CloudEndure
An AWS Company

aws

Affordable Enterprise-Grade Disaster Recovery Using AWS

Until recently, enterprise-grade disaster recovery had been prohibitively expensive for most organizations. Thanks to the rapid development of cloud infrastructure, organizations can now attain top-of-the-line DR capabilities at a fraction of the cost.

The Challenge

An enterprise-grade disaster recovery (DR) solution is no longer something that is "nice to have." Unexpected downtime not only can damage your hard-earned reputation, it can also have detrimental financial implications. In fact, Gartner estimates the average cost of IT downtime at $5,600 per minute, which can add up quickly.

If your organization is like other large-scale businesses today, you understand the critical need to recover rapidly from IT outages, application failures, or malicious attacks in order to ensure business resilience, stay competitive, and avoid regulatory penalties. Your global customers expect constant availability, while your employees need reliable access to company systems to keep your business going.

Despite the demand for availability, it turns out that for many organizations, traditional on-premises DR solutions — especially those that enable minimal recovery point objectives (RPOs) and recovery time objectives (RTOs) — are prohibitively expensive due to heavy capital expenditures and costly duplicate third-party software and services.

As a result, some organizations choose to take the risk of having only a backup system, which enables data retrieval, but does not prevent costly downtime because of its long data recovery times. Other companies choose to protect only the most essential servers, which leaves their businesses vulnerable. Many companies that lay out a substantial initial investment in DR later dedicate these resources to other pressing IT needs. But "set it and forget it" does not work for DR — it is a system that quickly becomes obsolete if not tested frequently.

The solution to this challenge lies in the cloud. By moving their DR to the cloud, businesses can now attain top-of-the-line IT resilience at a fraction of the cost.

Here we compare on-premises DR to cloud-based DR — with a focus on the expected costs of each strategy. We will also touch upon the benefits and risks of each strategy to help you decide if leveraging the cloud, and specifically Amazon Web Services (AWS), for DR is the right approach for your business.

The solution to this challenge lies in the cloud. By moving their DR to the cloud, businesses can now attain top-of-the-line IT resilience at a fraction of the cost.

On-Premises Disaster Recovery

Enterprises have traditionally handled IT disaster recovery internally. Keeping a robust on-premises DR solution in place and up to date requires a large investment of resources.

Hardware

Most on-premises DR solutions depend on the purchase of duplicate servers on site or at a secondary location to be used in the event of an outage. These servers incur both capital expenditures (CapEx) and ongoing IT operating expenses (OpEx), including power and cooling. Moreover, they typically require a hardware refresh every three to five years.

Software licenses

In order to launch recovery machines when source machines fail, on-premises DR solutions commonly require maintaining duplicate third-party software licenses and, in some cases, application- or DR-specific replication software. This can lead to high expenditures, especially for enterprises that use costly applications from vendors, such as Oracle, SAP, and Microsoft.

DR infrastructure & services

Any IT resilience solution must be able to restore entire systems to their pre-disaster state. On-premises DR solutions require the purchase of data protection software and, in certain cases, replication appliances. If the organization needs enterprise-grade RTOs and RPOs, they have to pay for duplicate compute and storage infrastructure at their DR site.

Management & monitoring

IT staff resources are necessary to continually manage and monitor the DR hardware, software, and infrastructure. There can be a lot of heavy lifting involved in this, such as converting machines from one infrastructure to another. And, in the event of a disaster, manual network configurations will need to be done, which is a time-consuming process.

Cloud-Based Disaster Recovery

In recent years, the DR industry has changed dramatically due to the advancement of cloud technology and the enormous growth of public cloud infrastructure, which allows you to pay only for the resources you use. In parallel, replication technologies have evolved to leverage cloud infrastructure in a cost-effective manner, forming a "perfect marriage" between DR and the cloud. Cloud-based solutions reduce DR total cost of ownership (TCO) by reducing both CapEx and OpEx.

Hardware

When using the cloud as your target DR infrastructure, no hardware is needed, and you pay for your fully provisioned cloud DR site only when required, such as during a disaster or drill. This means no CapEx investment or unnecessary duplicate provisioning of resources.

