



Consolidated Disaster Recovery Using Virtualization
Affordable Workload Protection and Recovery



Consolidated Disaster Recovery Using Virtualization Affordable Workload Protection and Recovery

“Consolidation was the primary driver that fueled the first wave of server virtualization adoption, and affordable resiliency will fuel the next wave. Virtualization has lowered the cost of providing resiliency to a low enough point that firms are all but obliged to consider deploying virtualization to support a much broader set of applications than they might have in the past.”

— Forrester Research Inc.,
“X86 Server Virtualization for High Availability and Disaster Recovery”, by Stephanie Balaouras and Christopher Voce, October 24, 2007

“The IT industry is clearly moving toward an environment in which application workloads are dynamically relocated from one server to another to meet service-level agreements, address uptime requirements, more effectively use underlying server hardware or some combination of all three. Server virtualization is a technology that accommodates this type of dynamic workload reassignment.”

—Gartner, Inc.,
“I/O Is the New Frontier of x86 Virtualization”, by John Enck, May 21, 2007

Executive Summary

Traditional disaster recovery infrastructures including tape backup, image capture, high-end replication and hardware clustering have failed to keep pace with business requirements for recovery speed and integrity at a reasonable cost. Budgetary constraints and the high cost and complexity of established recovery solutions mean that most organizations can afford to protect only a fraction of their total server infrastructure – typically only their most business-critical server workloads. This common protection scenario leaves the majority of the server network under-insured in the event of downtime or disaster.

While organizations can easily justify the expense of protecting business-critical server workloads such as customer-facing applications (e.g. web servers and online order processing), it is harder to find sufficient funds to protect workloads deemed less critical such as email servers, internal Web servers or batch reporting applications. Protection plans for these types of workloads, which constitute the majority of an organization’s infrastructure, might be described as “best effort.”

There is a fine line between downtime being merely a minor inconvenience to internal users and resulting in lost opportunities. When a system – any system – is down, business and employee productivity suffers. An internal web server failure, for instance, may impede the ability of a technology company to prepare sales proposals using existing documentation stored on an intranet. If repeated outages prevent employees from accessing corporate systems or completing tasks, the long-term negative effects can be significant. If a server workload is worth running in the first place, it is also worth protecting.

Charged with the task of extending disaster recovery capabilities to cover a broader spectrum of server workloads in the enterprise, IT departments are beginning to explore new disaster recovery alternatives. An emerging trend is for organizations to leverage server virtualization to achieve disaster recovery capabilities. Once confined to use primarily in software development, test and server consolidation scenarios, server virtualization and supporting technologies can afford significant cost and performance advantages over conventional recovery options.

Tape backups can be slow to recover while data replication requires an identically configured standby physical server with duplicate hardware and software. In contrast to these data-centric approaches, consolidated recovery allows organizations to replicate whole workloads (data, application and operating systems) to a warm standby virtualized environment and rapidly recover a workload in the event of a production outage simply by powering up a virtual machine. Multiple physical and virtual workloads can be consolidated onto a single recovery server or purpose-built appliance, allowing organizations to avoid the high cost of duplicate hardware and software.

This white paper discusses why a growing number of organizations are now leveraging virtual machines to protect the servers in their network that have commonly been left under-protected. It will consider how server virtualization, together with workload portability technologies, enables organizations to implement a disaster recovery plan that is more affordable and flexible than traditional options, while providing rapid restore times and enterprise-level workload protection.



Disaster Recovery by the Numbers

In the best of all possible worlds, organizations would implement backup and recovery processes for all server workloads regardless of their perceived criticality. Moreover, these recovery processes would be fast, have minimal impact on production operations and would be recoverable with a high level of data integrity. In reality, disaster recovery needs must always be weighed against the fiscal need for cost-effectiveness.

IT organizations, perennially tasked with doing more with less, often must sacrifice recovery performance due to budgetary constraints. More often than not, economic and technological factors conspire to force organizations to over-insure their business-critical workloads, while under-insuring the lion's share of their server infrastructure.

