


Architecting a Cloud Disaster Recovery Service



Introduction

Disaster recovery (DR) is an integral part of business operation continuity. A severe disruption of normal operating facilities, including hurricanes, floods, earthquakes, and human-made disasters, rarely stops customers, clients, and other business dependents from expecting service. Businesses are often expected to be available 24x7, and that means they need to plan for DR. Typically, in the event of a disaster, DR involves the restoration of business services by redeploying applications and infrastructure in new datacenters or locations.

Contents

Business Considerations When Choosing a DR Solution	2
Traditional Pre-Cloud DR Solutions	3
First-Generation Cloud-Based DR Solutions	4
Integrated Cloud-Based DR Solutions	5
The Best of Both Models	5

Prior to receiving and providing services, customers and internal business units often negotiate service-level agreements (SLAs) that must be met. One of the most important steps in DR planning is working out agreements with stakeholders. The agreements must meet the demands of the average customer or client yet still be feasible to carry out while meeting the SLA time requirements. Specifically, the SLAs will include a set of service-level objectives, such as time to recovery and point of recovery.

Without a functional DR plan, a business is at risk. The cost of services being unavailable includes lost revenue during the outage, as well as long-term impact on customer confidence in your business's ability to deliver services.

Ideally, DR plans should minimize the time to recovery but also keep the cost of the DR operation to a minimum.

When determining SLAs, it's important to keep in mind the opportunity cost of having your business systems down. Depending on your type of operation, you may incur lost revenue due to your sales system being down. Sales won't be made, and cash flow will halt. If you provide a long-term service or business solution, customer frustration will increase. This is especially the case if you provide support services or self-serve services customers depend on for their needs, such as checking the status of or reviewing personal accounts, that go down in a disaster.

Ideally, DR plans should minimize the time to recovery but also keep the cost of the DR operation to a minimum. However, these objectives can be in conflict when using traditional, dedicated backup data centers. The closer the configuration of the backup data center is to the production data center, the faster you'll likely restore services and avoid incidents while operating from a DR center.

Business Considerations When Choosing a DR Solution

While building a DR plan, businesses need to prioritize applications, workloads, and data. There must also be a set order in which applications and data will be restored to avoid further conflicts and interruptions.

First, a DR plan should include a complete inventory of applications and the data used by them. Customer-facing applications and services should be at the top of the priorities list, along with key operational systems that must work in order to keep other applications running.

Some back-end services that aren't customer-facing can be placed lower on the list of priorities. For example, it may be more important to get your order processing system up and running first and then tend to the inventory tracking system later. In this instance, sales will continue, and customers can continue to do business. Inventory tracking, while important, can be caught up on later after a more complete recovery has been made.

Additionally, business intelligence and analytics systems, such as machine learning workloads, depend on data from other systems. These workloads would likely have a lower priority than most online transaction processing systems.

When creating SLAs, businesses must define two objectives: recovery point objectives (RPOs) and recovery time objectives (RTOs). Your chosen recovery point will depend on your tolerance for data loss.

While building a DR plan, businesses need to prioritize applications, workloads, and data. There must also be a set order in which applications and data will be restored to avoid further conflicts and interruptions.

Different use cases will have different tolerances. For instance, accounting systems can't tolerate any data loss that will leave them in an inconsistent or incorrect state. On the other hand, an inventory system has some room for error. In a system that tracks the SKU of each item in a warehouse, some data loss may cause an item to be over-counted or undercounted by a small amount. This problem can be remedied later and isn't usually disastrous to business operations. In an even more extreme case, if data from a data mart were lost completely, administrators could rebuild the contents of the repository by re-running extraction, transformation, and load operations.

down to how similar a business needs a DR environment to be to their normal operation environment. A DR environment with small capacity and resource availability will have a low cost, but may lead to decreased performance or the elimination of low-priority systems and workloads.

On the other hand, a DR environment identical to the normal production environment will be able to carry out all normal business needs, but will have a higher cost.

