



Solution Brief

Simplify Jenkins application and data lifecycle management

Deliver multi-cloud data protection, restoration, and portability to Jenkins using Astra Control in the public cloud and on-premises

Diane Patton, NetApp
January 2022

TABLE OF CONTENTS

Introduction	3
Astra Control Overview	3
Managing Jenkins with Astra Control	4
Persisting Jenkins data	5
Cloning Jenkins to a different cluster	6
About NetApp.....	9
Legal Notice	9

Introduction

Enterprises are moving from a monolithic software development approach to an agile one that encompasses continuous integration and continuous delivery (CI/CD). The CI/CD method provides for smaller changes to code very frequently. Immediately after small changes are checked in and merged, new builds and testing are completed on the new code via automation. This process allows developers to identify any potential issues easier and faster, as well as to get new updates to customers sooner. Jenkins is an open source tool that can be used to automate many different tasks related to building, testing, delivering, and/or deploying software. It also stores build metadata, testing results and may also store build artifacts.

When Jenkins is deployed on a Kubernetes cluster, the volumes holding the pipeline and other artifact data and test results must always remain persistent. A new software build might not pass the automated tests, and developers might need additional information regarding test results, so all the data must be continuously available. Further, the pipeline, with its build metadata and artifacts, may need to be relocated to a new cluster for portability and/or disaster recovery. Astra Control provides automated persistence and cloning for these scenarios.

Use Astra Control to:

- Persist your Jenkins pipeline definitions, artifacts, test messages and console output.
- Restore all your Jenkins data and Kubernetes resources from a disaster scenario.
- Maintain your build data, pipeline configurations, and artifacts when you relocate to a new Kubernetes cluster or namespace in the same or different public cloud.

Astra Control Overview

Astra Control is a fully managed service or customer managed on-premises application that makes it easier for our customers to manage, protect, and move their data-rich containerized workloads running on Kubernetes within and across public clouds and on-premises. Astra Control provides persistent container storage that leverages NetApp's proven and expansive storage portfolio in the public cloud and on premises. It also offers a rich set of advanced application-aware data management functionality (like snapshot, revert, backup and restore, activity log, and active cloning) for data protection, disaster recovery, data audit, and migration use cases for your modern apps.

Managing Jenkins with Astra Control

Astra Control Service (ACS) provides management, protection, and cloning for Google Kubernetes Engine (GKE) or Azure Kubernetes Service (AKS) clusters located in the public cloud. Astra Control Center (ACC) provides the same experience and functionality for RedHat Openshift Container Platform (OCP), Rancher or Upstream clusters located on-premises.

Simply add your Kubernetes clusters in the cloud to ACS or on premises to ACC.

Upon adding a cluster, ACS:

- Installs NetApp Trident, NetApps open source Kubernetes storage orchestrator and three kubernetes storage classes using a NetApp backend in your cloud provider.
- Creates a bucket on the cloud object store for future backups for all your registered clusters. You can also add your own bucket.
- Creates a service account on your cluster for itself

ACC on premises (supported with RedHat Openshift Container Platform, Rancher and Upstream Kubernetes) uses your current Trident installation, Trident based storage classes, ONTAP backend, and allows you to import your own object storage bucket for backups and cloning.

As an example using ACS, Figure 1 shows two clusters, one GKE cluster located in the Google Cloud Platform (GCP) region us-west2 (Los Angeles) and one AKS cluster located in Azure region us-east (Virginia).

Figure 1) Registered clusters.

Clusters

Actions

+ Add

Search

1-2 of 2 entries


< >

<input type="checkbox"/>	Name	Ready	Type	Version	Location	Actions
<input type="checkbox"/>	cluster-1-patd	<div><div></div></div>	<div><div></div>Google Kubernetes Engine</div>	v1.21.5-gke.1302	us-west2-b	<div>Running</div>
<input type="checkbox"/>	cluster-patd-3	<div><div></div></div>	<div><div></div>Azure Kubernetes Service</div>	v1.21.7	eastus	<div>Running</div>

After your cluster is registered, install Jenkins on your cluster using the current Bitnami Helm chart at [bitnami/Jenkins](#) or a custom manifest. (Support for the Jenkins Kubernetes plug-in is coming soon. See [Cloning Jenkins Kubernetes plugin with Astra](#)). Trident automatically generates the Kubernetes Persistent Volumes (PVs) and NetApp backend volumes that Jenkins needs.

Astra Control discovers all the applications on your Kubernetes clusters. You can manage together just the application, all the resources in the entire namespace, or a custom group based on object labels. Just use the drop down menu to manage the application. Figure 2 shows that we are managing all resources in the entire Jenkins namespace as one unit.

Figure 2) Managing the Jenkins namespace.