Software licenses

When using the cloud as your target infrastructure, and an appropriate replication tool, you can eliminate the need for duplicate software licenses for your DR site since there are no duplicate standby systems or standby licenses. The DR solution keeps servers continuously in sync in the cloud, without running any OS or application licenses. In the event of a disaster or a test, you can launch your servers within minutes, and only then will you need these third-party licenses.

DR infrastructure & services

Whereas traditional solutions require duplicate compute and storage infrastructure provisioned in the DR site, the elasticity of the cloud allows you to replicate your workloads into a low-cost storage area, which means you do not need to pay for expensive compute during regular DR operation. During a disaster or drill, you can launch fully provisioned workloads, and only then do you need to pay for more comprehensive resources. With the cloud, you get the resilience of a highly available system with near-zero RPOs and RTOs at the cost of a cold standby solution.

Management & monitoring

Cloud-based DR solutions provide much better automation than traditional solutions, which means fewer IT resources are required to launch or maintain the service. Automated machine conversion ensures that the heavy lifting typically involved in converting machines from one infrastructure to another is rapid and simplified. As a result, machines can boot natively in the cloud, even if they originated from a dissimilar infrastructure. Moreover, a DR solution that offers automated orchestration of the application stacks, which can be performed in advance during the implementation stage, can eliminate the need for time-consuming, manual network configurations during a disaster.

Transdev Moves Recovery Site to AWS and Reduces DR Costs by 73%

With billions of passengers around the world relying on Transdev to get to schools, jobs, and vacation destinations, Transdev cannot take risks with their IT availability. This is why, after years of using a traditional on-premises DR strategy, Transdev decided to leverage AWS for DR. Transdev now uses CloudEndure Disaster Recovery, offered by AWS, to protect their business-critical workloads, including Microsoft Active Directory, SQL Servers, MSCS clustered SAP workloads, and niche applications. As a result, Transdev has reduced disaster recovery costs by 73%.

> “We estimate that building our disaster recovery in the cloud was 73% cheaper than if we built the same solution on-prem.”
>
> — Matthieu Traverse
> Lead IT Architect,
> Transdev

transdev
MOBILITY INSPIRED BY YOU

Additional Benefits of Cloud-Based Disaster Recovery

While cloud-based DR is clearly a less expensive approach than traditional on-premises DR, you may be wondering whether this option is as effective and robust. The answer is yes. Not only does cloud-based DR technology provide top-of-the-line DR, it also provides capabilities not available with traditional DR strategies.

Easy testing

Quickly spin up machines for your periodic DR drills without disrupting your source environment.

Self-service DR

Configure your cloud environment, replicate your servers, and perform DR drills whenever you want. Deployment is easy, and access to cloud resources is instantaneous.

Flexibility between different infrastructure

Protect physical, virtual, or cloud-based source machines by replicating them into a DR site in the cloud.

Geographic redundancy

Choose a region in your preferred cloud that is located in a different geographic region than your source environment in order to achieve geographic redundancy.

Cloud-Based Disaster Recovery vs. On-Premises Disaster Recovery

	Cloud-Based DR	On-Premises DR
Enterprise-grade recovery objectives	✓	✓
Total cost of ownership (TCO)	Low	High
Automated deployment & maintenance	✓	✗
Easy, non-disruptive testing	✓	✗
Easy scalability	✓	✗
Self-service DR	✓	✗
Flexibility between infrastructure	✓	✗

What to Consider Before Moving Your DR to the Cloud

For some enterprises, moving DR to the cloud may seem like a radical move. However, in recent years, more and more enterprises, government entities, hospitals, and flagship academic institutions have done so. AWS and other clouds offer enterprise-grade security, compliance, and data integrity. As such, many organizations have declared cloud-first initiatives to outsource infrastructure to the cloud wherever possible, with DR being one of the first candidates.

