Top 4 Ways to Reduce Disaster Recovery Costs Using the Cloud

aws

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 need to recover rapidly from IT disruptions, application failures, or malicious attacks, in order to maintain business continuity, 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 provide 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 allows data retrieval, but can come with 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.

By moving DR to the cloud, businesses can now attain top-of-the-line IT resilience at a fraction of the cost.

Here we present four ways that using the cloud can reduce your DR costs, compared to an on-premises DR strategy. We will also touch upon the benefits and risks of each strategy to help you decide if leveraging the cloud for DR, and specifically Amazon Web Services (AWS), is the right approach for your business.

By moving DR to the cloud, businesses can now attain top-of-the-line IT resilience at a fraction of the cost.

Cloud Disaster Recovery Site

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, particularly in the following four ways:

1
Hardware

When using the cloud as your 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 significant capital expenditures (CapEx) investment or unnecessary duplicate provisioning of resources.

2
DR infrastructure & services

Whereas traditional on-premises DR solutions require duplicate compute and storage infrastructure provisioned in the DR site, an appropriate cloud DR solution allows you to replicate your applications using low-cost cloud resources, which means you do not need to pay for expensive compute during regular DR operation. During a disaster or drill, you can launch your fully provisioned DR site, 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 minimal RPOs and RTOs, at the cost of a cold standby solution.

3
Software licenses

When using the cloud for DR, 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 OS or application licenses. In the event of a disaster or a DR drill, you can launch your servers within minutes, and only then will you need these third-party licenses.

4
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 server conversion minimizes the manual processes typically involved in converting servers from one infrastructure to another, which simplifies and speeds up recovery. As a result, servers 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.

On-Premises Disaster Recovery Site

Compared to using the cloud as a DR site, maintaining a robust, traditional on-premises DR site can require a large investment of resources in the same four areas:

1
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 CapEx and ongoing IT operating expenses (OpEx), including power and cooling. Moreover, they typically require a hardware refresh every three to five years.

2
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.

3
Software licenses

In order to launch recovery servers when source servers 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.

4
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 servers 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.

Olli Salumeria Moves Recovery Site to AWS and Reduces DR Costs by 80%

Maintaining the availability of SAP ERP (enterprise resource planning) applications is critical for Olli Salumeria, in order to meet the needs of its global customers. This is why Olli Salumeria decided to replace its legacy physical DR data center, and instead leverage AWS as a DR site for production applications. Olli Salumeria implemented AWS Elastic Disaster Recovery (AWS DRS) in less than six weeks. Using AWS DRS, Olli Salumeria can now meet required RPOs and RTOs, and has reduced DR costs by 80%.

"We established a secure and scalable platform that helped us solve technical challenges, address new business requirements, and keep our data secure."

Gregg Gilliam, Director of Information Technology and Supply Chain at Olli Salumeria

Additional Benefits of Cloud-Based Disaster Recovery

In addition to reducing DR costs, cloud-based DR technology also provides capabilities that are not available with traditional DR strategies.

Easy testing

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

Flexibility between different infrastructure

Replicate on-premises or cloud-based servers to a DR site in the cloud.

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.

Geographic redundancy

Choose a cloud region 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	✓	✗
Non-disruptive DR drills	✓	✗
Easy scalability	✓	✗
Self-service DR	✓	✗
Flexibility between infrastructure	✓	✗

What to Consider Before Moving Your DR Site 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 offers a deep set of security, compliance, and data integrity tools. 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. With the right DR technology, you can 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 organization.

What RPO can I achieve when using AWS as my DR infrastructure?
When using DR technologies that provide continuous data replication, you should expect RPOs of seconds, depending on the latency and network quality between your source servers and your AWS Region.

What RTO can I achieve when using AWS as my DR infrastructure?
Two key capabilities help you recover quickly on AWS:

1. Automated conversion of your source servers to Amazon Elastic Compute Cloud (Amazon EC2) instances
2. Automated large-scale DR orchestration

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

Can AWS support my physical and virtual servers? 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 environment to AWS?
Yes. When you use AWS as your DR infrastructure, you simply have a dormant copy of your servers on AWS, which can then be 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.

Is putting my DR in the cloud a security risk?
The DR solution you choose should use proper data-at-rest and data-in-transit encryption to help keep your data 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.

Will 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.

"The cloud on its weakest day is more secure than a client-server solution."

Sean Roche, Digital Innovation Directorate, Central Intelligence Agency (CIA)

How do I conduct DR drills with a cloud-based DR solution?
DR drills can be much easier when using AWS as a DR site. If you use an on-premises data center as a recovery site, you need to verify 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 recovery site, you can simply pay as you go for the AWS resources you use. Furthermore, you can launch your source servers on AWS in complete isolation, thereby performing DR drills without impacting or conflicting with your source applications.

 |

How time-consuming and costly will failback to my primary environment be, after 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 replicate your data from your recovery environment on AWS to your primary environment, so that your data is in sync and you can fail back whenever you're ready.

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.

About AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery (AWS DRS) minimizes downtime and data loss with fast, reliable recovery of on-premises and cloud-based applications. It uses cost-effective AWS resources to maintain an up-to-date copy of your source servers on AWS, thereby removing idle disaster recovery site resources.

During normal operation, use AWS Elastic Disaster Recovery to maintain disaster readiness by performing non-disruptive recovery and failback drills. In the event of an IT disruption, recover your applications on AWS within minutes, at their most up-to-date state or from a previous point in time. Point-in-time recovery is useful for recovery from data corruption events such as ransomware. After the issue is resolved in your primary environment, you can use AWS Elastic Disaster Recovery to fail back whenever you're ready.