Mastering Disaster Recovery: Business Continuity and Virtualization Best Practices

WHITE PAPER

vmware®

Table of Contents

Introduction

Imagine if you were able to slash both planned and unplanned downtime while protecting all of your important systems and applications, independent of specific hardware requirements. In this paper you will learn how to move beyond the challenges of traditional management approaches to disaster recovery (DR) and master how to build a true business continuity solution.

Challenges of Traditional Disaster Recovery

Traditional disaster recovery plans depend on a very complex set of processes and infrastructure: duplicate datacenters, duplicate server infrastructure, processes for getting data to a recovery site, processes for restarting servers, processes for reinstalling operating systems, and so on. Because disaster recovery can be complex, organizations often find themselves unable to deliver protection to no more than a few privileged production workloads, leaving other workloads (e.g., file/print servers, internal web servers, departmental applications) unprotected and poorly covered.

The complexity of traditional disaster recovery plans and data centers pressures organizations to be heavily dependent on extensive product and process training, on the accuracy and completeness of thick paper "runbooks" that document the recovery process, and on perfect execution of the recovery process when an outage does occur. Moreover, as testing can be disruptive and expensive, organizations are limited in their ability to ensure that all of their training, documentation, and execution is practiced and can successfully recover their IT services.

As a result of these challenges, tests of recovery plans often fail. Basic recovery of critical workloads – if successful at all – often takes days or weeks, and a significant amount of IT time and resources are consumed by managing and maintaining recovery plans. In short, most firms fail to meet the continuity requirements set by their organizations.

Requirements for Building Business Continuity Solutions

Mastering business continuity requires three key steps focusing on reliability, platform independence, and broad application support. First, it's important to build a disaster recovery solution using a reliable platform. After all, the more reliable the platform, the less likely it is that you'll experience downtime in the first place. Choose an infrastructure platform that's battle-tested, flexible, dynamic, and reliable.

Second, the solution needs to be independent of physical infrastructure. Less dependence on specific hardware and configurations means more flexibility in how you can protect and recover applications in the face of downtime. The IT environment is constantly in flux, and servers will continue to be upgraded and their configuration changed. None of this should impact your business continuity solution. A strong advantage for virtual machines in DR planning is that they always see the same set
of virtual hardware, which is transparently mapped to the underlying physical hardware. Virtualizing compute, network, and storage means that your applications don't care what kind of physical equipment they're running on. You can move them around and they'd never know it.

Third, instead of protecting only specific applications or operating systems, what if you could extend protection across the broadest range of applications and operating systems. An optimal solution won't require making trade-offs about which applications are worthy of protection. Virtualization such as that from VMware encapsulates any standard x86 operating system and the applications that run on it. No modifications or special drivers are required. And all of the protections that are part of the VMware platform, such as data protection, availability, and disaster recovery, apply across all virtual machines. This translates into the broadest degree of protection and also greatly reduces the cost of protection.

Reducing the Risk of Disaster Recovery Failures

A big part of what makes traditional disaster recovery challenging is how to effectively control the risks inherent in a complex IT environment. A number of these risk drivers have an impact
on the ability to carry out a rapid, reliable recovery process.

One of the biggest risks is that traditional DR plans don't actually match the dynamics of your current IT environment. Environments are constantly changing, so if you're not keeping your DR plans, which is hard to do with traditional DR, in sync with those changes, it's likely that any subsequent DR test or failover scenario will fail.

Another risk is human error. Recovery plans can be very complex, with a lot of moving parts. Even if the IT staff is trained on how to execute the process, it's likely that mistakes will be made along the way; It hard to manage all the required steps and application dependencies that is needed during an actual disaster. Also, in general, a lack of automation around the recovery process increases the recovery time.

In the event of a major disaster, will key DR staff be available? What if they are home with their families or out of contact? In that case, could another member of the IT staff perform the recovery process successfully? Dependency on certain people is a risk in a recovery scenario.

