tierpoint

2020 Edition

Whitepaper

Why Adopt DRaaS?

5 Considerations to Maximize Business Benefits

Disaster recovery as a service (DRaaS) is a crucial cloud service used by businesses of all sizes. The demand for DRaaS is projected to continue growing at a rate of 44% per year through 2023[1], with an even higher growth rate in some industries such as financial services and telecommunications.[2]

DRaaS enables organizations to safeguard their digital assets and significantly speed their recovery from a disaster or IT disruption. Managing production data is difficult enough for IT departments without the added burden of replicating data and workloads and maintaining additional equipment.

INTRODUCTION

DRaaS replicates an organization's primary IT systems, applications and data to a cloud environment. If disaster strikes, data and applications can then be restored from the cloud. Alternatively, the cloud environment can serve as a production environment for employees and customers until the primary site is restored.

DRaaS adoption continues to grow rapidly due to several important factors.

One factor is the escalating volume of data that companies are accumulating. IoT devices, the adoption of popular apps, increased connectivity from new networking tech (like 5G), and the growing use of AI for business and consumer applications, are all contributing to this huge data growth. A 2018 survey by the Enterprise Storage Forum found that approximately 38% of companies expect to need 100 terabytes or more of additional storage by the end of 2020, while 23% anticipate needing a petabyte or more of extra data storage.[3] At least half of that data will be going into cloud storage, as IT and business leaders have become comfortable trusting their data to the cloud.[4]

One third of organizations surveyed lost $1 million or more per hour of downtime.

Concern about the rising rate of cybercrime also contributes to the growth in DRaaS adoption. In 2019, ransomware attacks on businesses rose by 365%, encrypting or destroying data and bringing IT systems to a screeching halt. DRaaS provides a readily available backup to mitigate the damage from ransomware and other cyberattacks.[5]

Finally, businesses today need to maintain higher levels of uptime in their IT systems. The cost of downtime can be significant. And while hurricanes and cyber-attacks often get media coverage, by far the most frequent causes of downtime are power outages, network problems and IT system failures, according to Research from the Uptime Institute found that are the big three causes of downtime.[6] Those are the downtime events that happen all the time and rack up major expenses.

Disaster recovery is essential for minimizing disruption to the business and protecting its data. DRaaS is playing a part in many organizations' data loss prevention strategy. IT budgets are expected to increase by 50% in 2020 with over 7% of that budget dedicated to disaster recovery.

This paper examines the benefits of, and key issues involved in, DRaaS adoption.

DRaaS offers several advantages over on-premises DR at a secondary facility. These advantages include:

- Rapid implementation
- No hardware or maintenance costs
- Rapid recovery with minimal downtime
- Nearly complete recovery of data
- Predictable, operational costs with no capital expense
- The flexibility of provider managed or self-administered DRaaS

1. DRaaS Options for Multiple Needs

DRaaS first gained popularity among small and mid-size organizations with limited budgets and IT staff. Today, however, organizations of all sizes are moving away from traditional DR and using DRaaS for fast, enterprise-class, disaster recovery without the need to own and maintain an off-site data center. Brian Anderson, vice president of product management at TierPoint says there's been a "groundswell" of adoption over the past two years, partly due to the increased maturity level of DRaaS solutions.

> ***"DRaaS technology today is more mature and feature-rich than five or six years ago."***
>
> Brian Anderson
> Vice President of Product Management, TierPoint

