Disaster Recovery

OTAVA

1.0 Executive Summary

Mapping out your disaster recovery infrastructure, identifying key operations, and developing a testing strategy to efficiently recover and restore data and infrastructure is a complex, long-term project.

Why have disaster recovery?

Protecting your data is imperative, as it is critical to your business operations. Without the ability to protect your data, your business can't survive. Therefore, disaster recovery is a significant consideration for your organization. There are many parts to your production application that can fail at any time, including people, processes and technology. You must have a plan in place to address each of these components.

This white paper addresses at a high level the framework of disasters, components of disaster recovery infrastructure and steps to a successful disaster recovery. It's ideal for executives and IT decision-makers seeking an introduction and up-to-date information regarding disaster recovery best practices, including specific technology recommendations.

(source: https://www.otava.com/blog/why-you-should-value-your-data/)

2.0
Disaster Recovery and Business Continuity

While disaster recovery is an important part of the overall business continuity structure, it's not a replacement for business continuity. There are many components to a business continuity plan, including emergency contacts and roles, client and employee communication plans, and procedures and protocols after a disaster is declared. The best IT disaster recovery plans fit within an overall business continuity plan. disaster recovery plan tackles the challenge of coordinating efforts and navigating a complex communication and workflow model in the event of a disaster. The plan must identify and support the complex interdependencies typically found in a larger organization that all work to keep the business running.

3.0
Framework of Disasters

There are three types of disasters: **Natural**, **man-made** and **technical failures**. Disaster recovery initiatives often focus on natural disasters; however, it's important to understand that a variety of disasters can put an organization out of business if there is no recovery plan in place.

(source: https://zap.zerto.com/wp-content/uploads/2018/11/IDC_ITResilience_Info.pdf /)

Natural disasters

Natural disasters are perhaps the least likely to impact a production site yet get the most attention. Ice storms, hurricanes and flash floods are all examples of natural disasters that can destroy a business. To avoid as many of them as possible, find a recovery site that offers a safe geography. While the coastal regions have a higher population density that makes them attractive areas for a data center, they are subject to more risk from natural disasters as well as the usual risks from technical failures and man-made disasters. The Midwest, however, with its cooler temperatures and relatively low risk for natural disasters is one area that provides a safe haven for recovery and production sites alike.

Within the Midwest, tornadoes are the natural disaster most likely to affect the area. Tornadoes tend to move in an easterly direction, so placing your recovery and production sites in a north-south configuration reduces the risk of exposure.

Man-made disasters

Man-made disasters are far more likely to affect businesses than natural disasters, but often get less attention than natural disasters. These failures, however, can be as equally devastating as a flood or hurricane. Examples of man-made disasters include a terrorist attack, cyber attack, a car crash that triggers a power outage or digging that damages a fiber line.

Process failures

In our experience, process failures are the most common type of disaster failure, when poor administrative policies and procedures allow for human error. Examples of process failures include patch failures that allow for vulnerability, or a user who destroys data or a critical production system because an administrator failed to remove their login credentials. Process failures also include failure of quality control on manual functions with something as simple as a password change or firewall rule changes.

Software failures

When bad code affects your software, the repercussions can be severe enough to declare a disaster. Because of this bad code, software failures fall into man-made disasters.

Software failures can comprise:

- A corrupted database due to a SQL attack
- A memory leak, which can diminish an application's performance or cause it to fail
- A bug that destroys your data

To protect production and recovery sites from such failures as much as possible, software must be updated, patched and maintained on a regular basis.

Cyber attacks

Cyber attacks are becoming more and more prevalent in today's online environment. Many large organizations have been taken down by cyber attacks[1], and businesses are spending more than ever on IT security[2]. When designing a recovery plan, you must consider the risks of a cyber attack such as DDOS or phishing. The best recovery designs protect your data and help you stay in a production state when your site is under attack.

Technology failures

We've found that technology failures are the most common for our clients and get the least attention when planning for disaster recovery. Technology failures encompass your entire infrastructure, which include building, power, computing, storage, and network failures. All of these elements are important to consider when creating an IT disaster recovery plan and are discussed below.

