



A Guide to Modern IT Disaster Recovery

How to prepare for and mitigate the “Black Swan Events” in your data center



www.vmware.com

INSIDE:

INTRODUCTION: What are Black Swans, and what do they have to do with your data center?

CHAPTER 1: Expect the Unexpected

CHAPTER 2: Start with an *Intelligent and Virtual* Foundation

CHAPTER 4: Disaster Recovery Best Practices — Top 10

CONCLUSION: A Quick-Start Guide to Disaster Recovery

APPENDIX: Disaster Recovery 101—The Basics

INTRODUCTION

1

What are “Black Swans,” and what do they have to do with your data center?

You know VMware as the virtualization company that has been the market leader for the past 11 years. In fact, according to Gartner, more than 80% of all virtualized applications in the world run on VMware today. This ebook highlights the VMware perspective on disaster recovery in the data center. But let’s forget about IT for a moment.

The **Theory of Black Swan Events** is a metaphor that encapsulates the concept of *surprise events that have a major impact*. It refers to unexpected events of large magnitude and consequence and their dominant role in history. Such events, considered extreme outliers, play vastly larger roles than regular occurrences.

The Black Swan, a book by Nassim Nicholas Taleb, explains that while Black Swan Events

are unpredictable, a person or organization can plan for negative events, and by doing so strengthen their ability to respond, as well as exploit positive events. Taleb contends that people in general — and specifically enterprises — are very vulnerable to hazardous Black Swan Events and are exposed to high losses if unprepared.

There is an obvious parallel between the Theory of Black Swan Events and the need for disaster preparedness for your critical IT assets.

Deploying automated disaster recovery (DR) is the way to protect IT and business from unpredictable events — even Black Swan Events. The following chapters explain the basics of DR and the required infrastructure. They also offer DR hidden realities and best practices with real-world advice.

CHAPTER 1

2

Expect the Unexpected

Our hope is you never need to activate an IT disaster recovery plan. Our job is to provide automated protection if you do need it.

YOUR DATA CENTER IS YOUR CASTLE. That’s where all your critical IT components — hardware, data, and software — reside. You protect it with the latest bulletproof security solutions and make it reliable through redundant multiprocessing, highly scalable platforms, and super-fast optical networks.

And yet, it is not fully protected from forces beyond your control such as natural disasters; man-made events like road closures; security procedures or partner service interruption at a specific site.

Downtime and loss of data, even if temporary, can have long-lasting effects for business and can contribute to the demise of the otherwise well-lubed business:

- ▶ Loss of revenue from your customers’ inability to do business with you
- ▶ Diminished market credibility and customer trust, resulting in churn
- ▶ Penalties for violated SLAs with partners, suppliers, distributors, and franchisees
- ▶ Costs of recovering and repairing the lost data
- ▶ Legal costs of meeting internal and external compliance requirements

How do you balance the disaster recovery risk and investment equation? Is the potential risk greater than the investment? Let’s put it in perspective:

- ▶ 43% of companies experiencing disasters never reopen, and 29% close within two years.¹
- ▶ 93% of businesses that lost their data center for 10 days went bankrupt within one year.²
- ▶ 40% of all companies that experience a major disaster will go out of business if they cannot gain access to their data within 24 hours.³

“CIOs and IT organizations should consider scenarios in which normal operations could be disrupted and adopt/adapt practices and technologies that enable them to deal with potential disruption from hostile, external actions as well as internal system failures.”

—Gartner’s *Top Predictions for IT Organizations and Users, 2011 and Beyond*

These stakes are as high as your entire business, and it is within your power to mitigate the risk.

¹ McGladrey and Pullen

² National Archives & Records Administration

³ Gartner, December 2009

“DR is the IT industry’s way to prepare and fight the Black Swan Events.”

CHAPTER 2

CHAPTER 2

continued

Start with an *INTELLIGENT* and *VIRTUAL* Foundation

Reliable → Repeatable → Recoverable

UNTIL RELIABLE VIRTUALIZATION MANAGEMENT SOLUTIONS became available several years ago, DR solutions fell well short of satisfying business requirements due to the following:

- ▶ High Cost
- ▶ Complexity
- ▶ Lack of Reliability

With traditional manual DR solutions, the **high cost** came from the need to deploy a second failover site with dedicated infrastructure, software licenses, and human personnel. The **complexity** was high because, to ensure the recovery of entire business services, the recovery plans had to manipulate many individual components and moving parts: applications, hosts, network, and storage. The **lack of reliability** of these procedures was diminished by low automation and inability to test any recovery procedure.

