Ntirety™
Secure. Managed. Cloud.

How to Move On:

Exploring DR Business Continuity

TABLE OF CONTENTS

How to Move On:

Exploring DR Business Continuity

Business continuity is mission-critical for your organization and proactively safeguarding against lost data and downtime is a must. A disaster will hit your enterprise—it's only a matter of time. It's important to protect your organization's data and integrity by understanding the steps you need to take and the nature of the threats you face.

In this eBook, we'll cover the various types of disaster recovery that are available, and how the cloud can be a harbor in the storm. Plus, we'll provide insights into why security and disaster recovery are interlinked, a primer on readying your enterprise for the inevitable, and how to choose a provider that can keep you whole.

The Many Names and Faces of Disaster Recovery

When discussing disaster recovery, people often throw out a variety of words and terms to describe their strategy. Sometimes, these terms are used interchangeably, even when they mean very different things.

Disaster Recovery

This is a term that has been making the rounds since the mid- to late- '70s. Although the meaning has evolved slightly over time, the disaster recovery process generally focuses on preventing loss from natural and man-made disasters, such as floods, tornadoes, hazardous material spills, IT bugs, or bio-terrorism. Many times, a company's disaster recovery plan is to duplicate their bare metal infrastructure to create geographic redundancy.

Recovery Time Objective (RTO)

As you build your disaster recovery strategy, you must make two crucial determinations. First, figure out how much time you can afford to wait while your infrastructure works to get back up and running after a disaster. This number will be your RTO. Some businesses can only survive without a specific IT system for a few minutes. Others can tolerate a wait of an hour, a day, or a week. It all depends on the objectives of your business.

Recovery Point Objective (RPO)

The second determination an organization must make as they discuss disaster recovery is how much tolerance they have for losing data. For example, if your system goes down, can your business still operate if the data you recover is a week old? Perhaps you can only tolerate a data loss of a few days or hours. This figure will be your RPO.

IT Resilience

This term measures an organization's ability to adapt to both planned and unplanned failures, along with their capacity to maintain high availability. Maintaining IT resilience is unique from traditional disaster recovery in that it also encompasses planned events, such as cloud migrations, datacenter consolidations, and maintenance.

Load Balancing

To gain IT resilience and keep applications highly available, companies must engage in load balancing, which is the practice of

building an infrastructure that can distribute, manage, and shift workload traffic evenly across servers and data centers. With load balancing, a downed server is no concern because there are several other servers ready to pick up the slack.

Streaming giant Netflix often tests the load balancing ability of their network with a proprietary program called Chaos Monkey. Using this tool, they ensure that their infrastructure can sustain random failures by purposefully creating breakdowns throughout their environment. This is a great example for companies to follow. Ask yourself: What would happen if someone turned off my server or DDOSed my website? Would everything come crashing to a halt if an employee accidentally deleted a crucial file?

Backup

Backups are just one piece of the disaster recovery puzzle. Imagine if you took a snapshot of your entire workload and replicated it on a separate server or disc—that is a backup. With backups, you always have a point-in-time copy of your workload to revert back to if something happened to your environment; however, anytime you must revert to a backup, anything created or changed between the time the last snapshot was taken and the time the disaster occurred will be lost.

Failover Cluster

Another piece of the disaster recovery puzzle, failover clusters are groups of independent servers (often called nodes) that work together to increase the availability and scalability of clustered applications. Connected through networking and software, these servers "failover," or begin working, when one or more nodes fail.Which type of failover server you choose depends on how crucial the system is, along with the RPO and RTO objectives of the disaster recovery plan. Failover servers are classified as follows:

- **Cold Standby:** Receives data backups from the production system; is installed and configured only if production fails.
- **Warm Standby:** Receives backups from production and is up and running at all times; in the case of a failure, the processes and subsystems are started on the warm standby to take over the production role.
- **Hot Standby:** This configuration is up and running with up-to-date data and processes that are always ready; however, a hot standby will not process requests unless the production server fails.

Replication

This term represents the process of copying one server's application and database systems to another server as part of a disaster recovery plan. Sometimes, this means replacing schedule backups. In fact, replication happens closer to real-time than traditional backups, and therefore can typically yield an adherence to shorter RPO and RTO.