FIGURE 1: How Cloud-Based Disaster Recovery as a Service (DRaaS) Works

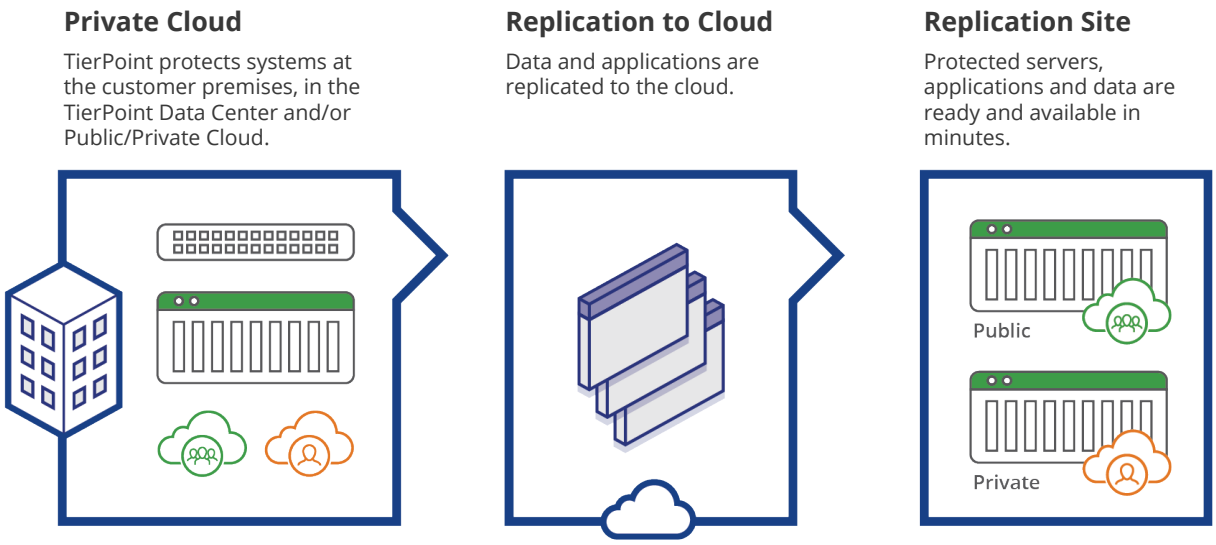

The scalability and flexibility of DRaaS also makes it an attractive alternative to do-it-yourself DR.

A customer will need to assess their workloads, apps and environments for their acceptable Recovery Point Objective (RPO) and Recovery Time Objective (RTO). This will help the customer find the right recovery strategy for their business and identify the amount of data loss the business is willing to tolerate, how rapidly it needs an application restored, and what the budget can afford. For example, customer-facing applications are usually high priority and need fast recovery. These would be in a high priority tier, as would network services and email. Less important applications would be on a lower, and less expensive, tier.[7]

"There are multiple flavors of DRaaS," says Anderson. "There's no one-size-fits-all."

DRaaS providers typically offer additional services beyond initial support and implementation, including testing, runbook creation, change management, business continuity planning, and day-to-day management. Those can be useful services for IT departments that lack DR expertise, especially in more complex hybrid and multicloud environments.

> ***"There are multiple flavors of DRaaS. There's no one-size-fits-all."***
>
> Brian Anderson
> Vice President of Product Management, TierPoint

More organizations today have, or will soon have, a hybrid IT environment, which might include multiple public cloud services, private cloud services, and on-premises or collocated servers for legacy systems. Managing disaster recovery for such complex environments requires expertise in cloud services and hybrid architectures, as well as in disaster recovery.

The bottom line is that DRaaS enables an organization to outsource some or all its DR implementation.

2. Understanding RTO and RPO

DRaaS allows for automated and almost instantaneous **failover** to one or more clouds. When the primary site fails, control is automatically switched to the cloud site with the replicated data. Later, when the outage is resolved, DRaaS lets the organization return control to the primary site through a process called **failback** that ensures data stays current.

How fast an organization needs to bring its systems back online and how much data it can afford to lose are measured by the recovery time objective (RTO) and the recovery point objective (RPO), respectively. It's important to understand these before designing a DR plan or evaluating DRaaS options.

Recovery Time Objective (RTO).

The RTO defines how long it will take an organization to get their servers back up. This measure is critical to business operations. Without access to applications and data, few organizations today can function.

In conventional disaster recovery implementations, an RTO in the 24- to 48-hour range is common. That's because conventional DR is based on deploying a secondary set of physical servers in a geographically dispersed location. Each night, the organization backs up to those servers. When disaster strikes the primary center, IT staff need time to bring up the servers from the secondary site and restore from backups.

That approach may work for businesses that don't need to operate 24x7 and that can afford to lose an hour or more of data. For businesses in finance, healthcare or ecommerce, or any customer-centric, 24-hour business, a 24- to 48-hour RTO is far too long.

The cloud-based automation behind DRaaS guarantees a shorter RTO. Without having to bring up servers and restore from backups, DRaaS can compress traditional disaster recovery processes from days down to an hour or less for hybrid environments, and faster for all cloud environments. With a cloud to cloud recovery, it can be a little as 10 minutes to get back up and running. DRaaS can have a similar impact with *recovery point objectives*.

Recovery Point Objective (RPO).

