AWS Whitepaper

Disaster Recovery of On-Premises Applications to AWS

Disaster Recovery of On-Premises Applications to AWS: AWS Whitepaper

Table of Contents

.... v

Abstract and introduction 1

   Abstract 1

   Are you Well-Architected? 1

   Introduction 1

      What is disaster recovery? 1

      What is a disaster? 2

      Why disaster recovery? 2

How does disaster recovery help business continuity? 4

Recovery objectives 5

Solutions and methodologies 6

   Various resilience solutions compared 6

   Disaster recovery compared to backup 7

   Disaster recovery compared to high availability 9

   Which applications require disaster recovery? 10

Using AWS for disaster recovery of on-premises applications 11

   Benefits of using AWS for disaster recovery 11

   AWS services related to disaster recovery 12

      AWS Elastic Disaster Recovery 12

      AWS DataSync 12

      Amazon Application Recovery Controller (ARC) 12

   Shared responsibility 13

      AWS responsibility: Resiliency of the Cloud 13

      Customer responsibility: Resiliency in the Cloud and Outside 13

   Business impact analysis and risk assessment 14

      How to match a disaster recovery solution to an application 16

   Disaster recovery implementation 18

      Grouping applications into waves 18

      Configuration of staging area 19

      Using multiple AWS accounts 19

      Networking considerations 19

      Mapping application structure 20

      General process for each wave 20

      First disaster recovery drill 21

This whitepaper is for historical reference only. Some content might be outdated and some links might not be available.

Disaster Recovery of On-Premises Applications to AWS

Publication date: **January 19, 2022** (*Document history*)

Abstract

This whitepaper outlines the best practices for planning, implementing, and maintaining disaster recovery for on-premises applications using AWS. It lays out the differences between disaster recovery and other resilience strategies, and describes the steps of building a disaster recovery plan. It also offers different approaches for mitigating risks and meeting recovery time objectives (RTO), and meeting recovery point objectives (RPO) by using AWS as the disaster recovery site. This whitepaper covers how to use AWS as a disaster recovery site for on-premises applications.

Refer to Disaster Recovery of Workloads on AWS: Recovery in the Cloud for information about disaster recovery for AWS-hosted workloads.

Are you Well-Architected?

The AWS Well-Architected Framework helps you understand the pros and cons of the decisions you make when building systems in the cloud. The six pillars of the Framework allow you to learn architectural best practices for designing and operating reliable, secure, efficient, cost-effective, and sustainable systems. Using the AWS Well-Architected Tool, available at no charge in the AWS Management Console, you can review your workloads against these best practices by answering a set of questions for each pillar.

For more expert guidance and best practices for your cloud architecture—reference architecture deployments, diagrams, and whitepapers—refer to the AWS Architecture Center.

Introduction

What is disaster recovery?

Disaster recovery is the process of preparing for and recovering from a disruptive event.

What is a disaster?

In the context of a company's IT environment, a disaster is an event that partially or completely disrupts the operations of one or more applications. A disaster normally requires human intervention to *fail over* to secondary copies of applications in order to maintain their functionality.

The four main categories of a disaster:

- **Human errors** – Unintentional actions leading to a security breach such as inadvertent misconfiguration of the software or a database

- **Malicious attacks** – Unauthorized actions that affect a victim's system such as a denial-of-service (DoS) or ransomware attack

- **Natural disasters** – Environmental factors that cause a system failure such as earthquakes or floods

- **Technical failures** – A malfunction of software, hardware, or a facility such as a power failure or a network connectivity failure

There are several factors to consider when planning your response to a specific disaster:

- **Expected duration of the disaster** – How soon will the application recover and how likely is the disaster to resolve on its own?

- **Size of impact (also known as blast radius)** – Which applications are affected and to what extent is their functionality impaired?

- **Geographic impact** – May be regional, national, continental, or global.

- **Tolerance of downtime** – How significant is the impact of the application not functioning?

Why disaster recovery?

A properly planned and implemented disaster recovery solution helps mitigate the following issues that can be caused by a disaster:

- **Direct and indirect financial loss** – The impact of direct financial loss is mostly relevant for applications that are critical for any revenue-generating processes. For example, external-facing IT systems that are provided to customers for a fee or internal IT systems that process data relevant for revenue generation. Indirect financial loss includes, for example, customers

switching to a competing product and the cost of work needed to resume normal operation after the disaster is over.

- **Reputational damage** – In addition to financial loss, downtime caused by unexpected incidents can significantly harm a company's reputation. A short recovery period aided by a disaster recovery solution can help avoid irreversible damage to the corporate image.

- **Failure to abide by compliance standards** – Multiple compliance standards, including System and Organization Controls (SOC), the Payment Card Industry (PCI) Data Security Standard, and the Health Insurance Portability and Accountability Act (HIPAA), require a disaster recovery plan. Some standards even add very specific requirements, such as minimal physical distance between the source site and the disaster recovery site.

How does disaster recovery help business continuity?

Disaster recovery is a component of the overall business continuity strategy of an organization. Business continuity is the ability of the organization and all the supporting applications to run critical business functions at all times, including during emergency events.

To achieve business continuity, you must implement various types of resilience mechanisms. *Resilience* is the ability of an application to recover from an outage, either automatically or with human intervention.

Disaster recovery (sometimes called *business continuity/disaster recovery*, *BC/DR*, or *DR*) is an important part of your resilience strategy and determines how you respond when a disaster strikes. This response varies between applications and should be based on your organization's business objectives for each application. These objectives should specify (among other things) the strategy for minimizing loss of data and reducing downtime when your applications are not available for use. This approach helps your organization maintain operations as part of business continuity planning (BCP).

Recovery objectives

As part of disaster recovery planning, you need to define an RTO and RPO for each application based on impact analysis and risk assessment.

RTO is the maximum acceptable delay between the interruption of an application and the restoration of its service. This objective determines what is considered an acceptable time window for an application to be unavailable.

RPO is the maximum acceptable gap between the data in the disaster recovery site and the latest data stored in the application when the disaster strikes. This objective determines what is considered the maximum amount of time acceptable for interruption/loss of data that can be caused by a disaster.
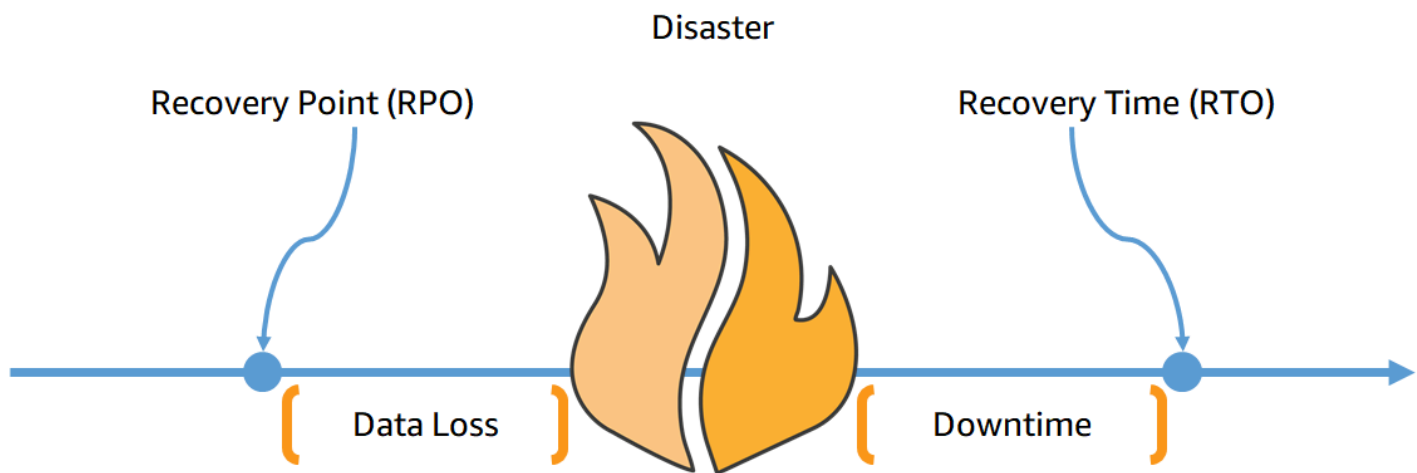
Recovery objectives