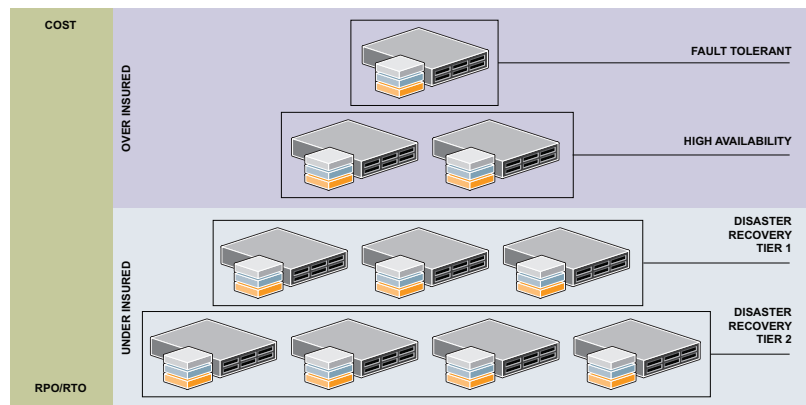


Diagram: Business continuity budgets over-insure top-tier servers while under-insuring the majority of the server network.

The 80/20 Problem

Typically, organizations allocate as much as 80 percent of their disaster recovery budget to safeguard only their most business-critical servers – usually as little as 20 percent of the total server network. A recovery budget allocated along these lines leaves the remaining 80 percent of servers under-protected should the business experience a server failure or catastrophic site disaster. While the loss of any one of these non-critical servers, which comprise the majority of an organization's IT enterprise, may not bring the business grinding to a halt, the loss would nonetheless impact business and employee productivity and cost the organization time and money. Most organizations would admit that an enterprise-wide disaster recovery plan that adheres to the 80/20 workload protection ratio is deficient. So why do so many businesses spend so much to ensure that only the most critical servers are protected by a disaster recovery solution?

The One-to-One Problem

The short answer to the question posed above is money. Traditional disaster recovery solutions such as server clustering or high-end data replication can be extremely expensive – so much so that they can only be implemented on a limited basis by large enterprises (to protect top-tier servers) and are completely cost-prohibitive for most small and medium-sized businesses.

"It's rare for a firm to build out a new data center that would act as an idle standby for a production data center in case something bad happens at the primary site. According to the Uptime Institute, a new 20,000-square-foot (1,800-square meter) tier three data center costs more than \$44 million."

— Forrester Research Inc., "Maximizing Data Center Investments for Disaster Recovery and Business Resiliency", by Stephanie Balaouras and Galen Schreck, October 5, 2007



“In DR, one of the biggest challenges is maintaining identical server configurations across both data centers.”

— Forrester Research Inc.,
“Maximizing Data Center Investments for Disaster Recovery and Business Resiliency”, by Stephanie Balaouras and Galen Schreck, October 5, 2007

“According to Forrester’s Enterprise and SMB Hardware Survey, North America and Europe, Q3 2007, 40% of enterprises never conduct a full test of their disaster recovery plan. This is due in part to the fact that DR testing is a very manual, complex, and potentially disruptive process.”

— Forrester Research Inc.,
“X86 Server Virtualization for High Availability and Disaster Recovery”, by Stephanie Balaouras and Christopher Voce, October 24, 2007

“The testing of disaster recovery solutions is very important and should be done at least yearly and when significant processes change.”

— Gartner, Inc.,
“Server Capacity on Demand Spans Many Capabilities”, by John R. Phelps, February 13, 2007

A factor that contributes greatly to the high cost of server clustering and data replication is that these solutions tend to require one-to-one hardware and software redundancy to protect data center assets, resulting in expensive hardware and a higher total cost of ownership. The one-to-one relationship means that organizations must maintain exactly the same server configuration at the recovery site as at the production site with precisely the same operating system versions, application licenses and patch levels installed. Essentially, a mirror-image of the production data center must be kept standing in reserve as a warm standby environment in the event of a downtime event.