RPO defines how much information is saved or lost once servers are restored. For example, does an organization recover to a point 15 minutes back (therefore losing only 15 minutes of data) or 24 hours back? Note that for organizations without a disaster recovery plan, the RPO is effectively never.

In traditional DR, the communication between the primary production environment and the secondary site typically happens on a set schedule, such as every night. That creates a 24-hour RPO.

In a DRaaS scenario, however, the two environments can stay in near-constant contact, with bandwidth availability as the only major constraint. As a result, not only can it take as little as l0 minutes to get back and up and running in a cloud-to-cloud scenario, it's possible to lose little to no data. A low RPO can be particularly important with certain workloads or when dealing with data storage requirements under compliance regulations.

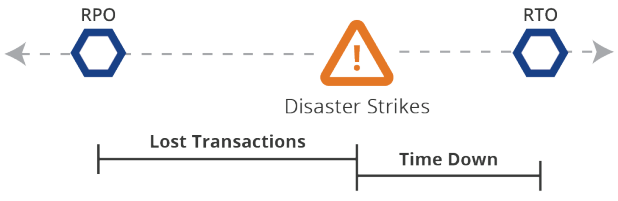

Short RPO (15 minutes)
Very little data loss is acceptable.

Short RTO (15 minutes - 4 hours)
Systems are restored quickly.

Long RPO (24 hours)
Data is updated less frequently.

Long RTO (1-2 days)
It's no longer necessary to wait this long.

Synchronous vs Asynchronous.

The level of recovery that an organization chooses, will determine the type of replication used in the DRaaS implementation--asynchronous or synchronous.

- Synchronous replication copies data to both primary storage and to DRaaS at the same time, so both have near-identical copies of data. Synchronous is used for situations requiring a rapid RTO and little to no data loss.

- Asynchronous replication copies data first to primary storage, then to the DR service. This is often used for DRaaS environments that can tolerate a slower recovery and greater data loss. It's also ideal for DR sites that are far from the primary site, due to the lesser bandwidth requirements of asynchronous replication.

"It's much simpler to budget for DRaaS than for other types of DR solutions."

Brian Anderson
Vice President of Product Management, TierPoint

3. Shifting to a Predictable and Inclusive Cost Model

When deciding what DR solution to implement, organizations should consider that many disruptions that cause downtime are, in fact, fairly common occurrences, such as power outages and human error. Others, such as a fire or flood, are usually far less common, though they can be significantly more disruptive and cause longer downtime.

A conventional off-site DR implementation is not cheap. It requires capital expenditures for the initial hardware and software, with regular intervals of hardware upgrades to meet growing volumes of data and workloads. Then there are operational expenses for software and hardware maintenance, IT staff costs, utilities, and other costs needed to keep a secondary IT system running. Using a colocation provider as a separate DR site lowers the staff and infrastructure costs, but still requires the customer to purchase and manage the software and hardware.

DRaaS shifts the total cost of ownership (TCO) into a commodified operational expense that focuses on performance instead of trying to predict hardware needs years in advance. Organizations can plan around a fixed monthly cost that might fluctuate only slightly if additional resources are needed to offset heavy demand.

"A benefit of DRaaS is that the costs are predictable and inclusive," said Anderson. "It's much simpler to budget for DRaaS than for other types of DR solutions."

DRaaS costs have also fallen significantly over the past several years and compare very favorably next to do-it-yourself DR.

4. Fully Managed or Self-Service DRaaS

An IT department will need to decide on the level of managed services it requires. Those can range from fully managed to completely self-service.

"Before a disaster, self-service usually makes more sense by giving companies immediate access and fine-tuned control," says Anderson. "But when disaster actually strikes, it can be invaluable to have access to managed services and enterprise-grade support staff as needed."

DRaaS offers that level of flexibility. The three main categories of DRaaS services are:

Self-service, as you might expect, is entirely managed by the customer. Gartner defines self-service DRaaS as being when a customer uses the DRaaS provider's infrastructure and tools to take on responsibilities such as failover and fallback procedures, handling recovery configuration activation and shutdown, and managing the virtual machine (VM) replication.[8] The DRaaS company provides the tools and the customer is responsible for using them.

Assisted recovery DRaaS gives the customer the responsibility for activating the recovery process in the event of a disaster but relies on the provider to handle the recovery infrastructure and manage data replication. The DRaaS provider takes a supporting role only.