Log in, configure, and use Jenkins as you normally would; you can create builds, configure pipelines, etc

Persisting Jenkins data

All Jenkins data can be automatically persisted by using snapshots and backups. You can add your own execution hooks if desired to provide application consistent snapshots. Astra Control snapshots and backups preserve the application, its metadata, and its volumes in one easily manageable unit. The PV snapshot is stored in NetApp backend service, and the application snapshot is stored in Astra. All PV backups are stored in the object store.

Both on-demand and scheduled snapshots and backups are supported. You can set up a snapshot and backup schedule for the volume and all the Kubernetes objects that are associated with Jenkins, as shown in Figure 3.

Figure 3) Configuring the protection policy.

The screenshot shows the 'Configure protection policy' interface. The title bar includes a shield icon, the text 'Configure protection policy', and 'STEP 1/2: DETAILS'. The main content area is divided into two sections: 'PROTECTION SCHEDULE' and 'BACKUP DESTINATION'. In the 'PROTECTION SCHEDULE' section, there is a 'Hourly' option selected, with a description: 'Every hour on the 0th minute, keep the last 4 snapshots, keep the last backup'. Below this, there are four radio buttons: 'Hourly' (selected), 'Daily', 'Weekly', and 'Monthly'. A red box highlights the 'Hourly' section, which includes a 'Time (optional)' dropdown set to 'On the hour', a 'Snapshots to keep' field set to 4, and a 'Backups to keep' field set to 1. The 'BACKUP DESTINATION' section shows a 'Bucket' dropdown with a long alphanumeric string and the text 'Available' and 'Default'. On the right side, there is a sidebar titled 'CONFIGURING PROTECTION POLICIES' with instructions and a list of resources: 'Application jenkins', 'Namespace jenkins', and 'Cluster cluster-patd-3'. At the bottom, there are 'Cancel' and 'Next' buttons.

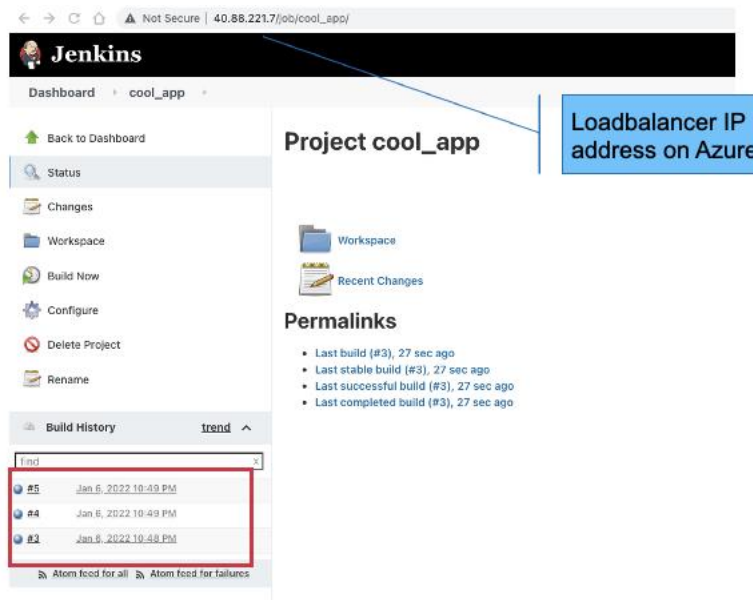
Astra automatically takes snapshots and backups based on the schedule you established. Of course, you can always take snapshots and backups on demand. You can even back up from a previous snapshot.

Cloning Jenkins to a different cluster

Now you are protected from a disaster, such as losing the cluster or accidentally deleting the namespace since you can restore from backup to a new cluster, namespace, or the same namespace on the same cluster. With ACS, you also have the ability to redeploy Jenkins to a new namespace or a new cluster on the same or different public cloud state with a direct clone. ACC allows you to redeploy/restore to the same or a new namespace on the same on premises cluster or different on premises cluster.

For example, suppose that you have a project located in the cloud being managed by ACS that will soon change ownership. The current team in Virginia uses Azure AKS and already has Jenkins configured with a link to the proper GitHub repository, artifacts, and a testing pipeline setup. They also have stored build metadata. You would like to redeploy Jenkins, along with all its data, closer to the new team in Los Angeles that use Google GKE. Jenkins is currently running on the *cluster-patd-3* cluster in the us-east (Virginia) Azure region and has three builds already made as shown in Figure 4.

