance. Basic management is available, with easy access to advanced setup by using the partner solution.

Currently, integrated security solutions include vulnerability assessment by Qualys and Rapid7.

> **NOTE**
> Defender for Cloud does not install the Log Analytics agent on partner virtual appliances because most security vendors prohibit external agents running on their appliances.

To learn more about the integration of vulnerability scanning tools from Qualys, including a built-in scanner available to customers who've enabled Microsoft Defender for servers, see Defender for Cloud's integrated Qualys vulnerability scanner for Azure and hybrid machines.

Defender for Cloud also offers vulnerability analysis for your:

- SQL databases - see Explore vulnerability assessment reports in the vulnerability assessment dashboard
- Azure Container Registry images - see Use Microsoft Defender for container registries to scan your images for vulnerabilities

How security solutions are integrated

Azure security solutions that are deployed from Defender for Cloud are automatically connected. You can also

connect other security data sources, including computers running on-premises or in other clouds.

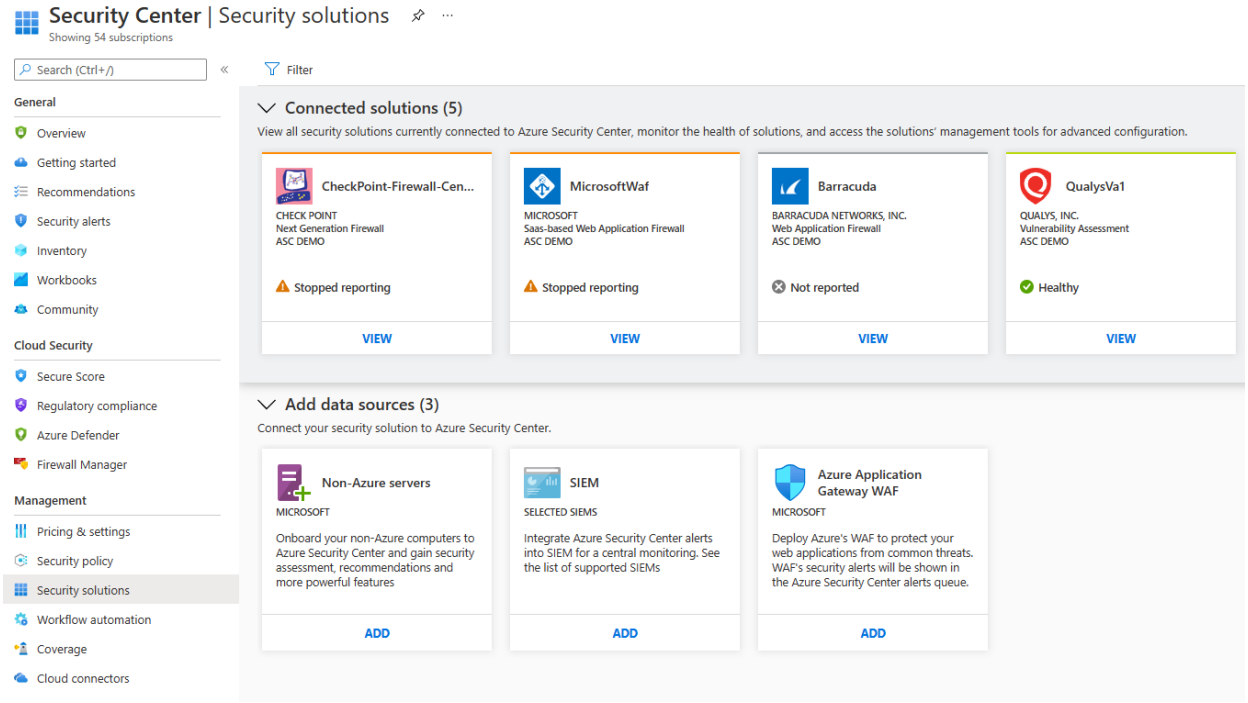

Manage integrated Azure security solutions and other data sources

1. From the Azure portal, open **Defender for Cloud**.
2. From Defender for Cloud's menu, select **Security solutions**.

From the **Security solutions** page, you can see the health of integrated Azure security solutions and run basic management tasks.

Connected solutions

The **Connected solutions** section includes security solutions that are currently connected to Defender for Cloud. It also shows the health status of each solution.

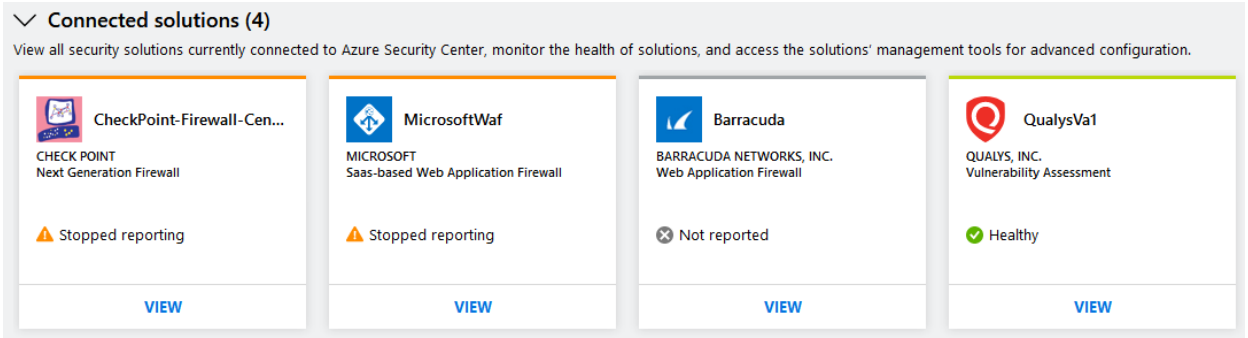

The status of a partner solution can be:

- **Healthy** (green) - no health issues.
- **Unhealthy** (red) - there's a health issue that requires immediate attention.
- **Stopped reporting** (orange) - the solution has stopped reporting its health.
- **Not reported** (gray) - the solution hasn't reported anything yet and no health data is available. A solution's status may be unreported if it was connected recently and is still deploying.

> **NOTE**
> If health status data is not available, Defender for Cloud shows the date and time of the last event received to indicate whether the solution is reporting or not. If no health data is available and no alerts were received within the last 14 days, Defender for Cloud indicates that the solution is unhealthy or not reporting.

Select **VIEW** for additional information and options such as:

- **Solution console** - Opens the management experience for this solution.
- **Link VM** - Opens the Link Applications page. Here you can connect resources to the partner solution.
- **Delete solution**
- **Configure**

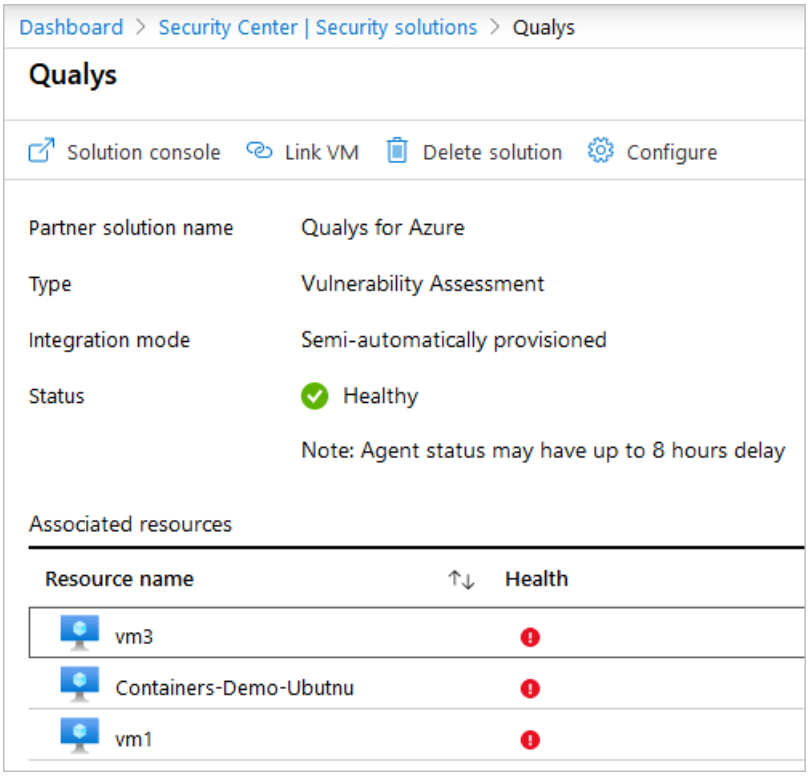

Discovered solutions

Defender for Cloud automatically discovers security solutions running in Azure but not connected to Defender for Cloud and displays the solutions in the **Discovered solutions** section. These solutions include Azure solutions, like Azure AD Identity Protection, and partner solutions.

> **NOTE**
> Enable **advanced protections** at the subscription level for the discovered solutions feature. Learn more in Quickstart: Enable enhanced security features.

Select **CONNECT** under a solution to integrate with Defender for Cloud and be notified of security alerts.

Add data sources

The **Add data sources** section includes other available data sources that can be connected. For instructions on adding data from any of these sources, click **ADD**.

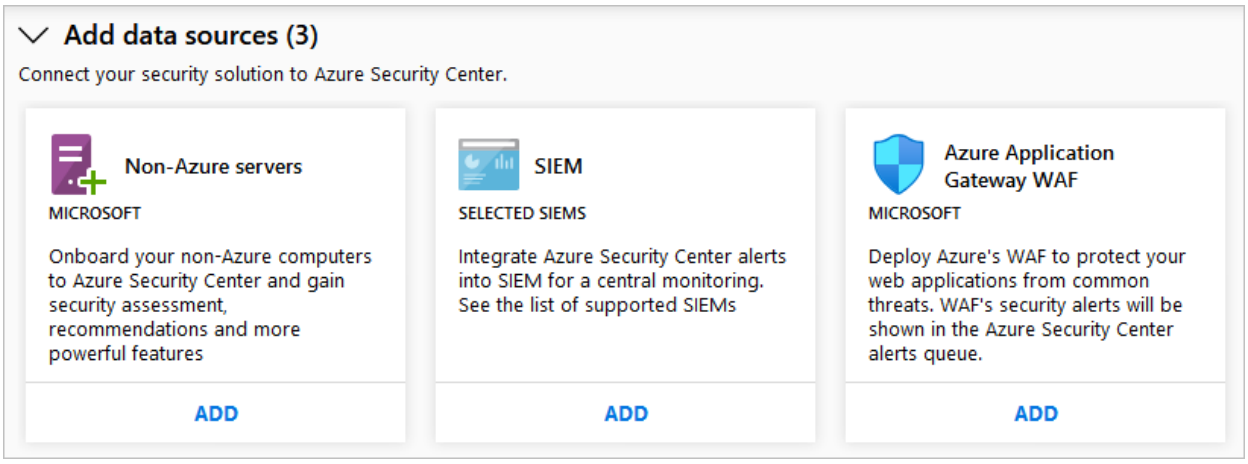

Next steps

In this article, you learned how to integrate partner solutions in Defender for Cloud. To learn how to setup an integration with Microsoft Sentinel, or any other SIEM, see Continuously export Defender for Cloud data.

Protect your network resources

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud continuously analyzes the security state of your Azure resources for network security best practices. When Defender for Cloud identifies potential security vulnerabilities, it creates recommendations that guide you through the process of configuring the needed controls to harden and protect your resources.

For a full list of the recommendations for Networking, see Networking recommendations.

This article addresses recommendations that apply to your Azure resources from a network security perspective. Networking recommendations center around next generation firewalls, Network Security Groups, JIT VM access, overly permissive inbound traffic rules, and more. For a list of networking recommendations and remediation actions, see Managing security recommendations in Microsoft Defender for Cloud.

The **Networking** features of Defender for Cloud include:

- Network map (requires Microsoft Defender for servers)
- Adaptive network hardening (requires Microsoft Defender for servers)
- Networking security recommendations

View your networking resources and their recommendations

From the asset inventory page, use the resource type filter to select the networking resources that you want to investigate:

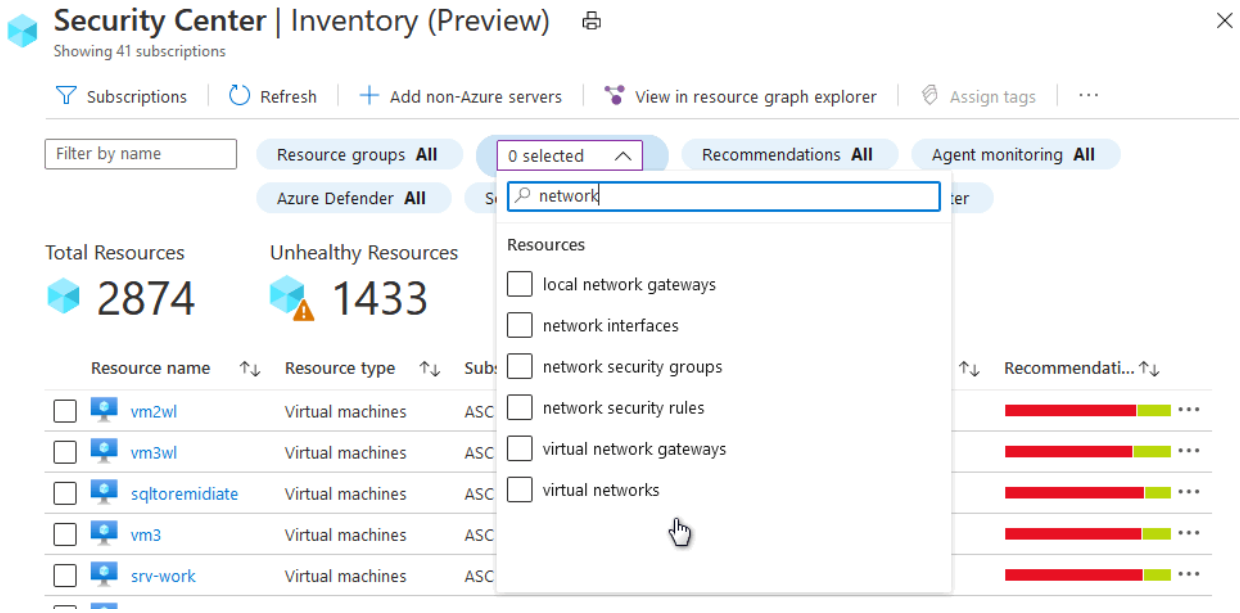

Network map

The interactive network map provides a graphical view with security overlays giving you recommendations and insights for hardening your network resources. Using the map you can see the network topology of your Azure workloads, connections between your virtual machines and subnets, and the capability to drill down from the map into specific resources and the recommendations for those resources.

To open the Network map:

1. From Defender for Cloud's menu, open the **Workload protections** dashboard.
2. Select **Network map**.

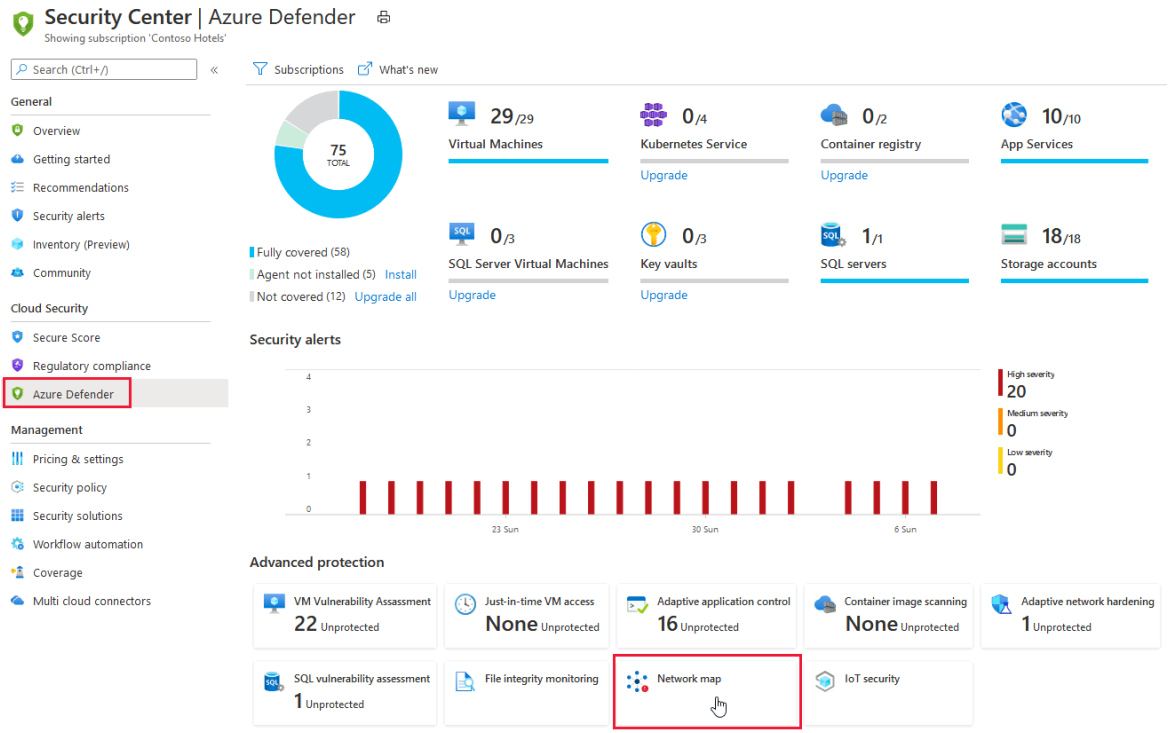

3. Select the **Layers** menu choose **Topology**.

The default view of the topology map displays:

- Currently selected subscriptions - The map is optimized for the subscriptions you selected in the portal. If you modify your selection, the map is regenerated with the new selections.
- VMs, subnets, and VNets of the Resource Manager resource type ("classic" Azure resources are not supported)
- Peered VNets
- Only resources that have network recommendations with a high or medium severity
- Internet-facing resources

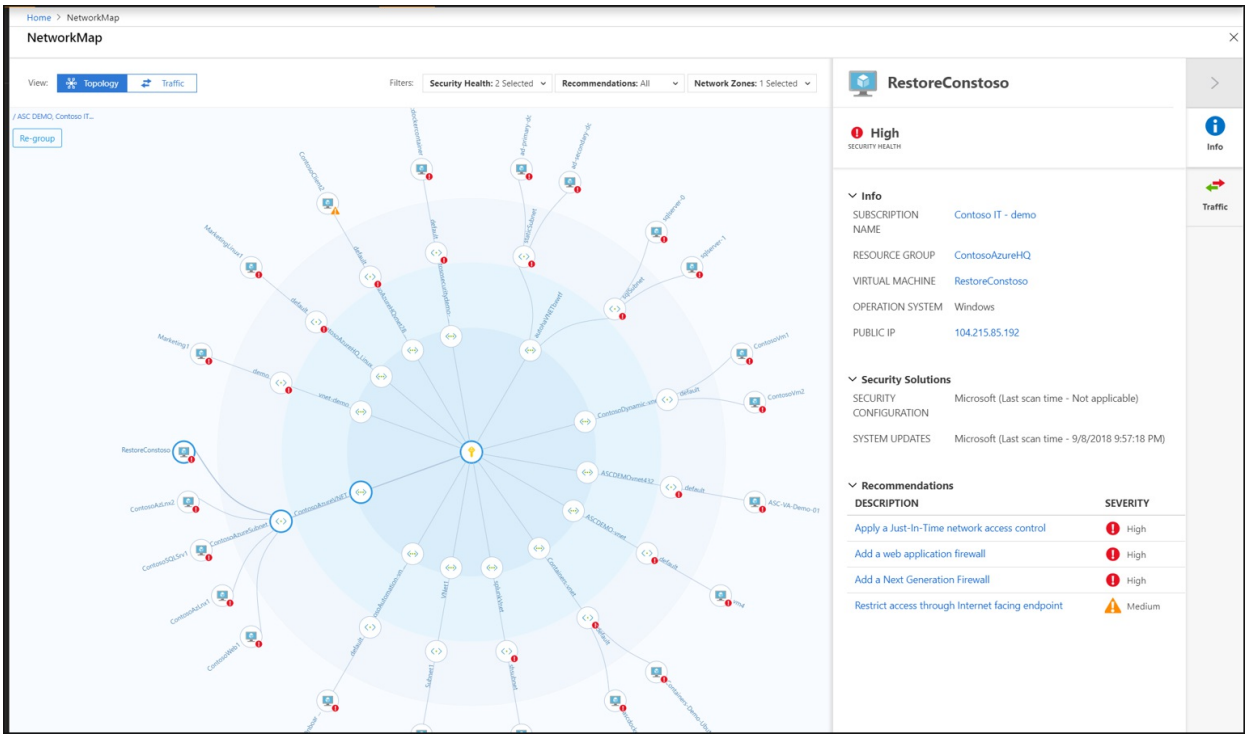

Understanding the network map

The network map can show you your Azure resources in a **Topology** view and a **Traffic** view.

The topology view

In the **Topology** view of the networking map, you can view the following insights about your networking resources:

- In the inner circle, you can see all the Vnets within your selected subscriptions, the next circle is all the subnets, the outer circle is all the virtual machines.
- The lines connecting the resources in the map let you know which resources are associated with each other, and how your Azure network is structured.
- Use the severity indicators to quickly get an overview of which resources have open recommendations from Defender for Cloud.
- You can click any of the resources to drill down into them and view the details of that resource and its recommendations directly, and in the context of the Network map.
- If there are too many resources being displayed on the map, Microsoft Defender for Cloud uses its proprietary algorithm to 'smart cluster' your resources, highlighting the ones that are in the most critical state, and have the most high severity recommendations.

Because the map is interactive and dynamic, every node is clickable, and the view can change based on the filters:

1. You can modify what you see on the network map by using the filters at the top. You can focus the map based on:
    - **Security health**: You can filter the map based on Severity (High, Medium, Low) of your Azure resources.
    - **Recommendations**: You can select which resources are displayed based on which recommendations are active on those resources. For example, you can view only resources for which Defender for Cloud recommends you enable Network Security Groups.
    - **Network zones**: By default, the map displays only Internet facing resources, you can select internal VMs as well.
2. You can click **Reset** in top left corner at any time to return the map to its default state.

To drill down into a resource:

1. When you select a specific resource on the map, the right pane opens and gives you general information about the resource, connected security solutions if there are any, and the recommendations relevant to the resource. It's the same type of behavior for each type of resource you select.
2. When you hover over a node in the map, you can view general information about the resource, including subscription, resource type, and resource group.
3. Use the link to zoom into the tool tip and refocus the map on that specific node.
4. To refocus the map away from a specific node, zoom out.

The Traffic view

The **Traffic** view provides you with a map of all the possible traffic between your resources. This provides you with a visual map of all the rules you configured that define which resources can communicate with whom. This enables you to see the existing configuration of the network security groups as well as quickly identify possible risky configurations within your workloads.

Uncover unwanted connections

The strength of this view is in its ability to show you these allowed connections together with the vulnerabilities that exist, so you can use this cross-section of data to perform the necessary hardening on your resources.

For example, you might detect two machines that you weren't aware could communicate, enabling you to better isolate the workloads and subnets.

Investigate resources

To drill down into a resource:

1. When you select a specific resource on the map, the right pane opens and gives you general information about the resource, connected security solutions if there are any, and the recommendations relevant to the resource. It's the same type of behavior for each type of resource you select.
2. Click **Traffic** to see the list of possible outbound and inbound traffic on the resource - this is a comprehensive list of who can communicate with the resource and who it can communicate with, and through which protocols and ports. For example, when you select a VM, all the VMs it can communicate with are shown, and when you select a subnet, all the subnets which it can communicate with are shown.

This data is based on analysis of the Network Security Groups as well as advanced machine learning algorithms that analyze multiple rules to understand their crossovers and interactions.

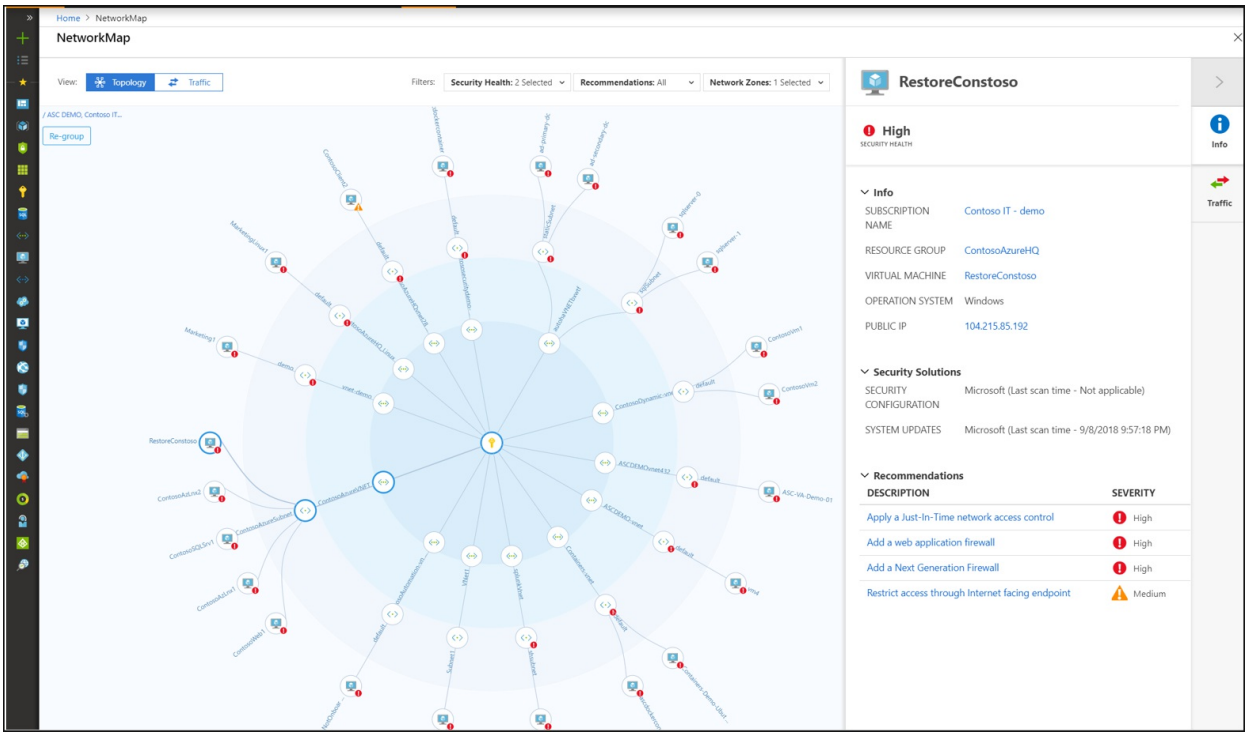

Next steps

To learn more about recommendations that apply to other Azure resource types, see the following:

- Protecting your machines and applications in Microsoft Defender for Cloud

Manage multi-factor authentication (MFA) enforcement on your subscriptions

2/15/2022 • 5 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

If you're only using passwords to authenticate your users, you're leaving an attack vector open. Users often use weak passwords or reuse them for multiple services. With MFA enabled, your accounts are more secure, and users can still authenticate to almost any application with single sign-on (SSO).

There are multiple ways to enable MFA for your Azure Active Directory (AD) users based on the licenses that your organization owns. This page provides the details for each in the context of Microsoft Defender for Cloud.

MFA and Microsoft Defender for Cloud

Defender for Cloud places a high value on MFA. The security control that contributes the most to your secure score is **Enable MFA**.

The recommendations in the Enable MFA control ensure you're meeting the recommended practices for users of your subscriptions:

- MFA should be enabled on accounts with owner permissions on your subscription
- MFA should be enabled on accounts with write permissions on your subscription

There are three ways to enable MFA and be compliant with the two recommendations in Defender for Cloud: security defaults, per-user assignment, conditional access (CA) policy. Each of these options is explained below.

Free option - security defaults

If you're using the free edition of Azure AD, use security defaults to enable multi-factor authentication on your tenant.

MFA for Microsoft 365 Business, E3, or E5 customers

Customers with Microsoft 365 can use **Per-user assignment**. In this scenario, Azure AD MFA is either enabled or disabled for all users, for all sign-in events. There is no ability to enable multi-factor authentication for a subset of users, or under certain scenarios, and management is through the Office 365 portal.

MFA for Azure AD Premium customers

For an improved user experience, upgrade to Azure AD Premium P1 or P2 for **conditional access (CA) policy** options. To configure a CA policy, you'll need Azure Active Directory (AD) tenant permissions.

Your CA policy must:

- enforce MFA
- include the Microsoft Azure Management app ID (797f4846-ba00-4fd7-ba43-dac1f8f63013) or all apps
- not exclude the Microsoft Azure Management app ID

Azure AD Premium P1 customers can use Azure AD CA to prompt users for multi-factor authentication

during certain scenarios or events to fit your business requirements. Other licenses that include this functionality: Enterprise Mobility + Security E3, Microsoft 365 F1, and Microsoft 365 E3.

Azure AD Premium P2 provides the strongest security features and an improved user experience. This license adds risk-based conditional access to the Azure AD Premium P1 features. Risk-based CA adapts to your users' patterns and minimizes multi-factor authentication prompts. Other licenses that include this functionality: Enterprise Mobility + Security E5 or Microsoft 365 E5.

Learn more in the Azure Conditional Access documentation.

Identify accounts without multi-factor authentication (MFA) enabled

You can view the list of user accounts without MFA enabled from either the Defender for Cloud recommendations details page, or using Azure Resource Graph.

View the accounts without MFA enabled in the Azure portal

From the recommendation details page, select a subscription from the **Unhealthy resources** list or select **Take action** and the list will be displayed.

View the accounts without MFA enabled using Azure Resource Graph

To see which accounts don't have MFA enabled, use the following Azure Resource Graph query. The query returns all unhealthy resources - accounts - of the recommendation "MFA should be enabled on accounts with owner permissions on your subscription".

1. Open **Azure Resource Graph Explorer**.

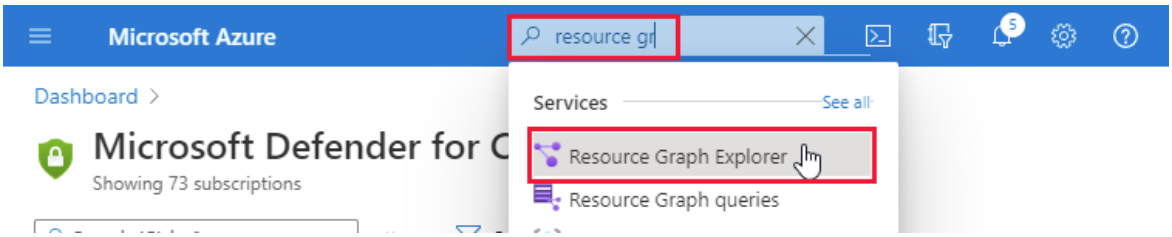

2. Enter the following query and select **Run query**.

```
securityresources
 | where type == "microsoft.security/assessments"
 | where properties.displayName == "MFA should be enabled on accounts with owner permissions on your
subscription"
 | where properties.status.code == "Unhealthy"
```

3. The `additionalData` property reveals the list of account object IDs for accounts that don't have MFA enforced.

> **NOTE**
> The accounts are shown as object IDs rather than account names to protect the privacy of the account holders.

> **TIP**
> Alternatively, you can use the Defender for Cloud REST API method Assessments - Get.

FAQ - MFA in Defender for Cloud

- We're already using CA policy to enforce MFA. Why do we still get the Defender for Cloud recommendations?

- We're using a third-party MFA tool to enforce MFA. Why do we still get the Defender for Cloud recommendations?
- Why does Defender for Cloud show user accounts without permissions on the subscription as "requiring MFA"?
- We're enforcing MFA with PIM. Why are PIM accounts shown as noncompliant?
- Can I exempt or dismiss some of the accounts?
- Are there any limitations to Defender for Cloud's identity and access protections?

We're already using CA policy to enforce MFA. Why do we still get the Defender for Cloud recommendations?

To investigate why the recommendations are still being generated, verify the following configuration options in your MFA CA policy:

- You've included the accounts in the **Users** section of your MFA CA policy (or one of the groups in the **Groups** section)
- The Azure Management app ID (797f4846-ba00-4fd7-ba43-dac1f8f63013), or all apps, are included in the **Apps** section of your MFA CA policy
- The Azure Management app ID isn't excluded in the **Apps** section of your MFA CA policy

We're using a third-party MFA tool to enforce MFA. Why do we still get the Defender for Cloud recommendations?

Defender for Cloud's MFA recommendations don't support third-party MFA tools (for example, DUO).

If the recommendations are irrelevant for your organization, consider marking them as "mitigated" as described in Exempting resources and recommendations from your secure score. You can also disable a recommendation.

Why does Defender for Cloud show user accounts without permissions on the subscription as "requiring MFA"?

Defender for Cloud's MFA recommendations refer to Azure RBAC roles and the Azure classic subscription administrators role. Verify that none of the accounts have such roles.

We're enforcing MFA with PIM. Why are PIM accounts shown as noncompliant?

Defender for Cloud's MFA recommendations currently don't support PIM accounts. You can add these accounts to a CA Policy in the Users/Group section.

Can I exempt or dismiss some of the accounts?

The capability to exempt some accounts that don’t use MFA isn't currently supported.

Are there any limitations to Defender for Cloud's identity and access protections?

There are some limitations to Defender for Cloud's identity and access protections:

- Identity recommendations aren't available for subscriptions with more than 600 accounts. In such cases, these recommendations will be listed under "unavailable assessments".
- Identity recommendations aren't available for Cloud Solution Provider (CSP) partner's admin agents.
- Identity recommendations don’t identify accounts that are managed with a privileged identity management (PIM) system. If you're using a PIM tool, you might see inaccurate results in the **Manage access and permissions** control.

Next steps

To learn more about recommendations that apply to other Azure resource types, see the following article:

- Protecting your network in Microsoft Defender for Cloud

Additional threat protections in Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

In addition to its built-in advanced protection plans, Microsoft Defender for Cloud also offers the following threat protection capabilities.

TIP

To enable Defender for Cloud's threat protection capabilities, you must enable enhanced security features on the subscription containing the applicable workloads.

Threat protection for Azure network layer

Defender for Cloud network-layer analytics are based on sample IPFIX data, which are packet headers collected by Azure core routers. Based on this data feed, Defender for Cloud uses machine learning models to identify and flag malicious traffic activities. Defender for Cloud also uses the Microsoft Threat Intelligence database to enrich IP addresses.

Some network configurations restrict Defender for Cloud from generating alerts on suspicious network activity. For Defender for Cloud to generate network alerts, ensure that:

- Your virtual machine has a public IP address (or is on a load balancer with a public IP address).
- Your virtual machine's network egress traffic isn't blocked by an external IDS solution.

For a list of the Azure network layer alerts, see the Reference table of alerts.

Threat protection for Azure Cosmos DB (Preview)

The Azure Cosmos DB alerts are generated by unusual and potentially harmful attempts to access or exploit Azure Cosmos DB accounts.

For more information, see:

- Advanced threat protection for Azure Cosmos DB (Preview)
- The list of threat protection alerts for Azure Cosmos DB (Preview)

Display recommendations in Microsoft Defender for Cloud Apps

Microsoft Defender for Cloud Apps (formerly known as Microsoft Cloud App Security) is a cloud access security broker (CASB) that supports various deployment modes including log collection, API connectors, and reverse proxy. It provides rich visibility, control over data travel, and sophisticated analytics to identify and combat cyberthreats across all your Microsoft and third-party cloud services.

If you've enabled Microsoft Defender for Cloud Apps, and selected the integration from within Defender for Cloud's settings, your hardening recommendations from Defender for Cloud will appear in Defender for Cloud Apps with no additional configuration needed.

> **NOTE**
>
> Defender for Cloud stores security-related customer data in the same geo as its resource. If Microsoft hasn't yet deployed Defender for Cloud in the resource's geo, then it stores the data in the United States. When Microsoft Defender for Cloud Apps is enabled, this information is stored in accordance with the geo location rules of Microsoft Defender for Cloud Apps. For more information, see Data storage for non-regional services.

Stream security alerts from other Microsoft services

Display Azure WAF alerts in Defender for Cloud

Azure Application Gateway offers a web application firewall (WAF) that provides centralized protection of your web applications from common exploits and vulnerabilities.

Web applications are increasingly targeted by malicious attacks that exploit commonly known vulnerabilities. The Application Gateway WAF is based on Core Rule Set 3.0 or 2.2.9 from the Open Web Application Security Project. The WAF is updated automatically to protect against new vulnerabilities.

If you have a license for Azure WAF, your WAF alerts are streamed to Defender for Cloud with no additional configuration needed. For more information on the alerts generated by WAF, see Web application firewall CRS rule groups and rules.

Display Azure DDoS Protection alerts in Defender for Cloud

Distributed denial of service (DDoS) attacks are known to be easy to execute. They've become a great security concern, particularly if you're moving your applications to the cloud. A DDoS attack attempts to exhaust an application's resources, making the application unavailable to legitimate users. DDoS attacks can target any endpoint that can be reached through the internet.

To defend against DDoS attacks, purchase a license for Azure DDoS Protection and ensure you're following application design best practices. DDoS Protection provides different service tiers. For more information, see Azure DDoS Protection overview.

If you have Azure DDoS Protection enabled, your DDoS alerts are streamed to Defender for Cloud with no additional configuration needed. For more information on the alerts generated by DDoS Protection, see Reference table of alerts.

Next steps

To learn more about the security alerts from these threat protection features, see the following articles:

- Reference table for all Defender for Cloud alerts
- Security alerts in Defender for Cloud
- Manage and respond to security alerts in Defender for Cloud
- Continuously export Defender for Cloud data

Organize subscriptions into management groups and assign roles to users

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how to manage your organization's security posture at scale by applying security policies to all Azure subscriptions linked to your Azure Active Directory tenant.

For visibility into the security posture of all subscriptions linked to an Azure AD tenant, you'll need an Azure role with sufficient read permissions assigned on the root management group.

Organize your subscriptions into management groups

Introduction to management groups

Use management groups to efficiently manage access, policies, and reporting on **groups of subscriptions**, as well as effectively manage the entire Azure estate by performing actions on the root management group. You can organize subscriptions into management groups and apply your governance policies to the management groups. All subscriptions within a management group automatically inherit the policies applied to the management group.

Each Azure AD tenant is given a single top-level management group called the **root management group**. This root management group is built into the hierarchy to have all management groups and subscriptions fold up to it. This group allows global policies and Azure role assignments to be applied at the directory level.

The root management group is created automatically when you do any of the following actions:

- Open **Management Groups** in the Azure portal.
- Create a management group with an API call.
- Create a management group with PowerShell. For PowerShell instructions, see Create management groups for resource and organization management.

Management groups aren't required to onboard Defender for Cloud, but we recommend creating at least one so that the root management group gets created. After the group is created, all subscriptions under your Azure AD tenant will be linked to it.

For a detailed overview of management groups, see the Organize your resources with Azure management groups article.

View and create management groups in the Azure portal

1. From the Azure portal, use the search box in the top bar to find and open **Management Groups**.

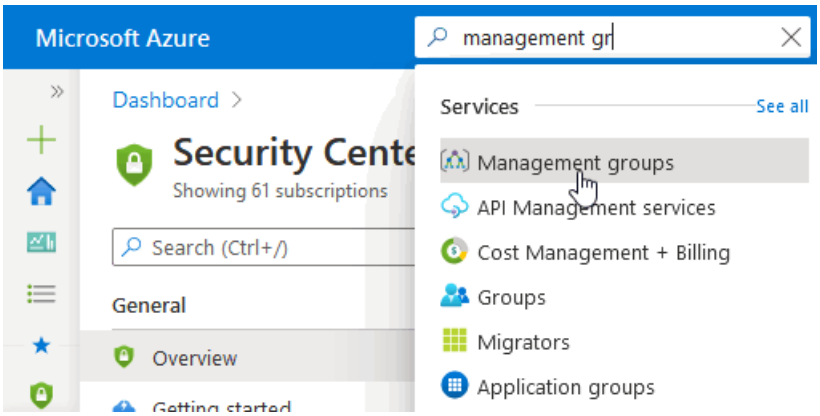

The list of your management groups appears.

2. To create a management group, select **Add management group**, enter the relevant details, and select **Save**.

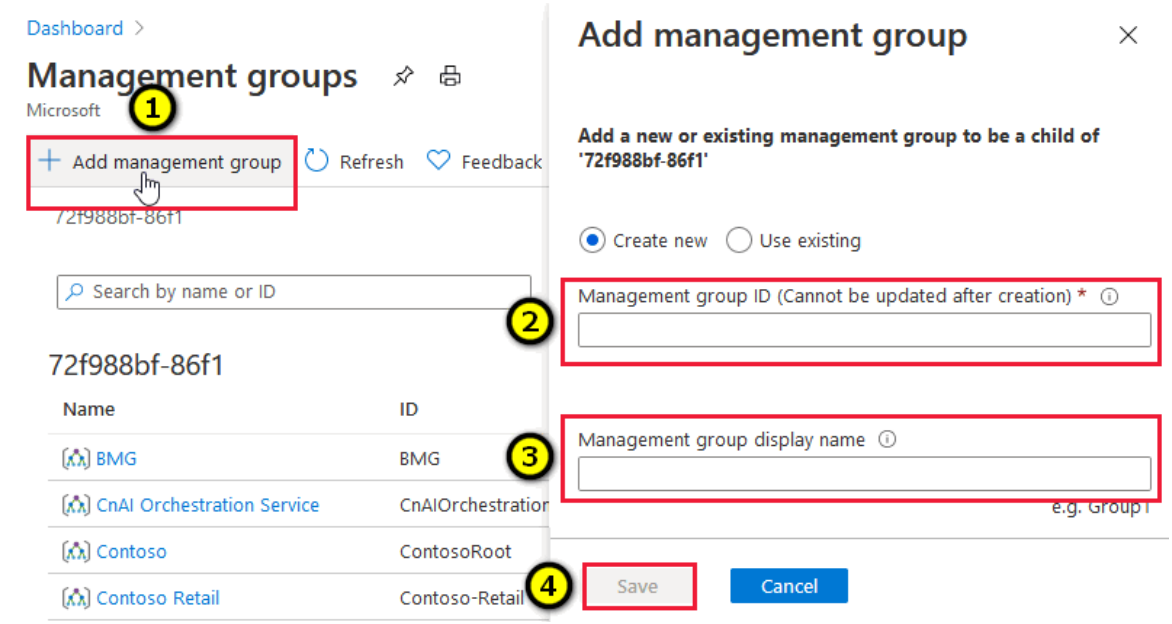

- The **Management Group ID** is the directory unique identifier that is used to submit commands on this management group. This identifier isn't editable after creation as it is used throughout the Azure system to identify this group.
- The display name field is the name that is displayed within the Azure portal. A separate display name is an optional field when creating the management group and can be changed at any time.

Add subscriptions to a management group

You can add subscriptions to the management group that you created.

1. From the Azure portal, open **Management Groups** and select the management group for your subscription.

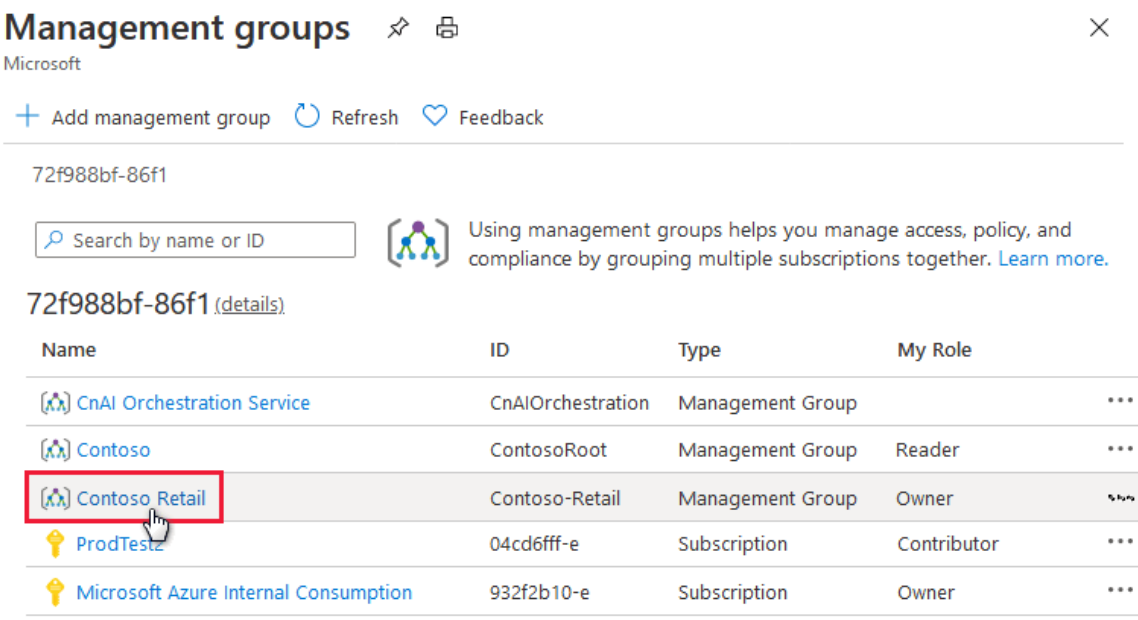

2. When the group's page opens, select **Subscriptions**.

3. From the subscriptions page, select **Add**, then select your subscriptions and select **Save**. Repeat until you've added all the subscriptions in the scope.

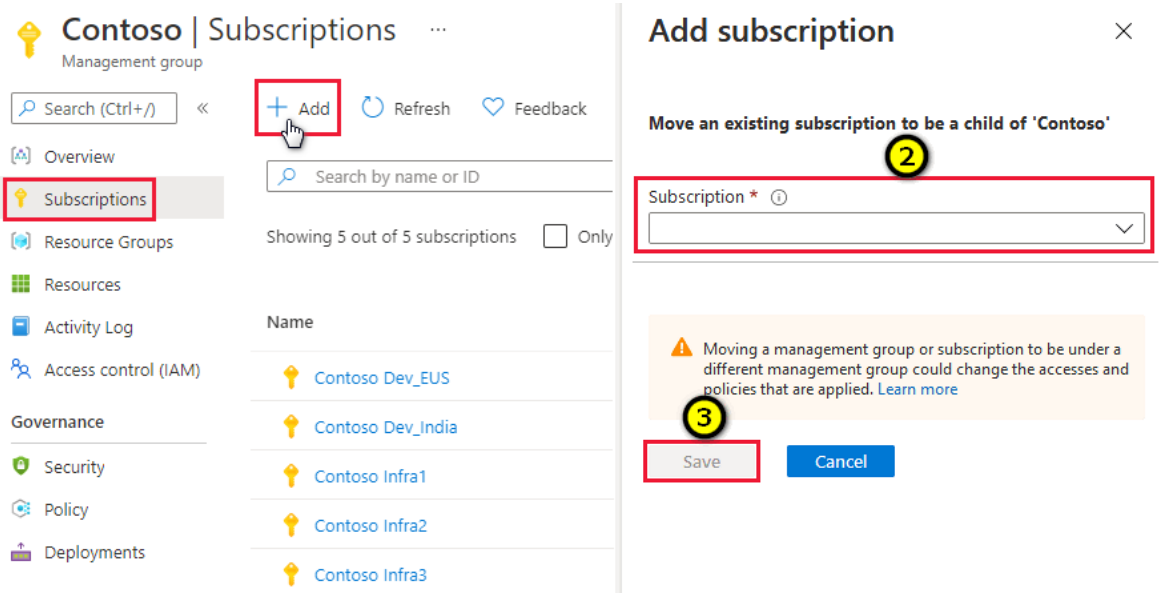

> **IMPORTANT**
>
> Management groups can contain both subscriptions and child management groups. When you assign a user an Azure role to the parent management group, the access is inherited by the child management group's subscriptions. Policies set at the parent management group are also inherited by the children.

Assign Azure roles to other users

Assign Azure roles to users through the Azure portal:

1. From the Azure portal, use the search box in the top bar to find and open **Management Groups**.

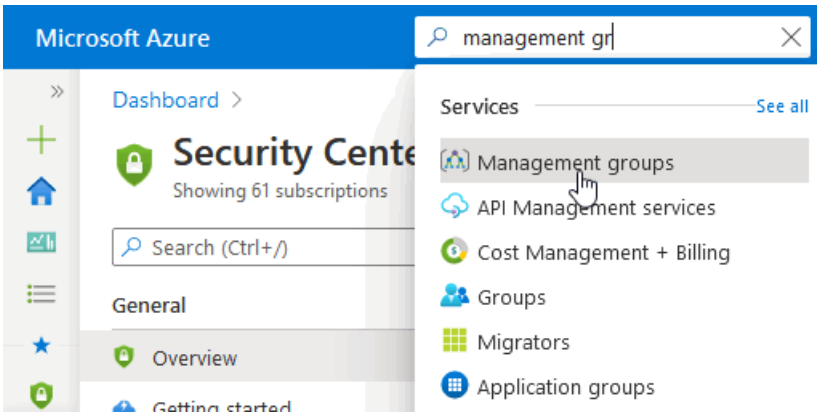

The list of your management groups appears.

2. Select the relevant management group.
3. Select **Access control (IAM)**, open the **Role assignments** tab and select **Add** > **Add role assignment**.

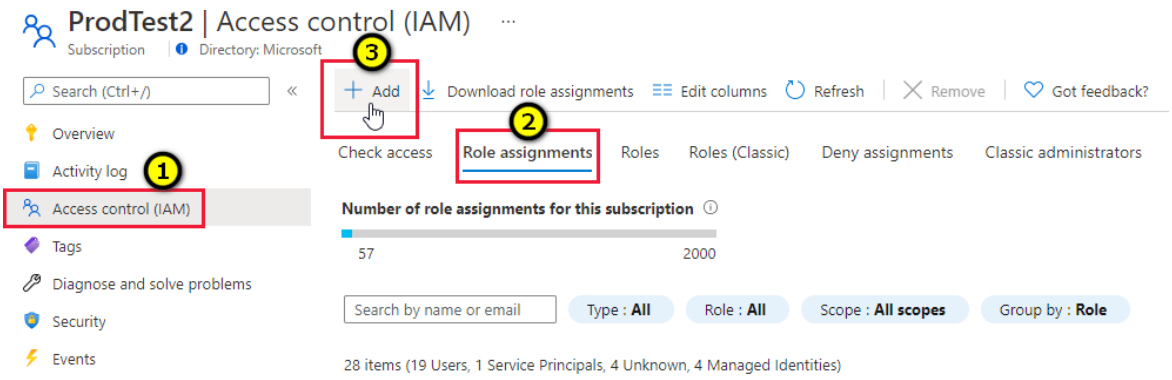

4. From the **Add role assignment** page, select the relevant role.

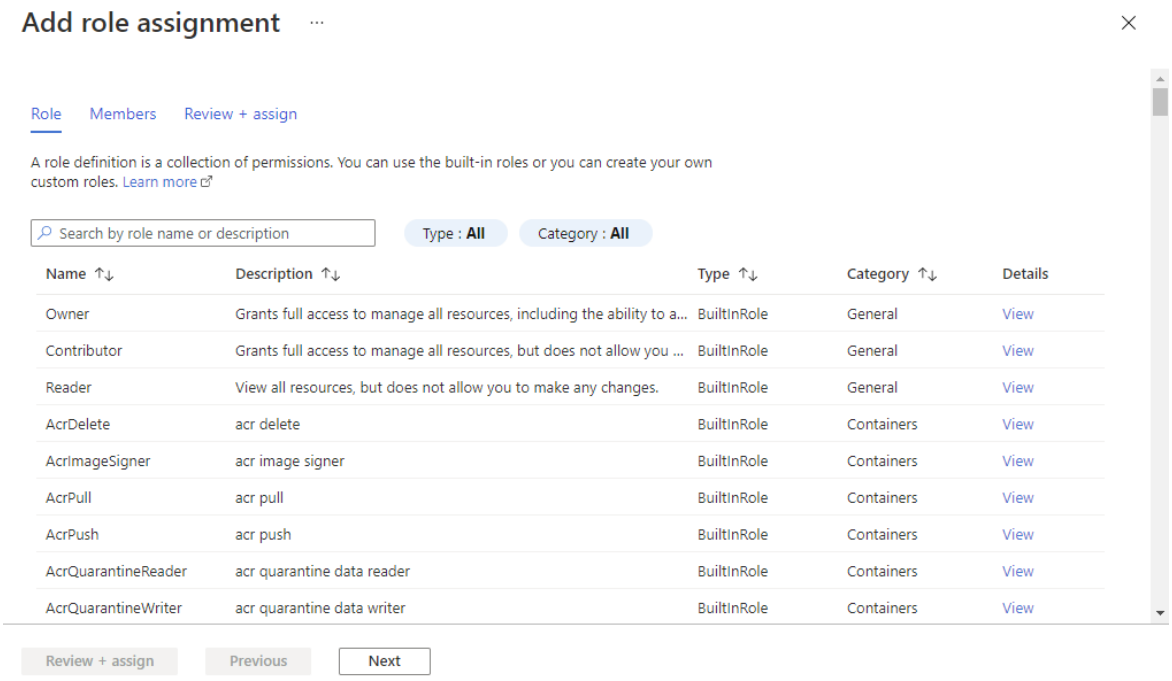

5. From the **Members** tab, select **+ Select members** and assign the role to the relevant members.
6. On the **Review + assign** tab, select **Review + assign** to assign the role.

Assign Azure roles to users with PowerShell:

1. Install Azure PowerShell.

2. Run the following commands:

```
# Login to Azure as a Global Administrator user
Connect-AzAccount
```

3. When prompted, sign in with global admin credentials.

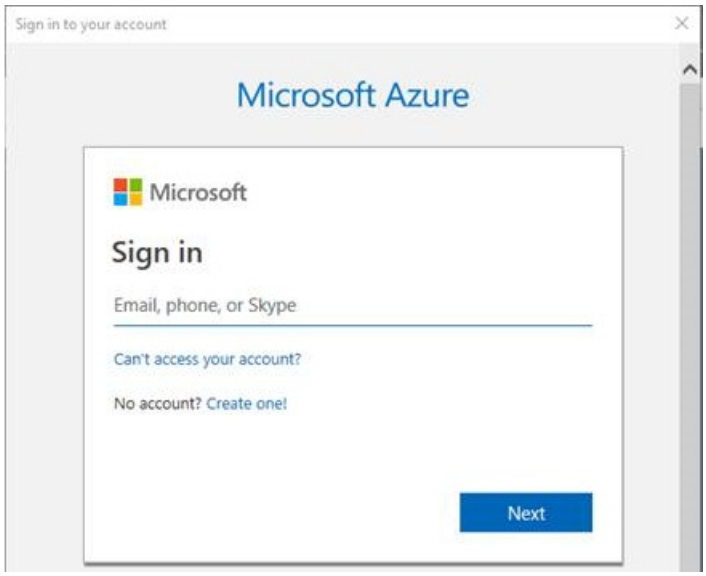

4. Grant reader role permissions by running the following command:

```
# Add Reader role to the required user on the Root Management Group
# Replace "user@domian.com" with the user to grant access to
New-AzRoleAssignment -SignInName "user@domain.com" -RoleDefinitionName "Reader" -Scope "/"
```

5. To remove the role, use the following command:

```
Remove-AzRoleAssignment -SignInName "user@domain.com" -RoleDefinitionName "Reader" -Scope "/"
```

Remove elevated access

Once the Azure roles have been assigned to the users, the tenant administrator should remove itself from the user access administrator role.

1. Sign in to the Azure portal or the Azure Active Directory admin center.
2. In the navigation list, select **Azure Active Directory** and then select **Properties**.
3. Under **Access management for Azure resources**, set the switch to **No**.
4. To save your setting, select **Save**.

Next steps

On this page, you learned how to organize subscriptions into management groups and assign roles to users. For related information, see:

- Permissions in Microsoft Defender for Cloud

Grant and request tenant-wide visibility

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

A user with the Azure Active Directory (AD) role of **Global Administrator** might have tenant-wide responsibilities, but lack the Azure permissions to view that organization-wide information in Microsoft Defender for Cloud. Permission elevation is required because Azure AD role assignments don't grant access to Azure resources.

Grant tenant-wide permissions to yourself

To assign yourself tenant-level permissions:

1. If your organization manages resource access with Azure AD Privileged Identity Management (PIM), or any other PIM tool, the global administrator role must be active for the user following the procedure below.

2. As a Global Administrator user without an assignment on the root management group of the tenant, open Defender for Cloud's **Overview** page and select the **tenant-wide visibility** link in the banner.

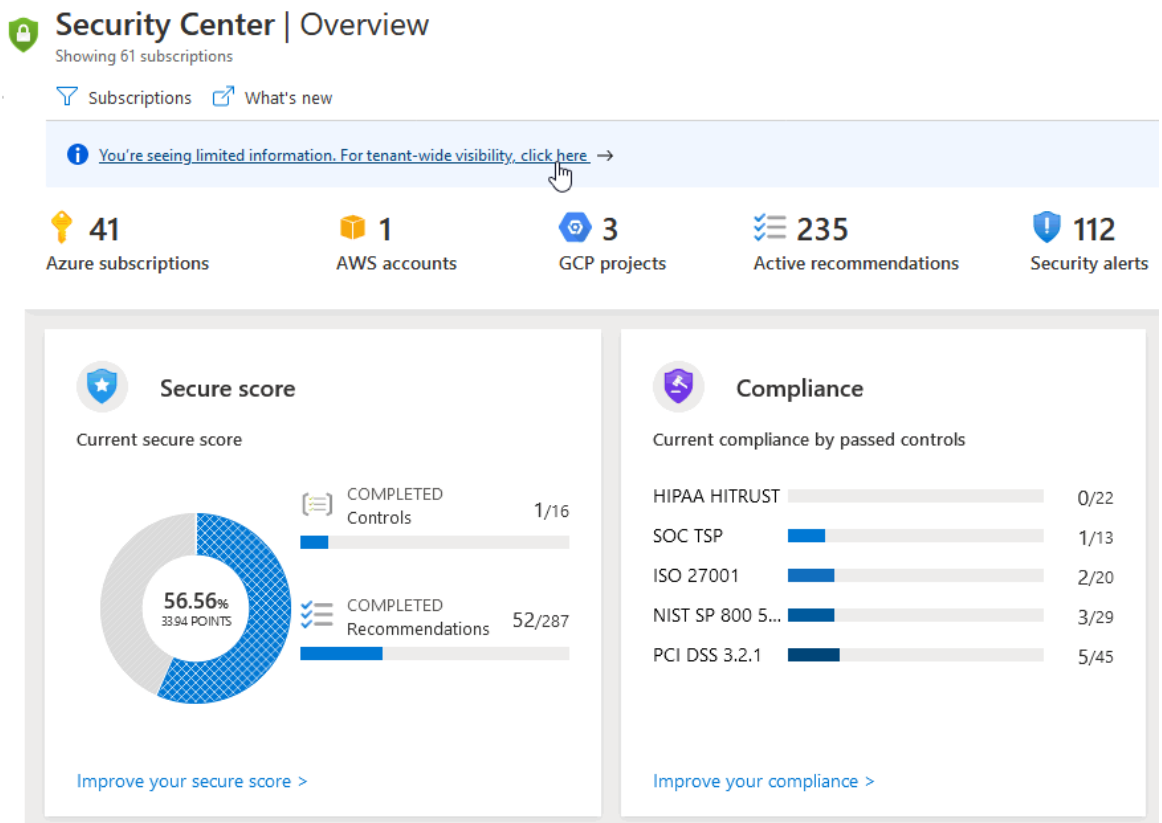

3. Select the new Azure role to be assigned.

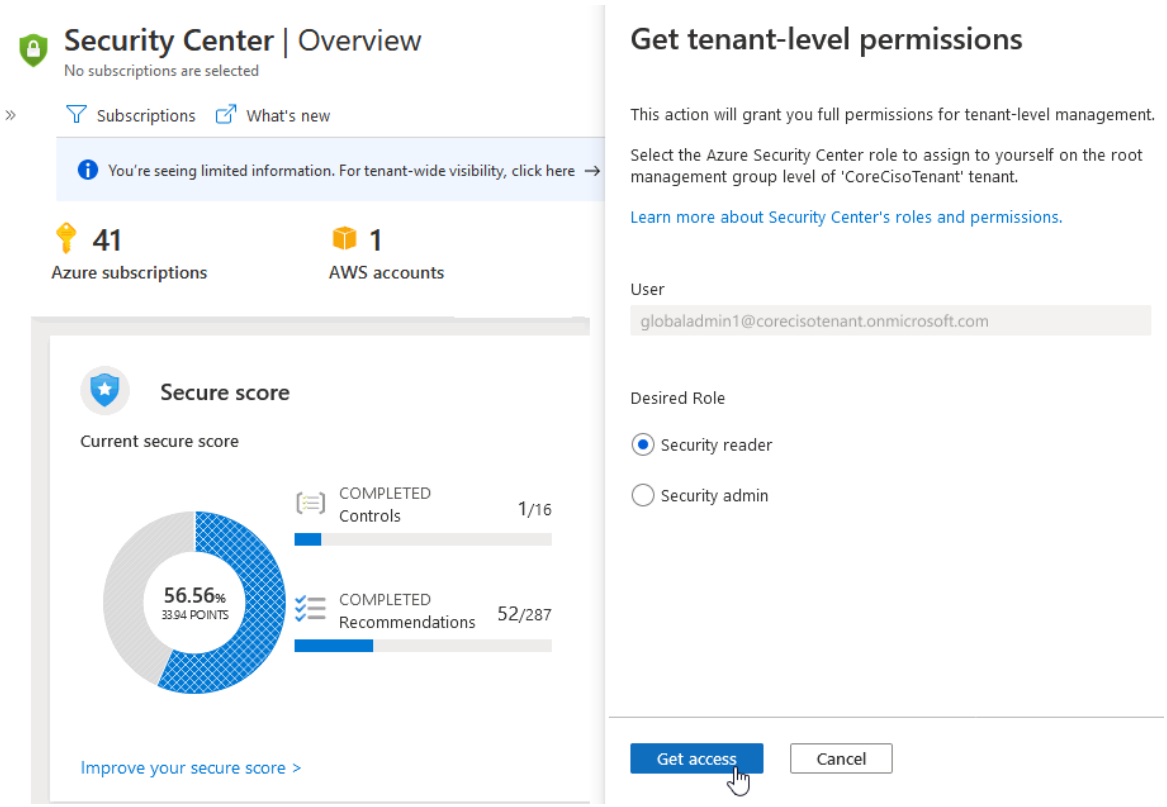

> **TIP**
>
> Generally, the Security Admin role is required to apply policies on the root level, while Security Reader will suffice to provide tenant-level visibility. For more information about the permissions granted by these roles, see the Security Admin built-in role description or the Security Reader built-in role description.
>
> For differences between these roles specific to Defender for Cloud, see the table in Roles and allowed actions.

The organizational-wide view is achieved by granting roles on the root management group level of the tenant.

4. Log out of the Azure portal, and then log back in again.
5. Once you have elevated access, open or refresh Microsoft Defender for Cloud to verify you have visibility into all subscriptions under your Azure AD tenant.

The simple process above performs a number of operations automatically for you:

1. The user's permissions are temporarily elevated.
2. Using the new permissions, the user is assigned to the desired Azure RBAC role on the root management group.
3. The elevated permissions are removed.

For more details of the Azure AD elevation process, see Elevate access to manage all Azure subscriptions and management groups.

Request tenant-wide permissions when yours are insufficient

If you login to Defender for Cloud and see a banner telling you that your view is limited, you can click through to send a request to the global administrator for your organization. In the request, you can include the role you'd like to be assigned and the global administrator will make a decision about which role to grant.

It's the global administrator's decision whether to accept or reject these requests.

IMPORTANT

You can only submit one request every seven days.

To request elevated permissions from your global administrator:

1. From the Azure portal, open Microsoft Defender for Cloud.
2. If you see the banner "You're seeing limited information." select it.

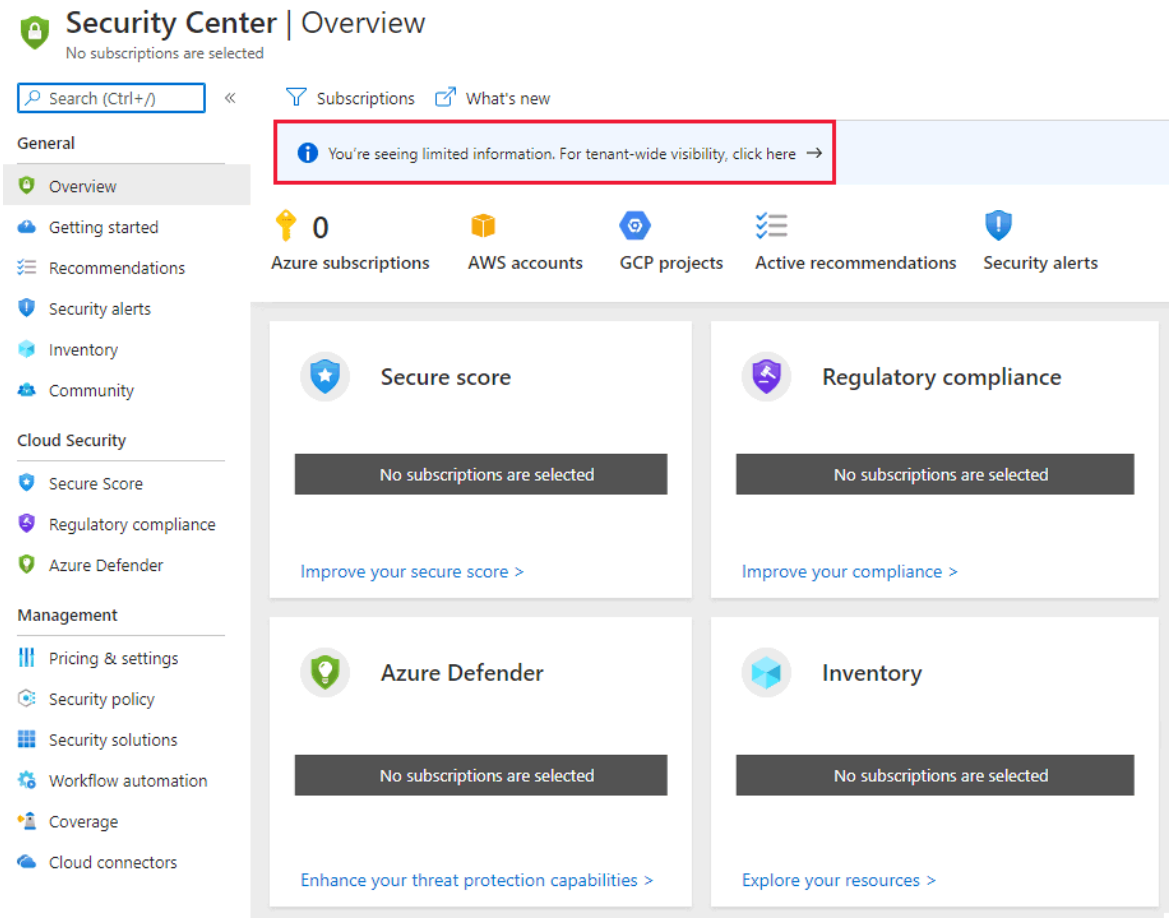

3. In the detailed request form, select the desired role and the justification for why you need these permissions.

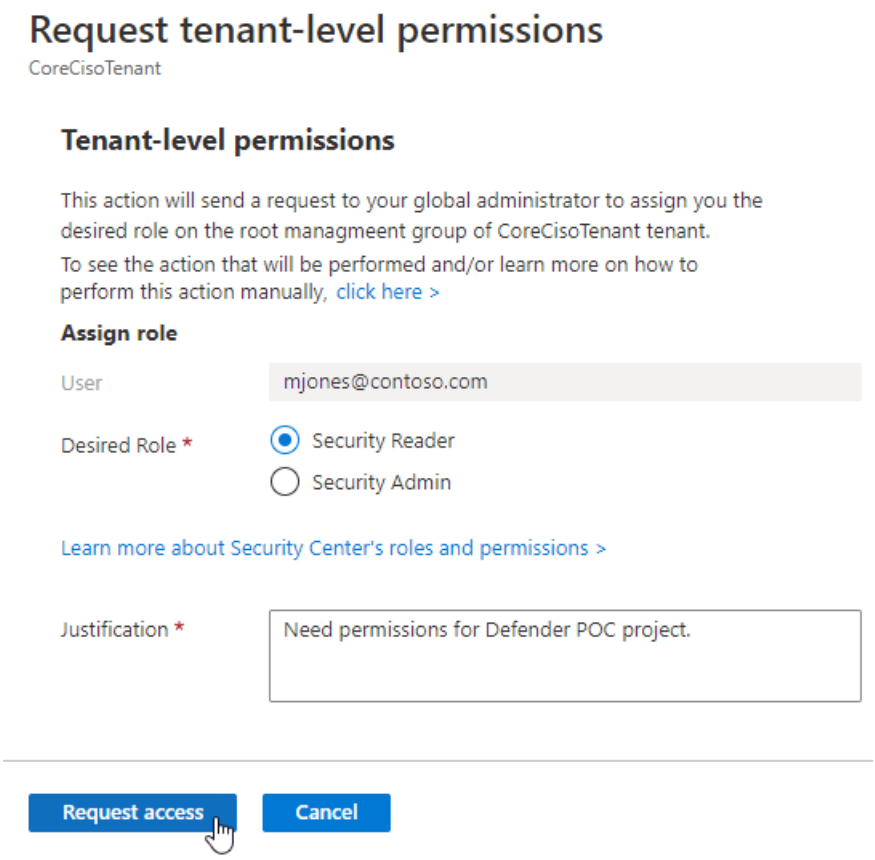

4. Select **Request access**.

   An email is sent to the global administrator. The email contains a link to Defender for Cloud where they can approve or reject the request.

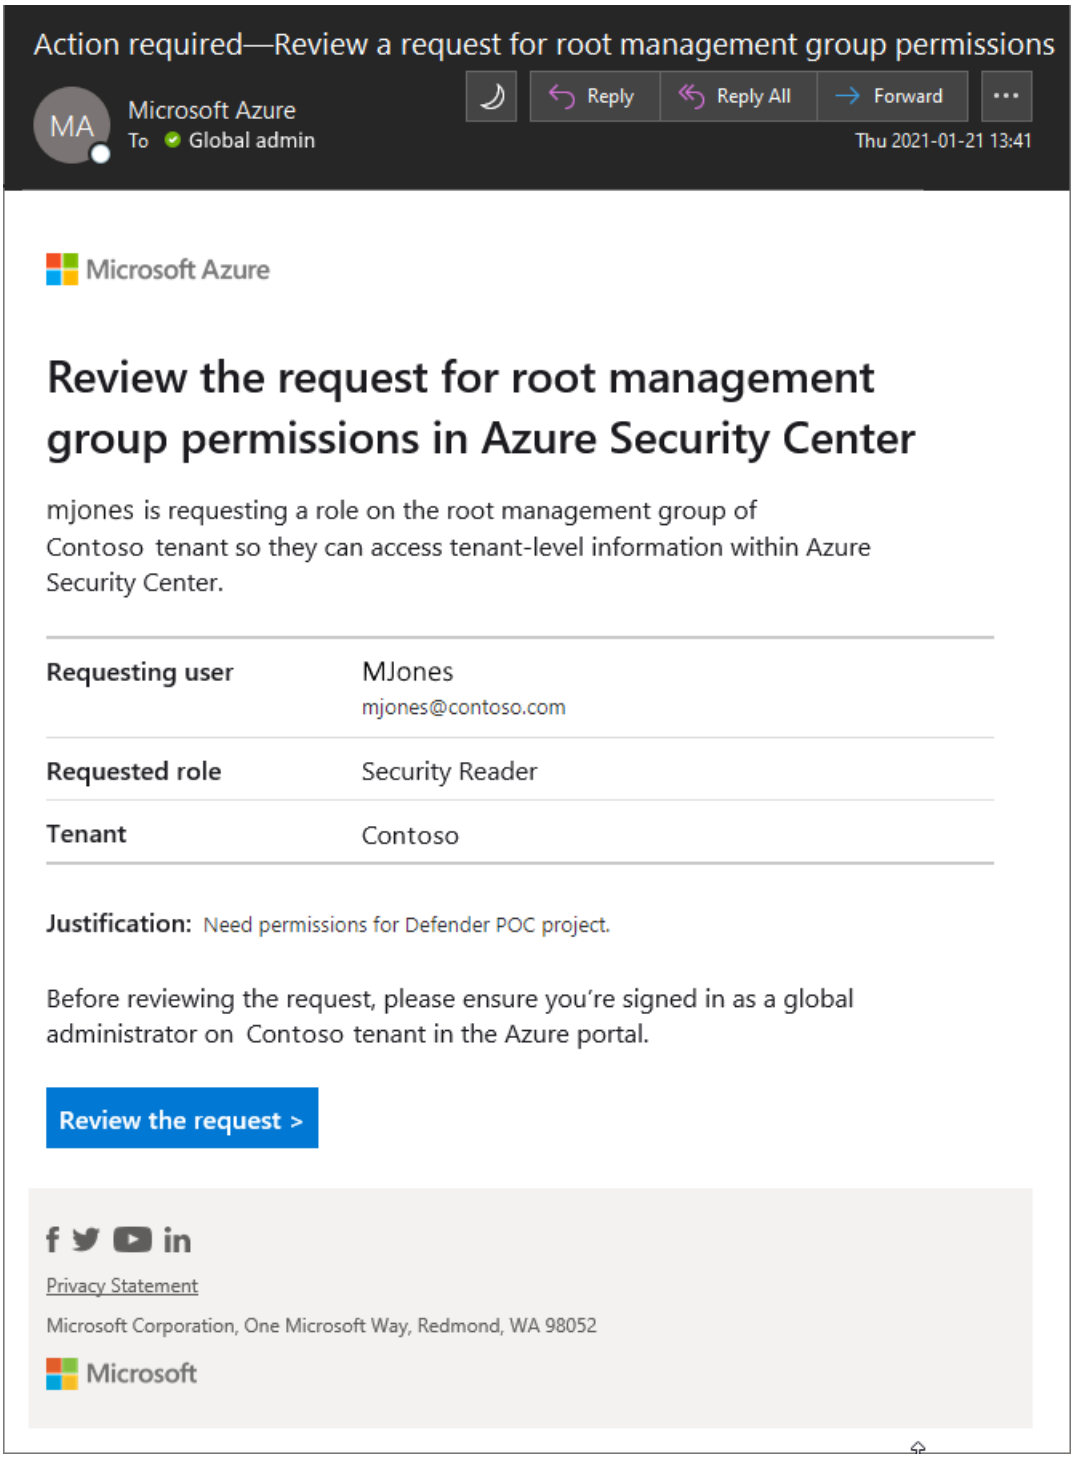

After the global administrator selects **Review the request** and completes the process, the decision is emailed to the requesting user.

Next steps

Learn more about Defender for Cloud permissions in the following related page:

- Permissions in Microsoft Defender for Cloud

Enable Defender for Cloud on all subscriptions in a management group

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

You can use Azure Policy to enable Microsoft Defender for Cloud on all the Azure subscriptions within the same management group (MG). This is more convenient than accessing them individually from the portal, and works even if the subscriptions belong to different owners.

To onboard a management group and all its subscriptions:

1. As a user with **Security Admin** permissions, open Azure Policy and search for the definition **Enable Defender for Cloud on your subscription**.

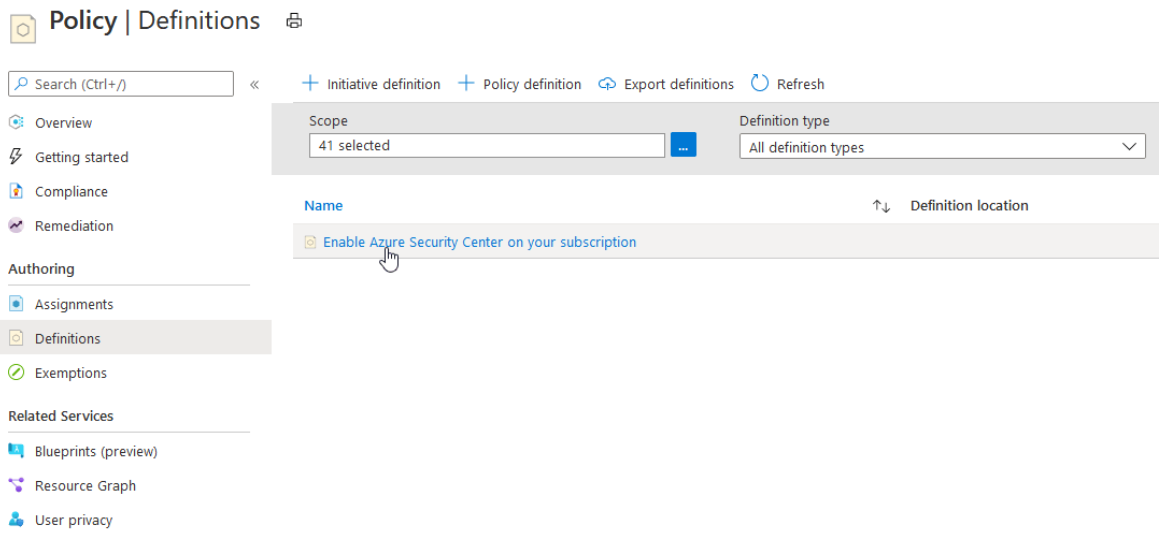

2. Select **Assign** and ensure you set the scope to the MG level.

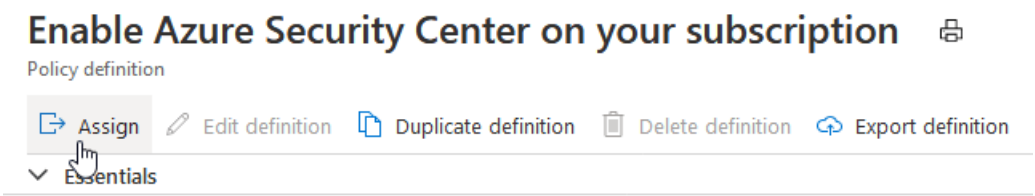

TIP

Other than the scope, there are no required parameters.