Infrastructure failure

Below are the seven five major causes of an infrastructure failure, as well as real-life examples of each. It's these types of failures that make disaster recovery planning such a time-consuming, complex project that can be extremely difficult to manage alone.

[1]http://www.nytimes.com/interactive/2015/02/05/technology/recent-cyberattacks.html
[2]http://www.gartner.com/newsroom/id/3135617

Building failures

Building failures comprise damage to the data center itself and can instantly cause your production site to go downproblems for your production site. A secure building is the first step to housing a safe infrastructure, and if your building is no longer secure, your equipment is at risk.

Examples of building failures include:

- Destroyed equipment from a collapsed roof
- Someone driving a backhoe into the building (oops)
- Broken pipes in the building

When choosing a data center building, make careful consideration as to the quality and strength of the materials used in its construction. Your building should be able to withstand any of the above examples of building failures as well as whatever Mother Nature can throw at it.

Power failures

The most common example of a power failure is, of course, a power outage. This can be caused by any number of factors, including storms or human error managing the electric grid. One such example is the power outage of August 2003 that caused 50 million people to lose power in much of the Northeast and parts of the Midwest as well as Canada for several days. That outage was in part due to a software bug that caused an alarm system failure.

Computing failures

Servers, whether virtual or physical, are one of the most basic tenets of any IT infrastructure, and a failure can have a huge impact on your your business. A computing failure is similar to a utility failure but is more specific to the technology it powers.

Computing failures include:

- Motherboard failures
- CPU overheating from a broken server fan
- A battery failure in a RAID card

Storage failures

Storage systems are essential to any site because that is where its data is kept. When a storage failure occurs, that directly affects your the data, and therefore your business operations of any organization. A lost encryption key, for example, is potentially devastating for any business. Without the key, data is inaccessible, you must consider encryption keys in your plan.

A different second type of storage failure is a bug within the storage unit that corrupts the data. And, a final example of storage failure is a subsequent disk failure during a rebuild of a server following a failed disk in a RAID group. If any of these circumstances happen to your production site, you should consider your data lost, and you must rely on your recovery plan to maintain business operations.

Network and security failures

The network comprises the Internet connectivity, equipment and physical fiber of your site, and it'is what keeps you connected to your users. If your the network goes down, it is impossible to communicate to your customers if you do not have athere is no backup site. Examples of network failures include a loss of iInternet connectivity due to a bandwidth overload, a cut fiber, a firmware bug or a network hardware failure.

Network security devices also have the possibility of failing, such as a Web Application Firewall (WAF), AV server, or log monitoring system. These types of failures are common enough that your recovery plan must be designed so you can remain in a production state should these systems go down.

The above failures are only some of the aspects one must plan for when setting up a recovery site. There is more to consider ahead in 4.0, Components of Disaster Recovery infrastructure.

(source: https://www.scientificamerican.com/article/2003-blackout-five-years-later/)

4.0 Components of Disaster Recovery Infrastructure

Now that the framework of disaster types has been discussed, let's consider the components of your infrastructure to help prevent or protect against as many of them as possible. Keep in mind that each component of your infrastructure, whether production or recovery, should have appropriate security and compliance protocols in place as required.

Creating a complete disaster recovery infrastructure is complex, and each component should factor into whether you decide to build your own recovery site or outsource it to a third-party provider. It's a painful expense to have a fully duplicated production stack that sits idle for months, but it sure is handy (and worth the expense) when the time comes. The following outlines the challenges behind each layer of disaster recovery infrastructure.

Dedicated hardware

Whether the hardware for your recovery site is cloud based or a dedicated physical server, a successful recovery site starts with hardware. If you use physical servers, you must be able to handle the associated costs and resources that come with it. This is where the benefits of having managed cloud servers come in—the cloud provider handles the expenses associated with upgrading and maintaining the hardware.

Network

For enterprise or more complex server configurations, more than just a server image is required for recovery. Firewall rules, VLANs, VPNs and the network must be fully replicated at the disaster recovery site before it can go live. It's imperative the network be replicated accurately and securely, because if the production site fails and proper settings aren't in place at the recovery site, it won't function properly.