RTO and RPO for each application depend on many factors (such as service level agreements (SLA) and external compliance requirements), but there are some common standards. Common figures for mission-critical applications (tier-1 applications) include an RTO of 15 minutes and a near-zero RPO. For important applications that are not mission critical (tier-2 applications), the RTO is typically four hours and the RPO is two hours. For all other applications (tier-3 applications), a typical RTO is 8 to 24 hours and RPO is four hours.

Solutions and methodologies

The following section provides an overview of common solutions and methodologies that are advertised as disaster recovery, and explains the difference between these solutions and actual disaster recovery.

Various resilience solutions compared

- **Backup** – Backup protects against data loss by storing historical data so that if any data is lost, it can be recovered from the backup. Backup solutions can store historical data locally, in a remote location, or in both. The advantage for local backups is recovery speed, and for remote backups, the advantage is more resiliency. Backup solutions often have relatively low total cost of ownership (TCO), as the only infrastructure needed is storage, and the performance requirements for that storage are low (for example, some companies still use tape-based backup because of the low cost of tapes).

- **Archiving** – A subcategory of backup solutions is archiving. Archives provide unchanged historical copies of data to meet legal and compliance requirements. Archives are normally kept for a longer term than standard backups. Unlike backup, which may provide quicker file restoration (normally measured in hours or days), archives are not utilized by routine business operations and can be stored in low-cost, off-site locations.

- **High availability (HA)** – High availability enables an application to continue operating uninterrupted if a component of that application malfunctions. Detecting the malfunction and ensuring that the application continues to work as normal is almost always an automated process. Ideally, a user of the application would not experience anything unusual in case of such a failure. A typical example is a multi-node database.

  Most modern multi-node databases continue operation uninterrupted if a single component fails. High availability is normally introduced as part of the design and implementation of the system, as it is much harder to add to an existing application that was not designed with high availability in mind. High availability solutions ensure minimal impact on users (ideally, no impact) in case of issues. However, they are only meant to deal with a small localized event (for example, failure of a single server or subnet). A high availability solution will not be able to handle a wider disaster, such as the failure of an entire data center or a corrupted software update.

- **Disaster recovery** – Disaster recovery helps ensure business continuity for applications in case of an issue that prevents the application from recovering automatically, or that requires a

significant amount of time until recovery is achieved. Disaster recovery includes the ability to use a secondary application in a secondary location that will serve the application's users until the original instance is fixed or recovered.

Switching users to the secondary location is not an automatic process, but is instead performed on the basis of an explicit decision by an authorized person or group of people in the organization, because there are costs associated with it. For example, there is some downtime while the failover is commencing, and there is the cost of the labor for people participating in the switch. These implications need to be weighed against the chances of the source site returning to normal operation in a timely manner.

Secondary location solutions usually have a higher TCO than backup, because the secondary site needs to be maintained at all times (during normal operation as well) and needs to be advanced enough to support the functionality of the application in case of a disaster.

These three resilience solutions are complementary of each other. Business requirements may dictate that workloads should apply a combination of these solutions, depending on the business resilience requirements of each application.

Disaster recovery compared to backup

There's an important distinction between backup and disaster recovery. Backup is the process of making an extra copy (or multiple copies) of data. You back up data to be able to restore it in case it is lost or corrupted. You might need to restore backup data if you encounter an accidental deletion, database corruption, or problem with a software upgrade. It is important to have a backup solution in place. Backup protects your data in case of theft of equipment storing data, employee accidents (deletion of an important file), a technical issue (crashed hard drive), or malicious tampering (ransomware). With this protection, you can access a copy of your data and restore it easily.

Disaster recovery, on the other hand, refers to the plan and processes for quickly reestablishing access to applications, data, and IT resources after an outage. This plan might involve switching over to a redundant set of servers and storage systems until your source data center is functional again. For example, a disaster can lead to a disruption of your entire network, resulting in your employees being unable to work for the entire day (or even longer). However, a proper disaster recovery solution would allow your employees to continue to work using the mirrored system, while your IT team fixes the problem in the original network.

Some organizations mistake backup for disaster recovery. But as they may discover after a serious outage, simply having copies of data doesn't mean you can keep your business running. To ensure business continuity, you need a robust, tested disaster recovery solution that enables maintaining normal operation until the disaster is resolved.

In terms of similarities, both backup and disaster recovery solutions maintain copies of historical data that may have changed in the source storage (often referred to as *snapshots* or *point-in-time copies*). In the case of backup solutions, this is a core part of the solution's value: to be able to restore a previous version of data in case it was incorrectly modified or corrupted. In the case of disaster recovery solutions, this is done to enable successful recovery if the latest state of the data prevents normal operation. Database corruptions, ransomware data encryption, and incorrect software configuration all fall under this category and would require the disaster recovery site to be based on a previous version of the data.

However, when backup and disaster recovery are compared, there are multiple distinct differences that exist between the two:

- **Purpose** — Backups work best when you need to gain access to a lost or damaged file or object, such as an email, PowerPoint presentation, or database. Backups are also used for long-term data archival, or for purposes such as data retention. However, if you want your business to quickly restore its functions after an unforeseen event, you should opt for disaster recovery. With both the disaster recovery site and solution in place, you can perform a failover to transfer applications to the disaster recovery site, and your business can continue to function as normal even if the production site is unavailable. On the other hand, restoring a single piece of data (such as a file) is much easier to do using a backup of that data, rather than recovering an entire server where that data was stored.

- **RTO and RPO** — Setting RTO and RPO is crucial for any business. Because restoring data from backups often does not help with business continuity, the concept of RTOs and RPOs is not applicable. Disaster recovery, on the other hand, implies replicating your critical applications with the aim of quickly performing failover if necessary to assure the business continuity of the affected applications.

- **Resource allocation** — Backups are usually stored in a compressed state and do not need to be restored quickly. Therefore, backups normally use low-cost and low-performance storage (frequently off site). Disaster recovery, on the other hand, requires a separate site with operational IT infrastructure that should always be ready for a possible failover at any time.

In recent years, the term *disaster recovery solution* has become very popular, with different meanings in different cases. Therefore, it's important to analyze each product to make sure it fulfills the business continuity needs of the organization, including RPO, RTO, and the ability to quickly continue running the application from the disaster recovery site in case the source site loses functionality.

Disaster recovery compared to high availability

High availability (HA) and disaster recovery rely on some of the same best practices, such as monitoring for failures, deploying to multiple locations, and failing over. However, high availability focuses on a single component failure, whereas disaster recovery focuses on continuity in case of a wider failure of the entire application or significant parts of the application.

