



OneNeck White Paper

Cloud Enabled Disaster Recovery

Maintaining an “Always Ready” Environment

Many organizations depend on the 24/7 availability of their mission critical IT systems. Yet how many companies have a well documented disaster recovery plan that enables operations to continue in the face of significant facility and/or human loss? And how many test the plan regularly, like once each quarter? Probably not too many.

We've seen a number of real disasters during the past several years. What comes to mind when thinking "disaster" are terrorist attacks, hurricanes, earthquakes, and the like. But, disasters also include events such as a significant vendor suddenly declaring bankruptcy or a system security breach.

Should your company experience a sudden and unrecoverable loss of an entire data center, for example, do you know what you need to do to resume operations in another data center? How much data will you lose in resuming operations? How long will it take for you to resume operations? Businesses must answer these core questions to define what qualifies as an acceptable response to a real disaster. Then, they must put the processes in place to support their answers.

Processes should include organization redundancy, such as moving to an alternate data center in a reasonable time frame, and data accessibility after a disaster. Cloud computing can serve as an important foundation of a rapid recovery, low data loss disaster recovery solution. A cloud solution enables you to rapidly recover without paying to run an entirely redundant data center 24/7.

While the concept of disaster recovery planning has been around for nearly 40 years, many companies find themselves coming up short despite their increased dependence on IT systems. Businesses typically agree that they want a disaster recovery solution. However many are daunted by the high cost to setup and maintain an "always ready" disaster recovery site. The complexities and costs of a disaster recovery solution historically have caused businesses to cut corners or risk no planning whatsoever. The cloud can remove these and many other barriers.

DISASTER RECOVERY PLANNING BASICS

Disaster recovery planning (DRP) is often considered a component or subset of business continuity planning (BCP). Disaster recovery generally refers to the processes and procedures used to recover after a disruptive event. It often focuses around the critical IT systems within an organization. Business continuity planning usually includes the business functions that need to continue after a disruption so your business continues to make money.

At a high level, disaster recovery planning usually involves the following steps:

Identify the scope and boundaries – This is typically the first step towards completing your disaster recovery plan. Identifying the scope involves prioritizing the critical systems for disaster recovery and assigning a value to the failures of those systems. This is also the phase in which you determine the RTO and RPO requirements from the business units for critical applications.

Establish the Budget – Budgeting for disaster recovery plans can be tricky. Often you will want to do an assessment of the costs to the business of suffering different disaster scenarios. Comparing different options for recovery can vary the costs of the disaster recovery plan. Reducing RPO and RTO requirements can soften the financial costs of the disaster recovery plan. But you should be realistic and ensure executive management understands the risks of data loss and system availability being stretched out. Both IT and executive management must come to an agreement on the budget and IT will work within the constraints of the budget that has been established.

Develop and Deploy the Plan – Developing and deploying the plan can be the most involved part of this process. Often the plan is actually a "script" of activities that occur in order and are executed by a recovery team made up of resources from IT. Roles and responsibilities are assigned in the plan as well. Deploying the plan involves choosing the tools and technologies needed to meet the RTO and RPO requirements established in the first step while still working within the constraints of the budget.

Test – Test, test, test. Disaster recovery plans are simply not effective if they are not properly and frequently tested. Test the systems you're going to use in recovery regularly to validate that all the pieces work. Always record your test results and update the disaster recovery plan to address any shortcomings. As your business environment changes, so should your disaster recovery plan.

The Recovery Time Objective (RTO) is the duration of time and a service level within which a business process must be restored after a disaster to avoid unacceptable consequences associated with a break in continuity.

The Recovery Point objective (RPO) is time of data loss that is acceptable. If the RPO of a company is 2 hours and the time it takes to get the system back into production is 5 hours, the RPO is still 2 hours.

DISASTER RECOVERY BEFORE THE CLOUD

In the past, disaster recovery planning involved infrastructure deployments that could cost nearly as much as your primary point of presence. Disaster recovery plans were very expensive and time consuming to create. The production infrastructure that was defined as mission critical essentially had to be replicated in the disaster recovery site. This duplication involved a second set of servers, network equipment and expensive replication software to maintain a reasonable

RTO and RPO for recovery. To add to the financial sting, this solution also meant a lot of wasted server and network infrastructure resources sitting idle until needed.

To survive challenging conditions and uncontrollable threats, you need to create a reliable, efficient infrastructure. You also need to develop a business continuity strategy that keeps your business running, protects your data, helps retain your customers and ultimately helps reduce your total operating costs over the long term.