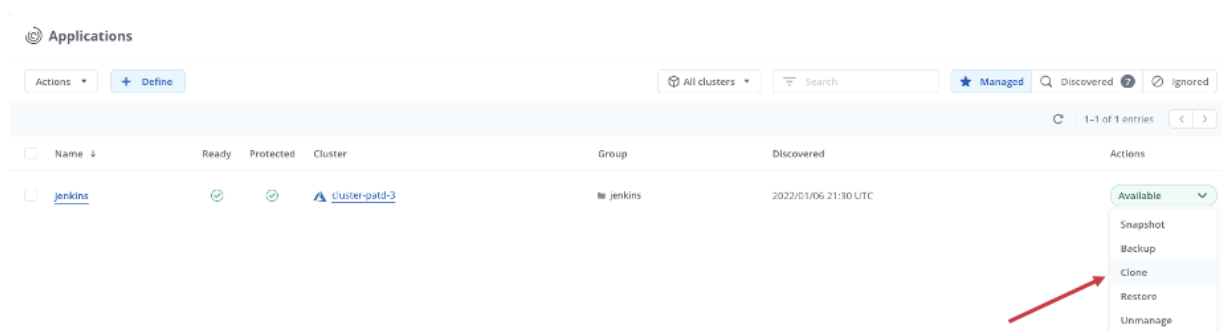
Figure 4) Current Jenkins state on the AKS cluster.



Ensure a new cluster is registered with ACS. Our example in Figure 1 shows *cluster-1-patd* registered in GCP us-west2 (Los Angeles) region.

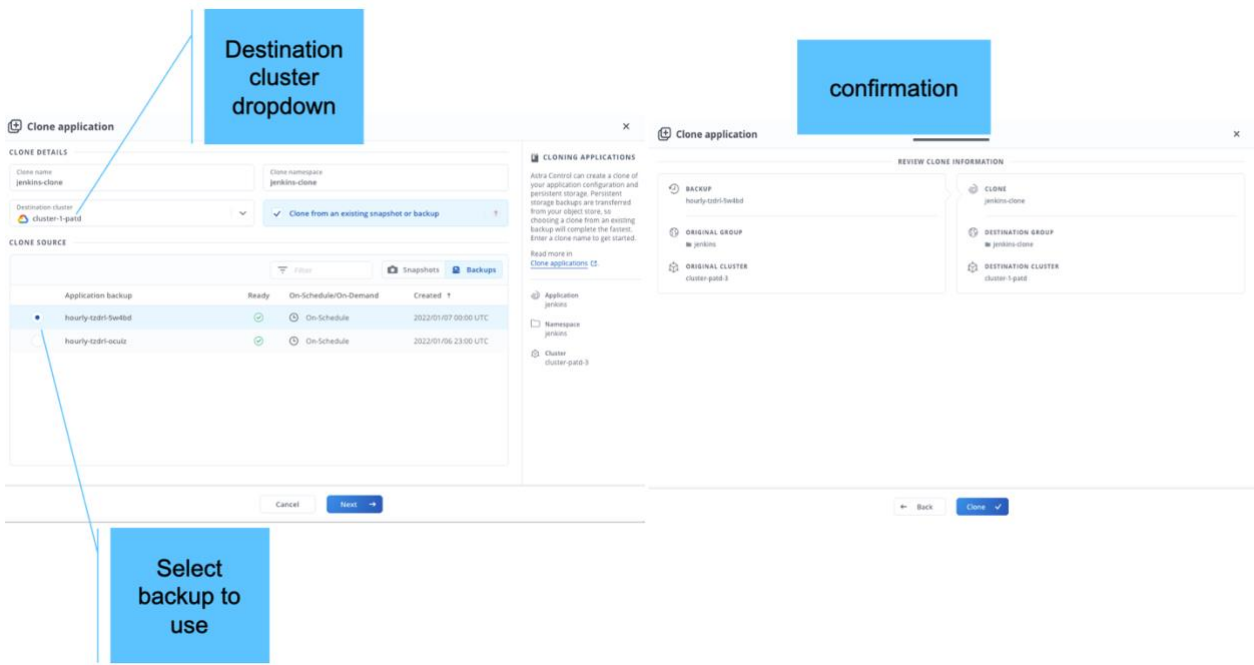
Next, use the drop down menu to clone Jenkins and its resources as depicted in Figure 5.

Figure 5) Clone Jenkins and its resources.




Clone Jenkins to the new cluster, *cluster-1-patd*, using a backup as shown in Figure 6. Cloning from backup brings up a new instance of Jenkins on the new cluster, running at the same state as when the backup was taken.

Figure 6) Cloning Jenkins using a backup.