As a result of these and other risks, some of the common costs involved in an unsuccessful or prolonged recovery process include:

- The cost of lost revenue and lost employee productivity while key systems are down
- Ballooning staff overtime costs, as people begin working around the clock to restore critical applications and infrastructure
- Effects on customers and end-users; for example, imagine if your website or mission critical apps were down for even 30 minutes, how much business could you lose?

Reducing and Managing Recovery Risk

Non-virtual environments create significant recovery risk. Physical DR can be complex – with a lot of duplicate infrastructure, lots of configuration settings to keep updated and in sync across sites, lots of processes and moving parts involved in a recovery scenario. Infrequent testing of DR plans leads to a "testing gap", where organizations can't really be sure that the DR plans on file actually match the current IT environment's configuration. After all, environments are constantly changing, which drives up the recovery risk.

Two key things help reduce and control recovery risk:

- Virtualization encapsulates entire servers, including all configuration data, so the risk of a failed recovery is greatly reduced; the current application/OS configuration is always a part of a virtual machine. In this way, virtualization greatly reduces the recovery risk for servers and applications, and can hold that risk level constant.
- Frequent and automated DR testing ensures DR plans are complete, accurate, and can be reliably executed every single time. Frequent testing reduces the testing gap to ensure that your DR plans match your current IT environment. This helps ensure that overall recovery risk stays low and that infrastructure changes don't drive up recovery risk.

Key Features of Virtualization for Disaster Recovery

- Hardware Independence. This means being able to recover onto any x86 hardware. So, you have the flexibility to buy different servers for your recovery site, or even fewer servers. Continuing to virtualize your production site will free up additional machines that can then be moved over to your recovery site.
- Encapsulation. Because virtualization captures everything about a server into just a few files on disk, the real benefit here is mobility and the ability to move your VMs wherever you want. You can back them up in the same way you currently protect your other files. You can also replicate them to your disaster recovery site, so that they will be available when you need to recover from an outage.
- Partitioning and Consolidation. Server consolidation means doing more with less. You're reducing your physical footprint, which also means you can streamline your disaster recovery plans and standardize your recovery process.
- Resource Pooling. Allows you to make better use of both your production and recovery infrastructure. In terms of leveraging your investment in recovery infrastructure, it also means that you can be running other workloads at your recovery site. Because you can map out recovery resources ahead of time, you can guarantee that your recovering VMs get access to the resources they need at your recovery site. This is how you are able to safely run other workloads without interfering with your ability to fully recover.

Preparing Your DR Program

There are several critical items to consider and prepare as you set up a disaster recovery program:

- What is the maximum downtime you can accept? What is your Recovery Time Objective (RTO)?
- How much data loss can you accept? What is your Recovery Plan Objective (RPO)?
- What kind of recovery plan do you need for a partial disaster, such as a storage array failure, compared to a complete disaster, such as an entire facility being unavailable because of fire damage?
- How do you classify your computing resources for restoration after a disaster? Are they critical, urgent, important, normal, or nonessential?

Why VMware Software for Business Continuity?

With VMware vCenter™ Site Recovery Manager, VMware has leveraged the disaster recovery features and capabilities of the VMware vSphere™ platform with a product developed specifically for disaster recovery. Site Recovery Manager simplifies and automates the key elements of disaster recovery: setting up disaster recovery plans, testing those plans, executing failover when a datacenter disaster occurs, and failing back to the primary datacenter.

Site Recovery Manager makes it possible for customers to provide faster, more reliable, and more affordable disaster recovery protection than previously possible. Although not a part of VMware vSphere, Site Recovery Manager works closely with VMware vSphere to manage and automate disaster recovery for virtual environments.