Disaster recovery has different objectives from high availability. Your disaster recovery strategy requires different approaches than those for high availability, focusing on deploying discrete systems (usually to multiple locations to minimize the impact of a local issue), so that you can fail over the entire application if necessary.

For example, an application that runs on a single virtual machine (VM) in a data center is not highly available. If a local flooding issue affects that data center, this scenario requires failover to another location to meet recovery objectives. Compare this scenario to a highly available application that is deployed across multiple active Availability Zones in the same AWS Region and all Availability Zones are serving production traffic. In this case, even in the localized event of one Availability Zone failing, the high availability strategy is accomplished by automatically routing all traffic to the remaining functional Availability Zones.

How you approach data resilience is also different between high availability and disaster recovery. Consider a storage solution that synchronously replicates to a nearby storage appliance to achieve the high availability of persistent data. If a file or files are mistakenly deleted or corrupted, those destructive changes will be replicated to the secondary storage device. In this scenario, despite the high availability of the storage itself, the ability to recover data in the case of data deletion or corruption is not present. When using a disaster recovery solution in the same scenario, normally a point-in-time-recovery capability is included that can be used.

Another difference between high availability and disaster recovery is how a failover is initiated. In high availability solutions, an event is initiated automatically when needed for high availability (normally within seconds), which results in little to no impact on the end user. In disaster recovery solutions, failing over often incurs additional financial or non-financial impact (for example, the

need to fail back all the new data after the disaster is over or the need to provision more resources in the disaster recovery site). Therefore, human intervention is required to initiate a failover event. Also, failing over is normally not instantaneous, and the application remains down until the failover is complete. A well-designed disaster recovery plan should define who is authorized to initiate a failover, how to reach these people, and what they need to consider when making the decision to fail over applications.

Lastly, in most cases, high availability solutions need to be selected at the time an application is designed (or refactored), as they are an integral part of the application. Disaster recovery solutions may be added to an existing application without significant re-architecture or modification work in the application itself.

Which applications require disaster recovery?

A malfunction of almost any application has a negative impact on the organization. No matter the size or role of the application, any malfunction on a key application or even a non-production application can have a negative impact; the more critical the application, the greater the impact. Therefore, all applications can benefit from a disaster recovery solution that can help quickly and easily mitigate any malfunction.

To determine whether to implement a disaster recovery solution, you need to consider the return on investment (ROI). On the one hand, each disaster recovery solution has direct and indirect costs such as software licenses and hardware, infrastructure, maintenance, and drills. On the other hand, every time a disaster strikes, it incurs costs of its own. To determine your maximum TCO for a disaster recovery solution for each of your applications, you'll need to perform a disaster risk analysis: what is the probability of a disaster happening and what are the direct and indirect financial consequences of the disaster?

Using AWS for disaster recovery of on-premises applications

The following section outlines the benefits of using AWS for disaster recovery, explores various AWS disaster recovery solutions, and explains how to assess and implement a disaster recovery solution.

Benefits of using AWS for disaster recovery

Using AWS for disaster recovery offers services that have the following benefits:

- **Elasticity** – Disaster recovery-related AWS services are normally billed per usage. This provides the flexibility to utilize AWS as a disaster recovery site by paying only for the resources used, rather than committing to a long-term contract or set number of servers.

- **TCO** – AWS offers a lower TCO than traditional, non-cloud solutions. AWS-based disaster recovery uses minimal resources in customers' AWS accounts – primarily low-cost storage for replicating data. Customers are billed only for fully provisioned servers when launched at the time of recovery or drill.

- **RTO and RPO** – Ability to achieve RTOs of minutes by launching the disaster recovery site on demand and RPOs of seconds using continuous data replication.

- **Source infrastructure support** – Supports most applications running on x86 architecture (including physical and virtual).

- **Hypervisor support** – Supports any hypervisor (including physical servers, when there is no hypervisor at all).

- **Wide OS support** – AWS supports a [large variety](#) of Linux distributions and Windows versions.

- **Environment isolation** – Ability to create the disaster recovery site in an isolated environment so that drills do not impact the source site.

- **Application support** – Atomicity, consistency, isolation, durability (ACID)-compliant applications are supported, including any ACID-compliant database, such as Microsoft SQL Server, Oracle Database, and SAP HANA.

- **Automation** – Ability to fully automate all of the disaster recovery-related operations.

- **Ease of use** – Ability to add the disaster recovery capabilities to working applications with no need for redesign or re-architecture work.

AWS services related to disaster recovery

AWS can be used as the infrastructure running your disaster recovery site. Additionally, AWS offers a number of services that can help you replicate your source applications into the AWS infrastructure and recover them in the case of a disaster.

AWS Elastic Disaster Recovery

AWS Elastic Disaster Recovery (AWS DRS) is designed for cloud-based disaster recovery of virtual and physical servers. Elastic Disaster Recovery continuously replicates applications and databases from any supported source to AWS using block-level replication of the underlying server. The service allows you to use AWS as a disaster recovery site for on-premises applications comprising servers (physical and virtual), including databases. During normal operations, Elastic Disaster Recovery continuously replicates changes to the source data to a staging area subnet in your AWS account. The staging area design uses affordable storage and minimal compute resources to maintain ongoing replication. When you initiate a failover or drill, the staged resources are used to automatically create a full-capacity deployment in your Amazon Virtual Private Cloud (VPC), which is used as the disaster recovery site. You can launch recovery instances on AWS within minutes, using the most up-to-date server state or a previous point in time. After the disaster, you can use Elastic Disaster Recovery to fail back to your primary site.

AWS DataSync

AWS DataSync is an online data transfer service that simplifies, automates, and accelerates moving data between on-premises storage systems and AWS storage services, and also between AWS storage services. DataSync can copy data between Network File System (NFS), Server Message Block (SMB) file servers, self-managed object storage, AWS Snowcone, Amazon Simple Storage Service (Amazon S3) buckets, Amazon Elastic File System (Amazon EFS), and Amazon FSx for Windows File Server file systems.

If you have large network-attached storage (NAS) appliances, AWS recommends using AWS DataSync for replicating them to Amazon S3, Amazon S3 Glacier, Amazon EFS, or Amazon FSx for Windows File Server.

Amazon Application Recovery Controller (ARC)

Amazon Application Recovery Controller (ARC) gives you insights into whether your applications and resources are ready for recovery, and helps you manage and coordinate failover using readiness

check and routing control features. These features continually monitor your application's ability to recover from failures, and enable you to control your application recovery across multiple AWS Regions, Availability Zones, as well as on premises. These capabilities make application recoveries simpler and more reliable by eliminating the manual steps required by traditional tools and processes.

Shared responsibility

Disaster recovery is a shared responsibility between AWS and you, the customer. It is important that you understand how disaster recovery and availability, as part of resiliency, operate under this shared model.

AWS responsibility: Resiliency of the Cloud

AWS is responsible for resiliency of the infrastructure that runs all of the services offered in the AWS Cloud. This infrastructure comprises the hardware, software, networking, and facilities that run AWS Cloud services. AWS uses commercially reasonable efforts to make these AWS Cloud services available, ensuring service availability meets or exceeds [AWS Service Level Agreements (SLAs)](#).

The [AWS Global Cloud Infrastructure](#) is designed to enable customers to build highly resilient workload architectures. Each AWS Region is fully isolated and consists of multiple Availability Zones, which are physically isolated partitions of infrastructure. Availability Zones isolate faults that could impact workload resiliency, preventing them from impacting other zones in the AWS Region. At the same time, all zones in an AWS Region are interconnected with high-bandwidth, low-latency networking, over fully redundant, dedicated metro fiber providing high-throughput, low-latency networking between zones. All traffic between zones is encrypted. The network performance is sufficient to accomplish synchronous replication between zones. Availability Zones simplify the process of partitioning applications for high availability.