Some businesses, simply constrained by their budget, would resort to disaster recovery plans that involved only a subset of their mission critical infrastructure and compromise on systems that shouldn't be compromised. Worse yet, some companies' disaster recovery plans involve procuring equipment and secondary data center space after the disaster had occurred and relying on offsite data backups to restore their systems. Recovery time and recovery point objectives were measured in days or weeks rather than minutes or hours.

CLOUD RECOVERY

Businesses need a better solution that meets critical system's aggressive RTO and RPO requirements and fits within their budget. The advent of cloud computing has provided an excellent platform to enable "Cloud Recovery." Cloud Recovery leverages the elasticity of the cloud to deliver cost effective disaster recovery solutions for businesses. Businesses no longer need to invest in or manage dedicated server and network hardware in their disaster recovery site. Cloud infrastructure inherently enables easy and frequent data replication between sites and systems – another key component to maintaining data integrity and reducing RTO and RPO times. Because the cloud resources are pooled, the costs to deliver Cloud Recovery solutions to businesses are a fraction of the cost normally associated with a disaster recovery environment.

Whether your existing production infrastructure is physical, virtual or in the cloud today, a Cloud Recovery solution could be an excellent way to execute your organization's DRP needs.

ONENECK'S CLOUD RECOVERY OFFERING

RTO and RPO to Fit Within Your Budget. OneNeck's Cloud Recovery solutions enable aggressive RTO and RPO options to fit within your disaster recovery plan budget. Solutions can be tailored per server depending on your needs. As a result, no reason exists that allow budget constraints to cause a compromise on which critical systems get left behind after a disaster.

Cloud Recovery in OneNeck's Data Centers. OneNeck hosts and manages its Cloud Recovery infrastructure in its state-of-the-art data centers in the Southwest U.S. The data centers, developed from industry best practices, allow many of its customers to leverage a disaster recovery facility that would otherwise be inaccessible. OneNeck's data centers meet the stringent requirements for Tier 3, SAS70 Type II backed operational facilities. The OneNeck data centers provide a secure, reliable environment for the Cloud Recovery infrastructure maintain by a 24/7 on-site team.

99.99% Guaranteed Availability. By leveraging OneNeck's cloud architecture and skilled experts, you have the comfort of knowing your disaster recovery operations are safe and available 24/7. During a disaster recovery event, the Cloud Recovery systems are backed by the same production hosting guarantee offered by OneNeck – 99.99% availability.

No More Hardware Life-Cycles to Manage. OneNeck's Cloud Recovery offering frees up capital. There is no need to purchase servers and network gear for your DRP. Businesses don't have to manage hardware lifecycles including support contracts and hardware maintenance agreements on these systems.

Disaster Recovery Replication. An investment in your disaster recovery plan would be useless without consistent replication of your critical data from production to the cloud recovery environment. OneNeck provides replication solutions for your intellectual property using a variety of replication solutions including SAN level replication, virtual machine replication and file system level replication.

Global Load Balancing and Automated Failover. Meeting aggressive RTO needs is easy with OneNeck's Cloud Recovery solution. The OneNeck Cloud Recovery includes global load balancing options to speed automated failover in the event of a disaster.

Flexibility and Scalability. Changing disaster recovery infrastructure to keep pace with production changes are now simple to manage in the OneNeck cloud. Adding replication target space or server resources such as processors, RAM and disk can be accomplished quickly, efficiently and often within 24 hours.

WAN Connectivity Solutions. OneNeck provides a carrier neutral environment and cross connects from the telecommunications demarc to the Cloud Recovery environment. In the event of a production failure, access to the Cloud Recovery environment can be accomplished over a dedicated connection to a customer's existing private WAN (i.e.: MPLS) or via site-to-site VPN links over the Internet.

Lower DRP Costs. Leveraging OneNeck's Cloud Recovery solution can reduce overall IT costs for businesses by providing a fixed, predictable spend for your DRP infrastructure needs. Businesses no longer incur the variable costs associated with dedicated physical DR assets.

THE “ANYWHERE, ANYTIME” APPROACH

Industry experts predict cloud computing will transform IT in a major way over the next decade. By providing rapid access to dynamically scalable and virtualized IT resources, companies can be more responsive to potential disruptions and data loss. Virtualization technologies combined with new cloud management techniques enable an “anywhere, anytime” approach for IT services. Cloud computing can provide a higher-quality, more efficient infrastructure while simultaneously providing significant cost savings.

To survive challenging conditions and uncontrollable threats, you need to create a reliable, efficient infrastructure. You also need to develop a business continuity strategy that keeps your business running, protects your data, helps retain your customers and ultimately helps reduce your total operating costs over the long term.