Disaster Recovery Scenarios with Site Recovery Manager

vCenter Site Recovery Manager can be used in a number of different failover scenarios:

- Active-Passive: Site Recovery Manager supports the traditional active-passive DR scenario, where a production site running applications is recovered at a second site that is idle until failover is required. Although the most common configuration, this scenario also means that you are paying a lot of money for a DR site that is idle most of the time.
- Active-Active: To make better use of the recovery site, Site Recovery Manager also enables you to leverage your recovery site for other workloads when you aren't using it for DR. Site Recovery Manager can be configured to automatically shutdown or suspend VMs at the recovery site as part of the failover process so that you can easily free up compute capacity for the workloads being recovered.
- Bidirectional: Site Recovery Manager can also provide bidirectional failover protection so that you can run active production workloads at both sites and failover to the other site in either direction. The spare capacity at the other site will be used to run the VMs that are failed over.
- Local Failover: Although less common, some IT teams need to be able to failover within a given "site" or campus, for example when a storage array failure occurs or when building maintenance forces you to move workloads to a different campus building. These customers are leveraging Site Recovery Manager to perform these failovers.

Mastering Disaster Recovery Planning

Let's take a look at how vCenter Site Recovery Manager helps with setup of disaster recovery plans.

- Create recovery plans. Site Recovery Manager allows users to create recovery plans for different scenarios and different parts of their infrastructure. Recovery plans could apply to a single VM, a single application, all VM's connected to a specific array, a business unit, or even to a whole site. Each recovery plan specifies the recovery process for the virtual machines covered by that recovery plan.
- Integrate with replication. Even in a vSphere environment, ensuring that replication is properly configured requires close coordination between application, server and storage teams and involves manually verifying that all data that needs to be protected is indeed being replicated. Even in a virtualized environment this is complicated by the fact that while array-based replication is generally configured for LUN's, multiple VM's usually share a single LUN and a single VM might span multiple LUN's. Using the storage replication adapters, Site Recovery Manager can directly query the replication software to determine which LUN's are replicated. Because it is closely connected to VMware vCenter Server, Site Recovery Manager can then determine which VM's are protected by replication and then displays that information to the user to make it easy to validate that replication is properly configured.
- Map recovery resources. Another part of DR setup requires specifying what resources will be used for recovery of the virtual machines at the recovery site. Site Recovery Manager helps users specify how to map computing resources used at the production site to recovery site resources, how to map network resources used at the production site to recovery site network resources, and how to map the management objects (e.g., folders and hierarchies) you've configured in vCenter Server on the primary site to how you want those to appear at the recovery site.
- Specify recovery process. Traditional disaster recovery plans often use a runbook to document the different steps and processes in recovery. This runbook specifies all of the steps required for recovery, which includes things like how to validate hardware configuration, the startup order to follow for different systems, etc. These runbooks, often a physical binder, are hard to make accurate, hard to keep up to date, and hard to train people on. Site Recovery Manager makes it possible to turn that physical runbook into software—within the Site Recovery plug-in in vCenter Server you specify the key steps in the recovery process. You can extend and customize the recovery process by adding scripts and call-outs into the recovery process.

Summary

Become a master of disaster recovery and achieve these key benefits:

1. Expanded protection. VMware software overcomes many of the cost and complexity challenges that have made it difficult for organizations to effectively protect more than a select few of their important x86 systems and applications.
2. Slashed planned downtime. Eliminating sources of downtime such as hardware maintenance and storage reconfiguration can significantly improve the availability of your IT infrastructure.
3. Minimized unplanned downtime. VMware software provides a reliable platform for your computing infrastructure, provides features like VMware HA that minimize downtime from server failures, and helps you to deliver rapid and reliable automated disaster recovery with Site Recovery Manager.

All of these benefits are independent of what x86 hardware you're using, what operating system, and what applications you have in your environment. So don't just manage your diaster recovery process. Master it.

VMware, Inc. 3401 Hillview Avenue Palo Alto CA 94304 USA Tel 877-486-9273 Fax 650-427-5001 www.vmware.com