Customer responsibility: Resiliency in the Cloud and Outside

Your responsibility is determined by the AWS Cloud services that you select. This determines the amount of configuration work you must perform as part of your resiliency responsibilities. You are responsible for managing resiliency of your data and workloads, whether on AWS or outside of it, including disaster recovery, high availability, backup, versioning, and replication strategies.

	SECURE DATA BACKUP	
CUSTOMER RESPONSIBILITY FOR RESILIENCY 'IN' THE CLOUD	WORKLOAD ARCHITECTURE	
	CHANGE MANAGEMENT	FAILURE MANAGEMENT
	NETWORKING, QUOTAS & CONSTRAINTS	

	HARDWARE & SERVICES			
AWS RESPONSIBILITY FOR RESILIENCY 'OF' THE CLOUD	COMPUTE	STORAGE	DATABASE	NETWORKING
	AWS GLOBAL INFRASTRUCTURE			
	REGIONS	AVAILABILITY ZONES		EDGE LOCATIONS
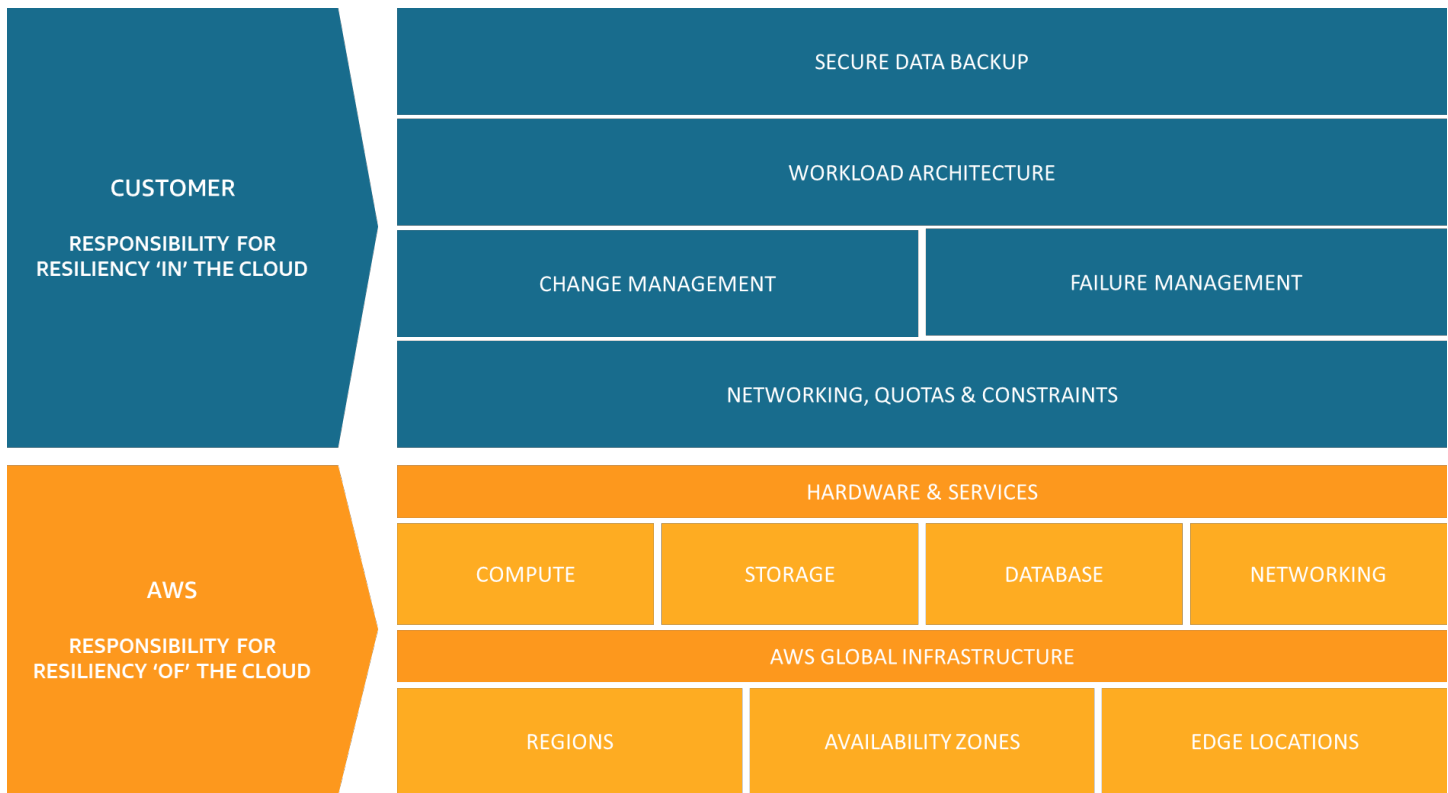
AWS shared responsibility model

Business impact analysis and risk assessment

One of the first steps to perform when planning the implementation of a disaster recovery solution is a business impact analysis for all relevant applications. A business impact analysis should quantify the business impact of a disruption to each application. It should identify the impact your internal and external customers will face from not being able to use your applications and the effect that will have on your business with regards to cost, reputation, and compliance. The analysis should help you determine how quickly the application needs to be made available (RTO) and how much data loss can be tolerated (RPO). However, recovery objectives should not be defined without also considering the likelihood of disruption and the cost of recovery when calculating the business value of providing disaster recovery for an application.

The business impact of a disaster may not be constant. For example, the impact might be dependent on the timing of the disaster — disruption to your payroll system is likely to have a very high impact to the business just before employees are supposed to be paid, but it may have a low impact just after employees have been paid.

With all of this information, you can document the threat, risk, and impact of different disaster scenarios and the associated recovery options. This information should be used to choose the best disaster recovery strategy and tools for each application and to match the risk and impact (financial, among other types) of a disaster for each application.

The following table is a sample risk analysis for disaster recovery planning for determining maximum TCO.

Table 1 - Sample risk analysis for disaster recovery planning

Risk analysis for application <APPLICATION NAME>

Disaster type	Likelihood (per year)	Consequences	Risk (likelihood x consequences)
Weather	10% (based on past statistics)	$100,000	$10,000
Power outage	5% (based on past statistics)	$150,000 (includes equipment that would need to be replaced)	$7,500
			$17,500

After you have determined your maximum TCO, identify whether there is a disaster recovery solution that has a lower TCO than the cost estimated in the risk analysis (taking the probability of a disaster into account). If such a solution exists, then it makes sense financially to use that solution for that application.

In addition to the financial calculations, the RTO and RPO for each application need to be considered during this process. The more critical the application is, the more aggressive its RTO and RPO requirements, resulting in a higher TCO of an overall disaster recovery solution. Therefore, solutions that don't offer the needed RTO and RPO should not be considered, even if they make sense based on the raw financial calculation.

How to match a disaster recovery solution to an application

There are a variety of disaster recovery solutions on the market. Currently available disaster recovery solutions have the following main differences:

- **Source application support limitations** – Different disaster recovery products and services support a different subset of operating systems, central processing unit (CPU) architectures, applications, and hypervisors. For applications that are not supported by a disaster recovery solution, the solution will either not be able to correctly (if at all) copy the data to the disaster recovery site or run a working copy of that application in the disaster recovery site when needed (the latter sometimes can only be discovered during a disaster recovery drill).