Many organizations had limited confidence in meeting their Recovery Point Objective (RPO) and Recovery Time Objective (RTO) in the event of a disaster. IT departments were hesitant to expand disaster protection, uncertain whether the quality of the insurance was really worth its cost.

Virtualization is fundamental and critical to the success of DR planning. Virtualization abstracts the complexity of hardware and software and allows standardization of processes, thus making planning and automation of the recovery procedures much more reliable and repeatable.

Physical Recovery Process — 40 Hours



Virtual Recovery Process — 4 Hours



In fact, in a recent IDG survey, 70% of customers achieved improved BC/DR with virtualization.¹

An intelligent virtual infrastructure based on VMware is the right foundation for the modern DR solution. Highly adaptive and scalable, it is optimized for business-critical workloads with built-in intelligence.

The VMware DR solution provides:

- ▶ The simplest way to replicate applications to a secondary site
- ▶ The simplest way to set up recovery and migration plans
- ▶ Fully automated, most reliable site recovery and migration

¹ IDG Research, *Benefits of Virtualizing Business Critical Applications*, March 2011

Cost-efficient DR: With the rapid adoption of virtualization and the evolution of replication technology, DR is becoming more cost-efficient. Virtualization enables infrastructure consolidation at the failover site. Less costly replication options are more broadly available, using lower-end storage appliances or stand-alone software solutions. With these advances, DR can protect large-scale mission-critical IT assets, as well as smaller sites and Tier 2 applications.


Automated DR: In virtual environments, end users are shielded from the complexity of managing each step in the recovery process. Now, a DR solution can automatically execute and coordinate all the steps required to ensure the desired level of protection. Traditional runbooks are no longer “good enough” to manage recovery plans and are replaced with software-driven recovery

plans. Setting up a recovery plan in a virtual environment is as simple as selecting RPOs and RTOs for each business service.

Reliable site recovery and migration: With virtualization, organizations get a much stronger assurance that they can meet their RPOs and RTOs. Virtualization provides the ability to test recovery plans frequently in a non-disruptive manner. Manual recovery processes are now replaced with automated recovery, eliminating the risk associated with user errors and ensuring predictable recovery.

The chart below shows how organizations with virtualized infrastructures utilize DR capabilities along with other benefits of virtualization.

How would you describe your organization's usage of the following virtualization features/capabilities with its **production** environment-based virtual machines? (Percent of respondents, N=119)



CHAPTER 3

CHAPTER 3

continued

5

Disaster Recovery Myths and Realities

Disaster Recovery is like an insurance policy that you can test without having an accident.

MYTH 1: Disaster Recovery is a luxury feature; it's expensive and resource consuming.

REALITY: VMware vCenter™ Site Recovery Manager (SRM) gives you the flexibility to define failover scenarios that meet your choice of coverage, speed, and cost of recovery. For example, while a dedicated recovery site is a robust solution (and yes, more expensive), in many cases it is sufficient to have an active bidirectional approach where two or more data centers are complementary with enough capacity to pick up critical applications. Therefore, no resources are wasted and business continuity is maintained.

Overall, SRM customers consistently report significant savings of money, resources, and time.

How it's done at ... [Challenger Limited](#)

Challenger Limited issues annuities and provides investment products and services. The organization runs two co-located data centers, supporting around 500 staff in Australia.

In order to meet business requirements for fast recovery and minimal data loss, Challenger Limited implemented a dual-cluster VMware infrastructure that was linked to networked storage devices in its two co-located data centers, at about one-third the cost of a physi-

cal disaster recovery environment. SRM has enabled the organization to dispense with most of the 50 tapes previously used to back up data, saving one person-day per week. In addition, Challenger Limited automated hundreds of steps in its disaster recovery processes.

Business results:

- ▶ Improved RPO from 24 hours to 90 minutes and recovery time from 24 hours to less than four hours
- ▶ Reduced the number of people needed to restore systems to one person
- ▶ Cut capital investment for disaster recovery to a third of the cost of a physical environment
- ▶ Eliminated the need to acquire 15 standby physical servers at a cost of \$200,000

MYTH 2: Architecting and properly managing a DR solution is a complex task requiring special skills and expensive resources.

REALITY: Not with VMware. Physical DR can be complex because of duplicate and siloed infrastructures and configuration synchronization issues across sites. Virtualization encapsulates servers, OS, and applications, including all configuration data, so the complexity is greatly reduced. Virtualization and automation ensure that recovery plans are simple, complete, and can be reliably executed by staff with no special skills required.

With SRM, setting up an automated recovery plan is effortless and can be done in a matter of minutes, instead of the weeks required to set up manual runbooks.