3. Select **Create a remediation task** to ensure all existing subscriptions that don't have Defender for Cloud enabled, will get onboarded.

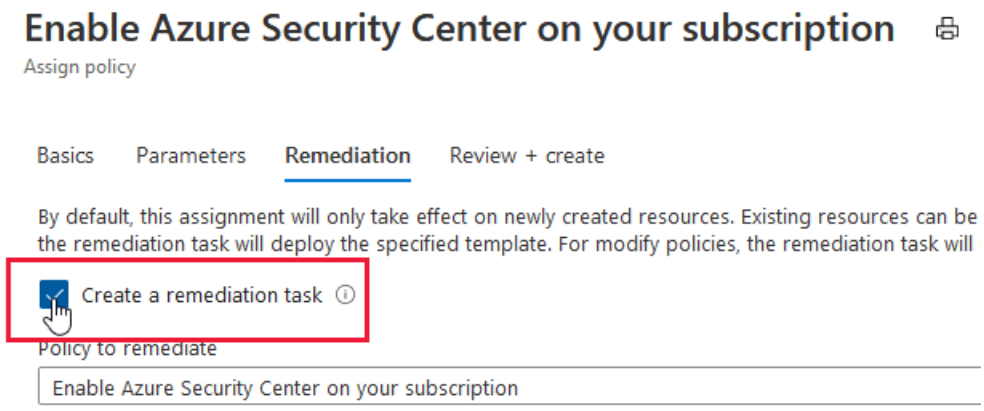

4. When the definition is assigned it will:

   a. Detect all subscriptions in the MG that aren't yet registered with Defender for Cloud.
   b. Mark those subscriptions as "non-compliant".
   c. Mark as "compliant" all registered subscriptions (regardless of whether they have Defender for Cloud's enhanced security features on or off).

   The remediation task will then enable Defender for Cloud, for free, on the non-compliant subscriptions.

> **IMPORTANT**
> The policy definition will only enable Defender for Cloud on **existing** subscriptions. To register newly created subscriptions, open the compliance tab, select the relevant non-compliant subscriptions, and create a remediation task.Repeat this step when you have one or more new subscriptions you want to monitor with Defender for Cloud.

Optional modifications

There are a variety of ways you might choose to modify the Azure Policy definition:

- **Define compliance differently** - The supplied policy classifies all subscriptions in the MG that aren't yet registered with Defender for Cloud as "non-compliant". You might choose to set it to all subscriptions without Defender for Cloud's enhanced security features enabled.

  The supplied definition, defines *either* of the 'pricing' settings below as compliant. Meaning that a subscription set to 'standard' or 'free' is compliant.

  > **TIP**
  > When any Microsoft Defender plan is enabled, it's described in a policy definition as being on the 'Standard' setting. When it's disabled, it's 'Free'. To learn about the differences between these plans, see Microsoft Defender for Cloud's enhanced security features.

```
"existenceCondition": {
    "anyof": [
        {
            "field": "microsoft.security/pricings/pricingTier",
            "equals": "standard"
        },
        {
            "field": "microsoft.security/pricings/pricingTier",
            "equals": "free"
        }
    ]
},
```

If you change it to the following, only subscriptions set to 'standard' would be classified as compliant:

```
"existenceCondition": {
      {
        "field": "microsoft.security/pricings/pricingTier",
        "equals": "standard"
      },
},
```

- **Define some Microsoft Defender plans to apply when enabling Defender for Cloud** - The supplied policy enables Defender for Cloud without any of the optional enhanced security features. You might choose to enable one or more of the Microsoft Defender plans.

  The supplied definition's `deployment` section has a parameter `pricingTier`. By default, this is set to `free`, but you can modify it.

Next steps:

Now that you've onboarded an entire management group, enable the enhanced security features.

Enable enhanced protections

Cross-tenant management in Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Cross-tenant management enables you to view and manage the security posture of multiple tenants in Defender for Cloud by leveraging Azure Lighthouse. Manage multiple tenants efficiently, from a single view, without having to sign in to each tenant's directory.

- Service providers can manage the security posture of resources, for multiple customers, from within their own tenant.
- Security teams of organizations with multiple tenants can view and manage their security posture from a single location.

Set up cross-tenant management

Azure delegated resource management is one of the key components of Azure Lighthouse. Set up cross-tenant management by delegating access to resources of managed tenants to your own tenant using these instructions from Azure Lighthouse's documentation: Onboard a customer to Azure Lighthouse.

How does cross-tenant management work in Defender for Cloud

You are able to review and manage subscriptions across multiple tenants in the same way that you manage multiple subscriptions in a single tenant.

From the top menu bar, click the filter icon, and select the subscriptions, from each tenant's directory, you'd like to view.

The views and actions are basically the same. Here are some examples:

- **Manage security policies**: From one view, manage the security posture of many resources with policies, take actions with security recommendations, and collect and manage security-related data.
- **Improve Secure Score and compliance posture**: Cross-tenant visibility enables you to view the overall security posture of all your tenants and where and how to best improve the secure score and compliance posture for each of them.
- **Remediate recommendations**: Monitor and remediate a recommendation for many resources from various tenants at one time. You can then immediately tackle the vulnerabilities that present the highest risk across all tenants.
- **Manage Alerts**: Detect alerts throughout the different tenants. Take action on resources that are out of

compliance with actionable remediation steps.

- **Manage advanced cloud defense features and more**: Manage the various threat protection services, such as just-in-time (JIT) VM access, Adaptive Network Hardening, adaptive application controls, and more.

Next steps

This article explains how cross-tenant management works in Defender for Cloud. To discover how Azure Lighthouse can simplify cross-tenant management within an enterprise which uses multiple Azure AD tenants, see Azure Lighthouse in enterprise scenarios.

Automate onboarding of Microsoft Defender for Cloud using PowerShell

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

You can secure your Azure workloads programmatically, using the Microsoft Defender for Cloud PowerShell module. Using PowerShell enables you to automate tasks and avoid the human error inherent in manual tasks. This is especially useful in large-scale deployments that involve dozens of subscriptions with hundreds and thousands of resources, all of which must be secured from the beginning.

Onboarding Microsoft Defender for Cloud using PowerShell enables you to programmatically automate onboarding and management of your Azure resources and add the necessary security controls.

This article provides a sample PowerShell script that can be modified and used in your environment to roll out Defender for Cloud across your subscriptions.

In this example, we will enable Defender for Cloud on a subscription with ID: d07c0080-170c-4c24-861d-9c817742786c and apply the recommended settings that provide a high level of protection, by enabling Microsoft Defender for Cloud's enhanced security features, which provides advanced threat protection and detection capabilities:

1. Enable the enhanced security in Microsoft Defender for Cloud.
2. Set the Log Analytics workspace to which the Log Analytics agent will send the data it collects on the VMs associated with the subscription – in this example, an existing user defined workspace (myWorkspace).
3. Activate Defender for Cloud's automatic agent provisioning which deploys the Log Analytics agent.
4. Set the organization's CISO as the security contact for Defender for Cloud alerts and notable events.
5. Assign Defender for Cloud's default security policies.

Prerequisites

These steps should be performed before you run the Defender for Cloud cmdlets:

1. Run PowerShell as admin.
2. Run the following commands in PowerShell:

```
Set-ExecutionPolicy -ExecutionPolicy AllSigned
```

```
Install-Module -Name Az.Security -Force
```

Onboard Defender for Cloud using PowerShell

1. Register your subscriptions to the Defender for Cloud Resource Provider:

```
Set-AzContext -Subscription "d07c0080-170c-4c24-861d-9c817742786c"
```

```
Register-AzResourceProvider -ProviderNamespace 'Microsoft.Security'
```

2. Optional: Set the coverage level (Microsoft Defender for Cloud's enhanced security features on/off) of the subscriptions. If undefined, these features are off:

```
Set-AzContext -Subscription "d07c0080-170c-4c24-861d-9c817742786c"
```

```
Set-AzSecurityPricing -Name "VirtualMachines" -PricingTier "Standard"
```

3. Configure a Log Analytics workspace to which the agents will report. You must have a Log Analytics workspace that you already created, that the subscription's VMs will report to. You can define multiple subscriptions to report to the same workspace. If not defined, the default workspace will be used.

```
Set-AzSecurityWorkspaceSetting -Name "default" -Scope "/subscriptions/d07c0080-170c-4c24-861d-
9c817742786c" -WorkspaceId"/subscriptions/d07c0080-170c-4c24-861d-
9c817742786c/resourceGroups/myRg/providers/Microsoft.OperationalInsights/workspaces/myWorkspace"
```

4. Auto-provision installation of the Log Analytics agent on your Azure VMs:

```
Set-AzContext -Subscription "d07c0080-170c-4c24-861d-9c817742786c"
```

```
Set-AzSecurityAutoProvisioningSetting -Name "default" -EnableAutoProvision
```

> **NOTE**
> We recommend that you enable auto provisioning to make sure that your Azure virtual machines are automatically protected by Microsoft Defender for Cloud.

5. Optional: It is highly recommended that you define the security contact details for the subscriptions you onboard, which will be used as the recipients of alerts and notifications generated by Defender for Cloud:

```
Set-AzSecurityContact -Name "default1" -Email "CISO@my-org.com" -AlertAdmin -NotifyOnAlert
```

6. Assign the default Defender for Cloud policy initiative:

```
Register-AzResourceProvider -ProviderNamespace 'Microsoft.PolicyInsights'
```

```
$Policy = Get-AzPolicySetDefinition | where {$_.Properties.displayName -EQ 'Azure Security
Benchmark'}

New-AzPolicyAssignment -Name 'ASC Default <d07c0080-170c-4c24-861d-9c817742786c>' -DisplayName
'Defender for Cloud Default <subscription ID>' -PolicySetDefinition $Policy -Scope
'/subscriptions/d07c0080-170c-4c24-861d-9c817742786c'
```

You've successfully onboarded Microsoft Defender for Cloud with PowerShell.

You can now use these PowerShell cmdlets with automation scripts to programmatically iterate across subscriptions and resources. This saves time and reduces the likelihood of human error. You can use this sample script as reference.

See also

To learn more about how you can use PowerShell to automate onboarding to Defender for Cloud, see the following article:

- Az.Security

To learn more about Defender for Cloud, see the following articles:

- Setting security policies in Microsoft Defender for Cloud -- Learn how to configure security policies for your Azure subscriptions and resource groups.
- Managing and responding to security alerts in Microsoft Defender for Cloud -- Learn how to manage and respond to security alerts.

Manage and respond to security alerts in Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This topic shows you how to view and process Defender for Cloud's alerts and protect your resources.

Advanced detections that trigger security alerts are only available with Microsoft Defender for Cloud's enhanced security features enabled. A free trial is available. To upgrade, see Enable enhanced protections.

What are security alerts?

Defender for Cloud automatically collects, analyzes, and integrates log data from your Azure resources, the network, and connected partner solutions - like firewall and endpoint protection solutions - to detect real threats and reduce false positives. A list of prioritized security alerts is shown in Defender for Cloud along with the information you need to quickly investigate the problem and steps to take to remediate an attack.

To learn about the different types of alerts, see Security alerts - a reference guide.

For an overview of how Defender for Cloud generates alerts, see How Microsoft Defender for Cloud detects and responds to threats.

Manage your security alerts

1. From Defender for Cloud's overview page, select the **Security alerts** tile at the top of the page, or the link from the sidebar..

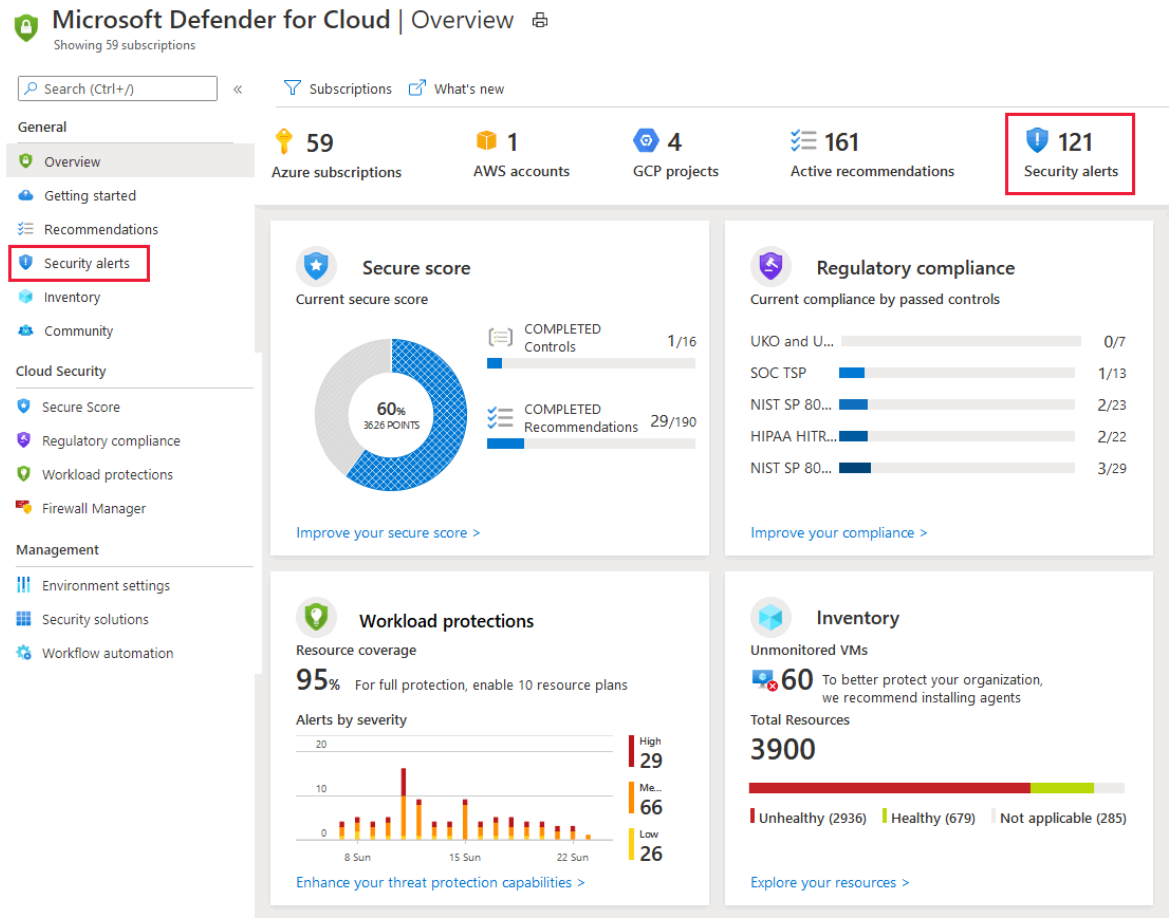

The security alerts page opens.

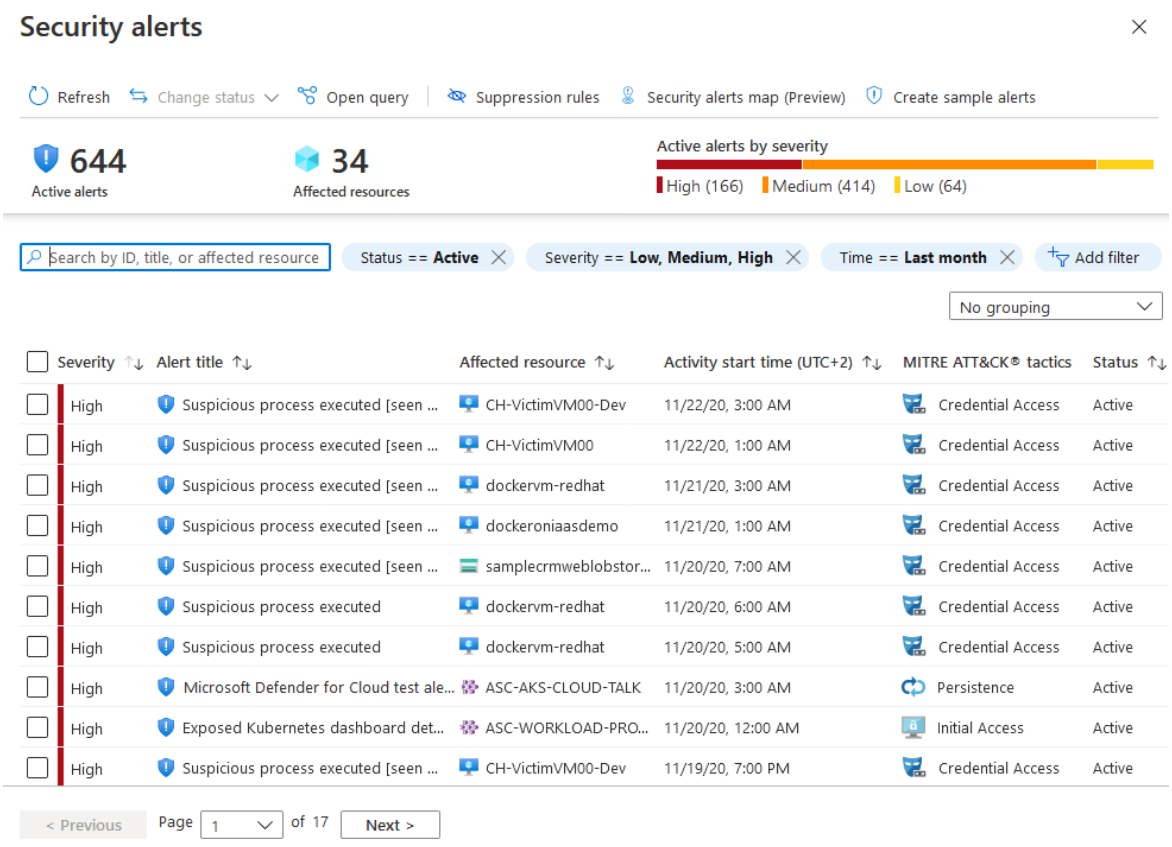

2. To filter the alerts list, select any of the relevant filters. You can optionally add further filters with the **Add filter** option.

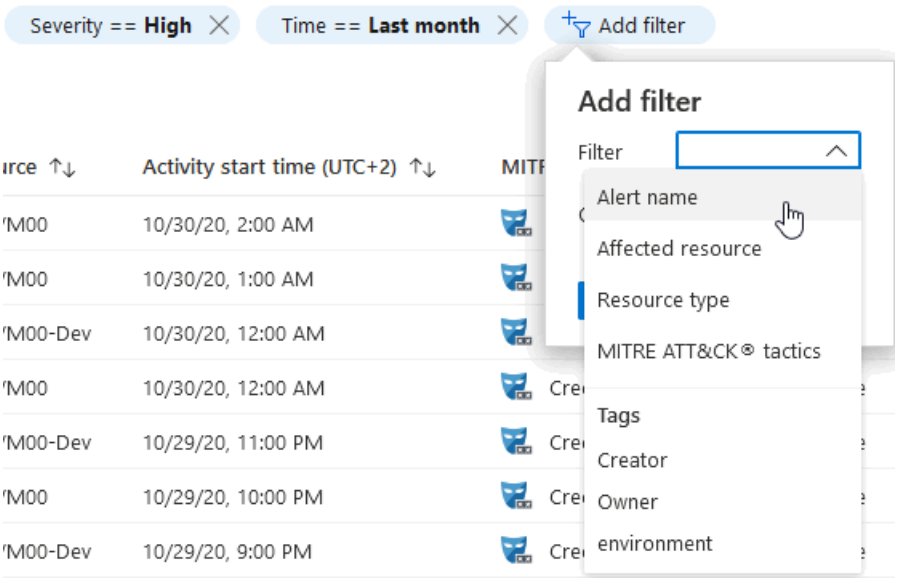

The list updates according to the filtering options you've selected. Filtering can be very helpful. For example, you might you want to address security alerts that occurred in the last 24 hours because you are investigating a potential breach in the system.

Respond to security alerts

1. From the **Security alerts** list, select an alert. A side pane opens and shows a description of the alert and all the affected resources.

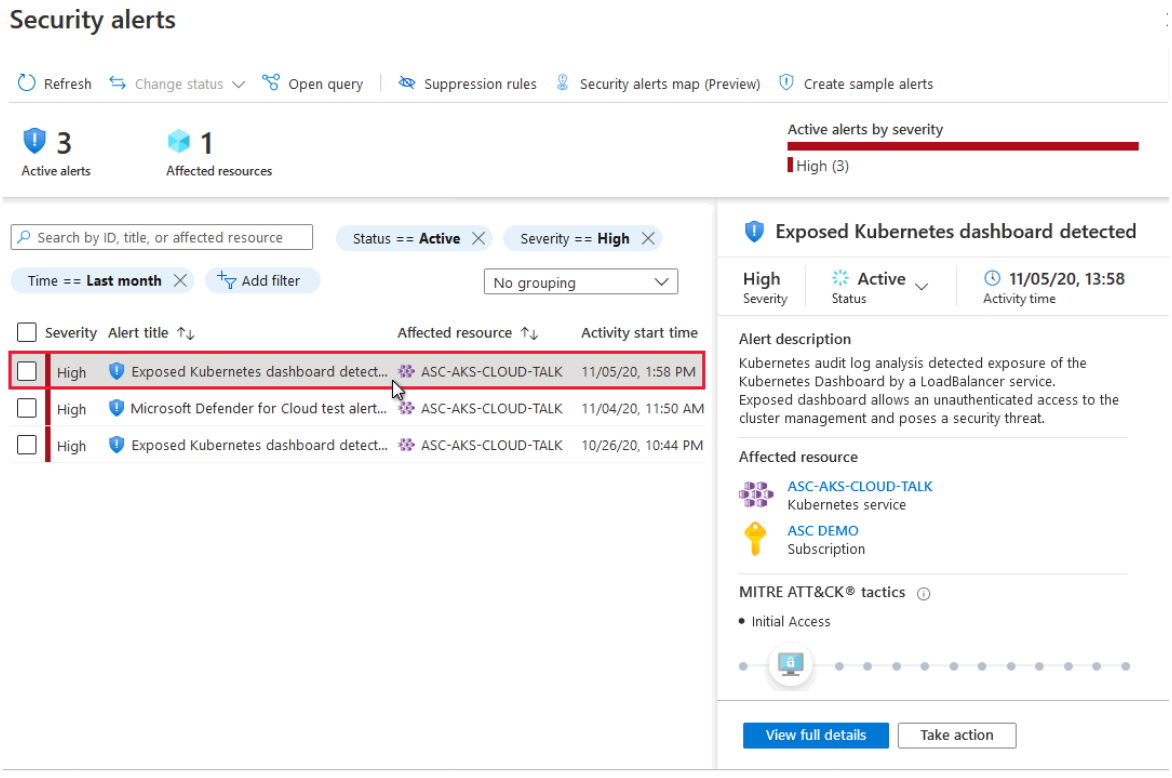

> **TIP**
>
> With this side pane open, you can quickly review the alerts list with the up and down arrows on your keyboard.

2. For further information, select **View full details**.

The left pane of the security alert page shows high-level information regarding the security alert: title, severity, status, activity time, description of the suspicious activity, and the affected resource. Alongside

the affected resource are the Azure tags relevant to the resource. Use these to infer the organizational context of the resource when investigating the alert.

The right pane includes the **Alert details** tab containing further details of the alert to help you investigate the issue: IP addresses, files, processes, and more.

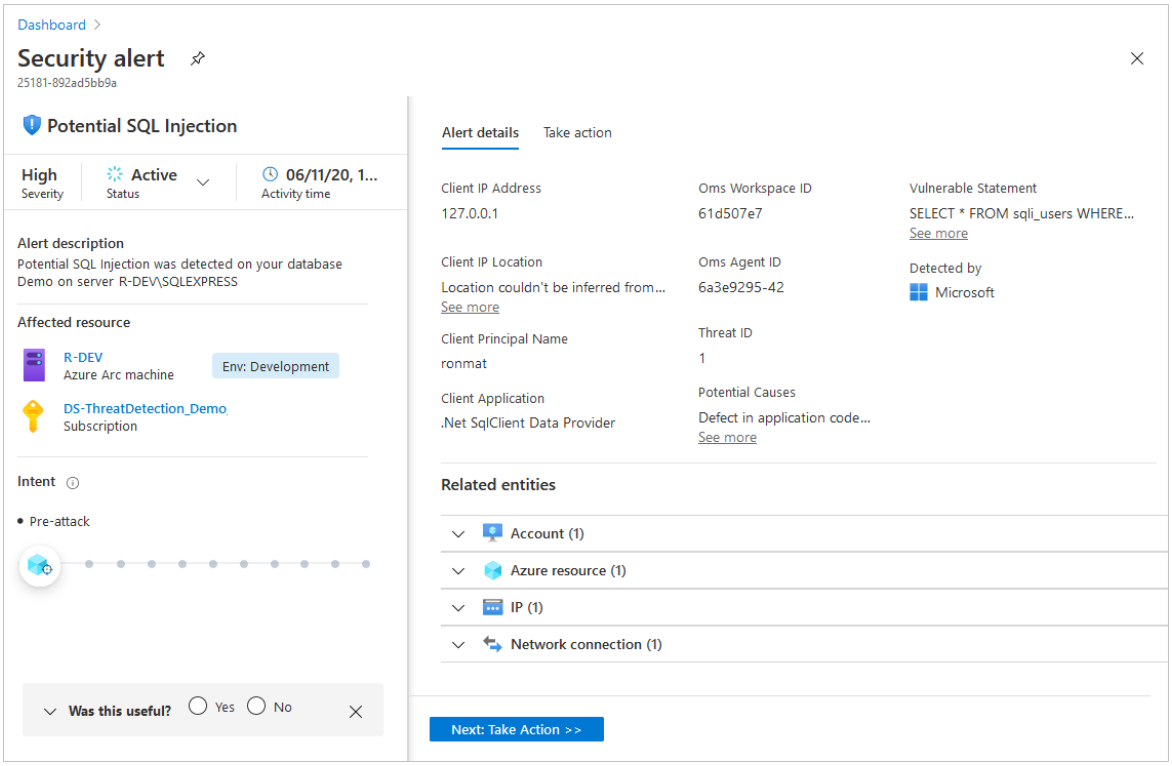

Also in the right pane is the **Take action** tab. Use this tab to take further actions regarding the security alert. Actions such as:

- *Mitigate the threat* - provides manual remediation steps for this security alert
- *Prevent future attacks* - provides security recommendations to help reduce the attack surface, increase security posture, and thus prevent future attacks
- *Trigger automated response* - provides the option to trigger a logic app as a response to this security alert
- *Suppress similar alerts* - provides the option to suppress future alerts with similar characteristics if the alert isn't relevant for your organization

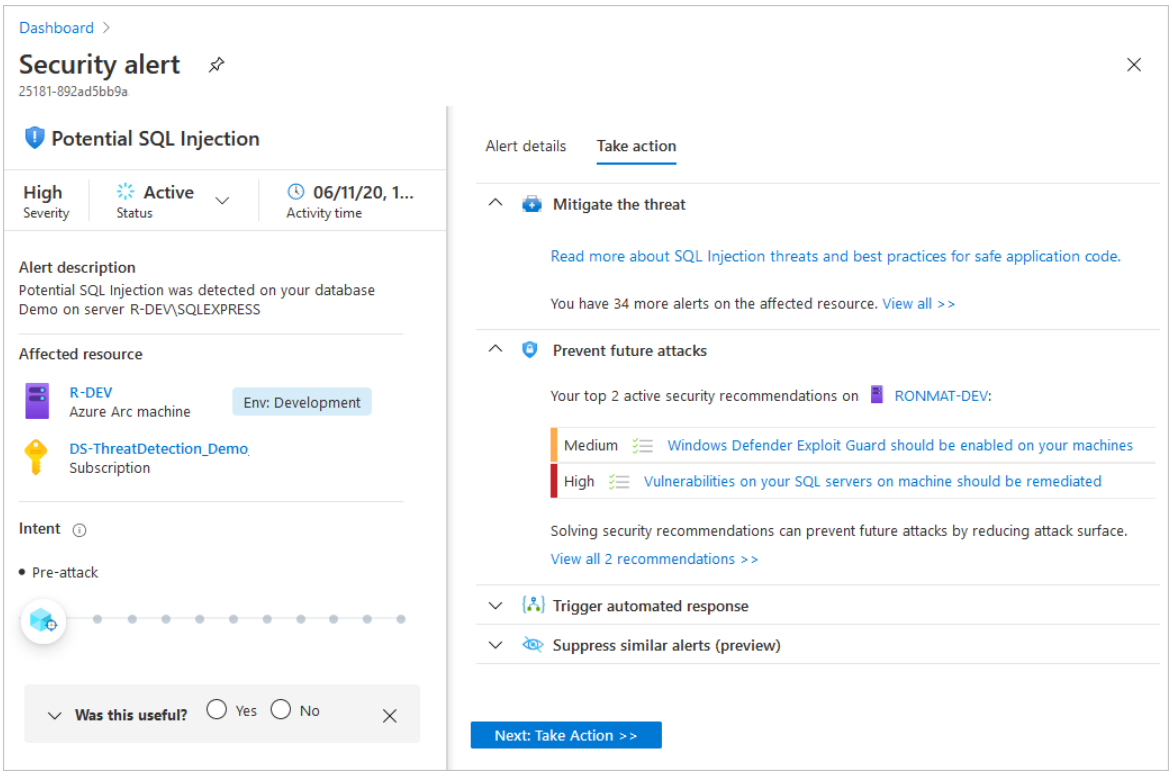

Change the status of multiple security alerts at once

The alerts list includes checkboxes so you can handle multiple alerts at once. For example, for triaging purposes you might decide to dismiss all informational alerts for a specific resource.

1. Filter according to the alerts you want to handle in bulk.

   In this example, we've selected all alerts with severity of 'Informational' for the resource 'ASC-AKS-CLOUD-TALK'.

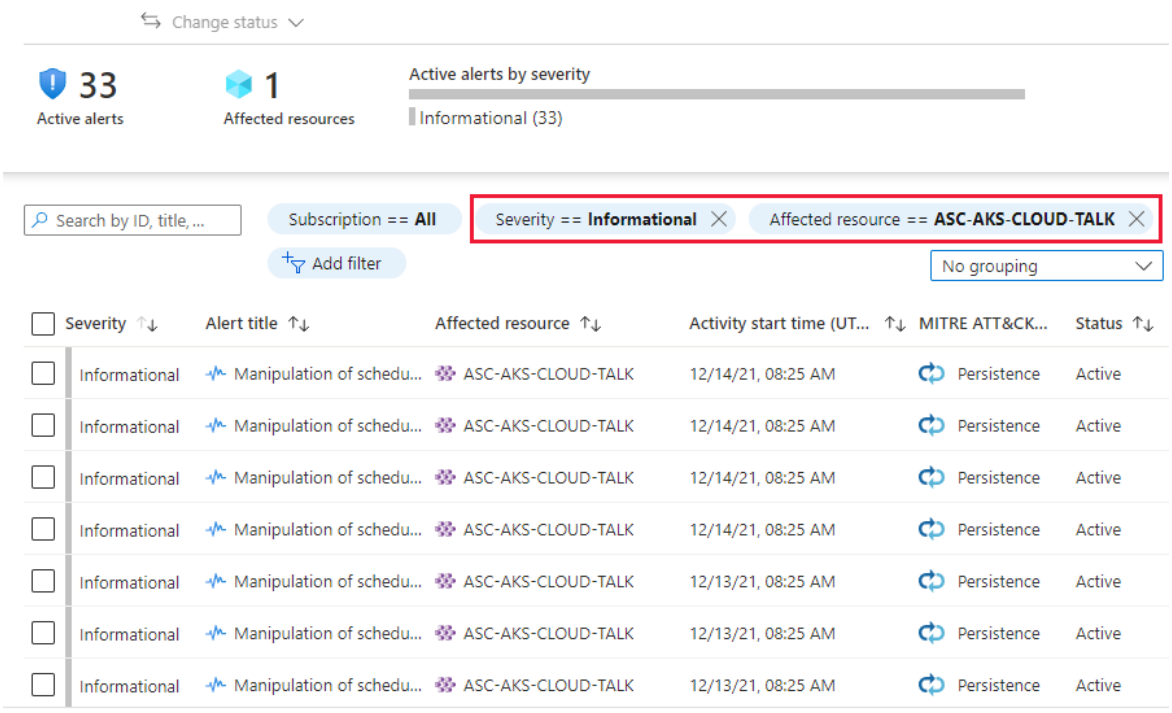

2. Use the checkboxes to select the alerts to be processed - or use the checkbox at the top of the list to select them all.

   In this example, we've selected all alerts. Notice that the **Change status** button is now available.

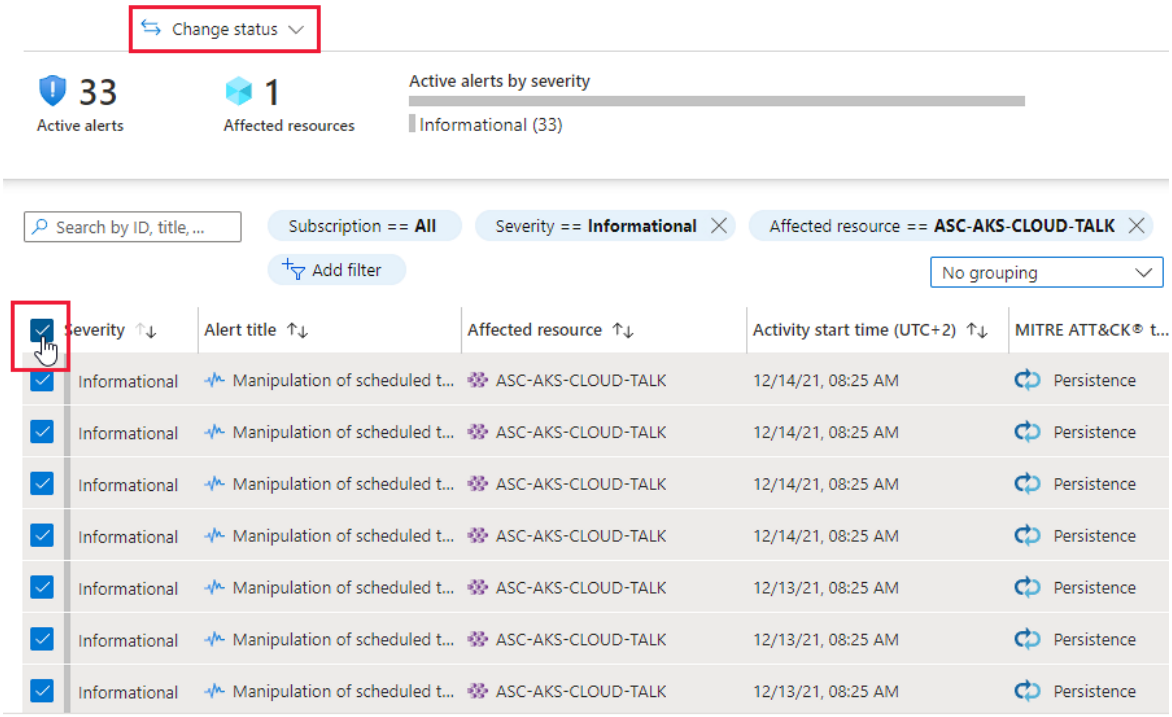

3. Use the **Change status** options to set the desired status.

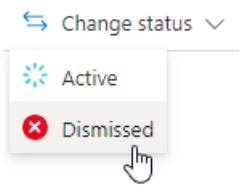

The alerts shown in the current page will have their status changed to the selected value.

See also

In this document, you learned how to view security alerts. See the following pages for related material:

- Configure alert suppression rules
- Automate responses to Defender for Cloud triggers
- Security alerts - a reference guide

Suppress alerts from Microsoft Defender for Cloud

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how you can use alerts suppression rules to suppress false positives or other unwanted security alerts from Defender for Cloud.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free (Most security alerts are only available with enhanced security features)
Required roles and permissions:	**Security admin** and **Owner** can create/delete rules. **Security reader** and **Reader** can view rules.
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

What are suppression rules?

The various Microsoft Defender plans detect threats in any area of your environment and generate security alerts.

When a single alert isn't interesting or relevant, you can manually dismiss it. Alternatively, use the suppression rules feature to automatically dismiss similar alerts in the future. Typically, you'd use a suppression rule to:

- Suppress alerts that you've identified as false positives
- Suppress alerts that are being triggered too often to be useful

Your suppression rules define the criteria for which alerts should be automatically dismissed.

Caution

Suppressing security alerts reduces the effectiveness of Defender for Cloud's threat protection. You should carefully check the potential impact of any suppression rule, and monitor it over time.

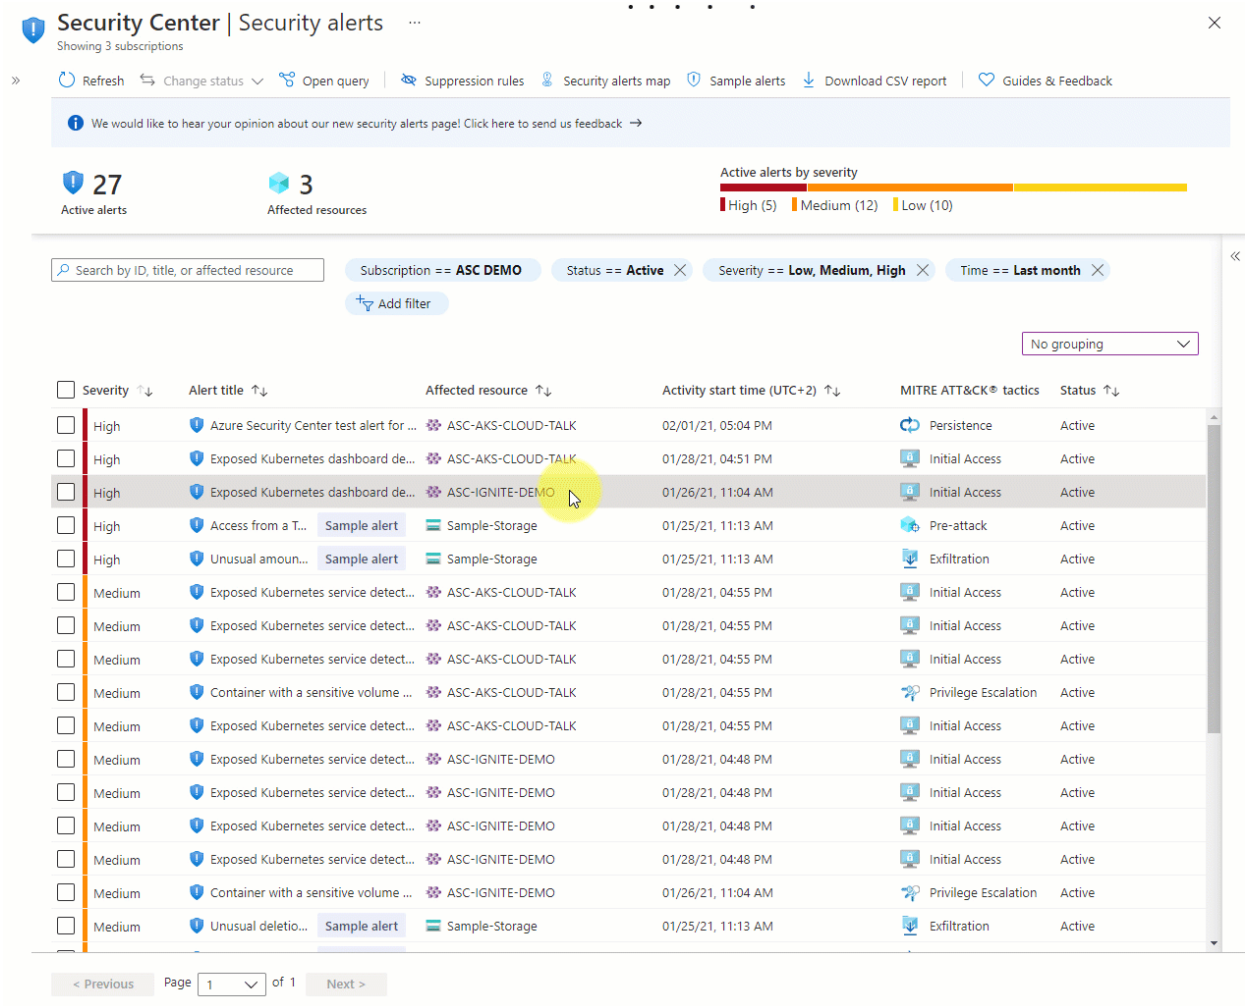

Create a suppression rule

There are a few ways you can create rules to suppress unwanted security alerts:

- To suppress alerts at the management group level, use Azure Policy
- To suppress alerts at the subscription level, you can use the Azure portal or the REST API as explained below

Suppression rules can only dismiss alerts that have already been triggered on the selected subscriptions.

To create a rule directly in the Azure portal:

1. From Defender for Cloud's security alerts page:

    - Select the specific alert you don't want to see anymore, and from the details pane, select **Take action**.

    - Or, select the **suppression rules** link at the top of the page, and from the suppression rules page select **Create new suppression rule**:

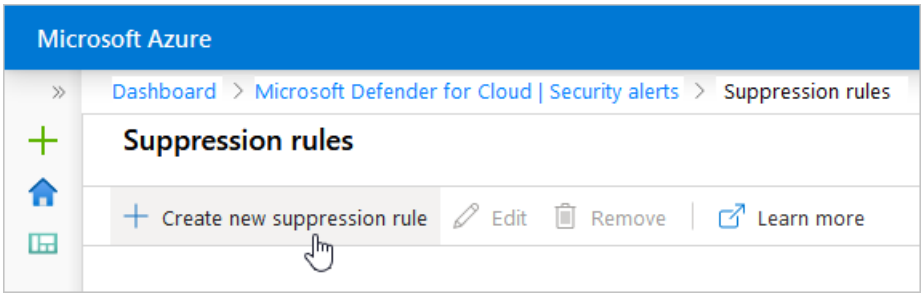

2. In the new suppression rule pane, enter the details of your new rule.

- Your rule can dismiss the alert on **all resources** so you don't get any alerts like this one in the future.
- Your rule can dismiss the alert **on specific criteria** - when it relates to a specific IP address, process name, user account, Azure resource, or location.

> **TIP**
> If you opened the new rule page from a specific alert, the alert and subscription will be automatically configured in your new rule. If you used the **Create new suppression rule** link, the selected subscriptions will match the current filter in the portal.

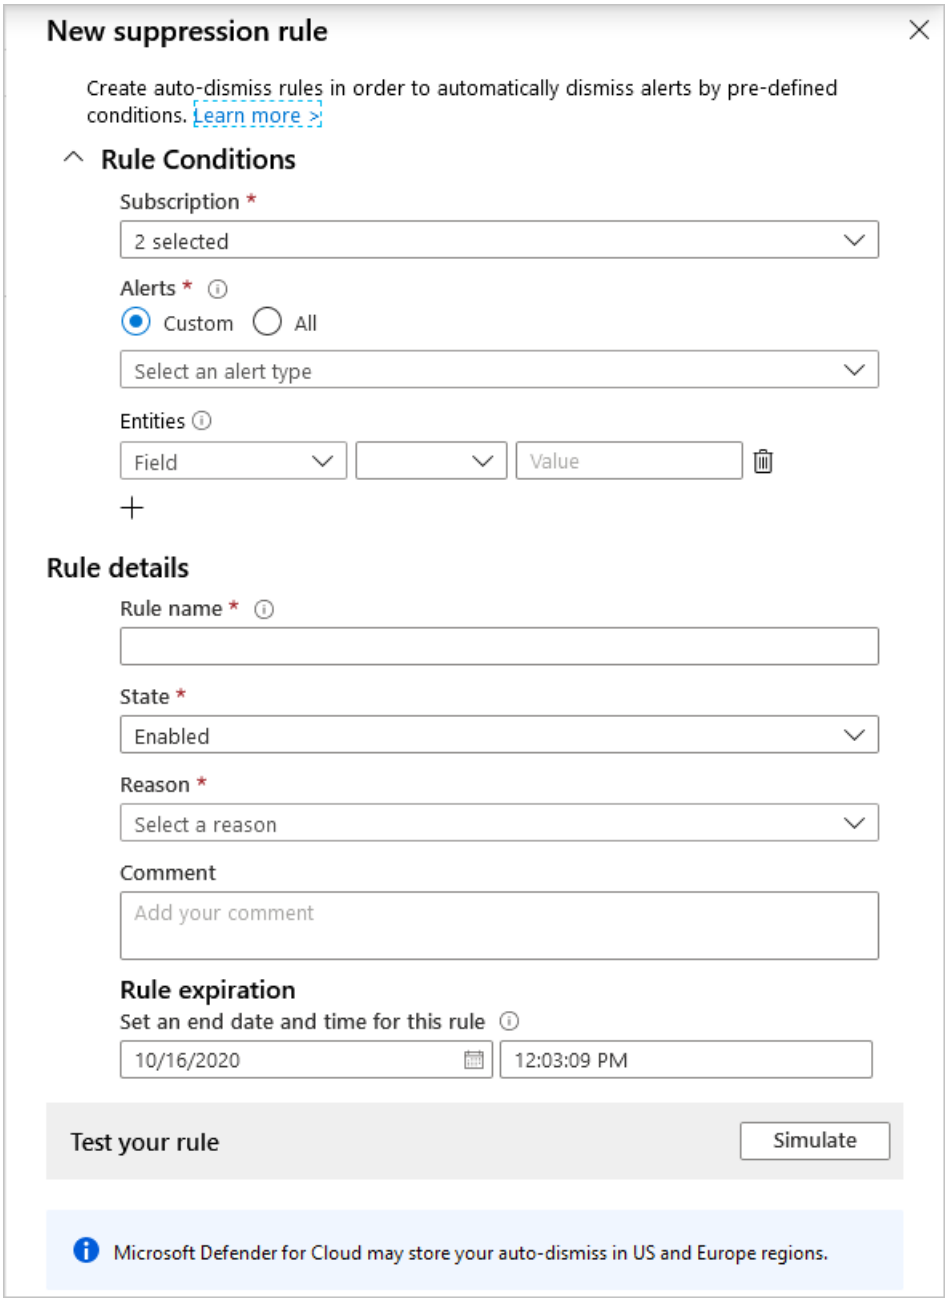

3. Enter details of the rule:

   - **Name** - A name for the rule. Rule names must begin with a letter or a number, be between 2 and 50 characters, and contain no symbols other than dashes (-) or underscores (_).
   - **State** - Enabled or disabled.
   - **Reason** - Select one of the built-in reasons or 'other' if they don't meet your needs.
   - **Expiration date** - An end date and time for the rule. Rules can run for up to six months.

4. Optionally, test the rule using the **Simulate** button to see how many alerts would have been dismissed if this rule had been active.

5. Save the rule.

Edit a suppression rule

To edit a rule you've created, use the suppression rules page.

1. From Defender for Cloud's security alerts page, select the **suppression rules** link at the top of the page.
2. The suppression rules page opens with all the rules for the selected subscriptions.

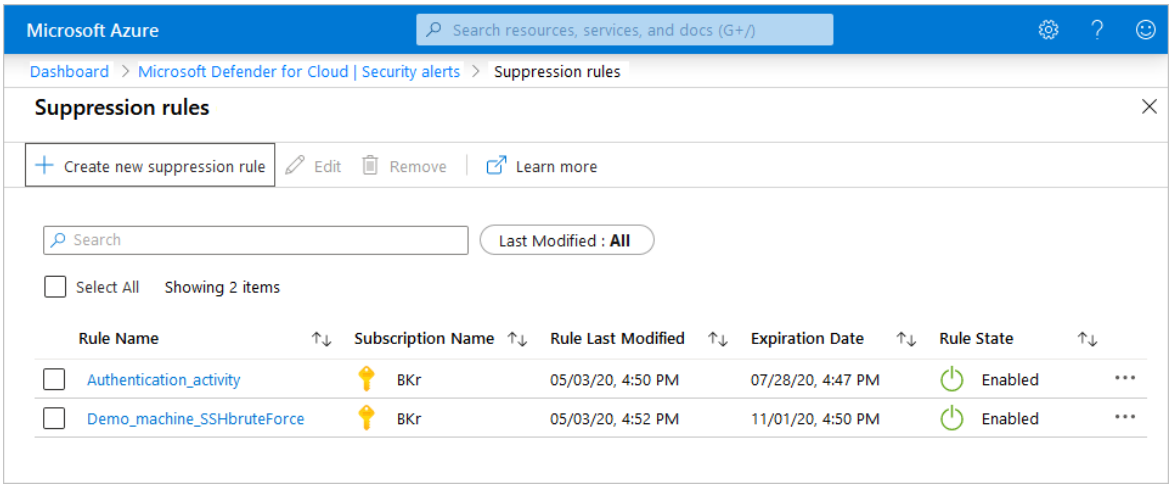

3. To edit a single rule, open the ellipsis menu (...) for the rule and select **Edit**.
4. Make the necessary changes and select **Apply**.

Delete a suppression rule

To delete one or more rules you've created, use the suppression rules page.

1. From Defender for Cloud's security alerts page, select the **suppression rules** link at the top of the page.
2. The suppression rules page opens with all the rules for the selected subscriptions.
3. To delete a single rule, open the ellipsis menu (...) for the rule and select **Delete**.
4. To delete multiple rules, select the check boxes for the rules to be deleted and select **Delete**.

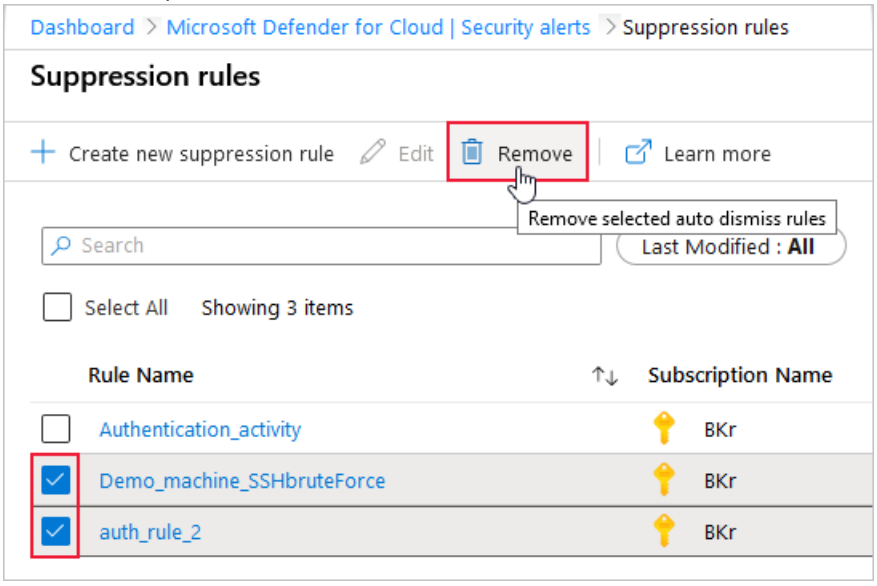

Create and manage suppression rules with the API

You can create, view, or delete alert suppression rules via Defender for Cloud's REST API.

The relevant HTTP methods for suppression rules in the REST API are:

- **PUT**: To create or update a suppression rule in a specified subscription.
- **GET**:
  - To list all rules configured for a specified subscription. This method returns an array of the applicable rules.
  - To get the details of a specific rule on a specified subscription. This method returns one suppression rule.
  - To simulate the impact of a suppression rule still in the design phase. This call identifies which of your existing alerts would have been dismissed if the rule had been active.
- **DELETE**: Deletes an existing rule (but doesn't change the status of alerts already dismissed by it).

For full details and usage examples, see the API documentation.

Next steps

This article described the suppression rules in Microsoft Defender for Cloud that automatically dismiss unwanted alerts.

For more information on security alerts, see the following pages:

- Security alerts and the intent kill chain - A reference guide to the security alerts you might get from Defender for Cloud.

Stream alerts to a SIEM, SOAR, or IT Service Management solution

2/15/2022 • 4 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud can stream your security alerts into the most popular Security Information and Event Management (SIEM), Security Orchestration Automated Response (SOAR), and IT Service Management (ITSM) solutions.

There are Azure-native tools for ensuring you can view your alert data in all of the most popular solutions in use today, including:

- **Microsoft Sentinel**
- **Splunk Enterprise and Splunk Cloud**
- **IBM's QRadar**
- **ServiceNow**
- **ArcSight**
- **Power BI**
- **Palo Alto Networks**

Stream alerts to Microsoft Sentinel

Defender for Cloud natively integrates with Microsoft Sentinel, Azure's cloud-native SIEM and SOAR solution.

Learn more about Microsoft Sentinel.

Microsoft Sentinel's connectors for Defender for Cloud

Microsoft Sentinel includes built-in connectors for Microsoft Defender for Cloud at the subscription and tenant levels:

- Stream alerts to Microsoft Sentinel at the subscription level
- Connect all subscriptions in your tenant to Microsoft Sentinel

When you connect Defender for Cloud to Microsoft Sentinel, the status of Defender for Cloud alerts that get ingested into Microsoft Sentinel is synchronized between the two services. So, for example, when an alert is closed in Defender for Cloud, that alert will display as closed in Microsoft Sentinel as well. Changing the status of an alert in Defender for Cloud "won't"* affect the status of any Microsoft Sentinel **incidents** that contain the synchronized Microsoft Sentinel alert, only that of the synchronized alert itself.

Enabling the preview feature, **bi-directional alert synchronization**, will automatically sync the status of the original Defender for Cloud alerts with Microsoft Sentinel incidents that contain the copies of those Defender for Cloud alerts. So, for example, when a Microsoft Sentinel incident containing a Defender for Cloud alert is closed, Defender for Cloud will automatically close the corresponding original alert.

Learn more in Connect alerts from Microsoft Defender for Cloud.

> **NOTE**
> The bi-directional alert synchronization feature isn't available in the Azure Government cloud.

Configure ingestion of all audit logs into Microsoft Sentinel

Another alternative for investigating Defender for Cloud alerts in Microsoft Sentinel is to stream your audit logs into Microsoft Sentinel: - Connect Windows security events - Collect data from Linux-based sources using Syslog - Connect data from Azure Activity log

> **TIP**
> Microsoft Sentinel is billed based on the volume of data ingested for analysis in Microsoft Sentinel and stored in the Azure Monitor Log Analytics workspace. Microsoft Sentinel offers a flexible and predictable pricing model. Learn more at the Microsoft Sentinel pricing page.

Stream alerts with Azure Monitor

To stream alerts into **ArcSight**, **Splunk**, **QRadar**, **SumoLogic**, **Syslog servers**, **LogRhythm**, **Logz.io Cloud Observability Platform**, and other monitoring solutions. connect Defender for Cloud with Azure monitor via Azure Event Hubs:

> **NOTE**
> To stream alerts at the tenant level, use this Azure policy and set the scope at the root management group (you'll need permissions for the root management group as explained in Defender for Cloud permissions): Deploy export to event hub for Microsoft Defender for Cloud alerts and recommendations.

1. Enable continuous export to stream Defender for Cloud alerts into a dedicated event hub at the subscription level. To do this at the Management Group level using Azure Policy, see Create continuous export automation configurations at scale
2. Connect the event hub to your preferred solution using Azure Monitor's built-in connectors.
3. Optionally, stream the raw logs to the event hub and connect to your preferred solution. Learn more in Monitoring data available.

To view the event schemas of the exported data types, visit the Event hub event schemas.

Other streaming options

As an alternative to Sentinel and Azure Monitor, you can use Defender for Cloud's built-in integration with Microsoft Graph Security API. No configuration is required and there are no additional costs.

You can use this API to stream alerts from your **entire tenant** (and data from many other Microsoft Security products) into third-party SIEMs and other popular platforms:

- **Splunk Enterprise and Splunk Cloud** - Use the Microsoft Graph Security API Add-On for Splunk
- **Power BI** - Connect to the Microsoft Graph Security API in Power BI Desktop
- **ServiceNow** - Follow the instructions to install and configure the Microsoft Graph Security API application from the ServiceNow Store
- **QRadar** - IBM's Device Support Module for Microsoft Defender for Cloud via Microsoft Graph API
- **Palo Alto Networks**, **Anomali**, **Lookout**, **InSpark**, and more - Microsoft Graph Security API

Next steps

This page explained how to ensure your Microsoft Defender for Cloud alert data is available in your SIEM, SOAR, or ITSM tool of choice. For related material, see:

- What is Microsoft Sentinel?
- Alert validation in Microsoft Defender for Cloud - Verify your alerts are correctly configured
- Continuously export Defender for Cloud data

Continuously export Microsoft Defender for Cloud data

2/15/2022 • 10 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud generates detailed security alerts and recommendations. You can view them in the portal or through programmatic tools. You might also need to export some or all of this information for tracking with other monitoring tools in your environment.

You fully customize *what* will be exported, and *where* it will go with **continuous export**. For example, you can configure it so that:

- All high severity alerts are sent to an Azure Event Hub
- All medium or higher severity findings from vulnerability assessment scans of your SQL servers are sent to a specific Log Analytics workspace
- Specific recommendations are delivered to an Event Hub or Log Analytics workspace whenever they're generated
- The secure score for a subscription is sent to a Log Analytics workspace whenever the score for a control changes by 0.01 or more

Even though the feature is called *continuous*, there's also an option to export weekly snapshots.

This article describes how to configure continuous export to Log Analytics workspaces or Azure Event Hubs.

> **NOTE**
> If you need to integrate Defender for Cloud with a SIEM, see Stream alerts to a SIEM, SOAR, or IT Service Management solution.

> **TIP**
> Defender for Cloud also offers the option to perform a one-time, manual export to CSV. Learn more in Manual one-time export of alerts and recommendations.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free

ASPECT	DETAILS

Required roles and permissions:	• **Security admin** or **Owner** on the resource group
	• Write permissions for the target resource.
	• If you're using the Azure Policy 'DeployIfNotExist' policies described below you'll also need permissions for assigning policies
	• To export data to Event Hub, you'll need Write permission on the Event Hub Policy.
	• To export to a Log Analytics workspace:
	  ◦ if it **has the SecurityCenterFree solution**, you'll need a minimum of read permissions for the workspace solution: `Microsoft.OperationsManagement/solutions/read`
	  ◦ if it **doesn't have the SecurityCenterFree solution**, you'll need write permissions for the workspace solution: `Microsoft.OperationsManagement/solutions/action`
	  ◦ Learn more about Azure Monitor and Log Analytics workspace solutions

Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

What data types can be exported?

Continuous export can export the following data types whenever they change:

- Security alerts.
- Security recommendations.
- Security findings. These can be thought of as 'sub' recommendations and belong to a 'parent' recommendation. For example:
  - The recommendations System updates should be installed on your machines (powered by Update Center) and System updates should be installed on your machines each has one 'sub' recommendation per outstanding system update.
  - The recommendation Machines should have vulnerability findings resolved has a 'sub' recommendation for every vulnerability identified by the vulnerability scanner.

  > **NOTE**
  > If you're configuring a continuous export with the REST API, always include the parent with the findings.

- Secure score per subscription or per control.
- Regulatory compliance data.

Set up a continuous export

You can configure continuous export from the Microsoft Defender for Cloud pages in Azure portal, via the REST API, or at scale using the supplied Azure Policy templates. Select the appropriate tab below for details of each.

- **Use the Azure portal**
- **Use the REST API**
- **Deploy at scale with Azure Policy**

Configure continuous export from the Defender for Cloud pages in Azure portal

The steps below are necessary whether you're setting up a continuous export to Log Analytics workspace or Azure Event Hubs.

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the specific subscription for which you want to configure the data export.
3. From the sidebar of the settings page for that subscription, select **Continuous Export**.

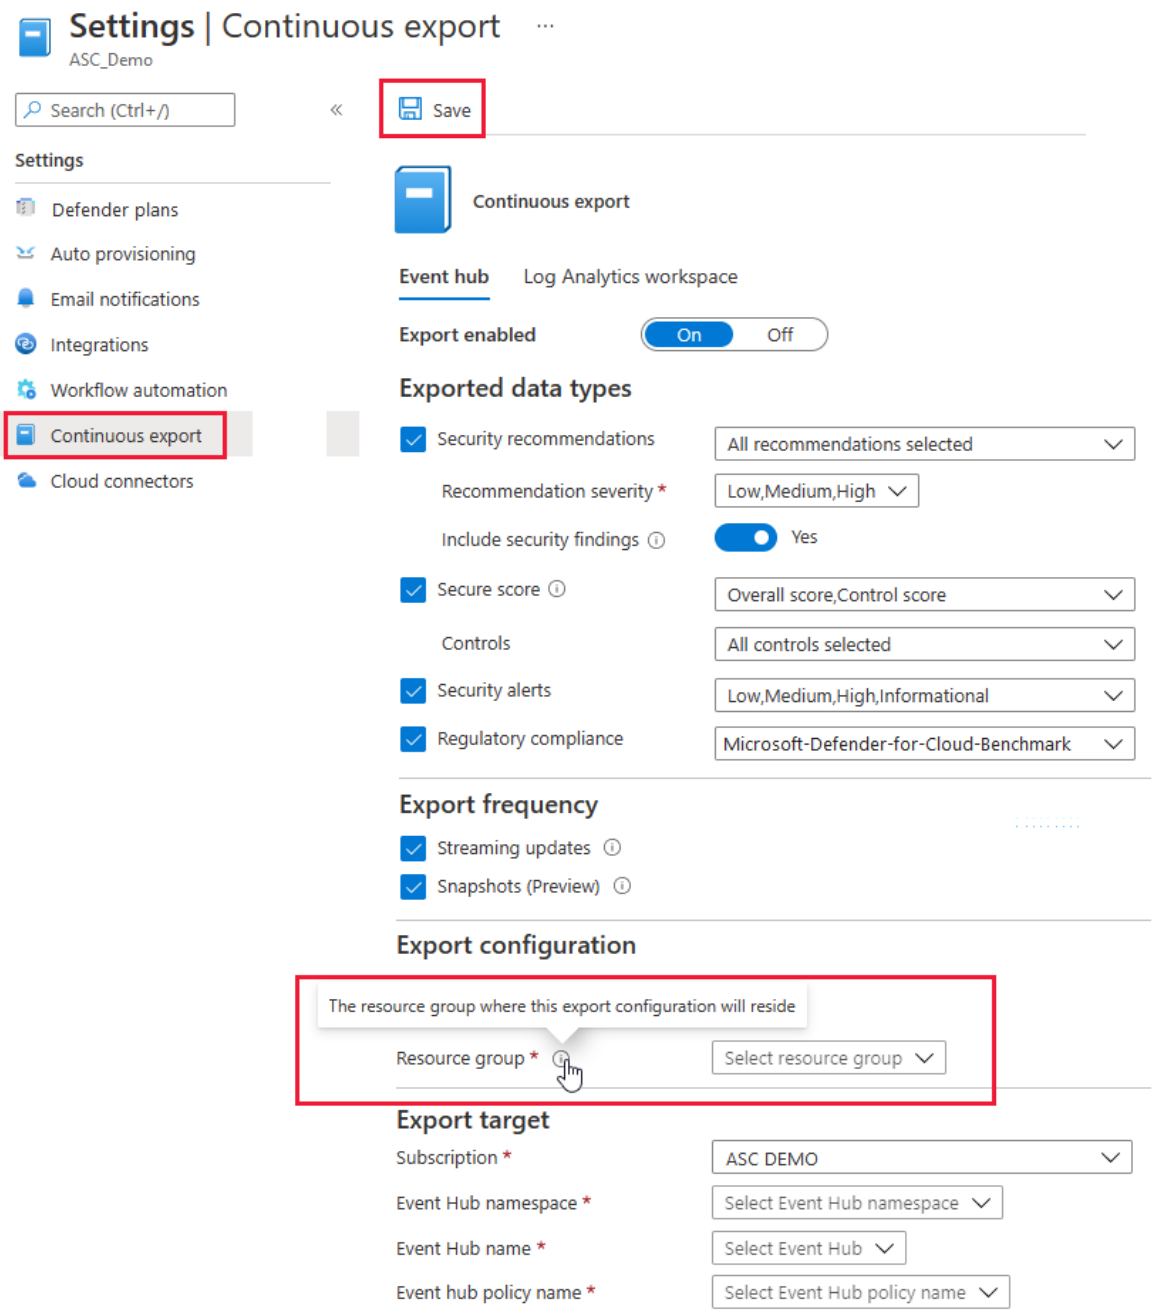

Here you see the export options. There's a tab for each available export target.

4. Select the data type you'd like to export and choose from the filters on each type (for example, export only high severity alerts).
5. Select the appropriate export frequency:
   - **Streaming** – assessments will be sent when a resource's health state is updated (if no updates occur, no data will be sent).
   - **Snapshots** – a snapshot of the current state of the selected data types will be sent once a week per subscription. To identify snapshot data, look for the field `IsSnapshot`.

6. Optionally, if your selection includes one of these recommendations, you can include the vulnerability assessment findings together with them:

   - SQL databases should have vulnerability findings resolved
   - SQL servers on machines should have vulnerability findings resolved
   - Container registry images should have vulnerability findings resolved (powered by Qualys)
   - Machines should have vulnerability findings resolved
   - System updates should be installed on your machines

   To include the findings with these recommendations, enable the **include security findings** option.

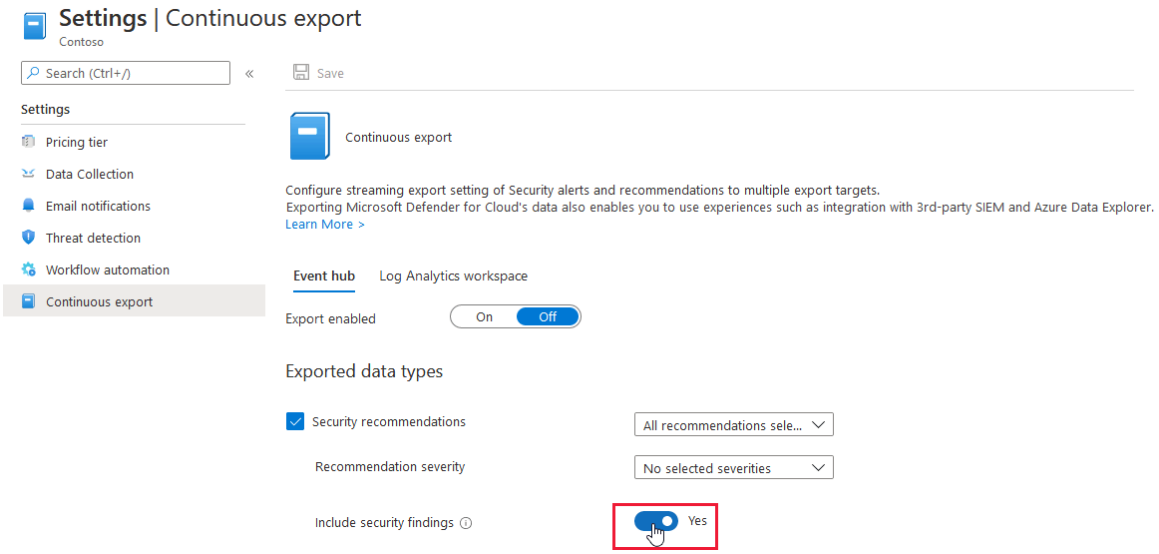

7. From the "Export target" area, choose where you'd like the data saved. Data can be saved in a target on a different subscription (for example on a Central Event Hub instance or a central Log Analytics workspace).

8. Select **Save**.

Information about exporting to a Log Analytics workspace

If you want to analyze Microsoft Defender for Cloud data inside a Log Analytics workspace or use Azure alerts together with Defender for Cloud alerts, set up continuous export to your Log Analytics workspace.

Log Analytics tables and schemas

Security alerts and recommendations are stored in the *SecurityAlert* and *SecurityRecommendation* tables respectively.

The name of the Log Analytics solution containing these tables depends on whether you have enabled the enhanced security features: Security ('Security and Audit') or SecurityCenterFree.

> **TIP**
> To see the data on the destination workspace, you must enable one of these solutions **Security and Audit** or **SecurityCenterFree**.

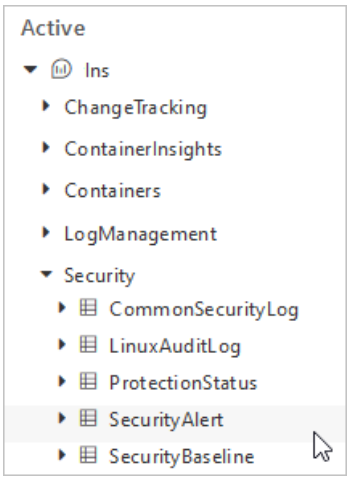

To view the event schemas of the exported data types, visit the Log Analytics table schemas.

View exported alerts and recommendations in Azure Monitor

You might also choose to view exported Security Alerts and/or recommendations in Azure Monitor.

Azure Monitor provides a unified alerting experience for a variety of Azure alerts including Diagnostic Log, Metric alerts, and custom alerts based on Log Analytics workspace queries.

To view alerts and recommendations from Defender for Cloud in Azure Monitor, configure an Alert rule based on Log Analytics queries (Log Alert):

1. From Azure Monitor's **Alerts** page, select **New alert rule**.

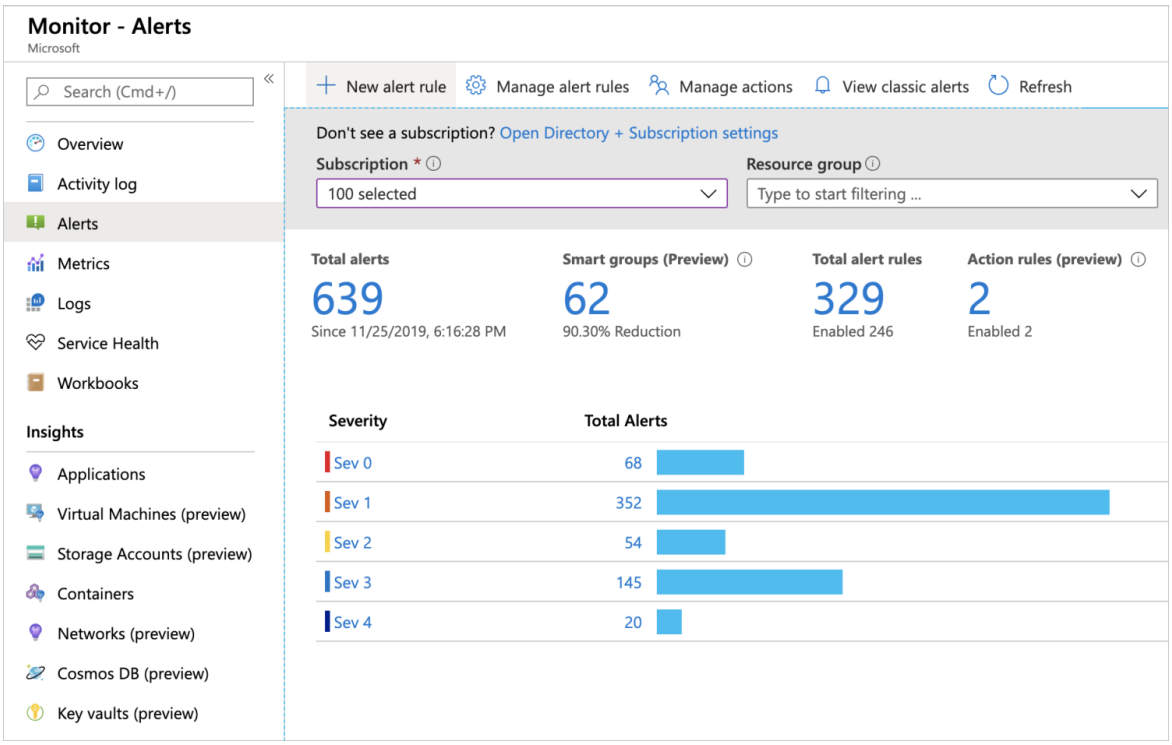

2. In the create rule page, configure your new rule (in the same way you'd configure a log alert rule in Azure Monitor):

   - For **Resource**, select the Log Analytics workspace to which you exported security alerts and recommendations.
   - For **Condition**, select **Custom log search**. In the page that appears, configure the query, lookback period, and frequency period. In the search query, you can type *SecurityAlert* or *SecurityRecommendation* to query the data types that Defender for Cloud continuously exports to

as you enable the Continuous export to Log Analytics feature.

- Optionally, configure the Action Group that you'd like to trigger. Action groups can trigger email sending, ITSM tickets, WebHooks, and more.

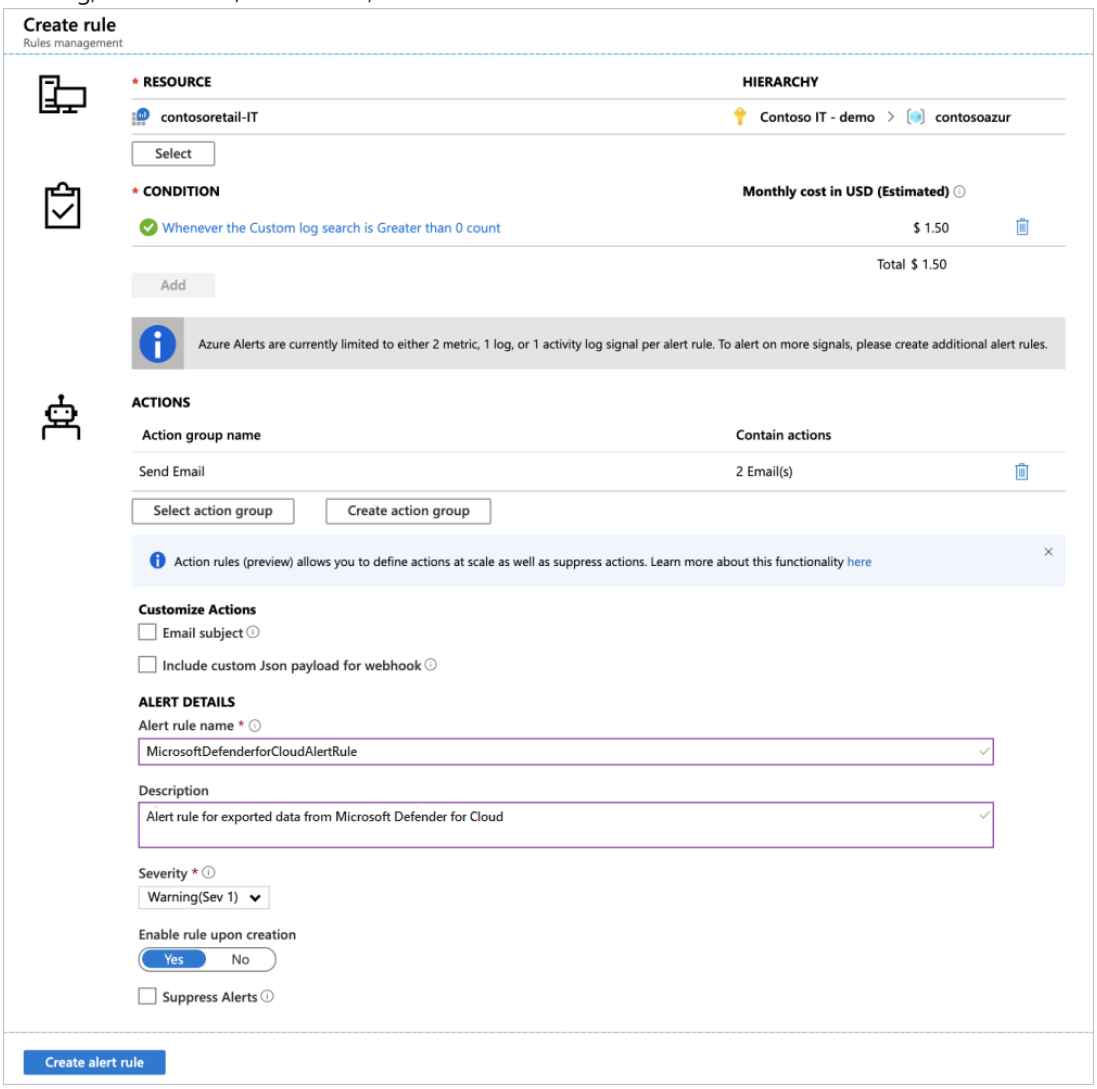

You'll now see new Microsoft Defender for Cloud alerts or recommendations (depending on your configured continuous export rules and the condition you defined in your Azure Monitor alert rule) in Azure Monitor alerts, with automatic triggering of an action group (if provided).

