8 Key considerations for planning Disaster Recovery on public cloud

WHITE PAPER

MICROLAND®
Extraordinary. Everyday.

WHITEPAPER

8 Key considerations for planning Disaster Recovery on public cloud

Abstract

In today's digital world a disaster can strike anytime, anywhere. In such situations a Disaster Recovery(DR) plan is critical. Smart IT leaders are putting in place plans that detail out the steps that can be taken to be better prepared in the likelihood that disaster strikes. With more and more organizations moving to the cloud, disaster recovery becomes easier (compared to on-premise) and is significantly faster. To ensure cloud-based DR plans are effective, this paper recommends 8 key areas to consider:

1. Defining your business's tolerance levels for downtime and data loss in the event of disaster, i.e. your Recovery Time Objective (RTO) and Recovery Point Objective (RPO)
2. Having the right hardware and software in place to help manage the crisis if disaster strikes
3. Identifying the right personnel to be involved in the recovery process (with clear roles and responsibilities)
4. Having in place an ongoing communication plan to keep all stakeholders abreast of the situation on a real-time basis
5. Putting in place backup systems that are ready and accessible quickly in case of an outage, so that critical work can continue while the recovery process is going on
6. Ensuring SLAs with third-party vendors are comprehensive enough to cover their role in the event of disaster
7. Identifying what data needs protection, and planning how to go about it
8. Testing the plan regularly to iron out any issues or challenges

8 Key considerations for planning Disaster Recovery on public cloud

Introduction

Disasters are unavoidable. From the storm of the century to the blackhole severing a power line at a local construction site, disasters come in many forms. However, even the most mundane of 'disasters' can have a devastating effect on your business if it keeps you from interacting with customers or destroys data.

A Business Continuity Institute poll conducted by risk experts found that 85% of the people who took part in the survey had concerns that their businesses were at risk of a cyber-attack within a period of 12 months from the time the poll was conducted. In today's digital world disasters can mean a significant disruption in business operations. Which is where the Disaster Recovery (DR) process comes in. The process encompasses a custom designed strategy for your business to resume normal computing capabilities in as little time as possible in the event of a disaster. It is important that the DR plan is carefully strategized and implemented with scalability in mind.

Most businesses confer to the National Institute of Standards and Technology (NIST) SP 800-34 defined IT contingency plans which outline a six-step planning process:

1. Develop the Contingency Planning Policy Statement
2. Conduct the Business Impact Analysis (BIA)
3. Identify Preventive Controls
4. Create Contingency Strategies
5. Plan Testing, Training and Exercises
6. Plan Maintenance

One of the most common challenges for companies today is to understand what Disaster Recovery (DR) really means and how to do it successfully. While companies trust a Managed Service Provider (MSP) to build it for them, but in most cases, they never fully test the solution built for them. At Microland, we have put together a strategy that is functionally akin to an offsite backup or even a high availability (HA) solution and has gone through a rigorous testing process to cause least bit of disruption to the client business.

From offsite to recovery in the Cloud

Enterprises of all sizes have started using the public cloud, primarily for backup, recovery and archiving. However, there are a series of questions that need to be considered. Can the public cloud be used as a DR site? What are the pros and cons of such an approach? What best practices should apply? Is it only for small businesses that do not have a secondary site of their own? Or is there relevance even for larger organizations?

Why cloud-based DR...

While virtualized cloud platforms are well suited to providing DR, under normal operating conditions, a cloud-based DR service may only need a small share of resources to synchronize from the primary site to the cloud. The complete set of resources needed to run the application can be provisioned (and paid for) in the event of an actual disaster.

Automated virtualization platforms for DR ensure that additional resources can be rapidly activated when needed, thus significantly reducing recovery time after a failure. Additionally, by its very nature, on-demand cloud computing provides significant cost-benefit when peak resource demands are much higher than average demand. Cloud allows IT teams to maintain the state of an application using low cost resources under ordinary operating conditions and bring to use the more powerful and high-cost resources only when disaster strikes.

Additionally, a cloud DR site can be triggered through a normal laptop or mobile device with a wireless internet connection even if the 'as-designed' connectivity fails. They are also easy to maintain while ensuring high levels of security in alignment with security, privacy and compliance standards and best practices.

Forward thinking IT leaders are continuously seeking an affordable, manageable solution to protect their operations from data loss. Which is where cloud-based DR can act as a boon by enabling them to replicate business data to a secondary off-site location without building and managing an additional IT data center and infrastructure; which is a great cost and management effort saving benefit. And, for smaller enterprises and businesses, the cloud provides the opportunity to implement a DR plan that would have not been otherwise possible.