Circuits

You must account for the telecommunications circuits between your two sites, as well as other locations. If you opt not to use circuits, a VPN must be created between your locations. You should have redundant circuits for your replications to minimize data loss. Make sure to buy the right size circuits—if your circuit is too small, data cannot be replicated properly and end users cannot be served during a disaster. If it's too large, you risk wasting money.

Encryption

Encryption, already required for many compliance regulations such as PCI, must be in place at the disaster recovery site. For some compliance regulations, data must be protected in transit and at rest, meaning end-to-end encryption is required. Depending on your application, encryption can sometimes present performance problems, and should be planned for accordingly.

Disaster recovery backup

Offsite backup ensures data is always available in the event of a disaster or corruption between recovery and replication sites, and is a separate process from disaster recovery. Therefore, whatever backup solution you have for your production site must be replicated to your recovery site. This includes a secure, compliant geographically safe location, as well as the staff, technology and resources to manage such a site. You must continue to back up your data even if you are running in your recovery site, or you risk permanent data loss.

Static public IP

The graphic to the right details the challenges behind a static public IP address and managed DNS when failing to a recovery site. If you have devices that have static IP addresses pointing to your servers, you'll need point them to your recovery site when you fail over. If you use a managed DNS, you'll need a process for who will change the DNS record so it points to the recovery site. This should be documented in your DR playbook who's responsible for changing the records.

As of now, the process is mostly manual, but there are some vendors who will help you automatically reconfigure your DNS records to point to your recovery site during a disaster. This comes in handy should your designated "DNS guy" be unavailable for any reason.

Failover Process For A Static Public IP

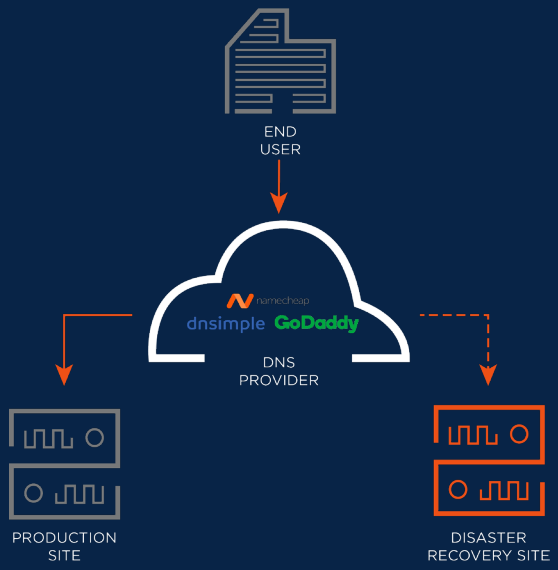

During a disaster, someone needs to log in to the DNS Provider and update it to point to the recovery site. This is a manual process.

Maintenance and operations in both locations

After you've allocated the resources needed for the technical layers of disaster recovery infrastructure, you must also budget for staff to oversee, maintain and operate the replication and additional infrastructure for the recovery site. Simply dividing your current staff among two sites is not enough to maintain a successful disaster recovery site. Recovering the production site during a disaster will keep your staff busy, so people will need to be dedicated specifically to your disaster recovery site.

Auxiliary services

Often, auxiliary services are required for your production site to operate. Your disaster recovery plan must take into account those services as well. Without them, you'll have a dead recovery site.

Typical auxiliary services include:

- Authentication devices such as an active directory or radius, and time-keeping systems (required for compliance) that help keep your devices in sync and communicating properly.
- Email, monitoring, file, print, and FTP servers.
- SSL certificates

All of the necessary supporting systems will need to be replicated, tested and managed in your disaster recovery site, and require maintenance and upgrading as necessary.

Data replication technology

You must take into consideration the technology you use to replicate your data so each site can communicate. This adds complexity and is another tool to manage and maintain that normally isn't needed in a single-site application.

4.1 Total Cost of Ownership

Managing, auditing and budgeting for each of the above infrastructure components can be quite expensive. For many organizations, the ability to outsource their infrastructure is ideal, because it alleviates some of those costs. However, some businesses are wary of extending their circle of trust to an outside party and prefer to host their recovery infrastructure themselves. If that is the case, the following section discusses the total cost of ownership (TCO) of such a project.