Manual one-time export of alerts and recommendations

To download a CSV report for alerts or recommendations, open the **Security alerts** or **Recommendations** page and select the **Download CSV report** button.

> **TIP**
> Due to Azure Resource Graph limitations, the reports are limited to a file size of 13K rows. If you're seeing errors related to too much data being exported, try limiting the output by selecting a smaller set of subscriptions to be exported.

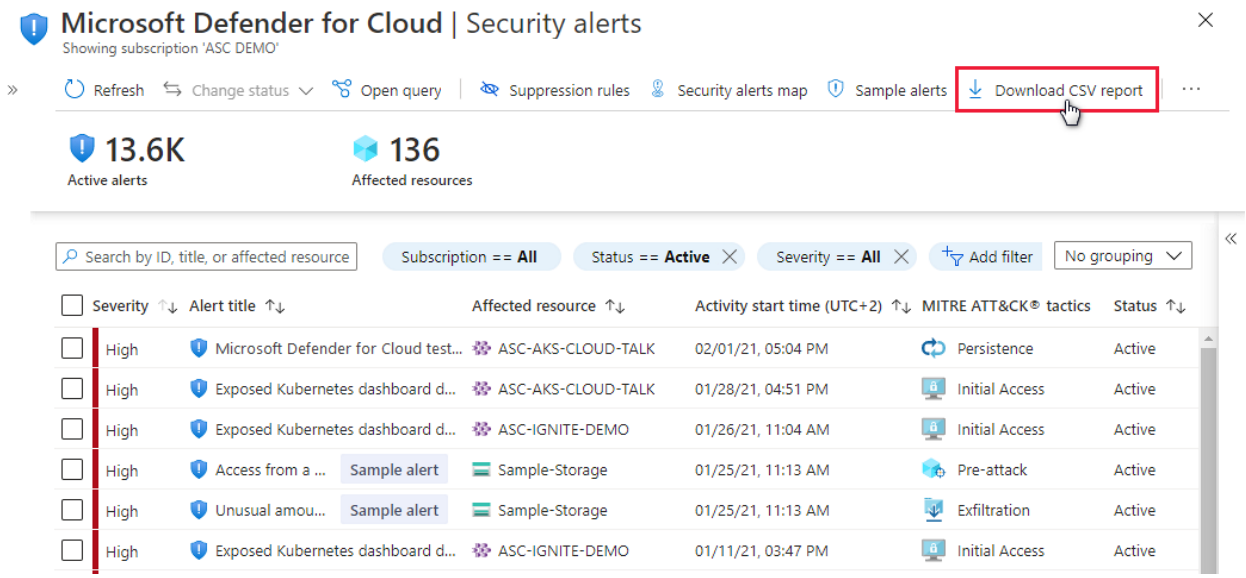

NOTE

These reports contain alerts and recommendations for resources from the currently selected subscriptions.

FAQ - Continuous export

What are the costs involved in exporting data?

There is no cost for enabling a continuous export. Costs might be incurred for ingestion and retention of data in your Log Analytics workspace, depending on your configuration there.

Learn more about Log Analytics workspace pricing.

Learn more about Azure Event Hub pricing.

Does the export include data about the current state of all resources?

No. Continuous export is built for streaming of **events**:

- **Alerts** received before you enabled export won't be exported.
- **Recommendations** are sent whenever a resource's compliance state changes. For example, when a resource turns from healthy to unhealthy. Therefore, as with alerts, recommendations for resources that haven't changed state since you enabled export won't be exported.
- **Secure score** per security control or subscription is sent when a security control's score changes by 0.01 or more.
- **Regulatory compliance status** is sent when the status of the resource's compliance changes.

Why are recommendations sent at different intervals?

Different recommendations have different compliance evaluation intervals, which can vary from a few minutes to every few days. Consequently, recommendations will differ in the amount of time it takes for them to appear in your exports.

Does continuous export support any business continuity or disaster recovery (BCDR) scenarios?

When preparing your environment for BCDR scenarios, where the target resource is experiencing an outage or other disaster, it's the organization's responsibility to prevent data loss by establishing backups according to the guidelines from Azure Event Hubs, Log Analytics workspace, and Logic App.

Learn more in Azure Event Hubs - Geo-disaster recovery.

Is continuous export available for free?

Yes! Note that many alerts are only provided when you've enabled advanced protections. A good way to preview the alerts you'll get in your exported data is to see the alerts shown in Defender for Cloud's pages in the Azure portal.

Next steps

In this article, you learned how to configure continuous exports of your recommendations and alerts. You also learned how to download your alerts data as a CSV file.

For related material, see the following documentation:

- Learn more about workflow automation templates.
- Azure Event Hubs documentation
- Microsoft Sentinel documentation
- Azure Monitor documentation
- Export data types schemas

Continuously export Microsoft Defender for Cloud data

2/15/2022 • 10 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud generates detailed security alerts and recommendations. You can view them in the portal or through programmatic tools. You might also need to export some or all of this information for tracking with other monitoring tools in your environment.

You fully customize *what* will be exported, and *where* it will go with **continuous export**. For example, you can configure it so that:

- All high severity alerts are sent to an Azure Event Hub
- All medium or higher severity findings from vulnerability assessment scans of your SQL servers are sent to a specific Log Analytics workspace
- Specific recommendations are delivered to an Event Hub or Log Analytics workspace whenever they're generated
- The secure score for a subscription is sent to a Log Analytics workspace whenever the score for a control changes by 0.01 or more

Even though the feature is called *continuous*, there's also an option to export weekly snapshots.

This article describes how to configure continuous export to Log Analytics workspaces or Azure Event Hubs.

> **NOTE**
> If you need to integrate Defender for Cloud with a SIEM, see Stream alerts to a SIEM, SOAR, or IT Service Management solution.

> **TIP**
> Defender for Cloud also offers the option to perform a one-time, manual export to CSV. Learn more in Manual one-time export of alerts and recommendations.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free

ASPECT	DETAILS

Required roles and permissions:	• **Security admin** or **Owner** on the resource group
	• Write permissions for the target resource.
	• If you're using the Azure Policy 'DeployIfNotExist' policies described below you'll also need permissions for assigning policies
	• To export data to Event Hub, you'll need Write permission on the Event Hub Policy.
	• To export to a Log Analytics workspace:
	  ◦ if it **has the SecurityCenterFree solution**, you'll need a minimum of read permissions for the workspace solution: `Microsoft.OperationsManagement/solutions/read`
	  ◦ if it **doesn't have the SecurityCenterFree solution**, you'll need write permissions for the workspace solution: `Microsoft.OperationsManagement/solutions/action`
	  ◦ Learn more about Azure Monitor and Log Analytics workspace solutions

Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

What data types can be exported?

Continuous export can export the following data types whenever they change:

- Security alerts.
- Security recommendations.
- Security findings. These can be thought of as 'sub' recommendations and belong to a 'parent' recommendation. For example:
  - The recommendations System updates should be installed on your machines (powered by Update Center) and System updates should be installed on your machines each has one 'sub' recommendation per outstanding system update.
  - The recommendation Machines should have vulnerability findings resolved has a 'sub' recommendation for every vulnerability identified by the vulnerability scanner.

  > **NOTE**
  > If you're configuring a continuous export with the REST API, always include the parent with the findings.

- Secure score per subscription or per control.
- Regulatory compliance data.

Set up a continuous export

You can configure continuous export from the Microsoft Defender for Cloud pages in Azure portal, via the REST API, or at scale using the supplied Azure Policy templates. Select the appropriate tab below for details of each.

- **Use the Azure portal**
- **Use the REST API**
- **Deploy at scale with Azure Policy**

Configure continuous export from the Defender for Cloud pages in Azure portal

The steps below are necessary whether you're setting up a continuous export to Log Analytics workspace or Azure Event Hubs.

1. From Defender for Cloud's menu, open **Environment settings**.
2. Select the specific subscription for which you want to configure the data export.
3. From the sidebar of the settings page for that subscription, select **Continuous Export**.

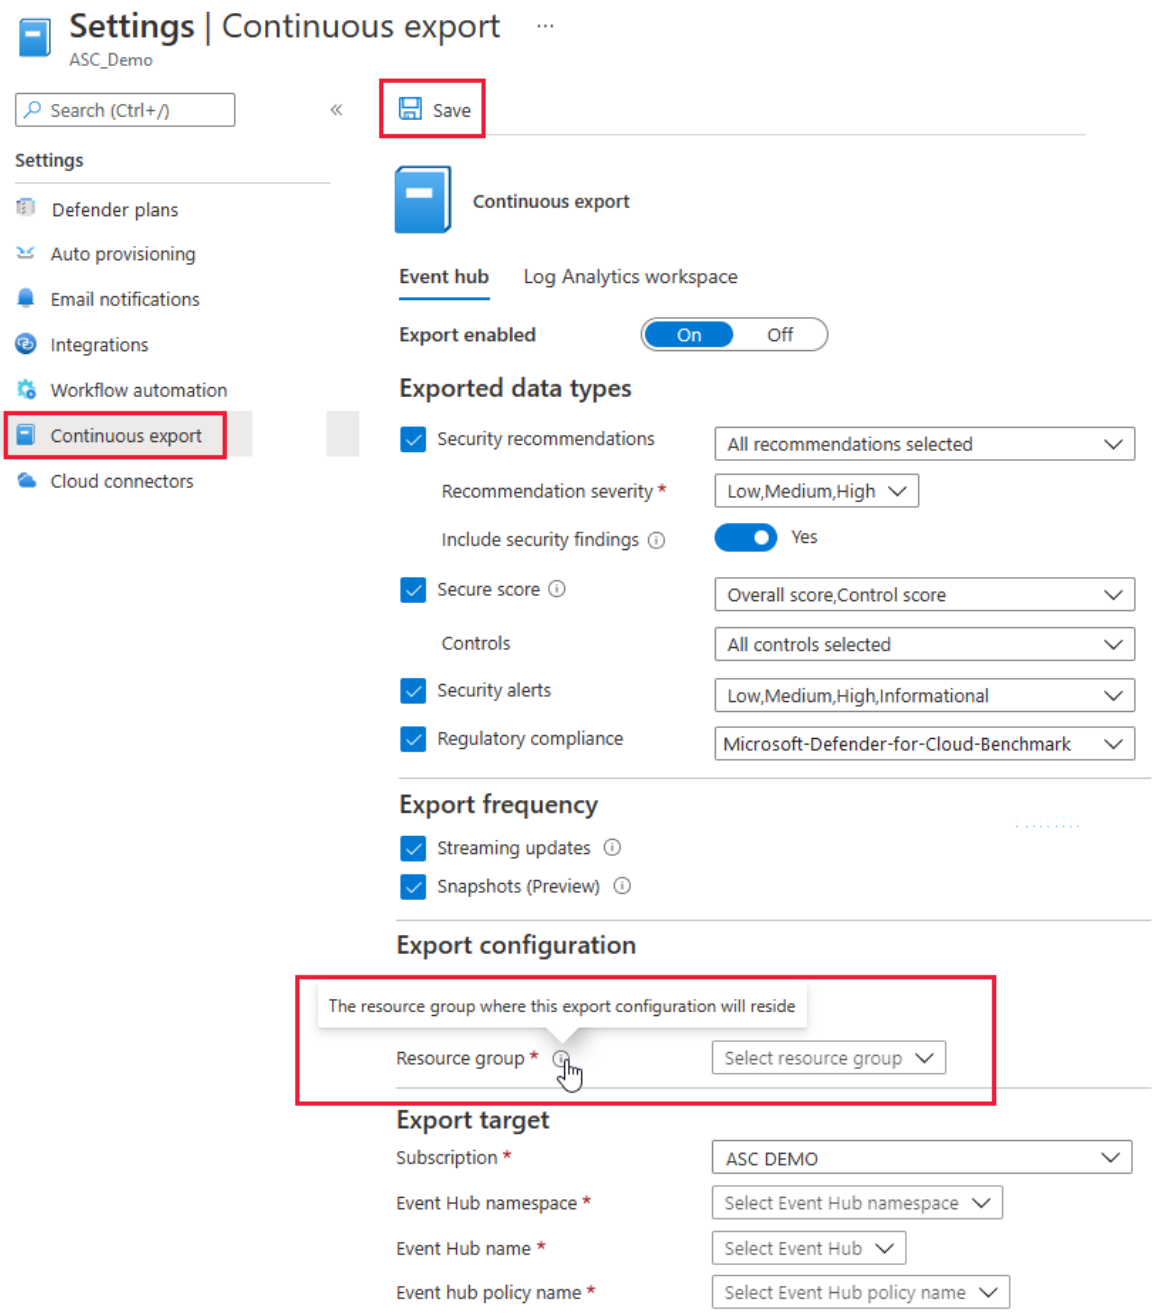

Here you see the export options. There's a tab for each available export target.

4. Select the data type you'd like to export and choose from the filters on each type (for example, export only high severity alerts).
5. Select the appropriate export frequency:
   - **Streaming** – assessments will be sent when a resource's health state is updated (if no updates occur, no data will be sent).
   - **Snapshots** – a snapshot of the current state of the selected data types will be sent once a week per subscription. To identify snapshot data, look for the field `IsSnapshot`.

6. Optionally, if your selection includes one of these recommendations, you can include the vulnerability assessment findings together with them:

   - SQL databases should have vulnerability findings resolved
   - SQL servers on machines should have vulnerability findings resolved
   - Container registry images should have vulnerability findings resolved (powered by Qualys)
   - Machines should have vulnerability findings resolved
   - System updates should be installed on your machines

   To include the findings with these recommendations, enable the **include security findings** option.

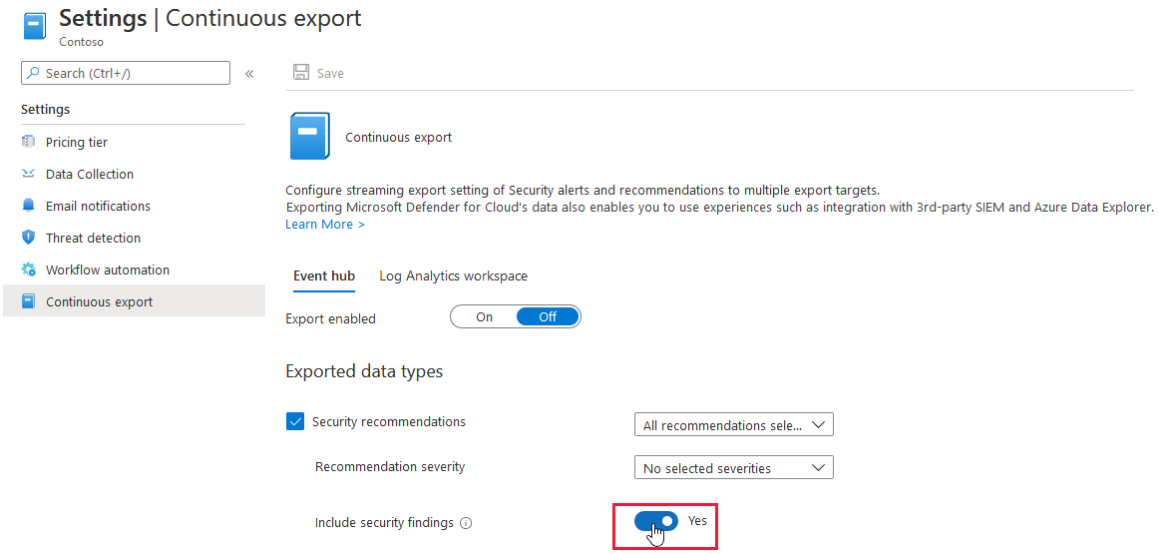

7. From the "Export target" area, choose where you'd like the data saved. Data can be saved in a target on a different subscription (for example on a Central Event Hub instance or a central Log Analytics workspace).

8. Select **Save**.

Information about exporting to a Log Analytics workspace

If you want to analyze Microsoft Defender for Cloud data inside a Log Analytics workspace or use Azure alerts together with Defender for Cloud alerts, set up continuous export to your Log Analytics workspace.

Log Analytics tables and schemas

Security alerts and recommendations are stored in the *SecurityAlert* and *SecurityRecommendation* tables respectively.

The name of the Log Analytics solution containing these tables depends on whether you have enabled the enhanced security features: Security ('Security and Audit') or SecurityCenterFree.

> **TIP**
> To see the data on the destination workspace, you must enable one of these solutions **Security and Audit** or **SecurityCenterFree**.

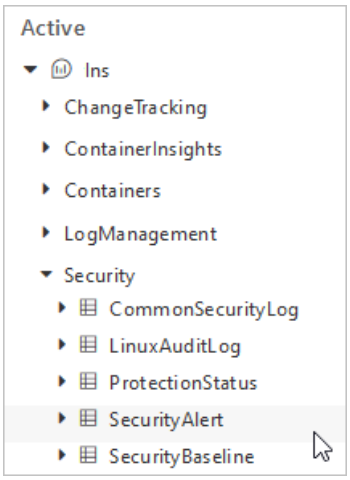

To view the event schemas of the exported data types, visit the Log Analytics table schemas.

View exported alerts and recommendations in Azure Monitor

You might also choose to view exported Security Alerts and/or recommendations in Azure Monitor.

Azure Monitor provides a unified alerting experience for a variety of Azure alerts including Diagnostic Log, Metric alerts, and custom alerts based on Log Analytics workspace queries.

To view alerts and recommendations from Defender for Cloud in Azure Monitor, configure an Alert rule based on Log Analytics queries (Log Alert):

1. From Azure Monitor's **Alerts** page, select **New alert rule**.

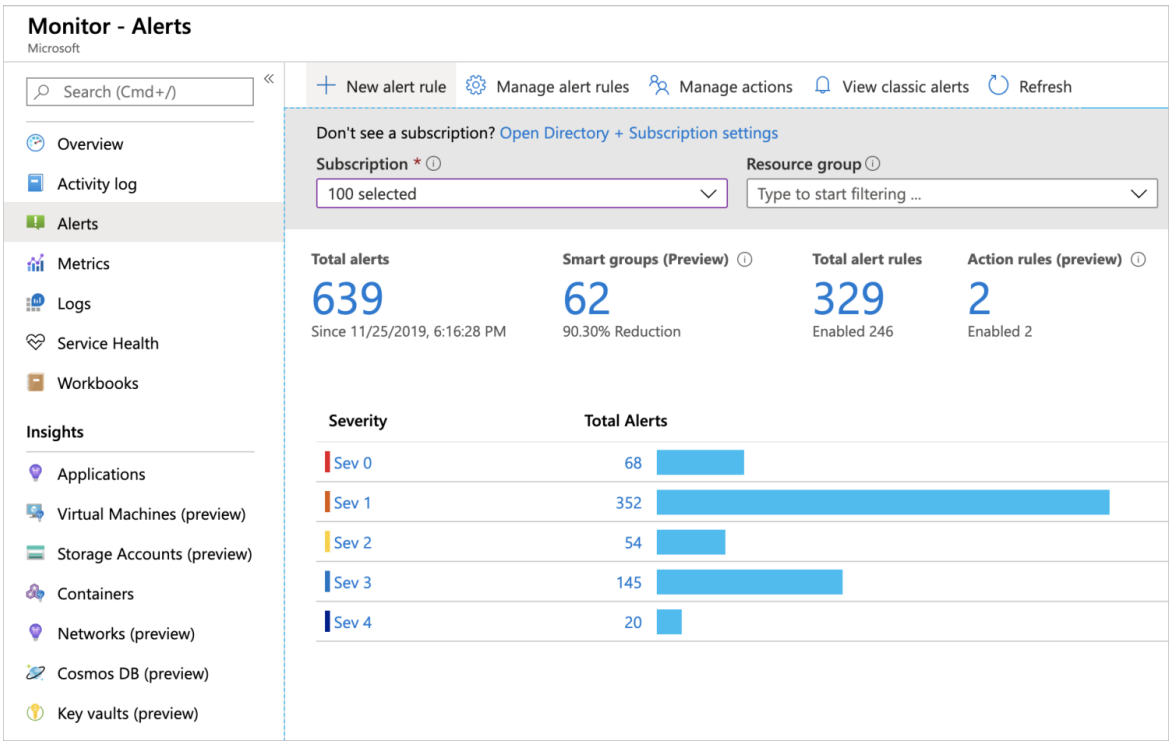

2. In the create rule page, configure your new rule (in the same way you'd configure a log alert rule in Azure Monitor):

   - For **Resource**, select the Log Analytics workspace to which you exported security alerts and recommendations.
   - For **Condition**, select **Custom log search**. In the page that appears, configure the query, lookback period, and frequency period. In the search query, you can type *SecurityAlert* or *SecurityRecommendation* to query the data types that Defender for Cloud continuously exports to

as you enable the Continuous export to Log Analytics feature.

- Optionally, configure the Action Group that you'd like to trigger. Action groups can trigger email sending, ITSM tickets, WebHooks, and more.

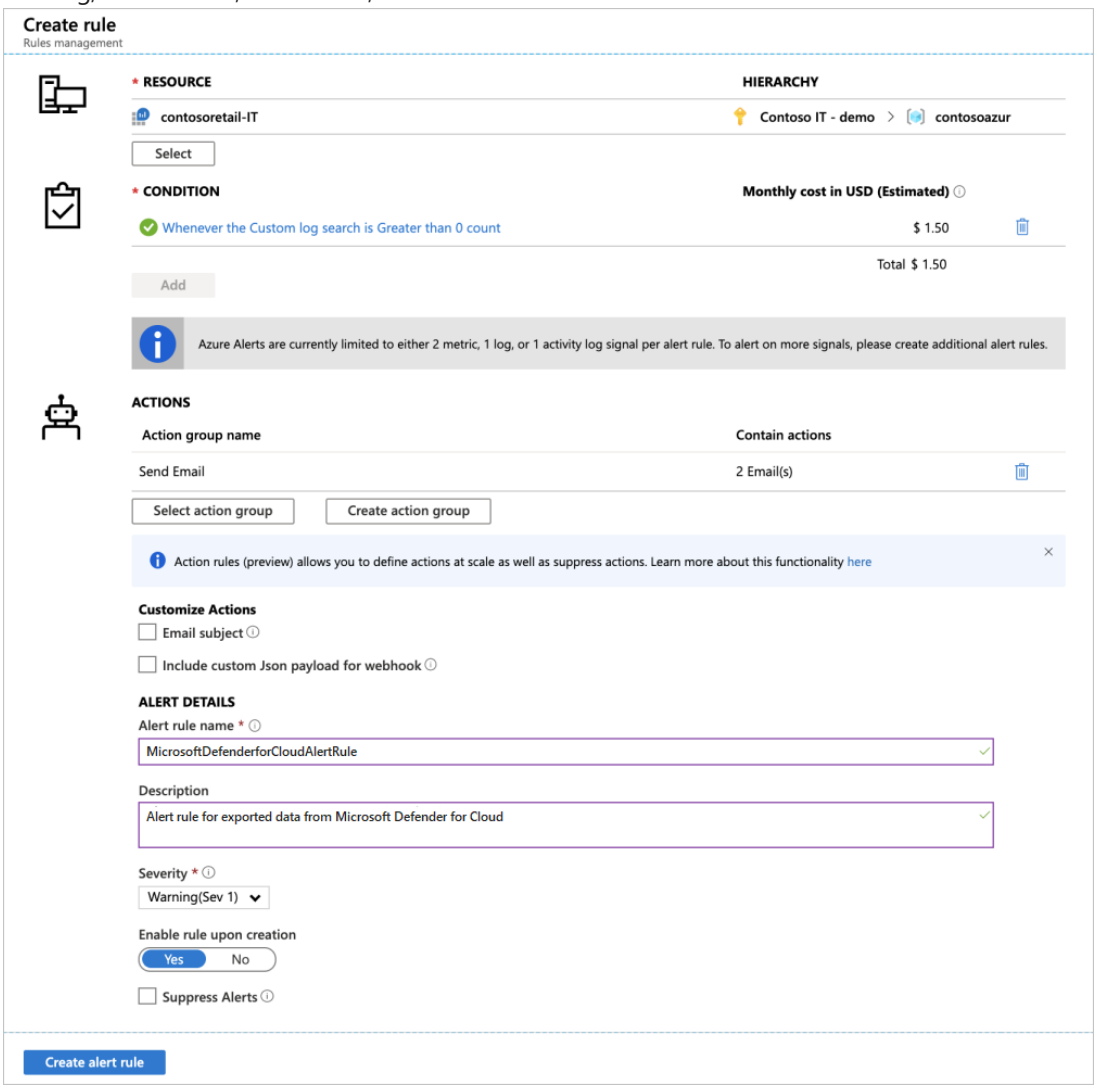

You'll now see new Microsoft Defender for Cloud alerts or recommendations (depending on your configured continuous export rules and the condition you defined in your Azure Monitor alert rule) in Azure Monitor alerts, with automatic triggering of an action group (if provided).

Manual one-time export of alerts and recommendations

To download a CSV report for alerts or recommendations, open the **Security alerts** or **Recommendations** page and select the **Download CSV report** button.

> **TIP**
> Due to Azure Resource Graph limitations, the reports are limited to a file size of 13K rows. If you're seeing errors related to too much data being exported, try limiting the output by selecting a smaller set of subscriptions to be exported.

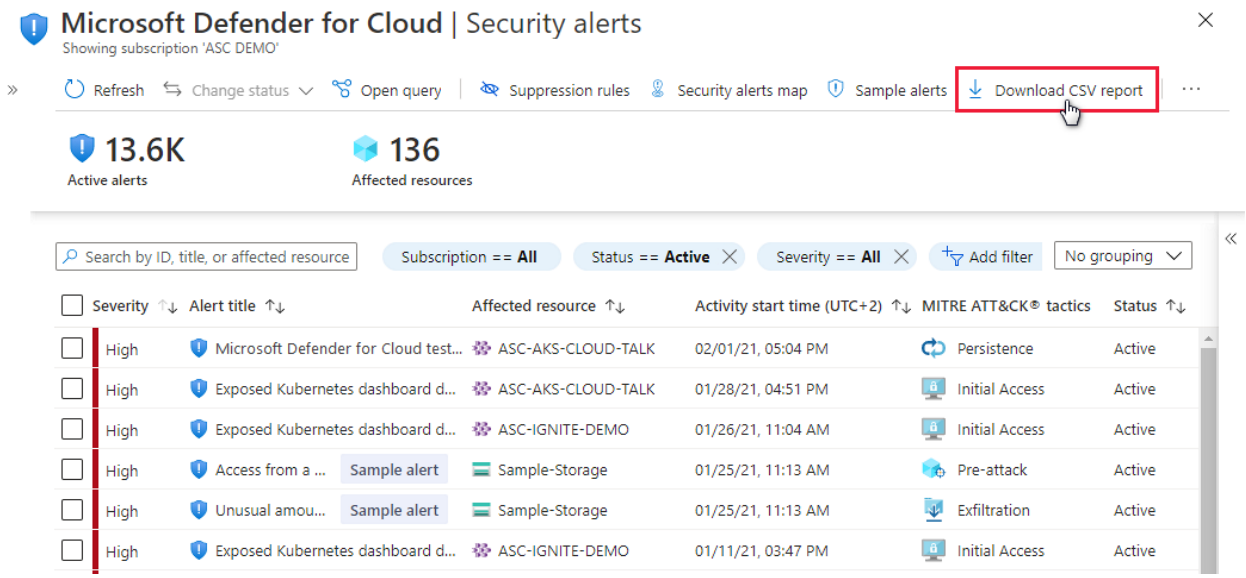

NOTE

These reports contain alerts and recommendations for resources from the currently selected subscriptions.

FAQ - Continuous export

What are the costs involved in exporting data?

There is no cost for enabling a continuous export. Costs might be incurred for ingestion and retention of data in your Log Analytics workspace, depending on your configuration there.

Learn more about Log Analytics workspace pricing.

Learn more about Azure Event Hub pricing.

Does the export include data about the current state of all resources?

No. Continuous export is built for streaming of **events**:

- **Alerts** received before you enabled export won't be exported.
- **Recommendations** are sent whenever a resource's compliance state changes. For example, when a resource turns from healthy to unhealthy. Therefore, as with alerts, recommendations for resources that haven't changed state since you enabled export won't be exported.
- **Secure score** per security control or subscription is sent when a security control's score changes by 0.01 or more.
- **Regulatory compliance status** is sent when the status of the resource's compliance changes.

Why are recommendations sent at different intervals?

Different recommendations have different compliance evaluation intervals, which can vary from a few minutes to every few days. Consequently, recommendations will differ in the amount of time it takes for them to appear in your exports.

Does continuous export support any business continuity or disaster recovery (BCDR) scenarios?

When preparing your environment for BCDR scenarios, where the target resource is experiencing an outage or other disaster, it's the organization's responsibility to prevent data loss by establishing backups according to the guidelines from Azure Event Hubs, Log Analytics workspace, and Logic App.

Learn more in Azure Event Hubs - Geo-disaster recovery.

Is continuous export available for free?

Yes! Note that many alerts are only provided when you've enabled advanced protections. A good way to preview the alerts you'll get in your exported data is to see the alerts shown in Defender for Cloud's pages in the Azure portal.

Next steps

In this article, you learned how to configure continuous exports of your recommendations and alerts. You also learned how to download your alerts data as a CSV file.

For related material, see the following documentation:

- Learn more about workflow automation templates.
- Azure Event Hubs documentation
- Microsoft Sentinel documentation
- Azure Monitor documentation
- Export data types schemas

Security alerts schemas

2/15/2022 • 6 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

If your subscription has enhanced security features enabled, you'll receive security alerts when Defender for Cloud detects threats to their resources.

You can view these security alerts in Microsoft Defender for Cloud's pages - overview dashboard, alerts, resource health pages, or workload protections dashboard - and through external tools such as:

- Microsoft Sentinel - Microsoft's cloud-native SIEM. The Sentinel Connector gets alerts from Microsoft Defender for Cloud and sends them to the Log Analytics workspace for Microsoft Sentinel.
- Third-party SIEMs - Send data to Azure Event Hubs. Then integrate your Event Hub data with a third-party SIEM. Learn more in Stream alerts to a SIEM, SOAR, or IT Service Management solution.
- The REST API - If you're using the REST API to access alerts, see the online Alerts API documentation.

If you're using any programmatic methods to consume the alerts, you'll need the correct schema to find the fields that are relevant to you. Also, if you're exporting to an Event Hub or trying to trigger Workflow Automation with generic HTTP connectors, use the schemas to properly parse the JSON objects.

> **IMPORTANT**
> The schema is slightly different for each of these scenarios, so make sure you select the relevant tab below.

The schemas

- Microsoft Sentinel
- Azure Activity Log
- Workflow automation
- Continuous export
- MS Graph API

The Sentinel Connector gets alerts from Microsoft Defender for Cloud and sends them to the Log Analytics Workspace for Microsoft Sentinel.

To create a Microsoft Sentinel case or incident using Defender for Cloud alerts, you'll need the schema for those alerts shown below.

Learn more in the Microsoft Sentinel documentation.

The data model of the schema

FIELD	DESCRIPTION
AlertName	Alert display name
AlertType	unique alert identifier
ConfidenceLevel	(Optional) The confidence level of this alert (High/Low)
ConfidenceScore	(Optional) Numeric confidence indicator of the security alert
Description	Description text for the alert
DisplayName	The alert's display name
EndTime	The impact end time of the alert (the time of the last event contributing to the alert)
Entities	A list of entities related to the alert. This list can hold a mixture of entities of diverse types
ExtendedLinks	(Optional) A bag for all links related to the alert. This bag can hold a mixture of links for diverse types
ExtendedProperties	A bag of additional fields which are relevant to the alert
IsIncident	Determines if the alert is an incident or a regular alert. An incident is a security alert that aggregates multiple alerts into one security incident
ProcessingEndTime	UTC timestamp in which the alert was created
ProductComponentName	(Optional) The name of a component inside the product which generated the alert.
ProductName	constant ('Azure Security Center')
ProviderName	unused
RemediationSteps	Manual action items to take to remediate the security threat
ResourceId	Full identifier of the affected resource
Severity	The alert severity (High/Medium/Low/Informational)
SourceComputerId	a unique GUID for the affected server (if the alert is generated on the server)
SourceSystem	unused
StartTime	The impact start time of the alert (the time of the first event contributing to the alert)
SystemAlertId	Unique identifier of this security alert instance

FIELD	DESCRIPTION
TenantId	the identifier of the parent Azure Active directory tenant of the subscription under which the scanned resource resides
TimeGenerated	UTC timestamp on which the assessment took place (Security Center's scan time) (identical to DiscoveredTimeUTC)
Type	constant ('SecurityAlert')
VendorName	The name of the vendor that provided the alert (e.g. 'Microsoft')
VendorOriginalId	unused
WorkspaceResourceGroup	in case the alert is generated on a VM, Server, Virtual Machine Scale Set or App Service instance that reports to a workspace, contains that workspace resource group name
WorkspaceSubscriptionId	in case the alert is generated on a VM, Server, Virtual Machine Scale Set or App Service instance that reports to a workspace, contains that workspace subscriptionId

Next steps

This article described the schemas that Microsoft Defender for Cloud's threat protection tools use when sending security alert information.

For more information on the ways to access security alerts from outside Defender for Cloud, see:

- Microsoft Sentinel - Microsoft's cloud-native SIEM
- Azure Event Hubs - Microsoft's fully managed, real-time data ingestion service
- Continuously export Defender for Cloud data
- Log Analytics workspaces - Azure Monitor stores log data in a Log Analytics workspace, a container that includes data and configuration information

Manage security incidents in Microsoft Defender for Cloud

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Triaging and investigating security alerts can be time consuming for even the most skilled security analysts. For many, it's hard to know where to begin.

Defender for Cloud uses analytics to connect the information between distinct security alerts. Using these connections, Defender for Cloud can provide a single view of an attack campaign and its related alerts to help you understand the attacker's actions and the affected resources.

This page provides an overview of incidents in Defender for Cloud.

What is a security incident?

In Defender for Cloud, a security incident is an aggregation of all alerts for a resource that align with kill chain patterns. Incidents appear in the Security alerts page. Select an incident to view the related alerts and get more information.

Managing security incidents

1. On Defender for Cloud's alerts page, use the **Add filter** button to filter by alert name to the alert name **Security incident detected on multiple resources**.

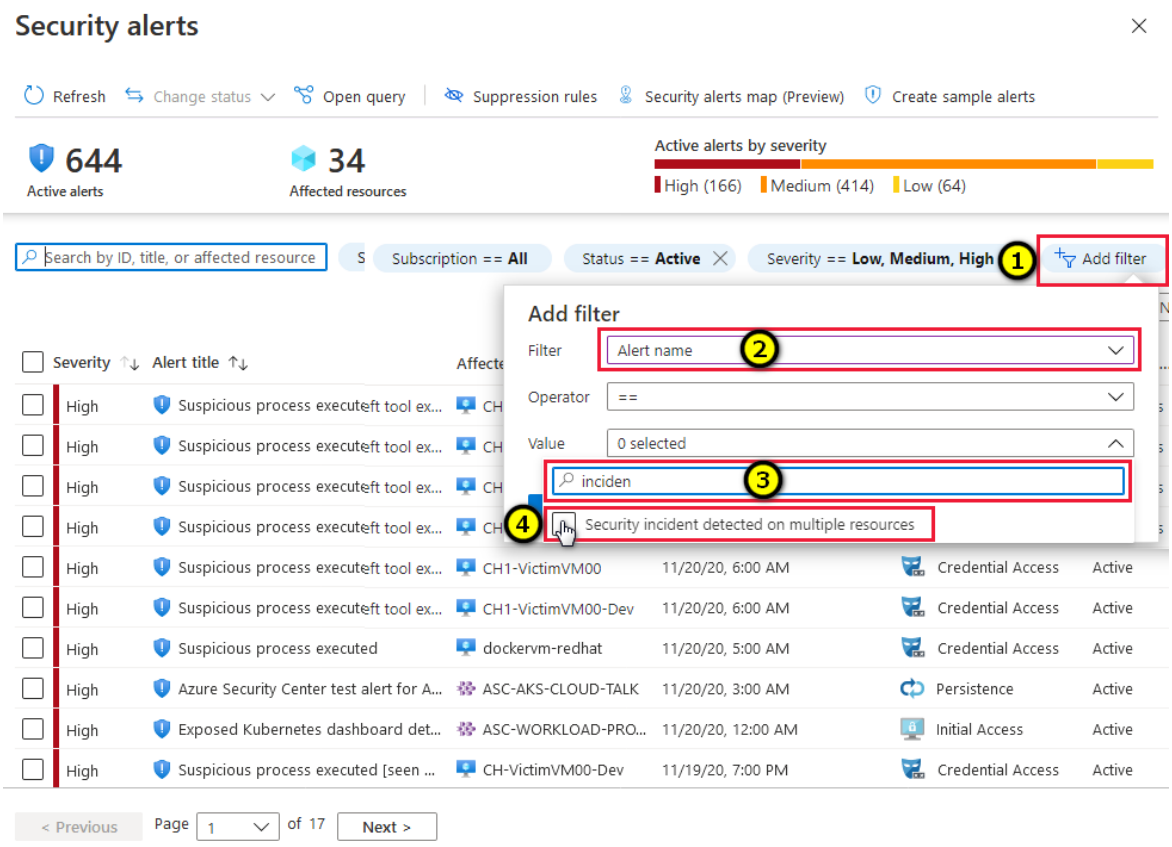

The list is now filtered to show only incidents. Notice that security incidents have a different icon to security alerts.

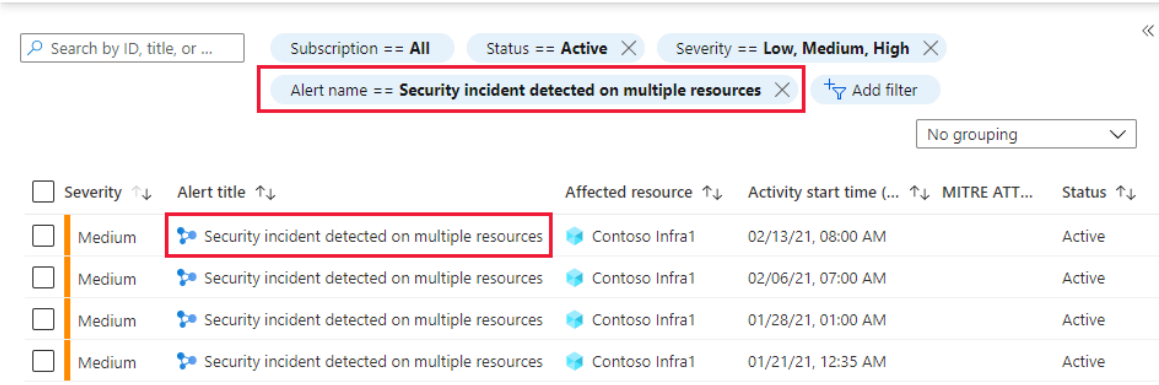

2. To view details of an incident, select one from the list. A side pane appears with more details about the incident.

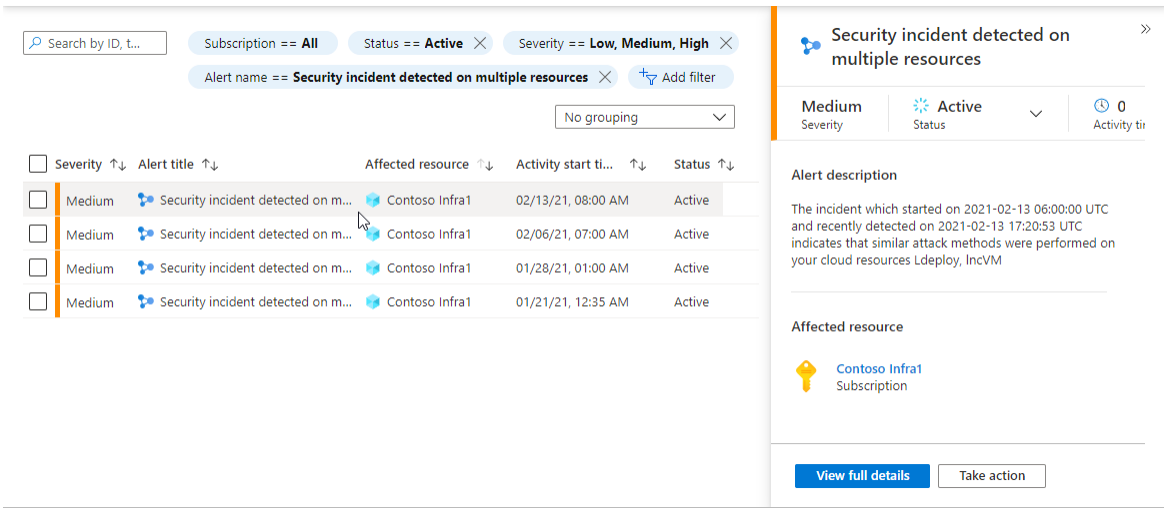

3. To view more details, select **View full details**.

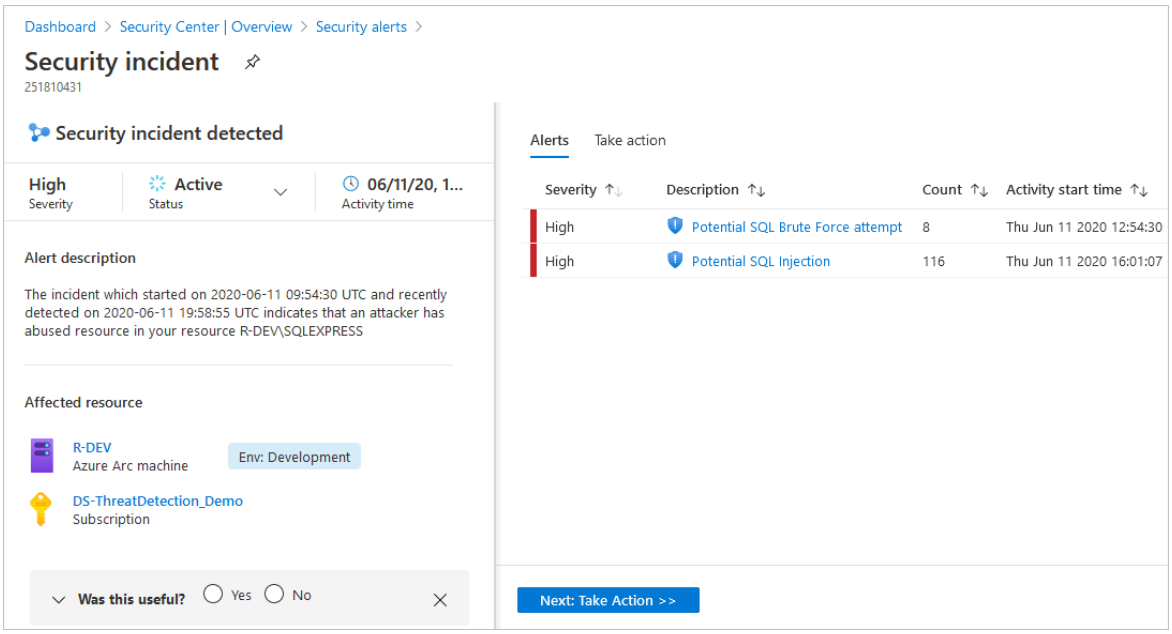

The left pane of the security incident page shows high-level information about the security incident: title, severity, status, activity time, description, and the affected resource. Next to the affected resource you can see the relevant Azure tags. Use these tags to infer the organizational context of the resource when investigating the alert.

The right pane includes the **Alerts** tab with the security alerts that were correlated as part of this incident.

> **TIP**
>
> For more information about a specific alert, select it.

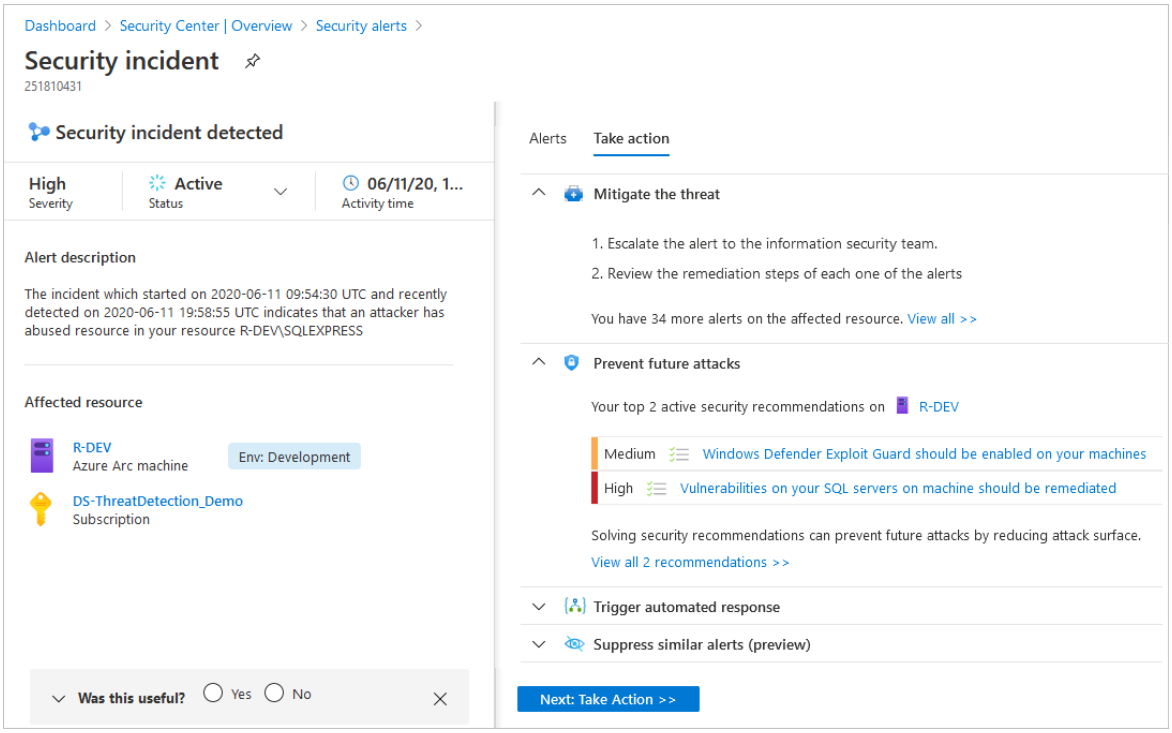

To switch to the **Take action** tab, select the tab or the button on the bottom of the right pane. Use this tab to take further actions such as:

- *Mitigate the threat* - provides manual remediation steps for this security incident
- *Prevent future attacks* - provides security recommendations to help reduce the attack surface, increase

security posture, and prevent future attacks

- *Trigger automated response* - provides the option to trigger a Logic App as a response to this security incident
- *Suppress similar alerts* - provides the option to suppress future alerts with similar characteristics if the alert isn't relevant for your organization

> **NOTE**
> The same alert can exist as part of an incident, as well as to be visible as a standalone alert.

4. To remediate the threats in the incident, follow the remediation steps provided with each alert.

Next steps

This page explained the security incident capabilities of Defender for Cloud. For related information, see the following pages:

- Security alerts in Defender for Cloud
- Manage and respond to security alerts

Microsoft Defender for Cloud threat intelligence report

2/15/2022 • 2 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page explains how Microsoft Defender for Cloud's threat intelligence reports can help you learn more about a threat that triggered a security alert.

What is a threat intelligence report?

Defender for Cloud's threat protection works by monitoring security information from your Azure resources, the network, and connected partner solutions. It analyzes this information, often correlating information from multiple sources, to identify threats. For more information, see How Microsoft Defender for Cloud detects and responds to threats.

When Defender for Cloud identifies a threat, it triggers a security alert, which contains detailed information regarding the event, including suggestions for remediation. To help incident response teams investigate and remediate threats, Defender for Cloud provides threat intelligence reports containing information about detected threats. The report includes information such as:

- Attacker's identity or associations (if this information is available)
- Attackers' objectives
- Current and historical attack campaigns (if this information is available)
- Attackers' tactics, tools, and procedures
- Associated indicators of compromise (IoC) such as URLs and file hashes
- Victimology, which is the industry and geographic prevalence to assist you in determining if your Azure resources are at risk
- Mitigation and remediation information

NOTE

The amount of information in any particular report will vary; the level of detail is based on the malware's activity and prevalence.

Defender for Cloud has three types of threat reports, which can vary according to the attack. The reports available are:

- **Activity Group Report**: provides deep dives into attackers, their objectives, and tactics.
- **Campaign Report**: focuses on details of specific attack campaigns.
- **Threat Summary Report**: covers all of the items in the previous two reports.

This type of information is useful during the incident response process, where there's an ongoing investigation to understand the source of the attack, the attacker's motivations, and what to do to mitigate this issue in the

future.

How to access the threat intelligence report?

1. From Defender for Cloud's menu, open the **Security alerts** page.
2. Select an alert.

   The alerts details page opens with more details about the alert. Below is the **Ransomware indicators detected** alert details page.

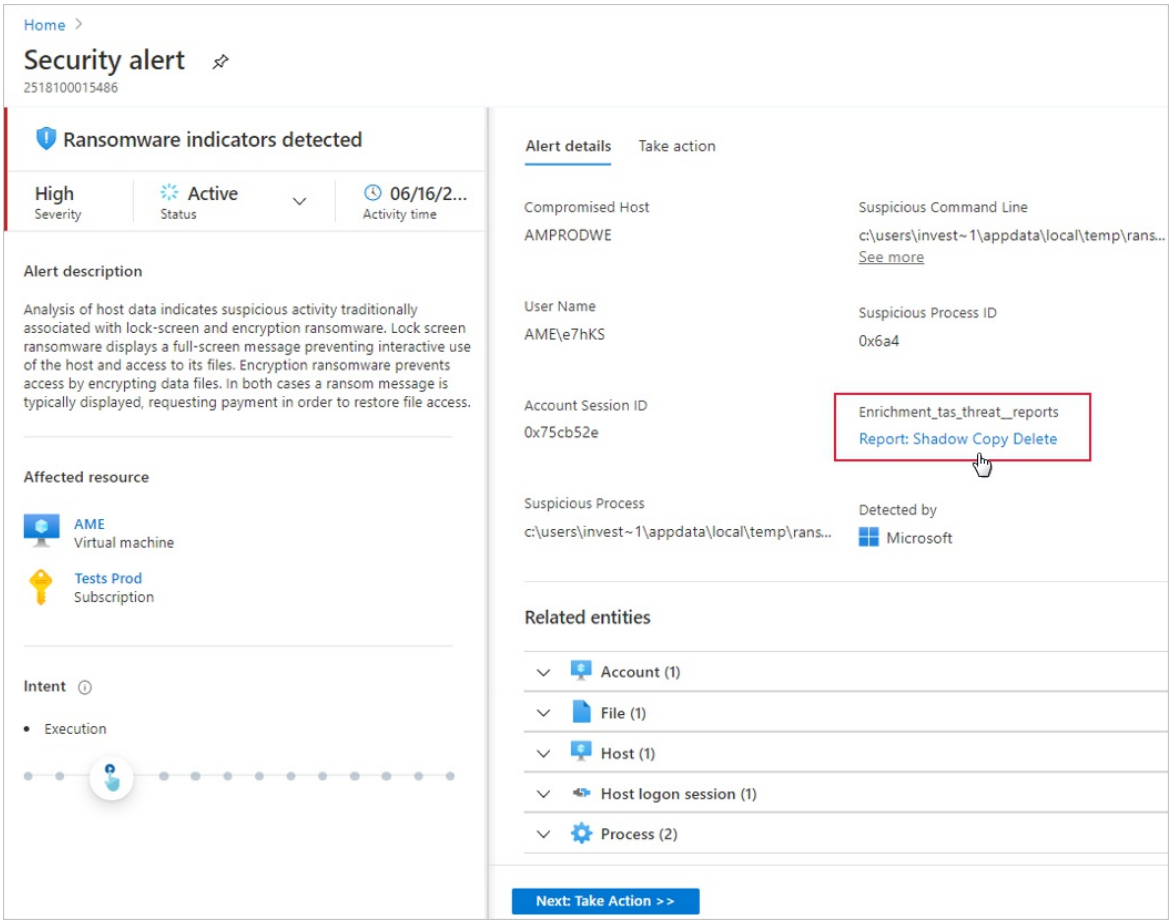

3. Select the link to the report, and a PDF will open in your default browser.

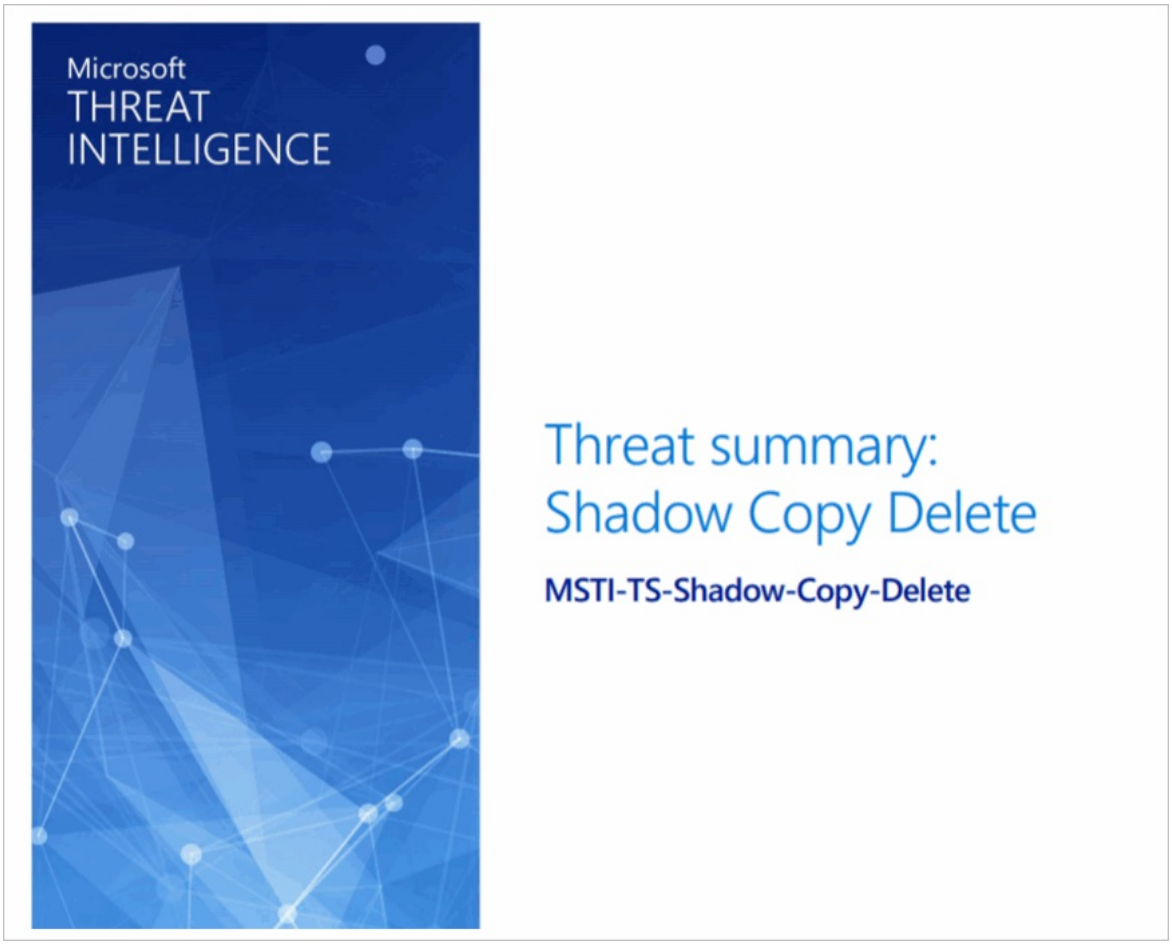

You can optionally download the PDF report.

> **TIP**
> The amount of information available for each security alert will vary according to the type of alert.

Next steps

This page explained how to open threat intelligence reports when investigating security alerts. For related information, see the following pages:

- Managing and responding to security alerts in Microsoft Defender for Cloud. Learn how to manage and respond to security alerts.
- Handling security incidents in Microsoft Defender for Cloud

Alert validation in Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
>
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This document helps you learn how to verify if your system is properly configured for Microsoft Defender for Cloud alerts.

What are security alerts?

Alerts are the notifications that Defender for Cloud generates when it detects threats on your resources. It prioritizes and lists the alerts along with the information needed to quickly investigate the problem. Defender for Cloud also provides recommendations for how you can remediate an attack. For more information, see Security alerts in Defender for Cloud and Managing and responding to security alerts

Generate sample security alerts

If you're using the new, preview alerts experience as described in Manage and respond to security alerts in Microsoft Defender for Cloud, you can create sample alerts in a few clicks from the security alerts page in the Azure portal.

Use sample alerts to:

- evaluate the value and capabilities of your Microsoft Defender plans
- validate any configurations you've made for your security alerts (such as SIEM integrations, workflow automation, and email notifications)

To create sample alerts:

1. As a user with the role **Subscription Contributor**, from the toolbar on the alerts page, select **Create sample alerts**.
2. Select the subscription.
3. Select the relevant Microsoft Defender plan/s for which you want to see alerts.
4. Select **Create sample alerts**.

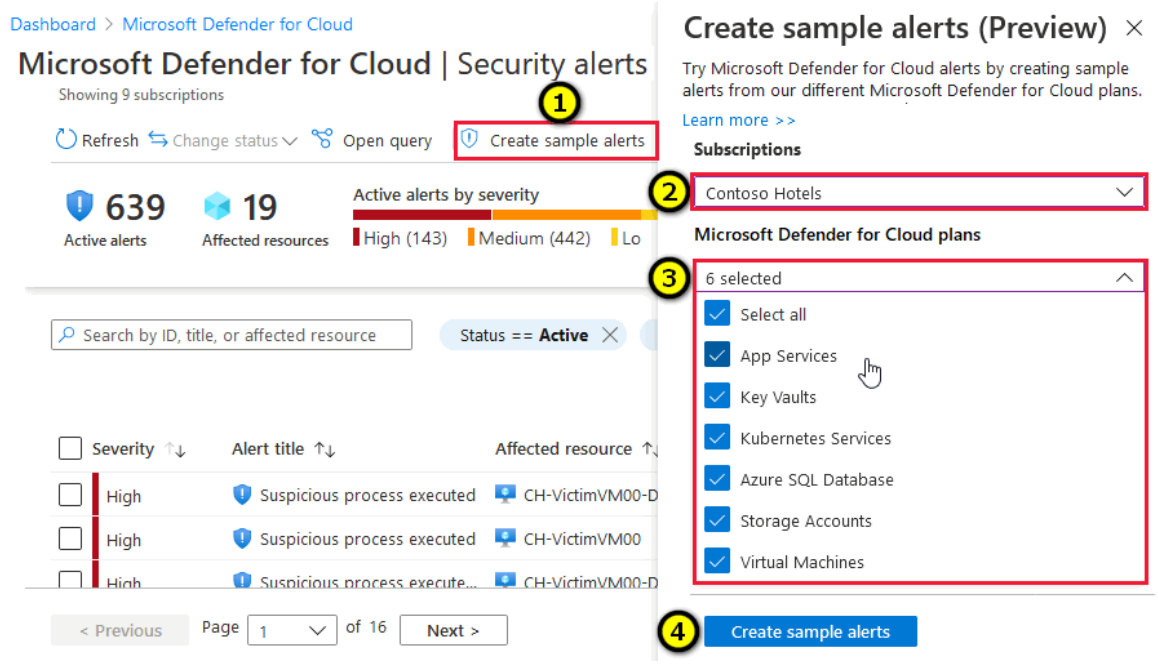

A notification appears letting you know that the sample alerts are being created:

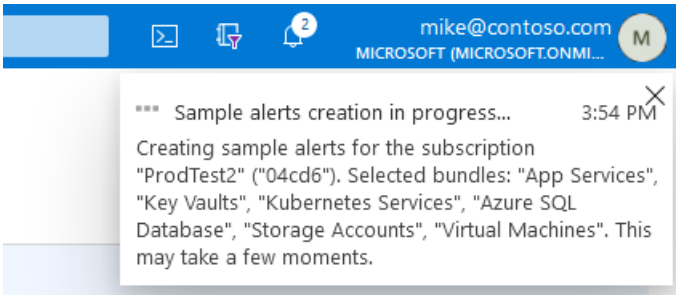

After a few minutes, the alerts appear in the security alerts page. They'll also appear anywhere else that you've configured to receive your Microsoft Defender for Cloud security alerts (connected SIEMs, email notifications, and so on).

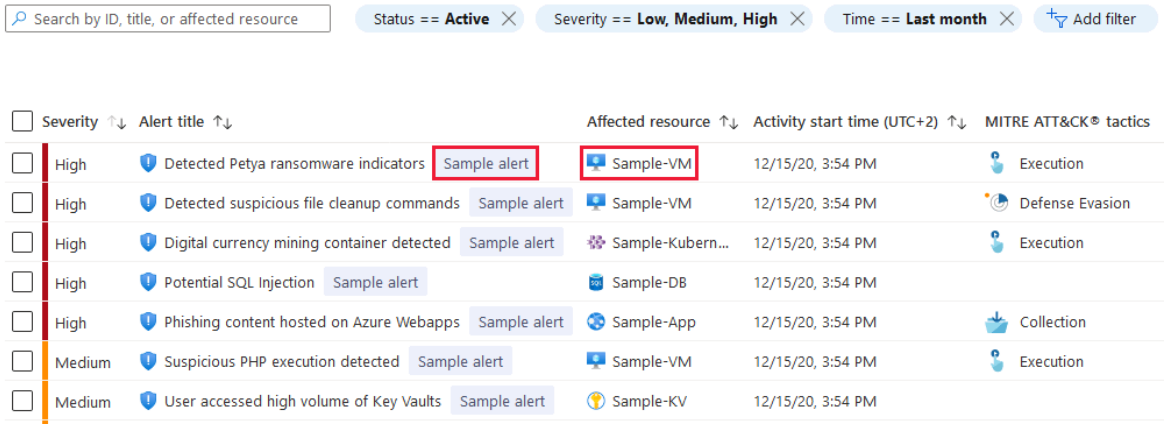

TIP

The alerts are for simulated resources.

Simulate alerts on your Azure VMs (Windows)

After the Log Analytics agent is installed on your machine, follow these steps from the computer where you want to be the attacked resource of the alert:

1. Copy an executable (for example **calc.exe**) to the computer's desktop, or other directory of your

convenience, and rename it as **ASC_AlertTest_662jfi039N.exe**.

2. Open the command prompt and execute this file with an argument (just a fake argument name), such as:
   `ASC_AlertTest_662jfi039N.exe -foo`
3. Wait 5 to 10 minutes and open Defender for Cloud Alerts. An alert should appear.

> **NOTE**
>
> When reviewing this test alert for Windows, make sure the field **Arguments Auditing Enabled** is **true**. If it is **false**, then you need to enable command-line arguments auditing. To enable it, use the following command:
>
> ```
> reg add "HKEY_LOCAL_MACHINE\SOFTWARE\Microsoft\Windows\CurrentVersion\policies\system\Audit" /f /v
> "ProcessCreationIncludeCmdLine_Enabled"
> ```

Simulate alerts on your Azure VMs (Linux)

After the Log Analytics agent is installed on your machine, follow these steps from the computer where you want to be the attacked resource of the alert:

1. Copy an executable to a convenient location and rename it to `./asc_alerttest_662jfi039n`. For example:

   ```
   cp /bin/echo ./asc_alerttest_662jfi039n
   ```

2. Open the command prompt and execute this file:

   ```
   ./asc_alerttest_662jfi039n testing eicar pipe
   ```

3. Wait 5 to 10 minutes and then open Defender for Cloud Alerts. An alert should appear.

Simulate alerts on Kubernetes

If you've integrated Azure Kubernetes Service with Defender for Cloud, you can test that your alerts are working with the following kubectl command:

```
kubectl get pods --namespace=asc-alerttest-662jfi039n
```

For more information about defending your Kubernetes nodes and clusters, see Introduction to Microsoft Defender for Containers

Next steps

This article introduced you to the alerts validation process. Now that you're familiar with this validation, try the following articles:

- Validating Azure Key Vault threat detection in Microsoft Defender for Cloud
- Managing and responding to security alerts in Microsoft Defender for Cloud - Learn how to manage alerts, and respond to security incidents in Defender for Cloud.
- Understanding security alerts in Microsoft Defender for Cloud - Learn about the different types of security alerts.

Automate responses to Microsoft Defender for Cloud triggers

2/15/2022 • 6 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Every security program includes multiple workflows for incident response. These processes might include notifying relevant stakeholders, launching a change management process, and applying specific remediation steps. Security experts recommend that you automate as many steps of those procedures as you can. Automation reduces overhead. It can also improve your security by ensuring the process steps are done quickly, consistently, and according to your predefined requirements.

This article describes the workflow automation feature of Microsoft Defender for Cloud. This feature can trigger Logic Apps on security alerts, recommendations, and changes to regulatory compliance. For example, you might want Defender for Cloud to email a specific user when an alert occurs. You'll also learn how to create Logic Apps using Azure Logic Apps.

Availability

ASPECT	DETAILS
Release state:	General availability (GA)
Pricing:	Free
Required roles and permissions:	**Security admin role** or **Owner** on the resource group Must also have write permissions for the target resource To work with Azure Logic Apps workflows, you must also have the following Logic Apps roles/permissions: - Logic App Operator permissions are required or Logic App read/trigger access (this role can't create or edit logic apps; only *run* existing ones) - Logic App Contributor permissions are required for Logic App creation and modification If you want to use Logic App connectors, you may need additional credentials to sign in to their respective services (for example, your Outlook/Teams/Slack instances)
Clouds:	✅ Commercial clouds ✅ National (Azure Government, Azure China 21Vianet)

Create a logic app and define when it should automatically run

1. From Defender for Cloud's sidebar, select **Workflow automation**.

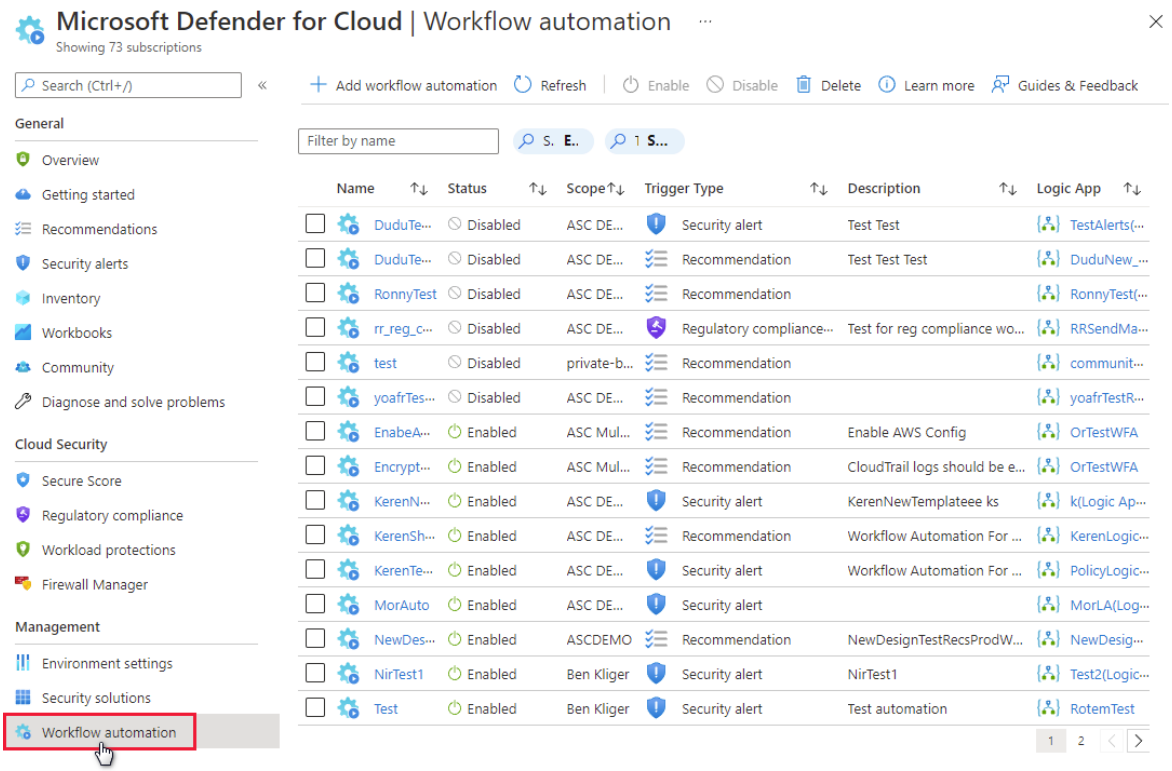

From this page you can create new automation rules, as well as enable, disable, or delete existing ones.

2. To define a new workflow, click **Add workflow automation**. The options pane for your new automation opens.

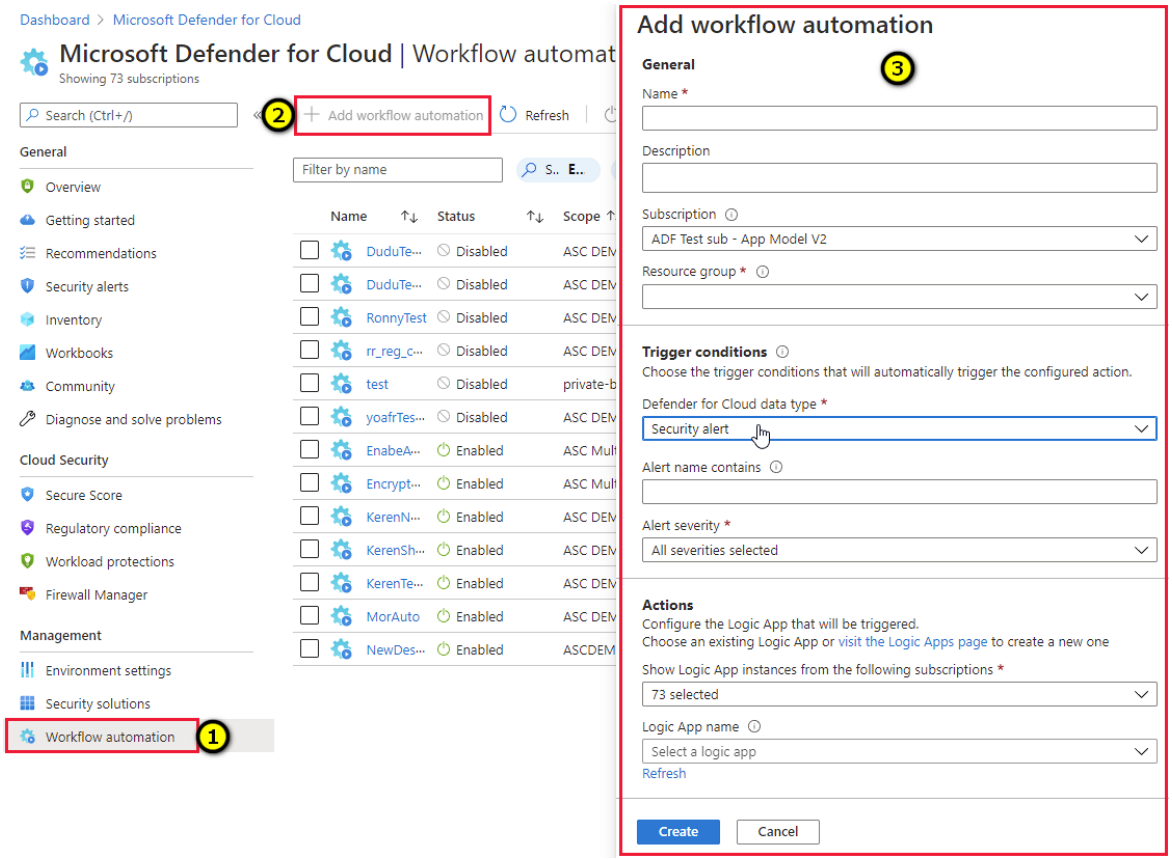

Here you can enter:

a. A name and description for the automation.

b. The triggers that will initiate this automatic workflow. For example, you might want your Logic App to run when a security alert that contains "SQL" is generated.

> **NOTE**
>
> If your trigger is a recommendation that has "sub-recommendations", for example **Vulnerability assessment findings on your SQL databases should be remediated**, the logic app will not trigger for every new security finding; only when the status of the parent recommendation changes.

c. The Logic App that will run when your trigger conditions are met.

3. From the Actions section, select **visit the Logic Apps page** to begin the Logic App creation process.

   You'll be taken to Azure Logic Apps.

4. Select **Add**.

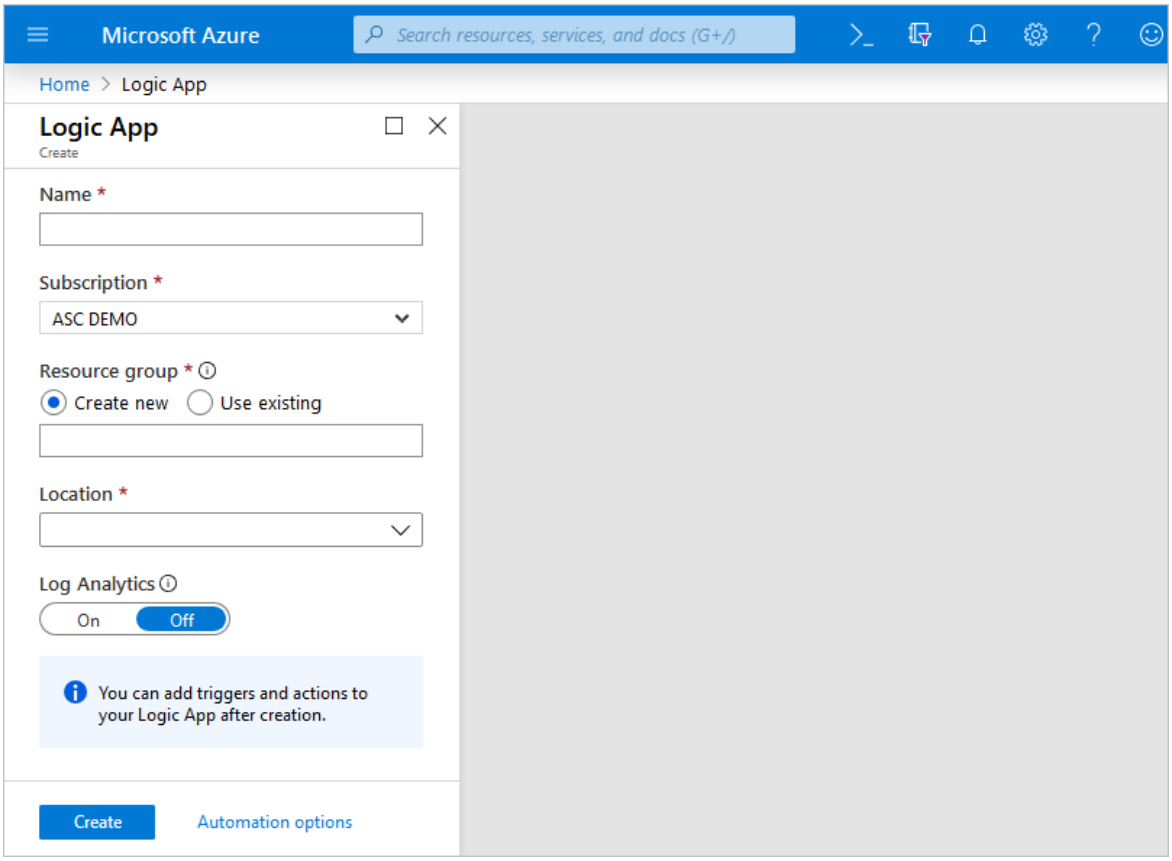

5. Enter a name, resource group, and location, and select **Review and create** > **Create**.

   The message **Deployment is in progress** appears. Wait for the deployment complete notification to appear and select **Go to resource** from the notification.

6. In your new logic app, you can choose from built-in, predefined templates from the security category. Or you can define a custom flow of events to occur when this process is triggered.

> **TIP**
>
> Sometimes in a logic app, parameters are included in the connector as part of a string and not in their own field. For an example of how to extract parameters, see step #14 of Working with logic app parameters while building Microsoft Defender for Cloud workflow automations.

   The logic app designer supports these Defender for Cloud triggers:

- **When a Microsoft Defender for Cloud Recommendation is created or triggered** - If your logic app relies on a recommendation that gets deprecated or replaced, your automation will stop

working and you'll need to update the trigger. To track changes to recommendations, use the release notes.

- **When a Defender for Cloud Alert is created or triggered** - You can customize the trigger so that it relates only to alerts with the severity levels that interest you.
- **When a Defender for Cloud regulatory compliance assessment is created or triggered** - Trigger automations based on updates to regulatory compliance assessments.

> **NOTE**
>
> If you are using the legacy trigger "When a response to a Microsoft Defender for Cloud alert is triggered", your logic apps will not be launched by the Workflow Automation feature. Instead, use either of the triggers mentioned above.