Given the high cost of maintaining a mirrored environment, and the onerous task of keeping the configuration of the two environments perfectly in sync, it is not hard to see why most organizations opt to protect only their top-tier servers and trust the rest of their server network to lower-cost alternatives such as tape backups.

Traditional Disaster Recovery Infrastructures

This section provides an overview of traditional recovery infrastructures and how organizations can use established metrics to evaluate the various technology options in the recovery continuum.

Recovery Metrics

The two most common metrics used to evaluate disaster recovery solutions are recovery time objective (RTO) and recovery point objective (RPO), but a third metric, test time objective (TTO), is also emerging as another key measure of the effectiveness of recovery alternatives.

Recovery Time Objective (RTO)

Recovery time objective measures the amount of time a computer system or application can stop functioning before it is considered intolerable to the organization. RTO can be determined to be from seconds to days, depending on how critical a given workload is to the organization. RTO is used to determine the type of backup and disaster recovery plans and processes that should be implemented to protect a specific workload.

Recovery Point Objective (RPO)

Recovery point objective describes a point in time to which data must be restored in order to be acceptable to the owner(s) of the processes supported by that data. This is often thought of as the time between the last available backup and the time a disruption could potentially occur. Representing the nearest historical point in time to which a workload can be recovered, RPO is a measure of data loss. The RPO is established for a given workload based on the organization’s degree of tolerance for loss of data or manual rekeying of data.

Test Time Objective (TTO)

Test time objective (TTO) measures the time and effort required to test a disaster recovery plan to ensure its effectiveness. For an organization to be confident in their recovery strategy, solutions and methodology, the recovery infrastructure must be thoroughly and regularly tested. Routine testing should be relatively easy, quick to implement and non-disruptive to business operations.

RTO	Recovery Time Objective	The Measure of Downtime
RPO	Recovery Point Objective	The Measure of Data Loss
TTO	Test Time Objective	The Measure of Testing Ease

Diagram: Disaster recovery metrics.



Total Cost of Ownership (TCO)

Although not strictly a recovery metric, total cost of ownership is another factor that must be weighed when selecting a recovery infrastructure. Some the questions organizations need to ask when selecting recovery technologies include:

- How easy to administer is the disaster recovery solution and how many IT staff are required to maintain the recovery infrastructure?
- What specialized IT knowledge is required to support and maintain the recovery infrastructure?
- Is the recovery infrastructure scalable to accommodate larger workloads and future growth?
- How flexible is the solution in terms of replication, recovery and restore processes?
- Does the recovery infrastructure support multiplatform data center environments that may exist today or result in the future as a result of mergers and acquisitions?

Together, RTO, RPO, TTO and TCO form the basis on which an organization's workload protection and recovery strategy can be developed.

Weighing the Options

To protect server workloads, organizations have a number of different traditional disaster planning and recovery solutions at their disposal. Conventional recovery approaches include tape backup, image capture, high-end replication and server clustering. This section discusses how these solutions stack up in terms of cost, RTO, RPO and TTO.

"Given the total cost of downtime – which includes lost revenue, lost worker productivity, and lost market share – older approaches to disaster recovery, such as cold site recovery or even recovery from a shared IT infrastructure, are no longer adequate. Typically, a shared site infrastructure can't support recovery time objectives of less than 24 hours because data and system restores must be done from tape – and first the tapes must be found and then shipped to the site. Increasingly, organizations prefer a stand-by, dedicated IT infrastructure (servers, storage and network) that mirrors their production IT configuration and is ready to take over production processing at any time..."