Managed DRaaS. Fully managed DRaaS is ideal for IT departments that lack the time or expertise to manage DR. The DRaaS vendor is responsible for implementing, managing and executing a DR plan, managing the failover and fallback processes and doing all the other activities that are part of running a DR environment.

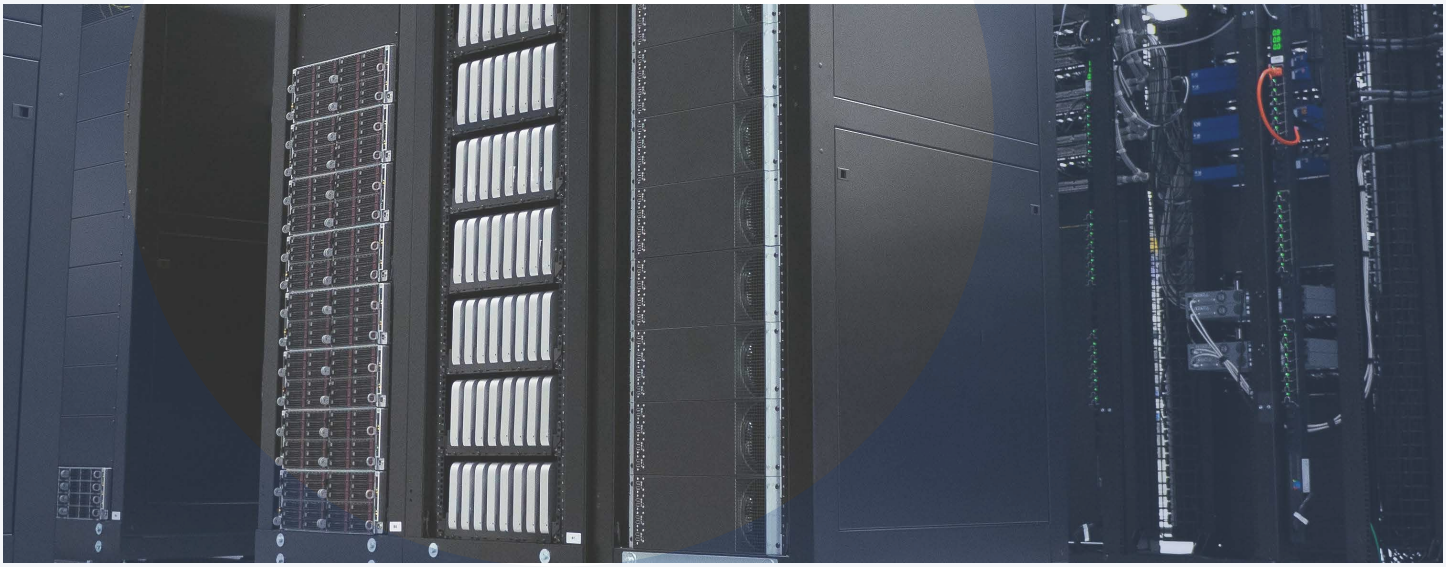

5. Getting Started with DRaaS: Next Steps

Draft your DR plan. Analyze the risks of common disruptions and how they might affect the organization. Take an inventory of applications and data, and the level of RTO and RPO each requires.

"Disaster recovery planning requires prioritizing to make sure the most important applications are restored fast enough," says Anderson.

Also consider business continuity. If employees can't get to the office or the data center, how will they do their work? Some DRaaS providers have business continuity services, such as temporary offices with phone service and internet. Ideally, your solution provider should be able to help you formulate a business continuity plan and prioritize your applications and data.

Evaluate solution providers. Flexibility is key to DRaaS. Every organization has different, and often complex, requirements. With the rise of mixed-platform, multicloud and hybrid environments, a DRaaS implementation must be able to accommodate a range of needs.

- Look for DRaaS providers that have partnerships with leading DR software vendors as well as multiple cloud providers. Having a menu of partners generally indicates they can deliver a more flexible, customizable solution. It also ensures you won't be locked into a single vendor choice, which may not fit your organization's needs.
- Check their infrastructure reliability. Do they have redundant backup generators, power feeds and connections to other carrier networks?
- Verify that compliance certifications meet your industry's requirements and that the provider has experience with regulatory compliance in your industry. Ask for references from organizations in your industry or a similar one.
- Scrutinize the service level agreement (SLA) to ensure it spells out your specific requirements. Have your lawyer and key department heads examine it as well. Find out how data will be transferred and where it will be stored, how long the recovery environment will be available to you, and what the risk of data loss is.