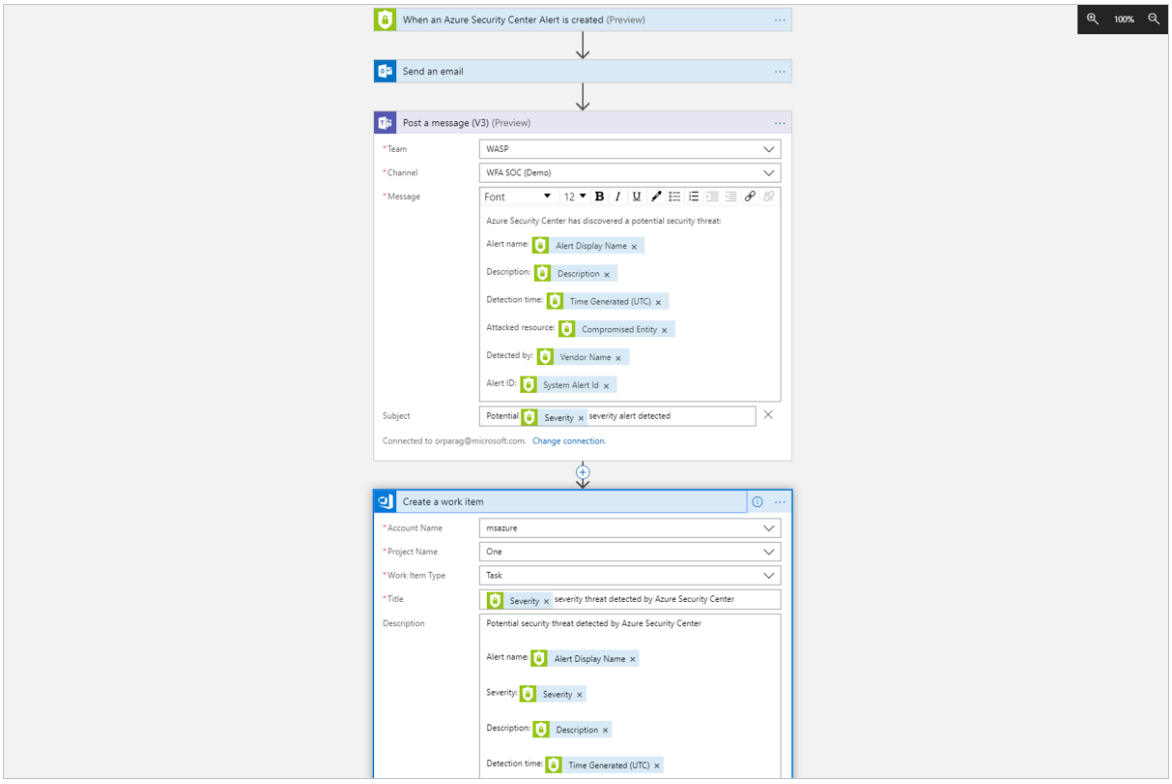

7. After you've defined your logic app, return to the workflow automation definition pane ("Add workflow automation"). Click **Refresh** to ensure your new Logic App is available for selection.

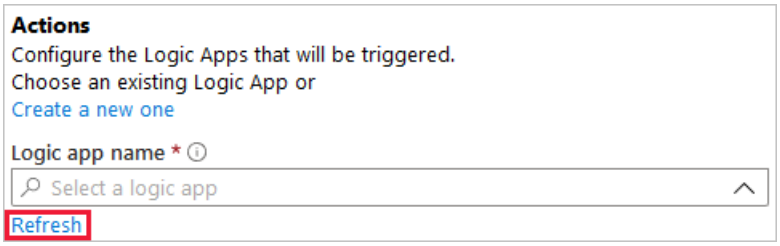

8. Select your logic app and save the automation. Note that the Logic App dropdown only shows Logic Apps with supporting Defender for Cloud connectors mentioned above.

Manually trigger a Logic App

You can also run Logic Apps manually when viewing any security alert or recommendation.

To manually run a Logic App, open an alert or a recommendation and click **Trigger Logic App**:

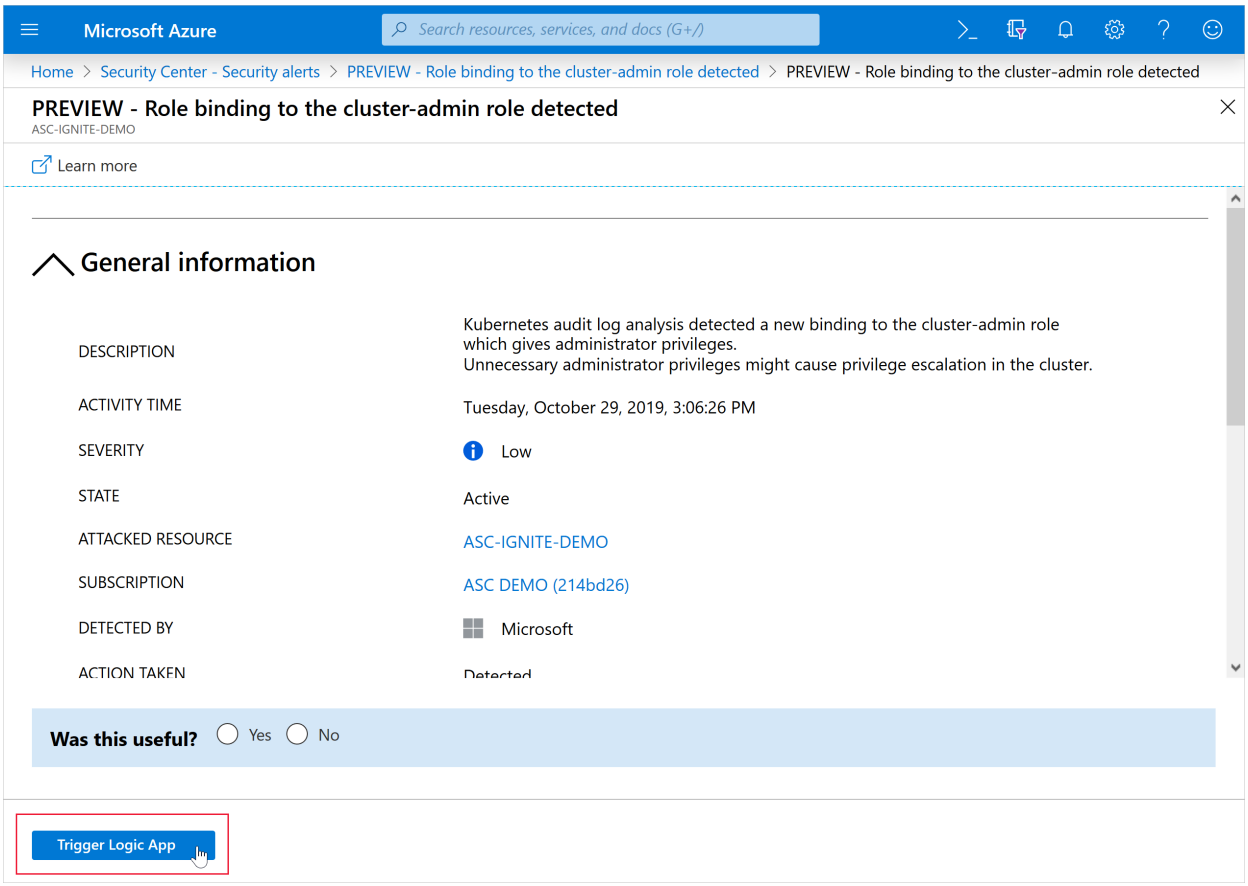

Configure workflow automation at scale using the supplied policies

Automating your organization's monitoring and incident response processes can greatly improve the time it takes to investigate and mitigate security incidents.

To deploy your automation configurations across your organization, use the supplied Azure Policy 'DeployIfNotExist' policies described below to create and configure workflow automation procedures.

Get started with workflow automation templates.

To implement these policies:

1. From the table below, select the policy you want to apply:

GOAL	POLICY	POLICY ID
Workflow automation for security alerts	Deploy Workflow Automation for Microsoft Defender for Cloud alerts	f1525828-9a90-4fcf-be48-268cdd02361e
Workflow automation for security recommendations	Deploy Workflow Automation for Microsoft Defender for Cloud recommendations	73d6ab6c-2475-4850-afd6-43795f3492ef
Workflow automation for regulatory compliance changes	Deploy Workflow Automation for Microsoft Defender for Cloud regulatory compliance	509122b9-ddd9-47ba-a5f1-d0dac20be63c

TIP

You can also find these by searching Azure Policy:

1. Open Azure Policy.

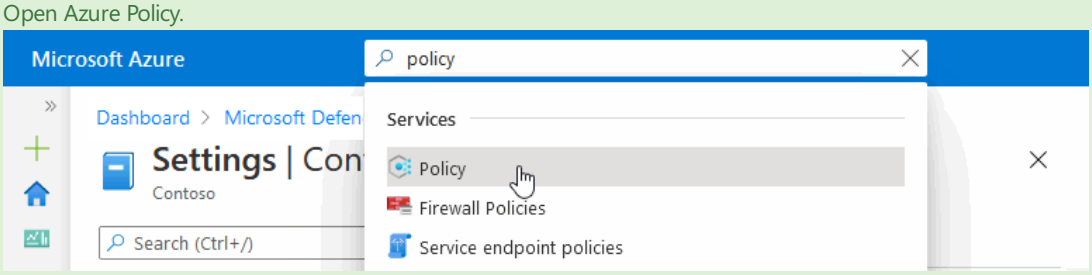

2. From the Azure Policy menu, select **Definitions** and search for them by name.

2. From the relevant Azure Policy page, select **Assign**.

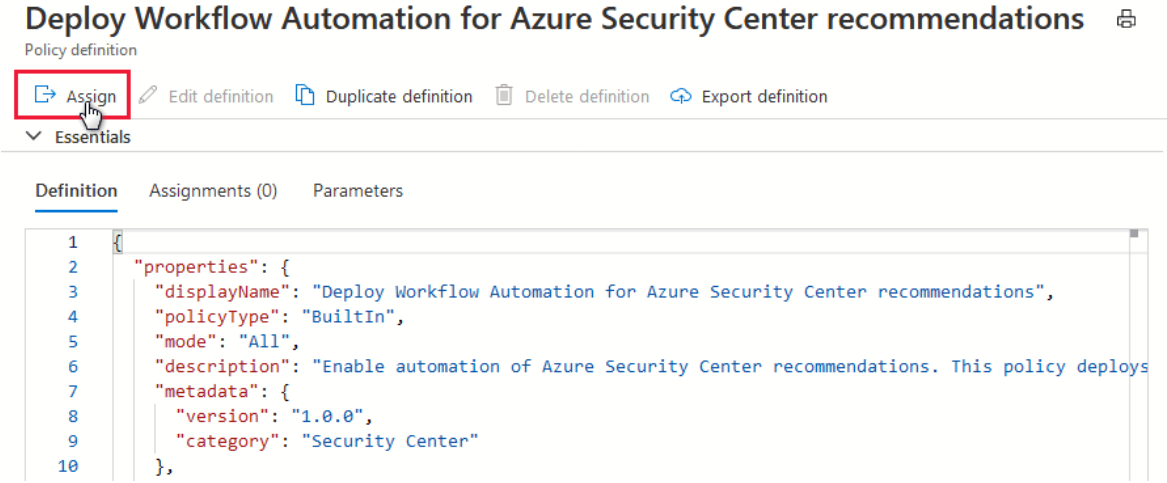

3. Open each tab and set the parameters as desired:

   a. In the **Basics** tab, set the scope for the policy. To use centralized management, assign the policy to the Management Group containing the subscriptions that will use the workflow automation configuration.

   b. In the **Parameters** tab, set the resource group and data type details.

TIP

Each parameter has a tooltip explaining the options available to you.

Azure Policy's parameters tab (1) provides access to similar configuration options as Defender for Cloud's workflow automation page (2).

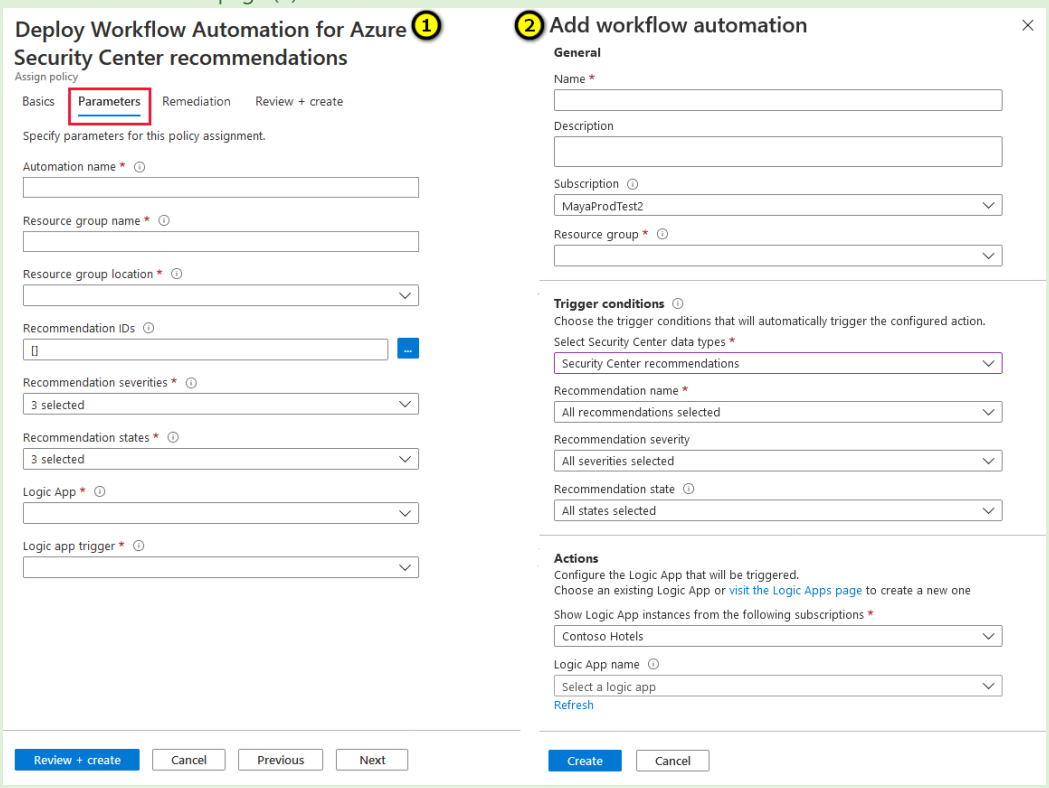

c. Optionally, to apply this assignment to existing subscriptions, open the **Remediation** tab and select the option to create a remediation task.

4. Review the summary page and select **Create**.

Data types schemas

To view the raw event schemas of the security alerts or recommendations events passed to the Logic App instance, visit the Workflow automation data types schemas. This can be useful in cases where you are not using Defender for Cloud's built-in Logic App connectors mentioned above, but instead are using Logic App's generic HTTP connector - you could use the event JSON schema to manually parse it as you see fit.

FAQ - Workflow automation

Does workflow automation support any business continuity or disaster recovery (BCDR) scenarios?

When preparing your environment for BCDR scenarios, where the target resource is experiencing an outage or other disaster, it's the organization's responsibility to prevent data loss by establishing backups according to the guidelines from Azure Event Hubs, Log Analytics workspace, and Logic App.

For every active automation, we recommend you create an identical (disabled) automation and store it in a different location. When there's an outage, you can enable these backup automations and maintain normal operations.

Learn more about Business continuity and disaster recovery for Azure Logic Apps.

Next steps

In this article, you learned about creating Logic Apps, automating their execution in Defender for Cloud, and running them manually.

For related material, see:

- The Microsoft Learn module on how to use workflow automation to automate a security response
- Security recommendations in Microsoft Defender for Cloud
- Security alerts in Microsoft Defender for Cloud
- About Azure Logic Apps
- Connectors for Azure Logic Apps
- Workflow automation data types schemas

Archive for what's new in Defender for Cloud?

2/15/2022 • 119 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

The primary What's new in Defender for Cloud? release notes page contains updates for the last six months, while this page contains older items.

This page provides you with information about:

- New features
- Bug fixes
- Deprecated functionality

August 2021

Updates in August include:

- Microsoft Defender for Endpoint for Linux now supported by Azure Defender for servers (in preview)
- Two new recommendations for managing endpoint protection solutions (in preview)
- Built-in troubleshooting and guidance for solving common issues
- Regulatory compliance dashboard's Azure Audit reports released for general availability (GA)
- Deprecated recommendation 'Log Analytics agent health issues should be resolved on your machines'
- Azure Defender for container registries now scans for vulnerabilities in registries protected with Azure Private Link
- Security Center can now auto provision the Azure Policy's Guest Configuration extension (in preview)
- Recommendations to enable Azure Defender plans now support "Enforce"
- CSV exports of recommendation data now limited to 20 MB
- Recommendations page now includes multiple views

Microsoft Defender for Endpoint for Linux now supported by Azure Defender for servers (in preview)

Azure Defender for servers includes an integrated license for Microsoft Defender for Endpoint. Together, they provide comprehensive endpoint detection and response (EDR) capabilities.

When Defender for Endpoint detects a threat, it triggers an alert. The alert is shown in Security Center. From Security Center, you can also pivot to the Defender for Endpoint console, and perform a detailed investigation to uncover the scope of the attack.

During the preview period, you'll deploy the Defender for Endpoint for Linux sensor to supported Linux machines in one of two ways depending on whether you've already deployed it to your Windows machines:

- Existing users with Defender for Cloud's enhanced security features enabled and Microsoft Defender for Endpoint for Windows
- New users who have never enabled the integration with Microsoft Defender for Endpoint for Windows

Learn more in Protect your endpoints with Security Center's integrated EDR solution: Microsoft Defender for

Endpoint.

Two new recommendations for managing endpoint protection solutions (in preview)

We've added two **preview** recommendations to deploy and maintain the endpoint protection solutions on your machines. Both recommendations include support for Azure virtual machines and machines connected to Azure Arc-enabled servers.

RECOMMENDATION	DESCRIPTION	SEVERITY
Endpoint protection should be installed on your machines	To protect your machines from threats and vulnerabilities, install a supported endpoint protection solution. Learn more about how Endpoint Protection for machines is evaluated. (Related policy: Monitor missing Endpoint Protection in Azure Security Center)	High
Endpoint protection health issues should be resolved on your machines	Resolve endpoint protection health issues on your virtual machines to protect them from latest threats and vulnerabilities. Azure Security Center supported endpoint protection solutions are documented here. Endpoint protection assessment is documented here. (Related policy: Monitor missing Endpoint Protection in Azure Security Center)	Medium

NOTE

The recommendations show their freshness interval as 8 hours, but there are some scenarios in which this might take significantly longer. For example, when an on premises machine is deleted, it takes 24 hours for Security Center to identify the deletion. After that, the assessment will take up to 8 hours to return the information. In that specific situation therefore, it may take 32 hours for the machine to be removed from the list of affected resources.

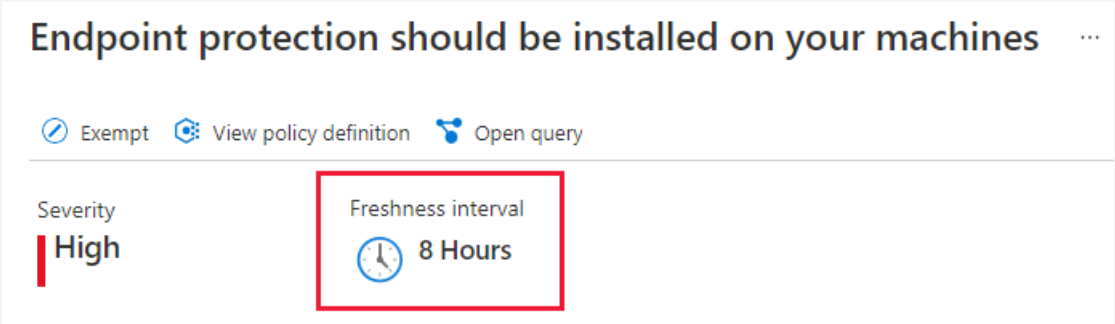

Built-in troubleshooting and guidance for solving common issues

A new, dedicated area of the Security Center pages in the Azure portal provides a collated, ever-growing set of self-help materials for solving common challenges with Security Center and Azure Defender.

When you're facing an issue, or are seeking advice from our support team, **Diagnose and solve problems** is another tool to help you find the solution:

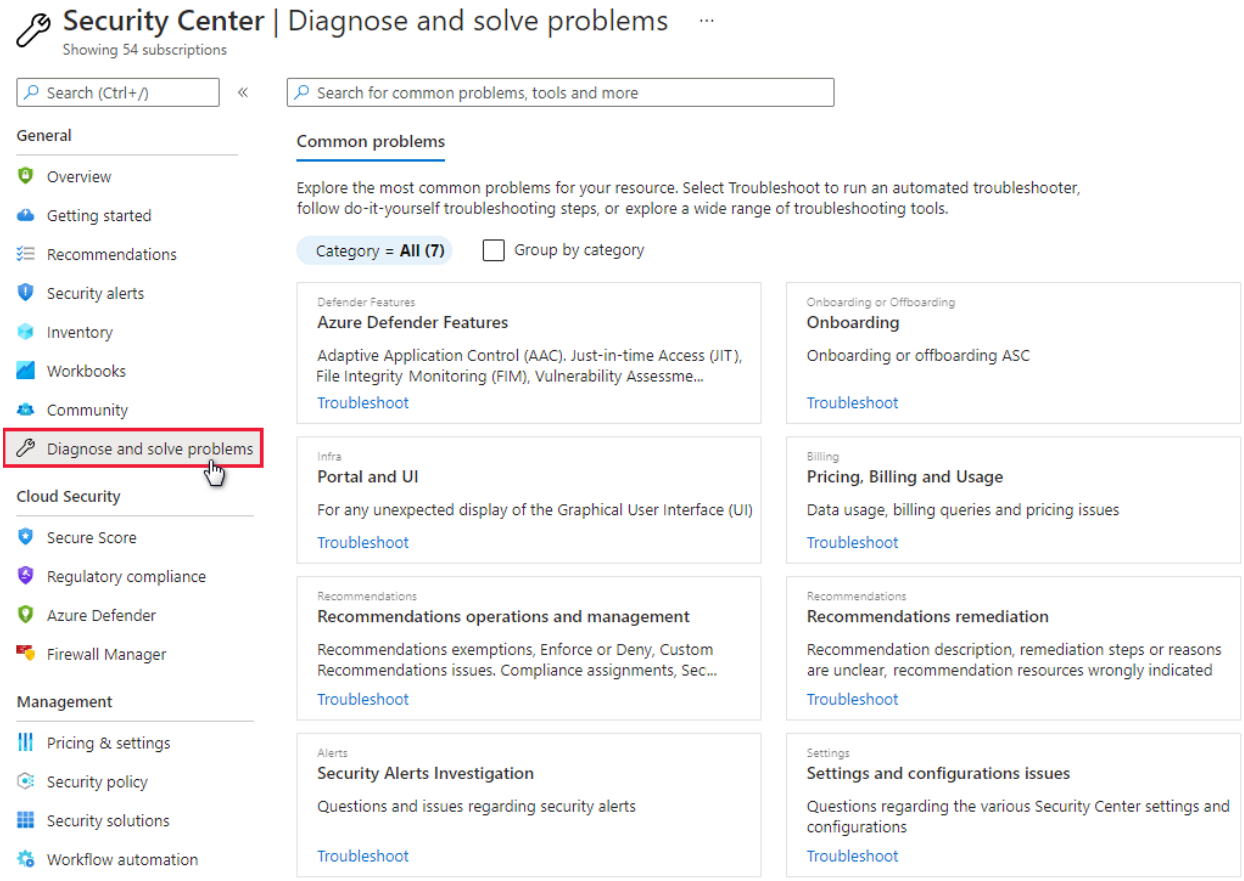

Regulatory compliance dashboard's Azure Audit reports released for general availability (GA)

The regulatory compliance dashboard's toolbar offers Azure and Dynamics certification reports for the standards applied to your subscriptions.

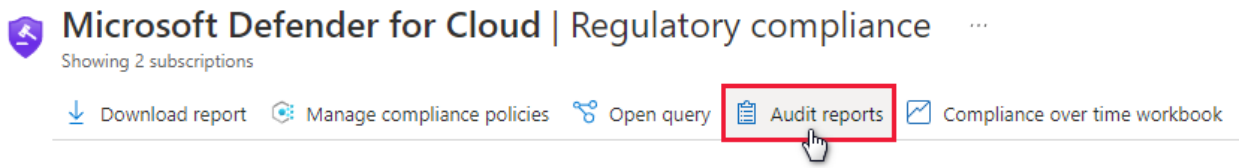

You can select the tab for the relevant reports types (PCI, SOC, ISO, and others) and use filters to find the specific reports you need.

For more information, see Generate compliance status reports and certificates.

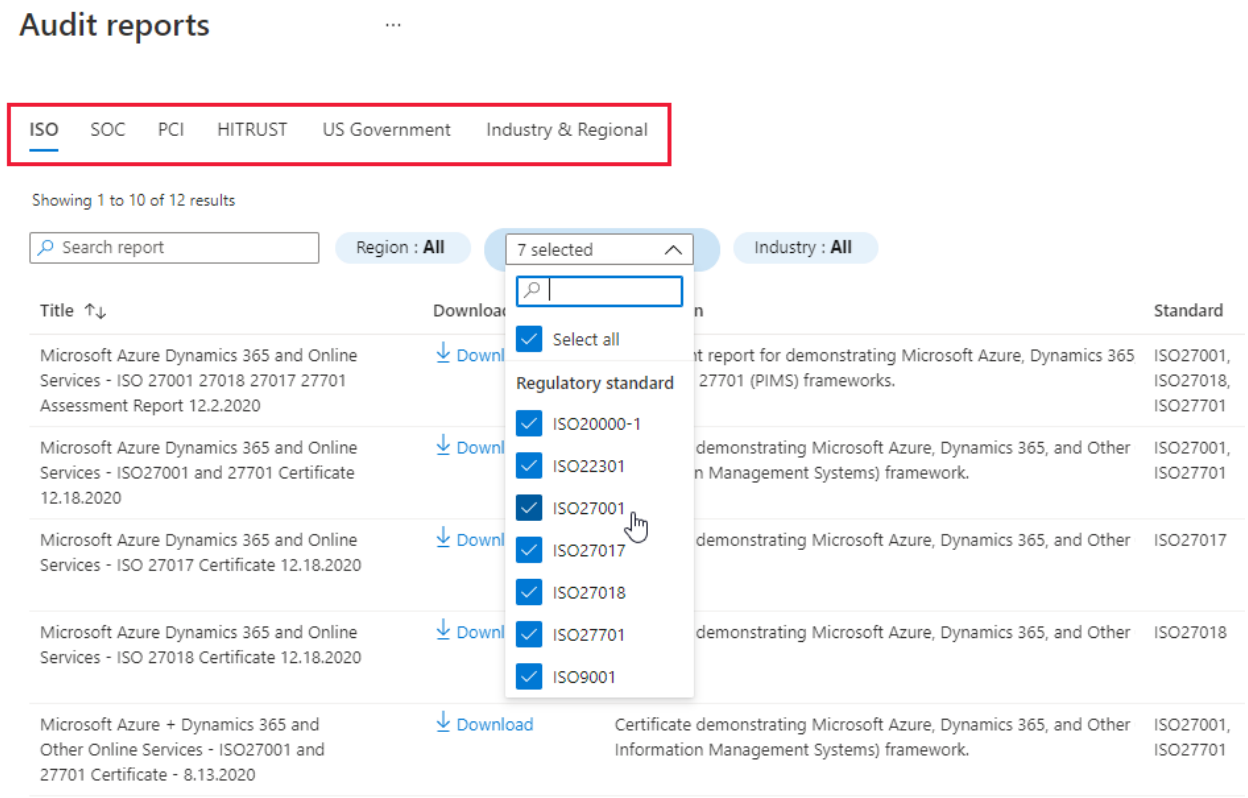

Deprecated recommendation 'Log Analytics agent health issues should be resolved on your machines'

We've found that recommendation **Log Analytics agent health issues should be resolved on your machines** impacts secure scores in ways that are inconsistent with Security Center's Cloud Security Posture Management (CSPM) focus. Typically, CSPM relates to identifying security misconfigurations. Agent health issues don't fit into this category of issues.

Also, the recommendation is an anomaly when compared with the other agents related to Security Center: this is the only agent with a recommendation related to health issues.

The recommendation has been deprecated.

As a result of this deprecation, we've also made minor changes to the recommendations for installing the Log Analytics agent (**Log Analytics agent should be installed on...**).

It's likely that this change will impact your secure scores. For most subscriptions, we expect the change to lead to an increased score, but it's possible the updates to the installation recommendation might result in decreased scores in some cases.

> **TIP**
>
> The asset inventory page was also affected by this change as it displays the monitored status for machines (monitored, not monitored, or partially monitored - a state which refers to an agent with health issues).

Azure Defender for container registries now scans for vulnerabilities in registries protected with Azure Private Link

Azure Defender for container registries includes a vulnerability scanner to scan images in your Azure Container Registry registries. Learn how to scan your registries and remediate findings in Use Azure Defender for container registries to scan your images for vulnerabilities.

To limit access to a registry hosted in Azure Container Registry, assign virtual network private IP addresses to the registry endpoints and use Azure Private Link as explained in Connect privately to an Azure container registry using Azure Private Link.

As part of our ongoing efforts to support additional environments and use cases, Azure Defender now also

scans container registries protected with Azure Private Link.

Security Center can now auto provision the Azure Policy's Guest Configuration extension (in preview)

Azure Policy can audit settings inside a machine, both for machines running in Azure and Arc connected machines. The validation is performed by the Guest Configuration extension and client. Learn more in Understand Azure Policy's Guest Configuration.

With this update, you can now set Security Center to automatically provision this extension to all supported machines.

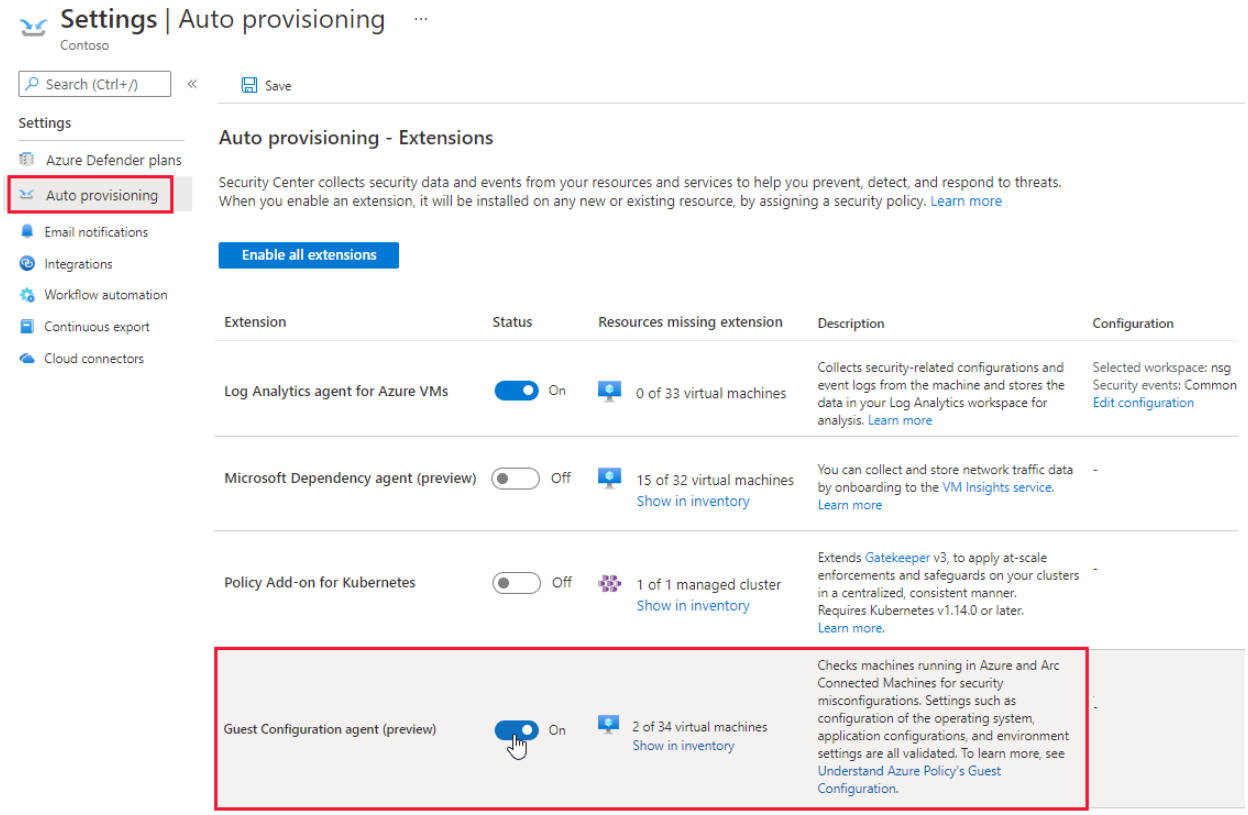

Learn more about how auto provisioning works in Configure auto provisioning for agents and extensions.

Recommendations to enable Azure Defender plans now support "Enforce"

Security Center includes two features that help ensure newly created resources are provisioned in a secure manner: **enforce** and **deny**. When a recommendation offers these options, you can ensure your security requirements are met whenever someone attempts to create a resource:

- **Deny** stops unhealthy resources from being created
- **Enforce** automatically remediates non-compliant resources when they're created

With this update, the enforce option is now available on the recommendations to enable Azure Defender plans (such as **Azure Defender for App Service should be enabled**, **Azure Defender for Key Vault should be enabled**, **Azure Defender for Storage should be enabled**).

Learn more about these options in Prevent misconfigurations with Enforce/Deny recommendations.

CSV exports of recommendation data now limited to 20 MB

We're instituting a limit of 20 MB when exporting Security Center recommendations data.

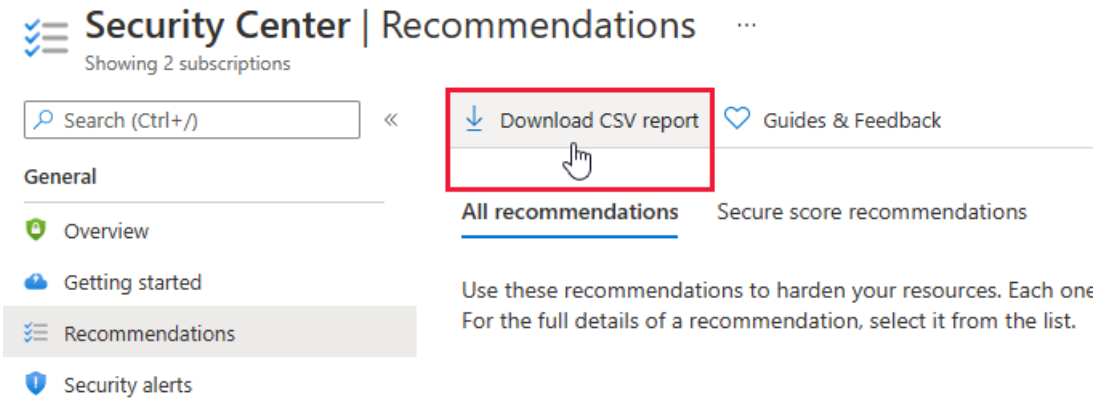

If you need to export larger amounts of data, use the available filters before selecting, or select subsets of your subscriptions and download the data in batches.

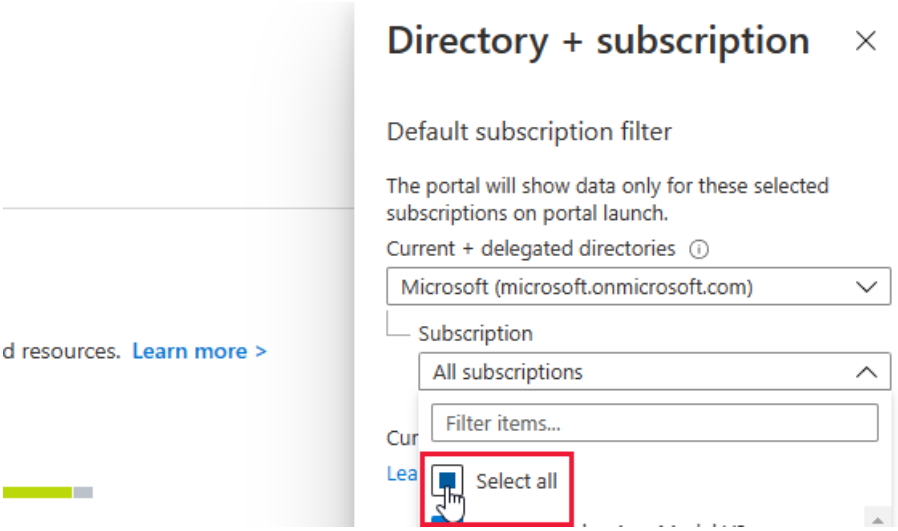

Learn more about performing a CSV export of your security recommendations.

Recommendations page now includes multiple views

The recommendations page now has two tabs to provide alternate ways to view the recommendations relevant to your resources:

- **Secure score recommendations** - Use this tab to view the list of recommendations grouped by security control. Learn more about these controls in Security controls and their recommendations.
- **All recommendations** - Use this tab to view the list of recommendations as a flat list. This tab is also great for understanding which initiative (including regulatory compliance standards) generated the recommendation. Learn more about initiatives and their relationship to recommendations in What are security policies, initiatives, and recommendations?.

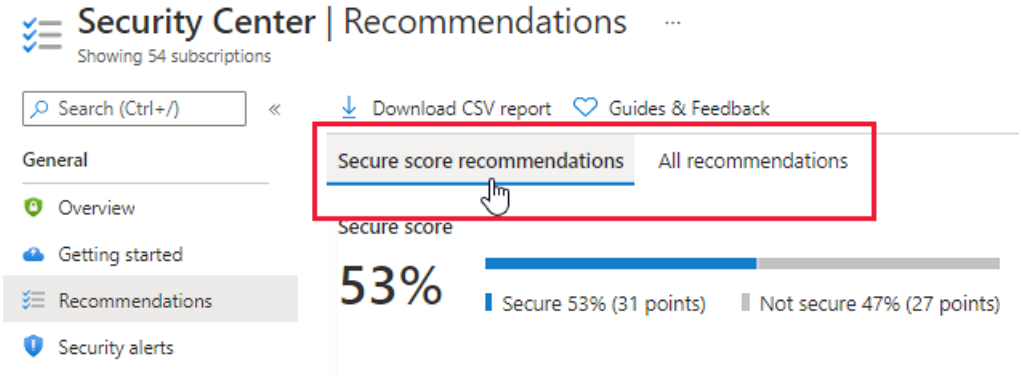

July 2021

Updates in July include:

- Azure Sentinel connector now includes optional bi-directional alert synchronization (in preview)
- Logical reorganization of Azure Defender for Resource Manager alerts
- Enhancements to recommendation to enable Azure Disk Encryption (ADE)
- Continuous export of secure score and regulatory compliance data released for general availability (GA)
- Workflow automations can be triggered by changes to regulatory compliance assessments (GA)
- Assessments API field 'FirstEvaluationDate' and 'StatusChangeDate' now available in workspace schemas and logic apps
- 'Compliance over time' workbook template added to Azure Monitor Workbooks gallery

Azure Sentinel connector now includes optional bi-directional alert synchronization (in preview)

Security Center natively integrates with Azure Sentinel, Azure's cloud-native SIEM and SOAR solution.

Azure Sentinel includes built-in connectors for Azure Security Center at the subscription and tenant levels. Learn more in Stream alerts to Azure Sentinel.

When you connect Azure Defender to Azure Sentinel, the status of Azure Defender alerts that get ingested into Azure Sentinel is synchronized between the two services. So, for example, when an alert is closed in Azure Defender, that alert will display as closed in Azure Sentinel as well. Changing the status of an alert in Azure Defender "won't"* affect the status of any Azure Sentinel **incidents** that contain the synchronized Azure Sentinel alert, only that of the synchronized alert itself.

Enabling this preview feature, **bi-directional alert synchronization**, will automatically sync the status of the original Azure Defender alerts with Azure Sentinel incidents that contain the copies of those Azure Defender alerts. So, for example, when an Azure Sentinel incident containing an Azure Defender alert is closed, Azure Defender will automatically close the corresponding original alert.

Learn more in Connect Azure Defender alerts from Azure Security Center.

Logical reorganization of Azure Defender for Resource Manager alerts

The alerts listed below were provided as part of the Azure Defender for Resource Manager plan.

As part of a logical reorganization of some of the Azure Defender plans, we've moved some alerts from **Azure Defender for Resource Manager** to **Azure Defender for servers**.

The alerts are organized according to two main principles:

- Alerts that provide control-plane protection - across many Azure resource types - are part of Azure Defender for Resource Manager
- Alerts that protect specific workloads are in the Azure Defender plan that relates to the corresponding workload

These are the alerts that were part of Azure Defender for Resource Manager, and which, as a result of this change, are now part of Azure Defender for servers:

- ARM_AmBroadFilesExclusion
- ARM_AmDisablementAndCodeExecution
- ARM_AmDisablement
- ARM_AmFileExclusionAndCodeExecution
- ARM_AmTempFileExclusionAndCodeExecution
- ARM_AmTempFileExclusion
- ARM_AmRealtimeProtectionDisabled
- ARM_AmTempRealtimeProtectionDisablement
- ARM_AmRealtimeProtectionDisablementAndCodeExec

- ARM_AmMalwareCampaignRelatedExclusion
- ARM_AmTemporarilyDisablement
- ARM_UnusualAmFileExclusion
- ARM_CustomScriptExtensionSuspiciousCmd
- ARM_CustomScriptExtensionSuspiciousEntryPoint
- ARM_CustomScriptExtensionSuspiciousPayload
- ARM_CustomScriptExtensionSuspiciousFailure
- ARM_CustomScriptExtensionUnusualDeletion
- ARM_CustomScriptExtensionUnusualExecution
- ARM_VMAccessUnusualConfigReset
- ARM_VMAccessUnusualPasswordReset
- ARM_VMAccessUnusualSSHReset

Learn more about the Azure Defender for Resource Manager and Azure Defender for servers plans.

Enhancements to recommendation to enable Azure Disk Encryption (ADE)

Following user feedback, we've renamed the recommendation **Disk encryption should be applied on virtual machines**.

The new recommendation uses the same assessment ID and is called **Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources**.

The description has also been updated to better explain the purpose of this hardening recommendation:

RECOMMENDATION	DESCRIPTION	SEVERITY
Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources	By default, a virtual machine's OS and data disks are encrypted-at-rest using platform-managed keys; temp disks and data caches aren't encrypted, and data isn't encrypted when flowing between compute and storage resources. For a comparison of different disk encryption technologies in Azure, see https://aka.ms/diskencryptioncomparison. Use Azure Disk Encryption to encrypt all this data. Disregard this recommendation if: (1) you're using the encryption-at-host feature, or (2) server-side encryption on Managed Disks meets your security requirements. Learn more in Server-side encryption of Azure Disk Storage.	High

Continuous export of secure score and regulatory compliance data released for general availability (GA)

Continuous export provides the mechanism for exporting your security alerts and recommendations for tracking with other monitoring tools in your environment.

When you set up your continuous export, you configure what is exported, and where it will go. Learn more in the overview of continuous export.

We've enhanced and expanded this feature over time:

- In November 2020, we added the **preview** option to stream changes to your **secure score**.

For full details, see Secure score is now available in continuous export (preview).

- In December 2020, we added the **preview** option to stream changes to your **regulatory compliance assessment data**.
  For full details, see Continuous export gets new data types (preview).

With this update, these two options are released for general availability (GA).

Workflow automations can be triggered by changes to regulatory compliance assessments (GA)

In February 2021, we added a **preview** third data type to the trigger options for your workflow automations: changes to regulatory compliance assessments. Learn more in Workflow automations can be triggered by changes to regulatory compliance assessments.

With this update, this trigger option is released for general availability (GA).

Learn how to use the workflow automation tools in Automate responses to Security Center triggers.

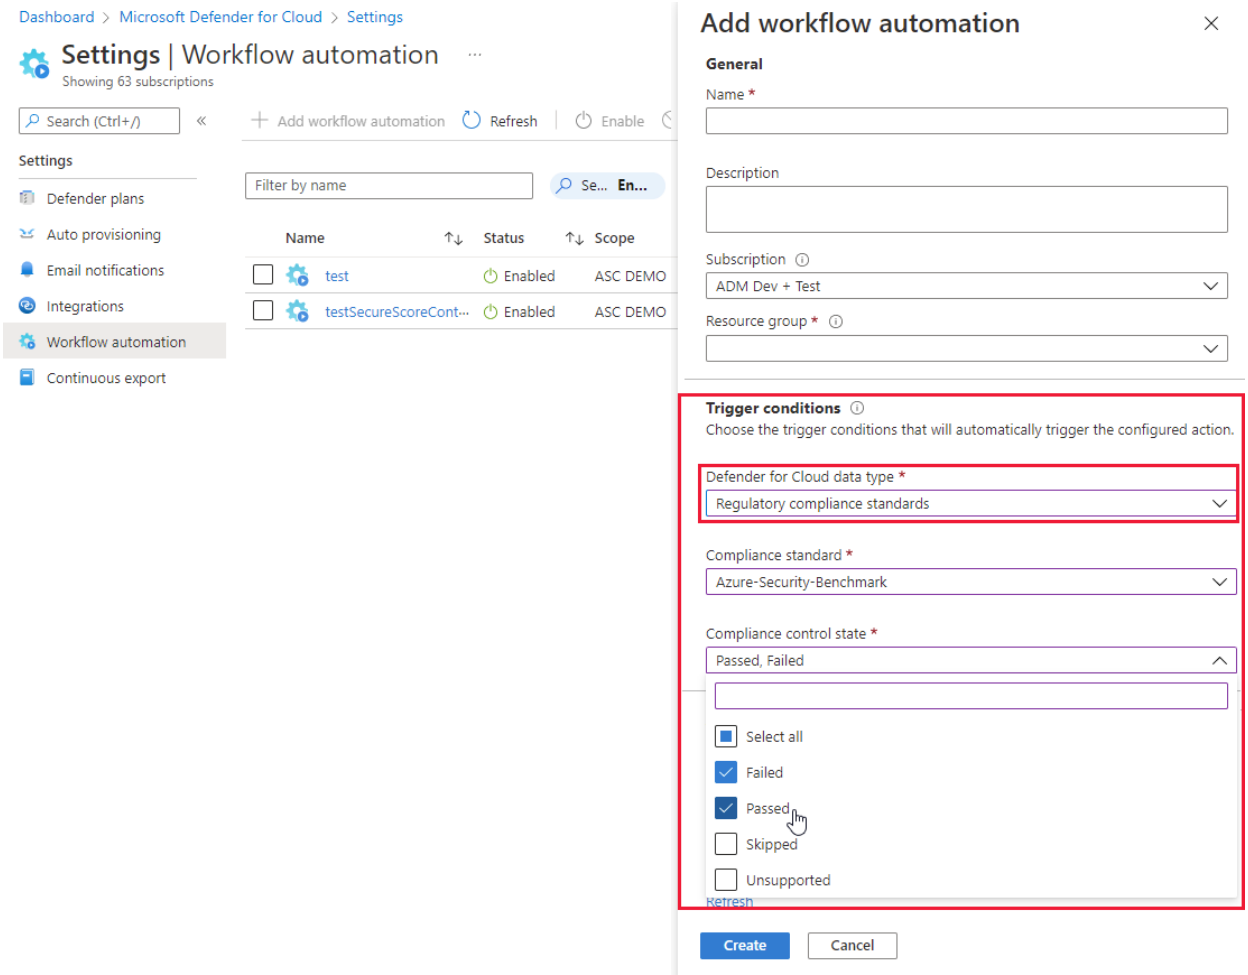

Assessments API field 'FirstEvaluationDate' and 'StatusChangeDate' now available in workspace schemas and logic apps

In May 2021, we updated the Assessment API with two new fields, **FirstEvaluationDate** and **StatusChangeDate**. For full details, see Assessments API expanded with two new fields.

Those fields were accessible through the REST API, Azure Resource Graph, continuous export, and in CSV exports.

With this change, we're making the information available in the Log Analytics workspace schema and from logic apps.

'Compliance over time' workbook template added to Azure Monitor Workbooks gallery

In March, we announced the integrated Azure Monitor Workbooks experience in Security Center (see Azure

Monitor Workbooks integrated into Security Center and three templates provided).

The initial release included three templates to build dynamic and visual reports about your organization's security posture.

We've now added a workbook dedicated to tracking a subscription's compliance with the regulatory or industry standards applied to it.

Learn about using these reports or building your own in Create rich, interactive reports of Security Center data.

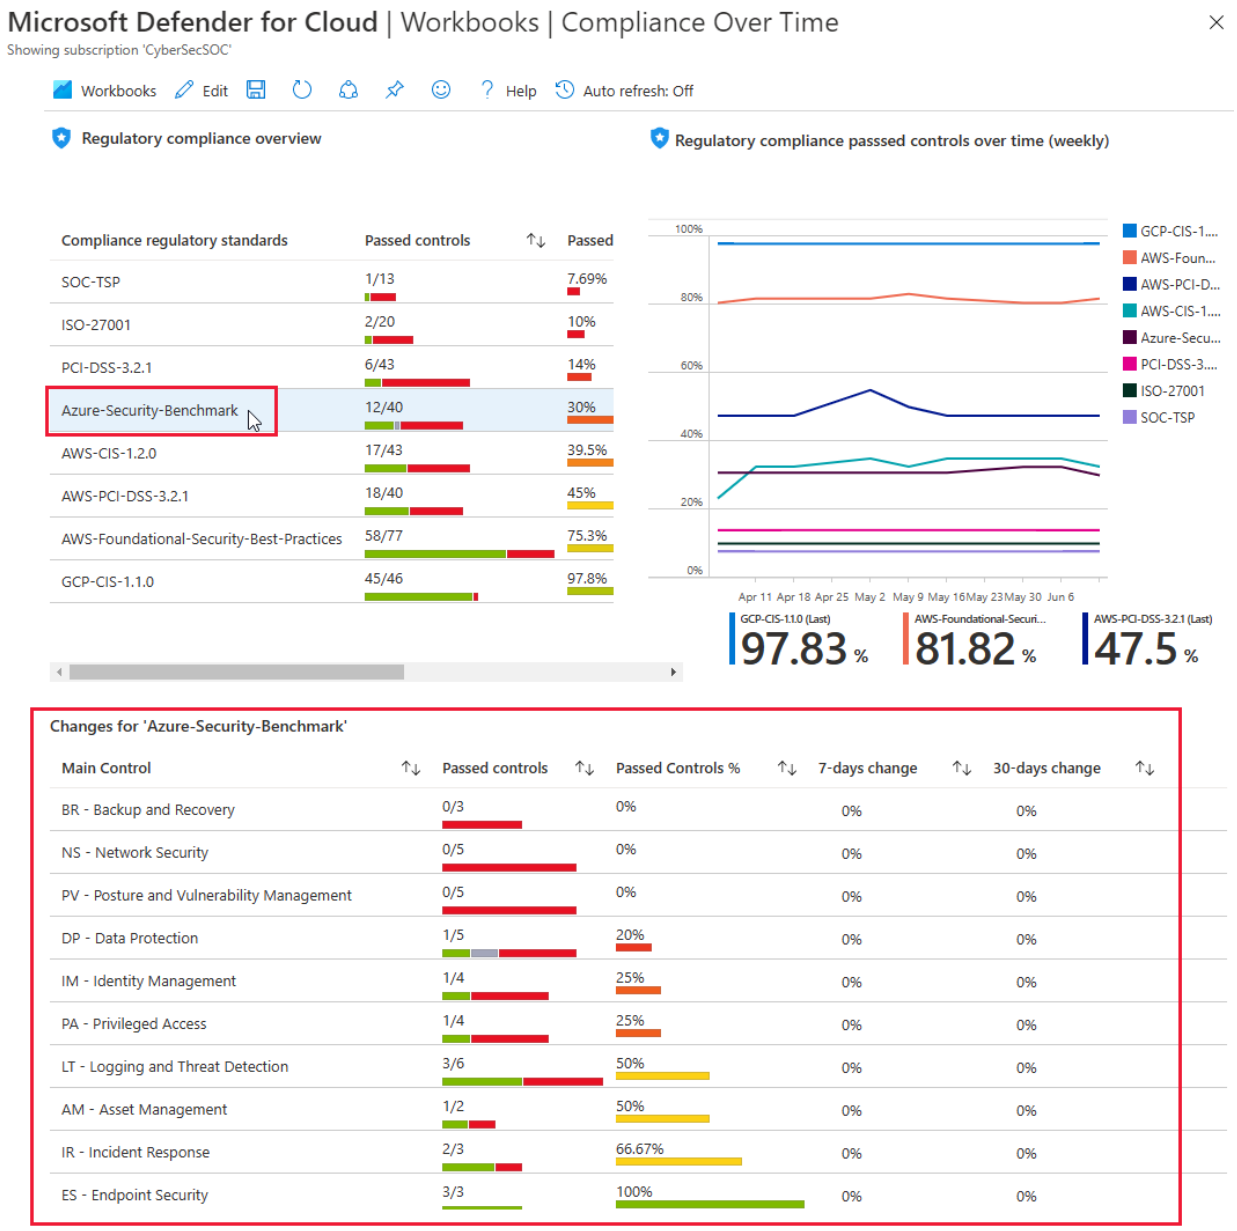

Changes for 'Azure-Security-Benchmark'

Main Control	Passed controls	Passed Controls %	7-days change	30-days change
BR - Backup and Recovery	0/3	0%	0%	0%
NS - Network Security	0/5	0%	0%	0%
PV - Posture and Vulnerability Management	0/5	0%	0%	0%
DP - Data Protection	1/5	20%	0%	0%
IM - Identity Management	1/4	25%	0%	0%
PA - Privileged Access	1/4	25%	0%	0%
LT - Logging and Threat Detection	3/6	50%	0%	0%
AM - Asset Management	1/2	50%	0%	0%
IR - Incident Response	2/3	66.67%	0%	0%
ES - Endpoint Security	3/3	100%	0%	0%

June 2021

Updates in June include:

- New alert for Azure Defender for Key Vault
- Recommendations to encrypt with customer-managed keys (CMKs) disabled by default
- Prefix for Kubernetes alerts changed from "AKS_" to "K8S_"
- Deprecated two recommendations from "Apply system updates" security control

New alert for Azure Defender for Key Vault

To expand the threat protections provided by Azure Defender for Key Vault, we've added the following alert:

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTIC	SEVERITY
Access from a suspicious IP address to a key vault (KV_SuspiciousIPAccess)	A key vault has been successfully accessed by an IP that has been identified by Microsoft Threat Intelligence as a suspicious IP address. This may indicate that your infrastructure has been compromised. We recommend further investigation. Learn more about Microsoft's threat intelligence capabilities.	Credential Access	Medium

For more information, see:

- Introduction to Azure Defender for Key Vault
- Respond to Azure Defender for Key Vault alerts
- List of alerts provided by Azure Defender for Key Vault

Recommendations to encrypt with customer-managed keys (CMKs) disabled by default

Security Center includes multiple recommendations to encrypt data at rest with customer-managed keys, such as:

- Container registries should be encrypted with a customer-managed key (CMK)
- Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest
- Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK)

Data in Azure is encrypted automatically using platform-managed keys, so the use of customer-managed keys should only be applied when required for compliance with a specific policy your organization is choosing to enforce.

With this change, the recommendations to use CMKs are now **disabled by default**. When relevant for your organization, you can enable them by changing the *Effect* parameter for the corresponding security policy to **AuditIfNotExists** or **Enforce**. Learn more in Enable a security policy.

This change is reflected in the names of the recommendation with a new prefix, **[Enable if required]**, as shown in the following examples:

- [Enable if required] Storage accounts should use customer-managed key to encrypt data at rest
- [Enable if required] Container registries should be encrypted with a customer-managed key (CMK)
- [Enable if required] Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest

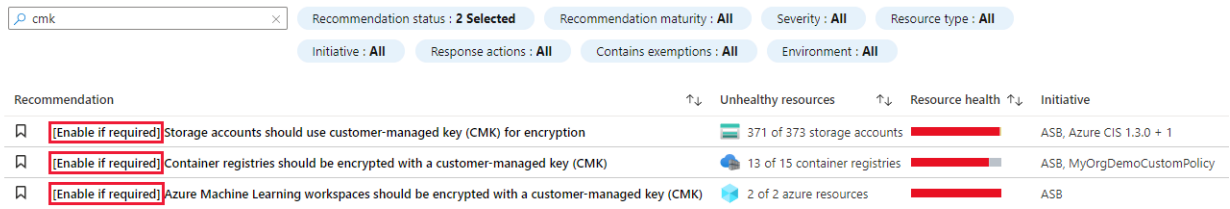

Prefix for Kubernetes alerts changed from "AKS_" to "K8S_"

Azure Defender for Kubernetes recently expanded to protect Kubernetes clusters hosted on-premises and in multi-cloud environments. Learn more in Use Azure Defender for Kubernetes to protect hybrid and multi-cloud Kubernetes deployments (in preview).

To reflect the fact that the security alerts provided by Azure Defender for Kubernetes are no longer restricted to clusters on Azure Kubernetes Service, we've changed the prefix for the alert types from "AKS_" to "K8S_". Where necessary, the names and descriptions were updated too. For example, this alert:

ALERT (ALERT TYPE)	DESCRIPTION
Kubernetes penetration testing tool detected (**AKS**_PenTestToolsKubeHunter)	Kubernetes audit log analysis detected usage of Kubernetes penetration testing tool in the **AKS** cluster. While this behavior can be legitimate, attackers might use such public tools for malicious purposes.

was changed to:

ALERT (ALERT TYPE)	DESCRIPTION
Kubernetes penetration testing tool detected (**K8S**_PenTestToolsKubeHunter)	Kubernetes audit log analysis detected usage of Kubernetes penetration testing tool in the **Kubernetes** cluster. While this behavior can be legitimate, attackers might use such public tools for malicious purposes.

Any suppression rules that refer to alerts beginning "AKS_" were automatically converted. If you've setup SIEM exports, or custom automation scripts that refer to Kubernetes alerts by alert type, you'll need to update them with the new alert types.

For a full list of the Kubernetes alerts, see Alerts for Kubernetes clusters.

Deprecated two recommendations from "Apply system updates" security control

The following two recommendations were deprecated:

- **OS version should be updated for your cloud service roles** - By default, Azure periodically updates your guest OS to the latest supported image within the OS family that you've specified in your service configuration (.cscfg), such as Windows Server 2016.
- **Kubernetes Services should be upgraded to a non-vulnerable Kubernetes version** - This recommendation's evaluations aren't as wide-ranging as we'd like them to be. We plan to replace the recommendation with an enhanced version that's better aligned with your security needs.

May 2021

Updates in May include:

- Azure Defender for DNS and Azure Defender for Resource Manager released for general availability (GA)
- Azure Defender for open-source relational databases released for general availability (GA)
- New alerts for Azure Defender for Resource Manager
- CI/CD vulnerability scanning of container images with GitHub workflows and Azure Defender (preview)
- More Resource Graph queries available for some recommendations
- SQL data classification recommendation severity changed
- New recommendations to enable trusted launch capabilities (in preview)
- New recommendations for hardening Kubernetes clusters (in preview)
- Assessments API expanded with two new fields
- Asset inventory gets a cloud environment filter

Azure Defender for DNS and Azure Defender for Resource Manager released for general availability (GA)

These two cloud-native breadth threat protection plans are now GA.

These new protections greatly enhance your resiliency against attacks from threat actors, and significantly increase the number of Azure resources protected by Azure Defender.

- **Azure Defender for Resource Manager** - automatically monitors all resource management operations performed in your organization. For more information, see:
  - Introduction to Azure Defender for Resource Manager
  - Respond to Azure Defender for Resource Manager alerts
  - List of alerts provided by Azure Defender for Resource Manager
- **Azure Defender for DNS** - continuously monitors all DNS queries from your Azure resources. For more information, see:
  - Introduction to Azure Defender for DNS
  - Respond to Azure Defender for DNS alerts
  - List of alerts provided by Azure Defender for DNS

To simplify the process of enabling these plans, use the recommendations:

- **Azure Defender for Resource Manager should be enabled**
- **Azure Defender for DNS should be enabled**

> **NOTE**
> Enabling Azure Defender plans results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center.

Azure Defender for open-source relational databases released for general availability (GA)

Azure Security Center expands its offer for SQL protection with a new bundle to cover your open-source relational databases:

- **Azure Defender for Azure SQL database servers** - defends your Azure-native SQL Servers
- **Azure Defender for SQL servers on machines** - extends the same protections to your SQL servers in hybrid, multi-cloud, and on-premises environments
- **Azure Defender for open-source relational databases** - defends your Azure Databases for MySQL, PostgreSQL, and MariaDB single servers

Azure Defender for open-source relational databases constantly monitors your servers for security threats and detects anomalous database activities indicating potential threats to Azure Database for MySQL, PostgreSQL, and MariaDB. Some examples are:

- **Granular detection of brute force attacks** - Azure Defender for open-source relational databases provides detailed information on attempted and successful brute force attacks. This lets you investigate and respond with a more complete understanding of the nature and status of the attack on your environment.
- **Behavioral alerts detection** - Azure Defender for open-source relational databases alerts you to suspicious and unexpected behaviors on your servers, such as changes in the access pattern to your database.
- **Threat intelligence-based detection** - Azure Defender applies Microsoft's threat intelligence and vast knowledge base to surface threat alerts so you can act against them.

Learn more in Introduction to Azure Defender for open-source relational databases.

New alerts for Azure Defender for Resource Manager

To expand the threat protections provided by Azure Defender for Resource Manager, we've added the following

alerts:

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS	SEVERITY
Permissions granted for an RBAC role in an unusual way for your Azure environment (Preview) (ARM_AnomalousRBACRoleAssignment)	Azure Defender for Resource Manager detected an RBAC role assignment that's unusual when compared with other assignments performed by the same assigner / performed for the same assignee / in your tenant due to the following anomalies: assignment time, assigner location, assigner, authentication method, assigned entities, client software used, assignment extent. This operation might have been performed by a legitimate user in your organization. Alternatively, it might indicate that an account in your organization was breached, and that the threat actor is trying to grant permissions to an additional user account they own.	Lateral Movement, Defense Evasion	Medium
Privileged custom role created for your subscription in a suspicious way (Preview) (ARM_PrivilegedRoleDefinitionCreation)	Azure Defender for Resource Manager detected a suspicious creation of privileged custom role definition in your subscription. This operation might have been performed by a legitimate user in your organization. Alternatively, it might indicate that an account in your organization was breached, and that the threat actor is trying to create a privileged role to use in the future to evade detection.	Lateral Movement, Defense Evasion	Low
Azure Resource Manager operation from suspicious IP address (Preview) (ARM_OperationFromSuspiciousIP)	Azure Defender for Resource Manager detected an operation from an IP address that has been marked as suspicious in threat intelligence feeds.	Execution	Medium

ALERT (ALERT TYPE)	DESCRIPTION	MITRE TACTICS	SEVERITY
Azure Resource Manager operation from suspicious proxy IP address (Preview) (ARM_OperationFromSuspiciousProxyIP)	Azure Defender for Resource Manager detected a resource management operation from an IP address that is associated with proxy services, such as TOR. While this behavior can be legitimate, it's often seen in malicious activities, when threat actors try to hide their source IP.	Defense Evasion	Medium

For more information, see:

- Introduction to Azure Defender for Resource Manager
- Respond to Azure Defender for Resource Manager alerts
- List of alerts provided by Azure Defender for Resource Manager

CI/CD vulnerability scanning of container images with GitHub workflows and Azure Defender (preview)

Azure Defender for container registries now provides DevSecOps teams observability into GitHub Action workflows.

The new vulnerability scanning feature for container images, utilizing Trivy, helps your developers scan for common vulnerabilities in their container images *before* pushing images to container registries.

Container scan reports are summarized in Azure Security Center, providing security teams better insight and understanding about the source of vulnerable container images and the workflows and repositories from where they originate.

Learn more in Identify vulnerable container images in your CI/CD workflows.

More Resource Graph queries available for some recommendations

All of Security Center's recommendations have the option to view the information about the status of affected resources using Azure Resource Graph from the **Open query**. For full details about this powerful feature, see Review recommendation data in Azure Resource Graph Explorer (ARG).

Security Center includes built-in vulnerability scanners to scan your VMs, SQL servers and their hosts, and container registries for security vulnerabilities. The findings are returned as recommendations with all the individual findings for each resource type gathered into a single view. The recommendations are:

- Vulnerabilities in Azure Container Registry images should be remediated (powered by Qualys)
- Vulnerabilities in your virtual machines should be remediated
- SQL databases should have vulnerability findings resolved
- SQL servers on machines should have vulnerability findings resolved

With this change, you can use the **Open query** button to also open the query showing the security findings.

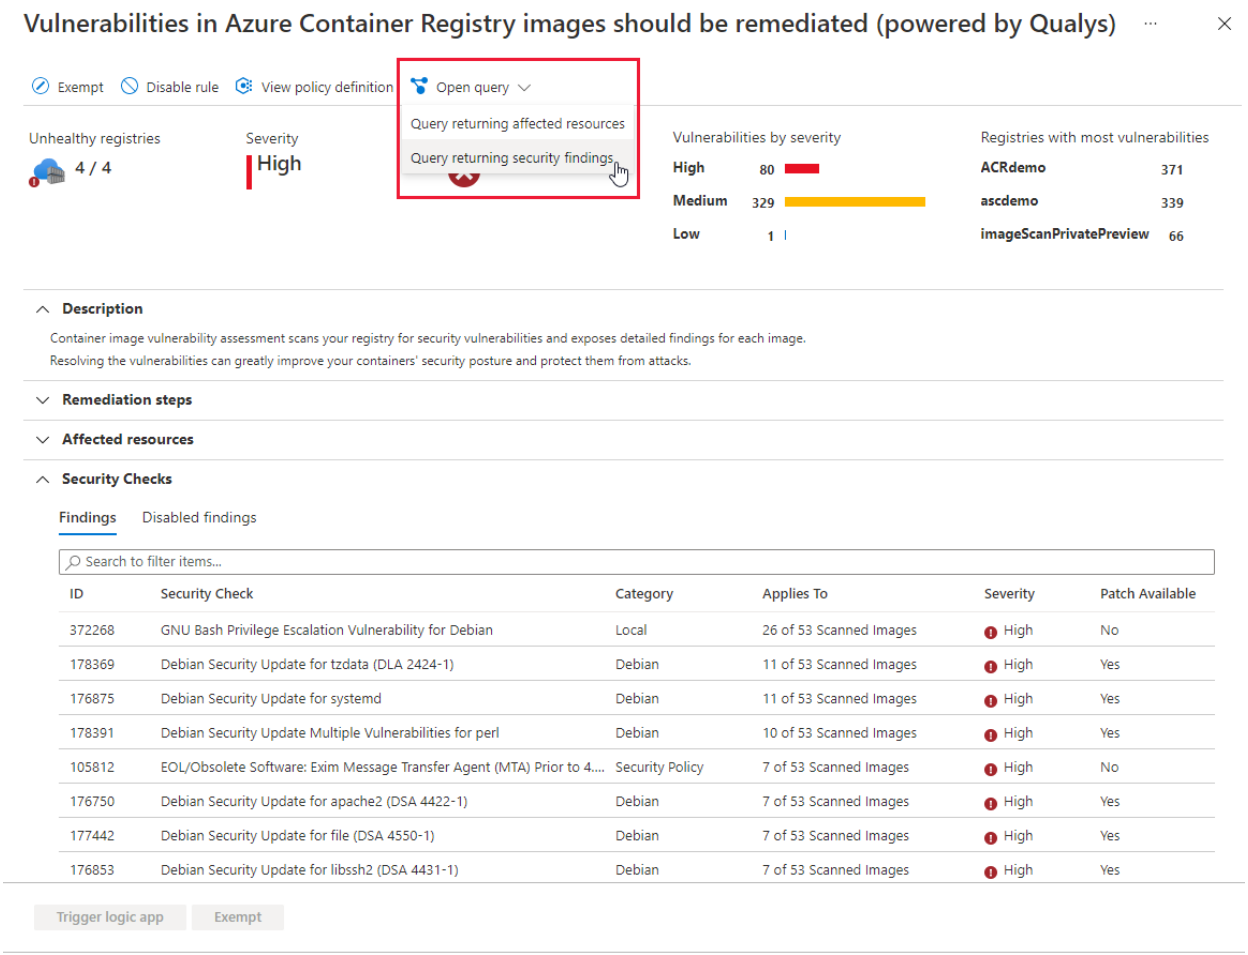

The **Open query** button offers additional options for some other recommendations where relevant.

Learn more about Security Center's vulnerability scanners:

- Azure Defender's integrated Qualys vulnerability scanner for Azure and hybrid machines
- Azure Defender's integrated vulnerability assessment scanner for SQL servers
- Azure Defender's integrated vulnerability assessment scanner for container registries

SQL data classification recommendation severity changed

The severity of the recommendation **Sensitive data in your SQL databases should be classified** has been changed from **High** to **Low**.

This is part of an ongoing change to this recommendation announced in our upcoming changes page.

New recommendations to enable trusted launch capabilities (in preview)

Azure offers trusted launch as a seamless way to improve the security of generation 2 VMs. Trusted launch protects against advanced and persistent attack techniques. Trusted launch is composed of several, coordinated infrastructure technologies that can be enabled independently. Each technology provides another layer of defense against sophisticated threats. Learn more in Trusted launch for Azure virtual machines.

> **IMPORTANT**
>
> Trusted launch requires the creation of new virtual machines. You can't enable trusted launch on existing virtual machines that were initially created without it.
>
> Trusted launch is currently in public preview. The preview is provided without a service level agreement, and it's not recommended for production workloads. Certain features might not be supported or might have constrained capabilities.

Security Center's recommendation, **vTPM should be enabled on supported virtual machines**, ensures

your Azure VMs are using a vTPM. This virtualized version of a hardware Trusted Platform Module enables attestation by measuring the entire boot chain of your VM (UEFI, OS, system, and drivers).

With the vTPM enabled, the **Guest Attestation extension** can remotely validate the secure boot. The following recommendations ensure this extension is deployed:

- **Secure Boot should be enabled on supported Windows virtual machines**
- **Guest Attestation extension should be installed on supported Windows virtual machines**
- **Guest Attestation extension should be installed on supported Windows virtual machine scale sets**
- **Guest Attestation extension should be installed on supported Linux virtual machines**
- **Guest Attestation extension should be installed on supported Linux virtual machine scale sets**

Learn more in Trusted launch for Azure virtual machines.

New recommendations for hardening Kubernetes clusters (in preview)

The following recommendations allow you to further harden your Kubernetes clusters

- **Kubernetes clusters should not use the default namespace** - To protect against unauthorized access for ConfigMap, Pod, Secret, Service, and ServiceAccount resource types, prevent usage of the default namespace in Kubernetes clusters.
- **Kubernetes clusters should disable automounting API credentials** - To prevent a potentially compromised Pod resource from running API commands against Kubernetes clusters, disable automounting API credentials.
- **Kubernetes clusters should not grant CAPSYSADMIN security capabilities**

Learn how Security Center can protect your containerized environments in Container security in Security Center.

Assessments API expanded with two new fields

We've added the following two fields to the Assessments REST API:

- **FirstEvaluationDate** – The time that the recommendation was created and first evaluated. Returned as UTC time in ISO 8601 format.
- **StatusChangeDate** – The time that the status of the recommendation last changed. Returned as UTC time in ISO 8601 format.

The initial default value for these fields - for all recommendations - is `2021-03-14T00:00:00+0000000Z`.

To access this information, you can use any of the methods in the table below.

TOOL	DETAILS
REST API call	GET https://management.azure.com/subscriptions/<SUBSCRIPTION_ID>/providers/Microsoft.Security/assessments?api-version=2019-01-01-preview&$expand=statusEvaluationDates
Azure Resource Graph	`securityresources` `where type == "microsoft.security/assessments"`
Continuous export	The two dedicated fields will be available the Log Analytics workspace data
CSV export	The two fields are included in the CSV files

TOOL	DETAILS

Learn more about the Assessments REST API.

Asset inventory gets a cloud environment filter

Security Center's asset inventory page offers many filters to quickly refine the list of resources displayed. Learn more in Explore and manage your resources with asset inventory.

A new filter offers the option to refine the list according to the cloud accounts you've connected with Security Center's multi-cloud features:

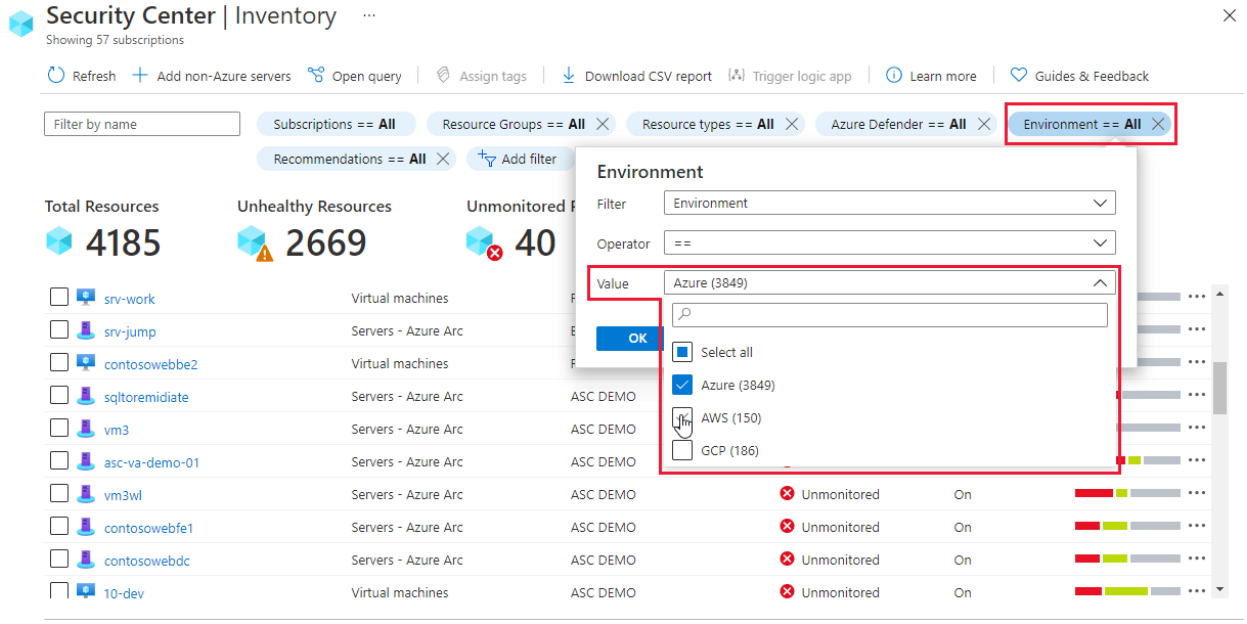

Learn more about the multi-cloud capabilities:

- Connect your AWS accounts to Azure Security Center
- Connect your GCP accounts to Azure Security Center

April 2021

Updates in April include:

- Refreshed resource health page (in preview)
- Container registry images that have been recently pulled are now rescanned weekly (released for general availability (GA))
- Use Azure Defender for Kubernetes to protect hybrid and multi-cloud Kubernetes deployments (in preview)
- Microsoft Defender for Endpoint integration with Azure Defender now supports Windows Server 2019 and Windows 10 Virtual Desktop (WVD) released for general availability (GA)
- Recommendations to enable Azure Defender for DNS and Resource Manager (in preview)
- Three regulatory compliance standards added: Azure CIS 1.3.0, CMMC Level 3, and New Zealand ISM Restricted
- Four new recommendations related to guest configuration (in preview)
- CMK recommendations moved to best practices security control
- 11 Azure Defender alerts deprecated
- Two recommendations from "Apply system updates" security control were deprecated
- Azure Defender for SQL on machine tile removed from Azure Defender dashboard
- 21 recommendations moved between security controls

Refreshed resource health page (in preview)

Security Center's resource health has been expanded, enhanced, and improved to provide a snapshot view of the overall health of a single resource.

You can review detailed information about the resource and all recommendations that apply to that resource. Also, if you're using the advanced protection plans of Microsoft Defender, you can see outstanding security alerts for that specific resource too.

To open the resource health page for a resource, select any resource from the asset inventory page.

This preview page in Security Center's portal pages shows:

1. **Resource information** - The resource group and subscription it's attached to, the geographic location, and more.
2. **Applied security feature** - Whether Azure Defender is enabled for the resource.
3. **Counts of outstanding recommendations and alerts** - The number of outstanding security recommendations and Azure Defender alerts.
4. **Actionable recommendations and alerts** - Two tabs list the recommendations and alerts that apply to the resource.