The technology you choose for your cloud-based DR solution can vary greatly from one vendor to another. Some DR solutions cannot guarantee consistency or support all of your applications, which would impact your implementation success rate. Other technologies may impact your server performance or deliver inadequate RPOs or RTOs. The right technology will enable you to achieve the enterprise-grade resilience and performance of on-premises solutions with the dramatic cost reduction of leveraging the cloud for DR.

We address some of the concerns you may have specific to cloud-based DR into AWS, as well as questions to consider when evaluating the right DR solution for your enterprise.

What RPO can I achieve when using AWS to protect my workloads?

When using DR technologies that provide continuous data replication, you should expect near-zero RPO (normally seconds), depending on the latency and network quality between your source servers and target AWS Region.

What RTO can I achieve when using AWS as a target infrastructure?

Two key capabilities enable quick recovery into AWS:

1) Automated machine conversion of your source servers into AWS instances
2) Automated large-scale DR orchestration

Cloud-based DR technologies that include these two capabilities deliver recovery times of minutes and can launch all of your target machines in parallel.

Can AWS support my physical and virtual machines? What about legacy applications?

A cloud-based DR solution that performs replication at the OS level (rather than at the hypervisor or SAN level) can support recovery of any type of infrastructure into AWS, including physical, virtual, and cloud-based servers. When the replication is conducted at the block level, any file system or application (including legacy applications) can be transparently supported. Companies commonly use cloud-based DR for mission-critical databases and applications from vendors, such as Oracle, SAP, and Microsoft.

Is it possible to use AWS for disaster recovery without moving my primary workloads to AWS?

Absolutely. When you use AWS as your DR target infrastructure, you simply have a dormant copy of your workloads as AWS instances, which can then be fully launched whenever you choose to do so. You can continue to use any infrastructure you choose for your production environment, as long as your DR solution performs block-level replication.

Isn't putting my DR in the cloud a security risk?

As long as your DR solution uses proper data-at-rest and data-in-transit encryption, your data is secure. If desired, request that your DR solution allow you to be in control of the data path for the replication traffic over your private networks. Ask AWS any specific questions you might have about meeting the regulatory compliance requirements applicable to your business.

> "The cloud on its weakest day is more secure than a client-server solution."
>
> — Sean Roche
> Digital Innovation Directorate,
> Central Intelligence Agency (CIA)

Won't setting up DR in AWS disrupt my source system?

This entirely depends on the DR solution. Some cloud-based DR solutions require rebooting your system or taking frequent snapshots and may impact system performance or require local storage at the expense of your primary applications. Others are designed to be non-intrusive.

How do I conduct DR drills with a cloud-based DR solution?

DR drills are much easier when using the cloud as a target. If you use an on-premises data center as a target, you need to ensure that the resources needed for the drill are provisioned and paid for in advance. In some cases, your source applications will be disrupted to avoid network conflicts. However, when using AWS as a DR target, you can simply request the resources when needed and only pay for them upon use. Furthermore, you can spin up your target AWS machines in complete isolation, thereby performing DR drills without any impact or conflict with your source applications.

If I run my DR servers on AWS, how time consuming and costly will the failback into my primary infrastructure be once the disaster is over?

With some DR solutions, this can be a cumbersome manual process of setting up your source servers and applications from scratch, moving the data, and then keeping it in sync until the point of failback. Other DR solutions allow you to simply reverse the replication direction and keep the data in real-time sync back to your primary site within minutes once the disaster is over and you are ready for failback.

What if my servers experience a disruption that requires me to recover to a previous point in time? Is point-in-time recovery possible?

Yes. With the appropriate cloud-based DR solution, you can recover back to previous consistent points in time. So whether it is a virus, hacker, or ransomware attack that compromises your data or corrupts your database, you can recover to a time prior to the disruption.

