WHITE PAPER: DISASTER RECOVERY

Confidence in a connected world.

Five Critical Recovery Flaws Your Last Disaster Recovery Test Missed

A look at why traditional disaster recovery testing fails to uncover serious disaster recovery gaps and what it can mean to your business

April 2009

Five Critical Recovery Flaws Your Last Disaster Recovery Test Missed

Contents

Introduction 4
The failure of disaster recovery testing 4
A closer look at what a disaster recovery test can miss 5
Replication inconsistencies (different RDF groups) 6
Missing network resources 7
Tampering risk 8
Point-in-time copies never tested 9
Insufficient dr configuration/resources 10
A few dr testing mistakes to avoid 11
The value of automated dr testing and monitoring 11
How disaster recovery advisor detects hidden dr vulnerabilities 12
Detecting optimization opportunities 12
Summary 13

Introduction

Businesses today are spending millions of dollars to develop and maintain disaster recovery (DR) infrastructures that will ensure business continuity. But despite such huge investments of time and resources, most IT professionals are still not completely confident in their ability to recover in an emergency. With industry analysts citing DR failure rates of at least 60 percent, there's good reason to be concerned.

Realists that they are, most IT managers understand that the complexity and scale of today's infrastructures, high change rates, the number of stakeholders tied to the change management process, and DR testing costs make recovery exceedingly difficult even in the best of circumstances. But the limitations of traditional DR testing are putting IT organizations at an even greater disadvantage. At a time when businesses are under more pressure than ever to ensure continuity and minimize data loss, IT organizations have no way to accurately measure if their DR plans will actually work when they need them.

This paper explores the reasons why periodic DR testing and manual auditing is not enough to ensure DR readiness. It takes a closer look at the challenges of traditional DR testing and explains how and why most tests will miss the serious data protection gaps and recovery vulnerabilities that are lurking in most environments.

In addition, the paper examines how automated DR testing and monitoring, a new approach to DR management, is helping companies around the world make up for the shortcomings of traditional DR testing. These solutions provide companies with the ability to reduce the cost and operational disruptions caused by traditional testing methods while delivering a consistent, up-to-date view of the environment. Automation enables vulnerabilities to be detected and resolved immediately to ensure the highest level of DR readiness and business continuity.

The failure of disaster recovery testing

The theory

A DR test should emulate how well business operations can be transferred to a remote facility to get the organization back online within a specified recovery time objective (RTO) and recovery point objective (RPO).

A good DR test requires considerable advance planning, along with a sizable investment in time and resources. Large numbers of people in the IT organization need to be involved. Network and storage resource mappings must be reconfigured not just once but twice, first for the test and then again to restore normal operations. And to simulate a real disaster – which is the only way to truly determine how well the DR strategy works – mission-critical applications or the whole production environment must be taken down during the test, a step which most businesses are loathe to take.

When a test doesn't work, the team must locate and fix the problems and then repeat the process.

The reality

DR tests are difficult, costly and complicated. Most companies run lean IT organizations that just don't have the time or resources to execute full, by-the-book DR tests. Plus, simulating a disaster can be dangerous: upon completion of a test, IT professionals

often hold their breath, hoping that production will be easily resumed. With such concerns and limitations, it's no wonder the scope of DR tests is minimized. Shortcuts include:

- Testing just a few key portions of the infrastructure, rather than testing the full DR environment. Companies may, for example, test very few business services and postpone the rest to a future test.
- Keeping storage/database/application management servers and/or domain/name servers or file servers online while performing the test.
- Conducting orderly system shutdowns to protect production systems, rather than simulating the abrupt cessation of operations that would occur in a disaster.
- Testing failover servers but not applications.
- Testing applications but not simulating the actual load the application must bear following a full site recovery.
- Neglecting to test dependencies, data inconsistencies and mapping errors that may exist between SAN devices and hosts, or any of the other errors that can cause a recovery to fail. This is important because most applications operate within a federated architecture that includes complex interrelationships between databases, applications, middleware, flat files and so forth. To ensure successful recovery and data consistency, a DR test should ensure that all components in the federated architecture can be recovered, or restarted, to the same point-in-time, while ensuring write-order fidelity. However, most businesses do not do this.

In the end, they have test results that are at best incomplete and at worst worthless.