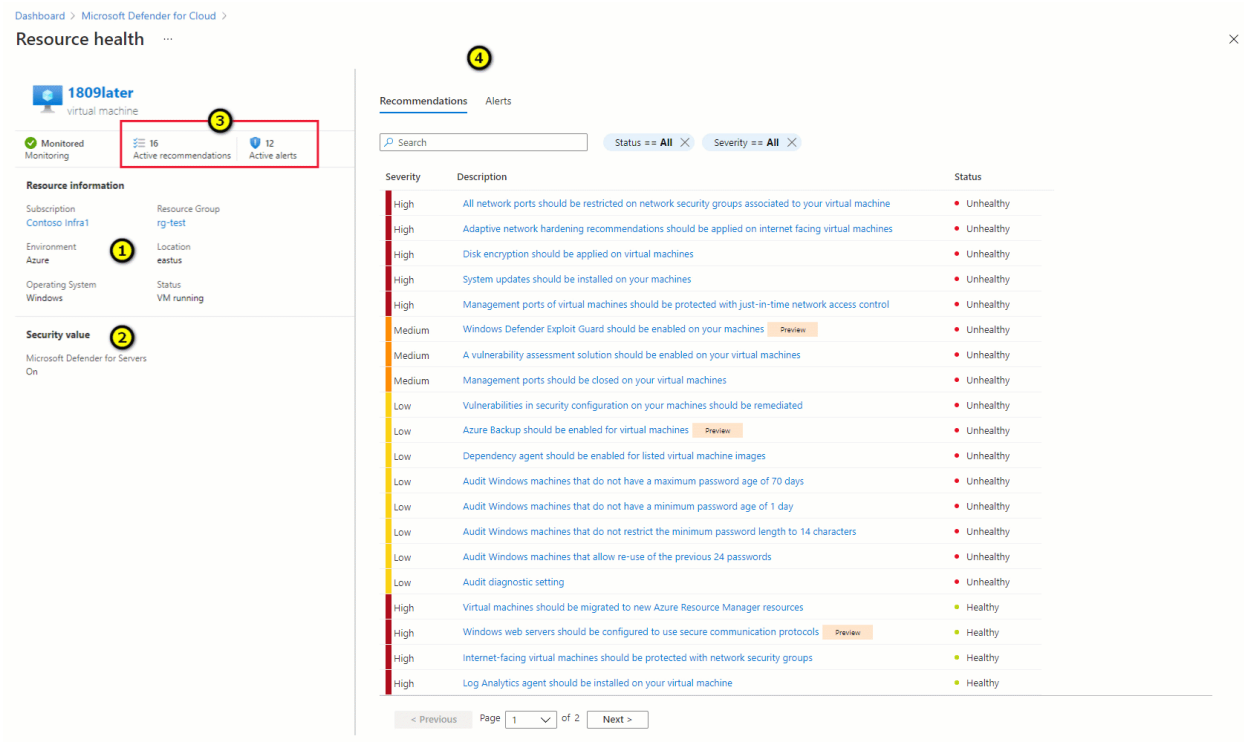

Learn more in Tutorial: Investigate the health of your resources.

Container registry images that have been recently pulled are now rescanned weekly (released for general availability (GA))

Azure Defender for container registries includes a built-in vulnerability scanner. This scanner immediately scans any image you push to your registry and any image pulled within the last 30 days.

New vulnerabilities are discovered every day. With this update, container images that were pulled from your registries during the last 30 days will be **rescanned** every week. This ensures that newly discovered vulnerabilities are identified in your images.

Scanning is charged on a per image basis, so there's no additional charge for these rescans.

Learn more about this scanner in Use Azure Defender for container registries to scan your images for vulnerabilities.

Use Azure Defender for Kubernetes to protect hybrid and multi-cloud Kubernetes deployments (in preview)

Azure Defender for Kubernetes is expanding its threat protection capabilities to defend your clusters wherever they're deployed. This has been enabled by integrating with Azure Arc-enabled Kubernetes and its new extensions capabilities.

When you've enabled Azure Arc on your non-Azure Kubernetes clusters, a new recommendation from Azure Security Center offers to deploy the Azure Defender extension to them with only a few clicks.

Use the recommendation (**Azure Arc-enabled Kubernetes clusters should have Azure Defender's extension installed**) and the extension to protect Kubernetes clusters deployed in other cloud providers, although not on their managed Kubernetes services.

This integration between Azure Security Center, Azure Defender, and Azure Arc-enabled Kubernetes brings:

- Easy provisioning of the Azure Defender extension to unprotected Azure Arc-enabled Kubernetes clusters (manually and at-scale)
- Monitoring of the Azure Defender extension and its provisioning state from the Azure Arc Portal
- Security recommendations from Security Center are reported in the new Security page of the Azure Arc Portal
- Identified security threats from Azure Defender are reported in the new Security page of the Azure Arc Portal
- Azure Arc-enabled Kubernetes clusters are integrated into the Azure Security Center platform and experience

Learn more in Use Azure Defender for Kubernetes with your on-premises and multi-cloud Kubernetes clusters.

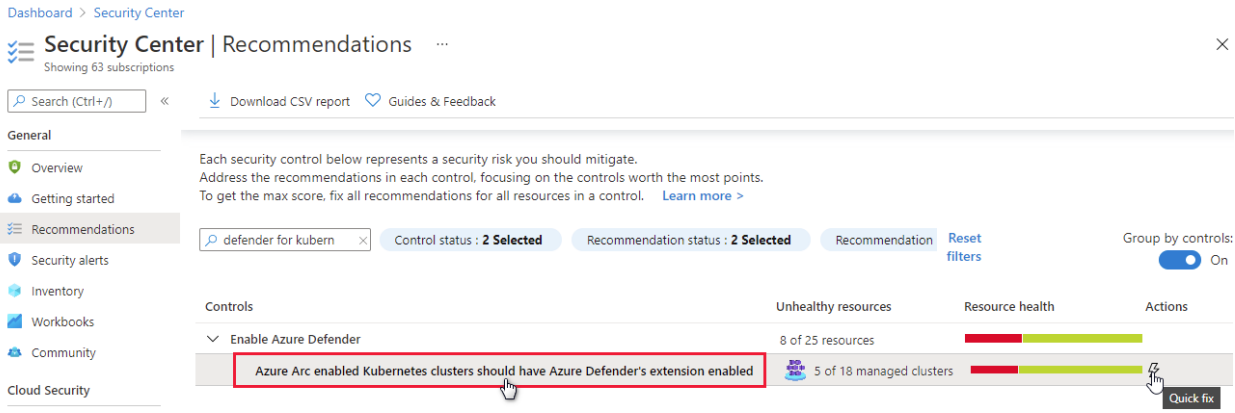

Microsoft Defender for Endpoint integration with Azure Defender now supports Windows Server 2019 and Windows 10 Virtual Desktop (WVD) released for general availability (GA)

Microsoft Defender for Endpoint is a holistic, cloud delivered endpoint security solution. It provides risk-based vulnerability management and assessment as well as endpoint detection and response (EDR). For a full list of the benefits of using Defender for Endpoint together with Azure Security Center, see Protect your endpoints with Security Center's integrated EDR solution: Microsoft Defender for Endpoint.

When you enable Azure Defender for servers on a Windows server, a license for Defender for Endpoint is included with the plan. If you've already enabled Azure Defender for servers and you have Windows 2019 servers in your subscription, they'll automatically receive Defender for Endpoint with this update. No manual action is required.

Support has now been expanded to include Windows Server 2019 and Windows Virtual Desktop (WVD).

> **NOTE**
>
> If you're enabling Defender for Endpoint on a Windows Server 2019 machine, ensure it meets the prerequisites described in Enable the Microsoft Defender for Endpoint integration.

Recommendations to enable Azure Defender for DNS and Resource Manager (in preview)

Two new recommendations have been added to simplify the process of enabling Azure Defender for Resource Manager and Azure Defender for DNS:

- **Azure Defender for Resource Manager should be enabled** - Defender for Resource Manager automatically monitors the resource management operations in your organization. Azure Defender detects threats and alerts you about suspicious activity.
- **Azure Defender for DNS should be enabled** - Defender for DNS provides an additional layer of protection for your cloud resources by continuously monitoring all DNS queries from your Azure resources. Azure Defender alerts you about suspicious activity at the DNS layer.

Enabling Azure Defender plans results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center.

> **TIP**
> Preview recommendations don't render a resource unhealthy, and they aren't included in the calculations of your secure score. Remediate them wherever possible, so that when the preview period ends they'll contribute towards your score. Learn more about how to respond to these recommendations in Remediate recommendations in Azure Security Center.

Three regulatory compliance standards added: Azure CIS 1.3.0, CMMC Level 3, and New Zealand ISM Restricted

We've added three standards for use with Azure Security Center. Using the regulatory compliance dashboard, you can now track your compliance with:

- CIS Microsoft Azure Foundations Benchmark 1.3.0
- CMMC Level 3
- New Zealand ISM Restricted

You can assign these to your subscriptions as described in Customize the set of standards in your regulatory compliance dashboard.

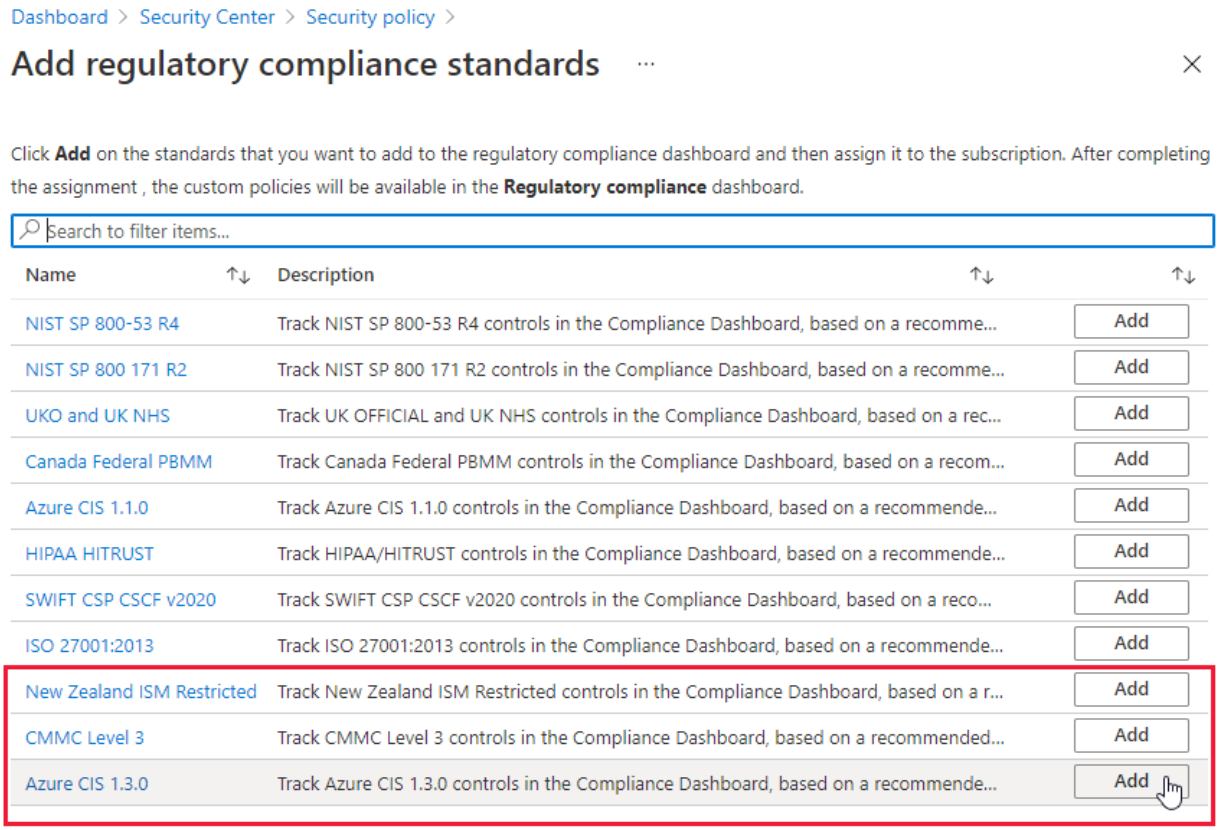

Learn more in:

- Customize the set of standards in your regulatory compliance dashboard
- Tutorial: Improve your regulatory compliance
- FAQ - Regulatory compliance dashboard

Four new recommendations related to guest configuration (in preview)

Azure's Guest Configuration extension reports to Security Center to help ensure your virtual machines' in-guest settings are hardened. The extension isn't required for Arc-enabled servers because it's included in the Arc Connected Machine agent. The extension requires a system-managed identity on the machine.

We've added four new recommendations to Security Center to make the most of this extension.

- Two recommendations prompt you to install the extension and its required system-managed identity:
  - **Guest Configuration extension should be installed on your machines**
  - **Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity**
- When the extension is installed and running, it will begin auditing your machines and you'll be prompted to harden settings such as configuration of the operating system and environment settings. These two recommendations will prompt you to harden your Windows and Linux machines as described:
  - **Windows Defender Exploit Guard should be enabled on your machines**
  - **Authentication to Linux machines should require SSH keys**

Learn more in Understand Azure Policy's Guest Configuration.

CMK recommendations moved to best practices security control

Every organization's security program includes data encryption requirements. By default, Azure customers' data is encrypted at rest with service-managed keys. However, customer-managed keys (CMK) are commonly required to meet regulatory compliance standards. CMKs let you encrypt your data with an Azure Key Vault key created and owned by you. This gives you full control and responsibility for the key lifecycle, including rotation and management.

Azure Security Center's security controls are logical groups of related security recommendations, and reflect your vulnerable attack surfaces. Each control has a maximum number of points you can add to your secure score if you remediate all of the recommendations listed in the control, for all of your resources. The **Implement security best practices** security control is worth zero points. So recommendations in this control don't affect your secure score.

The recommendations listed below are being moved to the **Implement security best practices** security control to better reflect their optional nature. This move ensures that these recommendations are in the most appropriate control to meet their objective.

- Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest
- Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK)
- Cognitive Services accounts should enable data encryption with a customer-managed key (CMK)
- Container registries should be encrypted with a customer-managed key (CMK)
- SQL managed instances should use customer-managed keys to encrypt data at rest
- SQL servers should use customer-managed keys to encrypt data at rest
- Storage accounts should use customer-managed key (CMK) for encryption

Learn which recommendations are in each security control in Security controls and their recommendations.

11 Azure Defender alerts deprecated

The 11 Azure Defender alerts listed below have been deprecated.

- New alerts will replace these two alerts and provide better coverage:

ALERTTYPE	ALERTDISPLAYNAME
ARM_MicroBurstDomainInfo	PREVIEW - MicroBurst toolkit "Get-AzureDomainInfo" function run detected
ARM_MicroBurstRunbook	PREVIEW - MicroBurst toolkit "Get-AzurePasswords" function run detected

- These nine alerts relate to an Azure Active Directory Identity Protection connector (IPC) that has already been deprecated:

ALERTTYPE	ALERTDISPLAYNAME
UnfamiliarLocation	Unfamiliar sign-in properties
AnonymousLogin	Anonymous IP address
InfectedDeviceLogin	Malware linked IP address
ImpossibleTravel	Atypical travel
MaliciousIP	Malicious IP address
LeakedCredentials	Leaked credentials
PasswordSpray	Password Spray
LeakedCredentials	Azure AD threat intelligence
AADAI	Azure AD AI

> **TIP**
> These nine IPC alerts were never Security Center alerts. They're part of the Azure Active Directory (AAD) Identity Protection connector (IPC) that was sending them to Security Center. For the last two years, the only customers who've been seeing those alerts are organizations who configured the export (from the connector to ASC) in 2019 or earlier. AAD IPC has continued to show them in its own alerts systems and they've continued to be available in Azure Sentinel. The only change is that they're no longer appearing in Security Center.

Two recommendations from "Apply system updates" security control were deprecated

The following two recommendations were deprecated and the changes might result in a slight impact on your secure score:

- **Your machines should be restarted to apply system updates**
- **Monitoring agent should be installed on your machines**. This recommendation relates to on-premises machines only and some of its logic will be transferred to another recommendation, **Log Analytics agent health issues should be resolved on your machines**

We recommend checking your continuous export and workflow automation configurations to see whether these recommendations are included in them. Also, any dashboards or other monitoring tools that might be using them should be updated accordingly.

Learn more about these recommendations in the security recommendations reference page.

Azure Defender for SQL on machine tile removed from Azure Defender dashboard

The Azure Defender dashboard's coverage area includes tiles for the relevant Azure Defender plans for your environment. Due to an issue with the reporting of the numbers of protected and unprotected resources, we've decided to temporarily remove the resource coverage status for **Azure Defender for SQL on machines** until the issue is resolved.

21 recommendations moved between security controls

The following recommendations were moved to different security controls. Security controls are logical groups of related security recommendations, and reflect your vulnerable attack surfaces. This move ensures that each of these recommendations is in the most appropriate control to meet its objective.

Learn which recommendations are in each security control in Security controls and their recommendations.

RECOMMENDATION	CHANGE AND IMPACT
Vulnerability assessment should be enabled on your SQL servers Vulnerability assessment should be enabled on your SQL managed instances Vulnerabilities on your SQL databases should be remediated new Vulnerabilities on your SQL databases in VMs should be remediated	Moving from Remediate vulnerabilities (worth 6 points) to Remediate security configurations (worth 4 points). Depending on your environment, these recommendations will have a reduced impact on your score.
There should be more than one owner assigned to your subscription Automation account variables should be encrypted IoT Devices - Auditd process stopped sending events IoT Devices - Operating system baseline validation failure IoT Devices - TLS cipher suite upgrade needed IoT Devices - Open Ports On Device IoT Devices - Permissive firewall policy in one of the chains was found IoT Devices - Permissive firewall rule in the input chain was found IoT Devices - Permissive firewall rule in the output chain was found Diagnostic logs in IoT Hub should be enabled IoT Devices - Agent sending underutilized messages IoT Devices - Default IP Filter Policy should be Deny IoT Devices - IP Filter rule large IP range IoT Devices - Agent message intervals and size should be adjusted IoT Devices - Identical Authentication Credentials IoT Devices - Audited process stopped sending events IoT Devices - Operating system (OS) baseline configuration should be fixed	Moving to **Implement security best practices**. When a recommendation moves to the Implement security best practices security control, which is worth no points, the recommendation no longer affects your secure score.

March 2021

Updates in March include:

- Azure Firewall management integrated into Security Center
- SQL vulnerability assessment now includes the "Disable rule" experience (preview)
- Azure Monitor Workbooks integrated into Security Center and three templates provided
- Regulatory compliance dashboard now includes Azure Audit reports (preview)
- Recommendation data can be viewed in Azure Resource Graph with "Explore in ARG"
- Updates to the policies for deploying workflow automation
- Two legacy recommendations no longer write data directly to Azure activity log
- Recommendations page enhancements

Azure Firewall management integrated into Security Center

When you open Azure Security Center, the first page to appear is the overview page.

This interactive dashboard provides a unified view into the security posture of your hybrid cloud workloads. Additionally, it shows security alerts, coverage information, and more.

As part of helping you view your security status from a central experience, we have integrated the Azure Firewall Manager into this dashboard. You can now check Firewall coverage status across all networks and centrally manage Azure Firewall policies starting from Security Center.

Learn more about this dashboard in Azure Security Center's overview page.

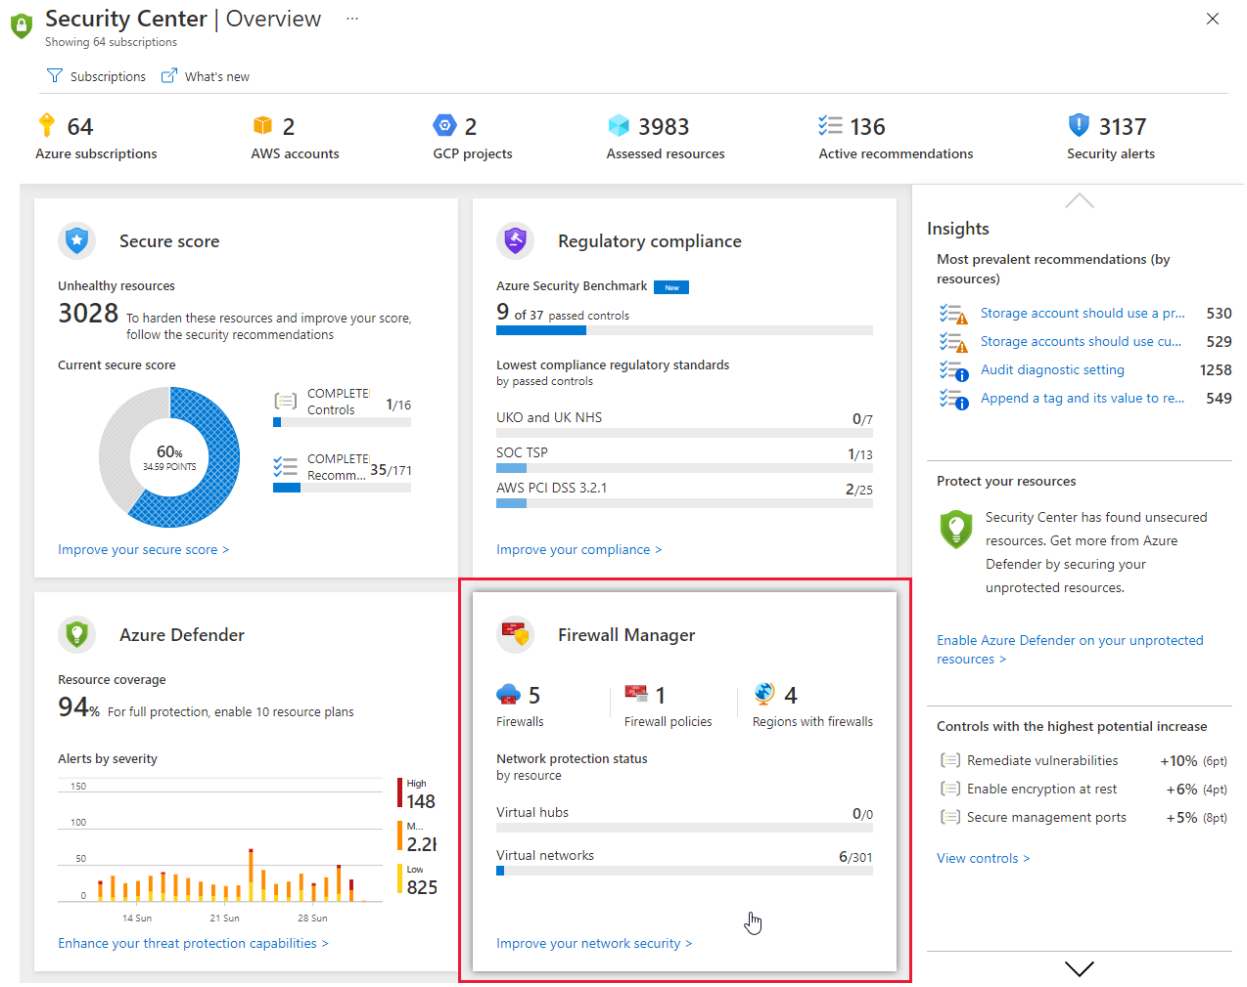

SQL vulnerability assessment now includes the "Disable rule" experience (preview)

Security Center includes a built-in vulnerability scanner to help you discover, track, and remediate potential database vulnerabilities. The results from your assessment scans provide an overview of your SQL machines' security state, and details of any security findings.

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

Learn more in Disable specific findings.

Azure Monitor Workbooks integrated into Security Center and three templates provided

As part of Ignite Spring 2021, we announced an integrated Azure Monitor Workbooks experience in Security Center.

You can use the new integration to start using the out-of-the-box templates from Security Center's gallery. By using workbook templates, you can access and build dynamic and visual reports to track your organization's security posture. Additionally, you can create new workbooks based on Security Center data or any other supported data types and quickly deploy community workbooks from Security Center's GitHub community.

Three templates reports are provided:

- **Secure Score Over Time** - Track your subscriptions' scores and changes to recommendations for your resources
- **System Updates** - View missing system updates by resources, OS, severity, and more
- **Vulnerability Assessment Findings** - View the findings of vulnerability scans of your Azure resources

Learn about using these reports or building your own in Create rich, interactive reports of Security Center data.

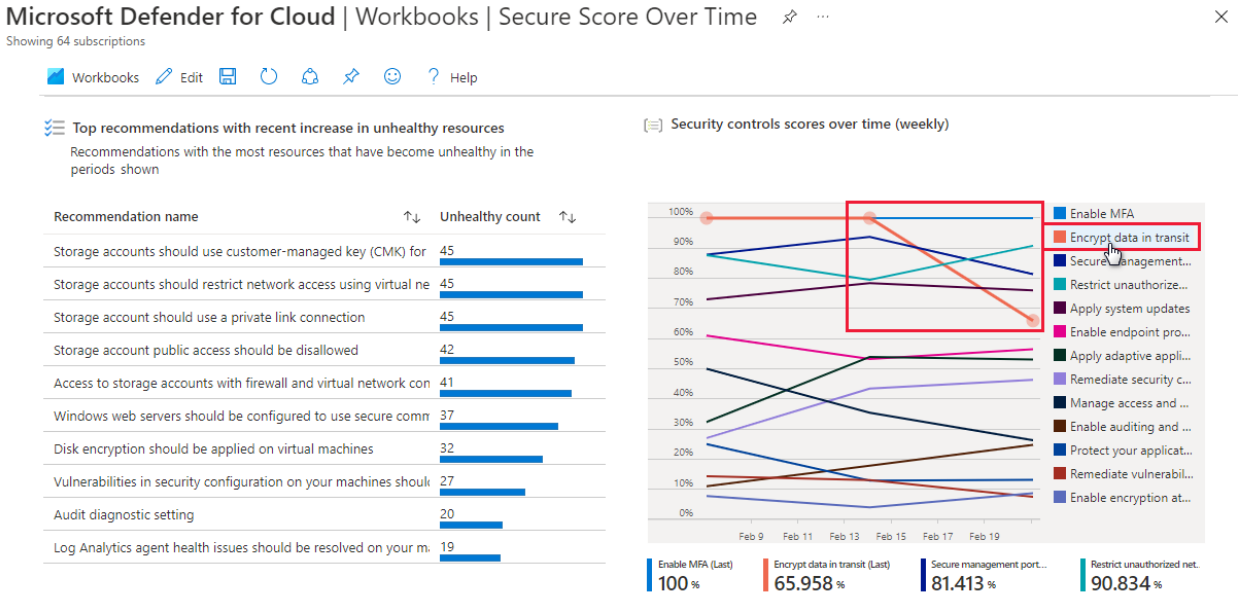

Regulatory compliance dashboard now includes Azure Audit reports (preview)

From the regulatory compliance dashboard's toolbar, you can now download Azure and Dynamics certification reports.

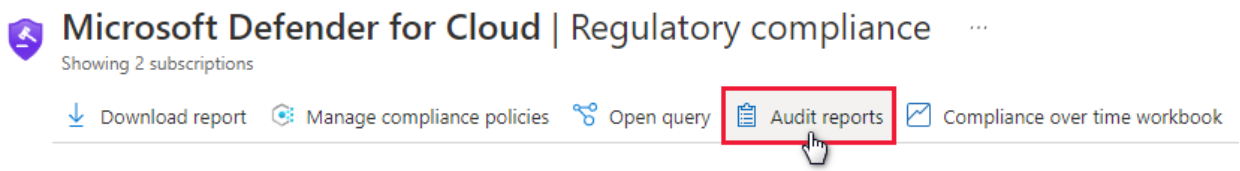

You can select the tab for the relevant reports types (PCI, SOC, ISO, and others) and use filters to find the specific reports you need.

Learn more about Managing the standards in your regulatory compliance dashboard.

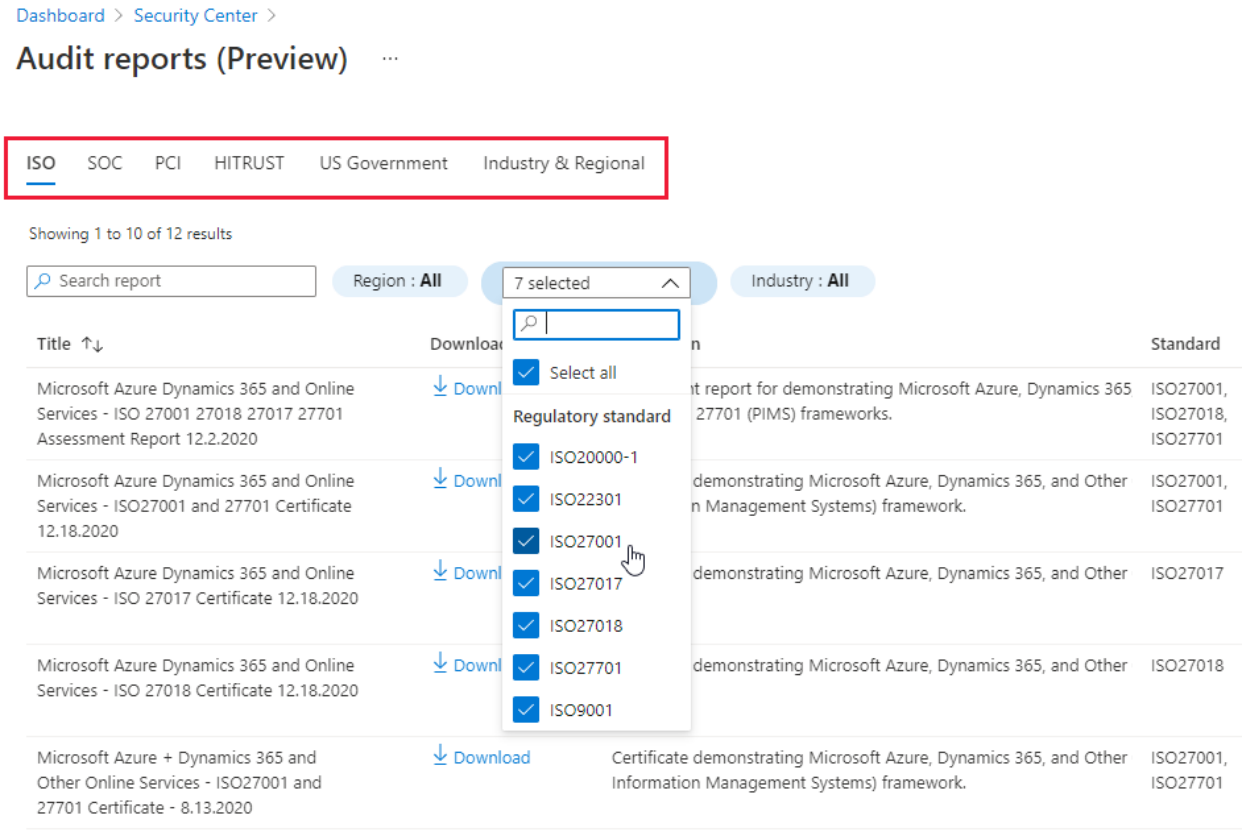

Recommendation data can be viewed in Azure Resource Graph with "Explore in ARG"

The recommendation details pages now include the "Explore in ARG" toolbar button. Use this button to open an Azure Resource Graph query and explore, export, and share the recommendation's data.

Azure Resource Graph (ARG) provides instant access to resource information across your cloud environments with robust filtering, grouping, and sorting capabilities. It's a quick and efficient way to query information across Azure subscriptions programmatically or from within the Azure portal.

Learn more about Azure Resource Graph (ARG).

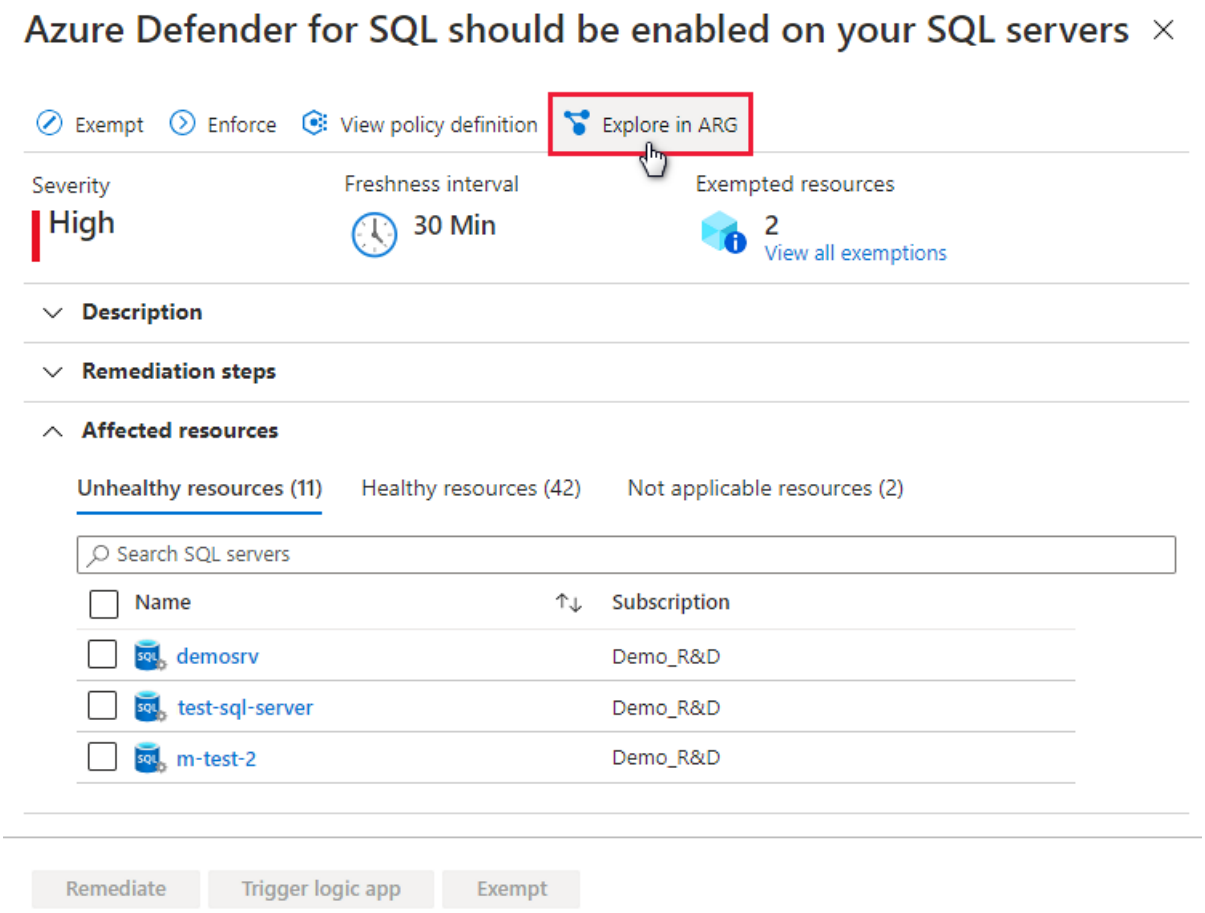

Updates to the policies for deploying workflow automation

Automating your organization's monitoring and incident response processes can greatly improve the time it takes to investigate and mitigate security incidents.

We provide three Azure Policy 'DeployIfNotExist' policies that create and configure workflow automation procedures so that you can deploy your automations across your organization:

GOAL	POLICY	POLICY ID
Workflow automation for security alerts	Deploy Workflow Automation for Azure Security Center alerts	f1525828-9a90-4fcf-be48-268cdd02361e
Workflow automation for security recommendations	Deploy Workflow Automation for Azure Security Center recommendations	73d6ab6c-2475-4850-afd6-43795f3492ef
Workflow automation for regulatory compliance changes	Deploy Workflow Automation for Azure Security Center regulatory compliance	509122b9-ddd9-47ba-a5f1-d0dac20be63c

There are two updates to the features of these policies:

- When assigned, they will remain enabled by enforcement.
- You can now customize these policies and update any of the parameters even after they have already been deployed. For example, if a user wants to add another assessment key, or edit an existing assessment key, they can do so.

Get started with workflow automation templates.

Learn more about how to Automate responses to Security Center triggers.

Two legacy recommendations no longer write data directly to Azure activity log

Security Center passes the data for almost all security recommendations to Azure Advisor, which in turn, writes it to Azure activity log.

For two recommendations, the data is simultaneously written directly to Azure activity log. With this change, Security Center stops writing data for these legacy security recommendations directly to activity Log. Instead, we're exporting the data to Azure Advisor as we do for all the other recommendations.

The two legacy recommendations are:

- Endpoint protection health issues should be resolved on your machines
- Vulnerabilities in security configuration on your machines should be remediated

If you've been accessing information for these two recommendations in activity log's "Recommendation of type TaskDiscovery" category, this is no longer available.

Recommendations page enhancements

We've released an improved version of the recommendations list to present more information at a glance.

Now on the page you'll see:

1. The maximum score and current score for each security control.
2. Icons replacing tags such as **Fix** and **Preview**.
3. A new column showing the Policy initiative related to each recommendation - visible when "Group by controls" is disabled.

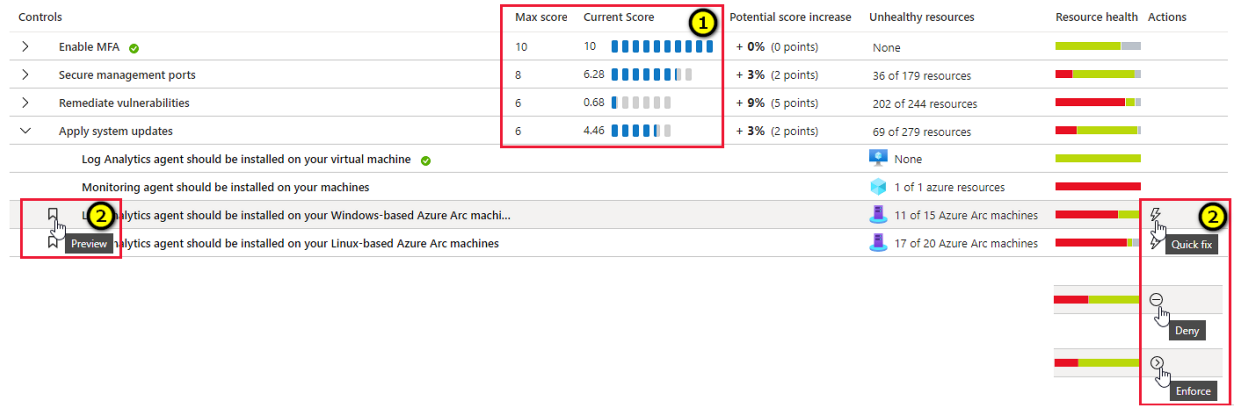

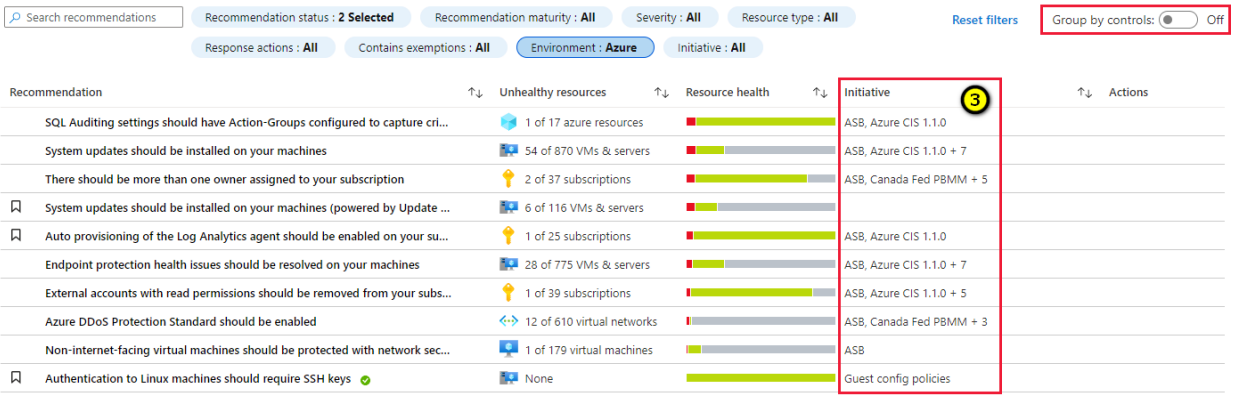

Learn more in Security recommendations in Azure Security Center.

February 2021

Updates in February include:

- New security alerts page in the Azure portal released for general availability (GA)
- Kubernetes workload protection recommendations released for general availability (GA)
- Microsoft Defender for Endpoint integration with Azure Defender now supports Windows Server 2019 and Windows 10 Virtual Desktop (WVD) (in preview)
- Direct link to policy from recommendation details page
- SQL data classification recommendation no longer affects your secure score
- Workflow automations can be triggered by changes to regulatory compliance assessments (in preview)
- Asset inventory page enhancements

New security alerts page in the Azure portal released for general availability (GA)

Azure Security Center's security alerts page has been redesigned to provide:

- **Improved triage experience for alerts** - helping to reduce alerts fatigue and focus on the most relevant threats easier, the list includes customizable filters and grouping options.
- **More information in the alerts list** - such as MITRE ATT&ACK tactics.
- **Button to create sample alerts** - to evaluate Azure Defender capabilities and test your alerts. configuration (for SIEM integration, email notifications, and workflow automations), you can create sample alerts from all Azure Defender plans.
- **Alignment with Azure Sentinel's incident experience** - for customers who use both products, switching between them is now a more straightforward experience and it's easy to learn one from the other.
- **Better performance** for large alerts lists.
- **Keyboard navigation** through the alert list.
- **Alerts from Azure Resource Graph** - you can query alerts in Azure Resource Graph, the Kusto-like API for all of your resources. This is also useful if you're building your own alerts dashboards. Learn more about Azure Resource Graph.
- **Create sample alerts feature** - To create sample alerts from the new alerts experience, see Generate sample Azure Defender alerts.

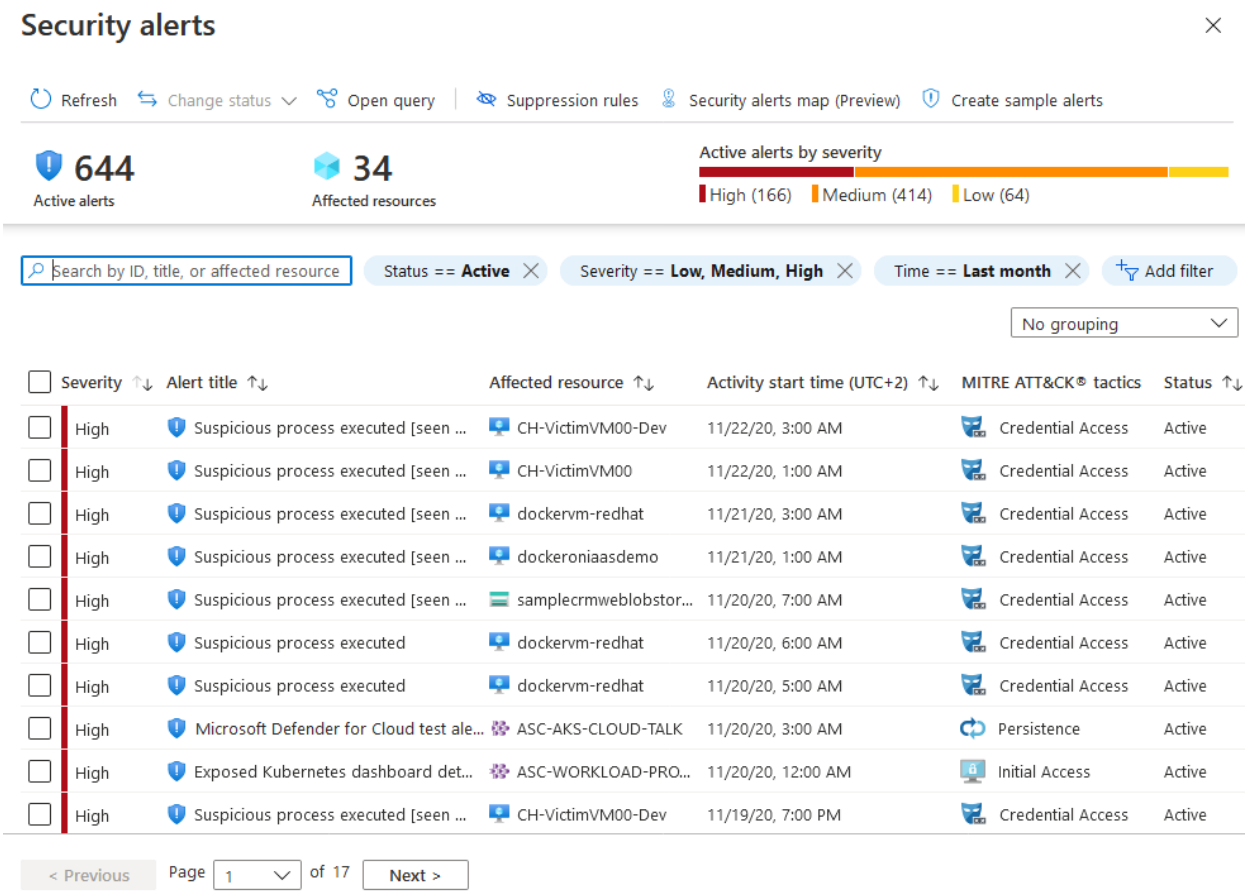

Kubernetes workload protection recommendations released for general availability (GA)

We're happy to announce the general availability (GA) of the set of recommendations for Kubernetes workload protections.

To ensure that Kubernetes workloads are secure by default, Security Center has added Kubernetes level hardening recommendations, including enforcement options with Kubernetes admission control.

When the Azure Policy add-on for Kubernetes is installed on your Azure Kubernetes Service (AKS) cluster, every request to the Kubernetes API server will be monitored against the predefined set of best practices - displayed as 13 security recommendations - before being persisted to the cluster. You can then configure to enforce the best practices and mandate them for future workloads.

For example, you can mandate that privileged containers shouldn't be created, and any future requests to do so will be blocked.

Learn more in Workload protection best-practices using Kubernetes admission control.

> **NOTE**
>
> While the recommendations were in preview, they didn't render an AKS cluster resource unhealthy, and they weren't included in the calculations of your secure score. with this GA announcement these will be included in the score calculation. If you haven't remediated them already, this might result in a slight impact on your secure score. Remediate them wherever possible as described in Remediate recommendations in Azure Security Center.

Microsoft Defender for Endpoint integration with Azure Defender now supports Windows Server 2019 and Windows 10 Virtual Desktop (WVD) (in preview)

Microsoft Defender for Endpoint is a holistic, cloud delivered endpoint security solution. It provides risk-based vulnerability management and assessment as well as endpoint detection and response (EDR). For a full list of the benefits of using Defender for Endpoint together with Azure Security Center, see Protect your endpoints with Security Center's integrated EDR solution: Microsoft Defender for Endpoint.

When you enable Azure Defender for servers on a Windows server, a license for Defender for Endpoint is included with the plan. If you've already enabled Azure Defender for servers and you have Windows 2019 servers in your subscription, they'll automatically receive Defender for Endpoint with this update. No manual action is required.

Support has now been expanded to include Windows Server 2019 and Windows Virtual Desktop (WVD).

> **NOTE**
>
> If you're enabling Defender for Endpoint on a Windows Server 2019 machine, ensure it meets the prerequisites described in Enable the Microsoft Defender for Endpoint integration.

Direct link to policy from recommendation details page

When you're reviewing the details of a recommendation, it's often helpful to be able to see the underlying policy. For every recommendation supported by a policy, there's a new link from the recommendation details page:

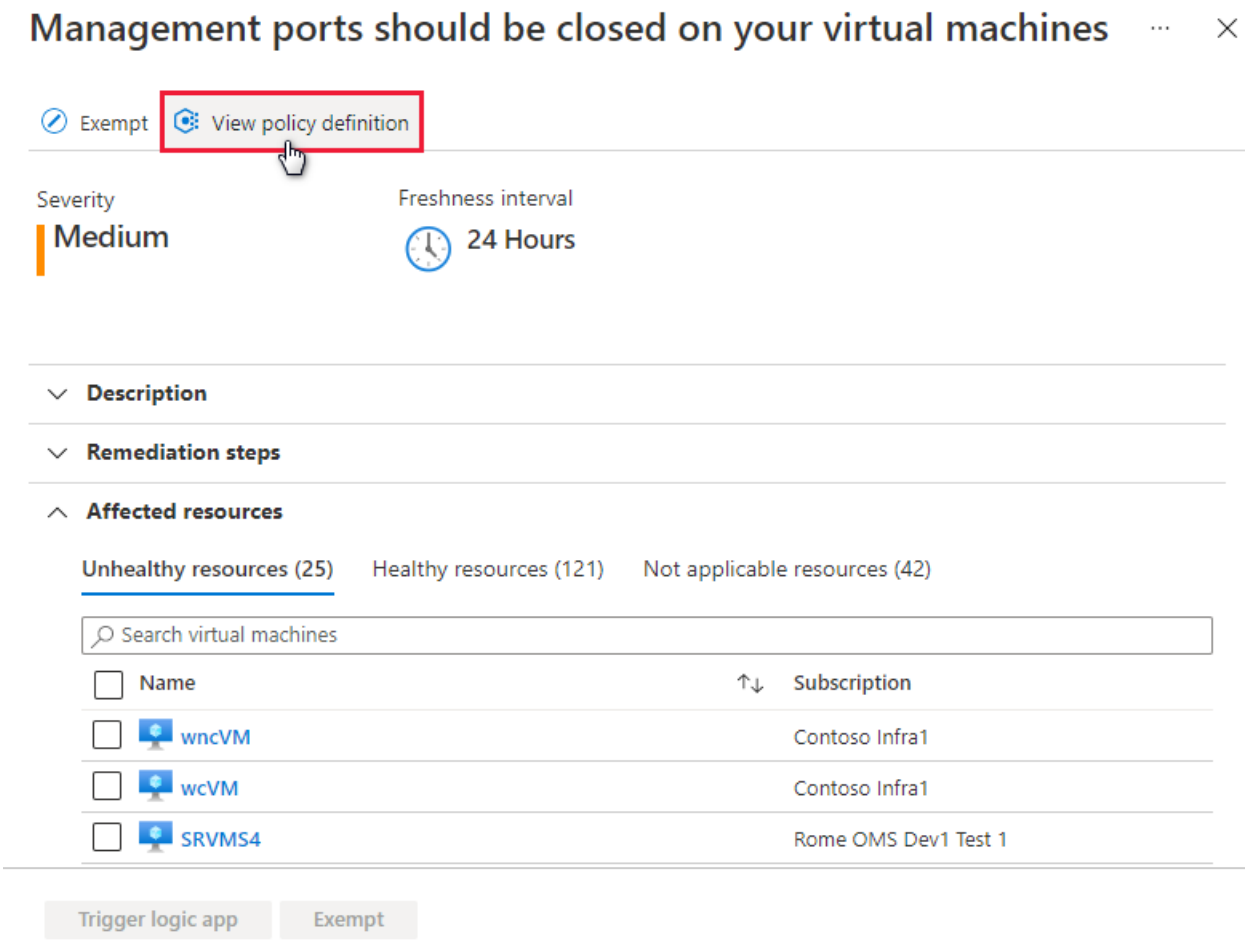

Use this link to view the policy definition and review the evaluation logic.

If you're reviewing the list of recommendations on our Security recommendations reference guide, you'll also see links to the policy definition pages:

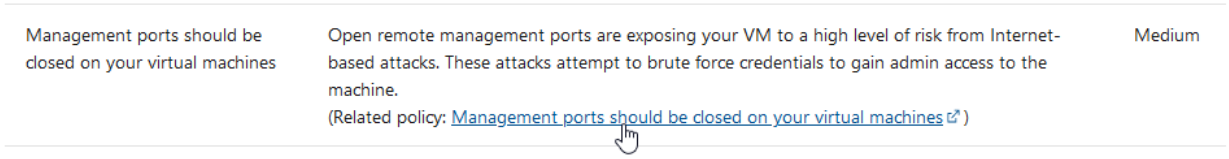

SQL data classification recommendation no longer affects your secure score

The recommendation **Sensitive data in your SQL databases should be classified** no longer affects your secure score. This is the only recommendation in the **Apply data classification** security control, so that control now has a secure score value of 0.

For a full list of all security controls in Security Center, together with their scores and a list of the recommendations in each, see Security controls and their recommendations.

Workflow automations can be triggered by changes to regulatory compliance assessments (in preview)

We've added a third data type to the trigger options for your workflow automations: changes to regulatory compliance assessments.

Learn how to use the workflow automation tools in Automate responses to Security Center triggers.

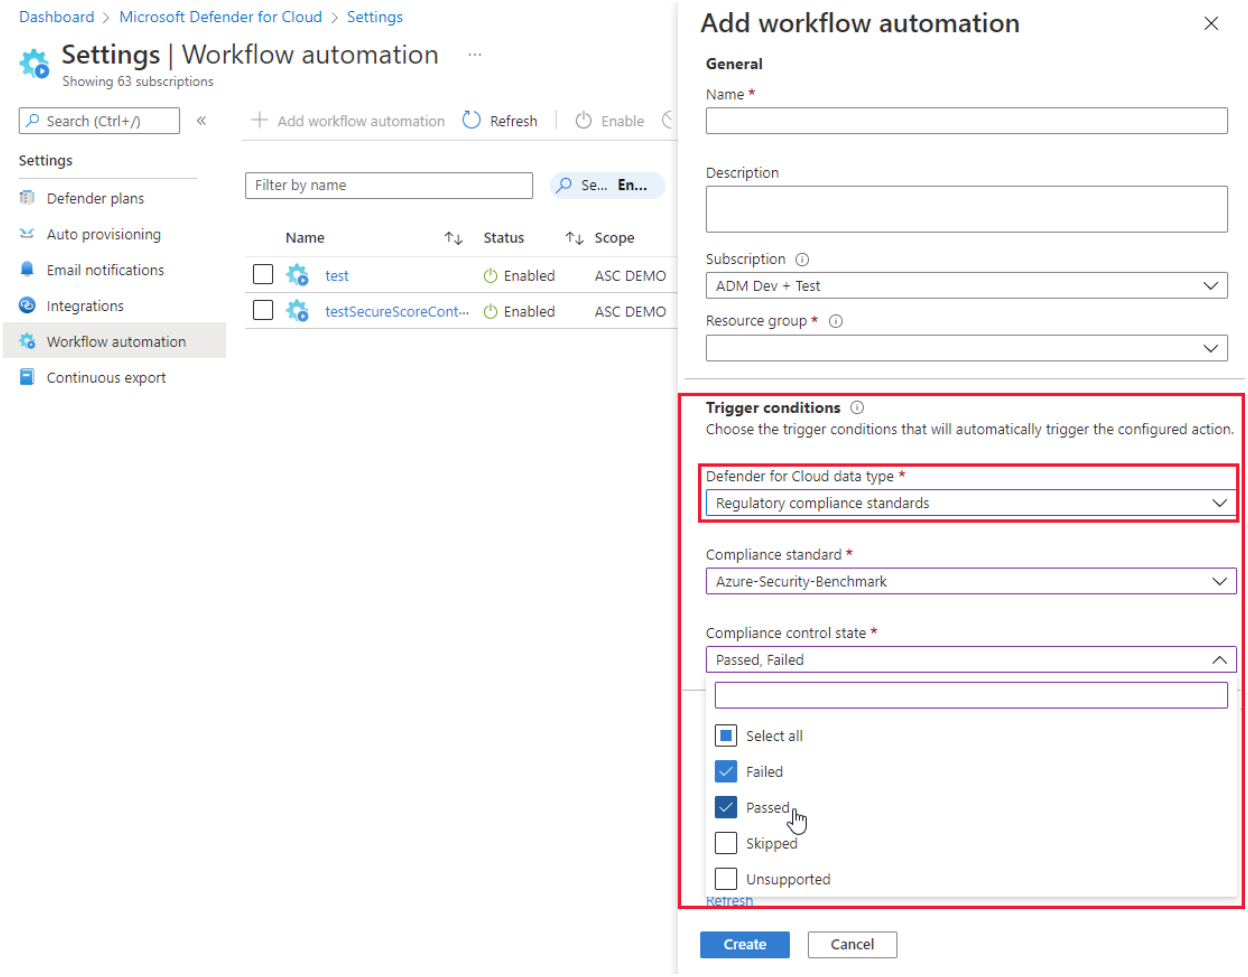

Asset inventory page enhancements

Security Center's asset inventory page has been improved in the following ways:

- Summaries at the top of the page now include **Unregistered subscriptions**, showing the number of subscriptions without Security Center enabled.

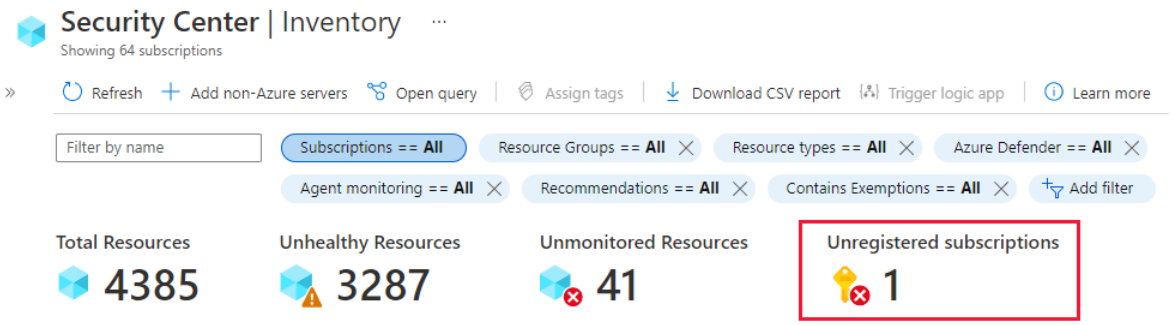

- Filters have been expanded and enhanced to include:
  - **Counts** - Each filter presents the number of resources that meet the criteria of each category

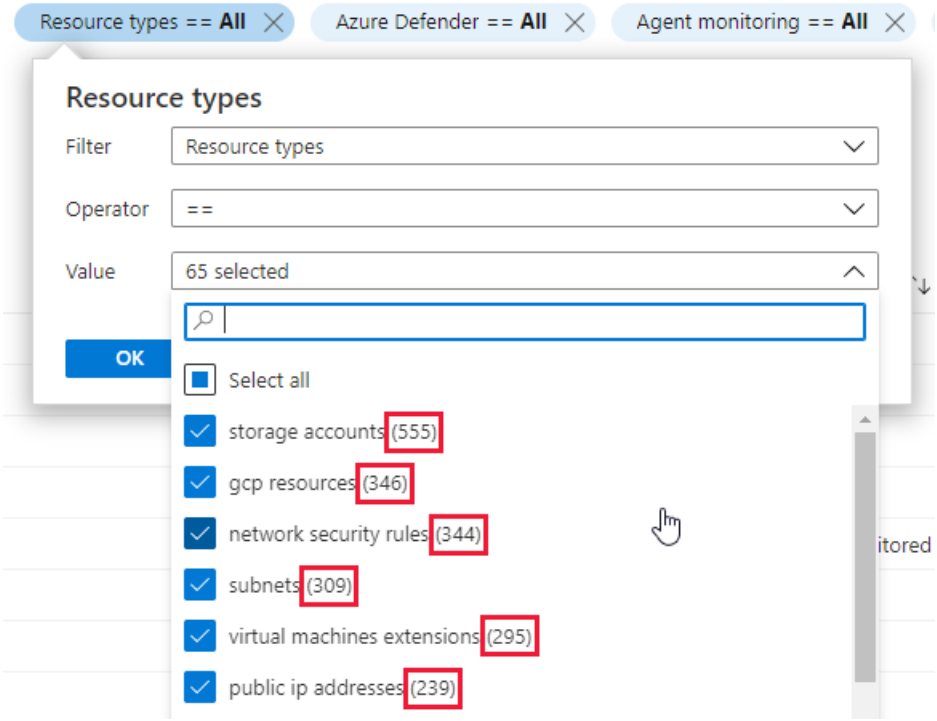

- **Contains exemptions filter** (Optional) - narrow the results to resources that have/haven't got exemptions. This filter isn't shown by default, but is accessible from the **Add filter** button.

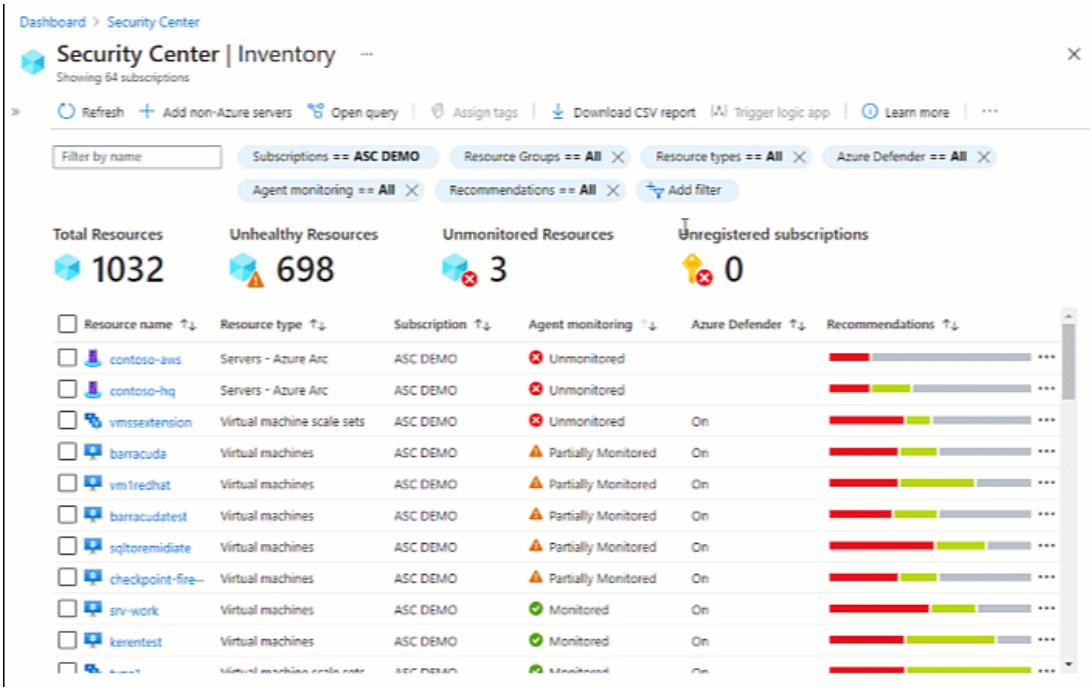

Learn more about how to Explore and manage your resources with asset inventory.

January 2021

Updates in January include:

- Azure Security Benchmark is now the default policy initiative for Azure Security Center
- Vulnerability assessment for on-premise and multi-cloud machines is released for general availability (GA)
- Secure score for management groups is now available in preview
- Secure score API is released for general availability (GA)
- Dangling DNS protections added to Azure Defender for App Service
- Multi-cloud connectors are released for general availability (GA)

- Exempt entire recommendations from your secure score for subscriptions and management groups
- Users can now request tenant-wide visibility from their global administrator
- 35 preview recommendations added to increase coverage of Azure Security Benchmark
- CSV export of filtered list of recommendations
- "Not applicable" resources now reported as "Compliant" in Azure Policy assessments
- Export weekly snapshots of secure score and regulatory compliance data with continuous export (preview)

Azure Security Benchmark is now the default policy initiative for Azure Security Center

Azure Security Benchmark is the Microsoft-authored, Azure-specific set of guidelines for security and compliance best practices based on common compliance frameworks. This widely respected benchmark builds on the controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security.

In recent months, Security Center's list of built-in security recommendations has grown significantly to expand our coverage of this benchmark.

From this release, the benchmark is the foundation for Security Center's recommendations and fully integrated as the default policy initiative.

All Azure services have a security baseline page in their documentation. These baselines are built on Azure Security Benchmark.

If you're using Security Center's regulatory compliance dashboard, you'll see two instances of the benchmark during a transition period:

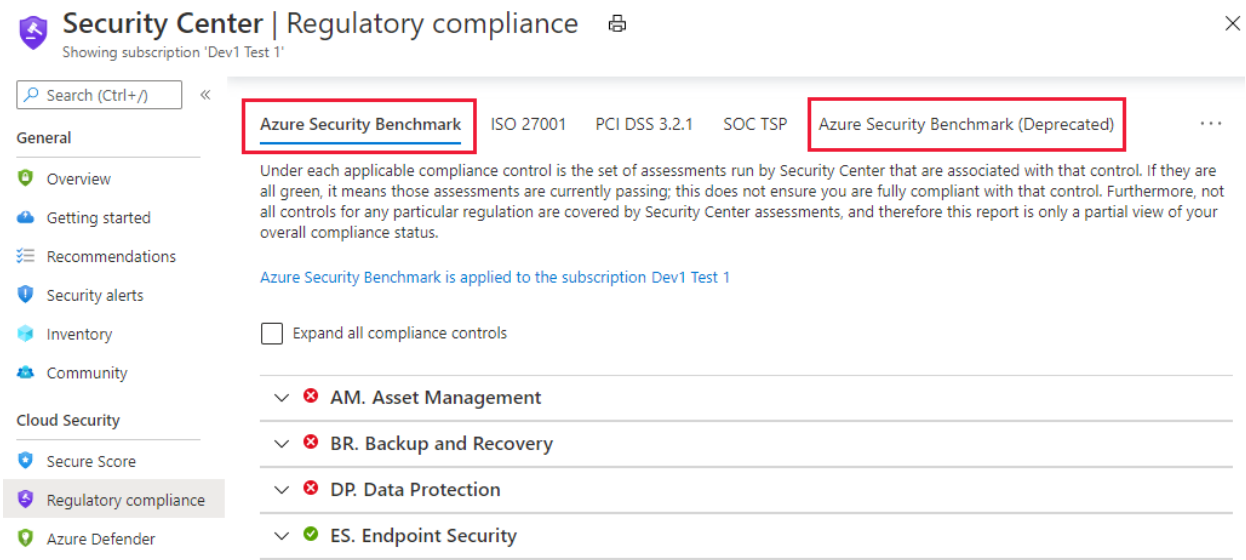

Existing recommendations are unaffected and as the benchmark grows, changes will automatically be reflected within Security Center.

To learn more, see the following pages:

- Learn more about Azure Security Benchmark
- Customize the set of standards in your regulatory compliance dashboard

Vulnerability assessment for on-premise and multi-cloud machines is released for general availability (GA)

In October, we announced a preview for scanning Azure Arc-enabled servers with Azure Defender for servers' integrated vulnerability assessment scanner (powered by Qualys).

It's now released for general availability (GA).

When you've enabled Azure Arc on your non-Azure machines, Security Center will offer to deploy the integrated vulnerability scanner on them - manually and at-scale.

With this update, you can unleash the power of **Azure Defender for servers** to consolidate your vulnerability management program across all of your Azure and non-Azure assets.

Main capabilities:

- Monitoring the VA (vulnerability assessment) scanner provisioning state on Azure Arc machines
- Provisioning the integrated VA agent to unprotected Windows and Linux Azure Arc machines (manually and at-scale)
- Receiving and analyzing detected vulnerabilities from deployed agents (manually and at-scale)
- Unified experience for Azure VMs and Azure Arc machines

Learn more about deploying the integrated Qualys vulnerability scanner to your hybrid machines.

Learn more about Azure Arc-enabled servers.

Secure score for management groups is now available in preview

The secure score page now shows the aggregated secure scores for your management groups in addition to the subscription level. So now you can see the list of management groups in your organization and the score for each management group.

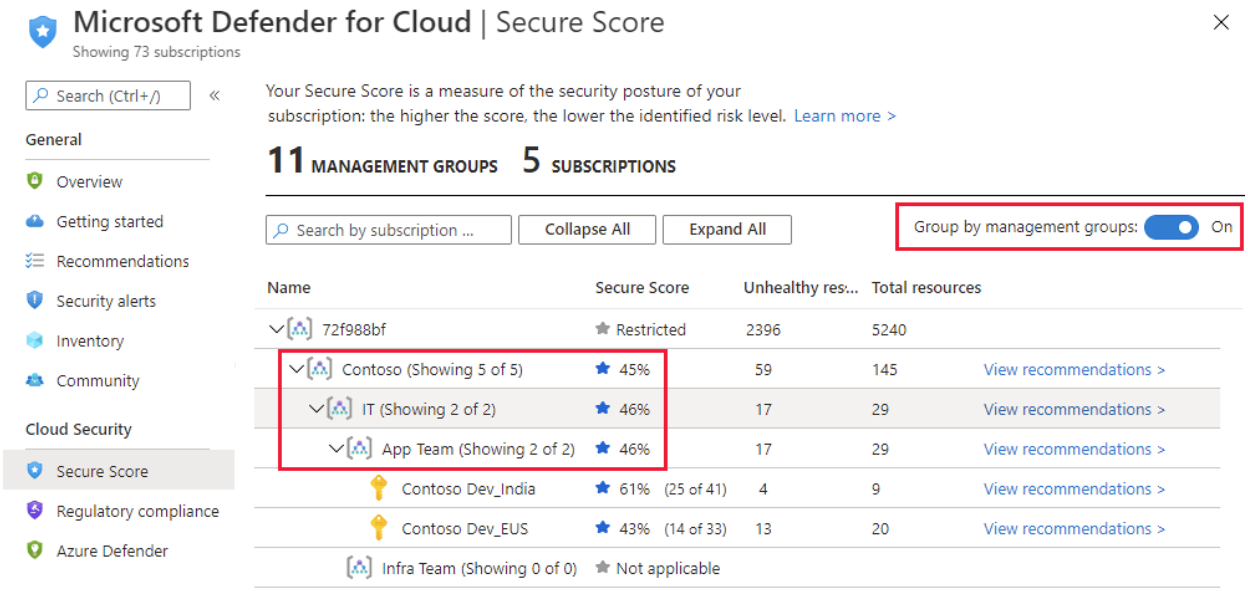

Learn more about secure score and security controls in Azure Security Center.

Secure score API is released for general availability (GA)

You can now access your score via the secure score API. The API methods provide the flexibility to query the data and build your own reporting mechanism of your secure scores over time. For example:

- use the **Secure Scores** API to get the score for a specific subscription
- use the **Secure Score Controls** API to list the security controls and the current score of your subscriptions

Learn about external tools made possible with the secure score API in the secure score area of our GitHub community.

Learn more about secure score and security controls in Azure Security Center.

Dangling DNS protections added to Azure Defender for App Service

Subdomain takeovers are a common, high-severity threat for organizations. A subdomain takeover can occur when you have a DNS record that points to a deprovisioned web site. Such DNS records are also known as "dangling DNS" entries. CNAME records are especially vulnerable to this threat.

Subdomain takeovers enable threat actors to redirect traffic intended for an organization's domain to a site

performing malicious activity.

Azure Defender for App Service now detects dangling DNS entries when an App Service website is decommissioned. This is the moment at which the DNS entry is pointing at a non-existent resource, and your website is vulnerable to a subdomain takeover. These protections are available whether your domains are managed with Azure DNS or an external domain registrar and applies to both App Service on Windows and App Service on Linux.

Learn more:

- App Service alert reference table - Includes two new Azure Defender alerts that trigger when a dangling DNS entry is detected
- Prevent dangling DNS entries and avoid subdomain takeover - Learn about the threat of subdomain takeover and the dangling DNS aspect
- Introduction to Azure Defender for App Service

Multi-cloud connectors are released for general availability (GA)

With cloud workloads commonly spanning multiple cloud platforms, cloud security services must do the same.

Azure Security Center protects workloads in Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP).

Connecting your AWS or GCP accounts integrates their native security tools like AWS Security Hub and GCP Security Command Center into Azure Security Center.

This capability means that Security Center provides visibility and protection across all major cloud environments. Some of the benefits of this integration:

- Automatic agent provisioning - Security Center uses Azure Arc to deploy the Log Analytics agent to your AWS instances
- Policy management
- Vulnerability management
- Embedded Endpoint Detection and Response (EDR)
- Detection of security misconfigurations
- A single view showing security recommendations from all cloud providers
- Incorporate all of your resources into Security Center's secure score calculations
- Regulatory compliance assessments of your AWS and GCP resources

From Defender for Cloud's menu, select **Multi-cloud connectors** and you'll see the options for creating new connectors:

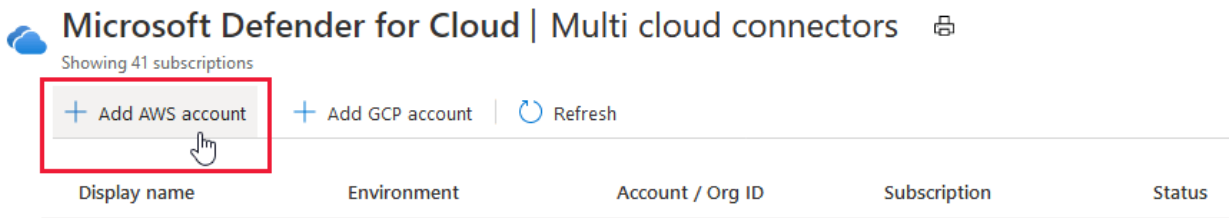

Learn more in:

- Connect your AWS accounts to Azure Security Center
- Connect your GCP accounts to Azure Security Center

Exempt entire recommendations from your secure score for subscriptions and management groups

We're expanding the exemption capability to include entire recommendations. Providing further options to fine-tune the security recommendations that Security Center makes for your subscriptions, management group, or

resources.

Occasionally, a resource will be listed as unhealthy when you know the issue has been resolved by a third-party tool which Security Center hasn't detected. Or a recommendation will show in a scope where you feel it doesn't belong. The recommendation might be inappropriate for a specific subscription. Or perhaps your organization has decided to accept the risks related to the specific resource or recommendation.

With this preview feature, you can now create an exemption for a recommendation to:

- **Exempt a resource** to ensure it isn't listed with the unhealthy resources in the future, and doesn't impact your secure score. The resource will be listed as not applicable and the reason will be shown as "exempted" with the specific justification you select.
- **Exempt a subscription or management group** to ensure that the recommendation doesn't impact your secure score and won't be shown for the subscription or management group in the future. This relates to existing resources and any you create in the future. The recommendation will be marked with the specific justification you select for the scope that you selected.

Learn more in Exempting resources and recommendations from your secure score.

Users can now request tenant-wide visibility from their global administrator

If a user doesn't have permissions to see Security Center data, they'll now see a link to request permissions from their organization's global administrator. The request includes the role they'd like and the justification for why it's necessary.

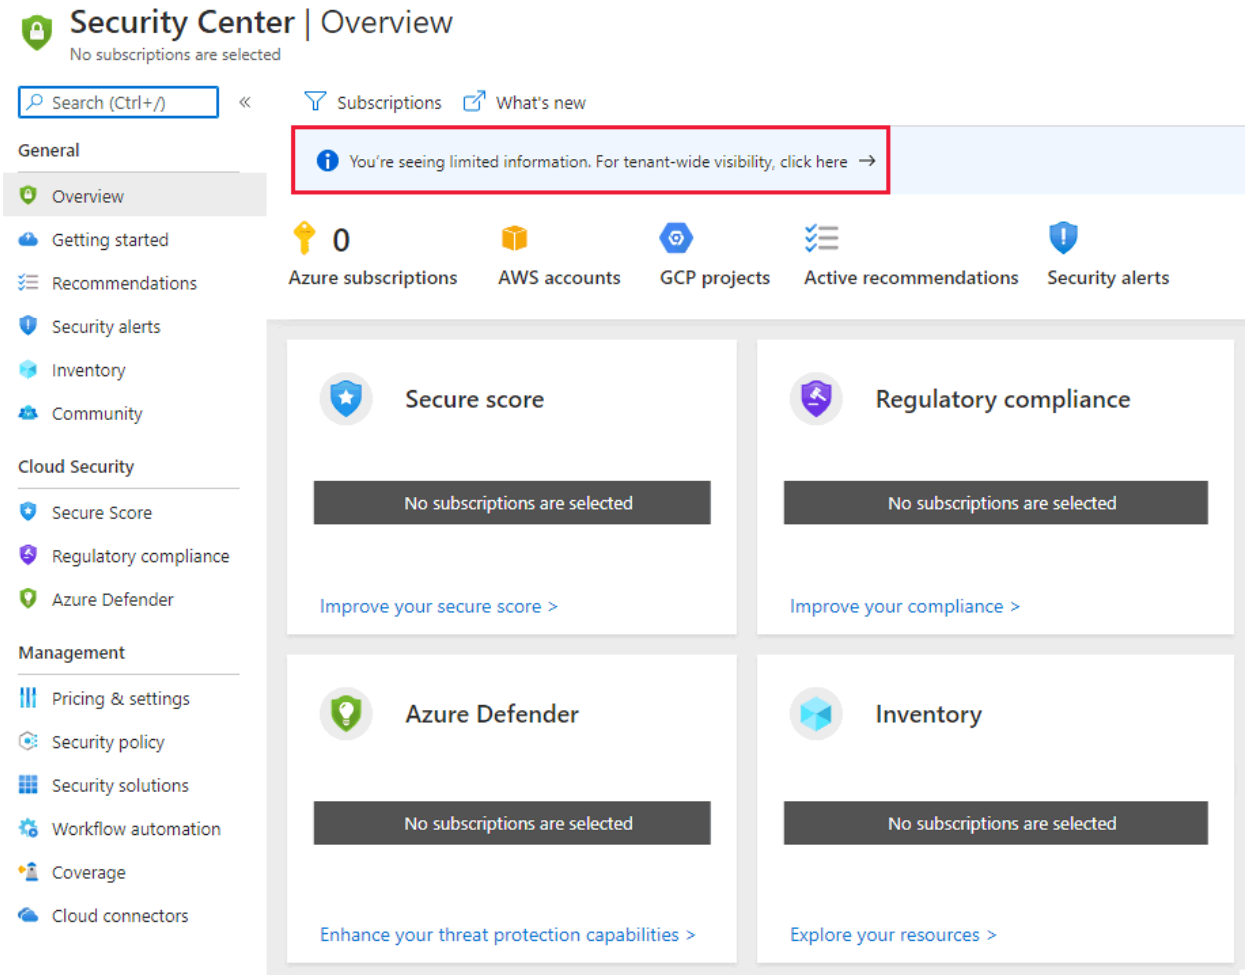

Learn more in Request tenant-wide permissions when yours are insufficient.