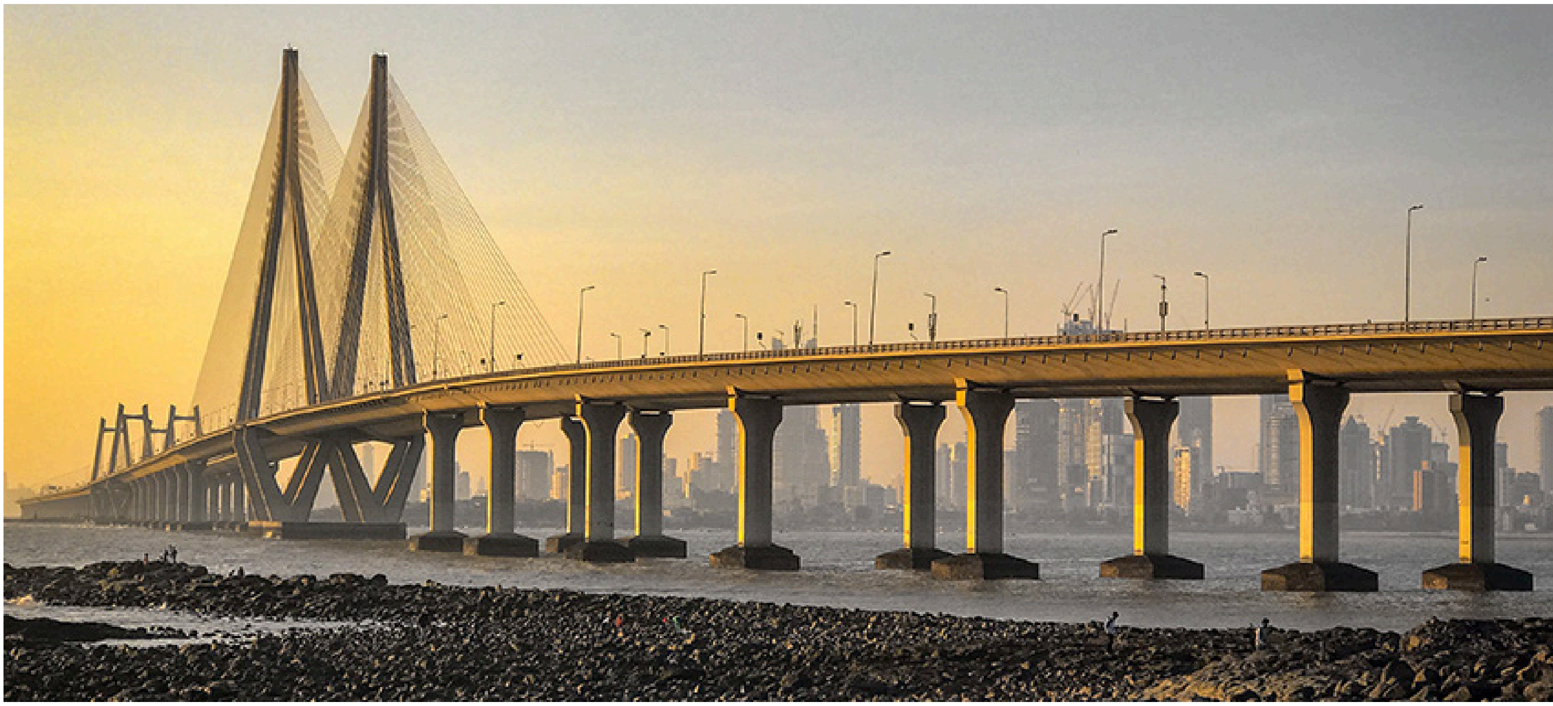

CloudEndure
An AWS Company

aws

About CloudEndure Disaster Recovery

CloudEndure Disaster Recovery minimizes downtime and data loss by enabling fast, reliable recovery of physical, virtual, and cloud-based servers into AWS in the event of IT disruption. Meet stringent recovery objectives and reduce your disaster recovery TCO with a single tool for your entire environment.

CloudEndure Disaster Recovery continuously replicates your workloads into a low-cost staging area in AWS, which reduces compute costs by 95% and eliminates the need to pay for duplicate OS and third-party application licenses. You only pay for fully provisioned workloads during a disaster or test. With CloudEndure Disaster Recovery, you can recover your environment in its most up-to-date state or a previous point in time for cases of data corruption, accidental system changes, or cyberattacks.

aws.amazon.com/cloudendure-disaster-recovery | cloudendure-info@amazon.com