Configuration changes: adding fuel to the fire

To complicate matters further, no one can test daily so a DR test only evaluates a company's recoverability at that instant in time. The moment a change is made to the infrastructure, the test results are thrown into question because there is no way to easily assess what impact that change may have on any other aspect of the environment.

Today's large datacenters are incredibly complex and often include hundreds of applications running on thousands of servers with multiple operating systems and databases. Clusters such as Veritas Cluster Server (VCS), HP ServiceGuard, Microsoft Cluster Server (MSCS), IBM HACMP, and Sun Clusters require complex resource configuration. In addition, GeoClusters store data in different geographically dispersed locations.

With such a setup, configuration changes are a fact of life. However, any small configuration change – such as the addition of a new volume/database file or reconfiguration of replication processes – can create a gap between the production and the DR environments. Even the smallest gap can cause a recovery operation to fail.

Adding to the challenge is that these errors may not impact normal operations. And, since companies are minimizing the scope of their DR test, there's a good chance the errors won't be caught when a test is run. Instead, they will remain undetected – until an emergency strikes and recovery is derailed.

A closer look at what a disaster recovery test can miss

Even when a test is conducted according to standard best practices, the number of gaps and errors it can miss is significant enough to pose a serious risk to the business. Given that most organizations have neither the time nor the resources to perform complete DR tests, the level of risk becomes downright frightening.

These risks fall into two categories:

Data protection risks

Application data, metadata and data links can be jeopardized by gaps in replication, setup, sequence of procedures, accessibility, mapping, zoning and more. Maintaining the completeness of the data and its internal structure consistency is a critical, but difficult, task. Direct impact: data loss and potential RPO violation if data is irrecoverable or recoverable but to a point in time that violates a required RPO.

Availability risks

Standby hosts, DR servers and cluster members may be unable to fulfill their role because of erroneous configuration, incorrect mapping of replicated storage to standby hosts, standby host configuration errors, and other issues. Direct impact: extended recovery time and potential RTO violation.

Hidden optimization opportunities

A side benefit of running a DR test is finding optimization opportunities, including:

Underutilized resources

Deployment gaps often result in excessive allocation of storage resources and inefficient use of Storage Area Networks (SAN) resources.

Best practices

Best-practice violations, or the inability to conform to established best practices, are frequently discovered.

Let's take a closer look at just five of the more common errors that often go undetected. We'll explore why they occur, why a DR test fails to catch them, and how they can impact operations.

Replication inconsistencies (different RDF groups)

Risk:

Data loss and increased time to recover

How does it happen?

This is a common gap found in large EMC SRDF/S and SRDF/A environments where multiple RDF groups are needed. It occurs most often when storage volumes from different RDF groups are provisioned to the host and used by the same database. The

provisioning tools do not alert or prevent this configuration. Each RDF group is associated with different replication adapters and

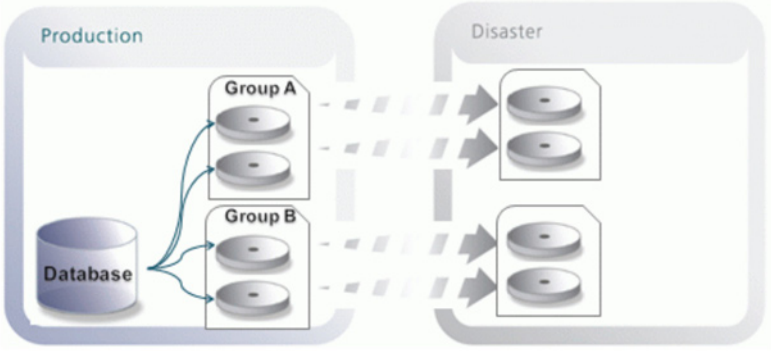

potentially different network infrastructures. Rolling disaster scenarios can result in corrupted replicas at the disaster recovery site.

Synchronous Replication Inconsistency—RDF Group
Result: data loss, increased time to recover

What is its impact?

A rolling disaster scenario is characterized by the gradual failure of hardware and network, as opposed to abrupt and immediate cessation. Most real-life disasters are rolling (for example, fire, flood, virus attacks, computer crime, etc.). In a rolling disaster, network components will not fail at exactly the same time, resulting in one RDF group being out of sync with the other RDF group. This will irreversibly corrupt, the database at the disaster recovery site. Data will need to be restored from a recent backup, increasing both the RTO and the RPO.

Why does the DR test miss this?