35 preview recommendations added to increase coverage of Azure Security Benchmark

Azure Security Benchmark is the default policy initiative in Azure Security Center.

To increase the coverage of this benchmark, the following 35 preview recommendations have been added to Security Center.

> **TIP**
> Preview recommendations don't render a resource unhealthy, and they aren't included in the calculations of your secure score. Remediate them wherever possible, so that when the preview period ends they'll contribute towards your score. Learn more about how to respond to these recommendations in Remediate recommendations in Azure Security Center.

SECURITY CONTROL	NEW RECOMMENDATIONS
Enable encryption at rest	- Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest - Azure Machine Learning workspaces should be encrypted with a customer-managed key (CMK) - Bring your own key data protection should be enabled for MySQL servers - Bring your own key data protection should be enabled for PostgreSQL servers - Cognitive Services accounts should enable data encryption with a customer-managed key(CMK) - Container registries should be encrypted with a customer-managed key (CMK) - SQL managed instances should use customer-managed keys to encrypt data at rest - SQL servers should use customer-managed keys to encrypt data at rest - Storage accounts should use customer-managed key (CMK) for encryption
Implement security best practices	- Subscriptions should have a contact email address for security issues - Auto provisioning of the Log Analytics agent should be enabled on your subscription - Email notification for high severity alerts should be enabled - Email notification to subscription owner for high severity alerts should be enabled - Key vaults should have purge protection enabled - Key vaults should have soft delete enabled
Manage access and permissions	- Function apps should have 'Client Certificates (Incoming client certificates)' enabled
Protect applications against DDoS attacks	- Web Application Firewall (WAF) should be enabled for Application Gateway - Web Application Firewall (WAF) should be enabled for Azure Front Door Serviceservice

SECURITY CONTROL	NEW RECOMMENDATIONS
Restrict unauthorized network access	- Firewall should be enabled on Key Vault - Private endpoint should be configured for Key Vault - App Configuration should use private link - Azure Cache for Redis should reside within a virtual network - Azure Event Grid domains should use private link - Azure Event Grid topics should use private link - Azure Machine Learning workspaces should use private link - Azure SignalR Service should use private link - Azure Spring Cloud should use network injection - Container registries should not allow unrestricted network access - Container registries should use private link - Public network access should be disabled for MariaDB servers - Public network access should be disabled for MySQL servers - Public network access should be disabled for PostgreSQL servers - Storage account should use a private link connection - Storage accounts should restrict network access using virtual network rules - VM Image Builder templates should use private link

Related links:

- Learn more about Azure Security Benchmark
- Learn more about Azure Database for MariaDB
- Learn more about Azure Database for MySQL
- Learn more about Azure Database for PostgreSQL

CSV export of filtered list of recommendations

In November 2020, we added filters to the recommendations page (Recommendations list now includes filters). In December, we expanded those filters (Recommendations page has new filters for environment, severity, and available responses).

With this announcement, we're changing the behavior of the **Download to CSV** button so that the CSV export only includes the recommendations currently displayed in the filtered list.

For example, in the image below you can see that the list has been filtered to two recommendations. The CSV file that is generated includes the status details for every resource affected by those two recommendations.

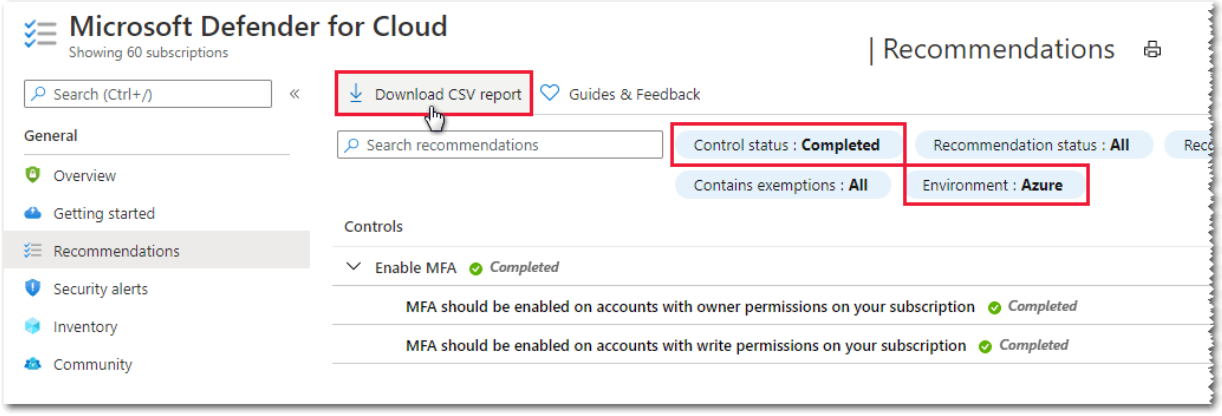

Learn more in Security recommendations in Azure Security Center.

"Not applicable" resources now reported as "Compliant" in Azure Policy assessments

Previously, resources that were evaluated for a recommendation and found to be **not applicable** appeared in Azure Policy as "Non-compliant". No user actions could change their state to "Compliant". With this change, they're reported as "Compliant" for improved clarity.

The only impact will be seen in Azure Policy where the number of compliant resources will increase. There will be no impact to your secure score in Azure Security Center.

Export weekly snapshots of secure score and regulatory compliance data with continuous export (preview)

We've added a new preview feature to the continuous export tools for exporting weekly snapshots of secure score and regulatory compliance data.

When you define a continuous export, set the export frequency:

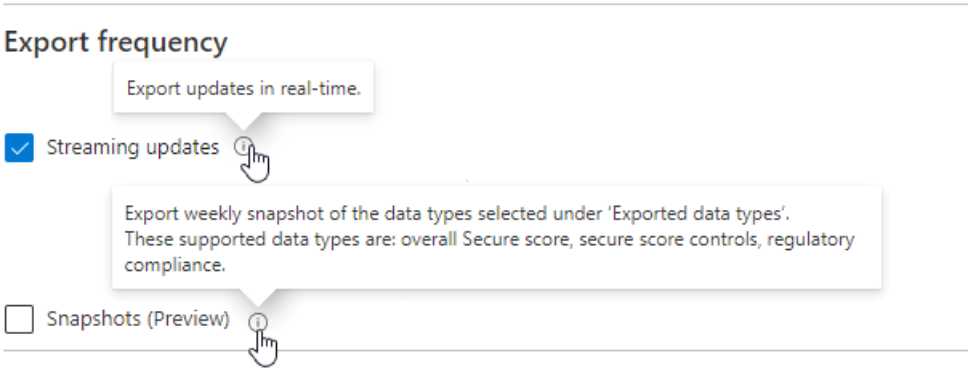

- **Streaming** – assessments will be sent when a resource's health state is updated (if no updates occur, no data will be sent).
- **Snapshots** – a snapshot of the current state of all regulatory compliance assessments will be sent every week (this is a preview feature for weekly snapshots of secure scores and regulatory compliance data).

Learn more about the full capabilities of this feature in Continuously export Security Center data.

December 2020

Updates in December include:

- Azure Defender for SQL servers on machines is generally available
- Azure Defender for SQL support for Azure Synapse Analytics dedicated SQL pool is generally available
- Global Administrators can now grant themselves tenant-level permissions
- Two new Azure Defender plans: Azure Defender for DNS and Azure Defender for Resource Manager (in preview)
- New security alerts page in the Azure portal (preview)
- Revitalized Security Center experience in Azure SQL Database & SQL Managed Instance
- Asset inventory tools and filters updated
- Recommendation about web apps requesting SSL certificates no longer part of secure score
- Recommendations page has new filters for environment, severity, and available responses
- Continuous export gets new data types and improved deployifnotexist policies

Azure Defender for SQL servers on machines is generally available

Azure Security Center offers two Azure Defender plans for SQL Servers:

- **Azure Defender for Azure SQL database servers** - defends your Azure-native SQL Servers
- **Azure Defender for SQL servers on machines** - extends the same protections to your SQL servers in

hybrid, multi-cloud, and on-premises environments

With this announcement, **Azure Defender for SQL** now protects your databases and their data wherever they're located.

Azure Defender for SQL includes vulnerability assessment capabilities. The vulnerability assessment tool includes the following advanced features:

- **Baseline configuration** (New!) to intelligently refine the results of vulnerability scans to those that might represent real security issues. After you've established your baseline security state, the vulnerability assessment tool only reports deviations from that baseline state. Results that match the baseline are considered as passing subsequent scans. This lets you and your analysts focus your attention where it matters.
- **Detailed benchmark information** to help you *understand* the discovered findings, and why they relate to your resources.
- **Remediation scripts** to help you mitigate identified risks.

Learn more about Azure Defender for SQL.

Azure Defender for SQL support for Azure Synapse Analytics dedicated SQL pool is generally available

Azure Synapse Analytics (formerly SQL DW) is an analytics service that combines enterprise data warehousing and big data analytics. Dedicated SQL pools are the enterprise data warehousing features of Azure Synapse. Learn more in What is Azure Synapse Analytics (formerly SQL DW)?.

Azure Defender for SQL protects your dedicated SQL pools with:

- **Advanced threat protection** to detect threats and attacks
- **Vulnerability assessment capabilities** to identify and remediate security misconfigurations

Azure Defender for SQL's support for Azure Synapse Analytics SQL pools is automatically added to Azure SQL databases bundle in Azure Security Center. You'll find a new "Azure Defender for SQL" tab in your Synapse workspace page in the Azure portal.

Learn more about Azure Defender for SQL.

Global Administrators can now grant themselves tenant-level permissions

A user with the Azure Active Directory role of **Global Administrator** might have tenant-wide responsibilities, but lack the Azure permissions to view that organization-wide information in Azure Security Center.

To assign yourself tenant-level permissions, follow the instructions in Grant tenant-wide permissions to yourself.

Two new Azure Defender plans: Azure Defender for DNS and Azure Defender for Resource Manager (in preview)

We've added two new cloud-native breadth threat protection capabilities for your Azure environment.

These new protections greatly enhance your resiliency against attacks from threat actors, and significantly increase the number of Azure resources protected by Azure Defender.

- **Azure Defender for Resource Manager** - automatically monitors all resource management operations performed in your organization. For more information, see:
  - Introduction to Azure Defender for Resource Manager
  - Respond to Azure Defender for Resource Manager alerts
  - List of alerts provided by Azure Defender for Resource Manager
- **Azure Defender for DNS** - continuously monitors all DNS queries from your Azure resources. For more information, see:
  - Introduction to Azure Defender for DNS

- Respond to Azure Defender for DNS alerts
- List of alerts provided by Azure Defender for DNS

New security alerts page in the Azure portal (preview)

Azure Security Center's security alerts page has been redesigned to provide:

- **Improved triage experience for alerts** - helping to reduce alerts fatigue and focus on the most relevant threats easier, the list includes customizable filters and grouping options
- **More information in the alerts list** - such as MITRE ATT&ACK tactics
- **Button to create sample alerts** - to evaluate Azure Defender capabilities and test your alerts configuration (for SIEM integration, email notifications, and workflow automations), you can create sample alerts from all Azure Defender plans
- **Alignment with Azure Sentinel's incident experience** - for customers who use both products, switching between them is now a more straightforward experience and it's easy to learn one from the other
- **Better performance** for large alerts lists
- **Keyboard navigation** through the alert list
- **Alerts from Azure Resource Graph** - you can query alerts in Azure Resource Graph, the Kusto-like API for all of your resources. This is also useful if you're building your own alerts dashboards. Learn more about Azure Resource Graph.

To access the new experience, use the 'try it now' link from the banner at the top of the security alerts page.

To create sample alerts from the new alerts experience, see Generate sample Azure Defender alerts.

Revitalized Security Center experience in Azure SQL Database & SQL Managed Instance

The Security Center experience within SQL provides access to the following Security Center and Azure Defender for SQL features:

- **Security recommendations** – Security Center periodically analyzes the security state of all connected Azure resources to identify potential security misconfigurations. It then provides recommendations on how to remediate those vulnerabilities and improve organizations' security posture.
- **Security alerts** – a detection service that continuously monitors Azure SQL activities for threats such as SQL injection, brute-force attacks, and privilege abuse. This service triggers detailed and action-oriented security alerts in Security Center and provides options for continuing investigations with Azure Sentinel, Microsoft's Azure-native SIEM solution.
- **Findings** – a vulnerability assessment service that continuously monitors Azure SQL configurations and helps remediate vulnerabilities. Assessment scans provide an overview of Azure SQL security states together with detailed security findings.

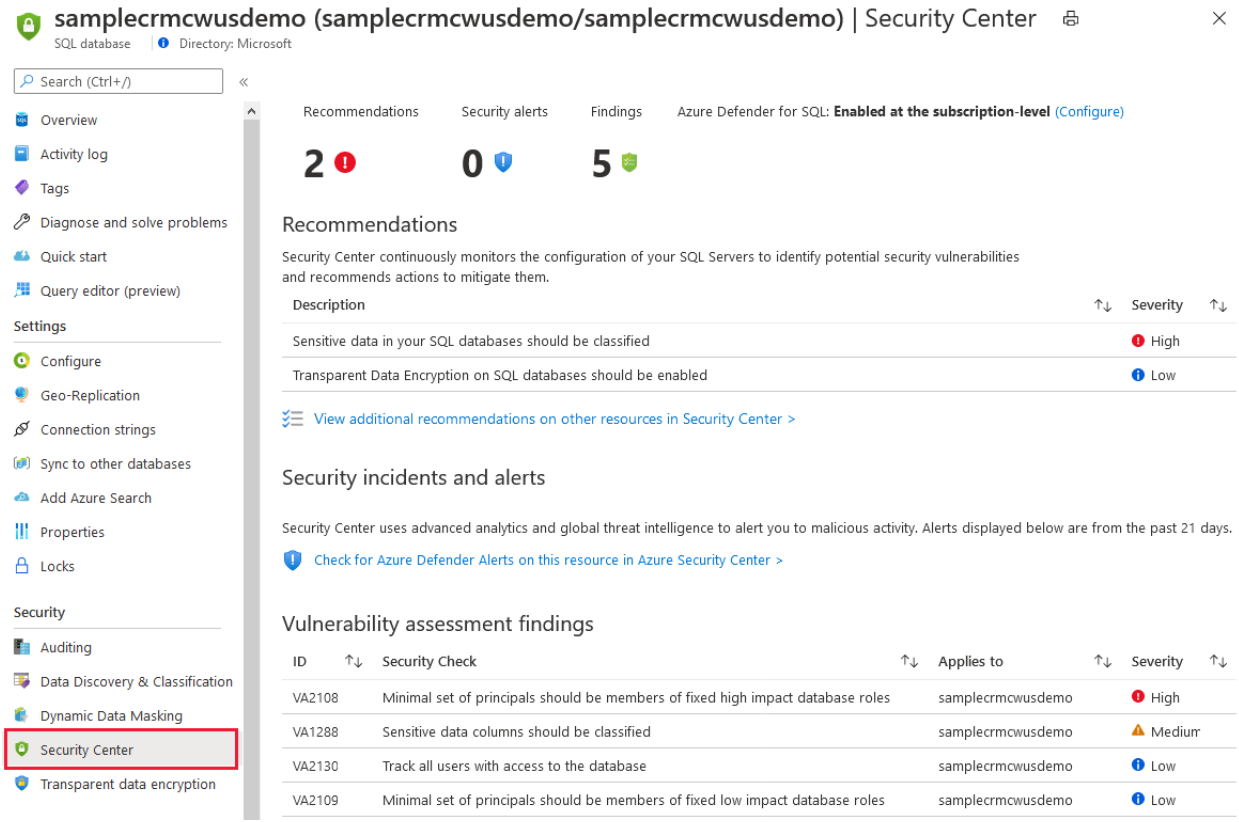

Asset inventory tools and filters updated

The inventory page in Azure Security Center has been refreshed with the following changes:

- **Guides and feedback** added to the toolbar. This opens a pane with links to related information and tools.
- **Subscriptions filter** added to the default filters available for your resources.
- **Open query** link for opening the current filter options as an Azure Resource Graph query (formerly called "View in resource graph explorer").
- **Operator options** for each filter. Now you can choose from more logical operators other than '='. For example, you might want to find all resources with active recommendations whose titles include the string 'encrypt'.

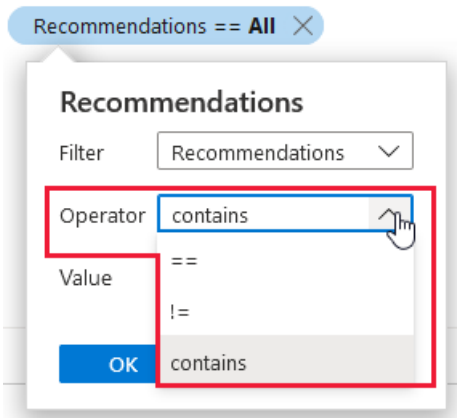

Learn more about inventory in Explore and manage your resources with asset inventory.

Recommendation about web apps requesting SSL certificates no longer part of secure score

The recommendation "Web apps should request an SSL certificate for all incoming requests" has been moved from the security control **Manage access and permissions** (worth a maximum of 4 pts) into **Implement security best practices** (which is worth no points).

Ensuring a web app requests a certificate certainly makes it more secure. However, for public-facing web apps it's irrelevant. If you access your site over HTTP and not HTTPS, you will not receive any client certificate. So if your application requires client certificates, you should not allow requests to your application over HTTP. Learn more in Configure TLS mutual authentication for Azure App Service.

With this change, the recommendation is now a recommended best practice that does not impact your score.

Learn which recommendations are in each security control in Security controls and their recommendations.

Recommendations page has new filters for environment, severity, and available responses

Azure Security Center monitors all connected resources and generates security recommendations. Use these recommendations to strengthen your hybrid cloud posture and track compliance with the policies and standards relevant to your organization, industry, and country.

As Security Center continues to expand its coverage and features, the list of security recommendations is growing every month. For example, see 29 preview recommendations added to increase coverage of Azure Security Benchmark.

With the growing list, there's a need to filter the recommendations to find the ones of greatest interest. In November, we added filters to the recommendations page (see Recommendations list now includes filters).

The filters added this month provide options to refine the recommendations list according to:

- **Environment** - View recommendations for your AWS, GCP, or Azure resources (or any combination)
- **Severity** - View recommendations according to the severity classification set by Security Center
- **Response actions** - View recommendations according to the availability of Security Center response options: Fix, Deny, and Enforce

  TIP

  The response actions filter replaces the **Quick fix available (Yes/No)** filter.

  Learn more about each of these response options:

  - Fix button
  - Prevent misconfigurations with Enforce/Deny recommendations

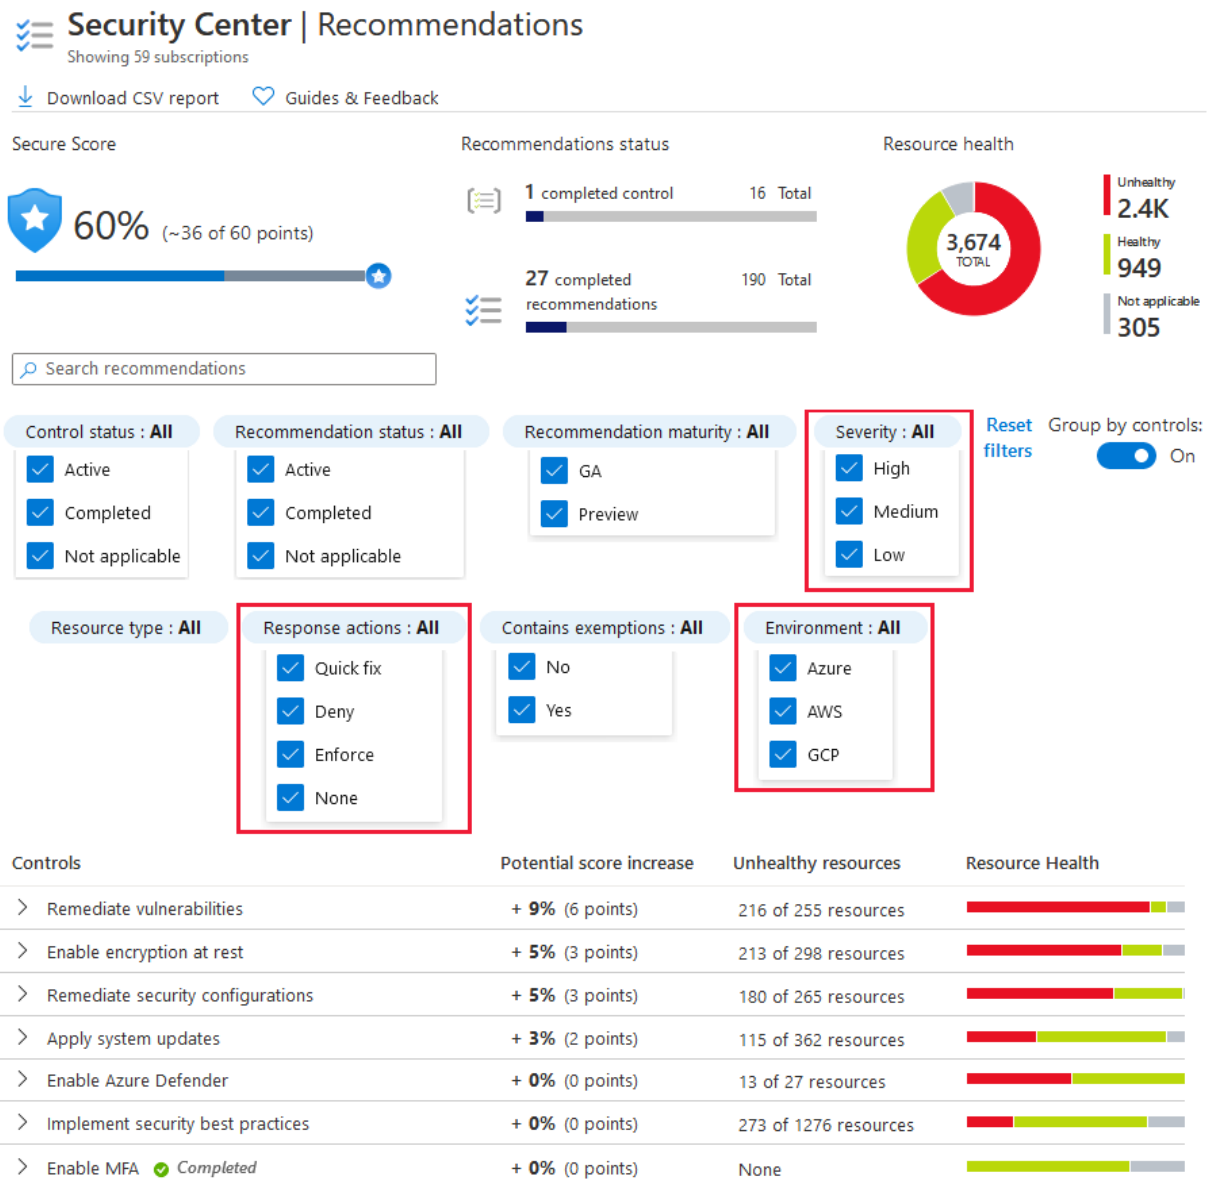

Continuous export gets new data types and improved deployifnotexist policies

Azure Security Center's continuous export tools enable you to export Security Center's recommendations and alerts for use with other monitoring tools in your environment.

Continuous export lets you fully customize what will be exported, and where it will go. For full details, see Continuously export Security Center data.

These tools have been enhanced and expanded in the following ways:

- **Continuous export's deployifnotexist policies enhanced**. The policies now:
  - **Check whether the configuration is enabled.** If it isn't, the policy will show as non-compliant and create a compliant resource. Learn more about the supplied Azure Policy templates in the "Deploy at scale with Azure Policy tab" in Set up a continuous export.
  - **Support exporting security findings.** When using the Azure Policy templates, you can configure your continuous export to include findings. This is relevant when exporting recommendations that have 'sub' recommendations, like findings from vulnerability assessment scanners or specific system updates for the 'parent' recommendation "System updates should be installed on your machines".
  - **Support exporting secure score data.**
- **Regulatory compliance assessment data added (in preview).** You can now continuously export

updates to regulatory compliance assessments, including for any custom initiatives, to a Log Analytics workspace or Event Hub. This feature is unavailable on national clouds.

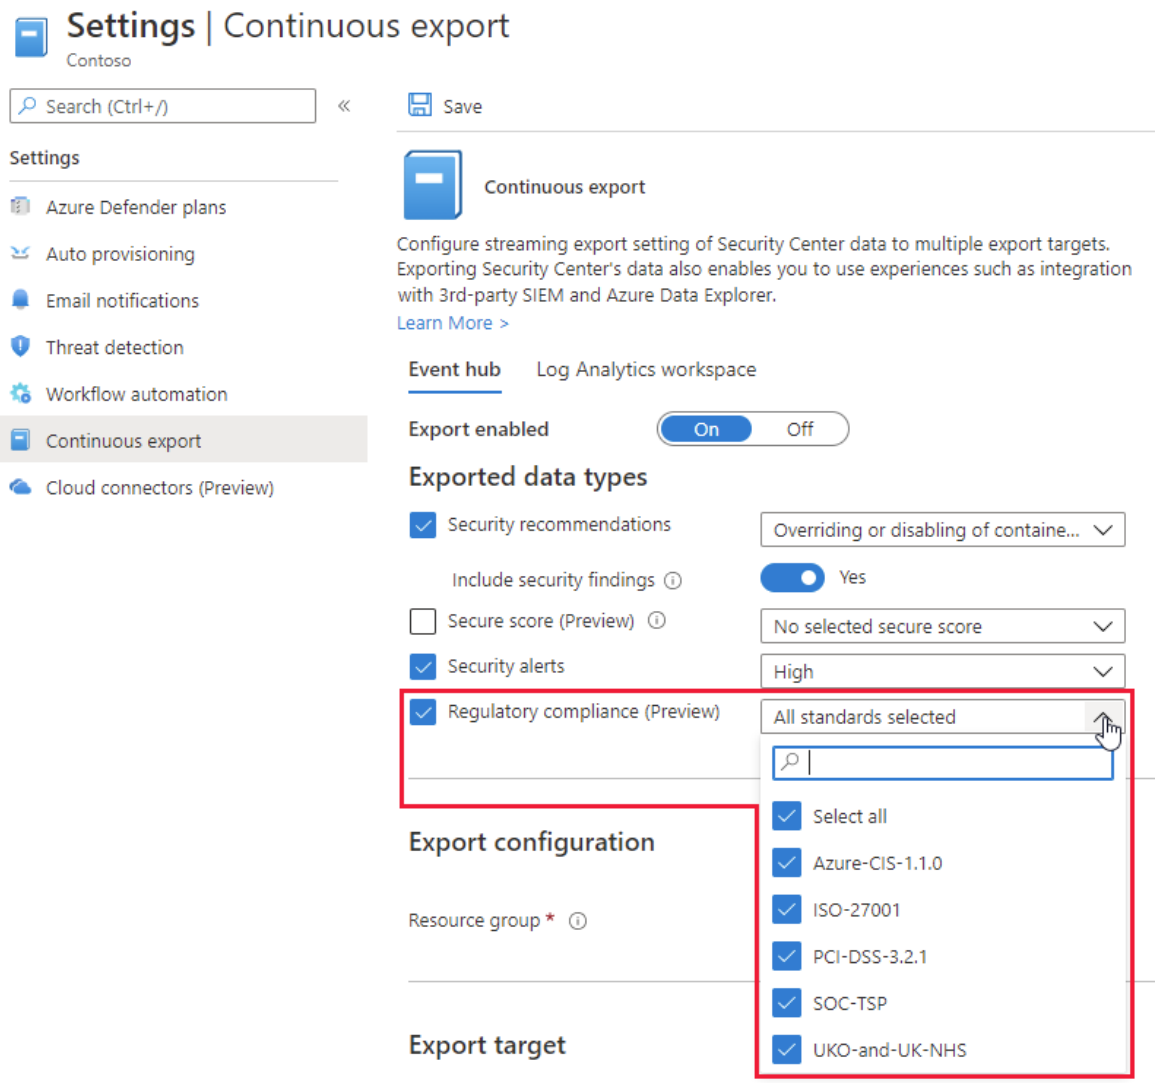

November 2020

Updates in November include:

- 29 preview recommendations added to increase coverage of Azure Security Benchmark
- NIST SP 800 171 R2 added to Security Center's regulatory compliance dashboard
- Recommendations list now includes filters
- Auto provisioning experience improved and expanded
- Secure score is now available in continuous export (preview)
- "System updates should be installed on your machines" recommendation now includes subrecommendations
- Policy management page in the Azure portal now shows status of default policy assignments

29 preview recommendations added to increase coverage of Azure Security Benchmark

Azure Security Benchmark is the Microsoft-authored, Azure-specific, set of guidelines for security and compliance best practices based on common compliance frameworks. Learn more about Azure Security Benchmark.

The following 29 preview recommendations have been added to Security Center to increase the coverage of this benchmark.

Preview recommendations don't render a resource unhealthy, and they aren't included in the calculations of

your secure score. Remediate them wherever possible, so that when the preview period ends they'll contribute towards your score. Learn more about how to respond to these recommendations in Remediate recommendations in Azure Security Center.

SECURITY CONTROL	NEW RECOMMENDATIONS
Encrypt data in transit	- Enforce SSL connection should be enabled for PostgreSQL database servers - Enforce SSL connection should be enabled for MySQL database servers - TLS should be updated to the latest version for your API app - TLS should be updated to the latest version for your function app - TLS should be updated to the latest version for your web app - FTPS should be required in your API App - FTPS should be required in your function App - FTPS should be required in your web App
Manage access and permissions	- Web apps should request an SSL certificate for all incoming requests - Managed identity should be used in your API App - Managed identity should be used in your function App - Managed identity should be used in your web App
Restrict unauthorized network access	- Private endpoint should be enabled for PostgreSQL servers - Private endpoint should be enabled for MariaDB servers - Private endpoint should be enabled for MySQL servers
Enable auditing and logging	- Diagnostic logs in App Services should be enabled
Implement security best practices	- Azure Backup should be enabled for virtual machines - Geo-redundant backup should be enabled for Azure Database for MariaDB - Geo-redundant backup should be enabled for Azure Database for MySQL - Geo-redundant backup should be enabled for Azure Database for PostgreSQL - PHP should be updated to the latest version for your API app - PHP should be updated to the latest version for your web app - Java should be updated to the latest version for your API app - Java should be updated to the latest version for your function app - Java should be updated to the latest version for your web app - Python should be updated to the latest version for your API app - Python should be updated to the latest version for your function app - Python should be updated to the latest version for your web app - Audit retention for SQL servers should be set to at least 90 days

Related links:

- Learn more about Azure Security Benchmark
- Learn more about Azure API apps
- Learn more about Azure function apps
- Learn more about Azure web apps
- Learn more about Azure Database for MariaDB
- Learn more about Azure Database for MySQL
- Learn more about Azure Database for PostgreSQL

NIST SP 800 171 R2 added to Security Center's regulatory compliance dashboard

The NIST SP 800-171 R2 standard is now available as a built-in initiative for use with Azure Security Center's regulatory compliance dashboard. The mappings for the controls are described in Details of the NIST SP 800-171 R2 Regulatory Compliance built-in initiative.

To apply the standard to your subscriptions and continuously monitor your compliance status, use the instructions in Customize the set of standards in your regulatory compliance dashboard.

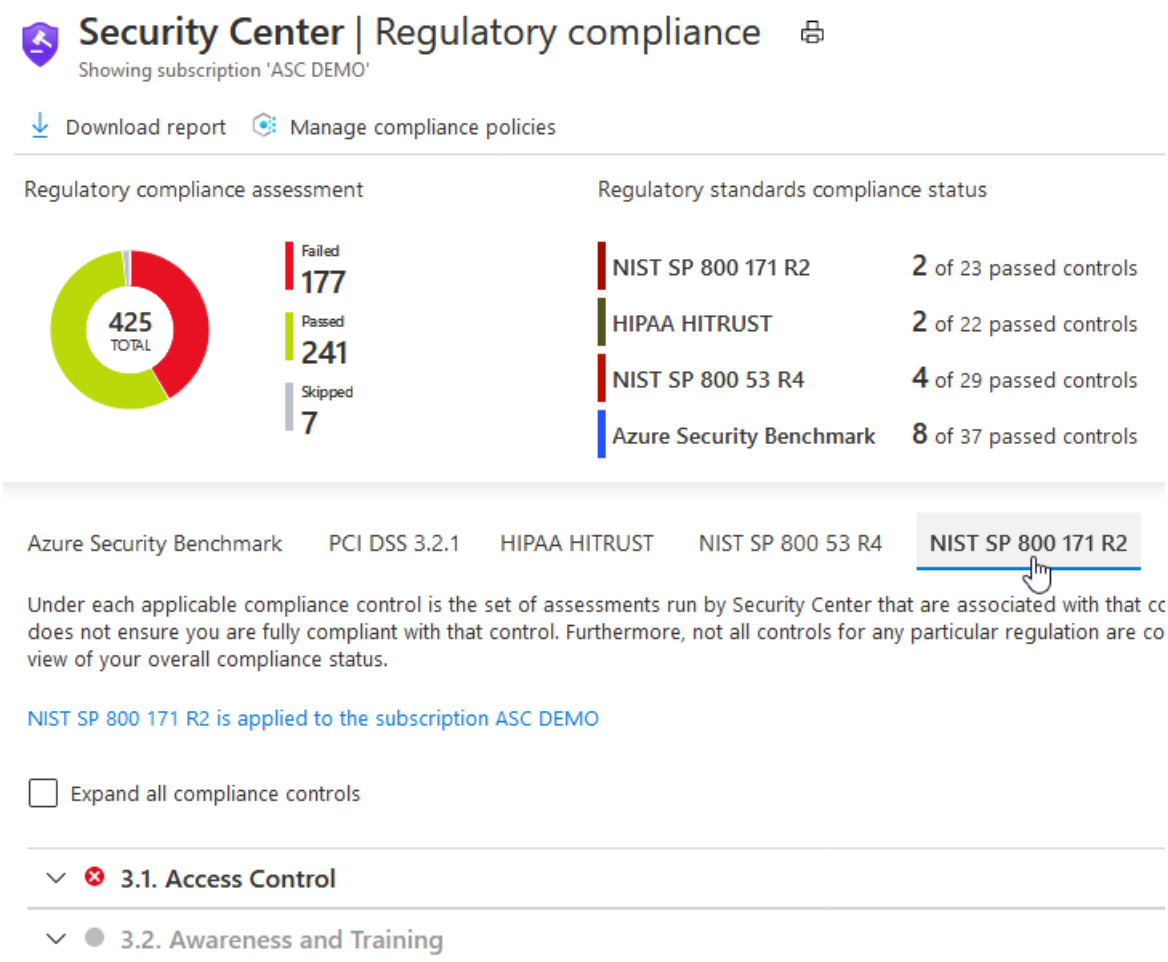

For more information about this compliance standard, see NIST SP 800-171 R2.

Recommendations list now includes filters

You can now filter the list of security recommendations according to a range of criteria. In the following example, the recommendations list has been filtered to show recommendations that:

- are **generally available** (that is, not preview)
- are for **storage accounts**
- support **quick fix** remediation

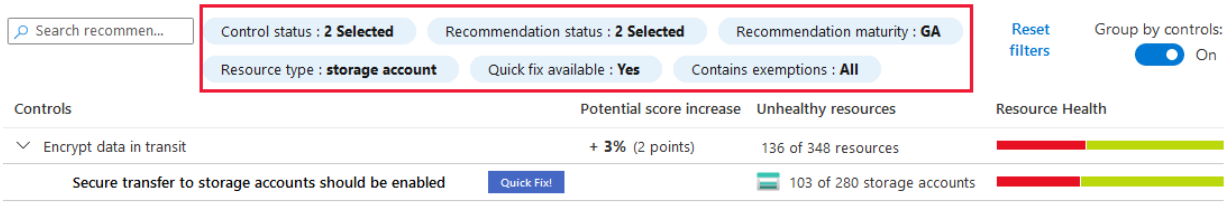

Auto provisioning experience improved and expanded

The auto provisioning feature helps reduce management overhead by installing the required extensions on new - and existing - Azure VMs so they can benefit from Security Center's protections.

As Azure Security Center grows, more extensions have been developed and Security Center can monitor a larger list of resource types. The auto provisioning tools have now been expanded to support other extensions and resource types by leveraging the capabilities of Azure Policy.

You can now configure the auto provisioning of:

- Log Analytics agent
- (New) Azure Policy Add-on for Kubernetes
- (New) Microsoft Dependency agent

Learn more in Auto provisioning agents and extensions from Azure Security Center.

Secure score is now available in continuous export (preview)

With continuous export of secure score, you can stream changes to your score in real-time to Azure Event Hubs or a Log Analytics workspace. Use this capability to:

- track your secure score over time with dynamic reports
- export secure score data to Azure Sentinel (or any other SIEM)
- integrate this data with any processes you might already be using to monitor secure score in your organization

Learn more about how to Continuously export Security Center data.

"System updates should be installed on your machines" recommendation now includes subrecommendations

The **System updates should be installed on your machines** recommendation has been enhanced. The new version includes subrecommendations for each missing update and brings the following improvements:

- A redesigned experience in the Azure Security Center pages of the Azure portal. The recommendation details page for **System updates should be installed on your machines** includes the list of findings as shown below. When you select a single finding, the details pane opens with a link to the remediation information and a list of affected resources.

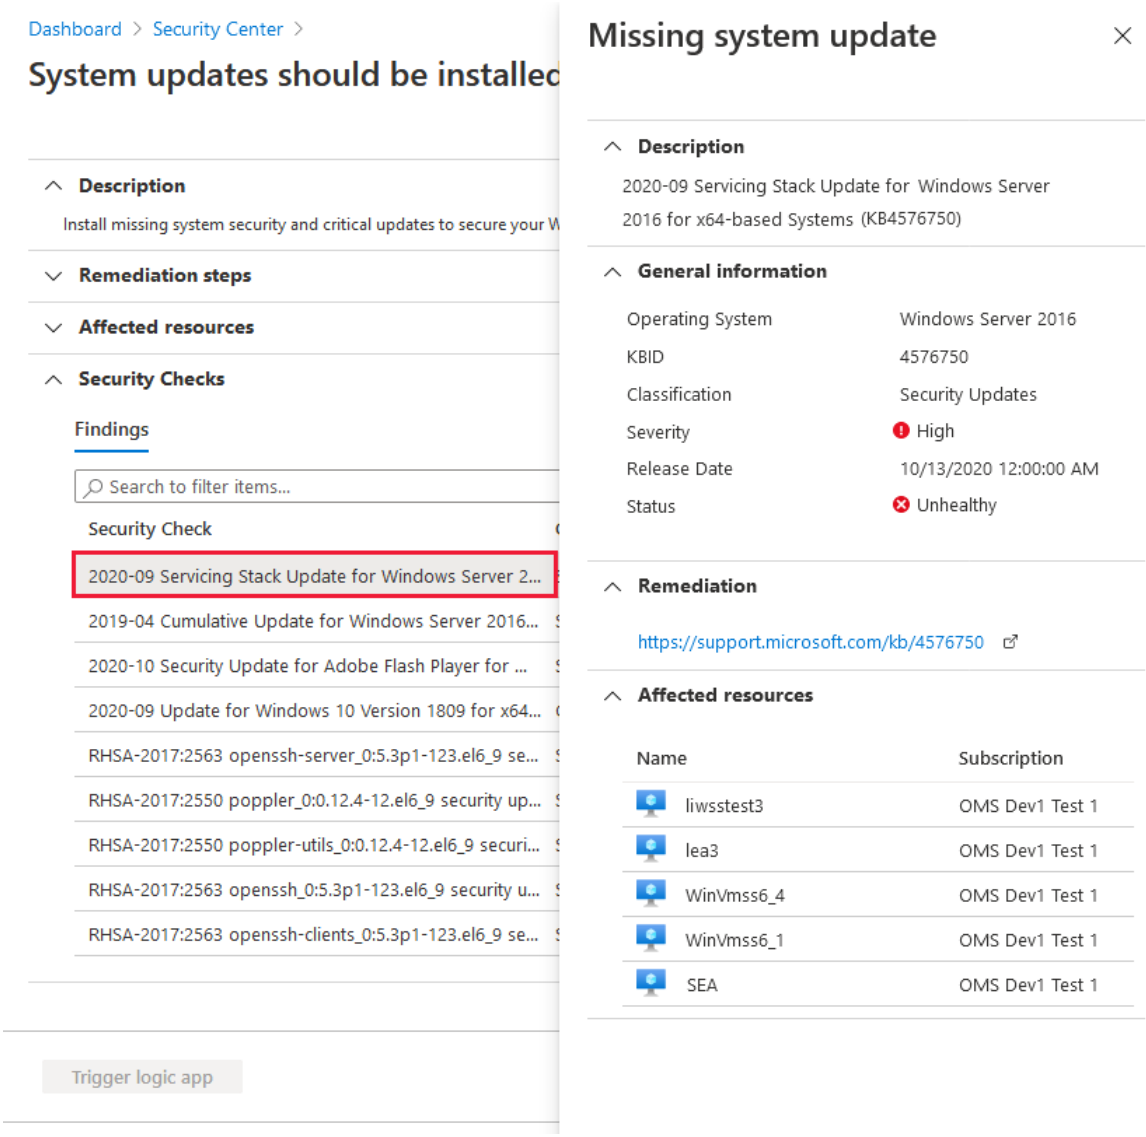

- Enriched data for the recommendation from Azure Resource Graph (ARG). ARG is an Azure service that's designed to provide efficient resource exploration. You can use ARG to query at scale across a given set of subscriptions so that you can effectively govern your environment.

  For Azure Security Center, you can use ARG and the Kusto Query Language (KQL) to query a wide range of security posture data.

  Previously, if you queried this recommendation in ARG, the only available information was that the recommendation needs to be remediated on a machine. The following query of the enhanced version will return each missing system updates grouped by machine.

```
securityresources
| where type =~ "microsoft.security/assessments/subassessments"
| where extract(@"(?i)providers/Microsoft.Security/assessments/([^/]*)", 1, id) == "4ab6e3c5-74dd-8b35-9ab9-f61b30875b27"
| where properties.status.code == "Unhealthy"
```

Policy management page in the Azure portal now shows status of default policy assignments

You can now see whether or not your subscriptions have the default Security Center policy assigned, in the Security Center's **security policy** page of the Azure portal.

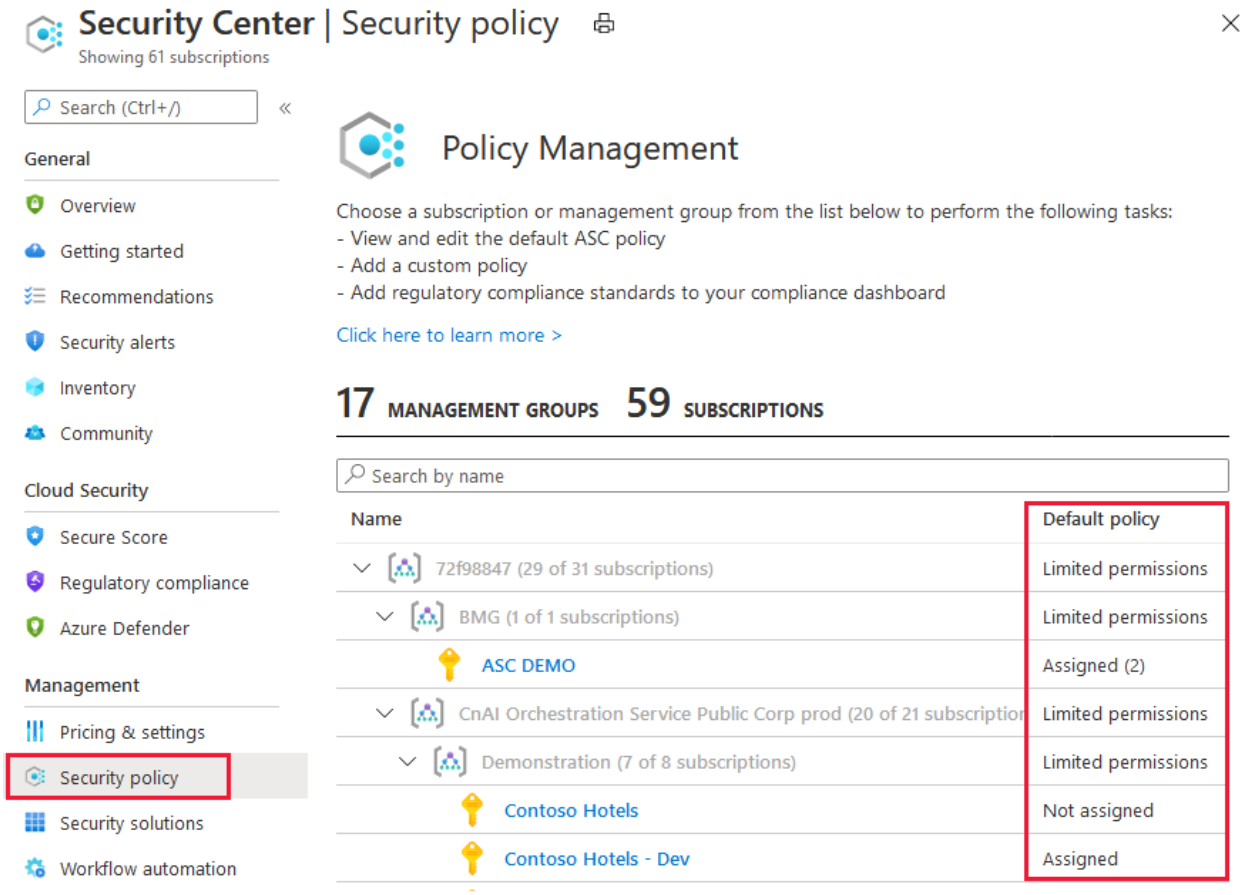

October 2020

Updates in October include:

- Vulnerability assessment for on-premise and multi-cloud machines (preview)
- Azure Firewall recommendation added (preview)
- Authorized IP ranges should be defined on Kubernetes Services recommendation updated with quick fix
- Regulatory compliance dashboard now includes option to remove standards
- Microsoft.Security/securityStatuses table removed from Azure Resource Graph (ARG)

Vulnerability assessment for on-premise and multi-cloud machines (preview)

Azure Defender for servers' integrated vulnerability assessment scanner (powered by Qualys) now scans Azure Arc-enabled servers.

When you've enabled Azure Arc on your non-Azure machines, Security Center will offer to deploy the integrated vulnerability scanner on them - manually and at-scale.

With this update, you can unleash the power of **Azure Defender for servers** to consolidate your vulnerability management program across all of your Azure and non-Azure assets.

Main capabilities:

- Monitoring the VA (vulnerability assessment) scanner provisioning state on Azure Arc machines
- Provisioning the integrated VA agent to unprotected Windows and Linux Azure Arc machines (manually and at-scale)
- Receiving and analyzing detected vulnerabilities from deployed agents (manually and at-scale)
- Unified experience for Azure VMs and Azure Arc machines

Learn more about deploying the integrated Qualys vulnerability scanner to your hybrid machines.

Learn more about Azure Arc-enabled servers.

Azure Firewall recommendation added (preview)

A new recommendation has been added to protect all your virtual networks with Azure Firewall.

The recommendation, **Virtual networks should be protected by Azure Firewall** advises you to restrict access to your virtual networks and prevent potential threats by using Azure Firewall.

Learn more about Azure Firewall.

Authorized IP ranges should be defined on Kubernetes Services recommendation updated with quick fix

The recommendation **Authorized IP ranges should be defined on Kubernetes Services** now has a quick fix option.

For more information about this recommendation and all other Security Center recommendations, see Security recommendations - a reference guide.

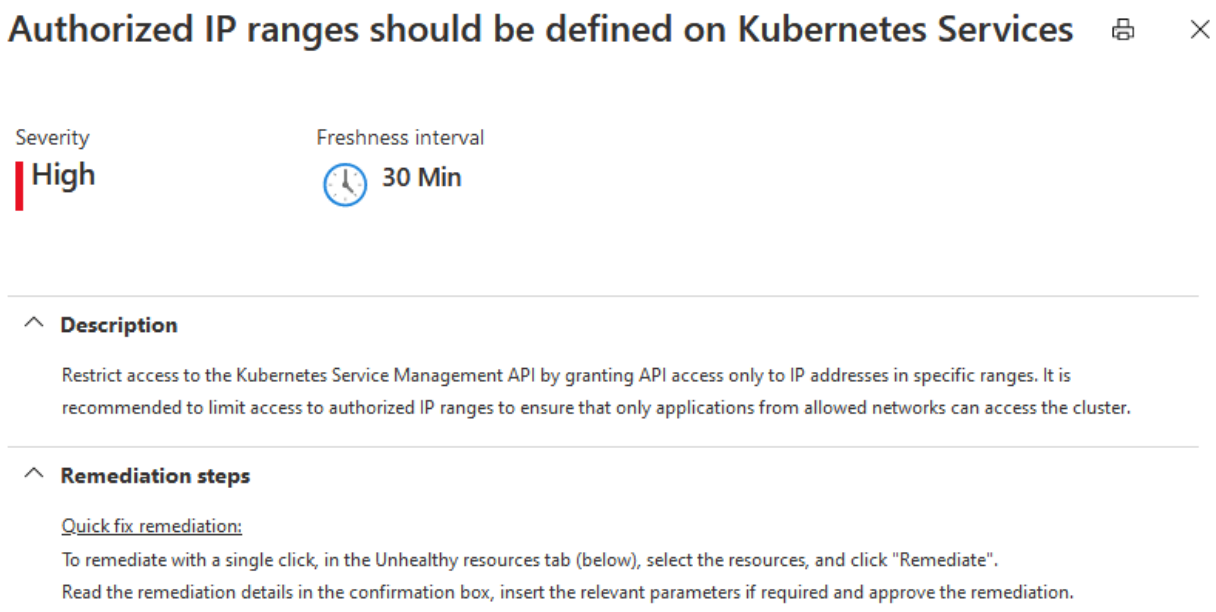

Regulatory compliance dashboard now includes option to remove standards

Security Center's regulatory compliance dashboard provides insights into your compliance posture based on how you're meeting specific compliance controls and requirements.

The dashboard includes a default set of regulatory standards. If any of the supplied standards isn't relevant to your organization, it's now a simple process to remove them from the UI for a subscription. Standards can be removed only at the *subscription* level; not the management group scope.

Learn more in Remove a standard from your dashboard.

Microsoft.Security/securityStatuses table removed from Azure Resource Graph (ARG)

Azure Resource Graph is a service in Azure that is designed to provide efficient resource exploration with the ability to query at scale across a given set of subscriptions so that you can effectively govern your environment.

For Azure Security Center, you can use ARG and the Kusto Query Language (KQL) to query a wide range of security posture data. For example:

- Asset inventory utilizes (ARG)
- We have documented a sample ARG query for how to Identify accounts without multifactor authentication (MFA) enabled

Within ARG, there are tables of data for you to use in your queries.

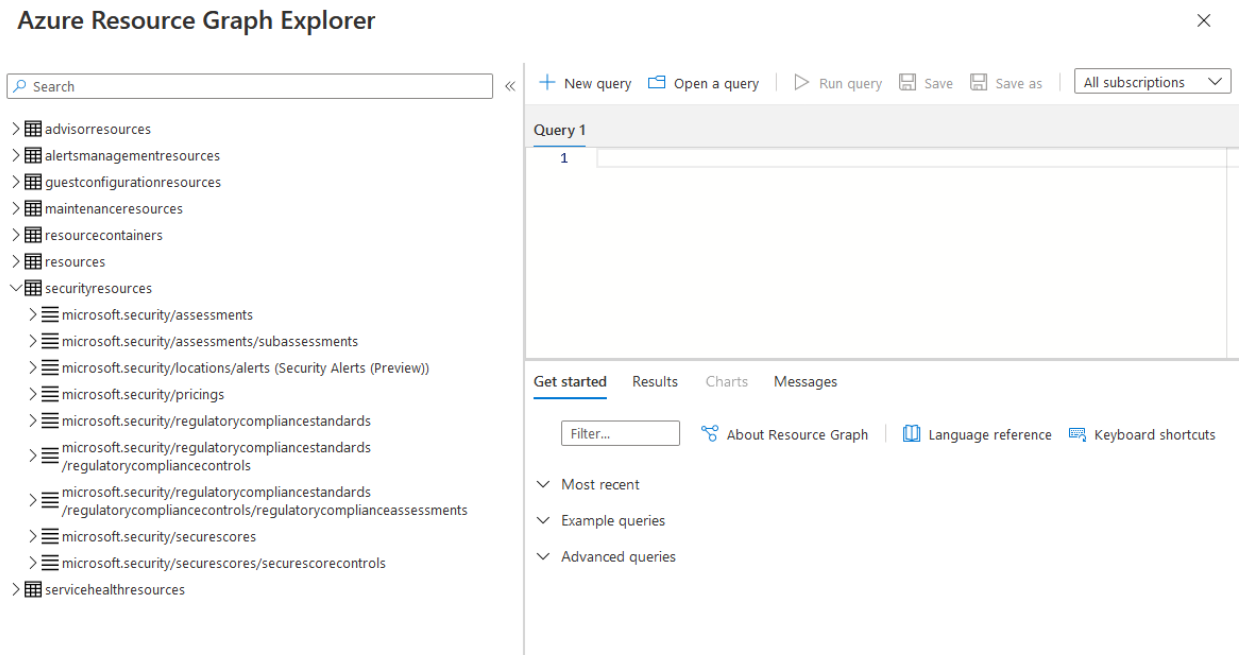

TIP

The ARG documentation lists all the available tables in Azure Resource Graph table and resource type reference.

From this update, the **Microsoft.Security/securityStatuses** table has been removed. The securityStatuses API is still available.

Data replacement can be used by Microsoft.Security/Assessments table.

The major difference between Microsoft.Security/securityStatuses and Microsoft.Security/Assessments is that while the first shows aggregation of assessments, the seconds holds a single record for each.

For example, Microsoft.Security/securityStatuses would return a result with an array of two policyAssessments:

```
{
id: "/subscriptions/449bcidd-3470-4804-ab56-2752595 felab/resourceGroups/mico-
rg/providers/Microsoft.Network/virtualNetworks/mico-rg-
vnet/providers/Microsoft.Security/securityStatuses/mico-rg-vnet",
name: "mico-rg-vnet",
type: "Microsoft.Security/securityStatuses",
properties:  {
    policyAssessments: [
        {assessmentKey: "e3deicce-f4dd-3b34-e496-8b5381bazd7e", category: "Networking", policyName: "Azure
DDOS Protection Standard should be enabled",...},
        {assessmentKey: "sefac66a-1ec5-b063-a824-eb28671dc527", category: "Compute", policyName: "",...}
    ],
    securitystateByCategory: [{category: "Networking", securityState: "None" }, {category: "Compute",...],
    name: "GenericResourceHealthProperties",
    type: "VirtualNetwork",
    securitystate: "High"
}
```

Whereas, Microsoft.Security/Assessments will hold a record for each such policy assessment as follows:

```
{
type: "Microsoft.Security/assessments",
id:  "/subscriptions/449bc1dd-3470-4804-ab56-2752595f01ab/resourceGroups/mico-rg/providers/Microsoft.
Network/virtualNetworks/mico-rg-vnet/providers/Microsoft.Security/assessments/e3delcce-f4dd-3b34-e496-
8b5381ba2d70",
name: "e3deicce-f4dd-3b34-e496-8b5381ba2d70",
properties:  {
    resourceDetails: {Source: "Azure", Id: "/subscriptions/449bc1dd-3470-4804-ab56-
2752595f01ab/resourceGroups/mico-rg/providers/Microsoft.Network/virtualNetworks/mico-rg-vnet"...},
    displayName: "Azure DDOS Protection Standard should be enabled",
    status: (code: "NotApplicable", cause: "VnetHasNOAppGateways", description: "There are no Application
Gateway resources attached to this Virtual Network"...}
}

{
type: "Microsoft.Security/assessments",
id:  "/subscriptions/449bc1dd-3470-4804-ab56-2752595f01ab/resourcegroups/mico-
rg/providers/microsoft.network/virtualnetworks/mico-rg-
vnet/providers/Microsoft.Security/assessments/80fac66a-1ec5-be63-a824-eb28671dc527",
name: "8efac66a-1ec5-be63-a824-eb28671dc527",
properties: {
    resourceDetails: (Source: "Azure", Id: "/subscriptions/449bc1dd-3470-4804-ab56-
2752595f01ab/resourcegroups/mico-rg/providers/microsoft.network/virtualnetworks/mico-rg-vnet"...),
    displayName: "Audit diagnostic setting",
    status:  {code: "Unhealthy"}
}
```

Example of converting an existing ARG query using securityStatuses to now use the assessments table:

Query that references SecurityStatuses:

```
SecurityResources
| where type == 'microsoft.security/securitystatuses' and properties.type == 'virtualMachine'
| where name in ({vmnames})
| project name, resourceGroup, policyAssesments = properties.policyAssessments, resourceRegion = location,
id, resourceDetails = properties.resourceDetails
```

Replacement query for the Assessments table:

```
securityresources
| where type == "microsoft.security/assessments" and id contains "virtualMachine"
| extend resourceName = extract(@"(?i)/([^/]*)/providers/Microsoft.Security/assessments", 1, id)
| extend source = tostring(properties.resourceDetails.Source)
| extend resourceId = trim(" ", tolower(tostring(case(source =~ "azure", properties.resourceDetails.Id,
source =~ "aws", properties.additionalData.AzureResourceId,
source =~ "gcp", properties.additionalData.AzureResourceId,
extract("^(.+)/providers/Microsoft.Security/assessments/.+$",1,id)))))
| extend resourceGroup = tolower(tostring(split(resourceId, "/")[4]))
| where resourceName in ({vmnames})
| project resourceName, resourceGroup, resourceRegion = location, id, resourceDetails =
properties.additionalData
```

Learn more at the following links:

- How to create queries with Azure Resource Graph Explorer
- Kusto Query Language (KQL)

September 2020

Updates in September include:

- Security Center gets a new look!
- Azure Defender released
- Azure Defender for Key Vault is generally available
- Azure Defender for Storage protection for Files and ADLS Gen2 is generally available
- Asset inventory tools are now generally available
- Disable a specific vulnerability finding for scans of container registries and virtual machines
- Exempt a resource from a recommendation
- AWS and GCP connectors in Security Center bring a multi-cloud experience
- Kubernetes workload protection recommendation bundle
- Vulnerability assessment findings are now available in continuous export
- Prevent security misconfigurations by enforcing recommendations when creating new resources
- Network security group recommendations improved
- Deprecated preview AKS recommendation "Pod Security Policies should be defined on Kubernetes Services"
- Email notifications from Azure Security Center improved
- Secure score doesn't include preview recommendations
- Recommendations now include a severity indicator and the freshness interval

Security Center gets a new look!

We've released a refreshed UI for Security Center's portal pages. The new pages include a new overview page and dashboards for secure score, asset inventory, and Azure Defender.

The redesigned overview page now has a tile for accessing the secure score, asset inventory, and Azure Defender dashboards. It also has a tile linking to the regulatory compliance dashboard.

Learn more about the overview page.

Azure Defender released

Azure Defender is the cloud workload protection platform (CWPP) integrated within Security Center for advanced, intelligent, protection of your Azure and hybrid workloads. It replaces Security Center's standard pricing tier option.

When you enable Azure Defender from the **Pricing and settings** area of Azure Security Center, the following Defender plans are all enabled simultaneously and provide comprehensive defenses for the compute, data, and service layers of your environment:

- Azure Defender for servers
- Azure Defender for App Service
- Azure Defender for Storage
- Azure Defender for SQL
- Azure Defender for Key Vault
- Azure Defender for Kubernetes
- Azure Defender for container registries

Each of these plans is explained separately in the Security Center documentation.

With its dedicated dashboard, Azure Defender provides security alerts and advanced threat protection for virtual machines, SQL databases, containers, web applications, your network, and more.

Learn more about Azure Defender

Azure Defender for Key Vault is generally available

Azure Key Vault is a cloud service that safeguards encryption keys and secrets like certificates, connection strings, and passwords.

Azure Defender for Key Vault provides Azure-native, advanced threat protection for Azure Key Vault, providing an additional layer of security intelligence. By extension, Azure Defender for Key Vault is consequently protecting many of the resources dependent upon your Key Vault accounts.

The optional plan is now GA. This feature was in preview as "advanced threat protection for Azure Key Vault".

Also, the Key Vault pages in the Azure portal now include a dedicated **Security** page for **Security Center** recommendations and alerts.

Learn more in Azure Defender for Key Vault.

Azure Defender for Storage protection for Files and ADLS Gen2 is generally available

Azure Defender for Storage detects potentially harmful activity on your Azure Storage accounts. Your data can be protected whether it's stored as blob containers, file shares, or data lakes.

Support for Azure Files and Azure Data Lake Storage Gen2 is now generally available.

From 1 October 2020, we'll begin charging for protecting resources on these services.

Learn more in Azure Defender for Storage.

Asset inventory tools are now generally available

The asset inventory page of Azure Security Center provides a single page for viewing the security posture of the resources you've connected to Security Center.

Security Center periodically analyzes the security state of your Azure resources to identify potential security vulnerabilities. It then provides you with recommendations on how to remediate those vulnerabilities.

When any resource has outstanding recommendations, they'll appear in the inventory.

Learn more in Explore and manage your resources with asset inventory.

Disable a specific vulnerability finding for scans of container registries and virtual machines

Azure Defender includes vulnerability scanners to scan images in your Azure Container Registry and your virtual machines.

If you have an organizational need to ignore a finding, rather than remediate it, you can optionally disable it. Disabled findings don't impact your secure score or generate unwanted noise.

When a finding matches the criteria you've defined in your disable rules, it won't appear in the list of findings.

This option is available from the recommendations details pages for:

- **Vulnerabilities in Azure Container Registry images should be remediated**
- **Vulnerabilities in your virtual machines should be remediated**

Learn more in Disable specific findings for your container images and Disable specific findings for your virtual machines.

Exempt a resource from a recommendation

Occasionally, a resource will be listed as unhealthy regarding a specific recommendation (and therefore lowering your secure score) even though you feel it shouldn't be. It might have been remediated by a process not tracked by Security Center. Or perhaps your organization has decided to accept the risk for that specific resource.

In such cases, you can create an exemption rule and ensure that resource isn't listed amongst the unhealthy resources in the future. These rules can include documented justifications as described below.

Learn more in Exempt a resource from recommendations and secure score.

AWS and GCP connectors in Security Center bring a multi-cloud experience

With cloud workloads commonly spanning multiple cloud platforms, cloud security services must do the same.

Azure Security Center now protects workloads in Azure, Amazon Web Services (AWS), and Google Cloud Platform (GCP).

Onboarding your AWS and GCP accounts into Security Center, integrates AWS Security Hub, GCP Security Command and Azure Security Center.

Learn more in Connect your AWS accounts to Azure Security Center and Connect your GCP accounts to Azure Security Center.

Kubernetes workload protection recommendation bundle

To ensure that Kubernetes workloads are secure by default, Security Center is adding Kubernetes level hardening recommendations, including enforcement options with Kubernetes admission control.

When you've installed the Azure Policy add-on for Kubernetes on your AKS cluster, every request to the Kubernetes API server will be monitored against the predefined set of best practices before being persisted to the cluster. You can then configure to enforce the best practices and mandate them for future workloads.

For example, you can mandate that privileged containers shouldn't be created, and any future requests to do so will be blocked.

Learn more in Workload protection best-practices using Kubernetes admission control.

Vulnerability assessment findings are now available in continuous export

Use continuous export to stream your alerts and recommendations to Azure Event Hubs, Log Analytics workspaces, or Azure Monitor. From there, you can integrate this data with SIEMs (such as Azure Sentinel, Power BI, Azure Data Explorer, and more.

Security Center's integrated vulnerability assessment tools return findings about your resources as actionable recommendations within a 'parent' recommendation such as "Vulnerabilities in your virtual machines should be remediated".

The security findings are now available for export through continuous export when you select recommendations and enable the **include security findings** option.

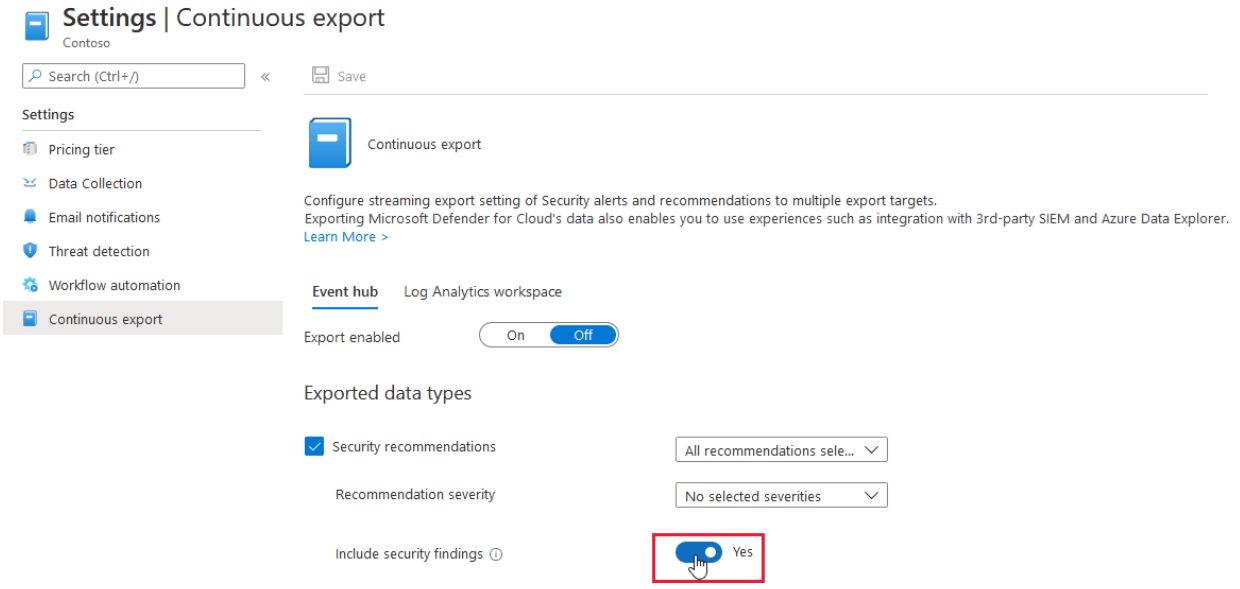

Related pages:

- Security Center's integrated Qualys vulnerability assessment solution for Azure virtual machines
- Security Center's integrated vulnerability assessment solution for Azure Container Registry images

- Continuous export

Prevent security misconfigurations by enforcing recommendations when creating new resources

Security misconfigurations are a major cause of security incidents. Security Center now has the ability to help *prevent* misconfigurations of new resources with regard to specific recommendations.

This feature can help keep your workloads secure and stabilize your secure score.

Enforcing a secure configuration, based on a specific recommendation, is offered in two modes:

- Using the **Deny** effect of Azure Policy, you can stop unhealthy resources from being created
- Using the **Enforce** option, you can take advantage of Azure Policy's **DeployIfNotExist** effect and automatically remediate non-compliant resources upon creation

This is available for selected security recommendations and can be found at the top of the resource details page.

Learn more in Prevent misconfigurations with Enforce/Deny recommendations.

Network security group recommendations improved

The following security recommendations related to network security groups have been improved to reduce some instances of false positives.

- All network ports should be restricted on NSG associated to your VM
- Management ports should be closed on your virtual machines
- Internet-facing virtual machines should be protected with Network Security Groups
- Subnets should be associated with a Network Security Group

Deprecated preview AKS recommendation "Pod Security Policies should be defined on Kubernetes Services"

The preview recommendation "Pod Security Policies should be defined on Kubernetes Services" is being deprecated as described in the Azure Kubernetes Service documentation.

The pod security policy (preview) feature, is set for deprecation and will no longer be available after October 15, 2020 in favor of Azure Policy for AKS.

After pod security policy (preview) is deprecated, you must disable the feature on any existing clusters using the deprecated feature to perform future cluster upgrades and stay within Azure support.

Email notifications from Azure Security Center improved

The following areas of the emails regarding security alerts have been improved:

- Added the ability to send email notifications about alerts for all severity levels
- Added the ability to notify users with different Azure roles on the subscription
- We're proactively notifying subscription owners by default on high-severity alerts (which have a high-probability of being genuine breaches)
- We've removed the phone number field from the email notifications configuration page

Learn more in Set up email notifications for security alerts.

Secure score doesn't include preview recommendations

Security Center continually assesses your resources, subscriptions, and organization for security issues. It then aggregates all the findings into a single score so that you can tell, at a glance, your current security situation: the higher the score, the lower the identified risk level.

As new threats are discovered, new security advice is made available in Security Center through new recommendations. To avoid surprise changes your secure score, and to provide a grace period in which you can explore new recommendations before they impact your scores, recommendations flagged as **Preview** are no

longer included in the calculations of your secure score. They should still be remediated wherever possible, so that when the preview period ends they'll contribute towards your score.

Also, **Preview** recommendations don't render a resource "Unhealthy".

An example of a preview recommendation:

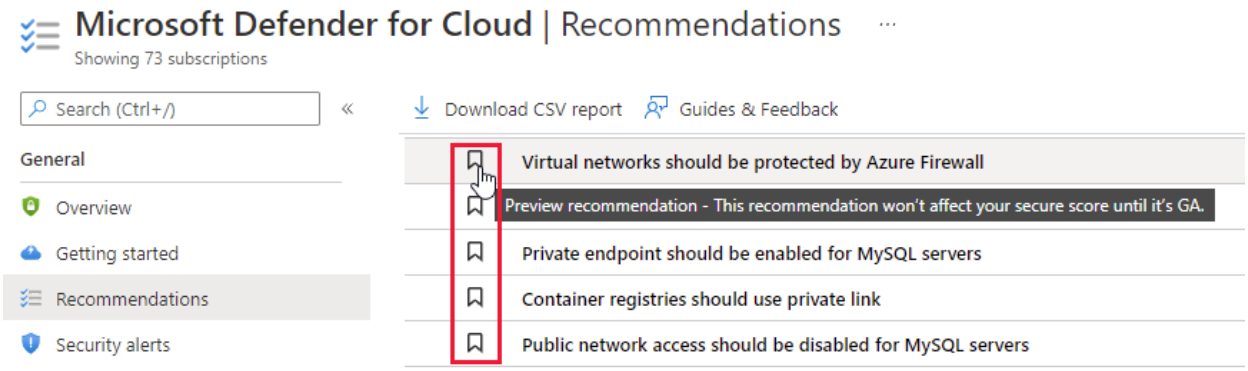

Learn more about secure score.

Recommendations now include a severity indicator and the freshness interval

The details page for recommendations now includes a freshness interval indicator (whenever relevant) and a clear display of the severity of the recommendation.

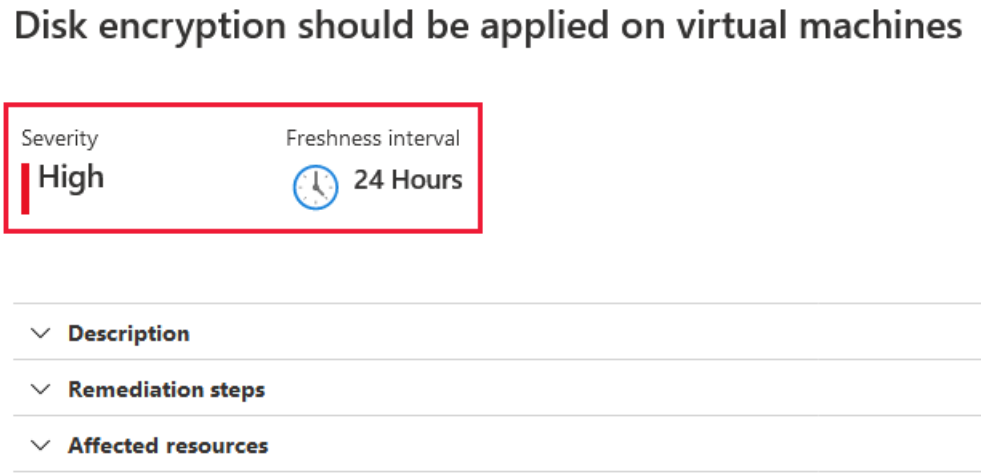

August 2020

Updates in August include:

- Asset inventory - powerful new view of the security posture of your assets
- Added support for Azure Active Directory security defaults (for multifactor authentication)
- Service principals recommendation added
- Vulnerability assessment on VMs - recommendations and policies consolidated
- New AKS security policies added to ASC_default initiative – for use by private preview customers only

Asset inventory - powerful new view of the security posture of your assets

Security Center's asset inventory (currently in preview) provides a way to view the security posture of the resources you've connected to Security Center.

Security Center periodically analyzes the security state of your Azure resources to identify potential security vulnerabilities. It then provides you with recommendations on how to remediate those vulnerabilities. When any resource has outstanding recommendations, they'll appear in the inventory.

You can use the view and its filters to explore your security posture data and take further actions based on your findings.

Learn more about asset inventory.

Added support for Azure Active Directory security defaults (for multifactor authentication)

Security Center has added full support for security defaults, Microsoft's free identity security protections.

Security defaults provide preconfigured identity security settings to defend your organization from common identity-related attacks. Security defaults already protecting more than 5 million tenants overall; 50,000 tenants are also protected by Security Center.

Security Center now provides a security recommendation whenever it identifies an Azure subscription without security defaults enabled. Until now, Security Center recommended enabling multifactor authentication using conditional access, which is part of the Azure Active Directory (AD) premium license. For customers using Azure AD free, we now recommend enabling security defaults.

Our goal is to encourage more customers to secure their cloud environments with MFA, and mitigate one of the highest risks that is also the most impactful to your secure score.

Learn more about security defaults.

Service principals recommendation added

A new recommendation has been added to recommend that Security Center customers using management certificates to manage their subscriptions switch to service principals.

The recommendation, **Service principals should be used to protect your subscriptions instead of Management Certificates** advises you to use Service Principals or Azure Resource Manager to more securely manage your subscriptions.

Learn more about Application and service principal objects in Azure Active Directory.

Vulnerability assessment on VMs - recommendations and policies consolidated

Security Center inspects your VMs to detect whether they're running a vulnerability assessment solution. If no vulnerability assessment solution is found, Security Center provides a recommendation to simplify the deployment.

When vulnerabilities are found, Security Center provides a recommendation summarizing the findings for you to investigate and remediate as necessary.

To ensure a consistent experience for all users, regardless of the scanner type they're using, we've unified four recommendations into the following two:

UNIFIED RECOMMENDATION	CHANGE DESCRIPTION
A vulnerability assessment solution should be enabled on your virtual machines	Replaces the following two recommendations: ***** Enable the built-in vulnerability assessment solution on virtual machines (powered by Qualys (now deprecated) (Included with standard tier) ***** Vulnerability assessment solution should be installed on your virtual machines (now deprecated) (Standard and free tiers)
Vulnerabilities in your virtual machines should be remediated	Replaces the following two recommendations: ***** Remediate vulnerabilities found on your virtual machines (powered by Qualys) (now deprecated) ***** Vulnerabilities should be remediated by a Vulnerability Assessment solution (now deprecated)

UNIFIED RECOMMENDATION	CHANGE DESCRIPTION

Now you'll use the same recommendation to deploy Security Center's vulnerability assessment extension or a privately licensed solution ("BYOL") from a partner such as Qualys or Rapid7.

Also, when vulnerabilities are found and reported to Security Center, a single recommendation will alert you to the findings regardless of the vulnerability assessment solution that identified them.