Replication can happen three different ways:

- Physical server to physical server
- Virtual server to virtual server
- Physical server to virtual server

Database Mirroring

As with backups and replication, database mirroring involves copying a set of data on two different pieces of hardware; however, with database mirroring, both copies run simultaneously. Anytime an update, insertion, or deletion is made on the principal database, it is also made on the mirror database so that your backup is always current.

Journaling

In the process of journaling, you create a log of every transaction that occurs within a backup or mirrored database. These logs are sometimes moved to another database for processing so that there is a warm standby failover configuration of the database.

At the end of the day, what you really need is business continuity.

A well-formed business continuity plan will use all of these methods to ensure your organization can overcome serious incidents or disasters. Going beyond availability, business continuity plans determine how your business will continue to run at times of trouble.

Can your business survive a systems failure? Can it survive a situation where your offices burn down? How quickly can you access your mission-critical data and mission-critical applications? How will people access your mission-critical applications while your primary servers are down? Do you need VPNs so employees can work from home or from a temporary space? Have you tested and retested your business continuity plan to ensure you can actually recover? Does your plan follow all relevant guidelines and regulations?

The right mix of solutions will depend on the way your business operates, the goals you're trying to achieve, and your RPO and RTO targets. In the end, the resilience of any IT infrastructure or business comes down to planning, design, and budget.

With the right partner to provide disaster recovery and business continuity management services, you can come up with a smart plan that proactively factors in all risk, TCO goals, and availability objectives.

Where the Business Cloud Weathers the Storm

Mission critical apps need special consideration

In 2017, two of the biggest storms in modern history battered cities across the country. The damages were estimated at $290 billion, which doesn't include the decreased or lost productivity that is still to come for offices and businesses across the regions.

It's impossible to predict everything, including the timing and scope of hurricanes and natural disasters, but ideally businesses prepare for situations just like these. It wasn't that long ago that disaster recovery consisted largely of the restoration of services and configurations from backup. This was a process that could take hours or even days.

We live in an era where the advantages of cloud and hosting have proven to be critical components to sustaining operations, communications, and recovery efforts for business, residents, and emergency services.

There is little doubt that those organizations affected by these events that hosted their own infrastructure in the geographies that were affected endured significant unfortunate interruptions. Meanwhile, professional hosting and cloud companies in those same geographies kept running for several reasons.

The modern data center is made purposefully for these types of situations—built in facilities designed to withstand disasters, with generators, emergency protocols, redundant communications, redundant (underground and satellite) internet connections, and more. Through each element of a hosted and cloud infrastructure, both best practices and uninterrupted services are the mission of providing for customers.

If you're concerned about continuity and minimizing the impact of unplanned IT events on business, these storms make it very difficult to justify hosting your own infrastructure any more.

Mission-critical? Not all costs are equal

Disaster preparedness and recovery only begin the discussion that stewards of IT and the leadership of organizations need to understand when it comes to their mission critical software and systems. It has been proven that a disaster-based outage for those systems is a business-impacting event with far greater cost than any perceived operational costs.

The loss of critical data, interruption of IT operations, loss of equipment, loss of communications, potential for injury, or even cyber-security incidents are risks that any business needs to guard against in general, but especially if it touches their mission-critical systems.

Looking for any glimmer of a silver lining from past hurricane seasons, these natural disasters offer an opportunity to reevaluate and update resiliency, backup, and recovery strategies and make the case to include cloud and hosting in the conversation when thinking about the most precious jewels of IT systems.

Think it through

Let's say you're ready to have this disaster readiness conversation. You then systematically consider all the ways mission-critical systems can fail. You quickly find that critical databases, DNS, networking, web applications, mobile, email, and other elements are not only fallible, they are likely interdependent. You diligently backup systems, but find that, if you're using an offsite tape system as your main channel for recovery, you must get ready for extended downtime. Tape recovery is slow, snapshots may not take place frequently enough, recovery isn't guaranteed, and it may take significant reconfiguration to return to operations once everything is restored.

In situations where the power is offline, your on-site generators may only allow for a graceful shutdown until utility services are restored. This can only add to the impact of a greater disaster because you can't even start to recover until the basics are stabilized. These and other factors mean that recovery can add up to a lot of contiguous hours and sleepless nights while the business breathlessly waits to get portions of functionality back online.

Natural events are not the only threat. Both human and systematic errors can impact a business, taking resources offline thanks to interdependencies, corruption, and cybersecurity events. An average business may not have the adequate resources to detect, contain, and react to an ever-widening range of potential threats.