— Forrester Research Inc.,
"Maximizing Data Center Investments for Disaster Recovery and Business Resiliency", by Stephanie Balaouras and Galen Schreck, October 5, 2007

Tape Backup

Tape backup is the workhorse of most disaster recovery plans. It refers to the process of using external tape drives and magnetic tape for storing duplicate copies of hard disk files. Server and desktop-based files are typically copied to the tapes using an automated backup utility that updates on a periodic schedule. Many companies use magnetic tape in combination with additional magnetic disks and optical disks in a backup management program that automatically moves data from one storage medium to another. Tape archives are usually stored offsite for recovery purposes and the regular pickup and storage of backup tapes may be managed by a third-party provider. Tape backup is the most economically prudent recovery alternative, however backup utilities and processes can be difficult to administer, as can the logistics of transporting, storing and retrieving tape archives in the event of an outage. It can frequently take days to restore a system from tape, a process that requires the manual rebuilding of systems (the reinstallation of operating systems, applications and patch levels) before the application data can be restored.

Image Capture

Image capture involves the scheduled conversion of a server workload to an image archive that can then be replicated to a remote location for disaster recovery. Many data centers maintain extensive libraries of backup image archives that they attempt to restore to new hardware in the event of a primary workload failure or disaster. Image capture is moderately more expensive than tape backup and maintains an adequate RPO, but RTO can be lengthy and error prone because images tend to be tied to the hardware from which they were originally captured and cannot easily be recovered to another server. A common problem is that, when a workload running on an older hardware configuration fails, the data center has no additional platforms of that server make and model to which they can restore the backup image. More flexible image-based solutions do exist that enable data centers to capture any image type and restore it to any hardware platform, reducing recovery time and minimizing the types of errors and delays mentioned above.



High-end Replication

There are a number of different replication methods. In the realm of database management, replication refers to the ability to keep distributed databases synchronized by routinely copying the entire database or subsets of the database to other servers in the network. Primary site replication maintains the master copy of the data in one site and sends read-only copies to the other sites. In a workflow environment, the master copy can move from one site to another. This is called "shared replication" or "transferred ownership replication." In symmetric replication, also called "update-anywhere" or "peer-to-peer replication," each site can receive updates, and all other sites are then updated. Failover replication, or hot backup, maintains an up-to-date copy of the data at a different site for backup. Although replication meets stringent recovery time and point objectives, the technology can be costly, complex and difficult to administer. This data-centric approach also requires one-to-one hardware and application redundancy.

Server Clustering

Server clustering generally refers to multiple servers that are linked together in order to handle variable workloads or to provide continued operation in the event that one server or node in the cluster fails. A cluster of servers provides fault tolerance and/or load balancing. If one server fails, one or more additional servers are still available. Load balancing distributes the workload over multiple systems. Clustering fully achieves recovery time and point objectives but at a very high cost. Because it can be prohibitively expensive and complicated to implement and maintain, clustering is typically a viable recovery option for only the most mission-critical server environments.

"The recovery infrastructure has rarely kept up with the current needs of the business. Although organizations continued to roll out additional servers, deploy new applications and provision additional storage, the backup/recovery environment typically lagged behind. Often, little or nothing was done to improve the backup process, outside of purchasing new tape volumes and occasionally additional or faster tape drives."

– Gartner, Inc.

"Recovery Will Move to Disk-Based, Manager of Mangers Approach by 2011," by Dave Russell, February 20, 2007

Solution	RPO	RTO	Cost	Weaknesses
Tape/manual rebuild	24 hr+	days	\$	- Difficult to administer - Slow, prone to errors
Image capture	24 hr	hours	\$\$\$	- Limited restore flexibility
High-end replication	minutes	minutes	\$\$\$\$	- Complicated configuration - Duplicate hardware
Server clustering	0	0	\$\$\$\$\$	- Duplicate hardware - Complicated setup

Diagram: Evaluating Traditional Disaster Recovery Alternatives

Demand for Newer Protection Options

Frustrated by the high cost and complexity of some of the traditional recovery options above and the lackluster performance of others, IT organizations have begun to seek out newer protection options that offer a better balance between cost and performance. There is growing interest in server virtualization as a disaster recovery technology platform. In the next section, we'll see how the same features that have made virtualization indispensable for accelerating server consolidation and other data center initiatives – namely workload encapsulation and workload portability – are bringing greater flexibility, cost-effectiveness and simplicity to disaster recovery.