Traditional Pre-Cloud DR Solutions

Traditional DR solutions tended to use either a cold site or active site recovery model. They each have benefits and drawbacks.

When using a cold site recovery model, a business contracts with a third party to provide needed hardware and network infrastructure during a disaster, which can lead to longer recovery times because teams will need to deploy a functional production environment on the infrastructure.

With cold site recovery, you have access to infrastructure, but your applications aren't deployed and data isn't up-to-date, and it will take time to deploy both. Cold site recovery is best for use cases with long recovery time objectives. Businesses that use this model should also

The more stringent the RTOs and RPOs, the greater the cost to a business.

RTOs depend on a business's tolerance for downtime. Your business's services and customers' expectations will determine how much downtime is tolerable.

The more stringent the RTOs and RPOs, the greater the cost to a business. Cost-benefit tradeoffs usually come

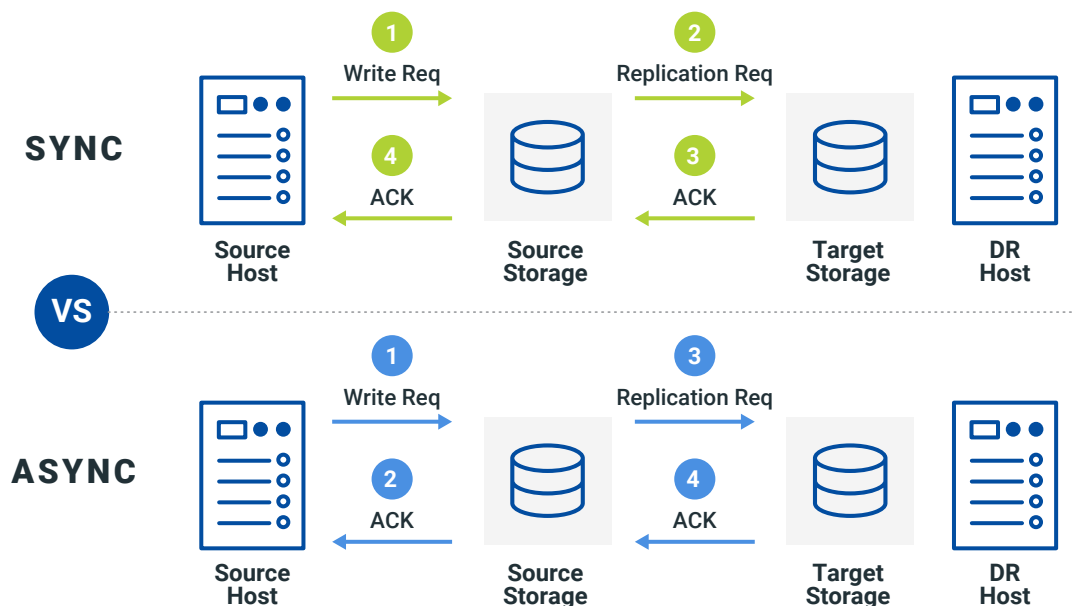


Figure 1: Synchronous vs. asynchronous replication.

have the personnel available to spend the time needed to install and configure all necessary systems and data.

The active site recovery model avoids some of the biggest drawbacks of the cold site recovery model, especially the longer time to recovery. This approach requires provisioning and maintaining dedicated hardware in a second facility. Applications are available and maintained in the active site even when the site isn't used for DR purposes.

Avoiding long recovery time comes with an ongoing cost of its own: With active recovery, you must keep replicating data because data is constantly changing in a production environment. If the DR site is to be ready at any time to take over production workloads, the data in the two sites must be kept up to date. This data replication may be done synchronously or asynchronously (see **Figure 1**).

Synchronous replication minimizes RTOs because each time a change is made in a production database, it's made in the corresponding database in the DR site. When a transaction is performed on the production database, the transaction isn't considered complete until both databases have been updated.

One of the considerations when choosing between synchronous and asynchronous is tolerance for increased latency.