The cost of implementing, managing and maintaining your own disaster recovery infrastructure is extensive. The equation to the right illustrates the TCO an organization should expect when adding DIY recovery infrastructure to your production TCO.

Whether you decide to host your infrastructure yourself or outsource it, every organization must have essential components to a successful recovery site. The following section discusses in more detail those elements to ensure a viable disaster recovery infrastructure for your business.

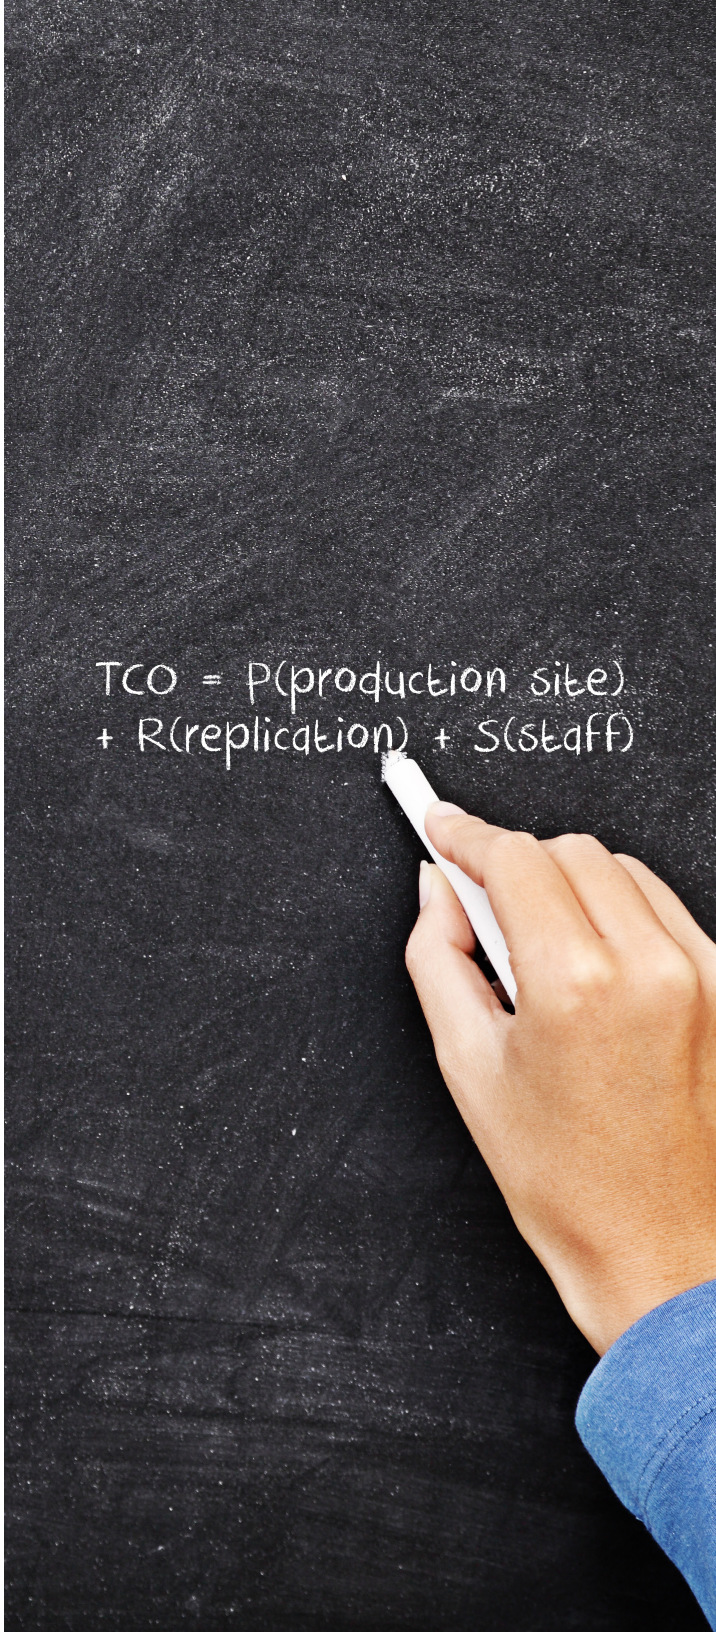

5.0. Components to a Successful Disaster Recovery Infrastructure

A successful disaster recovery infrastructure, whether provided as a service or built in-house, addresses all of the technical and economic challenges outlined in 3.0 and 4.0. Outsourcing your disaster recovery infrastructure to a service provider gives you a simple, technical solution to a complex problem and allows you to take advantage of resources and knowledge you might otherwise not have. No matter who provides your infrastructure, there are four main components to a successful recovery site: A simple environment, testing, compliance, and transparency.

Simple environment

Create a simple disaster recovery environment by virtualizing as many production servers as possible. With virtualization, the entire server, including the operating system, applications, patches and data can be replicated to another site. This virtual server can be copied or backed up to an offsite data center and spun up on a virtual host in minutes in the event of a disaster. This can reduce recovery times, compared to traditional disaster recovery approaches where physical servers need to be loaded with the OS and application software as well as patched to the last configuration used in production before the data can be restored.

Cloud servers can also be mirrored, or running in sync, at a remote site to ensure failover in the event that the original site should fail, ensuring complete data accuracy when recovering and restoring after an interruption.

Another aspect of cloud-based disaster recovery that improves recovery times is full network replication. As mentioned in 4.0, replicating the entire network and security configuration between the production and disaster recovery site as configuration changes are made saves you the time and trouble of configuring VLAN, firewall rules and VPNs before the disaster recovery site can go live.

Testing

Sometimes testing your disaster recovery strategy is so difficult and distracting that organizations can't even bother doing it. Therefore, testing must be easy to do. Ideally, testing should be done after any change to your production site, but that's not always practical. Testing must be done in a way that assumes a minimal amount of risk while keeping the recovery environment as up to date as possible.