How Virtualization is Transforming Disaster Recovery

By implementing virtualization technology, a single physical server can be configured to run multiple virtual machines in which each instance of the operating system runs its own applications as if it were the only OS on the server. The virtualization is accomplished by a layer of software called a “virtual machine monitor” (VMM) or “hypervisor” that resides between the hardware and the “guest” operating systems.

Virtualization first saw widespread use in development and test lab scenarios, which enabled organizations to rapidly provision different virtual operating environments in which to test new software prior to production. More recently, organizations have turned to virtualization for server consolidation as a means to reduce hardware expenditures and decrease the per-server cost of power, floor space and human resources. Virtualization allows increased infrastructure resource utilization and decreased hardware and maintenance costs, as well as accelerated server installation and configuration.

The rapid adoption of virtualization technologies has fundamentally changed the way organizations view the data center. By helping to dissolve the bonds between software and hardware, virtualization has encouraged organizations to see the data center not as a heterogeneous mix of different servers, operating systems, applications and data, but as a set of portable workload units. A workload encapsulates the data, applications and operating systems that reside on a physical or virtual host. The ability to profile, move, copy, protect and replicate entire server workloads as aggregated units between physical and virtual hosts is helping many organizations achieve new operational efficiencies and cost savings in the data center, and opening up new options for disaster recovery and consolidated workload protection.

Workload Portability

Virtualization allows workloads to be moved between similar virtual hosts. Workload Portability technologies make it possible to detach workloads from their native hardware configurations and move the entire software stack of a server to any physical or virtual host.

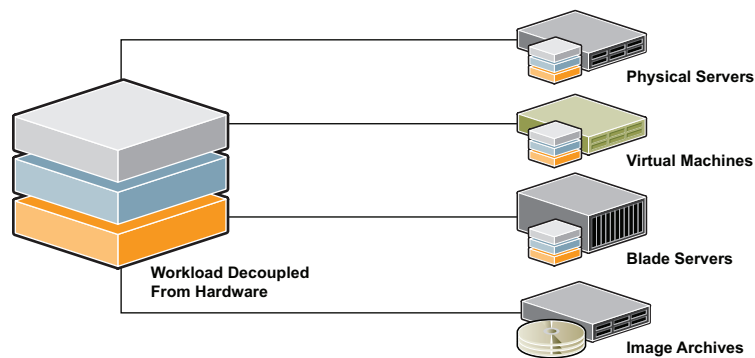


Diagram: Workload portability enables workloads to be migrated across different data center infrastructures.

This increase in the portability of server workloads empowers organizations to move and rebalance workloads in any direction between physical and virtual hosts – physical-to-virtual, virtual-to-physical, physical-to-physical, in and out of imaging formats and so on. Workload portability increases the flexibility, agility and overall efficiency of data centers. It also enables organizations to better address common challenges such as end-of-lease hardware migration, server consolidation, and more recently, disaster recovery.



Workload Protection

Recognizing the inherent versatility of the technology, organizations are beginning to extend their use of server virtualization to new areas of data center operations and harness the benefits of virtual infrastructures to more easily and cost-effectively protect server workloads running in physical or virtual environments. Workload protection is the task of copying or replicating physical or virtual server workloads to a secondary location for use as a warm standby environment in the event of a primary server outage or site-wide disaster. Virtualization enables organizations to achieve workload protection by creating a bootable archive of the workload on a virtual recovery platform where it can be rapidly recovered. As we will see in the next section, virtualization affords significant cost and performance advantages over more traditional disaster recovery options such as tape backup, imaging, replication or clustering.

Protecting Workloads with Consolidated Recovery

In a typical workload protection scenario using virtualization, organizations replicate multiple workloads to a single warm standby virtual recovery environment. Workloads are replicated remotely over a wide area network (WAN) to a geographically-dispersed recovery site. To automate the scheduled replication of workloads between primary and recovery sites, organizations use a workload portability solution for physical-to-virtual and virtual-to-virtual workload movement.