- **Recovery infrastructure support limitations** – Different disaster recovery solutions have different requirements for the disaster recovery site infrastructure. For example, not all recovery infrastructures support all of the operating systems of the source or can run with the required performance.

- **RPO** – Choose a solution that provides the RPO required for the application.

- **RTO** – Choose a solution that offers your required RTO.

- **TCO** – Usually the more features and capabilities a disaster recovery solution has, the higher its TCO will be. For example, more aggressive RTOs and RPOs would increase the TCO.

- **Compliance** – Applications that are covered by compliance certifications normally require their disaster recovery solutions to be covered by the requirements of the certifications.

- **Ease of operation** – Choose a solution that does not require significant effort or uncommon skills to operate and manage during normal operation (outside of drills and disasters).

- **Ease of disaster recovery drills** – Choose a solution that reduces the cost and simplifies the process for disaster recovery drills. This will encourage you to conduct the drills more frequently and improve your readiness for disasters.

- **Level of automation** – Disaster recovery solutions that offer a higher level of automation (for example, they have application programming interfaces (APIs) for integration with other solutions) can be used to extend their capabilities and meet more of your disaster recovery needs.

This list of technical criteria should be used when evaluating which disaster recovery solution best meets your needs. Keep in mind that different applications (or groups of applications) are likely to have different requirements. It's also important to pay attention to the small print and to confirm that the product will work as expected in real life scenarios.

Lastly, as a general best practice, narrow down your disaster recovery solutions to the ones you actually need. Operationally, each new solution will have separate overhead, an added learning curve, and additional implementation and maintenance efforts (such as monitoring the solution, upgrading versions, and applying security patches). Therefore, AWS recommends selecting a small number of disaster recovery solutions to cover all the required applications.

After you select the disaster recovery solution for each application, you need to plan the implementation project. Implementing disaster recovery for a large number of applications may be a lengthy and complex project.

To simplify and accelerate the project, include the following in your planning process:

- **Mapping** – Map the applications, the resources within the applications, and what needs to be recovered for the application to run. AWS recommends preparing a complete list of such resources, grouped by applications. These resources include servers, appliances (such as NAS appliances and firewalls), and networks. As a part of this mapping, dependencies within and between applications need to be established and documented.

  For each application, go through the risk analysis, define the disaster recovery tier, and determine the RTO and RPO.

- **Timeline** – For both the planning and the implementation stages, a realistic timeline needs to be defined based on the number of applications and their complexity and also the skillset and experience of the people involved in the project.

- **People** – The number of people that will be allocated for each stage of the project (including planning and implementation), who these people are, and the skillsets, roles, and responsibilities of each person.

- **Budget** – How much the entire project is estimated to cost. The budget should take into account the cost of the licenses of the AWS services used as well as other costs, such as data transfer from the source site to AWS.

- **Solution** – Choose solutions for each disaster recovery tier. Because this paper focuses on using AWS for disaster recovery, the prescribed suggestions are as follows:

  - For servers, use Elastic Disaster Recovery. Edge cases for which Elastic Disaster Recovery cannot be used include unsupported operating systems, non-ACID applications (such as MyISAM-backed MySQL databases), applications using shared storage that multiple nodes write to in parallel (such as Oracle RAC), systems that have distributed databases with multiple nodes that need to be in sync with each other (such as a Hadoop Cluster), and servers that the AWS Replication Agent can't be installed on (for example, third-party appliances).

- For NAS appliances, such as NetApp, use DataSync whenever possible.

- For components of source applications that are not supported by either Elastic Disaster Recovery or DataSync, an application-level solution needs to be used. For example, Oracle has application-level solutions for replicating Oracle RAC databases, NAS appliances can be replicated by periodically copying all the changed files to AWS, and Hadoop clusters may require duplicating each node with another one running in AWS. Application-level disaster recovery solutions have multiple disadvantages, including scope (only relevant for specific applications or application components), cost, and ability to replicate over long distances. Therefore, we recommend trying other disaster recovery solution options prior to opting for an application-level solution.

- **Success criteria** – For a disaster recovery implementation project, success is normally signified by a disaster recovery drill that achieves a failover according to the requirements of the disaster recovery plan.

- **Team resources** – Assign team members who are skilled in the resources you mapped for replication. For example, if you have a mix of Linux and Windows operating systems, then you need people who understand both operating systems.

- **Timelines** – Assign timelines to the implementation phase and to at least a single successful drill per application that will conclude the implementation. The disaster recovery implementation is done when all of the applications have had at least one successful drill.

Disaster recovery implementation

The following section explains how to implement an AWS disaster recovery solution, including grouping the application into waves, performing drills, and maintaining the solution.

Grouping applications into waves

Divide your applications into waves. Each wave should include applications that are designed to work together so that every wave can be tested individually without being obstructed by dependencies on resources that are allocated to future waves. By implementing disaster recovery in waves, any issues that arise will have limited impact, and the implementation project will be easier to manage.

Make sure that the first several waves are smaller in size (10-20 servers per wave) as the implementation of these waves may take more effort and be slower due to lack of familiarity with

the solutions. Once the team feels comfortable with the methodology and tools, the size of the waves can be increased according to the desired rate of progress and available resources. Large-scale disaster recovery projects usually include 100-200 servers in each wave.

Configuration of staging area

Before the implementation of the first wave, the staging area subnets should be created for all of the waves. We recommend using dedicated subnets for the staging areas. This allows setting up the needed connectivity without interfering with any other processes. This also helps prevent existing resources from competing over private IP addresses with the resources automatically provisioned by the disaster recovery solutions.

Using multiple AWS accounts

You can use multiple AWS accounts for staging areas and recovery sites (used for failover and drills). AWS accounts have API throttling, so replicating more than 300 servers to a staging area subnet in a single AWS account is not recommended. In these cases, multiple AWS accounts should be used. You can fail over servers that are using different AWS accounts for staging areas into the same recovery AWS account.

Due to AWS API throttling and based on AWS best practices, we recommend using dedicated AWS accounts for the staging areas as any additional activity in the staging AWS accounts (especially automated tools) may compete with incoming automated disaster recovery activities and may cause API throttling.

In general, the replication subnets (staging area) and the recovery subnets (disaster recovery site) of an application don't have to be in the same AWS account. This provides flexibility when designing the AWS account architecture for disaster recovery.

Networking considerations

While performing a disaster recovery drill, AWS recommends that you isolate the recovery subnets to prevent potential conflicts with the source location. Isolation can be done by using route rules, network access control list (network ACL) rules, or by associating restrictive security groups when launching the applications for drills.

If you are using Elastic Disaster Recovery, there are three points of interaction between the service components:

- The AWS Replication Agent needs to communicate with the Elastic Disaster Recovery and Amazon S3 endpoints in your recovery AWS Region.

- The AWS Replication Agent needs to communicate with the Elastic Disaster Recovery Replication Servers in the staging area subnet.

- The Elastic Disaster Recovery Replication Servers need to communicate with the Elastic Disaster Recovery and Amazon S3 endpoints in your recovery AWS Region.

It's important to define how the applications will be divided between different AWS accounts, whether the applications will be in the organization's main AWS account or in different AWS accounts due to organizational policies.

Mapping application structure

For each wave, the following components need to be analyzed:

- **Networks** – The networks that the components are located in need to be analyzed and documented. This includes the classless inter-domain routing (CIDR) of each network, a map of the components of the networks, the IP addresses of each component in each network, and the security group configuration that each failed over server should be using.