Updating dependencies

If you have scripts, queries, or automations referring to the previous recommendations or policy keys/names, use the tables below to update the references:

Before August 2020

RECOMMENDATION	SCOPE
Enable the built-in vulnerability assessment solution on virtual machines (powered by Qualys) Key: 550e890b-e652-4d22-8274-60b3bdb24c63	Built-in
Remediate vulnerabilities found on your virtual machines (powered by Qualys) Key: 1195afff-c881-495e-9bc5-1486211ae03f	Built-in
Vulnerability assessment solution should be installed on your virtual machines Key: 01b1ed4c-b733-4fee-b145-f23236e70cf3	BYOL
Vulnerabilities should be remediated by a Vulnerability Assessment solution Key: 71992a2a-d168-42e0-b10e-6b45fa2ecddb	BYOL

POLICY	SCOPE
Vulnerability assessment should be enabled on virtual machines Policy ID: 501541f7-f7e7-4cd6-868c-4190fdad3ac9	Built-in
Vulnerabilities should be remediated by a vulnerability assessment solution Policy ID: 760a85ff-6162-42b3-8d70-698e268f648c	BYOL

From August 2020

RECOMMENDATION	SCOPE
A vulnerability assessment solution should be enabled on your virtual machines Key: ffff0522-1e88-47fc-8382-2a80ba848f5d	Built-in + BYOL
Vulnerabilities in your virtual machines should be remediated Key: 1195afff-c881-495e-9bc5-1486211ae03f	Built-in + BYOL

POLICY	SCOPE
Vulnerability assessment should be enabled on virtual machines Policy ID: 501541f7-f7e7-4cd6-868c-4190fdad3ac9	Built-in + BYOL

New AKS security policies added to ASC_default initiative – for use by private preview customers only

To ensure that Kubernetes workloads are secure by default, Security Center is adding Kubernetes level policies and hardening recommendations, including enforcement options with Kubernetes admission control.

The early phase of this project includes a private preview and the addition of new (disabled by default) policies to the ASC_default initiative.

You can safely ignore these policies and there will be no impact on your environment. If you'd like to enable them, sign up for the preview at https://aka.ms/SecurityPrP and select from the following options:

1. **Single Preview** – To join only this private preview. Explicitly mention "ASC Continuous Scan" as the preview you would like to join.
2. **Ongoing Program** – To be added to this and future private previews. You'll need to complete a profile and privacy agreement.

July 2020

Updates in July include:

- Vulnerability assessment for virtual machines is now available for non-marketplace images
- Threat protection for Azure Storage expanded to include Azure Files and Azure Data Lake Storage Gen2 (preview)
- Eight new recommendations to enable threat protection features
- Container security improvements - faster registry scanning and refreshed documentation
- Adaptive application controls updated with a new recommendation and support for wildcards in path rules
- Six policies for SQL advanced data security deprecated

Vulnerability assessment for virtual machines is now available for non-marketplace images

When deploying a vulnerability assessment solution, Security Center previously performed a validation check before deploying. The check was to confirm a marketplace SKU of the destination virtual machine.

From this update, the check has been removed and you can now deploy vulnerability assessment tools to 'custom' Windows and Linux machines. Custom images are ones that you've modified from the marketplace defaults.

Although you can now deploy the integrated vulnerability assessment extension (powered by Qualys) on many more machines, support is only available if you're using an OS listed in Deploy the integrated vulnerability scanner to standard tier VMs

Learn more about the integrated vulnerability scanner for virtual machines (requires Azure Defender).

Learn more about using your own privately-licensed vulnerability assessment solution from Qualys or Rapid7 in Deploying a partner vulnerability scanning solution.

Threat protection for Azure Storage expanded to include Azure Files and Azure Data Lake Storage Gen2 (preview)

Threat protection for Azure Storage detects potentially harmful activity on your Azure Storage accounts. Security Center displays alerts when it detects attempts to access or exploit your storage accounts.

Your data can be protected whether it's stored as blob containers, file shares, or data lakes.

Eight new recommendations to enable threat protection features

Eight new recommendations have been added to provide a simple way to enable Azure Security Center's threat protection features for the following resource types: virtual machines, App Service plans, Azure SQL Database servers, SQL servers on machines, Azure Storage accounts, Azure Kubernetes Service clusters, Azure Container Registry registries, and Azure Key Vault vaults.

The new recommendations are:

- **Advanced data security should be enabled on Azure SQL Database servers**
- **Advanced data security should be enabled on SQL servers on machines**
- **Advanced threat protection should be enabled on Azure App Service plans**
- **Advanced threat protection should be enabled on Azure Container Registry registries**
- **Advanced threat protection should be enabled on Azure Key Vault vaults**
- **Advanced threat protection should be enabled on Azure Kubernetes Service clusters**
- **Advanced threat protection should be enabled on Azure Storage accounts**
- **Advanced threat protection should be enabled on virtual machines**

These new recommendations belong to the **Enable Azure Defender** security control.

The recommendations also include the quick fix capability.

> **IMPORTANT**
>
> Remediating any of these recommendations will result in charges for protecting the relevant resources. These charges will begin immediately if you have related resources in the current subscription. Or in the future, if you add them at a later date.
>
> For example, if you don't have any Azure Kubernetes Service clusters in your subscription and you enable the threat protection, no charges will be incurred. If, in the future, you add a cluster on the same subscription, it will automatically be protected and charges will begin at that time.

Learn more about each of these in the security recommendations reference page.

Learn more about threat protection in Azure Security Center.

Container security improvements - faster registry scanning and refreshed documentation

As part of the continuous investments in the container security domain, we are happy to share a significant performance improvement in Security Center's dynamic scans of container images stored in Azure Container Registry. Scans now typically complete in approximately two minutes. In some cases, they might take up to 15 minutes.

To improve the clarity and guidance regarding Azure Security Center's container security capabilities, we've also refreshed the container security documentation pages.

Learn more about Security Center's container security in the following articles:

- Overview of Security Center's container security features
- Details of the integration with Azure Container Registry
- Details of the integration with Azure Kubernetes Service
- How-to scan your registries and harden your Docker hosts
- Security alerts from the threat protection features for Azure Kubernetes Service clusters
- Security recommendations for containers

Adaptive application controls updated with a new recommendation and support for wildcards in path rules

The adaptive application controls feature has received two significant updates:

- A new recommendation identifies potentially legitimate behavior that hasn't previously been allowed. The new recommendation, **Allowlist rules in your adaptive application control policy should be updated**, prompts you to add new rules to the existing policy to reduce the number of false positives in adaptive application controls violation alerts.
- Path rules now support wildcards. From this update, you can configure allowed path rules using wildcards. There are two supported scenarios:
  - Using a wildcard at the end of a path to allow all executables within this folder and sub-folders
  - Using a wildcard in the middle of a path to enable a known executable name with a changing folder name (e.g. personal user folders with a known executable, automatically generated folder names, etc.).

Learn more about adaptive application controls.

Six policies for SQL advanced data security deprecated

Six policies related to advanced data security for SQL machines are being deprecated:

- Advanced threat protection types should be set to 'All' in SQL managed instance advanced data security settings
- Advanced threat protection types should be set to 'All' in SQL server advanced data security settings
- Advanced data security settings for SQL managed instance should contain an email address to receive security alerts
- Advanced data security settings for SQL server should contain an email address to receive security alerts
- Email notifications to admins and subscription owners should be enabled in SQL managed instance advanced data security settings
- Email notifications to admins and subscription owners should be enabled in SQL server advanced data security settings

Learn more about built-in policies.

June 2020

Updates in June include:

- Secure score API (preview)
- Advanced data security for SQL machines (Azure, other clouds, and on-premises) (preview)
- Two new recommendations to deploy the Log Analytics agent to Azure Arc machines (preview)
- New policies to create continuous export and workflow automation configurations at scale
- New recommendation for using NSGs to protect non-internet-facing virtual machines
- New policies for enabling threat protection and advanced data security

Secure score API (preview)

You can now access your score via the secure score API (currently in preview). The API methods provide the flexibility to query the data and build your own reporting mechanism of your secure scores over time. For example, you can use the **Secure Scores** API to get the score for a specific subscription. In addition, you can use the **Secure Score Controls** API to list the security controls and the current score of your subscriptions.

For examples of external tools made possible with the secure score API, see the secure score area of our GitHub community.

Learn more about secure score and security controls in Azure Security Center.

Advanced data security for SQL machines (Azure, other clouds, and on-premises) (preview)

Azure Security Center's advanced data security for SQL machines now protects SQL Servers hosted in Azure, on other cloud environments, and even on-premises machines. This extends the protections for your Azure-native SQL Servers to fully support hybrid environments.

Advanced data security provides vulnerability assessment and advanced threat protection for your SQL machines wherever they're located.

Set up involves two steps:

1. Deploying the Log Analytics agent to your SQL Server's host machine to provide the connection to Azure account.
2. Enabling the optional bundle in Security Center's pricing and settings page.

Learn more about advanced data security for SQL machines.

Two new recommendations to deploy the Log Analytics agent to Azure Arc machines (preview)

Two new recommendations have been added to help deploy the Log Analytics Agent to your Azure Arc machines and ensure they're protected by Azure Security Center:

- **Log Analytics agent should be installed on your Windows-based Azure Arc machines (Preview)**
- **Log Analytics agent should be installed on your Linux-based Azure Arc machines (Preview)**

These new recommendations will appear in the same four security controls as the existing (related) recommendation, **Monitoring agent should be installed on your machines**: remediate security configurations, apply adaptive application control, apply system updates, and enable endpoint protection.

The recommendations also include the Quick fix capability to help speed up the deployment process.

Learn more about these two new recommendations in the Compute and app recommendations table.

Learn more about how Azure Security Center uses the agent in What is the Log Analytics agent?.

Learn more about extensions for Azure Arc machines.

New policies to create continuous export and workflow automation configurations at scale

Automating your organization's monitoring and incident response processes can greatly improve the time it takes to investigate and mitigate security incidents.

To deploy your automation configurations across your organization, use these built-in 'DeployIfdNotExist' Azure policies to create and configure continuous export and workflow automation procedures:

The policy definitions can be found in Azure Policy:

GOAL	POLICY	POLICY ID
Continuous export to Event Hub	Deploy export to Event Hub for Azure Security Center alerts and recommendations	cdfcce10-4578-4ecd-9703-530938e4abcb
Continuous export to Log Analytics workspace	Deploy export to Log Analytics workspace for Azure Security Center alerts and recommendations	ffb6f416-7bd2-4488-8828-56585fef2be9
Workflow automation for security alerts	Deploy Workflow Automation for Azure Security Center alerts	f1525828-9a90-4fcf-be48-268cdd02361e

GOAL	POLICY	POLICY ID
Workflow automation for security recommendations	Deploy Workflow Automation for Azure Security Center recommendations	73d6ab6c-2475-4850-afd6-43795f3492ef

Get started with workflow automation templates.

Learn more about using the two export policies in Configure workflow automation at scale using the supplied policies and Set up a continuous export.

New recommendation for using NSGs to protect non-internet-facing virtual machines

The "implement security best practices" security control now includes the following new recommendation:

- **Non-internet-facing virtual machines should be protected with network security groups**

An existing recommendation, **Internet-facing virtual machines should be protected with network security groups**, didn't distinguish between internet-facing and non-internet facing VMs. For both, a high-severity recommendation was generated if a VM wasn't assigned to a network security group. This new recommendation separates the non-internet-facing machines to reduce the false positives and avoid unnecessary high-severity alerts.

Learn more in the Network recommendations table.

New policies for enabling threat protection and advanced data security

The new policy definitions below were added to the ASC Default initiative and are designed to assist with enabling threat protection or advanced data security for the relevant resource types.

The policy definitions can be found in Azure Policy:

POLICY	POLICY ID
Advanced data security should be enabled on Azure SQL Database servers	7fe3b40f-802b-4cdd-8bd4-fd799c948cc2
Advanced data security should be enabled on SQL servers on machines	6581d072-105e-4418-827f-bd446d56421b
Advanced threat protection should be enabled on Azure Storage accounts	308fbb08-4ab8-4e67-9b29-592e93fb94fa
Advanced threat protection should be enabled on Azure Key Vault vaults	0e6763cc-5078-4e64-889d-ff4d9a839047
Advanced threat protection should be enabled on Azure App Service plans	2913021d-f2fd-4f3d-b958-22354e2bdbcb
Advanced threat protection should be enabled on Azure Container Registry registries	c25d9a16-bc35-4e15-a7e5-9db606bf9ed4
Advanced threat protection should be enabled on Azure Kubernetes Service clusters	523b5cd1-3e23-492f-a539-13118b6d1e3a
Advanced threat protection should be enabled on Virtual Machines	4da35fc9-c9e7-4960-aec9-797fe7d9051d

POLICY	POLICY ID

Learn more about Threat protection in Azure Security Center.

May 2020

Updates in May include:

- Alert suppression rules (preview)
- Virtual machine vulnerability assessment is now generally available
- Changes to just-in-time (JIT) virtual machine (VM) access
- Custom recommendations have been moved to a separate security control
- Toggle added to view recommendations in controls or as a flat list
- Expanded security control "Implement security best practices"
- Custom policies with custom metadata are now generally available
- Crash dump analysis capabilities migrating to fileless attack detection

Alert suppression rules (preview)

This new feature (currently in preview) helps reduce alert fatigue. Use rules to automatically hide alerts that are known to be innocuous or related to normal activities in your organization. This lets you focus on the most relevant threats.

Alerts that match your enabled suppression rules will still be generated, but their state will be set to dismissed. You can see the state in the Azure portal or however you access your Security Center security alerts.

Suppression rules define the criteria for which alerts should be automatically dismissed. Typically, you'd use a suppression rule to:

- suppress alerts that you've identified as false positives
- suppress alerts that are being triggered too often to be useful

Learn more about suppressing alerts from Azure Security Center's threat protection.

Virtual machine vulnerability assessment is now generally available

Security Center's standard tier now includes an integrated vulnerability assessment for virtual machines for no additional fee. This extension is powered by Qualys but reports its findings directly back to Security Center. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Security Center.

The new solution can continuously scan your virtual machines to find vulnerabilities and present the findings in Security Center.

To deploy the solution, use the new security recommendation:

"Enable the built-in vulnerability assessment solution on virtual machines (powered by Qualys)"

Learn more about Security Center's integrated vulnerability assessment for virtual machines.

Changes to just-in-time (JIT) virtual machine (VM) access

Security Center includes an optional feature to protect the management ports of your VMs. This provides a defense against the most common form of brute force attacks.

This update brings the following changes to this feature:

- The recommendation that advises you to enable JIT on a VM has been renamed. Formerly, "Just-in-time network access control should be applied on virtual machines" it's now: "Management ports of virtual

machines should be protected with just-in-time network access control".

- The recommendation is triggered only if there are open management ports.

Learn more about the JIT access feature.

Custom recommendations have been moved to a separate security control

One security control introduced with the enhanced secure score was "Implement security best practices". Any custom recommendations created for your subscriptions were automatically placed in that control.

To make it easier to find your custom recommendations, we've moved them into a dedicated security control, "Custom recommendations". This control has no impact on your secure score.

Learn more about security controls in Enhanced secure score (preview) in Azure Security Center.

Toggle added to view recommendations in controls or as a flat list

Security controls are logical groups of related security recommendations. They reflect your vulnerable attack surfaces. A control is a set of security recommendations, with instructions that help you implement those recommendations.

To immediately see how well your organization is securing each individual attack surface, review the scores for each security control.

By default, your recommendations are shown in the security controls. From this update, you can also display them as a list. To view them as simple list sorted by the health status of the affected resources, use the new toggle 'Group by controls'. The toggle is above the list in the portal.

The security controls - and this toggle - are part of the new secure score experience. Remember to send us your feedback from within the portal.

Learn more about security controls in Enhanced secure score (preview) in Azure Security Center.

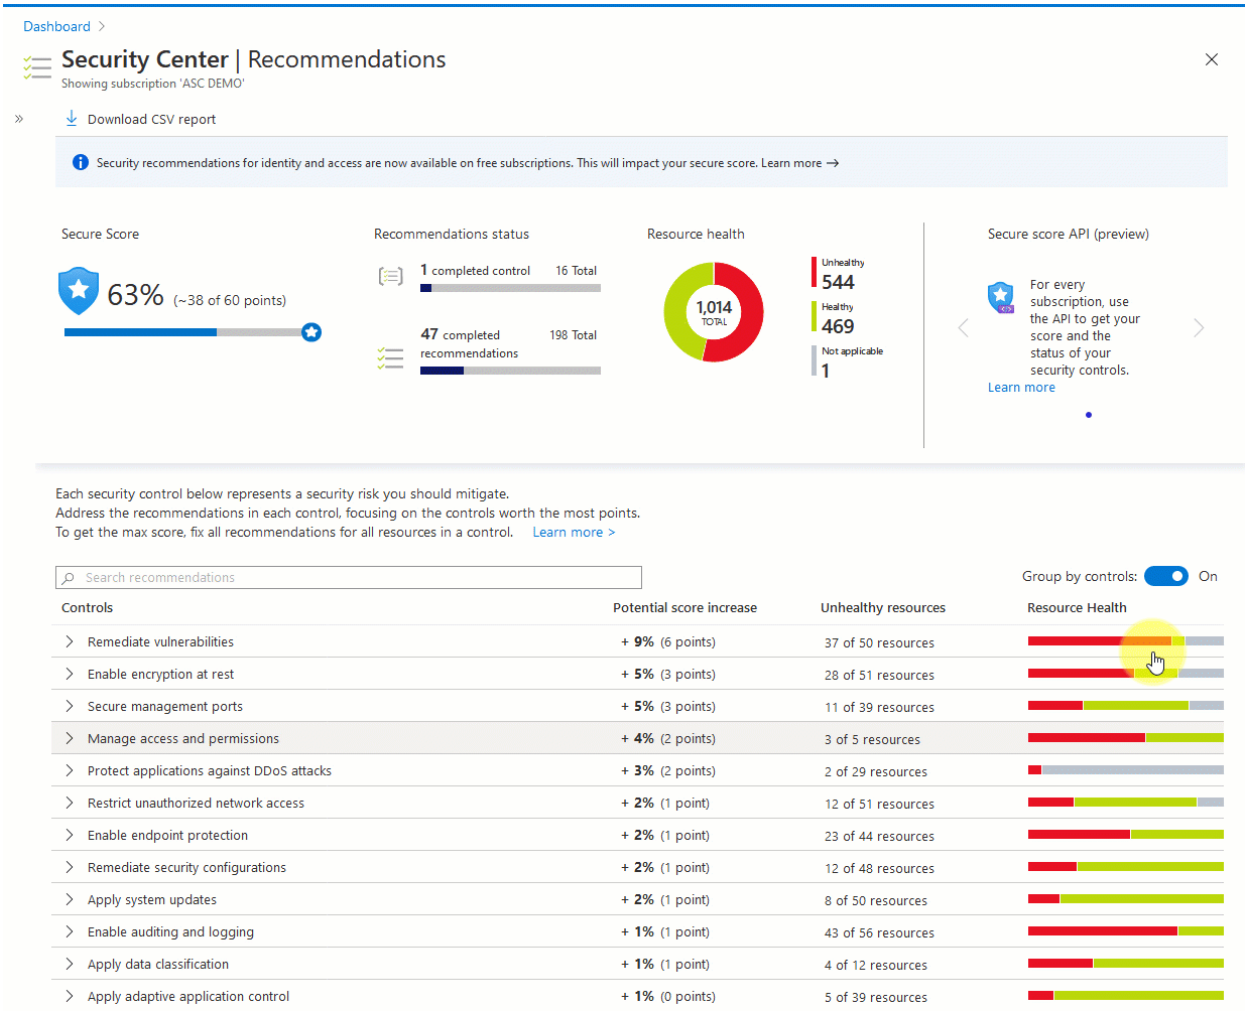

Expanded security control "Implement security best practices"

One security control introduced with the enhanced secure score is "Implement security best practices". When a recommendation is in this control, it doesn't impact the secure score.

With this update, three recommendations have moved out of the controls in which they were originally placed, and into this best practices control. We've taken this step because we've determined that the risk of these three recommendations is lower than was initially thought.

In addition, two new recommendations have been introduced and added to this control.

The three recommendations that moved are:

- **MFA should be enabled on accounts with read permissions on your subscription** (originally in the "Enable MFA" control)
- **External accounts with read permissions should be removed from your subscription** (originally in the "Manage access and permissions" control)
- **A maximum of 3 owners should be designated for your subscription** (originally in the "Manage access and permissions" control)

The two new recommendations added to the control are:

- **Guest configuration extension should be installed on Windows virtual machines (Preview)** - Using Azure Policy Guest Configuration provides visibility inside virtual machines to server and application settings (Windows only).

- **Windows Defender Exploit Guard should be enabled on your machines (Preview)** - Windows Defender Exploit Guard leverages the Azure Policy Guest Configuration agent. Exploit Guard has four components that are designed to lock down devices against a wide variety of attack vectors and block behaviors commonly used in malware attacks while enabling enterprises to balance their security risk

and productivity requirements (Windows only).

Learn more about Windows Defender Exploit Guard in Create and deploy an Exploit Guard policy.

Learn more about security controls in Enhanced secure score (preview).

Custom policies with custom metadata are now generally available

Custom policies are now part of the Security Center recommendations experience, secure score, and the regulatory compliance standards dashboard. This feature is now generally available and allows you to extend your organization's security assessment coverage in Security Center.

Create a custom initiative in Azure Policy, add policies to it and onboard it to Azure Security Center, and visualize it as recommendations.

We've now also added the option to edit the custom recommendation metadata. Metadata options include severity, remediation steps, threats information, and more.

Learn more about enhancing your custom recommendations with detailed information.

Crash dump analysis capabilities migrating to fileless attack detection

We are integrating the Windows crash dump analysis (CDA) detection capabilities into fileless attack detection. Fileless attack detection analytics brings improved versions of the following security alerts for Windows machines: Code injection discovered, Masquerading Windows Module Detected, Shell code discovered, and Suspicious code segment detected.

Some of the benefits of this transition:

- **Proactive and timely malware detection** - The CDA approach involved waiting for a crash to occur and then running analysis to find malicious artifacts. Using fileless attack detection brings proactive identification of in-memory threats while they are running.
- **Enriched alerts** - The security alerts from fileless attack detection include enrichments that aren't available from CDA, such as the active network connections information.
- **Alert aggregation** - When CDA detected multiple attack patterns within a single crash dump, it triggered multiple security alerts. Fileless attack detection combines all of the identified attack patterns from the same process into a single alert, removing the need to correlate multiple alerts.
- **Reduced requirements on your Log Analytics workspace** - Crash dumps containing potentially sensitive data will no longer be uploaded to your Log Analytics workspace.

April 2020

Updates in April include:

- Dynamic compliance packages are now generally available
- Identity recommendations now included in Azure Security Center free tier

Dynamic compliance packages are now generally available

The Azure Security Center regulatory compliance dashboard now includes **dynamic compliance packages** (now generally available) to track additional industry and regulatory standards.

Dynamic compliance packages can be added to your subscription or management group from the Security Center security policy page. When you've onboarded a standard or benchmark, the standard appears in your regulatory compliance dashboard with all associated compliance data mapped as assessments. A summary report for any of the standards that have been onboarded will be available to download.

Now, you can add standards such as:

- NIST SP 800-53 R4
- SWIFT CSP CSCF-v2020
- UK Official and UK NHS
- Canada Federal PBMM
- **Azure CIS 1.1.0 (new)** (which is a more complete representation of Azure CIS 1.1.0)

In addition, we've recently added the Azure Security Benchmark, the Microsoft-authored Azure-specific guidelines for security and compliance best practices based on common compliance frameworks. Additional standards will be supported in the dashboard as they become available.

Learn more about customizing the set of standards in your regulatory compliance dashboard.

Identity recommendations now included in Azure Security Center free tier

Security recommendations for identity and access on the Azure Security Center free tier are now generally available. This is part of the effort to make the cloud security posture management (CSPM) features free. Until now, these recommendations were only available on the standard pricing tier.

Examples of identity and access recommendations include:

- "Multifactor authentication should be enabled on accounts with owner permissions on your subscription."
- "A maximum of three owners should be designated for your subscription."
- "Deprecated accounts should be removed from your subscription."

If you have subscriptions on the free pricing tier, their secure scores will be impacted by this change because they were never assessed for their identity and access security.

Learn more about identity and access recommendations.

Learn more about Managing multifactor authentication (MFA) enforcement on your subscriptions.

March 2020

Updates in March include:

- Workflow automation is now generally available
- Integration of Azure Security Center with Windows Admin Center
- Protection for Azure Kubernetes Service
- Improved just-in-time experience
- Two security recommendations for web applications deprecated

Workflow automation is now generally available

The workflow automation feature of Azure Security Center is now generally available. Use it to automatically trigger Logic Apps on security alerts and recommendations. In addition, manual triggers are available for alerts and all recommendations that have the quick fix option available.

Every security program includes multiple workflows for incident response. These processes might include notifying relevant stakeholders, launching a change management process, and applying specific remediation steps. Security experts recommend that you automate as many steps of those procedures as you can. Automation reduces overhead and can improve your security by ensuring the process steps are done quickly, consistently, and according to your predefined requirements.

For more information about the automatic and manual Security Center capabilities for running your workflows, see workflow automation.

Learn more about creating Logic Apps.

Integration of Azure Security Center with Windows Admin Center

It's now possible to move your on-premises Windows servers from the Windows Admin Center directly to the Azure Security Center. Security Center then becomes your single pane of glass to view security information for all your Windows Admin Center resources, including on-premises servers, virtual machines, and additional PaaS workloads.

After moving a server from Windows Admin Center to Azure Security Center, you'll be able to:

- View security alerts and recommendations in the Security Center extension of the Windows Admin Center.
- View the security posture and retrieve additional detailed information of your Windows Admin Center managed servers in the Security Center within the Azure portal (or via an API).

Learn more about how to integrate Azure Security Center with Windows Admin Center.

Protection for Azure Kubernetes Service

Azure Security Center is expanding its container security features to protect Azure Kubernetes Service (AKS).

The popular, open-source platform Kubernetes has been adopted so widely that it's now an industry standard for container orchestration. Despite this widespread implementation, there's still a lack of understanding regarding how to secure a Kubernetes environment. Defending the attack surfaces of a containerized application requires expertise to ensuring the infrastructure is configured securely and constantly monitored for potential threats.

The Security Center defense includes:

- **Discovery and visibility** - Continuous discovery of managed AKS instances within the subscriptions registered to Security Center.
- **Security recommendations** - Actionable recommendations to help you comply with security best-practices for AKS. These recommendations are included in your secure score to ensure they're viewed as a part of your organization's security posture. An example of an AKS-related recommendation you might see is "Role-based access control should be used to restrict access to a Kubernetes service cluster".
- **Threat protection** - Through continuous analysis of your AKS deployment, Security Center alerts you to threats and malicious activity detected at the host and AKS cluster level.

Learn more about Azure Kubernetes Services' integration with Security Center.

Learn more about the container security features in Security Center.

Improved just-in-time experience

The features, operation, and UI for Azure Security Center's just-in-time tools that secure your management ports have been enhanced as follows:

- **Justification field** - When requesting access to a virtual machine (VM) through the just-in-time page of the Azure portal, a new optional field is available to enter a justification for the request. Information entered into this field can be tracked in the activity log.
- **Automatic cleanup of redundant just-in-time (JIT) rules** - Whenever you update a JIT policy, a cleanup tool automatically runs to check the validity of your entire ruleset. The tool looks for mismatches between rules in your policy and rules in the NSG. If the cleanup tool finds a mismatch, it determines the cause and, when it's safe to do so, removes built-in rules that aren't needed anymore. The cleaner never deletes rules that you've created.

Learn more about the JIT access feature.

Two security recommendations for web applications deprecated

Two security recommendations related to web applications are being deprecated:

- The rules for web applications on IaaS NSGs should be hardened. (Related policy: The NSGs rules for web applications on IaaS should be hardened)

- Access to App Services should be restricted. (Related policy: Access to App Services should be restricted [preview])

These recommendations will no longer appear in the Security Center list of recommendations. The related policies will no longer be included in the initiative named "Security Center Default".

Learn more about security recommendations.

February 2020

Fileless attack detection for Linux (preview)

As attackers increasing employ stealthier methods to avoid detection, Azure Security Center is extending fileless attack detection for Linux, in addition to Windows. Fileless attacks exploit software vulnerabilities, inject malicious payloads into benign system processes, and hide in memory. These techniques:

- minimize or eliminate traces of malware on disk
- greatly reduce the chances of detection by disk-based malware scanning solutions

To counter this threat, Azure Security Center released fileless attack detection for Windows in October 2018, and has now extended fileless attack detection on Linux as well.

January 2020

Enhanced secure score (preview)

An enhanced version of the secure score feature of Azure Security Center is now available in preview. In this version, multiple recommendations are grouped into Security Controls that better reflect your vulnerable attack surfaces (for example, restrict access to management ports).

Familiarize yourself with the secure score changes during the preview phase and determine other remediations that will help you to further secure your environment.

Learn more about enhanced secure score (preview).

November 2019

Updates in November include:

- Threat Protection for Azure Key Vault in North America regions (preview)
- Threat Protection for Azure Storage includes Malware Reputation Screening
- Workflow automation with Logic Apps (preview)
- Quick Fix for bulk resources generally available
- Scan container images for vulnerabilities (preview)
- Additional regulatory compliance standards (preview)
- Threat Protection for Azure Kubernetes Service (preview)
- Virtual machine vulnerability assessment (preview)
- Advanced data security for SQL servers on Azure Virtual Machines (preview)
- Support for custom policies (preview)
- Extending Azure Security Center coverage with platform for community and partners
- Advanced integrations with export of recommendations and alerts (preview)
- Onboard on-prem servers to Security Center from Windows Admin Center (preview)

Threat Protection for Azure Key Vault in North America Regions (preview)

Azure Key Vault is an essential service for protecting data and improving performance of cloud applications by offering the ability to centrally manage keys, secrets, cryptographic keys and policies in the cloud. Since Azure

Key Vault stores sensitive and business critical data, it requires maximum security for the key vaults and the data stored in them.

Azure Security Center's support for Threat Protection for Azure Key Vault provides an additional layer of security intelligence that detects unusual and potentially harmful attempts to access or exploit key vaults. This new layer of protection allows customers to address threats against their key vaults without being a security expert or manage security monitoring systems. The feature is in public preview in North America Regions.

Threat Protection for Azure Storage includes Malware Reputation Screening

Threat protection for Azure Storage offers new detections powered by Microsoft Threat Intelligence for detecting malware uploads to Azure Storage using hash reputation analysis and suspicious access from an active Tor exit node (an anonymizing proxy). You can now view detected malware across storage accounts using Azure Security Center.

Workflow automation with Logic Apps (preview)

Organizations with centrally managed security and IT/operations implement internal workflow processes to drive required action within the organization when discrepancies are discovered in their environments. In many cases, these workflows are repeatable processes and automation can greatly streamline processes within the organization.

Today we are introducing a new capability in Security Center that allows customers to create automation configurations leveraging Azure Logic Apps and to create policies that will automatically trigger them based on specific ASC findings such as Recommendations or Alerts. Azure Logic App can be configured to do any custom action supported by the vast community of Logic App connectors, or use one of the templates provided by Security Center such as sending an email or opening a ServiceNow™ ticket.

For more information about the automatic and manual Security Center capabilities for running your workflows, see workflow automation.

To learn about creating Logic Apps, see Azure Logic Apps.

Quick Fix for bulk resources generally available

With the many tasks that a user is given as part of Secure Score, the ability to effectively remediate issues across a large fleet can become challenging.

To simplify remediation of security misconfigurations and to be able to quickly remediate recommendations on a bulk of resources and improve your secure score, use Quick Fix remediation.

This operation will allow you to select the resources you want to apply the remediation to and launch a remediation action that will configure the setting on your behalf.

Quick fix is generally available today customers as part of the Security Center recommendations page.

See which recommendations have quick fix enabled in the reference guide to security recommendations.

Scan container images for vulnerabilities (preview)

Azure Security Center can now scan container images in Azure Container Registry for vulnerabilities.

The image scanning works by parsing the container image file, then checking to see whether there are any known vulnerabilities (powered by Qualys).

The scan itself is automatically triggered when pushing new container images to Azure Container Registry. Found vulnerabilities will surface as Security Center recommendations and included in the secure score together with information on how to patch them to reduce the attack surface they allowed.

Additional regulatory compliance standards (preview)

The Regulatory Compliance dashboard provides insights into your compliance posture based on Security Center

assessments. The dashboard shows how your environment complies with controls and requirements designated by specific regulatory standards and industry benchmarks and provides prescriptive recommendations for how to address these requirements.

The regulatory compliance dashboard has thus far supported four built-in standards: Azure CIS 1.1.0, PCI-DSS, ISO 27001, and SOC-TSP. We are now announcing the public preview release of additional supported standards: NIST SP 800-53 R4, SWIFT CSP CSCF v2020, Canada Federal PBMM and UK Official together with UK NHS. We are also releasing an updated version of Azure CIS 1.1.0, covering more controls from the standard and enhancing extensibility.

Learn more about customizing the set of standards in your regulatory compliance dashboard.

Threat Protection for Azure Kubernetes Service (preview)

Kubernetes is quickly becoming the new standard for deploying and managing software in the cloud. Few people have extensive experience with Kubernetes and many only focuses on general engineering and administration and overlook the security aspect. Kubernetes environment needs to be configured carefully to be secure, making sure no container focused attack surface doors are not left open is exposed for attackers. Security Center is expanding its support in the container space to one of the fastest growing services in Azure - Azure Kubernetes Service (AKS).

The new capabilities in this public preview release include:

- **Discovery & Visibility** - Continuous discovery of managed AKS instances within Security Center's registered subscriptions.
- **Secure Score recommendations** - Actionable items to help customers comply with security best practices for AKS, and increase their secure score. Recommendations include items such as "Role-based access control should be used to restrict access to a Kubernetes Service Cluster".
- **Threat Detection** - Host and cluster-based analytics, such as "A privileged container detected".

Virtual machine vulnerability assessment (preview)

Applications that are installed in virtual machines could often have vulnerabilities that could lead to a breach of the virtual machine. We are announcing that the Security Center standard tier includes built-in vulnerability assessment for virtual machines for no additional fee. The vulnerability assessment, powered by Qualys in the public preview, will allow you to continuously scan all the installed applications on a virtual machine to find vulnerable applications and present the findings in the Security Center portal's experience. Security Center takes care of all deployment operations so that no extra work is required from the user. Going forward we are planning to provide vulnerability assessment options to support our customers' unique business needs.

Learn more about vulnerability assessments for your Azure Virtual Machines.

Advanced data security for SQL servers on Azure Virtual Machines (preview)

Azure Security Center's support for threat protection and vulnerability assessment for SQL DBs running on IaaS VMs is now in preview.

Vulnerability assessment is an easy to configure service that can discover, track, and help you remediate potential database vulnerabilities. It provides visibility into your security posture as part of secure score and includes the steps to resolve security issues and enhance your database fortifications.

Advanced threat protection detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit your SQL server. It continuously monitors your database for suspicious activities and provides action-oriented security alerts on anomalous database access patterns. These alerts provide the suspicious activity details and recommended actions to investigate and mitigate the threat.

Support for custom policies (preview)

Azure Security Center now supports custom policies (in preview).

Our customers have been wanting to extend their current security assessments coverage in Security Center with their own security assessments based on policies that they create in Azure Policy. With support for custom policies, this is now possible.

These new policies will be part of the Security Center recommendations experience, Secure Score, and the regulatory compliance standards dashboard. With the support for custom policies, you're now able to create a custom initiative in Azure Policy, then add it as a policy in Security Center and visualize it as a recommendation.

Extending Azure Security Center coverage with platform for community and partners

Use Security Center to receive recommendations not only from Microsoft but also from existing solutions from partners such as Check Point, Tenable, and CyberArk with many more integrations coming. Security Center's simple onboarding flow can connect your existing solutions to Security Center, enabling you to view your security posture recommendations in a single place, run unified reports and leverage all of Security Center's capabilities against both built-in and partner recommendations. You can also export Security Center recommendations to partner products.

Learn more about Microsoft Intelligent Security Association.

Advanced integrations with export of recommendations and alerts (preview)

In order to enable enterprise level scenarios on top of Security Center, it's now possible to consume Security Center alerts and recommendations in additional places except the Azure portal or API. These can be directly exported to an Event Hub and to Log Analytics workspaces. Here are a few workflows you can create around these new capabilities:

- With export to Log Analytics workspace, you can create custom dashboards with Power BI.
- With export to Event Hub, you'll be able to export Security Center alerts and recommendations to your third-party SIEMs, to a third-party solution, or Azure Data Explorer.

Onboard on-prem servers to Security Center from Windows Admin Center (preview)

Windows Admin Center is a management portal for Windows Servers who are not deployed in Azure offering them several Azure management capabilities such as backup and system updates. We have recently added an ability to onboard these non-Azure servers to be protected by ASC directly from the Windows Admin Center experience.

With this new experience users will be to onboard a WAC server to Azure Security Center and enable viewing its security alerts and recommendations directly in the Windows Admin Center experience.

September 2019

Updates in September include:

- Managing rules with adaptive application controls improvements
- Control container security recommendation using Azure Policy

Managing rules with adaptive application controls improvements

The experience of managing rules for virtual machines using adaptive application controls has improved. Azure Security Center's adaptive application controls help you control which applications can run on your virtual machines. In addition to a general improvement to rule management, a new benefit enables you to control which file types will be protected when you add a new rule.

Learn more about adaptive application controls.

Control container security recommendation using Azure Policy

Azure Security Center's recommendation to remediate vulnerabilities in container security can now be enabled or disabled via Azure Policy.

To view your enabled security policies, from Security Center open the Security Policypage.

August 2019

Updates in August include:

- Just-in-time (JIT) VM access for Azure Firewall
- Single click remediation to boost your security posture (preview)
- Cross-tenant management

Just-in-time (JIT) VM access for Azure Firewall

Just-in-time (JIT) VM access for Azure Firewall is now generally available. Use it to secure your Azure Firewall protected environments in addition to your NSG protected environments.

JIT VM access reduces exposure to network volumetric attacks by providing controlled access to VMs only when needed, using your NSG and Azure Firewall rules.

When you enable JIT for your VMs, you create a policy that determines the ports to be protected, how long the ports are to remain open, and approved IP addresses from where these ports can be accessed. This policy helps you stay in control of what users can do when they request access.

Requests are logged in the Azure Activity Log, so you can easily monitor and audit access. The just-in-time page also helps you quickly identify existing VMs that have JIT enabled and VMs where JIT is recommended.

Learn more about Azure Firewall.

Single click remediation to boost your security posture (preview)

Secure score is a tool that helps you assess your workload security posture. It reviews your security recommendations and prioritizes them for you, so you know which recommendations to perform first. This helps you find the most serious security vulnerabilities to prioritize investigation.

In order to simplify remediation of security misconfigurations and help you to quickly improve your secure score, we've added a new capability that allows you to remediate a recommendation on a bulk of resources in a single click.

This operation will allow you to select the resources you want to apply the remediation to and launch a remediation action that will configure the setting on your behalf.

See which recommendations have quick fix enabled in the reference guide to security recommendations.

Cross-tenant management

Security Center now supports cross-tenant management scenarios as part of Azure Lighthouse. This enables you to gain visibility and manage the security posture of multiple tenants in Security Center.

Learn more about cross-tenant management experiences.

July 2019

Updates to network recommendations

Azure Security Center (ASC) has launched new networking recommendations and improved some existing ones. Now, using Security Center ensures even greater networking protection for your resources.

Learn more about network recommendations.

June 2019

Adaptive Network Hardening - generally available

One of the biggest attack surfaces for workloads running in the public cloud are connections to and from the public Internet. Our customers find it hard to know which Network Security Group (NSG) rules should be in place to make sure that Azure workloads are only available to required source ranges. With this feature, Security Center learns the network traffic and connectivity patterns of Azure workloads and provides NSG rule recommendations, for Internet facing virtual machines. This helps our customer better configure their network access policies and limit their exposure to attacks.

Learn more about adaptive network hardening.

Microsoft Defender for Cloud data security

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

To help customers prevent, detect, and respond to threats, Microsoft Defender for Cloud collects and processes security-related data, including configuration information, metadata, event logs, and more. Microsoft adheres to strict compliance and security guidelines—from coding to operating a service.

This article explains how data is managed and safeguarded in Defender for Cloud.

Data sources

Defender for Cloud analyzes data from the following sources to provide visibility into your security state, identify vulnerabilities and recommend mitigations, and detect active threats:

- **Azure services**: Uses information about the configuration of Azure services you have deployed by communicating with that service's resource provider.
- **Network traffic**: Uses sampled network traffic metadata from Microsoft's infrastructure, such as source/destination IP/port, packet size, and network protocol.
- **Partner solutions**: Uses security alerts from integrated partner solutions, such as firewalls and antimalware solutions.
- **Your machines**: Uses configuration details and information about security events, such as Windows event and audit logs, and syslog messages from your machines.

Data protection

Data segregation

Data is kept logically separate on each component throughout the service. All data is tagged per organization. This tagging persists throughout the data lifecycle, and it is enforced at each layer of the service.

Data access

To provide security recommendations and investigate potential security threats, Microsoft personnel may access information collected or analyzed by Azure services, including process creation events, and other artifacts, which may unintentionally include customer data or personal data from your machines.

We adhere to the Microsoft Online Services Data Protection Addendum, which states that Microsoft will not use Customer Data or derive information from it for any advertising or similar commercial purposes. We only use Customer Data as needed to provide you with Azure services, including purposes compatible with providing those services. You retain all rights to Customer Data.

Data use

Microsoft uses patterns and threat intelligence seen across multiple tenants to enhance our prevention and detection capabilities; we do so in accordance with the privacy commitments described in our Privacy Statement.

Manage data collection from machines

When you enable Defender for Cloud in Azure, data collection is turned on for each of your Azure subscriptions. You can also enable data collection for your subscriptions in Defender for Cloud. When data collection is enabled, Defender for Cloud provisions the Log Analytics agent on all existing supported Azure virtual machines and any new ones that are created.

The Log Analytics agent scans for various security-related configurations and events it into Event Tracing for Windows (ETW) traces. In addition, the operating system will raise event log events during the course of running the machine. Examples of such data are: operating system type and version, operating system logs (Windows event logs), running processes, machine name, IP addresses, logged in user, and tenant ID. The Log Analytics agent reads event log entries and ETW traces and copies them to your workspace(s) for analysis. The Log Analytics agent also enables process creation events and command line auditing.

If you aren't using Microsoft Defender for Cloud's enhanced security features, you can also disable data collection from virtual machines in the Security Policy. Data Collection is required for subscriptions that are protected by enhanced security features. VM disk snapshots and artifact collection will still be enabled even if data collection has been disabled.

You can specify the workspace and region where data collected from your machines is stored. The default is to store data collected from your machines in the nearest workspace as shown in the following table:

VM GEO	WORKSPACE GEO
United States, Brazil, South Africa	United States
Canada	Canada
Europe (Excluding United Kingdom)	Europe
United Kingdom	United Kingdom
Asia (Excluding India, Japan, Korea, China)	Asia Pacific
Korea	Asia Pacific
India	India
Japan	Japan
China	China
Australia	Australia

NOTE

Microsoft Defender for Storage stores artifacts regionally according to the location of the related Azure resource. Learn more in Introduction to Microsoft Defender for Storage.

Data consumption

Customers can access Defender for Cloud related data from the following data streams:

STREAM	DATA TYPES
Azure Activity log	All security alerts, approved Defender for Cloud just-in-time access requests, and all alerts generated by adaptive application controls.
Azure Monitor logs	All security alerts.
Azure Resource Graph	Security alerts, security recommendations, vulnerability assessment results, secure score information, status of compliance checks, and more.
Microsoft Defender for Cloud REST API	Security alerts, security recommendations, and more.

Next steps

In this document, you learned how data is managed and safeguarded in Microsoft Defender for Cloud.

To learn more about Microsoft Defender for Cloud, see What is Microsoft Defender for Cloud?

Azure Policy built-in definitions for Microsoft Defender for Cloud

2/15/2022 • 86 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This page is an index of Azure Policy built-in policy definitions related to Microsoft Defender for Cloud. The following groupings of policy definitions are available:

- The initiatives group lists the Azure Policy initiative definitions in the "Defender for Cloud" category.
- The default initiative group lists all the Azure Policy definitions that are part of Defender for Cloud's default initiative, Azure Security Benchmark. This Microsoft-authored, widely respected benchmark builds on controls from the Center for Internet Security (CIS) and the National Institute of Standards and Technology (NIST) with a focus on cloud-centric security.
- The category group lists all the Azure Policy definitions in the "Defender for Cloud" category.

For more information about security policies, see Working with security policies. For additional Azure Policy built-ins for other services, see Azure Policy built-in definitions.

The name of each built-in policy definition links to the policy definition in the Azure portal. Use the link in the **Version** column to view the source on the Azure Policy GitHub repo.

Microsoft Defender for Cloud initiatives

To learn about the built-in initiatives that are monitored by Defender for Cloud, see the following table:

NAME	DESCRIPTION	POLICIES	VERSION
Azure Security Benchmark	The Azure Security Benchmark initiative represents the policies and controls implementing security recommendations defined in Azure Security Benchmark v2, see https://aka.ms/azsecbm. This also serves as the Azure Security Center default policy initiative. You can directly assign this initiative, or manage its policies and compliance results within Azure Security Center.	205	45.0.0

NAME	DESCRIPTION	POLICIES	VERSION
Configure Advanced Threat Protection to be enabled on open-source relational databases	Enable Advanced Threat Protection on your non-Basic tier open-source relational databases to detect anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. See https://aka.ms/AzDforOpenSourceDBsDocu.	3	1.0.0
Configure Azure Defender to be enabled on SQL Servers and SQL Managed Instances	Enable Azure Defender on your SQL Servers and SQL Managed Instances to detect anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases.	2	2.0.0

Defender for Cloud's default initiative (Azure Security Benchmark)

To learn about the built-in policies that are monitored by Defender for Cloud, see the following table:

POLICY NAME (AZURE PORTAL)	DESCRIPTION	EFFECT(S)	VERSION (GITHUB)
[Preview]: All Internet traffic should be routed via your deployed Azure Firewall	Azure Security Center has identified that some of your subnets aren't protected with a next generation firewall. Protect your subnets from potential threats by restricting access to them with Azure Firewall or a supported next generation firewall	AuditIfNotExists, Disabled	3.0.0-preview
[Preview]: Azure Arc enabled Kubernetes clusters should have Azure Defender's extension installed	Azure Defender's extension for Azure Arc provides threat protection for your Arc enabled Kubernetes clusters. The extension collects data from nodes in the cluster and sends it to the Azure Defender for Kubernetes backend in the cloud for further analysis. Learn more in https://docs.microsoft.com/azure/security-center/defender-for-kubernetes-azure-arc.	AuditIfNotExists, Disabled	4.0.0-preview

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: Azure Arc enabled Kubernetes clusters should have the Azure Policy extension installed	The Azure Policy extension for Azure Arc provides at-scale enforcements and safeguards on your Arc enabled Kubernetes clusters in a centralized, consistent manner. Learn more at https://aka.ms/akspolicydoc.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: Azure Key Vault should disable public network access	Disable public network access for your key vault so that it's not accessible over the public internet. This can reduce data leakage risks. Learn more at: https://aka.ms/akvprivatelink.	Audit, Deny, Disabled	2.0.0-preview
[Preview]: Azure Kubernetes Service clusters should have Defender profile enabled	Microsoft Defender for Containers provides cloud-native Kubernetes security capabilities including environment hardening, workload protection, and run-time protection. When you enable the SecurityProfile.AzureDefender on your Azure Kubernetes Service cluster, an agent is deployed to your cluster to collect security event data. Learn more about Microsoft Defender for Containers in https://docs.microsoft.com/azure/security-center/defender-for-kubernetes-introduction	Audit, Disabled	1.0.1-preview
[Preview]: Certificates should have the specified maximum validity period	Manage your organizational compliance requirements by specifying the maximum amount of time that a certificate can be valid within your key vault.	audit, deny, disabled	2.1.0-preview
[Preview]: Guest Attestation extension should be installed on supported Linux virtual machines	Install Guest Attestation extension on supported Linux virtual machines to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machines.	AuditIfNotExists, Disabled	5.0.0-preview

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: Guest Attestation extension should be installed on supported Linux virtual machines scale sets	Install Guest Attestation extension on supported Linux virtual machines scale sets to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machine scale sets.	AuditIfNotExists, Disabled	4.0.0-preview
[Preview]: Guest Attestation extension should be installed on supported Windows virtual machines	Install Guest Attestation extension on supported virtual machines to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machines.	AuditIfNotExists, Disabled	3.0.0-preview
[Preview]: Guest Attestation extension should be installed on supported Windows virtual machines scale sets	Install Guest Attestation extension on supported virtual machines scale sets to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machine scale sets.	AuditIfNotExists, Disabled	2.0.0-preview

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: Kubernetes clusters should gate deployment of vulnerable images	Protect your Kubernetes clusters and container workloads from potential threats by restricting deployment of container images with vulnerable software components. Use Azure Defender CI/CD scanning (https://aka.ms/AzureDefenderCICDscanning) and Azure defender for container registries (https://aka.ms/AzureDefenderForContainerRegistries) to identify and patch vulnerabilities prior to deployment. Evaluation prerequisite: Policy Addon and Azure Defender Profile. Only applicable for private preview customers.	Audit, Deny, Disabled	1.0.2-preview
[Preview]: Log Analytics extension should be installed on your Linux Azure Arc machines	This policy audits Linux Azure Arc machines if the Log Analytics extension is not installed.	AuditIfNotExists, Disabled	1.0.1-preview
[Preview]: Log Analytics extension should be installed on your Windows Azure Arc machines	This policy audits Windows Azure Arc machines if the Log Analytics extension is not installed.	AuditIfNotExists, Disabled	1.0.1-preview
[Preview]: Network traffic data collection agent should be installed on Linux virtual machines	Security Center uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations and specific network threats.	AuditIfNotExists, Disabled	1.0.2-preview
[Preview]: Network traffic data collection agent should be installed on Windows virtual machines	Security Center uses the Microsoft Dependency agent to collect network traffic data from your Azure virtual machines to enable advanced network protection features such as traffic visualization on the network map, network hardening recommendations and specific network threats.	AuditIfNotExists, Disabled	1.0.2-preview

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: Private endpoint should be configured for Key Vault	Private link provides a way to connect Key Vault to your Azure resources without sending traffic over the public internet. Private link provides defense in depth protection against data exfiltration.	Audit, Deny, Disabled	1.1.0-preview
[Preview]: Secure Boot should be enabled on supported Windows virtual machines	Enable Secure Boot on supported Windows virtual machines to mitigate against malicious and unauthorized changes to the boot chain. Once enabled, only trusted bootloaders, kernel and kernel drivers will be allowed to run. This assessment only applies to trusted launch enabled Windows virtual machines.	Audit, Disabled	3.0.0-preview
[Preview]: Storage account public access should be disallowed	Anonymous public read access to containers and blobs in Azure Storage is a convenient way to share data but might present security risks. To prevent data breaches caused by undesired anonymous access, Microsoft recommends preventing public access to a storage account unless your scenario requires it.	audit, deny, disabled	3.0.1-preview
[Preview]: vTPM should be enabled on supported virtual machines	Enable virtual TPM device on supported virtual machines to facilitate Measured Boot and other OS security features that require a TPM. Once enabled, vTPM can be used to attest boot integrity. This assessment only applies to trusted launch enabled virtual machines.	Audit, Disabled	2.0.0-preview
A maximum of 3 owners should be designated for your subscription	It is recommended to designate up to 3 subscription owners in order to reduce the potential for breach by a compromised owner.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
A vulnerability assessment solution should be enabled on your virtual machines	Audits virtual machines to detect whether they are running a supported vulnerability assessment solution. A core component of every cyber risk and security program is the identification and analysis of vulnerabilities. Azure Security Center's standard pricing tier includes vulnerability scanning for your virtual machines at no extra cost. Additionally, Security Center can automatically deploy this tool for you.	AuditIfNotExists, Disabled	3.0.0
Adaptive application controls for defining safe applications should be enabled on your machines	Enable application controls to define the list of known-safe applications running on your machines, and alert you when other applications run. This helps harden your machines against malware. To simplify the process of configuring and maintaining your rules, Security Center uses machine learning to analyze the applications running on each machine and suggest the list of known-safe applications.	AuditIfNotExists, Disabled	3.0.0
Adaptive network hardening recommendations should be applied on internet facing virtual machines	Azure Security Center analyzes the traffic patterns of Internet facing virtual machines and provides Network Security Group rule recommendations that reduce the potential attack surface	AuditIfNotExists, Disabled	3.0.0
All network ports should be restricted on network security groups associated to your virtual machine	Azure Security Center has identified some of your network security groups' inbound rules to be too permissive. Inbound rules should not allow access from 'Any' or 'Internet' ranges. This can potentially enable attackers to target your resources.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Allowlist rules in your adaptive application control policy should be updated	Monitor for changes in behavior on groups of machines configured for auditing by Azure Security Center's adaptive application controls. Security Center uses machine learning to analyze the running processes on your machines and suggest a list of known-safe applications. These are presented as recommended apps to allow in adaptive application control policies.	AuditIfNotExists, Disabled	3.0.0
An Azure Active Directory administrator should be provisioned for SQL servers	Audit provisioning of an Azure Active Directory administrator for your SQL server to enable Azure AD authentication. Azure AD authentication enables simplified permission management and centralized identity management of database users and other Microsoft services	AuditIfNotExists, Disabled	1.0.0
API App should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks.	Audit, Disabled	1.0.0
API Management services should use a virtual network	Azure Virtual Network deployment provides enhanced security, isolation and allows you to place your API Management service in a non-internet routable network that you control access to. These networks can then be connected to your on-premises networks using various VPN technologies, which enables access to your backend services within the network and/or on-premises. The developer portal and API gateway, can be configured to be accessible either from the Internet or only within the virtual network.	Audit, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
App Configuration should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your app configuration instances instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/appconfig/private-endpoint.	AuditIfNotExists, Disabled	1.0.2
Audit usage of custom RBAC rules	Audit built-in roles such as 'Owner, Contributer, Reader' instead of custom RBAC roles, which are error prone. Using custom roles is treated as an exception and requires a rigorous review and threat modeling	Audit, Disabled	1.0.0
Auditing on SQL server should be enabled	Auditing on your SQL Server should be enabled to track database activities across all databases on the server and save them in an audit log.	AuditIfNotExists, Disabled	2.0.0
Authentication to Linux machines should require SSH keys	Although SSH itself provides an encrypted connection, using passwords with SSH still leaves the VM vulnerable to brute-force attacks. The most secure option for authenticating to an Azure Linux virtual machine over SSH is with a public-private key pair, also known as SSH keys. Learn more: https://docs.microsoft.com/azure/virtual-machines/linux/create-ssh-keys-detailed.	AuditIfNotExists, Disabled	2.2.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Authorized IP ranges should be defined on Kubernetes Services	Restrict access to the Kubernetes Service Management API by granting API access only to IP addresses in specific ranges. It is recommended to limit access to authorized IP ranges to ensure that only applications from allowed networks can access the cluster.	Audit, Disabled	2.0.1
Auto provisioning of the Log Analytics agent should be enabled on your subscription	To monitor for security vulnerabilities and threats, Azure Security Center collects data from your Azure virtual machines. Data is collected by the Log Analytics agent, formerly known as the Microsoft Monitoring Agent (MMA), which reads various security-related configurations and event logs from the machine and copies the data to your Log Analytics workspace for analysis. We recommend enabling auto provisioning to automatically deploy the agent to all supported Azure VMs and any new ones that are created.	AuditIfNotExists, Disabled	1.0.1
Automation account variables should be encrypted	It is important to enable encryption of Automation account variable assets when storing sensitive data	Audit, Deny, Disabled	1.1.0
Azure Backup should be enabled for Virtual Machines	Ensure protection of your Azure Virtual Machines by enabling Azure Backup. Azure Backup is a secure and cost effective data protection solution for Azure.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Cache for Redis should reside within a virtual network	Azure Virtual Network deployment provides enhanced security and isolation for your Azure Cache for Redis, as well as subnets, access control policies, and other features to further restrict access.When an Azure Cache for Redis instance is configured with a virtual network, it is not publicly addressable and can only be accessed from virtual machines and applications within the virtual network.	Audit, Deny, Disabled	1.0.3
Azure Cosmos DB accounts should have firewall rules	Firewall rules should be defined on your Azure Cosmos DB accounts to prevent traffic from unauthorized sources. Accounts that have at least one IP rule defined with the virtual network filter enabled are deemed compliant. Accounts disabling public access are also deemed compliant.	Audit, Deny, Disabled	2.0.0
Azure Cosmos DB accounts should use customer-managed keys to encrypt data at rest	Use customer-managed keys to manage the encryption at rest of your Azure Cosmos DB. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more at https://aka.ms/cosmosdb-cmk.	audit, deny, disabled	1.0.2
Azure DDoS Protection Standard should be enabled	DDoS protection standard should be enabled for all virtual networks with a subnet that is part of an application gateway with a public IP.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for App Service should be enabled	Azure Defender for App Service leverages the scale of the cloud, and the visibility that Azure has as a cloud provider, to monitor for common web app attacks.	AuditIfNotExists, Disabled	1.0.3
Azure Defender for Azure SQL Database servers should be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	AuditIfNotExists, Disabled	1.0.2
Azure Defender for DNS should be enabled	Azure Defender for DNS provides an additional layer of protection for your cloud resources by continuously monitoring all DNS queries from your Azure resources. Azure Defender alerts you about suspicious activity at the DNS layer. Learn more about the capabilities of Azure Defender for DNS at https://aka.ms/defender-for-dns . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	AuditIfNotExists, Disabled	1.0.0
Azure Defender for Key Vault should be enabled	Azure Defender for Key Vault provides an additional layer of protection and security intelligence by detecting unusual and potentially harmful attempts to access or exploit key vault accounts.	AuditIfNotExists, Disabled	1.0.3

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for open-source relational databases should be enabled	Azure Defender for open-source relational databases detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Learn more about the capabilities of Azure Defender for open-source relational databases at https://aka.ms/AzDforOpenSourceDBsDocu. Important: Enabling this plan will result in charges for protecting your open-source relational databases. Learn about the pricing on Security Center's pricing page: https://aka.ms/pricing-security-center	AuditIfNotExists, Disabled	1.0.0
Azure Defender for Resource Manager should be enabled	Azure Defender for Resource Manager automatically monitors the resource management operations in your organization. Azure Defender detects threats and alerts you about suspicious activity. Learn more about the capabilities of Azure Defender for Resource Manager at https://aka.ms/defender-for-resource-manager . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	AuditIfNotExists, Disabled	1.0.0
Azure Defender for servers should be enabled	Azure Defender for servers provides real-time threat protection for server workloads and generates hardening recommendations as well as alerts about suspicious activities.	AuditIfNotExists, Disabled	1.0.3

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for SQL servers on machines should be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	AuditIfNotExists, Disabled	1.0.2
Azure Defender for SQL should be enabled for unprotected Azure SQL servers	Audit SQL servers without Advanced Data Security	AuditIfNotExists, Disabled	2.0.1
Azure Defender for SQL should be enabled for unprotected SQL Managed Instances	Audit each SQL Managed Instance without advanced data security.	AuditIfNotExists, Disabled	1.0.2
Azure Defender for Storage should be enabled	Azure Defender for Storage provides detections of unusual and potentially harmful attempts to access or exploit storage accounts.	AuditIfNotExists, Disabled	1.0.3
Azure Event Grid domains should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The Private Link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your Event Grid domain instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/privateendpoints.	Audit, Disabled	1.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Event Grid topics should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The Private Link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your Event Grid topic instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/privateendpoints.	Audit, Disabled	1.0.2
Azure Machine Learning workspaces should be encrypted with a customer-managed key	Manage encryption at rest of Azure Machine Learning workspace data with customer-managed keys. By default, customer data is encrypted with service-managed keys, but customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more at https://aka.ms/azureml-workspaces-cmk.	Audit, Deny, Disabled	1.0.3

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Machine Learning workspaces should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The Private Link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to Azure Machine Learning workspaces, data leakage risks are reduced. Learn more about private links at: https://docs.microsoft.com/azure/machine-learning/how-to-configure-private-link.	Audit, Deny, Disabled	1.1.0
Azure Policy Add-on for Kubernetes service (AKS) should be installed and enabled on your clusters	Azure Policy Add-on for Kubernetes service (AKS) extends Gatekeeper v3, an admission controller webhook for Open Policy Agent (OPA), to apply at-scale enforcements and safeguards on your clusters in a centralized, consistent manner.	Audit, Disabled	1.0.2
Azure SignalR Service should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your Azure SignalR Service resource instead of the entire service, you'll reduce your data leakage risks. Learn more about private links at: https://aka.ms/asrs/privatelink.	Audit, Deny, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Spring Cloud should use network injection	Azure Spring Cloud instances should use virtual network injection for the following purposes: 1. Isolate Azure Spring Cloud from Internet. 2. Enable Azure Spring Cloud to interact with systems in either on premises data centers or Azure service in other virtual networks. 3. Empower customers to control inbound and outbound network communications for Azure Spring Cloud.	Audit, Disabled, Deny	1.1.0
Azure Web Application Firewall should be enabled for Azure Front Door entry-points	Deploy Azure Web Application Firewall (WAF) in front of public facing web applications for additional inspection of incoming traffic. Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities such as SQL injections, Cross-Site Scripting, local and remote file executions. You can also restrict access to your web applications by countries, IP address ranges, and other http(s) parameters via custom rules.	Audit, Deny, Disabled	1.0.2
Cognitive Services accounts should disable public network access	Disabling public network access improves security by ensuring that Cognitive Services account isn't exposed on the public internet. Creating private endpoints can limit exposure of Cognitive Services account. Learn more at: https://go.microsoft.com/fwlink/?linkid=2129800.	Audit, Deny, Disabled	2.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Cognitive Services accounts should enable data encryption with a customer-managed key	Customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data stored in Cognitive Services to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more about customer-managed keys at https://go.microsoft.com/fwlink/?linkid=2121321.	Audit, Deny, Disabled	2.0.0
Cognitive Services accounts should restrict network access	Network access to Cognitive Services accounts should be restricted. Configure network rules so only applications from allowed networks can access the Cognitive Services account. To allow connections from specific internet or on-premises clients, access can be granted to traffic from specific Azure virtual networks or to public internet IP address ranges.	Audit, Deny, Disabled	2.0.0
Container registries should be encrypted with a customer-managed key	Use customer-managed keys to manage the encryption at rest of the contents of your registries. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management. Learn more at https://aka.ms/acr/CMK.	Audit, Deny, Disabled	1.1.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Container registries should not allow unrestricted network access	Azure container registries by default accept connections over the internet from hosts on any network. To protect your registries from potential threats, allow access from only specific public IP addresses or address ranges. If your registry doesn't have an IP/firewall rule or a configured virtual network, it will appear in the unhealthy resources. Learn more about Container Registry network rules here: https://aka.ms/acr/portal/public-network and here https://aka.ms/acr/vnet.	Audit, Deny, Disabled	1.1.0
Container registries should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The private link platform handles the connectivity between the consumer and services over the Azure backbone network.By mapping private endpoints to your container registries instead of the entire service, you'll also be protected against data leakage risks. Learn more at: https://aka.ms/acr/private-link.	Audit, Disabled	1.0.1
Container registry images should have vulnerability findings resolved	Container image vulnerability assessment scans your registry for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks.	AuditIfNotExists, Disabled	2.0.1
CORS should not allow every resource to access your API App	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your API app. Allow only required domains to interact with your API app.	AuditIfNotExists, Disabled	1.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
CORS should not allow every resource to access your Function Apps	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your Function app. Allow only required domains to interact with your Function app.	AuditIfNotExists, Disabled	1.0.0
CORS should not allow every resource to access your Web Applications	Cross-Origin Resource Sharing (CORS) should not allow all domains to access your web application. Allow only required domains to interact with your web app.	AuditIfNotExists, Disabled	1.0.0
Deprecated accounts should be removed from your subscription	Deprecated accounts should be removed from your subscriptions. Deprecated accounts are accounts that have been blocked from signing in.	AuditIfNotExists, Disabled	3.0.0
Deprecated accounts with owner permissions should be removed from your subscription	Deprecated accounts with owner permissions should be removed from your subscription. Deprecated accounts are accounts that have been blocked from signing in.	AuditIfNotExists, Disabled	3.0.0
Email notification for high severity alerts should be enabled	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, enable email notifications for high severity alerts in Security Center.	AuditIfNotExists, Disabled	1.0.1
Email notification to subscription owner for high severity alerts should be enabled	To ensure your subscription owners are notified when there is a potential security breach in their subscription, set email notifications to subscription owners for high severity alerts in Security Center.	AuditIfNotExists, Disabled	2.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Endpoint protection health issues should be resolved on your machines	Resolve endpoint protection health issues on your virtual machines to protect them from latest threats and vulnerabilities. Azure Security Center supported endpoint protection solutions are documented here - https://docs.microsoft.com/azure/security-center/security-center-services?tabs=features-windows#supported-endpoint-protection-solutions. Endpoint protection assessment is documented here - https://docs.microsoft.com/azure/security-center/security-center-endpoint-protection.	AuditIfNotExists, Disabled	1.0.0
Endpoint protection should be installed on your machines	To protect your machines from threats and vulnerabilities, install a supported endpoint protection solution.	AuditIfNotExists, Disabled	1.0.0
Endpoint protection solution should be installed on virtual machine scale sets	Audit the existence and health of an endpoint protection solution on your virtual machines scale sets, to protect them from threats and vulnerabilities.	AuditIfNotExists, Disabled	3.0.0
Enforce SSL connection should be enabled for MySQL database servers	Azure Database for MySQL supports connecting your Azure Database for MySQL server to client applications using Secure Sockets Layer (SSL). Enforcing SSL connections between your database server and your client applications helps protect against 'man in the middle' attacks by encrypting the data stream between the server and your application. This configuration enforces that SSL is always enabled for accessing your database server.	Audit, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Enforce SSL connection should be enabled for PostgreSQL database servers	Azure Database for PostgreSQL supports connecting your Azure Database for PostgreSQL server to client applications using Secure Sockets Layer (SSL). Enforcing SSL connections between your database server and your client applications helps protect against 'man in the middle' attacks by encrypting the data stream between the server and your application. This configuration enforces that SSL is always enabled for accessing your database server.	Audit, Disabled	1.0.1
Ensure API app has 'Client Certificates (Incoming client certificates)' set to 'On'	Client certificates allow for the app to request a certificate for incoming requests. Only clients that have a valid certificate will be able to reach the app.	Audit, Disabled	1.0.0
Ensure that 'Java version' is the latest, if used as a part of the API app	Periodically, newer versions are released for Java either due to security flaws or to include additional functionality. Using the latest Python version for API apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	2.0.0
Ensure that 'Java version' is the latest, if used as a part of the Function app	Periodically, newer versions are released for Java software either due to security flaws or to include additional functionality. Using the latest Java version for Function apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	2.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Ensure that 'Java version' is the latest, if used as a part of the Web app	Periodically, newer versions are released for Java software either due to security flaws or to include additional functionality. Using the latest Java version for web apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	2.0.0
Ensure that 'PHP version' is the latest, if used as a part of the API app	Periodically, newer versions are released for PHP software either due to security flaws or to include additional functionality. Using the latest PHP version for API apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	2.1.0
Ensure that 'PHP version' is the latest, if used as a part of the WEB app	Periodically, newer versions are released for PHP software either due to security flaws or to include additional functionality. Using the latest PHP version for web apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	2.1.0
Ensure that 'Python version' is the latest, if used as a part of the API app	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for API apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Ensure that 'Python version' is the latest, if used as a part of the Function app	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for Function apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	3.0.0
Ensure that 'Python version' is the latest, if used as a part of the Web app	Periodically, newer versions are released for Python software either due to security flaws or to include additional functionality. Using the latest Python version for web apps is recommended in order to take advantage of security fixes, if any, and/or new functionalities of the latest version. Currently, this policy only applies to Linux web apps.	AuditIfNotExists, Disabled	3.0.0
Ensure WEB app has 'Client Certificates (Incoming client certificates)' set to 'On'	Client certificates allow for the app to request a certificate for incoming requests. Only clients that have a valid certificate will be able to reach the app.	Audit, Disabled	1.0.0
External accounts with owner permissions should be removed from your subscription	External accounts with owner permissions should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0
External accounts with read permissions should be removed from your subscription	External accounts with read privileges should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0
External accounts with write permissions should be removed from your subscription	External accounts with write privileges should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
FTPS only should be required in your API App	Enable FTPS enforcement for enhanced security	AuditIfNotExists, Disabled	2.0.0
FTPS only should be required in your Function App	Enable FTPS enforcement for enhanced security	AuditIfNotExists, Disabled	2.0.0
FTPS should be required in your Web App	Enable FTPS enforcement for enhanced security	AuditIfNotExists, Disabled	2.0.0
Function App should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks.	Audit, Disabled	1.0.0
Function apps should have 'Client Certificates (Incoming client certificates)' enabled	Client certificates allow for the app to request a certificate for incoming requests. Only clients with valid certificates will be able to reach the app.	Audit, Disabled	1.0.1
Geo-redundant backup should be enabled for Azure Database for MariaDB	Azure Database for MariaDB allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery option in case of a region failure. Configuring geo-redundant storage for backup is only allowed during server create.	Audit, Disabled	1.0.1
Geo-redundant backup should be enabled for Azure Database for MySQL	Azure Database for MySQL allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery option in case of a region failure. Configuring geo-redundant storage for backup is only allowed during server create.	Audit, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Geo-redundant backup should be enabled for Azure Database for PostgreSQL	Azure Database for PostgreSQL allows you to choose the redundancy option for your database server. It can be set to a geo-redundant backup storage in which the data is not only stored within the region in which your server is hosted, but is also replicated to a paired region to provide recovery option in case of a region failure. Configuring geo-redundant storage for backup is only allowed during server create.	Audit, Disabled	1.0.1
Guest Configuration extension should be installed on your machines	To ensure secure configurations of in-guest settings of your machine, install the Guest Configuration extension. In-guest settings that the extension monitors include the configuration of the operating system, application configuration or presence, and environment settings. Once installed, in-guest policies will be available such as 'Windows Exploit guard should be enabled'. Learn more at https://aka.ms/gcpol.	AuditIfNotExists, Disabled	1.0.2
Internet-facing virtual machines should be protected with network security groups	Protect your virtual machines from potential threats by restricting access to them with network security groups (NSG). Learn more about controlling traffic with NSGs at https://aka.ms/nsg-doc	AuditIfNotExists, Disabled	3.0.0
IP Forwarding on your virtual machine should be disabled	Enabling IP forwarding on a virtual machine's NIC allows the machine to receive traffic addressed to other destinations. IP forwarding is rarely required (e.g., when using the VM as a network virtual appliance), and therefore, this should be reviewed by the network security team.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Key Vault keys should have an expiration date	Cryptographic keys should have a defined expiration date and not be permanent. Keys that are valid forever provide a potential attacker with more time to compromise the key. It is a recommended security practice to set expiration dates on cryptographic keys.	Audit, Deny, Disabled	1.0.2
Key Vault secrets should have an expiration date	Secrets should have a defined expiration date and not be permanent. Secrets that are valid forever provide a potential attacker with more time to compromise them. It is a recommended security practice to set expiration dates on secrets.	Audit, Deny, Disabled	1.0.2
Key vaults should have purge protection enabled	Malicious deletion of a key vault can lead to permanent data loss. A malicious insider in your organization can potentially delete and purge key vaults. Purge protection protects you from insider attacks by enforcing a mandatory retention period for soft deleted key vaults. No one inside your organization or Microsoft will be able to purge your key vaults during the soft delete retention period.	Audit, Deny, Disabled	2.0.0
Key vaults should have soft delete enabled	Deleting a key vault without soft delete enabled permanently deletes all secrets, keys, and certificates stored in the key vault. Accidental deletion of a key vault can lead to permanent data loss. Soft delete allows you to recover an accidentally deleted key vault for a configurable retention period.	Audit, Deny, Disabled	2.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes cluster containers CPU and memory resource limits should not exceed the specified limits	Enforce container CPU and memory resource limits to prevent resource exhaustion attacks in a Kubernetes cluster. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	7.0.1
Kubernetes cluster containers should not share host process ID or host IPC namespace	Block pod containers from sharing the host process ID namespace and host IPC namespace in a Kubernetes cluster. This recommendation is part of CIS 5.2.2 and CIS 5.2.3 which are intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	3.0.2
Kubernetes cluster containers should only use allowed AppArmor profiles	Containers should only use allowed AppArmor profiles in a Kubernetes cluster. This recommendation is part of Pod Security Policies which are intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.3

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes cluster containers should only use allowed capabilities	Restrict the capabilities to reduce the attack surface of containers in a Kubernetes cluster. This recommendation is part of CIS 5.2.8 and CIS 5.2.9 which are intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.2
Kubernetes cluster containers should only use allowed images	Use images from trusted registries to reduce the Kubernetes cluster's exposure risk to unknown vulnerabilities, security issues and malicious images. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	7.0.4
Kubernetes cluster containers should run with a read only root file system	Run containers with a read only root file system to protect from changes at run-time with malicious binaries being added to PATH in a Kubernetes cluster. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes cluster pod hostPath volumes should only use allowed host paths	Limit pod HostPath volume mounts to the allowed host paths in a Kubernetes Cluster. This recommendation is part of Pod Security Policies which are intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.2
Kubernetes cluster pods and containers should only run with approved user and group IDs	Control the user, primary group, supplemental group and file system group IDs that pods and containers can use to run in a Kubernetes Cluster. This recommendation is part of Pod Security Policies which are intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.3
Kubernetes cluster pods should only use approved host network and port range	Restrict pod access to the host network and the allowable host port range in a Kubernetes cluster. This recommendation is part of CIS 5.2.4 which is intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes cluster services should listen only on allowed ports	Restrict services to listen only on allowed ports to secure access to the Kubernetes cluster. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	6.1.2
Kubernetes cluster should not allow privileged containers	Do not allow privileged containers creation in a Kubernetes cluster. This recommendation is part of CIS 5.2.1 which is intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	7.0.1
Kubernetes clusters should be accessible only over HTTPS	Use of HTTPS ensures authentication and protects data in transit from network layer eavesdropping attacks. This capability is currently generally available for Kubernetes Service (AKS), and in preview for AKS Engine and Azure Arc enabled Kubernetes. For more info, visit https://aka.ms/kubepolicydoc	audit, deny, disabled	6.0.1
Kubernetes clusters should disable automounting API credentials	Disable automounting API credentials to prevent a potentially compromised Pod resource to run API commands against Kubernetes clusters. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	2.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Kubernetes clusters should not allow container privilege escalation	Do not allow containers to run with privilege escalation to root in a Kubernetes cluster. This recommendation is part of CIS 5.2.5 which is intended to improve the security of your Kubernetes environments. This policy is generally available for Kubernetes Service (AKS), and preview for AKS Engine and Azure Arc enabled Kubernetes. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	4.0.1
Kubernetes clusters should not grant CAP_SYS_ADMIN security capabilities	To reduce the attack surface of your containers, restrict CAP_SYS_ADMIN Linux capabilities. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	3.0.2
Kubernetes clusters should not use the default namespace	Prevent usage of the default namespace in Kubernetes clusters to protect against unauthorized access for ConfigMap, Pod, Secret, Service, and ServiceAccount resource types. For more information, see https://aka.ms/kubepolicydoc.	audit, deny, disabled	2.1.2
Latest TLS version should be used in your API App	Upgrade to the latest TLS version	AuditIfNotExists, Disabled	1.0.0
Latest TLS version should be used in your Function App	Upgrade to the latest TLS version	AuditIfNotExists, Disabled	1.0.0
Latest TLS version should be used in your Web App	Upgrade to the latest TLS version	AuditIfNotExists, Disabled	1.0.0
Linux machines should have Log Analytics agent installed on Azure Arc	Machines are non-compliant if Log Analytics agent is not installed on Azure Arc enabled Linux server.	AuditIfNotExists, Disabled	1.1.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Linux machines should meet requirements for the Azure compute security baseline	Requires that prerequisites are deployed to the policy assignment scope. For details, visit https://aka.ms/gcpol. Machines are non-compliant if the machine is not configured correctly for one of the recommendations in the Azure compute security baseline.	AuditIfNotExists, Disabled	1.3.0
Log Analytics agent should be installed on your virtual machine for Azure Security Center monitoring	This policy audits any Windows/Linux virtual machines (VMs) if the Log Analytics agent is not installed which Security Center uses to monitor for security vulnerabilities and threats	AuditIfNotExists, Disabled	1.0.0
Log Analytics agent should be installed on your virtual machine scale sets for Azure Security Center monitoring	Security Center collects data from your Azure virtual machines (VMs) to monitor for security vulnerabilities and threats.	AuditIfNotExists, Disabled	1.0.0
Managed identity should be used in your API App	Use a managed identity for enhanced authentication security	AuditIfNotExists, Disabled	2.0.0
Managed identity should be used in your Function App	Use a managed identity for enhanced authentication security	AuditIfNotExists, Disabled	2.0.0
Managed identity should be used in your Web App	Use a managed identity for enhanced authentication security	AuditIfNotExists, Disabled	2.0.0
Management ports of virtual machines should be protected with just-in-time network access control	Possible network Just In Time (JIT) access will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
Management ports should be closed on your virtual machines	Open remote management ports are exposing your VM to a high level of risk from Internet-based attacks. These attacks attempt to brute force credentials to gain admin access to the machine.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
MFA should be enabled accounts with write permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with write privileges to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0
MFA should be enabled on accounts with owner permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with owner permissions to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0
MFA should be enabled on accounts with read permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with read privileges to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0
Microsoft Defender for Containers should be enabled	Microsoft Defender for Containers provides hardening, vulnerability assessment and run-time protections for your Azure, hybrid, and multi-cloud Kubernetes environments.	AuditIfNotExists, Disabled	1.0.0
Monitor missing Endpoint Protection in Azure Security Center	Servers without an installed Endpoint Protection agent will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
MySQL servers should use customer-managed keys to encrypt data at rest	Use customer-managed keys to manage the encryption at rest of your MySQL servers. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management.	AuditIfNotExists, Disabled	1.0.4