In contrast to data-centric recovery approaches that require system rebuilding, this scenario provides whole workload protection of physical or virtual workloads. The ability to replicate the entire workload (data, applications and operating systems), rather than just application data, to a virtual machine environment means that everything needed to rapidly recover in the event of an outage is available in a bootable virtual machine on the recovery server.




Diagram: Whole workload protection of physical machines using a virtual recovery environment.

Since multiple workloads, whether physical or virtual, can be consolidated onto a single virtualization-equipped recovery server, there is no need for costly investments in duplicate hardware and software for one-to-one redundancy. Virtualization allows organizations to utilize virtual host capacity as an affordable consolidated recovery platform to protect a greater percentage of their workloads. The consolidated approach to workload protection offers a significant cost advantage over one-to-one solutions such as high-end replication or clustering.

When evaluating workload portability solutions, organizations should look for live workload replication capabilities or the ability to transfer a workload into an off-line virtual machine without taking the production workload offline. Live replication reduces backup window disruptions and ensures business continuity.



“Many companies are successfully using server virtualization to improve the availability of applications at the primary site as well as facilitate a rapid restart of applications at the recovery site. When servers running virtual machines are connected to network storage, you can use array-based replication technologies to replicate virtual machine configuration files from the primary data center to the alternative site automatically. This reduces the need to fastidiously monitor system configurations between data centers. The system configurations are always in sync because any changes are continuously replicated to the alternative site. When it comes to recovery at the alternate site, it’s less “recovery” and more “restart” because the data and system images are already there.”

— Forrester Research Inc., “Maximizing Data Center Investments for Disaster Recovery and Business Resiliency”, by Stephanie Balaouras and Galen Schreck, October 5, 2007

Rapid Recovery

In the event of a primary server outage or site disaster, the virtualized recovery server can be activated to take over the running of the workload immediately. The workload can be kept running in the virtual recovery environment until the primary server has been restored, at which point the workload can be moved back to the restored server via virtual-to-physical workload replication. Recovery time and point objectives are achieved without the need for high-cost, complex clustering environments.

When planning a virtualized recovery environment and sizing the physical server that will host the recovery environment, organizations must ensure sufficient compute power and capacity to run the recovery workloads as normal for as long as it takes to restore the production server. An insufficiently sized and provisioned physical host will mean that the workloads may have to run in a degraded state, impacting business operations.

A workload profiling and analysis solution can be indispensable at the recovery planning stage for accurately sizing the required virtual recovery environment. Such a tool can be used to monitor workloads and resource utilization trends over a period of time to account for peaks and valleys in resource utilization and ensure that the recovery environment is properly balanced to handle current and future workload requirements on a production scale.



Diagram: A workload profile captures information that can prove invaluable for planning and right-sizing the recovery environment

Ease of Testing

For an organization to be confident in their disaster recovery strategy, solution and methodology, they must be thoroughly and regularly tested. However, as we saw above, conventional recovery infrastructures are prohibitively complex and invasive to business operations, making them infeasible to test on a regular basis. The first time many organizations’ recovery solutions are ever initiated at production scale is when downtime strikes. Despite having made significant investments in disaster recovery, these organizations have no way of knowing for certain how quickly or successfully their server workloads can be restored. This “cross-our-fingers-and-pray-it-works” approach to disaster preparedness isn’t really preparedness at all.

In addition to being quicker to restore than more costly solutions, virtualized recovery solutions are also much easier to test. The inherent testing capabilities afforded by virtualization make test time objectives (TTO) a meaningful recovery metric. The rapid test restore capabilities unique to virtualization allow organizations to rapidly and easily run test failure scenarios to ensure the integrity of their workload protection plan. Moreover, routine testing can be completed with absolutely no disruption to business operations.