Cloud answers

This is one of the most compelling benefits of cloud and hosting solutions. The cloud has answers that are built into not only the robustness of data centers themselves but the very nature of the platform that critical business systems run on.

Switching over to a cloud or hosting service offers expanded options during and after a disaster. From backup to redundant infrastructure, to ever-present monitoring, disaster planning, and availability—they are the foundation that hosted services and infrastructure are built on from day one.

Data is designed to be safe within a cloud and in many cases, the ability to recover can be as simple as a few clicks. In the right environment, there may not even be a recognized impact to information systems. The human capital is also something that is easily overlooked, but these companies have the trained 24x7 staff to sustain these disasters for very prolonged periods.

Catastrophic natural events present an opportunity to re-evaluate and (if needed) create a formal disaster recovery plan for every mission-critical app, especially for health care and financial institutions where loss of data mounts to heavy penalties on top of the more obvious business risks.

Starting out, IT professionals need to find out where everything is, evaluate any weaknesses. Appraise existing backups and recovery procedures. Test. Plan on recovery from every angle. Test this too. Script out what personnel will do. Run through this in detail. Make sure people are ready to leave their loved ones behind to take care of the IT infrastructure. Draft communications. Get commitments from people. Validate that every piece works and everyone knows what to do. Then test it again. Calmly and objectively assess your gaps. They will be many. That is where cloud and hosting come in.

Finally, when it comes to the tools and elements of an infrastructure, the abilities that spring from the spectrum of cloud and hosting-enabled features can help build a better, more resilient, more redundant disaster solution for a fraction of the cost that it can be built internally. Cloud provides assurance and capabilities in disaster situations that no other solution can deliver. Make it a part of your conversation.

When Security is Your DR Strategy

Both Disaster Recovery (DR) and Security/Compliance are critical aspects of your business, but many fail to see how these strategies relate to each other. For example, many security/compliance standards have requirements that speak to the state of the backups you're running – i.e. Are they encrypted? How long are they kept?

Beyond this simple example, many of the tools and processes used in one area directly apply to the other, and the best way to achieve your goal – 100% application uptime – is to use them in conjunction with each other.

There are alternative mitigation techniques, which take advantage of Disaster Recovery services, that you should consider along with these common suggestions:

1. **Regular backups of your data, including the OS**
   If the ransomware is trying to hold your data hostage, it won't matter if you have other copies of that data. And all the better if you have the ability to perform bare metal restores as well.
2. **Replication technology with journaling**
   Some technology used for replicating servers and virtual machines, such as Zerto, has a journaling feature that allows you to fail over to your DR site in a historic state rather than a current one. Similar to using backups, but potentially much less downtime if you have a complete DR solution that is regularly tested.In both examples there is one caveat – you have to accept the fact that you're not using the most absolute recent version of your data, but potentially something hours (or days, or weeks) old in order to failback to a time before the ransomware was installed.

In both examples there is one caveat – you have to accept the fact that you're not using the most absolute recent version of your data, but potentially something hours (or days, or weeks) old in order to failback to a time before the ransomware was installed.

If DR can be applied to enhance Security, is the converse statement true? Absolutely. While having a DR site and the ability to failover to it is nice, there is always a price to pay. Whether in loss of data (which could be minimal, admittedly) or in the time it takes to complete the failover, which grows in exponential proportion to the complexity of the environment, DR is not likely the whole solution.

Case in point, most companies address DR by considering only potential catastrophes such as natural disasters and by ensuring they only need to failover in an absolute emergency, e.g. a natural disaster at the primary datacenter.

But you know what's not a natural disaster? Being hacked, that's what. And though it's not natural, it can have the same disastrous potential. That's where security, applied with DR, provides some good news – if you're addressing both as part of your cloud DR strategy, you're covering as many bases as you can.

To sum up, both categories of service should go hand in hand, and significant investments in one should lead to similar investments in the other, or else you're just shifting your risk from one bucket to another one, and hackers don't care about your buckets.

Too often, companies put artificial limits on their spend based on categories that are meaningless – i.e. "We can spend 10% of revenue on Security, but only 5% on DR". That's like choosing whether to leave your front door unlocked or your windows wide open. In either case, be prepared to lose some furniture.