8 Key considerations for planning Disaster Recovery on public cloud

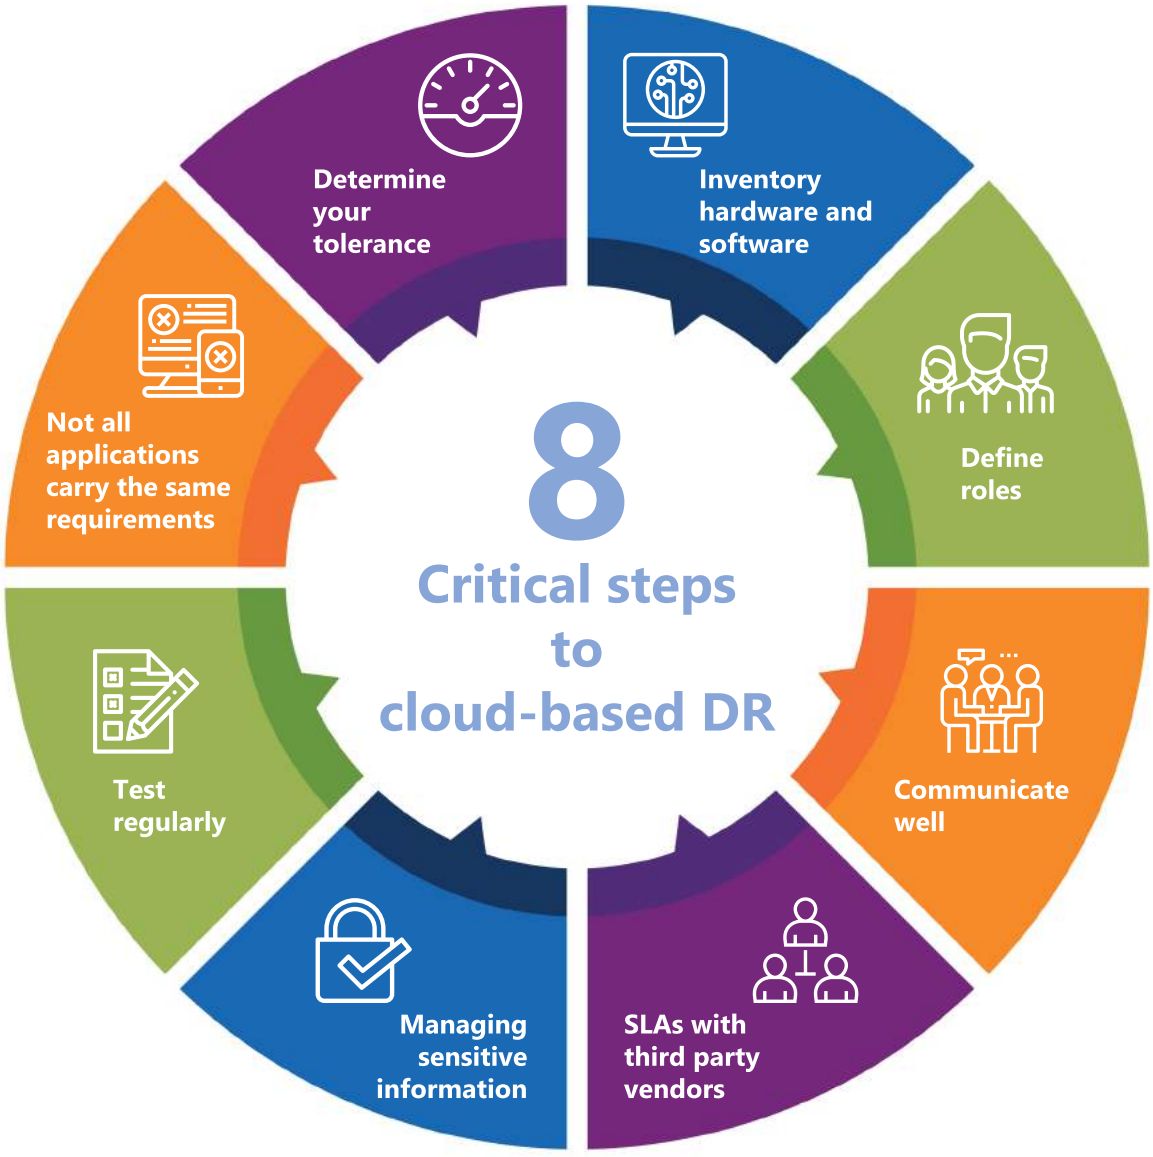

MICROLAND®

8 Critical steps to cloud-based DR

Now that we have established a business case for cloud-based DR, let us look at the 8 key considerations for planning DR on public cloud.

1. Determine your tolerance

Best practice organizations prioritize their applications before defining any DR services. And this is the most crucial step. When evaluating your options, Microland recommends considering two key metrics:

Recovery Point Objective (RPO): RPO is the degree to which a business can tolerate data loss. For example, 'immediate RPO' indicates zero tolerance for data loss. A 24-hour RPO would indicate that restoring data as of yesterday's backup is sufficient for your business. A thorough evaluation is needed to identify and define what would work for your business.

Recovery Time Objective (RTO): This metric determines the maximum tolerable time for recovering the data and bringing the application back online. Here, it's important to note that the RTO should include the time to restart systems, databases and applications and re-route communications. Simply restoring the data alone is not enough.