With virtualized recovery, users can run recovery “fire drills” to test recovery plans and assess actual versus target RTO and RPO simply by booting up a replicated copy of a production workload on the virtual recovery server. Because the test workload snapshot is fenced off from the production network, testing can be performed with no concerns about the integrity of the production environment. The speed and ease with which recovery plans and processes can be tested using a virtualized recovery environment bring auditable testing to recovery procedures and peace-of-mind that the solution will work in the event of an actual production failure or disaster.

Virtualized Recovery Extends DR Capabilities

Given the relative affordability for a consolidated recovery solution that leverages virtualization, organizations can easily justify using virtualization to augment and extend their current disaster recovery capabilities to a broader range of workloads. Virtualized recovery enables organizations of all sizes to protect a greater share of their server infrastructure with minimal capital investment and fewer barriers to implementation.

Different types of workloads require different levels of resiliency and different disaster recovery technologies. Analysts recommend a multi-tier disaster recovery approach that combines different availability and recovery technologies that overlap like shingles on a roof. Organization should opt for a range of recovery capabilities to address applications ranging from low to high criticality. Virtualized recovery adds another viable option to an organization’s recovery arsenal and enables more thorough protection for workloads previously deemed too non-critical to warrant the expense of a traditional disaster recovery solution.

Planning and Implementing a Virtualized Recovery Solution

In this section, we will walk through the steps involved in planning and successfully implementing a virtualized recovery solution for consolidating and protecting production workloads.

Planning

The implementation of a successful virtualized recovery solution begins with the creation of a detailed plan. Organizations must identify the workloads they wish to protect. They must also monitor workloads over time to ensure that the disaster recovery environment they implement has adequate capacity to support current and future requirements as workloads and resource utilization can grow significantly over time. The planning phase includes the following steps:

1. Discover Server Inventory

Discover and inventory the IT environment including physical and virtual servers and the data, applications and operating systems installed on them in order to identify the workloads you wish to include in your recovery plan.

2. Monitor Server Utilization

Monitor workload information such as CPU, disk, memory and network utilization rates over a period of time. The information collected will provide invaluable workload profiling and capacity planning data that can be used as the basis for generating a recovery plan.

3. Build the Disaster Recovery Plan

Create recovery plans to determine the most appropriate virtual recovery site capacity. Build enough headroom into your target virtualized recovery environment to ensure sufficient capacity for consolidated workloads to run with sufficient resources in the event of a disaster.

4. Configure the Virtual Recovery Environment

Match physical production servers with virtual recovery machines and configure your virtual recovery environment.



Implementation

Once the recovery plan has been developed, you will need to perform an initial full replication of physical workloads into the virtual recovery environment. Using a workload portability solution, a single, live physical-to-virtual migration is performed to copy the entire workload and stream it over the network into a virtual machine warm standby environment. After the initial full replication, you will schedule incremental workload replications to maintain synchronicity between the production server and the virtual recovery environment. The individual RPO you determined for each protected workload will determine the interval between scheduled workload synchronizations. Any file that has changed since the previous incremental transfer will be copied over to the virtual restore system on the next synchronization.

Once the synchronization schedule is determined, you can test the integrity of your disaster recovery plan by powering on the virtual recovery machine within an isolated, internal network to ensure the environment is intact.

In the event that a failure or disaster occurs, operations can be transferred from the failed production server to the virtual recovery environment. After this transfer, workloads will run as normal off the virtualized recovery server. This scenario offers several options for restoring workloads. If the original production server failure is repaired and the hardware intact, users can move the workload from the virtual recovery environment back to the original platform by using a workload portability solution to perform a virtual-to-physical workload transfer. If the original hardware cannot be repaired, users can restore the workload with a V2P transfer onto new, dissimilar hardware.

To recap, the steps involved in implementing the virtualized recovery solution are as follows:

1. Initial System Backup

Perform an automated full system backup by transferring whole server workloads (data, applications and operating systems) to the target virtual recovery environment using a workload portability solution.

2. Ongoing Incremental Backups

Automatically propagate all source changes at user-defined increments to the target virtual recovery environment to ensure that your recovery server contains an exact and up-to-date copy of your production environment.