When a company conducts an orderly shutdown of applications, databases and hosts, it leaves data in a consistent state. Gradual/rolling disasters that bring systems or network elements down one by one are extremely difficult to emulate in a DR test.

Note: Many companies actually experience this problem but incorrectly assume it is the result of some network abnormality. However, unless the issue is properly diagnosed and corrected, it will reoccur.

Missing network resources

Risk:

Extended recovery time

How does it happen?

This risk can generally be traced to a configuration mistake which occurs when DR is not considered during the configuration process. The source host is accessing network file systems (CIFS/NFS). The network file systems are stored on a production server/array/NAS device. The target DR server also accesses the network file systems from the same production server on the production

site. During a DR test, the production file server is not brought offline and the test succeeds. During a real disaster the production server will not be available.

What is its impact?

If the network file systems were not replicated to a DR site, data loss will result. If the systems are replicated, recovery time will be extended while the administrator locates corresponding file systems on the DR site and mounts them on the DR standby server. This assumes the organization has excellent site documentation. Without it, however, data loss will occur.

Why does the DR test miss this?

When running the DR test for a specific business service or application, most companies do not shut down the entire production datacenter. The DR test will be successful because the other assets are accessible and responding. Therefore, the DR site will use the production system unknowingly.

Tampering risk

Risk:

DR failure and data corruption

How does it happen?

This hidden risk is the result of an unauthorized host at the DR site erroneously configured with access to one or more storage devices. This is a very common error, and, much to the surprise of many organizations, there are dozens of reasons why it can happen. In each case, however, it can remain dormant during normal operations and is only revealed during an actual full-blown disaster. Here are just a few reasons why this error may occur:

- When performing a storage migration, the storage administrator forgets to remove old device mappings to the host. After repurposing the old devices to new hosts, some are still visible by the original, now unauthorized host.
- From time to time, extra mapping may be added to increase performance or resiliency of access to the disk. If zoning and masking are not controlled and managed from a central point, one of the paths might actually go “astray.”
- Sometimes HBAs are replaced not because they are faulty but because greater bandwidth is required. If soft-zoning is used and is not updated accordingly, an old HBA still retains permission to access the original storage devices. Once

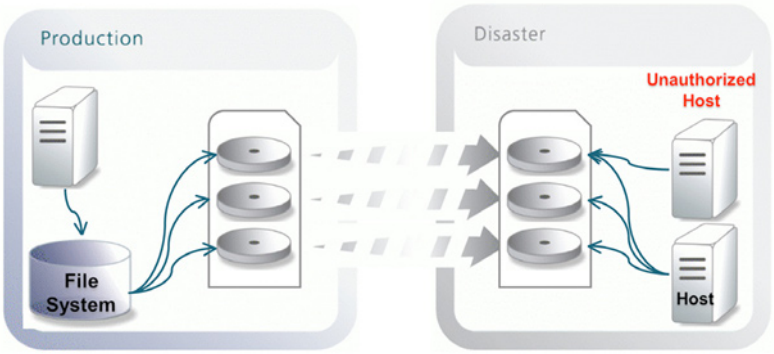

the HBA is reused on a different host (which can occur months after the upgrade) this host will actually get access rights to the SAN devices which belong to the original host.

Tampering Risk
Result: DR failure and data corruption

What is the impact?
During a disaster, a racing condition, with several unpleasant scenarios, will develop.

Scenario 1
The unauthorized host might gain exclusive access to the erroneously mapped disk. In this case, the designated standby will be unable to mount and use the locked devices, and it could take some time to isolate and fix the problem. There is also the risk of the unauthorized host actually using the erroneously mapped disk, thereby corrupting the data and rendering recovery impossible.

Scenario 2
Both the standby and the unauthorized hosts get concurrent access to the disk. If the unauthorized host attempts to use the erroneously mapped disk, not only will the data be corrupted instantly, but the now-active standby may unexpectedly crash.

Why does the DR test miss this?
Simply put, because all hosts are rarely brought up at the same time. As already explained, many organizations choose to test only one subset of the environment at a time. During a test, both the original and unauthorized server would not be started at the same time, but in a real event they would, and will wreak havoc on the data.