> ***"Disaster recovery planning requires prioritizing to make sure the most important applications are restored fast enough."***
>
> Brian Anderson
> Vice President of Product Management, TierPoint

- Ask about testing and updating of the DR environment. Simply passing a DR audit is not enough. You should update your DR plan regularly and incorporate any changes to your environment. You should also be able to test your DR plan at once a year, preferably more often, to ensure it is ready and that changes to the IT environment have not been overlooked.

Leverage your DRaaS environment. IT departments can leverage their DRaaS environments for other uses, such as application development and the testing.[9] Before making a change to a production environment, developers can first test it in a DRaaS environment. DRaaS can provide a sandbox for safely testing updates, security patches and new application development.

Alternatively, if the production systems need repair or major upgrades, DRaaS can enable the work to be done without any downtime for end-users. IT can failover the production applications to the DRaaS environment, allowing users to work from the cloud. The production environment can be failed back after work is completed.

Conclusion

Today's organizations have diverse and increasingly complex IT environments. A DRaaS implementation must support a range of IT environments as well as meet regulatory compliance, provide strong data security, and, of course, provide the recovery time that the business needs. It's not a job that most organizations can do on their own, not without great investments in hardware, IT staff and continuous maintenance and testing.

That's where DRaaS comes in. A DRaaS provider should be a trusted partner that can help you map out and execute a disaster recovery strategy as well as provide services and support where you most need them.

Almost no organization today can afford to be without access to its IT applications and data for very long. Many can't afford any downtime at all. But disasters and disruptions do happen, on a regular basis. Which is why disaster recovery and DRaaS is a vital component of every IT operation.

Whether your organization is faced with a ransomware attack, a power outage, fire, flood or hardware failure, it needs a backup plan to protect valuable data and IT systems. A DRaaS partner can help minimize the impact of unexpected disruptions and ensure maximum resiliency for your business.

Learn more about TierPoint Disaster Recovery Services. Contact us today.

References

[1] Disaster Recovery as a Service Market (Market Research Future, January 2019), https://www.marketresearchfuture.com/reports/disaster-recovery-service-market-3230

[2] Disaster Recovery as A Service (DRaaS) Market | Global Market Outlook 2018-2026 (Inkwood Research), https://www.inkwoodresearch.com/reports/disaster-recovery-as-a-service-market/

[3] Taylor, C. "Survey Reveals Tech Trends Reshaping Data Storage" (Enterprise Storage Forum, August 2018) https://www.enterprisestorageforum.com/storage-management/survey-reveals-tech-trends-reshaping-data-storage.html

[4] Ibid.

[5] Hines, D. "How Disaster Recovery Changes the Ransomware Game" (TierPoint Blog, September 28, 2018) https://blog.tierpoint.com/how-disaster-recovery-changes-the-ransomware-game

[6] Risk & Resiliency 2018 Report, Uptime Institute, https://uptimeinstitute.com/data-center-outages-are-common-costly-and-preventable

[7] The Strategic Guide to Disaster Recovery and DRaaS, 2019, TierPoint, https://www.tierpoint.com/the-strategic-guide-to-disaster-recovery-and-draas/

[8] Magic Quadrant for Disaster Recovery as a Service (Gartner, July 2018) http://web.tierpoint.com/gartner-magic-quadrant-disaster-recovery-2018-blog

[9] Hines, D. "Mitigate Cyber Threats with DRaaS" (TierPoint Blog, September 20, 2018) https://blog.tierpoint.com/mitigate-cyber-threats-with-draas

Learn more about how TierPoint can help you with your DRaaS Strategy.

Contact us today.

Call: 877.859.TIER (8437)
E-mail: sales@tierpoint.com
Visit: tierpoint.com

About TierPoint

A leading national provider of hybrid IT solutions, TierPoint helps organizations drive performance and manage risk. No U.S. provider comes close to matching TierPoint's unique combination of thousands of clients; more than 40 edge-capable data centers and 8 multitenant cloud pods coast to coast; and a comprehensive portfolio of cloud solutions, colocation, disaster recovery, security and other managed IT services. With white-glove customer service, TierPoint professionals customize and manage agile solutions that address each client's unique needs.