- **Server configuration** – Every server that participates in a wave needs its central processing unit (CPU) and random-access memory (RAM) information recorded to validate that the Amazon Elastic Compute Cloud (Amazon EC2) instances that Elastic Disaster Recovery launches during drills and recovery will have enough resources to work. Elastic Disaster Recovery selects a right-sized EC2 instance type based on your source server's operating system, CPU, and RAM. However, it is recommended that you verify the selection logic.

Using the results of the mapping described previously, you can build the disaster recovery site environment that will be used to recover your servers and other application components in case of a disaster or drill.

General process for each wave

AWS suggests sprints of one to two weeks for each wave. The recommended process for each wave is as follows:

1. **Define replication settings –** Your replication settings include configuring your staging area subnets and security groups (both based on the networking analysis described earlier), and

additional settings that determine how data will be replicated from your source servers to AWS. You can make changes to these settings at any time, for individual servers or a group of servers.

2. **Agent installation –** Install the AWS Replication Agent on all the servers and confirm that replication has started successfully.

3. **Launch settings configuration –** After replication has started successfully for all the servers in the wave, configure launch settings for each server to define how your drill and recovery instances will be launched on AWS. These settings include configuring the subnet within which instances will be launched, Amazon EC2 instance types, and license transfers.

4. **Testing –** Every time a server or a group of servers has reached Continuous Data Protection (indicated in the console as "Ready for recovery" and "Healthy" data replication status), run a sanity launch for the group to make sure that each drill instance has launched and booted successfully in AWS (the server has reached 2/2 status check in the Amazon EC2 Console).

   It's important that all the drill launches (tests conducted during implementation as well as periodic drills) are launched into isolated subnets to make sure they don't interfere with the source applications.

   Note that for every staging area subnet that you set up, the first wave may uncover configuration issues and the effort required to correct these issues should be considered when planning the wave.

5. **Clean up –** By default, Elastic Disaster Recovery removes any resources created during drills, either when requested by the user or when a new drill instance is launched.

First disaster recovery drill

After the implementation of all the waves is finished, run a formal disaster recovery drill to validate the drill section of the disaster recovery plan and to make sure nothing was missed during implementation. Finding and correcting issues during a full disaster recovery drill is more expensive and labor intensive than fixing issues during the implementation phase. Therefore, it's important to run the first comprehensive disaster recovery drill only after the implementation is finished — including test launches of all servers.

AWS recommends that you isolate the recovery subnets before performing any drill (including the first one), to avoid potential conflicts with the source environment.

Isolating launched applications is a drill methodology that enables you to continue operating the production site uninterrupted during drills, to lower risks in case the drill does not go as expected.

A successful first disaster recovery drill is the beginning of the maintenance phase.

Maintenance

A disaster recovery solution requires regular maintenance after implementation in order to ensure that it will work as expected during a disaster.

The two main parts of maintaining a disaster recovery solution are monitoring and periodic drills that should be performed as frequently as possible while taking the business goals and limitations into account.

Monitoring

There are multiple methods for monitoring your disaster recovery solution:

Built-in monitoring capabilities

Some services, such as Elastic Disaster Recovery, have several basic monitoring features. For example, after you add source servers to Elastic Disaster Recovery, you can monitor and interact with them from the Source Servers page. The Source Servers page is the default view in the Elastic Disaster Recovery Console. On the Source Servers page, you can view all of your source servers, monitor their recovery readiness and data replication state, view the last recovery result, view any pending actions, and sort your servers by a variety of categories. These built-in monitoring capabilities provide a general impression of the health of your disaster recovery solution.

Using AWS services for AWS Elastic Disaster Recovery monitoring

Ongoing monitoring for replication metrics, including receiving alerts and status updates, should be a high priority for limiting downtime and maintaining your disaster readiness.

One of the ways you can monitor Elastic Disaster Recovery is by configuring an Amazon CloudWatch dashboard to monitor the AWS resources you use for disaster recovery. Amazon CloudWatch is a monitoring and observability service that provides you with data and actionable insights to monitor your applications, respond to system-wide performance changes, and optimize resource utilization. Creating a CloudWatch dashboard can give you a single operational overview of your Elastic Disaster Recovery operations.

You can use built-in CloudWatch widgets to populate your disaster recovery dashboard, including metrics to visualize bandwidth consumption, network throughput, security, and server replication metrics. For example, you can use CloudWatch to:

- Monitor source server replication metrics such as lag duration and backlog. This can help you identify and remediate replication issues, so that your RPO remains as expected.

- Monitor Replication Server metrics such as CPU utilization, input/output (I/O) characteristics, and network throughput. These metrics can help you determine whether to optimize Replication Server EC2 instance sizes or [Amazon Elastic Block Store](#) (Amazon EBS) disk types.

- Monitor drill and recovery launch metrics, such as the amount of time since the last drill or recovery instance was launched. This helps validate that you are performing drills with the frequency defined by your disaster runbooks.
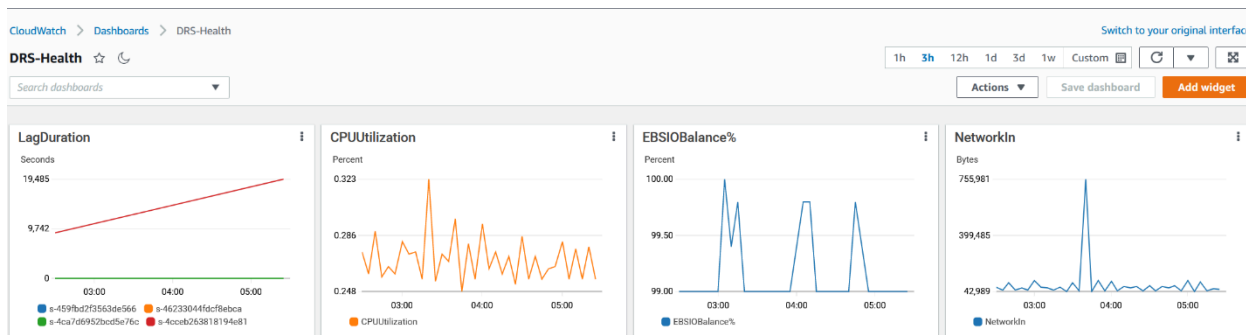
Example CloudWatch dashboard for Elastic Disaster Recovery monitoring

You can also use [Amazon EventBridge](#) to create rules for Elastic Disaster Recovery events you want to monitor, and configure notification alerts using [Amazon Simple Notification Service](#) (Amazon SNS). For example, you can configure Amazon SNS notifications for an EventBridge rule to notify you in the event that replication has stalled for any source servers. Doing so can help your teams quickly identify and address replication issues that could affect RPO and RTO if you need to recover these servers.

For a detailed walkthrough of how to create a CloudWatch dashboard for Elastic Disaster Recovery monitoring and how to use EventBridge and Amazon SNS to receive alerts, refer to [Disaster recovery monitoring of AWS Elastic Disaster Recovery](#).

Implementing monitoring using the APIs of AWS services

For the best customization and granularity of monitoring, you can use the APIs of the AWS services you are using for disaster recovery (for example, [Elastic Disaster Recovery APIs](#)). Using these APIs, most notably the DescribeSourceServers API that returns all the data about replicating servers, you can build your own monitoring solution to track any data that the Elastic Disaster Recovery service contains. Using APIs for monitoring is meant for advanced user customizations because it enables

you to choose the exact data you want to monitor and define triggers and alerts for your specific needs.