How it's done at ... [Swedbank](#)

Swedbank is one of the largest financial institutions in Scandinavia and the Baltic, with 362 branches in Sweden and 222 branches in Estonia, Latvia, and Lithuania. The bank serves 9.5 million private customers and 534,000 corporate customers, with 18,000 employees.

Preventing service disruption is critical to Swedbank. Swedbank had to meet recovery targets for its legacy applications through traditional means of backup and recovery, which were complex and time-consuming. Swedbank deployed SRM to simplify and automate the recovery process, management, and testing of recovery plans. Since the SRM implementation, Swedbank tests its DR capabilities at least twice a year. It shuts down one data center completely, moving workloads to the surviving data center. It runs everything in the backup data center for 24 hours, and then fails over again to the original data center.

Mart Nael, head of Core Infrastructure, Group IT at Swedbank, states, "Our recovery time is below 30 minutes for mission-critical workloads and less than four hours for the whole data center."

Business results:

- ▶ Positive ROI within one year from hardware-cost avoidance
- ▶ IT operational costs reduced 14% year-to-year
- ▶ 1,000 virtual machines managed by two full-time-equivalent staffers
- ▶ 30x faster server provisioning

"VMware Site Recovery Manager makes managing and testing our recovery plans as easy as pushing a button."

— KENNETH NEWBALL
SENIOR DISASTER RECOVERY ADMINISTRATOR
AHS-IS

MYTH 3: After all the planning, you never know if recovery will be successful in a real disaster.

REALITY: A recovery plan is not a complete plan without testing. In fact, the recovery plan can and should be tested with sufficient failures, and retested to ensure validity. SRM enables frequent non-disruptive testing of recovery plans.

How it's done at ... [Adventist Health System](#)

Adventist Health System (AHS), a healthcare organization in the U.S., supports 37 hospitals and cares for roughly four million patients annually. AHS Information Services (AHS-IS) serves hospitals in nine states and employs more than 500 people.

To ensure AHS-IS can provide excellent care, "Mission Zero" initiative aims to provide the highest levels of service and minimum downtime for critical healthcare systems like Cerner's charting and electronic medical record applications.

Adding SRM to its VMware infrastructure allowed AHS-IS to streamline operations even further by automating DR planning and testing. "VMware SRM makes managing and testing our recovery plans as easy as pushing

6

CHAPTER 3

continued

a button. The fact that we can run tests as often as we want gives us a high degree of confidence in the recoverability of our systems,” says Kenneth Newball, senior disaster recovery administrator at AHS-IS.

Business results:

- ▶ Reduced RTO by 75%, from 48 hours to less than one hour
- ▶ Eliminated the cost of flying a team of seven people to test remote DR
- ▶ Cut hardware purchases by 84.5%, maintenance by 93.1%, and power consumption by 90%

MYTH 4: DR expense is a sunk cost, like a protection plan that’s most likely never used.

REALITY: Even if the big disaster never happens, the recovery plan can be used as a migration plan with similar steps, helping you during planned downtimes such as site migrations. In addition, DR planning helps to fulfill compliance where disaster recovery plans are required. The outcome of recovery testing proves disaster preparedness and the ability to meet RTOs.

How it’s done at ... [Ohio Department of Developmental Disabilities](#)

The Ohio Department of Developmental Disabilities (DODD) runs a statewide system of support services for some 80,000 people with developmental disabilities. A disaster causing a systemwide failure would have very real human impact.

“In addition to our 10 developmental centers, we’re also responsible for making sure providers throughout the state get the support they need to receive funding from the federal government,” says Brian Brothers, network administrator manager. “If our services were to go down and we couldn’t ensure reimbursement for Medicaid funds, it would have a severe impact on the providers and the developmentally disabled they serve.” Some providers might shut their doors.

At DODD, SRM is responsible for a reliable, verifiable DR enablement that can be tested and audited. The agency has tested its disaster recovery solution twice. The second test involved 50 production servers, which were successfully failed over to the remote site in about 90 minutes. “If we do have a true disaster someday, our DR site becomes our production site. We expect to be up and running in less than two hours,” notes Kipp Bertke, IT manager for infrastructure and operations, Ohio Department of Developmental Disabilities.

DODD’s disaster recovery site doesn’t just sit idle. Instead, the backup site actively supports the application development team on a daily basis.