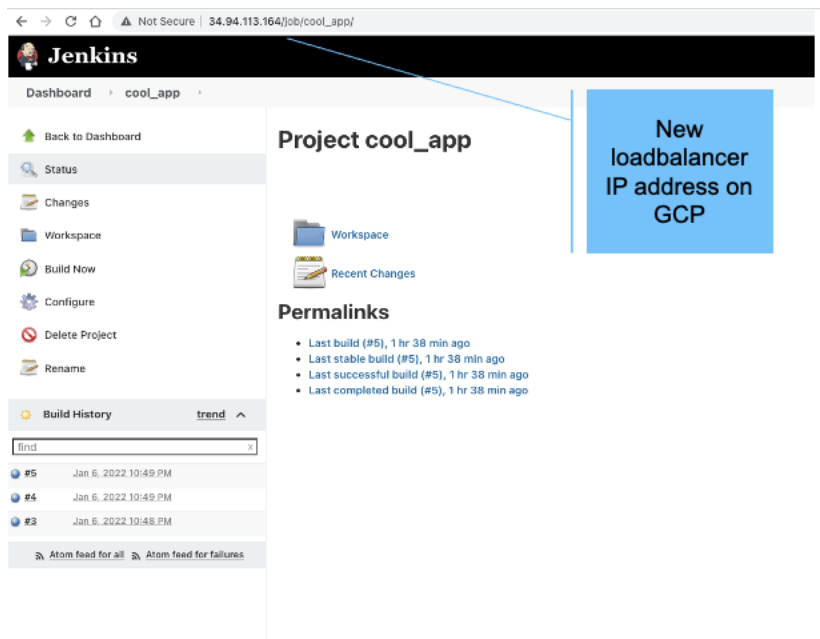
You could also clone at the current state, which integrates the backup and cloning together. The new Jenkins clone on the new cluster will be provisioned and managed automatically by Astra Control, as shown in Figure 7.

Figure 7) Provisioning the new Jenkins instance on a new cluster.



The new Jenkins instance will have all the data that was there with the cloned instance as shown in Figure 8. You can run a new build on the clone on the GKE cluster without having to reconfigure the Jenkins project from the beginning.

Figure 8) New Jenkins instance in a new cluster.



About NetApp

In a world full of generalists, NetApp is a specialist. We're focused on one thing, helping your business get the most out of your data. NetApp brings the enterprise-grade data services you rely on into the cloud, and the simple flexibility of cloud into the data center. Our industry-leading solutions work across diverse customer environments and the world's biggest public clouds.

As a cloud-led, data-centric software company, only NetApp can help build your unique data fabric, simplify and connect your cloud, and securely deliver the right data, services and applications to the right people—anytime, anywhere.

Legal Notice

Refer to the [Interoperability Matrix Tool \(IMT\)](#) on the NetApp Support site to validate that the exact product and feature versions described in this document are supported for your specific environment. The NetApp IMT defines the product components and versions that can be used to construct configurations that are supported by NetApp. Specific results depend on each customer's installation in accordance with published specifications.

Copyright Information

Copyright © 2022 NetApp, Inc. All Rights Reserved. Printed in the U.S. No part of this document covered by copyright may be reproduced in any form or by any means—graphic, electronic, or mechanical, including photocopying, recording, taping, or storage in an electronic retrieval system—without prior written permission of the copyright owner.

Software derived from copyrighted NetApp material is subject to the following license and disclaimer:

THIS SOFTWARE IS PROVIDED BY NETAPP “AS IS” AND WITHOUT ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY AND FITNESS FOR A PARTICULAR PURPOSE, WHICH ARE HEREBY DISCLAIMED. IN NO EVENT SHALL NETAPP BE LIABLE FOR ANY DIRECT, INDIRECT, INCIDENTAL, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES (INCLUDING, BUT NOT LIMITED TO, PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY OUT OF THE USE OF THIS SOFTWARE, EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

NetApp reserves the right to change any products described herein at any time, and without notice. NetApp assumes no responsibility or liability arising from the use of products described herein, except as expressly agreed to in writing by NetApp. The use or purchase of this product does not convey a license under any patent rights, trademark rights, or any other intellectual property rights of NetApp.

The product described in this manual may be protected by one or more U.S. patents, foreign patents, or pending applications.

Data contained herein pertains to a commercial item (as defined in FAR 2.101) and is proprietary to NetApp, Inc. The U.S. Government has a non-exclusive, non-transferrable, non-sublicensable, worldwide, limited irrevocable license to use the Data only in connection with and in support of the U.S. Government contract under which the Data was delivered. Except as provided herein, the Data may not be used, disclosed, reproduced, modified, performed, or displayed without the prior written approval of NetApp, Inc. United States Government license rights for the Department of Defense are limited to those rights identified in DFARS clause 252.227-7015(b).

Trademark Information

NETAPP, the NETAPP logo, and the marks listed at <http://www.netapp.com/TM> are trademarks of NetApp, Inc. Other company and product names may be trademarks of their respective owners.