Drills

Performing periodic drills is the only way to gain confidence that a disaster recovery solution will provide business continuity within the desired RTO and RPO for each application and that the disaster run books are accurate. The more frequently drills are performed, the higher the level of confidence. However, every drill has costs associated with it, including workforce costs and the cost of launched resources on AWS. Therefore, every organization needs to decide for itself the frequency of drills for each application. Industry best practice is to conduct a disaster recovery drill no less than once per year and for the more business-critical applications, no less than once per quarter.

Performing drills

AWS disaster recovery-related services facilitate frequent drills because they provide a simple mechanism for testing the recovery of your source environment at scale, without performance impact on your source applications.

How Elastic Disaster Recovery drills work

During normal business operation, Elastic Disaster Recovery continuously replicates data on your source servers to a low-cost staging area in your AWS account. When you launch source servers for drills or recovery, Elastic Disaster Recovery automates source server conversion, so your recovered applications run natively on AWS. Drills are non-disruptive and do not impact your source servers or ongoing data replication. The instances you launch during Elastic Disaster Recovery drills will operate in the same way on AWS as the instances you launch for recovery on AWS.

When you launch drill instances for disaster recovery drills, you are launching a copy of your servers in your recovery AWS Region from a point in time that you select. You can perform secure drills in isolated subnets configured in your Elastic Disaster Recovery launch settings, which define the settings and configuration of your launched drill instances. Your launch settings and EC2 launch template enable you to isolate your drill instances by using a separate subnet or different security groups to avoid network conflicts. You can also launch drill instances in a separate AWS account to further isolate your drill and production environments. Elastic Disaster Recovery automates provisioning the resources needed to launch your instances on AWS.

You can use Elastic Disaster Recovery to run a virtually unlimited number of drills, as often as you choose. There are [no additional fees for drills](#), beyond payment for the provisioned resources generated. You can minimize costs by performing some disaster recovery drills using smaller Amazon EC2 instances, instead of fully provisioning resources at scale.

Familiarity with drill and recovery processes enables your organization to verify that you can respond quickly if you must recover applications on AWS. You can facilitate disaster recovery drills at scale by automating drill and recovery processes.

How to respond in the event of a disaster

The following section explains how to respond in the event of a disaster, including how to perform a failover and how to perform a failback and return to normal operations.

Performing a failover

An actual failover is very similar to a drill. The two main differences are that during an actual failover, your users will be redirected to the disaster recovery site, and that a failback may be needed when the disaster is over. In the event of planned or unplanned downtime, begin the failover process by using the Elastic Disaster Recovery console to launch recovery instances on AWS from the latest state or a point in time you select.

1. Launch recovery instances for a single source server or multiple source servers. This action is recorded in [AWS CloudTrail](#).
2. Perform the actual failover (directing traffic to your recovery instances) using the tool or the means you use for directing traffic. For a Domain Name System (DNS) redirect, AWS recommends using [Amazon Route 53 Application Recovery Controller](#).
3. After a successful failover, and after the downtime is over, you can prepare for failback.

Performing a failback and returning to normal operations

After performing a successful failover, verify that any data that was written to your recovery systems is replicated back to your original systems before you perform the actual failback and redirect users to your primary systems. This can be data from changes that occurred while the disaster recovery site was active and that needs to merged back into the source applications, or all the data may need to be copied back in case the data of the source site cannot be recovered after the disaster is over.

To fail back to the source servers (or to new servers, if the original source servers are no longer available after the disaster is over), follow these steps:

1. Start by replicating the data from the disaster recovery site back to the source site (when using Elastic Disaster Recovery, replicate the data from your recovery instances on AWS back to your source servers using the Failback Client).

2. Continue using the disaster recovery site while the data is being replicated. When using Elastic Disaster Recovery, you can track failback replication progress from the console.

3. Choose a failback window that minimizes the impact on users (because during failback a short amount of downtime may occur).

4. During the failback window, go to the disaster recovery site and stop the application.

5. Make sure that all the data finishes replicating to the source site.

6. Start the application in the source site.

7. Make sure the application launches correctly.

8. Once you have verified that the application is running properly, redirect users back to the application in the source site.

Conclusion

A properly planned and implemented disaster recovery solution is an important part of your resilience strategy and is crucial to mitigating the financial loss and reputational damage that a disaster can cause. This whitepaper outlined the various considerations of how to choose a disaster recovery solution and how AWS can be utilized in such scenarios.

Organizations with on-premises workloads have a variety of options when it comes to using AWS for disaster recovery. This whitepaper covered the various services (including AWS Elastic Disaster Recovery, AWS DataSync, and Amazon Route 53) that can help you replicate your source applications to AWS and recover them in the case of a disaster. This whitepaper also outlined how to best match a disaster recovery solution to an application and offered guidance for the entire disaster recovery implementation process, from mapping applications to performing drills, and finally to performing a disaster recovery failover and failback.

Appendix A: Glossary

- **ACID** – atomicity, consistency, isolation, durability (ACID) is a set of software properties (often applied to databases) needed to guarantee data validity despite errors, power failures, and other mishaps. Application-consistent recovery has no benefits over crash-consistent recovery for ACID-compliant applications and crash-consistent recovery is possible without any performance impact on the application. This makes it easier to perform recovery of ACID-compliant applications. Almost all modern software, including databases (such as Microsoft SQL Server, Oracle, and SAP HANA) and file systems are ACID-compliant.

- **Application consistency** – Application consistency is a method used to help ensure that a recovered application functions correctly. This method requires the application to be aware of the replication process and participate in it. For example, an attempt to take an application-consistent snapshot of a database will cause the database to pause all pending transactions, complete all the transactions in-flight, commit all the changes from them to the disk, take the snapshot, and then resume normal operation. Application consistency is the oldest consistency type; however, application-consistent snapshots are rarely needed today as the vast majority of modern applications are ACID-compliant (refer to Appendix A: Glossary ACID entry) and recover properly from a crash-consistent snapshot (which has significantly less performance impact when taken).

- **Application-level disaster recovery solution –** A disaster recovery solution that operates on the application level. Both the data replication and recovery processes are application aware. The advantage of such solutions is that application awareness may offer extra functionality, such as application-consistent snapshots or replication from a larger server to a smaller server. The disadvantages of such solutions are that there are not many options available for popular applications because they only function for the application they are designed for and they are usually expensive.

- **Business continuity** – The ability of an organization to continue operating correctly following a disruptive incident.

- **Cold site** – A cold site is a disaster recovery site maintained for disaster preparation that provides power, network connectivity, air conditioning, and other elements necessary to maintain data. In the case of a disaster, you need to install the hardware, software, and data before recovery can take place. Cold sites are among the most cost-effective disaster recovery solutions, but they are also the slowest and least reliable.

- **Continuous data protection** – Real-time (or near real-time) protection of data as it changes. Each change made in the source application is copied to the disaster recovery site in real time.

With continuous data protection, you can achieve minimal RPO and ensure the data on your disaster recovery site is up to date with your production site. There are two main categories of continuous data protection: synchronous and asynchronous.

Synchronous data replication first commits changes to the disaster recovery site and only then commits the data to persistent storage after it commits changes to the disaster recovery site. This method achieves true zero RPO. However, because the network latency between the two sites is added to each disk write operation, it is only applicable for situations where that latency is very low. Therefore, synchronous data replication is mostly used for high availability within the same network, rather than for long-distance disaster recovery. Because synchronous replication has a significant impact on the I/O performance of the application, this needs to be considered when designing the application and defining its performance goals.