For example, if an inventory record is updated in the production database, the database management system would then initiate another update operation on the DR database. When the update is completed on the DR database, that database signals the production database that the data has been successfully updated and the transaction on the production system can be completed.

One of the considerations when choosing between synchronous and asynchronous is tolerance for increased latency. Synchronous replication can entail longer transaction times because data has to be transmitted between data centers during the transaction.

With asynchronous replication, data is updated in the DR database, but the transaction on the production system doesn't need to wait for the DR database transaction to be completed first. This enables longer distances between sites because transactions in the production system will be completed once data is sent to the backup site, without having to wait for the backup database to write and confirm the transaction.

This approach allows for longer latency when writing to the DR database without slowing the production system. In cases where the DR database can be out of synchronization for longer periods of time, a business may even choose to use batch processes to update the disaster recovery database.

Whichever form of DR you choose, you will need to maintain and test the DR infrastructure and systems. In addition to the maintenance costs, a DR environment is essentially unused capacity most of the time. Fortunately, virtualization and cloud technologies offer a better solution.

First-Generation Cloud-Based DR Solutions

Virtualization is redefining DR options. With virtualized infrastructure, there's no need to have dedicated hardware already provisioned in an active site or to wait for a team of system administrations and software engineers to deploy the latest versions of applications and data to a cold site.

Instead, virtual machines (VMs) running on-premises can be replicated in the cloud, essentially creating an active recovery site on demand. The key advantage to this approach is that the business doesn't have to pay for unused capacity; it only pays for the recovery site infrastructure while it's in use.

Many tasks remain the same as pre-cloud solutions. Network engineers, for example, would have to create virtual private networks and implement firewall rules in the recovery site. Authentication and authorization controls would also have to be established in the DR environment, and are especially important to keep up to

date. Someone who had a privilege revoked a week ago shouldn't have that privilege available in a recovery site because authorization controls weren't kept in sync.

Maintaining up-to-date data in the recovery site, however, is still a challenge. One option is to maintain a hot standby in the cloud, much like a synchronously updated recovery database in the active site model.

Alternatively, a business may be able to take advantage of snapshotting disk images and storing them in lower-cost object storage. From there, they can be used to initialize disks on VMs that are started only in response to a disaster at the production site.

Integrated Cloud-Based DR Solutions

DR doesn't have to be a time-consuming, manual, and sometimes error-prone activity. Automation can help. IT engineers can automate many of the tasks of defining and creating a DR environment.

Adopt the practice of using policy-driven operations and have well-defined policies in place so that engineers don't have to make guesses on the fly.

For example, system administrators can define VM templates or golden images so that they can be started rapidly. Database administrators can replicate data continuously so that it's up to date and available in the cloud. Network configurations can be scripted. The added advantage is that configuration code can be treated like any other application code and kept under version control.

Adopt the practice of using policy-driven operations and have well-defined policies in place so that engineers don't have to make guesses on the fly. For example, there should be no question about what AWS region should be used for a disaster, or the order in which applications are brought online. Policies should clearly specify where to establish DR resources. Replication policies should also be in place.

Keep in mind that different systems may have different replication policies. High-priority systems may require synchronous replication while lower-priority applications can be maintained using batch updates. Policies and plans should be tested regularly to verify that images and scripts function as expected in the DR environment.

Ideally, system administrators and application owners should be able to manage recovery operations from a single pane of glass. This kind of platform should give administrators everything they need in one place, so they don't have to jump from one dashboard to the next trying to piece together a comprehensive picture of the state of a DR operation.

The Best of Both Models

DR has traditionally been costly and difficult to implement. Businesses often had to balance the need for rapid recovery with the costs of maintaining unused infrastructure in an active recovery site.

Virtualization and cloud technologies now enable businesses to have the best of both models. VMs and applications can be started in the cloud when they're needed. Defining infrastructure and network configuration as code enables more automation, which reduces the time to deploy, as well as the risk of error.

Combine that with proper administration tools, such as a single pane of glass for managing DR operations, and DR is no longer the costly and difficult process it has been.