Business results:

- ▶ A reliable disaster recovery site that can be up and running in less than two hours
- ▶ Fully tested, active disaster recovery solution implemented for an agile, private cloud infrastructure
- ▶ Online systems providing faster, more reliable service

CHAPTER 4

Disaster Recovery Best Practices — Top 10

As shared by VMware’s 5,000+ SRM customers

1. Virtualize. Virtual environments are much more agile and easier to migrate. Virtualization hides the complexity by shielding the individual components and moving parts, thus simplifying the planning and increasing the visibility into the DR process. It also allows you to use hypervisor-based replication that is far more flexible and cost-efficient than storage-based replication.

2. Automate. Don’t let human error stand in your way. Use automated recovery plans, not a stack of notes in a binder. With the proper automation, a recovery plan can be done in a matter of minutes instead of weeks. Automation shields users from having to manage many of the recovery steps, and automatically coordinates activities such as preconfiguring networks and virtual machines, configuring the recovery infrastructure, and restarting applications.

3. Verify and test. Test your DR plans often. Use non-disruptive testing of your recovery and failback plans. Study a detailed report of the test outcomes, including the RTO achieved. With this information, you can gain the confidence that your disaster protection plan meets the business objectives. It will also provide the necessary training to the staff and show any possible issues early so they can be addressed.

4. Set achievable goals. Automated disaster recovery can be very powerful, but it’s not magic. For example, 100 virtual machines containing Exchange, Oracle SQL, and SAP cannot be failed over and started in 30 minutes. Set your RTO realistically. To set your baseline, run a test under different conditions and see what you can achieve.

5. Act early if you can. If you have warnings, use them! Act early to execute your well-tested DR plan before an actual disaster strikes to avoid a DR event altogether. IT confidence is a byproduct of a good, solid DR plan that has been tested. Examples are a forecasted storm, a possible tsunami, or a potential network outage threat.

6. Be proactive when at risk. Most outages are not caused by actual disasters, but by planned procedures gone wrong. Examples: software or network upgrades, data maintenance, facilities repairs, etc. By proactively migrating your critical applications, you can mitigate the risk and greatly reduce a possibility of outage or service degradation.

CHAPTER 4

continued

7. Assign responsibilities. Assign everybody involved in the DR plan a specific task. Don't expect the relevant personnel to always be at the disaster site or to be in control immediately. Implement necessary duplication and redundancy for people, just like you would do with computers.

8. Keep your recovery data as current as possible. It is a good practice to prepopulate your failover site with the data that doesn't change often, or by much. This will allow you to focus only on the fast-changing critical data at the time of failover, and ultimately meet your RTO with less effort.

9. Prepare for failback. Create and test a failback recovery plan, set up replication in reverse, and know when to trigger it. Agree on what to consider the "end" of the disaster so your business can go back to normal.

10. Don't just throw money at DR. Utilize cheaper, commodity failover site assets, or use the repurposed hardware left over after your primary data center has gone virtual. Consider bidirectional or shared failover sites, use more software in the cloud (SaaS), and also look at non-IT DR means (UPS or backup generators, fuel reserves, better fire protection, etc.).



"If your organization is already taking advantage of virtualization, then adding Site Recovery Manager to handle disaster recovery is a no-brainer."

— Jerry Wilkin
Senior Systems Administrator, Dayton Superior Corp

CONCLUSION

A Quick-Start Guide to Disaster Recovery

It can be done. It must be done. VMware can help you do it.

While your data center is critical to your ability to conduct business, events beyond your control (or even planned ones) can make IT services unavailable or highly limited. This situation, however rare, could be very damaging to the integrity of your business, your market credibility, and your customers' satisfaction and loyalty.

You can mitigate this risk by implementing a DR solution to protect your critical IT assets. A well-designed DR solution built on an intelligent virtual infrastructure can provide the required RTO and RPO while keeping the costs in check. Your DR plans can be tested in a nondisruptive way and benefit your IT department in areas beyond the typical DR needs.

Your IT infrastructure plays the most critical role for the feasibility and the ultimate success of your DR plans. Virtualized infrastructure proved to be the most reliable and cost-effective platform for DR by allowing you to abstract the moving parts and components of your data center, simplifying the replication architecture and requiring fewer resources overall.

So how do you start the journey to protect your IT assets? Use this quick-start list as your guide:

1. Identify your most critical applications and data.

What applications directly generate revenue, maintain safety, or are otherwise critical to business continuity? What data is absolutely critical for your customers, your internal accounting and finances, or compliance?

2. If you have not yet, consider virtualizing your key applications. This will not only cut much of the operational and maintenance costs by removing unnecessary complexity and operational cost, but it will also make your environment better suited for effective DR planning.

3. Agree on the target RTO and RPO. What data can you lose? For how long? When do you want to be back online with your critical applications? Make sure your goals are realistic.