Point-in-time copies never tested

Risk:
Data loss and increased time to recover

How does it happen?
Point-in-time copies like snapshots and BCVs are the second line of defense to protect against human errors, viruses and outages as well. The DR configuration for applications typically includes:

- Multiple local point-in-time copies such as EMC TimeFinder, HDS ShadowImage/Snapshot, NetApp FlexClone/Snapshot, or CLARiiON SnapView;
- Remote synchronous replication such as EMC SRDF, Hitachi TrueCopy, CLARiiON MirrorView, and NetApp SnapMirror;
- Local point-in-time copies on the remote site.

In addition, the copies could be mapped to the target DR servers, configured with multi-path software such as EMC PowerPath, Veritas DMP and MPIO, and defined in logical volumes such as Veritas VxVM.

Point-in-time copies can easily become corrupt, without being discovered, unless the application is fully started and the data

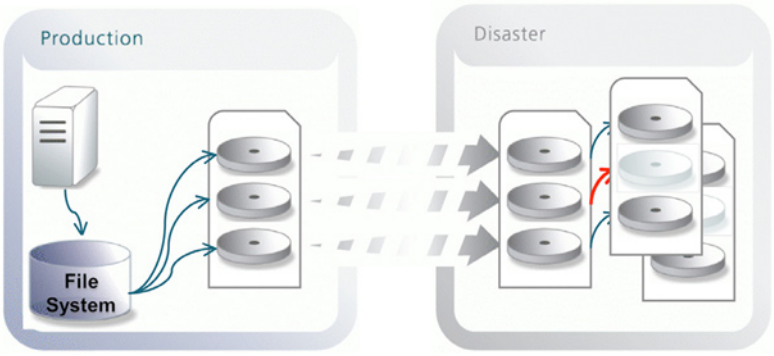

integrity is thoroughly tested. The diagram below illustrates a file system replica corruption caused by not ensuring the same data-age for all devices comprising the replica.

Local Replication with Remote Replication
Replication Age Inconsistency
Result: Data corruption

There are numerous scenarios that can lead to such a corruption, such as when the replica devices do not all belong to the same consistency group.

What is the impact?

The replica is corrupt and unusable. The file system will need to be recreated at the disaster recovery site and data restored from a recent backup, thereby increasing the time to recovery. All data created since the last backup will be lost.

Corrupted file systems may still be usable in many cases, and only a close inspection of the content can reveal the fact that the data is meaningless.

Why does the DR test miss this?

This gap can be missed if the specific business service is not tested for DR or if the DR test only includes turning on the DR server without actually running the applications.

Insufficient DR configuration/resources

Risk:

Extended recovery times, Recovery Time Objective violation

How does it happen?

DR and production infrastructures are usually not the same. When building a DR data center, organizations tend to assign fewer resources than their production environments have. If the DR configuration includes significantly fewer resources than production, there is a good chance it will be unable to assume production properly upon failover. It is not unusual, for example, to find a

production environment that has multiple paths to storage or software, but the DR environment has too few, or even just a single path. It is also common to find DR sites with misconfigured kernel parameters or insufficient memory or CPU to support full production load.

What is the impact?

When the DR site cannot assume production as planned, business operations cannot resume in accordance with the company's established SLA. In the best case scenario, IT must devote additional resources to execute the unplanned configuration of servers and storage. In the worst case scenario, the company will need to incur additional unplanned capital expenses.

Why does the DR test miss it?

Most DR tests do not simulate full production load, so these errors remain undetected. Since DR is mostly offline, this issue never comes to life until an emergency occurs.

A few DR testing mistakes to avoid

The list of potential configuration errors is huge. Veritas CommandCentral Disaster Recovery Advisor (DRA) maintains a comprehensive database of thousands of data protection and DR risk signatures and adds new ones every day. We've highlighted a few of the more common gaps in the preceding pages, but here are some additional risks to be aware of:

Keeping ECC server available in production in DR tests

Storage management tools such as ECC are the main tools used by system and storage administrators to understand and configure the relationship between servers and storage devices. It is common practice not to map all replica devices to the DR servers during normal operations. However, if a valid and current DR ECC environment is not maintained, there may be no easy way to tell how to map thousands of unmapped devices to the appropriate DR servers. In the stress and confusion that accompanies disaster events, this may lead to significantly extended recovery time.

Starting, but not using, applications

Often a company will just confirm the application started or, at best, run one or two transactions before returning to production. The danger with this shortcut is that real system usage is not simulated, so it is impossible to determine if there are underlying problems – database dependencies, for example, or the ability to support the true production load – that could have an impact in a real failover event.