Asynchronous data replication does not affect the performance of the source application because it commits the changes to the disaster recovery site and persistent storage in parallel. This makes it the most common method of continuous data protection for disaster recovery. However, the RPO that normally can be achieved with this method is *near zero* (less than a second) rather that absolute zero.

- **Disaster** – An event that is partially or completely disruptive to the functioning of one or more applications and cannot be resolved automatically.

- **Disaster recovery** – The process of preparing for and recovering from a disruptive event.

- **Disaster Recovery as a Service (DRaaS)** – A method of disaster recovery whereby infrastructure, personnel, equipment, drills, and other associated components are outsourced to a third-party provider.

- **Disaster recovery drill** – An implementation of the section of the disaster recovery plan dealing with response to a disaster for testing purposes rather than in case of a real disaster event. By following the exact steps in the plan and verifying that the disaster recovery site is functioning and is able to provide the required business continuity within the required RTO and RPO, you can confirm that this would also be the case if a real disaster strikes. Disaster recovery drills are performed periodically. The frequency of such a drill is based on multiple factors, such as requirements by compliance certifications and the cost of each drill for the organization.

- **Disaster recovery plan** – A plan of action (that should be as detailed as possible) to get your IT systems back online in the case of a disaster. The disaster recovery plan (DRP) should provide clear instructions and documentation for how to decide on and implement a recovery in the event of a disaster as well as how to return to normal operation after the disaster is over.

- **Disaster recovery site** – A site, data center, or cloud environment that is used to run recovered applications instead of the production environment in order to provide business continuity in case of a disaster.

- **Failback** – The process of returning to your source site from your disaster recovery site following a failover after your original applications have been restored.

- **Failover** – The process of switching from your production site to your disaster recovery site in the event of a disaster. After your production site is operational again, you implement a failback to return to normal operations.

- **Fault tolerance** – The ability of a system to remain in operation even if some of the components used to build the system fail.

- **High availability** – A system's ability to continue functioning even if a component in the system experiences a malfunction. High availability is a part of the overall resilience strategy that is complementary to disaster recovery.

- **Hot site** – A remote, fully operational copy of your source site that is ready to become active at any moment. In the event of a disaster, you can switch your operations to your hot site without additional steps. A hot site would typically provide the shortest RTO; however, its TCO is normally significantly higher than the alternatives.

- **Hypervisor-level disaster recovery solution –** A disaster recovery solution that runs on a hypervisor level and allows disaster recovery of VMs without installing agents on the VMs themselves. The disadvantage of such solutions is that they are hypervisor specific for both source and recovery site.

- **Point-in-time recovery** – A function of a disaster recovery solution that enables an organization to restore or recover data that was saved during a specific time from a specific point-in-time snapshot, which is a copy of all data saved to the disaster recovery site.

- **Server-level disaster recovery solution** – A disaster recovery solution that operates from within every server (usually by an agent running on the server) that can be used to replicate various disks and partitions on the server regardless of the applications using them. The advantage of such a solution is that it is both hardware/hypervisor and application agnostic. The disadvantage is that you need to install an agent on each server, which in turn can introduce operational overhead due to the need to maintain all the agents as operational.

- **Source site** – The original production site where your applications are running and their data is stored. A source site can be on premises or cloud based. You can have more than one source site.

- **Staging area subnet** – A part of the disaster recovery site dedicated to storing the resources used during normal operation for replication purposes. This way such resources can be

segregated from the section of the disaster recovery site where applications are launched during a failover or drill.

- **Recovery point objective (RPO)** – RPO is the maximum acceptable gap between the data in the disaster recovery site and the latest data stored in the application when the disaster strikes. This objective determines what is considered an acceptable loss of data (measured in time units) that can be caused by a disaster.

- **Recovery time objective (RTO)** – RTO is the maximum acceptable delay between the interruption of an application and the restoration of its service. This objective determines what is considered an acceptable time window for an application to be unavailable**.**

- **Server-level disaster recovery solution –** A disaster recovery solution that operates from within every server (usually by an agent running on the server) that can be used to replicate various disks and partitions on the server regardless of the applications using them. The advantage of such a solution is that it is both hardware/hypervisor and application agnostic. The disadvantage is that you need to install an agent on each server, which in turn can introduce operational overhead due to the need to maintain all the agents as operational.

- **Snapshot shipping –** A common method for data replication. It consists of taking a snapshot or a group of snapshots of the servers or disks on the source site and sending them over the network to the disaster recovery site, where they are stored, ready to be utilized to create new volumes or servers.

- **Storage-level disaster recovery solution –** One of the most expensive and oldest types of disaster recovery solutions. In these solutions, replication and recovery are performed on the storage appliance level and the data is replicated into a secondary storage appliance in the disaster recovery site. Such solutions tend to be very well integrated into the appliances and provide very good RTOs and RPOs. The disadvantages include that usually every storage vendor has only one relevant disaster recovery solution and no third-party alternatives. Another disadvantage is that they require the disaster recovery site to include a storage appliance as well, usually of the same manufacturer as the primary one, which makes this solution unsuitable for public clouds where customers don't have access to the configuration of the low-level block storage.

- **Total cost of ownership (TCO) –** The sum of all direct and indirect costs incurred as part of implementation and maintenance of a disaster recovery solution and plan. The TCO may have a fixed component (capital expenditures (CapEx)), a variable per-use component (operating expenses (OpEx)), or a combination of both.

- **Warm site/pilot light –** A warm site is a disaster recovery site where you have a small but functional copy of the source site. In the event of a disaster, the site is increased to be able to

withstand the entire load of the source site. The idea of the pilot light is an analogy that comes from the gas heater. In a gas heater, a small idle flame that's always on can quickly ignite the entire furnace to heat up a house as needed.

Contributors

Contributors to this document include:

- Leonid Feinberg, Head of Migration and Disaster Recovery Strategy, CloudEndure, AWS

- Pavel Rubin, Senior Technical Writer, CloudEndure, AWS

- Daniel Covey, Disaster Recovery Specialist Solution Architect, CloudEndure, AWS

- Sarah Michaels, Product Marketing Manager, CloudEndure, AWS

Further reading

For additional information, refer to:

- AWS Elastic Disaster Recovery product page
- Disaster Recovery of Workloads on AWS: Recovery in the Cloud
- Disaster Recovery Plan Checklist
- Affordable Enterprise-Grade Disaster Recovery Using AWS

Document history

Change	Description	Date
Whitepaper updated	Removed CloudEndure Disaster Recovery, added AWS Elastic Disaster Recovery (AWS DRS), added best practices for AWS DRS networking, wave planning, drills, failover, and failback, and added information and best practices for monitoring a disaster recovery solution.	November 14, 2022
Initial publication	First publication.	January 7, 2021

AWS Glossary

For the latest AWS terminology, see the AWS glossary in the *AWS Glossary Reference*.

Notices

Customers are responsible for making their own independent assessment of the information in this document. This document: (a) is for informational purposes only, (b) represents current AWS product offerings and practices, which are subject to change without notice, and (c) does not create any commitments or assurances from AWS and its affiliates, suppliers or licensors. AWS products or services are provided "as is" without warranties, representations, or conditions of any kind, whether express or implied. The responsibilities and liabilities of AWS to its customers are controlled by AWS agreements, and this document is not part of, nor does it modify, any agreement between AWS and its customers.