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Network Watcher should be enabled	Network Watcher is a regional service that enables you to monitor and diagnose conditions at a network scenario level in, to, and from Azure. Scenario level monitoring enables you to diagnose problems at an end to end network level view. It is required to have a network watcher resource group to be created in every region where a virtual network is present. An alert is enabled if a network watcher resource group is not available in a particular region.	AuditIfNotExists, Disabled	3.0.0
Non-internet-facing virtual machines should be protected with network security groups	Protect your non-internet-facing virtual machines from potential threats by restricting access with network security groups (NSG). Learn more about controlling traffic with NSGs at https://aka.ms/nsg-doc	AuditIfNotExists, Disabled	3.0.0
Only secure connections to your Azure Cache for Redis should be enabled	Audit enabling of only connections via SSL to Azure Cache for Redis. Use of secure connections ensures authentication between the server and the service and protects data in transit from network layer attacks such as man-in-the-middle, eavesdropping, and session-hijacking	Audit, Deny, Disabled	1.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
PostgreSQL servers should use customer-managed keys to encrypt data at rest	Use customer-managed keys to manage the encryption at rest of your PostgreSQL servers. By default, the data is encrypted at rest with service-managed keys, but customer-managed keys are commonly required to meet regulatory compliance standards. Customer-managed keys enable the data to be encrypted with an Azure Key Vault key created and owned by you. You have full control and responsibility for the key lifecycle, including rotation and management.	AuditIfNotExists, Disabled	1.0.4
Private endpoint connections on Azure SQL Database should be enabled	Private endpoint connections enforce secure communication by enabling private connectivity to Azure SQL Database.	Audit, Disabled	1.1.0
Private endpoint should be enabled for MariaDB servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for MariaDB. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure.	AuditIfNotExists, Disabled	1.0.2
Private endpoint should be enabled for MySQL servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for MySQL. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure.	AuditIfNotExists, Disabled	1.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Private endpoint should be enabled for PostgreSQL servers	Private endpoint connections enforce secure communication by enabling private connectivity to Azure Database for PostgreSQL. Configure a private endpoint connection to enable access to traffic coming only from known networks and prevent access from all other IP addresses, including within Azure.	AuditIfNotExists, Disabled	1.0.2
Public network access on Azure SQL Database should be disabled	Disabling the public network access property improves security by ensuring your Azure SQL Database can only be accessed from a private endpoint. This configuration denies all logins that match IP or virtual network based firewall rules.	Audit, Deny, Disabled	1.1.0
Public network access should be disabled for MariaDB servers	Disable the public network access property to improve security and ensure your Azure Database for MariaDB can only be accessed from a private endpoint. This configuration strictly disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules.	Audit, Disabled	1.0.2
Public network access should be disabled for MySQL servers	Disable the public network access property to improve security and ensure your Azure Database for MySQL can only be accessed from a private endpoint. This configuration strictly disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules.	Audit, Disabled	1.0.2

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Public network access should be disabled for PostgreSQL servers	Disable the public network access property to improve security and ensure your Azure Database for PostgreSQL can only be accessed from a private endpoint. This configuration disables access from any public address space outside of Azure IP range, and denies all logins that match IP or virtual network-based firewall rules.	Audit, Disabled	1.0.2
Remote debugging should be turned off for API Apps	Remote debugging requires inbound ports to be opened on API apps. Remote debugging should be turned off.	AuditIfNotExists, Disabled	1.0.0
Remote debugging should be turned off for Function Apps	Remote debugging requires inbound ports to be opened on function apps. Remote debugging should be turned off.	AuditIfNotExists, Disabled	1.0.0
Remote debugging should be turned off for Web Applications	Remote debugging requires inbound ports to be opened on a web application. Remote debugging should be turned off.	AuditIfNotExists, Disabled	1.0.0
Resource logs in App Services should be enabled	Audit enabling of resource logs on the app. This enables you to recreate activity trails for investigation purposes if a security incident occurs or your network is compromised.	AuditIfNotExists, Disabled	1.0.0
Resource logs in Azure Data Lake Store should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Azure Kubernetes Service should be enabled	Azure Kubernetes Service's resource logs can help recreate activity trails when investigating security incidents. Enable it to make sure the logs will exist when needed	AuditIfNotExists, Disabled	1.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Resource logs in Azure Stream Analytics should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Batch accounts should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Data Lake Analytics should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Event Hub should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in IoT Hub should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	3.0.1
Resource logs in Key Vault should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Logic Apps should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Resource logs in Search services should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Service Bus should be enabled	Audit enabling of resource logs. This enables you to recreate activity trails to use for investigation purposes; when a security incident occurs or when your network is compromised	AuditIfNotExists, Disabled	5.0.0
Resource logs in Virtual Machine Scale Sets should be enabled	It is recommended to enable Logs so that activity trail can be recreated when investigations are required in the event of an incident or a compromise.	AuditIfNotExists, Disabled	2.1.0
Role-Based Access Control (RBAC) should be used on Kubernetes Services	To provide granular filtering on the actions that users can perform, use Role-Based Access Control (RBAC) to manage permissions in Kubernetes Service Clusters and configure relevant authorization policies.	Audit, Disabled	1.0.2
Running container images should have vulnerability findings resolved	Container image vulnerability assessment scans container images running on your Kubernetes clusters for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks.	AuditIfNotExists, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Secure transfer to storage accounts should be enabled	Audit requirement of Secure transfer in your storage account. Secure transfer is an option that forces your storage account to accept requests only from secure connections (HTTPS). Use of HTTPS ensures authentication between the server and the service and protects data in transit from network layer attacks such as man-in-the-middle, eavesdropping, and session-hijacking	Audit, Deny, Disabled	2.0.0
Service Fabric clusters should have the ClusterProtectionLevel property set to EncryptAndSign	Service Fabric provides three levels of protection (None, Sign and EncryptAndSign) for node-to-node communication using a primary cluster certificate. Set the protection level to ensure that all node-to-node messages are encrypted and digitally signed	Audit, Deny, Disabled	1.1.0
Service Fabric clusters should only use Azure Active Directory for client authentication	Audit usage of client authentication only via Azure Active Directory in Service Fabric	Audit, Deny, Disabled	1.1.0
Service principals should be used to protect your subscriptions instead of management certificates	Management certificates allow anyone who authenticates with them to manage the subscription(s) they are associated with. To manage subscriptions more securely, use of service principals with Resource Manager is recommended to limit the impact of a certificate compromise.	AuditIfNotExists, Disabled	1.0.0
SQL databases should have vulnerability findings resolved	Monitor vulnerability assessment scan results and recommendations for how to remediate database vulnerabilities.	AuditIfNotExists, Disabled	4.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
SQL managed instances should use customer-managed keys to encrypt data at rest	Implementing Transparent Data Encryption (TDE) with your own key provides you with increased transparency and control over the TDE Protector, increased security with an HSM-backed external service, and promotion of separation of duties. This recommendation applies to organizations with a related compliance requirement.	Audit, Deny, Disabled	2.0.0
SQL servers on machines should have vulnerability findings resolved	SQL vulnerability assessment scans your database for security vulnerabilities, and exposes any deviations from best practices such as misconfigurations, excessive permissions, and unprotected sensitive data. Resolving the vulnerabilities found can greatly improve your database security posture.	AuditIfNotExists, Disabled	1.0.0
SQL servers should use customer-managed keys to encrypt data at rest	Implementing Transparent Data Encryption (TDE) with your own key provides increased transparency and control over the TDE Protector, increased security with an HSM-backed external service, and promotion of separation of duties. This recommendation applies to organizations with a related compliance requirement.	Audit, Deny, Disabled	2.0.1
SQL servers with auditing to storage account destination should be configured with 90 days retention or higher	For incident investigation purposes, we recommend setting the data retention for your SQL Server' auditing to storage account destination to at least 90 days. Confirm that you are meeting the necessary retention rules for the regions in which you are operating. This is sometimes required for compliance with regulatory standards.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Storage accounts should be migrated to new Azure Resource Manager resources	Use new Azure Resource Manager for your storage accounts to provide security enhancements such as: stronger access control (RBAC), better auditing, Azure Resource Manager based deployment and governance, access to managed identities, access to key vault for secrets, Azure AD-based authentication and support for tags and resource groups for easier security management	Audit, Deny, Disabled	1.0.0
Storage accounts should restrict network access	Network access to storage accounts should be restricted. Configure network rules so only applications from allowed networks can access the storage account. To allow connections from specific internet or on-premises clients, access can be granted to traffic from specific Azure virtual networks or to public internet IP address ranges	Audit, Deny, Disabled	1.1.1
Storage accounts should restrict network access using virtual network rules	Protect your storage accounts from potential threats using virtual network rules as a preferred method instead of IP-based filtering. Disabling IP-based filtering prevents public IPs from accessing your storage accounts.	Audit, Deny, Disabled	1.0.1
Storage accounts should use customer-managed key for encryption	Secure your blob and file storage account with greater flexibility using customer-managed keys. When you specify a customer-managed key, that key is used to protect and control access to the key that encrypts your data. Using customer-managed keys provides additional capabilities to control rotation of the key encryption key or cryptographically erase data.	Audit, Disabled	1.0.3

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Storage accounts should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The Private Link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your storage account, data leakage risks are reduced. Learn more about private links at - https://aka.ms/azureprivatelinkoverview	AuditIfNotExists, Disabled	2.0.0
Subnets should be associated with a Network Security Group	Protect your subnet from potential threats by restricting access to it with a Network Security Group (NSG). NSGs contain a list of Access Control List (ACL) rules that allow or deny network traffic to your subnet.	AuditIfNotExists, Disabled	3.0.0
Subscriptions should have a contact email address for security issues	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, set a security contact to receive email notifications from Security Center.	AuditIfNotExists, Disabled	1.0.1
System updates on virtual machine scale sets should be installed	Audit whether there are any missing system security updates and critical updates that should be installed to ensure that your Windows and Linux virtual machine scale sets are secure.	AuditIfNotExists, Disabled	3.0.0
System updates should be installed on your machines	Missing security system updates on your servers will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	4.0.0
There should be more than one owner assigned to your subscription	It is recommended to designate more than one subscription owner in order to have administrator access redundancy.	AuditIfNotExists, Disabled	3.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Transparent Data Encryption on SQL databases should be enabled	Transparent data encryption should be enabled to protect data-at-rest and meet compliance requirements	AuditIfNotExists, Disabled	2.0.0
Virtual machines should be migrated to new Azure Resource Manager resources	Use new Azure Resource Manager for your virtual machines to provide security enhancements such as: stronger access control (RBAC), better auditing, Azure Resource Manager based deployment and governance, access to managed identities, access to key vault for secrets, Azure AD-based authentication and support for tags and resource groups for easier security management	Audit, Deny, Disabled	1.0.0
Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources	By default, a virtual machine's OS and data disks are encrypted-at-rest using platform-managed keys. Temp disks, data caches and data flowing between compute and storage aren't encrypted. Disregard this recommendation if: 1. using encryption-at-host, or 2. server-side encryption on Managed Disks meets your security requirements. Learn more in: Server-side encryption of Azure Disk Storage: https://aka.ms/disksse, Different disk encryption offerings: https://aka.ms/diskencryptioncomparison	AuditIfNotExists, Disabled	2.0.3
Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity	The Guest Configuration extension requires a system assigned managed identity. Azure virtual machines in the scope of this policy will be non-compliant when they have the Guest Configuration extension installed but do not have a system assigned managed identity. Learn more at https://aka.ms/gcpol	AuditIfNotExists, Disabled	1.0.1

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
VM Image Builder templates should use private link	Azure Private Link lets you connect your virtual network to Azure services without a public IP address at the source or destination. The Private Link platform handles the connectivity between the consumer and services over the Azure backbone network. By mapping private endpoints to your VM Image Builder building resources, data leakage risks are reduced. Learn more about private links at: https://docs.microsoft.com/azure/virtual-machines/linux/image-builder-networking#deploy-using-an-existing-vnet.	Audit, Disabled, Deny	1.1.0
Vulnerabilities in container security configurations should be remediated	Audit vulnerabilities in security configuration on machines with Docker installed and display as recommendations in Azure Security Center.	AuditIfNotExists, Disabled	3.0.0
Vulnerabilities in security configuration on your machines should be remediated	Servers which do not satisfy the configured baseline will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
Vulnerabilities in security configuration on your virtual machine scale sets should be remediated	Audit the OS vulnerabilities on your virtual machine scale sets to protect them from attacks.	AuditIfNotExists, Disabled	3.0.0
Vulnerability assessment should be enabled on SQL Managed Instance	Audit each SQL Managed Instance which doesn't have recurring vulnerability assessment scans enabled. Vulnerability assessment can discover, track, and help you remediate potential database vulnerabilities.	AuditIfNotExists, Disabled	1.0.1
Vulnerability assessment should be enabled on your SQL servers	Audit Azure SQL servers which do not have recurring vulnerability assessment scans enabled. Vulnerability assessment can discover, track, and help you remediate potential database vulnerabilities.	AuditIfNotExists, Disabled	2.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Web Application Firewall (WAF) should be enabled for Application Gateway	Deploy Azure Web Application Firewall (WAF) in front of public facing web applications for additional inspection of incoming traffic. Web Application Firewall (WAF) provides centralized protection of your web applications from common exploits and vulnerabilities such as SQL injections, Cross-Site Scripting, local and remote file executions. You can also restrict access to your web applications by countries, IP address ranges, and other http(s) parameters via custom rules.	Audit, Deny, Disabled	2.0.0
Web Application should only be accessible over HTTPS	Use of HTTPS ensures server/service authentication and protects data in transit from network layer eavesdropping attacks.	Audit, Disabled	1.0.0
Windows Defender Exploit Guard should be enabled on your machines	Windows Defender Exploit Guard uses the Azure Policy Guest Configuration agent. Exploit Guard has four components that are designed to lock down devices against a wide variety of attack vectors and block behaviors commonly used in malware attacks while enabling enterprises to balance their security risk and productivity requirements (Windows only).	AuditIfNotExists, Disabled	1.1.1
Windows machines should have Log Analytics agent installed on Azure Arc	Machines are non-compliant if Log Analytics agent is not installed on Azure Arc enabled windows server.	AuditIfNotExists, Disabled	1.0.0

POLICY NAME	DESCRIPTION	EFFECT(S)	VERSION
Windows machines should meet requirements of the Azure compute security baseline	Requires that prerequisites are deployed to the policy assignment scope. For details, visit https://aka.ms/gcpol. Machines are non-compliant if the machine is not configured correctly for one of the recommendations in the Azure compute security baseline.	AuditIfNotExists, Disabled	1.0.1
Windows web servers should be configured to use secure communication protocols	To protect the privacy of information communicated over the Internet, your web servers should use the latest version of the industry-standard cryptographic protocol, Transport Layer Security (TLS). TLS secures communications over a network by using security certificates to encrypt a connection between machines.	AuditIfNotExists, Disabled	3.0.0

Microsoft Defender for Cloud category

NAME (AZURE PORTAL)	DESCRIPTION	EFFECT(S)	VERSION (GITHUB)
[Preview]: [Preview]: Azure Security agent should be installed on your Linux Arc machines	Install the Azure Security agent on your Linux Arc machines in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Azure Security agent should be installed on your Linux virtual machine scale sets	Install the Azure Security agent on your Linux virtual machine scale sets in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Azure Security agent should be installed on your Linux virtual machines	Install the Azure Security agent on your Linux virtual machines in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Azure Security agent should be installed on your Windows Arc machines	Install the Azure Security agent on your Windows Arc machines in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Azure Security agent should be installed on your Windows virtual machine scale sets	Install the Azure Security agent on your Windows virtual machine scale sets in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Azure Security agent should be installed on your Windows virtual machines	Install the Azure Security agent on your Windows virtual machines in order to monitor your machines for security configurations and vulnerabilities. Results of the assessments can seen and managed in Azure Security Center.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: ChangeTracking extension should be installed on your Linux Arc machine	Install ChangeTracking Extension on Linux Arc machines to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: ChangeTracking extension should be installed on your Linux virtual machine	Install ChangeTracking Extension on Linux virtual machines to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: ChangeTracking extension should be installed on your Linux virtual machine scale sets	Install ChangeTracking Extension on Linux virtual machine scale sets to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: ChangeTracking extension should be installed on your Windows Arc machine	Install ChangeTracking Extension on Windows Arc machines to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: ChangeTracking extension should be installed on your Windows virtual machine	Install ChangeTracking Extension on Windows virtual machines to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: ChangeTracking extension should be installed on your Windows virtual machine scale sets	Install ChangeTracking Extension on Windows virtual machine scale sets to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitoring Agent.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure Association to link virtual machines to default Azure Security Center Data Collection Rule	Configure machines to automatically create an association with the default data collection rule for Azure Security Center. Deleting this association will break the detection of security vulnerabilities for this virtual machine. Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure Azure Defender for SQL agent on virtual machine	Configure Windows machines to automatically install the Azure Defender for SQL agent where the Azure Monitor Agent is installed. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Creates a resource group and Log Analytics workspace in the same region as the machine. Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure ChangeTracking Extension for Linux Arc machines	Configure Linux Arc machines to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure ChangeTracking Extension for Linux virtual machine scale sets	Configure Linux virtual machine scale sets to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure ChangeTracking Extension for Linux virtual machines	Configure Linux virtual machines to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure ChangeTracking Extension for Windows Arc machines	Configure Windows Arc machines to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure ChangeTracking Extension for Windows virtual machine scale sets	Configure Windows virtual machine scale sets to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure ChangeTracking Extension for Windows virtual machines	Configure Windows virtual machines to automatically install the ChangeTracking Extension to enable File Integrity Monitoring(FIM) in Azure Security Center. FIM examines operating system files, Windows registries, application software, Linux system files, and more, for changes that might indicate an attack. The extension can be installed in virtual machines and locations supported by Azure Monitor Agent.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure machines to automatically create the Azure Security Center pipeline for Azure Monitor Agent	Configure machines to automatically create the Azure Security Center pipeline for Azure Monitor Agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Create a resource group and Log Analytics workspace in the same region as the machine to store audit records. Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	4.0.0-preview
[Preview]: [Preview]: Configure machines to receive a vulnerability assessment provider	Azure Defender includes vulnerability scanning for your machines at no extra cost. You don't need a Qualys license or even a Qualys account - everything's handled seamlessly inside Security Center. When you enable this policy, Azure Defender automatically deploys the Qualys vulnerability assessment provider to all supported machines that don't already have it installed.	DeployIfNotExists, Disabled	2.2.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure supported Linux Arc machines to automatically install the Azure Security agent	Configure supported Linux Arc machines to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target Linux Arc machines must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure supported Linux virtual machine scale sets to automatically install the Azure Security agent	Configure supported Linux virtual machine scale sets to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure supported Linux virtual machine scale sets to automatically install the Guest Attestation extension	Configure supported Linux virtual machines scale sets to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	5.0.0-preview
[Preview]: [Preview]: Configure supported Linux virtual machines to automatically enable Secure Boot	Configure supported Linux virtual machines to automatically enable Secure Boot to mitigate against malicious and unauthorized changes to the boot chain. Once enabled, only trusted bootloaders, kernel and kernel drivers will be allowed to run.	DeployIfNotExists, Disabled	5.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure supported Linux virtual machines to automatically install the Azure Security agent	Configure supported Linux virtual machines to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	6.0.0-preview
[Preview]: [Preview]: Configure supported Linux virtual machines to automatically install the Guest Attestation extension	Configure supported Linux virtual machines to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	6.0.0-preview
[Preview]: [Preview]: Configure supported virtual machines to automatically enable vTPM	Configure supported virtual machines to automatically enable vTPM to facilitate Measured Boot and other OS security features that require a TPM. Once enabled, vTPM can be used to attest boot integrity.	DeployIfNotExists, Disabled	2.0.0-preview
[Preview]: [Preview]: Configure supported Windows Arc machines to automatically install the Azure Security agent	Configure supported Windows Arc machines to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target Windows Arc machines must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure supported Windows machines to automatically install the Azure Security agent	Configure supported Windows machines to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target virtual machines must be in a supported location.	DeployIfNotExists, Disabled	4.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure supported Windows virtual machine scale sets to automatically install the Azure Security agent	Configure supported Windows virtual machine scale sets to automatically install the Azure Security agent. Security Center collects events from the agent and uses them to provide security alerts and tailored hardening tasks (recommendations). Target Windows virtual machine scale sets must be in a supported location.	DeployIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Configure supported Windows virtual machine scale sets to automatically install the Guest Attestation extension	Configure supported Windows virtual machines scale sets to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	3.0.0-preview
[Preview]: [Preview]: Configure supported Windows virtual machines to automatically enable Secure Boot	Configure supported Windows virtual machines to automatically enable Secure Boot to mitigate against malicious and unauthorized changes to the boot chain. Once enabled, only trusted bootloaders, kernel and kernel drivers will be allowed to run.	DeployIfNotExists, Disabled	3.0.0-preview
[Preview]: [Preview]: Configure supported Windows virtual machines to automatically install the Guest Attestation extension	Configure supported Windows virtual machines to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	4.0.0-preview
[Preview]: [Preview]: Configure VMs created with Shared Image Gallery images to install the Guest Attestation extension	Configure virtual machines created with Shared Image Gallery images to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	2.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Configure VMSS created with Shared Image Gallery images to install the Guest Attestation extension	Configure VMSS created with Shared Image Gallery images to automatically install the Guest Attestation extension to allow Azure Security Center to proactively attest and monitor the boot integrity. Boot integrity is attested via Remote Attestation.	DeployIfNotExists, Disabled	2.0.0-preview
[Preview]: [Preview]: Guest Attestation extension should be installed on supported Linux virtual machines	Install Guest Attestation extension on supported Linux virtual machines to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machines.	AuditIfNotExists, Disabled	5.0.0-preview
[Preview]: [Preview]: Guest Attestation extension should be installed on supported Linux virtual machines scale sets	Install Guest Attestation extension on supported Linux virtual machines scale sets to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled Linux virtual machine scale sets.	AuditIfNotExists, Disabled	4.0.0-preview
[Preview]: [Preview]: Guest Attestation extension should be installed on supported Windows virtual machines	Install Guest Attestation extension on supported virtual machines to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machines.	AuditIfNotExists, Disabled	3.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Guest Attestation extension should be installed on supported Windows virtual machines scale sets	Install Guest Attestation extension on supported virtual machines scale sets to allow Azure Security Center to proactively attest and monitor the boot integrity. Once installed, boot integrity will be attested via Remote Attestation. This assessment only applies to trusted launch enabled virtual machine scale sets.	AuditIfNotExists, Disabled	2.0.0-preview
[Preview]: [Preview]: Linux virtual machines should use Secure Boot	To protect against the installation of malware-based rootkits and boot kits, enable Secure Boot on supported Linux virtual machines. Secure Boot ensures that only signed operating systems and drivers will be allowed to run. This assessment only applies to Linux virtual machines that have the Azure Monitor Agent installed.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Machines should have ports closed that might expose attack vectors	Azure's Terms Of Use prohibit the use of Azure services in ways that could damage, disable, overburden, or impair any Microsoft server, or the network. The exposed ports identified by this recommendation need to be closed for your continued security. For each identified port, the recommendation also provides an explanation of the potential threat.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: Secure Boot should be enabled on supported Windows virtual machines	Enable Secure Boot on supported Windows virtual machines to mitigate against malicious and unauthorized changes to the boot chain. Once enabled, only trusted bootloaders, kernel and kernel drivers will be allowed to run. This assessment only applies to trusted launch enabled Windows virtual machines.	Audit, Disabled	3.0.0-preview

NAME	DESCRIPTION	EFFECT(S)	VERSION
[Preview]: [Preview]: Sensitive data in your SQL databases should be classified	Azure Security Center monitors the data discovery and classification scan results for your SQL databases and provides recommendations to classify the sensitive data in your databases for better monitoring and security	AuditIfNotExists, Disabled	3.0.0-preview
[Preview]: [Preview]: Virtual machines guest attestation status should be healthy	Guest attestation is performed by sending a trusted log (TCGLog) to an attestation server. The server uses these logs to determine whether boot components are trustworthy. This assessment is intended to detect compromises of the boot chain which might be the result of a bootkit or rootkit infection. This assessment only applies to Trusted Launch enabled virtual machines that have Guest Attestation extension installed.	AuditIfNotExists, Disabled	1.0.0-preview
[Preview]: [Preview]: vTPM should be enabled on supported virtual machines	Enable virtual TPM device on supported virtual machines to facilitate Measured Boot and other OS security features that require a TPM. Once enabled, vTPM can be used to attest boot integrity. This assessment only applies to trusted launch enabled virtual machines.	Audit, Disabled	2.0.0-preview
A maximum of 3 owners should be designated for your subscription	It is recommended to designate up to 3 subscription owners in order to reduce the potential for breach by a compromised owner.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
A vulnerability assessment solution should be enabled on your virtual machines	Audits virtual machines to detect whether they are running a supported vulnerability assessment solution. A core component of every cyber risk and security program is the identification and analysis of vulnerabilities. Azure Security Center's standard pricing tier includes vulnerability scanning for your virtual machines at no extra cost. Additionally, Security Center can automatically deploy this tool for you.	AuditIfNotExists, Disabled	3.0.0
Adaptive application controls for defining safe applications should be enabled on your machines	Enable application controls to define the list of known-safe applications running on your machines, and alert you when other applications run. This helps harden your machines against malware. To simplify the process of configuring and maintaining your rules, Security Center uses machine learning to analyze the applications running on each machine and suggest the list of known-safe applications.	AuditIfNotExists, Disabled	3.0.0
Adaptive network hardening recommendations should be applied on internet facing virtual machines	Azure Security Center analyzes the traffic patterns of Internet facing virtual machines and provides Network Security Group rule recommendations that reduce the potential attack surface	AuditIfNotExists, Disabled	3.0.0
All network ports should be restricted on network security groups associated to your virtual machine	Azure Security Center has identified some of your network security groups' inbound rules to be too permissive. Inbound rules should not allow access from 'Any' or 'Internet' ranges. This can potentially enable attackers to target your resources.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Allowlist rules in your adaptive application control policy should be updated	Monitor for changes in behavior on groups of machines configured for auditing by Azure Security Center's adaptive application controls. Security Center uses machine learning to analyze the running processes on your machines and suggest a list of known-safe applications. These are presented as recommended apps to allow in adaptive application control policies.	AuditIfNotExists, Disabled	3.0.0
Authorized IP ranges should be defined on Kubernetes Services	Restrict access to the Kubernetes Service Management API by granting API access only to IP addresses in specific ranges. It is recommended to limit access to authorized IP ranges to ensure that only applications from allowed networks can access the cluster.	Audit, Disabled	2.0.1
Auto provisioning of the Log Analytics agent should be enabled on your subscription	To monitor for security vulnerabilities and threats, Azure Security Center collects data from your Azure virtual machines. Data is collected by the Log Analytics agent, formerly known as the Microsoft Monitoring Agent (MMA), which reads various security-related configurations and event logs from the machine and copies the data to your Log Analytics workspace for analysis. We recommend enabling auto provisioning to automatically deploy the agent to all supported Azure VMs and any new ones that are created.	AuditIfNotExists, Disabled	1.0.1
Azure DDoS Protection Standard should be enabled	DDoS protection standard should be enabled for all virtual networks with a subnet that is part of an application gateway with a public IP.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for App Service should be enabled	Azure Defender for App Service leverages the scale of the cloud, and the visibility that Azure has as a cloud provider, to monitor for common web app attacks.	AuditIfNotExists, Disabled	1.0.3
Azure Defender for Azure SQL Database servers should be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	AuditIfNotExists, Disabled	1.0.2
Azure Defender for DNS should be enabled	Azure Defender for DNS provides an additional layer of protection for your cloud resources by continuously monitoring all DNS queries from your Azure resources. Azure Defender alerts you about suspicious activity at the DNS layer. Learn more about the capabilities of Azure Defender for DNS at https://aka.ms/defender-for-dns . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	AuditIfNotExists, Disabled	1.0.0
Azure Defender for Key Vault should be enabled	Azure Defender for Key Vault provides an additional layer of protection and security intelligence by detecting unusual and potentially harmful attempts to access or exploit key vault accounts.	AuditIfNotExists, Disabled	1.0.3

NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for open-source relational databases should be enabled	Azure Defender for open-source relational databases detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Learn more about the capabilities of Azure Defender for open-source relational databases at https://aka.ms/AzDforOpenSourceDBsDocu. Important: Enabling this plan will result in charges for protecting your open-source relational databases. Learn about the pricing on Security Center's pricing page: https://aka.ms/pricing-security-center	AuditIfNotExists, Disabled	1.0.0
Azure Defender for Resource Manager should be enabled	Azure Defender for Resource Manager automatically monitors the resource management operations in your organization. Azure Defender detects threats and alerts you about suspicious activity. Learn more about the capabilities of Azure Defender for Resource Manager at https://aka.ms/defender-for-resource-manager . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	AuditIfNotExists, Disabled	1.0.0
Azure Defender for servers should be enabled	Azure Defender for servers provides real-time threat protection for server workloads and generates hardening recommendations as well as alerts about suspicious activities.	AuditIfNotExists, Disabled	1.0.3

NAME	DESCRIPTION	EFFECT(S)	VERSION
Azure Defender for SQL servers on machines should be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	AuditIfNotExists, Disabled	1.0.2
Azure Defender for Storage should be enabled	Azure Defender for Storage provides detections of unusual and potentially harmful attempts to access or exploit storage accounts.	AuditIfNotExists, Disabled	1.0.3
Cloud Services (extended support) role instances should be configured securely	Protect your Cloud Service (extended support) role instances from attacks by ensuring they are not expolosed to any OS vulnerabilities.	AuditIfNotExists, Disabled	1.0.0
Cloud Services (extended support) role instances should have an endpoint protection solution installed	Protect your Cloud Services (extended support) role instances from threats and vulnerabilities by ensuring an endpoint protection solution is installed on them.	AuditIfNotExists, Disabled	1.0.0
Cloud Services (extended support) role instances should have system updates installed	Secure your Cloud Services (extended support) role instances by ensuring the latest security and critical updates are installed on them.	AuditIfNotExists, Disabled	1.0.0
Configure Azure Defender for App Service to be enabled	Azure Defender for App Service leverages the scale of the cloud, and the visibility that Azure has as a cloud provider, to monitor for common web app attacks.	DeployIfNotExists, Disabled	1.0.1
Configure Azure Defender for Azure SQL database to be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	DeployIfNotExists, Disabled	1.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Configure Azure Defender for DNS to be enabled	Azure Defender for DNS provides an additional layer of protection for your cloud resources by continuously monitoring all DNS queries from your Azure resources. Azure Defender alerts you about suspicious activity at the DNS layer. Learn more about the capabilities of Azure Defender for DNS at https://aka.ms/defender-for-dns . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	DeployIfNotExists, Disabled	1.0.1
Configure Azure Defender for Key Vaults to be enabled	Azure Defender for Key Vault provides an additional layer of protection and security intelligence by detecting unusual and potentially harmful attempts to access or exploit key vault accounts.	DeployIfNotExists, Disabled	1.0.1
Configure Azure Defender for open-source relational databases to be enabled	Azure Defender for open-source relational databases detects anomalous activities indicating unusual and potentially harmful attempts to access or exploit databases. Learn more about the capabilities of Azure Defender for open-source relational databases at https://aka.ms/AzDforOpenSourceDBsDocu. Important: Enabling this plan will result in charges for protecting your open-source relational databases. Learn about the pricing on Security Center's pricing page: https://aka.ms/pricing-security-center	DeployIfNotExists, Disabled	1.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Configure Azure Defender for Resource Manager to be enabled	Azure Defender for Resource Manager automatically monitors the resource management operations in your organization. Azure Defender detects threats and alerts you about suspicious activity. Learn more about the capabilities of Azure Defender for Resource Manager at https://aka.ms/defender-for-resource-manager . Enabling this Azure Defender plan results in charges. Learn about the pricing details per region on Security Center's pricing page: https://aka.ms/pricing-security-center .	DeployIfNotExists, Disabled	1.0.1
Configure Azure Defender for servers to be enabled	Azure Defender for servers provides real-time threat protection for server workloads and generates hardening recommendations as well as alerts about suspicious activities.	DeployIfNotExists, Disabled	1.0.0
Configure Azure Defender for SQL servers on machines to be enabled	Azure Defender for SQL provides functionality for surfacing and mitigating potential database vulnerabilities, detecting anomalous activities that could indicate threats to SQL databases, and discovering and classifying sensitive data.	DeployIfNotExists, Disabled	1.0.0
Configure Azure Defender for Storage to be enabled	Azure Defender for Storage provides detections of unusual and potentially harmful attempts to access or exploit storage accounts.	DeployIfNotExists, Disabled	1.0.0
Configure Microsoft Defender for Containers to be enabled	Microsoft Defender for Containers provides hardening, vulnerability assessment and run-time protections for your Azure, hybrid, and multi-cloud Kubernetes environments.	DeployIfNotExists, Disabled	1.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Container registry images should have vulnerability findings resolved	Container image vulnerability assessment scans your registry for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks.	AuditIfNotExists, Disabled	2.0.1
Deploy - Configure suppression rules for Azure Security Center alerts	Suppress Azure Security Center alerts to reduce alerts fatigue by deploying suppression rules on your management group or subscription.	deployIfNotExists	1.0.0
Deploy export to Event Hub for Azure Security Center data	Enable export to Event Hub of Azure Security Center data. This policy deploys an export to Event Hub configuration with your conditions and target Event Hub on the assigned scope. To deploy this policy on newly created subscriptions, open the Compliance tab, select the relevant non-compliant assignment and create a remediation task.	deployIfNotExists	4.0.0
Deploy export to Log Analytics workspace for Azure Security Center data	Enable export to Log Analytics workspace of Azure Security Center data. This policy deploys an export to Log Analytics workspace configuration with your conditions and target workspace on the assigned scope. To deploy this policy on newly created subscriptions, open the Compliance tab, select the relevant non-compliant assignment and create a remediation task.	deployIfNotExists	4.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Deploy Workflow Automation for Azure Security Center alerts	Enable automation of Azure Security Center alerts. This policy deploys a workflow automation with your conditions and triggers on the assigned scope. To deploy this policy on newly created subscriptions, open the Compliance tab, select the relevant non-compliant assignment and create a remediation task.	deployIfNotExists	4.0.0
Deploy Workflow Automation for Azure Security Center recommendations	Enable automation of Azure Security Center recommendations. This policy deploys a workflow automation with your conditions and triggers on the assigned scope. To deploy this policy on newly created subscriptions, open the Compliance tab, select the relevant non-compliant assignment and create a remediation task.	deployIfNotExists	4.0.0
Deploy Workflow Automation for Azure Security Center regulatory compliance	Enable automation of Azure Security Center regulatory compliance. This policy deploys a workflow automation with your conditions and triggers on the assigned scope. To deploy this policy on newly created subscriptions, open the Compliance tab, select the relevant non-compliant assignment and create a remediation task.	deployIfNotExists	4.0.0
Deprecated accounts should be removed from your subscription	Deprecated accounts should be removed from your subscriptions. Deprecated accounts are accounts that have been blocked from signing in.	AuditIfNotExists, Disabled	3.0.0
Deprecated accounts with owner permissions should be removed from your subscription	Deprecated accounts with owner permissions should be removed from your subscription. Deprecated accounts are accounts that have been blocked from signing in.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Email notification for high severity alerts should be enabled	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, enable email notifications for high severity alerts in Security Center.	AuditIfNotExists, Disabled	1.0.1
Email notification to subscription owner for high severity alerts should be enabled	To ensure your subscription owners are notified when there is a potential security breach in their subscription, set email notifications to subscription owners for high severity alerts in Security Center.	AuditIfNotExists, Disabled	2.0.0
Enable Azure Security Center on your subscription	Identifies existing subscriptions that are not monitored by Azure Security Center (ASC). Subscriptions not monitored by ASC will be registered to the free pricing tier. Subscriptions already monitored by ASC (free or standard), will be considered compliant. To register newly created subscriptions, open the compliance tab, select the relevant non-compliant assignment and create a remediation task. Repeat this step when you have one or more new subscriptions you want to monitor with Security Center.	deployIfNotExists	1.0.0
Enable Security Center's auto provisioning of the Log Analytics agent on your subscriptions with custom workspace.	Allow Security Center to auto provision the Log Analytics agent on your subscriptions to monitor and collect security data using a custom workspace.	DeployIfNotExists, Disabled	1.0.0
Enable Security Center's auto provisioning of the Log Analytics agent on your subscriptions with default workspace.	Allow Security Center to auto provision the Log Analytics agent on your subscriptions to monitor and collect security data using ASC default workspace.	DeployIfNotExists, Disabled	1.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Endpoint protection health issues should be resolved on your machines	Resolve endpoint protection health issues on your virtual machines to protect them from latest threats and vulnerabilities. Azure Security Center supported endpoint protection solutions are documented here - https://docs.microsoft.com/azure/security-center/security-center-services?tabs=features-windows#supported-endpoint-protection-solutions. Endpoint protection assessment is documented here - https://docs.microsoft.com/azure/security-center/security-center-endpoint-protection.	AuditIfNotExists, Disabled	1.0.0
Endpoint protection should be installed on your machines	To protect your machines from threats and vulnerabilities, install a supported endpoint protection solution.	AuditIfNotExists, Disabled	1.0.0
Endpoint protection solution should be installed on virtual machine scale sets	Audit the existence and health of an endpoint protection solution on your virtual machines scale sets, to protect them from threats and vulnerabilities.	AuditIfNotExists, Disabled	3.0.0
External accounts with owner permissions should be removed from your subscription	External accounts with owner permissions should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0
External accounts with read permissions should be removed from your subscription	External accounts with read privileges should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0
External accounts with write permissions should be removed from your subscription	External accounts with write privileges should be removed from your subscription in order to prevent unmonitored access.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Guest Configuration extension should be installed on your machines	To ensure secure configurations of in-guest settings of your machine, install the Guest Configuration extension. In-guest settings that the extension monitors include the configuration of the operating system, application configuration or presence, and environment settings. Once installed, in-guest policies will be available such as 'Windows Exploit guard should be enabled'. Learn more at https://aka.ms/gcpol.	AuditIfNotExists, Disabled	1.0.2
Internet-facing virtual machines should be protected with network security groups	Protect your virtual machines from potential threats by restricting access to them with network security groups (NSG). Learn more about controlling traffic with NSGs at https://aka.ms/nsg-doc	AuditIfNotExists, Disabled	3.0.0
IP Forwarding on your virtual machine should be disabled	Enabling IP forwarding on a virtual machine's NIC allows the machine to receive traffic addressed to other destinations. IP forwarding is rarely required (e.g., when using the VM as a network virtual appliance), and therefore, this should be reviewed by the network security team.	AuditIfNotExists, Disabled	3.0.0
Kubernetes Services should be upgraded to a non-vulnerable Kubernetes version	Upgrade your Kubernetes service cluster to a later Kubernetes version to protect against known vulnerabilities in your current Kubernetes version. Vulnerability CVE-2019-9946 has been patched in Kubernetes versions 1.11.9+, 1.12.7+, 1.13.5+, and 1.14.0+	Audit, Disabled	1.0.2
Log Analytics agent should be installed on your Cloud Services (extended support) role instances	Security Center collects data from your Cloud Services (extended support) role instances to monitor for security vulnerabilities and threats.	AuditIfNotExists, Disabled	2.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Log Analytics agent should be installed on your virtual machine for Azure Security Center monitoring	This policy audits any Windows/Linux virtual machines (VMs) if the Log Analytics agent is not installed which Security Center uses to monitor for security vulnerabilities and threats	AuditIfNotExists, Disabled	1.0.0
Log Analytics agent should be installed on your virtual machine scale sets for Azure Security Center monitoring	Security Center collects data from your Azure virtual machines (VMs) to monitor for security vulnerabilities and threats.	AuditIfNotExists, Disabled	1.0.0
Management ports of virtual machines should be protected with just-in-time network access control	Possible network Just In Time (JIT) access will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
Management ports should be closed on your virtual machines	Open remote management ports are exposing your VM to a high level of risk from Internet-based attacks. These attacks attempt to brute force credentials to gain admin access to the machine.	AuditIfNotExists, Disabled	3.0.0
MFA should be enabled accounts with write permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with write privileges to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0
MFA should be enabled on accounts with owner permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with owner permissions to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0
MFA should be enabled on accounts with read permissions on your subscription	Multi-Factor Authentication (MFA) should be enabled for all subscription accounts with read privileges to prevent a breach of accounts or resources.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Microsoft Defender for Containers should be enabled	Microsoft Defender for Containers provides hardening, vulnerability assessment and run-time protections for your Azure, hybrid, and multi-cloud Kubernetes environments.	AuditIfNotExists, Disabled	1.0.0
Monitor missing Endpoint Protection in Azure Security Center	Servers without an installed Endpoint Protection agent will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
Non-internet-facing virtual machines should be protected with network security groups	Protect your non-internet-facing virtual machines from potential threats by restricting access with network security groups (NSG). Learn more about controlling traffic with NSGs at https://aka.ms/nsg-doc	AuditIfNotExists, Disabled	3.0.0
Role-Based Access Control (RBAC) should be used on Kubernetes Services	To provide granular filtering on the actions that users can perform, use Role-Based Access Control (RBAC) to manage permissions in Kubernetes Service Clusters and configure relevant authorization policies.	Audit, Disabled	1.0.2
Running container images should have vulnerability findings resolved	Container image vulnerability assessment scans container images running on your Kubernetes clusters for security vulnerabilities and exposes detailed findings for each image. Resolving the vulnerabilities can greatly improve your containers' security posture and protect them from attacks.	AuditIfNotExists, Disabled	1.0.1
Security Center standard pricing tier should be selected	The standard pricing tier enables threat detection for networks and virtual machines, providing threat intelligence, anomaly detection, and behavior analytics in Azure Security Center	Audit, Disabled	1.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Service principals should be used to protect your subscriptions instead of management certificates	Management certificates allow anyone who authenticates with them to manage the subscription(s) they are associated with. To manage subscriptions more securely, use of service principals with Resource Manager is recommended to limit the impact of a certificate compromise.	AuditIfNotExists, Disabled	1.0.0
SQL databases should have vulnerability findings resolved	Monitor vulnerability assessment scan results and recommendations for how to remediate database vulnerabilities.	AuditIfNotExists, Disabled	4.0.0
SQL servers on machines should have vulnerability findings resolved	SQL vulnerability assessment scans your database for security vulnerabilities, and exposes any deviations from best practices such as misconfigurations, excessive permissions, and unprotected sensitive data. Resolving the vulnerabilities found can greatly improve your database security posture.	AuditIfNotExists, Disabled	1.0.0
Subnets should be associated with a Network Security Group	Protect your subnet from potential threats by restricting access to it with a Network Security Group (NSG). NSGs contain a list of Access Control List (ACL) rules that allow or deny network traffic to your subnet.	AuditIfNotExists, Disabled	3.0.0
Subscriptions should have a contact email address for security issues	To ensure the relevant people in your organization are notified when there is a potential security breach in one of your subscriptions, set a security contact to receive email notifications from Security Center.	AuditIfNotExists, Disabled	1.0.1
System updates on virtual machine scale sets should be installed	Audit whether there are any missing system security updates and critical updates that should be installed to ensure that your Windows and Linux virtual machine scale sets are secure.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
System updates should be installed on your machines	Missing security system updates on your servers will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	4.0.0
There should be more than one owner assigned to your subscription	It is recommended to designate more than one subscription owner in order to have administrator access redundancy.	AuditIfNotExists, Disabled	3.0.0
Virtual machines should encrypt temp disks, caches, and data flows between Compute and Storage resources	By default, a virtual machine's OS and data disks are encrypted-at-rest using platform-managed keys. Temp disks, data caches and data flowing between compute and storage aren't encrypted. Disregard this recommendation if: 1. using encryption-at-host, or 2. server-side encryption on Managed Disks meets your security requirements. Learn more in: Server-side encryption of Azure Disk Storage: https://aka.ms/disksse, Different disk encryption offerings: https://aka.ms/diskencryptioncomparison	AuditIfNotExists, Disabled	2.0.3
Virtual machines' Guest Configuration extension should be deployed with system-assigned managed identity	The Guest Configuration extension requires a system assigned managed identity. Azure virtual machines in the scope of this policy will be non-compliant when they have the Guest Configuration extension installed but do not have a system assigned managed identity. Learn more at https://aka.ms/gcpol	AuditIfNotExists, Disabled	1.0.1
Vulnerabilities in container security configurations should be remediated	Audit vulnerabilities in security configuration on machines with Docker installed and display as recommendations in Azure Security Center.	AuditIfNotExists, Disabled	3.0.0

NAME	DESCRIPTION	EFFECT(S)	VERSION
Vulnerabilities in security configuration on your machines should be remediated	Servers which do not satisfy the configured baseline will be monitored by Azure Security Center as recommendations	AuditIfNotExists, Disabled	3.0.0
Vulnerabilities in security configuration on your virtual machine scale sets should be remediated	Audit the OS vulnerabilities on your virtual machine scale sets to protect them from attacks.	AuditIfNotExists, Disabled	3.0.0

Next steps

In this article, you learned about Azure Policy security policy definitions in Defender for Cloud. To learn more about initiatives, policies, and how they relate to Defender for Cloud's recommendations, see What are security policies, initiatives, and recommendations?.

Endpoint protection assessment and recommendations in Microsoft Defender for Cloud

2/15/2022 • 4 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

Microsoft Defender for Cloud provides health assessments of supported versions of Endpoint protection solutions. This article explains the scenarios that lead Defender for Cloud to generate the following two recommendations:

- Endpoint protection should be installed on your machines
- Endpoint protection health issues should be resolved on your machines

TIP

At the end of 2021, we revised the recommendation that installs endpoint protection. One of the changes affects how the recommendation displays machines that are powered off. In the previous version, machines that were turned off appeared in the 'Not applicable' list. In the newer recommendation, they don't appear in any of the resources lists (healthy, unhealthy, or not applicable).

Windows Defender

- Defender for Cloud recommends **Endpoint protection should be installed on your machines** when Get-MpComputerStatus runs and the result is **AMServiceEnabled: False**
- Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when Get-MpComputerStatus runs and any of the following occurs:
  - Any of the following properties are false:
    - **AMServiceEnabled**
    - **AntispywareEnabled**
    - **RealTimeProtectionEnabled**
    - **BehaviorMonitorEnabled**
    - **IoavProtectionEnabled**
    - **OnAccessProtectionEnabled**
  - If one or both of the following properties are 7 or more:
    - **AntispywareSignatureAge**
    - **AntivirusSignatureAge**

Microsoft System Center endpoint protection

- Defender for Cloud recommends **Endpoint protection should be installed on your machines**

when importing **SCEPMpModule ("$env:ProgramFiles\Microsoft Security Client\MpProvider\MpProvider.psd1")** and running **Get-MProtComputerStatus** results in **AMServiceEnabled = false**.

- Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when **Get-MprotComputerStatus** runs and any of the following occurs:
  - At least one of the following properties is false:
    - **AMServiceEnabled**
    - **AntispywareEnabled**
    - **RealTimeProtectionEnabled**
    - **BehaviorMonitorEnabled**
    - **IoavProtectionEnabled**
    - **OnAccessProtectionEnabled**
  - If one or both of the following Signature Updates are greater or equal to 7:
    - **AntispywareSignatureAge**
    - **AntivirusSignatureAge**

Trend Micro

- Defender for Cloud recommends **Endpoint protection should be installed on your machines** when any of the following checks aren't met:
  - **HKLM:\SOFTWARE\TrendMicro\Deep Security Agent** exists
  - **HKLM:\SOFTWARE\TrendMicro\Deep Security Agent\InstallationFolder** exists
  - The **dsa_query.cmd** file is found in the Installation Folder
  - Running **dsa_query.cmd** results with **Component.AM.mode: on - Trend Micro Deep Security Agent detected**

Symantec endpoint protection

Defender for Cloud recommends **Endpoint protection should be installed on your machines** when any of the following checks aren't met:

- **HKLM:\Software\Symantec\Symantec Endpoint Protection\CurrentVersion\PRODUCTNAME = "Symantec Endpoint Protection"**
- **HKLM:\Software\Symantec\Symantec Endpoint Protection\CurrentVersion\public-opstate\ASRunningStatus = 1**

Or

- **HKLM:\Software\Wow6432Node\Symantec\Symantec Endpoint Protection\CurrentVersion\PRODUCTNAME = "Symantec Endpoint Protection"**
- **HKLM:\Software\Wow6432Node\Symantec\Symantec Endpoint Protection\CurrentVersion\public-opstate\ASRunningStatus = 1**

Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when any of the following checks aren't met:

- Check Symantec Version >= 12: Registry location: **HKLM:\Software\Symantec\Symantec Endpoint Protection\CurrentVersion" -Value "PRODUCTVERSION"**
- Check Real-Time Protection status: **HKLM:\Software\Wow6432Node\Symantec\Symantec Endpoint Protection\AV\Storages\Filesystem\RealTimeScan\OnOff == 1**

- Check Signature Update status: **HKLM\Software\Symantec\Symantec Endpoint Protection\CurrentVersion\public-opstate\LatestVirusDefsDate <= 7 days**
- Check Full Scan status: **HKLM:\Software\Symantec\Symantec Endpoint Protection\CurrentVersion\public-opstate\LastSuccessfulScanDateTime <= 7 days**
- Find signature version number Path to signature version for Symantec 12: **Registry Paths+ "CurrentVersion\SharedDefs" -Value "SRTSP"**
- Path to signature version for Symantec 14: **Registry Paths+ "CurrentVersion\SharedDefs\SDSDefs" -Value "SRTSP"**

Registry Paths:

- **"HKLM:\Software\Symantec\Symantec Endpoint Protection" + $Path;**
- **"HKLM:\Software\Wow6432Node\Symantec\Symantec Endpoint Protection" + $Path**

McAfee endpoint protection for Windows

Defender for Cloud recommends **Endpoint protection should be installed on your machines** when any of the following checks aren't met:

- **HKLM:\SOFTWARE\McAfee\Endpoint\AV\ProductVersion** exists
- **HKLM:\SOFTWARE\McAfee\AVSolution\MCSHIELDGLOBAL\GLOBAL\enableoas = 1**

Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when any of the following checks aren't met:

- McAfee Version: **HKLM:\SOFTWARE\McAfee\Endpoint\AV\ProductVersion >= 10**
- Find Signature Version: **HKLM:\Software\McAfee\AVSolution\DS\DS -Value "dwContentMajorVersion"**
- Find Signature date: **HKLM:\Software\McAfee\AVSolution\DS\DS -Value "szContentCreationDate" >= 7 days**
- Find Scan date: **HKLM:\Software\McAfee\Endpoint\AV\ODS -Value "LastFullScanOdsRunTime" >= 7 days**

McAfee Endpoint Security for Linux Threat Prevention

Defender for Cloud recommends **Endpoint protection should be installed on your machines** when any of the following checks aren't met:

- File **/opt/isec/ens/threatprevention/bin/isecav** exists
- **"/opt/isec/ens/threatprevention/bin/isecav --version"** output is: **McAfee name = McAfee Endpoint Security for Linux Threat Prevention and McAfee version >= 10**

Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when any of the following checks aren't met:

- **"/opt/isec/ens/threatprevention/bin/isecav --listtask"** returns **Quick scan, Full scan** and both of the scans <= 7 days
- **"/opt/isec/ens/threatprevention/bin/isecav --listtask"** returns **DAT and engine Update time** and both of them <= 7 days
- **"/opt/isec/ens/threatprevention/bin/isecav --getoasconfig --summary"** returns **On Access Scan** status

Sophos Antivirus for Linux

Defender for Cloud recommends **Endpoint protection should be installed on your machines** when any of the following checks aren't met:

- File **/opt/sophos-av/bin/savdstatus** exits or search for customized location **"readlink $(which savscan)"**
- **"/opt/sophos-av/bin/savdstatus --version"** returns Sophos name = **Sophos Anti-Virus and Sophos version >= 9**

Defender for Cloud recommends **Endpoint protection health issues should be resolved on your machines** when any of the following checks aren't met:

- **"/opt/sophos-av/bin/savlog --maxage=7 | grep -i "Scheduled scan .* completed" | tail -1"**, returns a value
- **"/opt/sophos-av/bin/savlog --maxage=7 | grep "scan finished"** | tail -1", returns a value
- **"/opt/sophos-av/bin/savdstatus --lastupdate"** returns lastUpdate, which should be <= 7 days
- **"/opt/sophos-av/bin/savdstatus -v"** is equal to **"On-access scanning is running"**
- **"/opt/sophos-av/bin/savconfig get LiveProtection"** returns enabled

Troubleshoot and support

Troubleshoot

Microsoft Antimalware extension logs are available at: **%Systemdrive%\WindowsAzure\Logs\Plugins\Microsoft.Azure.Security.IaaSAntimalware(Or PaaSAntimalware)\1.5.5.x(version#)\CommandExecution.log**

Support

For more help, contact the Azure experts on the MSDN Azure and Stack Overflow forums. Or file an Azure support incident. Go to the Azure support site and select Get support. For information about using Azure Support, read the Microsoft Azure support FAQ.

Manage user data in Microsoft Defender for Cloud

2/15/2022 • 3 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This article provides information about how you can manage the user data in Microsoft Defender for Cloud. Managing user data includes the ability to access, delete, or export data.

> **NOTE**
> This article provides steps about how to delete personal data from the device or service and can be used to support your obligations under the GDPR. For general information about GDPR, see the GDPR section of the Microsoft Trust Center and the GDPR section of the Service Trust portal.

A Defender for Cloud user assigned the role of Reader, Owner, Contributor, or Account Administrator can access customer data within the tool. To learn more about the Account Administrator role, see Built-in roles for Azure role-based access control to learn more about the Reader, Owner, and Contributor roles. See Azure subscription administrators.

Searching for and identifying personal data

A Defender for Cloud user can view their personal data through the Azure portal. Defender for Cloud only stores security contact details such as email addresses and phone numbers. For more information, see Provide security contact details in Microsoft Defender for Cloud.

In the Azure portal, a user can view allowed IP configurations using Defender for Cloud's just-in-time VM access feature. For more information, see Manage virtual machine access using just-in-time.

In the Azure portal, a user can view security alerts provided by Defender for Cloud including IP addresses and attacker details. For more information, see Managing and responding to security alerts in Microsoft Defender for Cloud.

Classifying personal data

You don't need to classify personal data found in Defender for Cloud's security contact feature. The data saved is an email address (or multiple email addresses) and a phone number. Contact data is validated by Defender for Cloud.

You don't need to classify the IP addresses and port numbers saved by Defender for Cloud's just-in-time feature.

Only a user assigned the role of Administrator can classify personal data by viewing alerts in Defender for Cloud.

Securing and controlling access to personal data

A Defender for Cloud user assigned the role of Reader, Owner, Contributor, or Account Administrator can access

security contact data.

A Defender for Cloud user assigned the role of Reader, Owner, Contributor, or Account Administrator can access their just-in-time policies.

A Defender for Cloud user assigned the role of Reader, Owner, Contributor, or Account Administrator can view their alerts.

Updating personal data

A Defender for Cloud user assigned the role of Owner, Contributor, or Account Administrator can update security contact data via the Azure portal.

A Defender for Cloud user assigned the role of Owner, Contributor, or Account Administrator can update their just-in-time policies.

An Account Administrator can't edit alert incidents. An alert incident is considered security data and is read only.

Deleting personal data

A Defender for Cloud user assigned the role of Owner, Contributor, or Account Administrator can delete security contact data via the Azure portal.

A Defender for Cloud user assigned the role of Owner, Contributor, or Account Administrator can delete the just-in-time policies via the Azure portal.

A Defender for Cloud user can't delete alert incidents. For security reasons, an alert incident is considered read-only data.

Exporting personal data

A Defender for Cloud user assigned the role of Reader, Owner, Contributor, or Account Administrator can export security contact data by:

- Copying from the Azure portal
- Executing the Azure REST API call, GET HTTP:

```
GET https://<endpoint>/subscriptions/{subscriptionId}/providers/Microsoft.Security/securityContacts?
api-version={api-version}
```

A Defender for Cloud user assigned the role of Account Administrator can export the just-in-time policies containing the IP addresses by:

- Copying from the Azure portal
- Executing the Azure REST API call, GET HTTP:

```
GET
https://<endpoint>/subscriptions/{subscriptionId}/resourceGroups/{resourceGroup}/providers/Microsoft.
Security/locations/{location}/jitNetworkAccessPolicies/default?api-version={api-version}
```

An Account Administrator can export the alert details by:

- Copying from the Azure portal
- Executing the Azure REST API call, GET HTTP:

```
GET https://<endpoint>/subscriptions/{subscriptionId}/providers/microsoft.Security/alerts?api-version={api-version}
```

For more information, see Get Security Alerts (GET Collection).

Restricting the use of personal data for profiling or marketing without consent

A Defender for Cloud user can choose to opt out by deleting their security contact data.

Just-in-time data is considered non-identifiable data and is retained for a period of 30 days.

Alert data is considered security data and is retained for a period of two years.

Auditing and reporting

Audit logs of security contact, just-in-time, and alert updates are maintained in Azure Activity Logs.

Defender for Cloud readiness roadmap

2/15/2022 • 2 minutes to read • Edit Online

> **NOTE**
> Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This document provides you a readiness roadmap that will assist you to get started with Defender for Cloud.

Understanding Defender for Cloud

Defender for Cloud provides unified security management and advanced threat protection for workloads running in Azure, on-premises, and in other clouds.

Use the following resources to get started with Defender for Cloud.

Articles

- Introduction to Defender for Cloud
- Defender for Cloud quickstart guide

Videos

- Quick Introduction Video
- Overview of Defender for Cloud Prevention, Detection and Response Capabilities

Planning and operations

To take full advantage of Defender for Cloud, it is important to understand how different individuals or teams in your organization use the service to meet secure operations, monitoring, governance, and incident response needs.

Use the following resources to assist you during the planning and operations processes.

- Defender for Cloud planning and operations guide

Onboarding computers to Defender for Cloud

Defender for Cloud automatically detects any Azure subscriptions or workspaces not protected by Microsoft Defender for Cloud. This includes Azure subscriptions using Defender for Cloud Free and workspaces that do not have the security solution enabled.

Use the following resources to assist you during the onboarding processes.

- Onboard non-Azure computers
- Defender for Cloud Hybrid - Overview

Mitigating security issues using Defender for Cloud

Defender for Cloud automatically collects, analyzes, and integrates log data from your Azure resources, the network, and connected partner solutions, like firewall and endpoint protection solutions, to detect real threats

and reduce false positives.

Use the following resources to assist you to manage security alerts and protect your resources.

Articles

- Protecting your network in Defender for Cloud
- Protecting Azure SQL service and data in Defender for Cloud

Defender for Cloud for incident response

To reduce costs and damage, it's important to have an incident response plan in place before an attack takes place. You can use Defender for Cloud in different stages of an incident response.

Use the following resources to understand how Defender for Cloud can be incorporated in your incident response process.

Videos

- Respond quickly to threats with next-generation security operation, and investigation

Articles

- Using Defender for Cloud for an incident response
- Use automation to respond to Defender for Cloud triggers

Advanced cloud defense

Azure VMs can take advantage of advanced cloud defense capabilities in Defender for Cloud. These capabilities include just-in-time virtual machine (VM) access, and adaptive application controls.

Use the following resources to learn how to use these capabilities in Defender for Cloud.

Videos

- Defender for Cloud – Just-in-time VM Access
- Defender for Cloud - Adaptive Application Controls

Articles

- Manage virtual machine access using just-in-time
- Adaptive Application Controls in Defender for Cloud

Hands-on activities

- Defender for Cloud hands-on lab
- Web Application Firewall (WAF) recommendation playbook in Defender for Cloud
- Defender for Cloud Playbook: Security Alerts

Additional resources

- Defender for Cloud Documentation Page
- Defender for Cloud REST API Documentation Page
- Defender for Cloud frequently asked questions (FAQ)
- Pricing page
- Identity security best practices
- Network security best practices
- PaaS recommendations

- Compliance
- Log analytics customers can now use Defender for Cloud to protect their hybrid cloud workloads

Microsoft Defender for Cloud Troubleshooting Guide

2/15/2022 • 6 minutes to read • Edit Online

NOTE

Azure Security Center and Azure Defender are now called **Microsoft Defender for Cloud**. We've also renamed *Azure Defender* plans to *Microsoft Defender* plans. For example, Azure Defender for Storage is now Microsoft Defender for Storage. Learn more about the recent renaming of Microsoft security services.

This guide is for information technology (IT) professionals, information security analysts, and cloud administrators whose organizations need to troubleshoot Defender for Cloud related issues.

Defender for Cloud uses the Log Analytics agent to collect and store data. See Microsoft Defender for Cloud Platform Migration to learn more. The information in this article represents Defender for Cloud functionality after transition to the Log Analytics agent.

TIP

A dedicated area of the Defender for Cloud pages in the Azure portal provides a collated, ever-growing set of self-help materials for solving common challenges with Defender for Cloud.

When you're facing an issue, or are seeking advice from our support team, **Diagnose and solve problems** is good place to look for solutions:

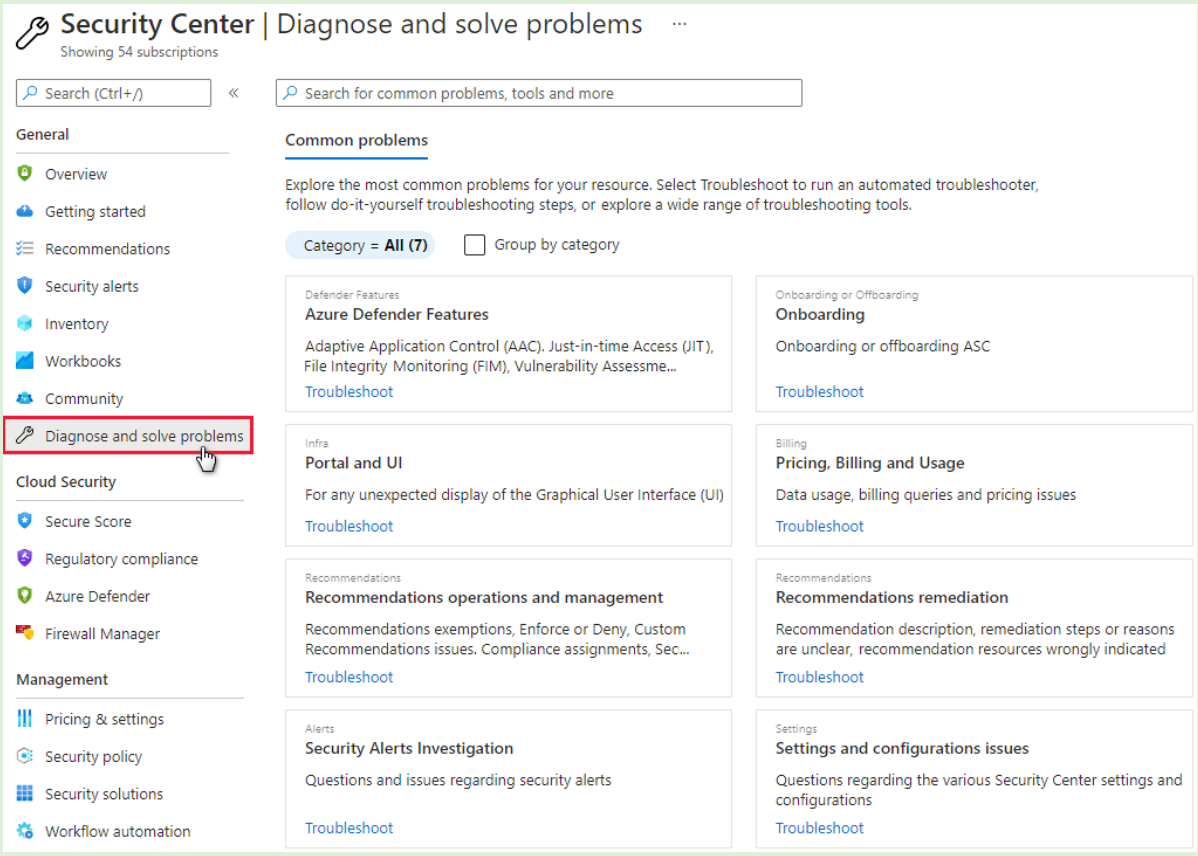

Troubleshooting guide

This guide explains how to troubleshoot Defender for Cloud related issues.

Alert types:

- Virtual Machine Behavioral Analysis (VMBA)
- Network Analysis
- SQL Database and Azure Synapse Analytics Analysis
- Contextual Information

Depending on the alert types, customers can gather the necessary information to investigate the alert by using the following resources:

- Security logs in the Virtual Machine (VM) event viewer in Windows
- AuditD in Linux
- The Azure activity logs, and the enable diagnostic logs on the attack resource.

Customers can share feedback for the alert description and relevance. Navigate to the alert itself, select the **Was This Useful** button, select the reason, and then enter a comment to explain which explains the feedback. We consistently monitor this feedback channel to improve our alerts.

Audit log

Most of the troubleshooting done in Defender for Cloud takes place by first looking at the Audit Log records for the failed component. Through audit logs, you can determine:

- Which operations were taken place
- Who initiated the operation
- When the operation occurred
- The status of the operation
- The values of other properties that might help you research the operation

The audit log contains all write operations (PUT, POST, DELETE) performed on your resources, however it does not include read operations (GET).

Log Analytics agent

Defender for Cloud uses the Log Analytics agent – this is the same agent used by the Azure Monitor service – to collect security data from your Azure virtual machines. After data collection is enabled and the agent is correctly installed in the target machine, the process below should be in execution:

- HealthService.exe

If you open the services management console (services.msc), you will also see the Log Analytics agent service running as shown below:

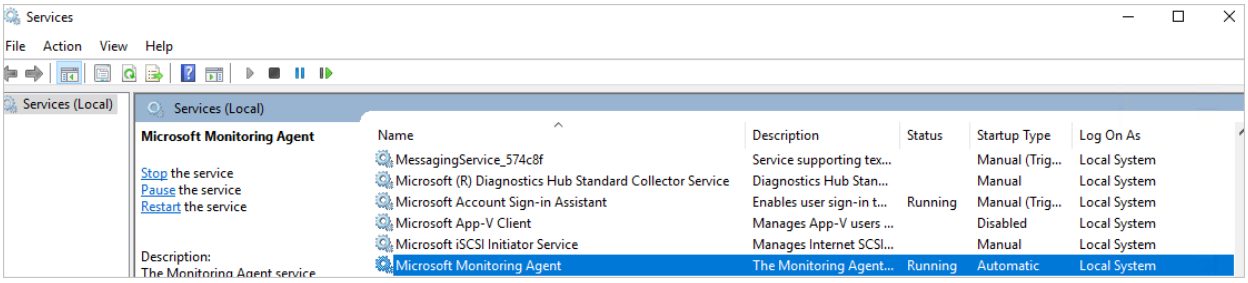

To see which version of the agent you have, open **Task Manager**, in the **Processes** tab locate the **Log Analytics agent Service**, right-click on it and click **Properties**. In the **Details** tab, look the file version as shown below:

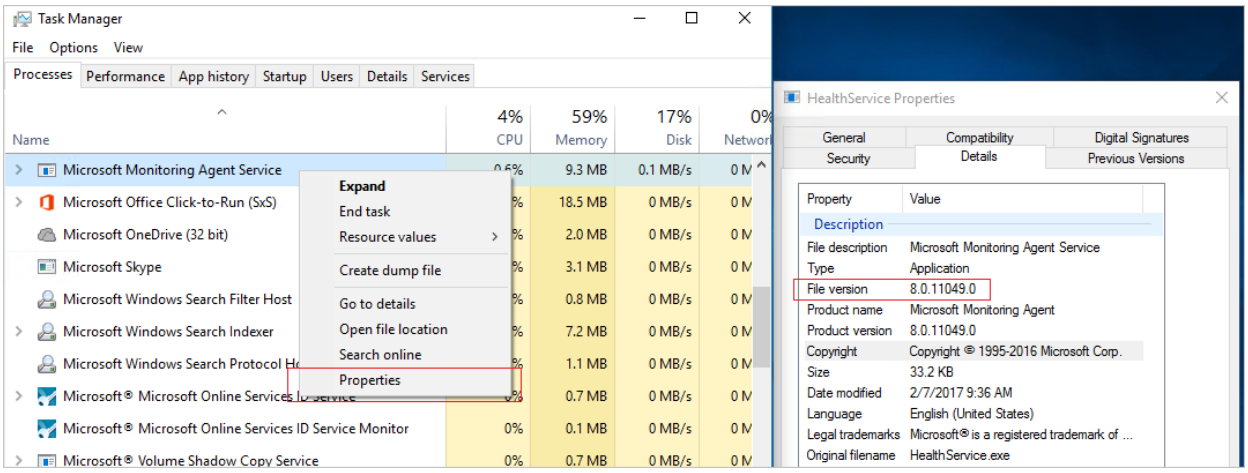

Log Analytics agent installation scenarios

There are two installation scenarios that can produce different results when installing the Log Analytics agent on your computer. The supported scenarios are:

- **Agent installed automatically by Defender for Cloud**: in this scenario you will be able to view the alerts in both locations, Defender for Cloud and Log search. You will receive email notifications to the email address that was configured in the security policy for the subscription the resource belongs to.
- **Agent manually installed on a VM located in Azure**: in this scenario, if you are using agents downloaded and installed manually prior to February 2017, you can view the alerts in the Defender for Cloud portal only if you filter on the subscription the workspace belongs to. If you filter on the subscription the resource belongs to, you won't see any alerts. You'll receive email notifications to the email address that was configured in the security policy for the subscription the workspace belongs to.

> **NOTE**
> To avoid the behavior explained in the second scenario, make sure you download the latest version of the agent.

Troubleshooting monitoring agent network requirements

For agents to connect to and register with Defender for Cloud, they must have access to network resources, including the port numbers and domain URLs.

- For proxy servers, you need to ensure that the appropriate proxy server resources are configured in agent settings. Read this article for more information on how to change the proxy settings.
- For firewalls that restrict access to the Internet, you need to configure your firewall to permit access to Log Analytics. No action is needed in agent settings.

The following table shows resources needed for communication.

AGENT RESOURCE	PORTS	BYPASS HTTPS INSPECTION
*.ods.opinsights.azure.com	443	Yes
*.oms.opinsights.azure.com	443	Yes
*.blob.core.windows.net	443	Yes
*.azure-automation.net	443	Yes

If you encounter onboarding issues with the agent, make sure to read the article How to troubleshoot Operations Management Suite onboarding issues.

Troubleshooting endpoint protection not working properly

The guest agent is the parent process of everything the Microsoft Antimalware extension does. When the guest agent process fails, the Microsoft Antimalware that runs as a child process of the guest agent may also fail. In scenarios like that is recommended to verify the following options:

- If the target VM is a custom image and the creator of the VM never installed guest agent.
- If the target is a Linux VM instead of a Windows VM then installing the Windows version of the antimalware extension on a Linux VM will fail. The Linux guest agent has specific requirements in terms of OS version and required packages, and if those requirements are not met the VM agent will not work there either.
- If the VM was created with an old version of guest agent. If it was, you should be aware that some old agents could not auto-update itself to the newer version and this could lead to this problem. Always use the latest version of guest agent if creating your own images.
- Some third-party administration software may disable the guest agent, or block access to certain file locations. If you have third-party installed on your VM, make sure that the agent is on the exclusion list.
- Certain firewall settings or Network Security Group (NSG) may block network traffic to and from guest agent.
- Certain Access Control List (ACL) may prevent disk access.
- Lack of disk space can block the guest agent from functioning properly.

By default the Microsoft Antimalware User Interface is disabled, read Enabling Microsoft Antimalware User Interface on Azure Resource Manager VMs Post Deployment for more information on how to enable it if you need.

Troubleshooting problems loading the dashboard

If you experience issues loading the workload protection dashboard, ensure that the user that registers the subscription to Defender for Cloud (i.e. the first user one who opened Defender for Cloud with the subscription) and the user who would like to turn on data collection should be *Owner* or *Contributor* on the subscription. From that moment on also users with *Reader* on the subscription can see the dashboard/alerts/recommendation/policy.

Contacting Microsoft Support

Some issues can be identified using the guidelines provided in this article, others you can also find documented at the Defender for Cloud public Microsoft Q&A page. However if you need further troubleshooting, you can open a new support request using **Azure portal** as shown below:

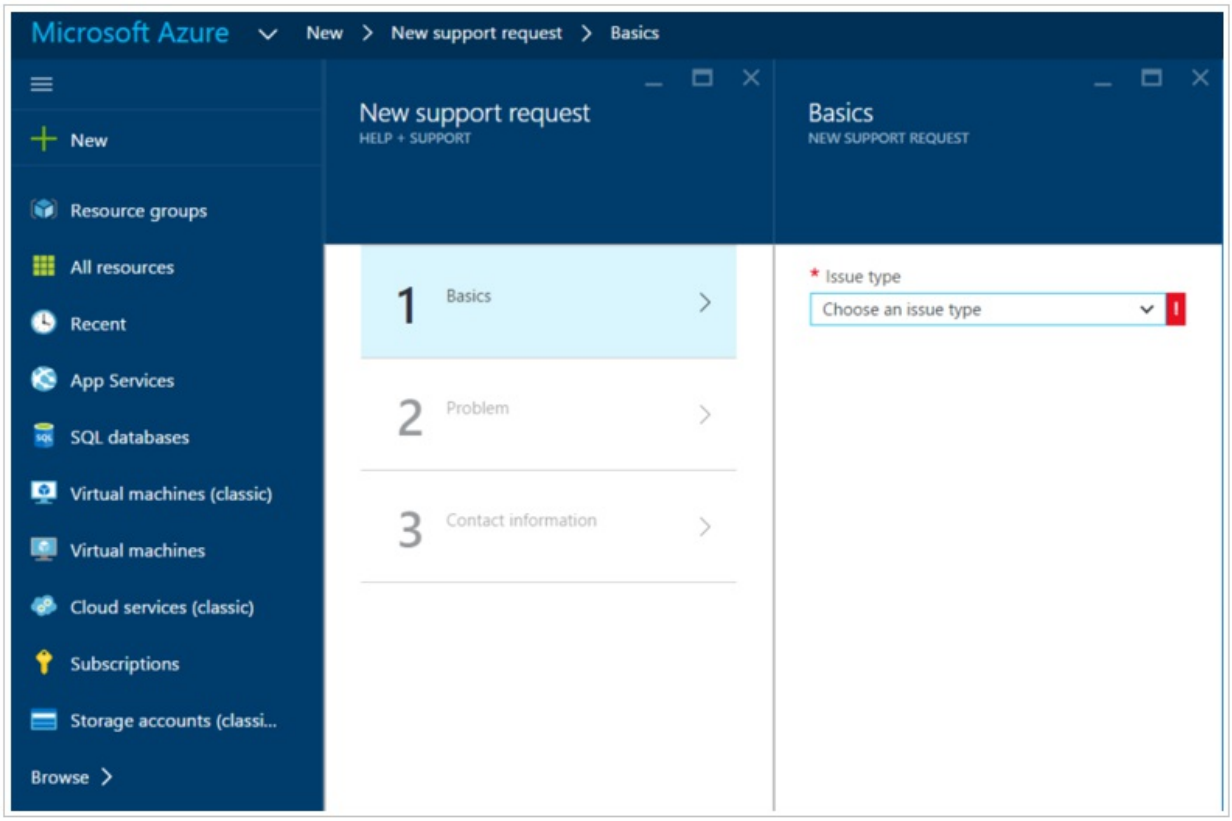

See also

In this page, you learned how to configure security policies in Microsoft Defender for Cloud. To learn more about Microsoft Defender for Cloud, see the following:

- Managing and responding to security alerts in Microsoft Defender for Cloud — Learn how to manage and respond to security alerts
- Alerts Validation in Microsoft Defender for Cloud
- Microsoft Defender for Cloud FAQ — Find frequently asked questions about using the service