The value of automated DR testing and monitoring

Clearly, periodic DR tests and manual audits will always be an important part of any DR strategy. They can uncover important flaws in processes, procedures or technology that could impact readiness. However, it is foolhardy to ignore the serious drawbacks of traditional testing that are leaving critical applications and data unprotected.

That's why many organizations are looking for a way to augment their DR testing methods with new solutions that can automatically check for vulnerabilities and identify problems before they impact business operations. This approach provides more robust protection between DR tests and ensures a higher level of readiness by expanding the overall effectiveness of periodic testing.

Automated DR monitoring technology is able to penetrate deeper into the environment to ensure the infrastructure status is always aligned with the protection goals. For instance, Disaster Recovery Advisor can quickly analyze dependencies between IT assets and the business services they support because it maintains the most comprehensive and constantly updated documentation of IT resources and dependencies in the production and DR sites.
DR management software can perform tasks that are simply too cumbersome or complex for humans to perform, such as assessing the accuracy of intricate mappings, making sure all replicas exist and are consistent, and identifying RTO/RPO violations.

Technology like Disaster Recovery Advisor can also provide incredible visibility into IT's ability to meet its Service Level Agreements. Using Disaster Recovery Advisor, the user can see the age of the most updated copies per server, both local and remote, thus revealing potential RPO violations. Data retention and number of copies can also be measured through Disaster Recovery Advisor, as it also provides visibility into the oldest replicas. The dashboard provides a high level view of replication and recoverability status.

How disaster recovery advisor detects hidden DR vulnerabilities

Using its powerful detection and analysis tools, Disaster Recovery Advisor scans storage, databases, servers and replication configurations for vulnerabilities such as unprotected databases or database partitions, noncompliant replication configurations, data that cannot be recovered to a valid consistency point, and much more.

The Disaster Recovery Advisor Data Collection Engine automatically collects configuration data from key IT assets, including storage management frameworks, servers and databases. This information is used to correlate the applications to the underlying infrastructure. The discovery and scanning process can be set to perform periodic rescans to detect potential configuration problems when they occur so that they can be addressed before they impact production availability. This allows Disaster Recovery Advisor to detect and assess changes over time and allows for ongoing data protection monitoring.

Once the data is collected, Disaster Recovery Advisor performs a comprehensive dependency analysis and builds a detailed disaster recovery topology map, which illustrates the dependencies of the applications through the infrastructure, including those at multiple sites. This topology serves as the foundation for the analysis by Disaster Recovery Advisor's Gap Detection Engine.

The Gap Detection Engine uses a gap signature knowledgebase of thousands of potential data protection gaps to automatically detect potential gaps and best-practice violations in your DR configuration. This is the equivalent of performing millions to tens of millions of manual comparisons.

When a gap match is identified, Disaster Recovery Advisor issues a ticket with a detailed description of the risk, its impact and a suggested remediation approach. This allows administrators to be proactive in minimizing risk to the organization by resolving issues before they escalate into major problems that require a protracted effort to remove.

Detecting optimization opportunities

In addition to uncovering recoverability risks, Disaster Recovery Advisor uses the configuration information it collects to identify optimization or other fine-tuning opportunities within the infrastructure. These improvements can include orphan storage elements and replication configuration issues that are impacting bandwidth utilization or I/O performance.

Summary

The need to ensure business continuity has never been more important, and enterprises are making significant investments to design and build solid disaster recovery systems. But the real challenge is not in building the environment – it is ensuring that it will be constantly ready to resume operation in the event of a disaster.

While traditional DR testing can provide the IT organization with valuable insights, it cannot ensure recoverability because it cannot detect many of the configuration gaps that can derail a DR effort. Only automated DR testing and monitoring solutions, like Disaster Recovery Advisor, can enable true DR readiness.

About Symantec

Symantec is a global leader in infrastructure software, enabling businesses and consumers to have confidence in a connected world. The company helps customers protect their infrastructure, information, and interactions by delivering software and services that address risks to security, availability, compliance, and performance. Headquartered in Cupertino, Calif., Symantec has operations in 40 countries. More information is available at www.symantec.com.

For specific country offices and contact numbers, please visit our Web site. For product information in the U.S., call toll-free 1 (800) 745 6054.

Symantec Corporation
World Headquarters
20330 Stevens Creek Boulevard
Cupertino, CA 95014 USA
+1 (408) 517 8000
1 (800) 721 3934
www.symantec.com

04/09 XXXXXXXX