Once you pick a testing schedule, automate adherence to it to maintain regular testing of your environment. This will ensure you are always ready in case of a disaster. We recommend testing your recovery site twice a year to make sure it's running optimally.

Compliance

If your production environment requires audits for compliance with operating standards such as HIPAA or PCI, you must also audit your DR environment. If you fail over to a non-compliant site, you risk fines or possible criminal liabilities, not to mention a severe security risk to your data.

Transparency

When you have a transparent view into your recovery infrastructure and receive automatic alerts regarding the success or failure of various components, it's much easier to spot problems and be proactive in tackling any potential issues. This allows you to manage your resources more efficiently and worry less about the manual processes behind such notifications. We strongly recommend your disaster recovery strategy has appropriate notifications and dashboards to make sure replication happens and the recovery site is always available.

6.0. Conclusion

Disaster recovery is a complex, expensive task for any organization, despite the advancements in technology to streamline the process. There are three major types of disasters to prepare for, and there are also many components to each of those disasters that must be considered. However, having a recovery site has enormous benefits and is imperative for any organization seeking a solid overall business continuity plan. Leveraging the capabilities of a full service provider allows an organization to realize these benefits, including a simple, virtualized environment and cost-effective, efficient testing. A transparent environment that allows you to automate routine notifications and upgrade infrastructure as necessary are also key elements to a successful disaster recovery infrastructure. All of these steps, combined with an overall business continuity plan, will allow you to focus more on your core mission and worry less about your IT infrastructure.

READY TO GET STARTED WITH SIMPLE DR?

Contact us today. We got this!

OTAVA

EXPECT EXCEPTIONAL

OTAVA provides secure, compliant hybrid cloud solutions for service providers, channel partners and enterprise clients. By actively aggregating best-of-breed cloud companies and investing in people, tools, and processes, Otava's global footprint continues to expand. The company provides its customers with a clear path to transformation through its highly effective solutions and broad portfolio of hybrid cloud, data protection, disaster recover, security and colocation services, all championed by its exceptional support team. Learn more at www.otava.com.

OTAVA.com (877) 740-5028 solutions@otava.com