The IT teams need to assign both RTO and RPO to the business' applications and data, and they're both driven by the cost of downtime. Some key considerations to keep in mind when driving this process: Which application can wait for four hours to be restored but cannot afford data loss? Which application can be rolled back to a 30-minute recovery point, but RTO is under an hour? Which applications need to come up in the first 12 hours and others in 36 hours or more?

Remember, the cost of downtime can include actual loss of revenue, loss of employee productivity, loss of customer goodwill and loss of reputation. The focus shouldn't be only on tangible financial losses.

2. Not all applications carry the same requirements

Not only do different businesses have different RPO and RTO targets, applications within a business will have different requirements as well. For instance, customer-facing applications usually demand a shorter RPO and a lower RTO because the loss of data and downtime can have a severe impact on the business. Whereas administrative applications may be able to withstand more downtime or a higher level of data loss.

Setting RTO too long or the RPO too high can put the organization at unacceptable levels of risk. Conversely, setting RPO and RTO at levels that are too aggressive leads to over-investment and ties up capital that could be spent in more productive ways.

Microland experts advise to divide your applications into three tiers. Tier 1 should include the applications you need immediately. These are the mission-critical apps you can't do business without. Tier 2 covers applications you need within eight to 10 hours, even up to 24 hours. They're essential, but you don't need them right away. Tier 3 applications can be comfortably recovered within a few days.

3. Inventory hardware and software

Your DR plan should include a complete inventory of hardware and applications in order of priority. Each application should have the vendor technical support contract information and contact numbers, so you can get back up and running quickly.

4. Define roles

All DR plans should clearly define the parties to be involved, their roles and responsibilities during a DR event. Among these responsibilities must be the decision to declare a disaster. Having clearly identified roles will garner a universal understanding of what tasks need to be completed and sets role clarity. This is especially critical when working with third-party vendors or providers. Protocols for a DR plan must include who and how to contact the appropriate individuals on the DR team, and in what order, to get systems up and running as soon as possible

5. Communicate well

Perhaps one of the most overlooked components of a DR plan is having a good communication plan. Communication is critical when responding to and recovering from any emergency, crisis event or disaster. So, having a clear communication strategy that identifies effective and reliable methods for communicating with employees, vendors, suppliers and customers is key. It is not enough to communicate in a timely manner just during the initial notification of an emergency. Having a written process in place to reference ensures efficient action post-disaster and alignment between organizations, employees and partners.

Communication also includes ensuring that the staff knows where to go (an alternate site if the primary office is unavailable), where to sit and how to access the systems from that site. It is a good idea to provide a map to the alternate site and make sure there are seating assignments in place.

6. SLAs with third party vendors

In case your business has outsourced the technology requirements, or if the systems are stored in a data center/co-location facility, it is important to ensure that there is a binding agreement with the vendors that defines their level of service in the event of a disaster. This will ensure that they start working on resolving the issue within a specified time.

7. Managing sensitive information

Defining operational and technical procedures to ensure the protection of sensitive information is a critical component of DR plan. These procedures should address how sensitive information will be maintained and accessed when a DR plan has been activated.

8. Test regularly

If you're not testing your DR process, you don't have one. Your backup may have failed, your supply chain may rely on someone incapable of dealing with disaster, your internet connection may be too slow to restore your data in the expected amount of time, the DR key employee may have changed. There are a lot of things that may break a perfect plan. The only way to find them is to test them.

Conclusion

Virtualization of businesses has made it imperative that IT leaders put their DR plan in the cloud. Best practice organizations are leveraging cloud computing to develop DR capabilities that are both not cost intensive and simpler to deploy, compared to traditional methods.

A DR plan that is detailed, flexible and scalable to meet ongoing business needs is key. However, since this process can be expensive, time-consuming and complex, it is important to plan and determine what your business needs are. When planning it is advisable to not focus only on the technical process, but make it holistic by building people, process and technology to make it an effective DR plan. Bringing on board an able partner who clearly understands these nuances and can add value to the client organization could be a real game changer for organizations.

https://www.drj.com/myblog/7-emerging-trends-in-disaster-recovery-industry.html

About the author

Sobha Chandana Areti
Architect - Cloud & Datacenter

Sobha Chandana Areti is a leading architect for Cloud & Datacenter services at Microland. She comes with more than 10 years of experience in the IT industry with specialized skills in AWS Cloud infrastructure services. An MBA graduate from MKU, Sobha is a cloud enthusiast and a constant learner.

For further information
Contact us at: **+1 646-254-3598** or Email us at : **info@microland.com**

About Microland

Microland accelerates the digital transformation journey for global enterprises enabling them to deliver high-value business outcomes and superior customer experience. Headquartered in Bangalore, India, Microland has more than 3,800 professionals across its offices in Australia, Europe, India, Middle East and North America. Microland partners with global enterprises to help them become more agile and innovative by integrating emerging technologies and applying automation, analytics and predictive intelligence to business processes.

Learn more about us at:

www.microland.com

MICROLAND®