4. Define the triggers for DR to bring all the planned activities to action. This can be a business decision based on the data you are getting, or a technical event automatically triggering your DR.

5. Identify what DR replication, failover, and failback options you want to implement. The resulting solution will be a compromise among level of protection, speed of recovery, and costs.

6. Select the solution vendor. Beware of vendors pushing specific hardware, operating system, or other limiting choices that don't align well with your environment. Study the level of your team's expertise required to maintain the solution or the amount of resources you need to allocate. Make sure you can test the solution without waiting for an actual disaster.

And finally, good luck. We hope you will never experience a Black Swan Event and never need to use your DR solution to recover from the unexpected. But if you ever do, VMware is here to make sure you are well prepared.

For more information on VMware vCenter Site Recovery Manager, or free 60-day evaluation, please visit [VMware SRM](#).

For details and hints on SRM implementation, and thoughts on delivering data protection, high availability, business continuity, and disaster recovery with VMware, we welcome you to read our [blog](#).

APPENDIX

APPENDIX

continued

11

Disaster Recovery 101 — The Basics

DISASTER RECOVERY IS A KEY PART of a company's business continuity initiative to ensure the availability of integral IT-dependent business processes and prevent any long-term negative effects of both planned and unplanned disruptions. The goal of DR is to restore critical IT services as quickly as possible and minimize business disruption.

Nothing impacts your ability to recover more than the agility of your IT and applications infrastructure. Just like fire safety must be built into the building before the fire occurs, and a car's safety features are engineered to reduce crash impact, the design of your IT infrastructure can make or break the success of your DR program.

IT AND APPLICATIONS INFRASTRUCTURE

Your data center's infrastructure plays an instrumental role in the effectiveness of your DR solution. The infrastructure can make DR very complex, hard to implement, and sometimes even impossible; or it can help to make your IT reliable, verifiable, and effective. The next section explains how.


Two key processes for simple and reliable disaster recovery:

FAILOVER

Failover is the capability to switch over to a redundant or standby server, system, or network upon the failure or termination of an existing asset. Failover should happen without any kind of human intervention or warning.

FAILBACK


Failback is the process of restoring a system or another asset that is in a failover state back to its original state. Effective failback returns the system to the state of operation before the disruption.



Key metrics for planning and measuring success of the procedures.

RPO

Recovery Point Objective (RPO) is the point in time to which you must recover data as defined by your organization, generally called an "acceptable loss" in a disaster situation. It allows an organization to define a window of time before a disaster when data may be lost and is tightly dependent on the type of data replication used. The higher granularity of data replication, the shorter the RPO.



RTO

While RTO is purely a technical metric, the decision to trigger the failover is a business one, and RTO can often take much longer than the actual DR itself. Whether initiated by humans or by an automatic trigger, the lead time to start DR should be also accounted for and included in RTO. Replication is a key element of any DR process in most cases, usually provided by the specific DR solution used.

REPLICATION

In the context of preparing for a failover, replication provides intentionally architected redundancy of your IT resources: hardware, data, software, networks, or all of them together. There are several factors in determining the depth and amount of replication needed: type of services to be protected, criticality of different components, technology, and cost.

DISASTER RECOVERY SCENARIOS

Various DR scenarios and techniques are available to meet your specific requirements and cost objectives. The right architecture can make your DR procedures more efficient, cost-effective, and predictable. Here are a few commonly used configurations from which to choose:

▶ **Active-Passive:** This is a more traditional DR scenario, where a production site running applications is recovered at a second site that is idle until failover is required. In this scenario you are paying for a DR site that is idle most of the time.

▶ **Active-Active:** Use your DR site for non-critical workloads when you are not using it for DR. Configure it to automatically shut down or suspend the virtual environment as part of the failover process so that you can easily free up compute capacity for the workloads being recovered.

▶ **Bidirectional:** Provide bidirectional failover protection so that you can run active production workloads at both sites and fail over in either direction. The spare capacity at the other site will be used to run the virtual environments that are failed over.

▶ **Local Failover:** Some workloads need to be able to fail over within a given "site" or campus; for example, when a storage failure occurs or when maintenance forces you to move workloads to a different local lab.

▶ **Shared Recovery Sites:** In the standard one-to-one deployment, a single data center is protected by a single recovery site. You may also choose to protect multiple data centers using a "shared" recovery site. All protected sites are visible and manageable within this single instance of the DR solution at the shared recovery site. Companies that have several sites that need protection will find this feature appealing. This topology can be implemented using the shared recovery site feature.

12