3. Run Fire Drills to Test DR Readiness

Perform a disaster recovery "fire drill" to check application integrity and recovery time. Run a "Test Restore" on the backup virtual machine to create a snapshot of the virtual disk file associated with the virtual machine. While in Test Restore mode, incremental jobs are suspended and will resume upon the next scheduled incremental transfer or once you shut down the test virtual machine.

4. Initiate One-Click Failover in the Event of System Outage

Should your production server fail, you can rapidly initiate a system failover in which the virtual recovery machine will rapidly start up – just reconnect sessions and the virtual recovery environment takes over the production workload.

5. Rapidly Restore Systems

Once your production system is repaired, you can easily perform a Virtual-to-Physical workload transfer to restore workloads to their original host server or to a new server. A workload portability solution enables you to move data, applications and operating systems from the virtual recovery server to any physical hardware.



A Word on Next Generation Recovery Appliances

With growing demand for virtualized recovery solutions to augment and extend resiliency and recovery capabilities, solution providers are beginning to bring to market innovative packaged recovery offerings in an effort to help customers streamline implementation, configuration and ongoing administration. One of the most compelling of these next-generation recovery options is a purpose-built consolidated recovery appliance that provides significantly lower total cost of ownership by pre-packaging all the hardware, software, storage and virtualization components required to implement a virtualized recovery environment to protect server workloads.

Essentially a complete recovery environment in a box, a virtualized recovery appliance offers many benefits to customers:

- Simplified and accelerated implementation and setup
- Packaged, balanced and right-sized configuration based on the size and criticality of workloads
- Simple, integrated management of the recovery environment as a single entity
- Guaranteed performance
- Fewer IT resources required to manage and administer the solution
- Lower total cost of acquisition and ownership
- Fewer barriers to entry, particularly for SMBs or departmental use within larger enterprises

In the near future, we believe the market will increasingly move toward an appliance-based model for workload protection. A “plug in and protect” appliance model offers a dramatically simplified approach to disaster recovery, especially when compared to conventional disaster recovery infrastructures such as replication and clustering. Purpose-built for protecting server-based consolidated workloads, a pre-packaged, pre-configured and right-sized recovery hardware appliance would reduce or alleviate many of the planning and implementation challenges associated with planning, implementing and maintaining a recovery environment.

Summing Up

After years of stagnation in disaster recovery infrastructures, server virtualization is providing enterprises of all sizes with exciting new options for protecting a greater percentage of their workloads without breaking the bank. A significant number of organizations worldwide have already implemented successful workload protection and recovery plans using virtualization. Given the compelling cost and performance benefits of virtualized recovery (affordable consolidated workload protection, rapid recovery capabilities, flexible restore options and ease of testing), a growing number of enterprises will soon harness the benefits of virtualization to protect a broad spectrum of workloads that have previously gone under-insured.



About PlateSpin Ltd.

PlateSpin provides a unified suite of solutions to help enterprises adopt, manage and extend their use of server virtualization in the data center. PlateSpin's patent-pending Workload Portability™ technology liberates software from hardware and streams server workloads (data, applications and operating systems) over the network between any physical or virtual host with a simple drag and drop. By providing greater visibility into workloads and how they use available resources, PlateSpin solutions bring greater flexibility and new efficiencies to the virtualized data center. Global 2000 companies use PlateSpin solutions to lower costs, improve service levels and solve today's most critical data center challenges including server consolidation, disaster recovery and hardware migration. For more information please visit www.platespin.com



PlateSpin Ltd.
200 - 340 King Street East
Toronto, Ontario, M5A 1K8
Phone: 416 203 6565
Toll Free: 1 877 528 3774
www.platespin.com

© 2007 PlateSpin Ltd. All rights reserved. PlateSpin and the PlateSpin logo are registered trademarks and PowerConvert, PowerRecon, Server Sync, Workload Portability and PowerSDK are trademarks of PlateSpin Ltd. PlateSpin conversion and optimization technology and related products are patent pending. All other marks and names mentioned herein may be trademarks of their respective companies.