WHITE PAPER

Making Your Business Disaster Ready with Virtual Infrastructure

vmware®

Making Your Business Disaster Ready with Virtual Infrastructure

The Need for Business Continuity

Business continuity and disaster recovery (DR) planning are critical to managing risks in a successful business. Between 60-90% of companies that don't have a proactive disaster plan find themselves out of business within 24 months of experiencing a major disaster[1]. However, implementation of a reliable recovery strategy with fast time to recovery is expensive largely because it involves maintaining duplicate equipment that mirrors the equipment in the primary datacenter and upgrading both primary and recovery-target equipment in lock-step, hence many companies forgo the process.

Yet, companies that make compromises in disaster recovery strategy such as limiting the disaster coverage to only the most critical applications, employing manual processes to recover on dissimilar equipment, or outsourcing to discount DR centers experience insufficient disaster protection in terms of application coverage, acceptable downtime, and reliability of recovery.

This compromise is not necessary. In this paper we will discuss how to make disaster recovery cost effective with virtual infrastructure.

Assessing the Impact of an Outage

Business operations that are heavily dependent on information systems may be significantly impacted even by a brief application outage. Impact of a data loss is even more drastic. IDC estimates that in a disaster situation the average loss is $3 million per incident and $381,000 per hour.

To develop an effective business continuity strategy, businesses must assess how critical each of the IT applications is to the business viability.

This assessment may include:

- Hourly cost of outage (especially critical in revenue generating systems such as e-commerce and CRM)
- Reliability of recovery (especially critical in financial systems)
- Existence of alternative or manual processes that can be temporarily used in case of disaster

A comprehensive business continuity strategy includes a combination of disaster protection methods for different applications ranging from disaster prevention to hot sites and data backups, as well as a staffing plan for disaster situations, a well documented plan of action, and audit and testing processes.

Specifically, the cost of including a specific application or system in the disaster protection plan (under the umbrella of the larger business continuity plan) should be gauged against the potential risk and impact of the outage of this application. It is only commercially feasible to implement disaster recovery for the applications where such solution costs are lower than the impact of the outage.

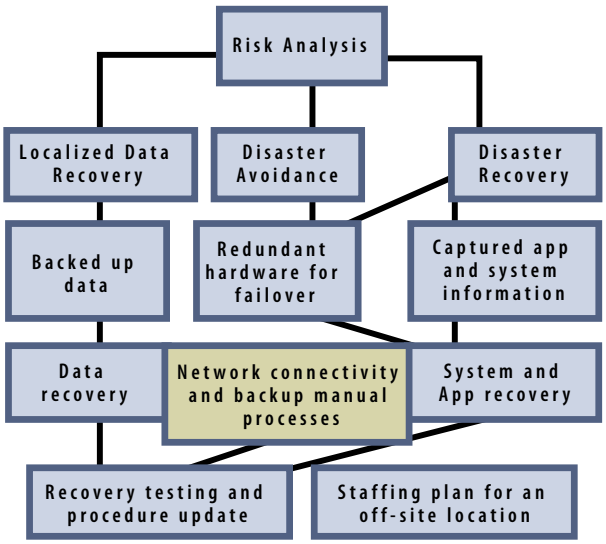

Commercially Feasible Business Continuity Strategy

There are several metrics commonly used in planning for disaster mitigation. Two of the most commonly used metrics are recovery point objective (RPO) and recovery time objective (RTO). Both are measured in minutes and hours. RPO describes how far the recovered data are out of synch with the production data at the time of the disaster. RTO describes how fast operations can be restored.

Other issues to consider are whether partial restoration of the IT systems such as restoration with decreased performance, decreased failure tolerance, or a partially incomplete set of data is sufficient to resume business operations after a disaster.

Methodologies for mitigating problems are as varied as the problems themselves.

Let us look at several possible Information System (IS) designs for business continuity.

- **Continuous availability**
  In this architecture, the workloads are load balanced over several – often geographically distributed – platforms. Each platform is provisioned to have spare capacity. When failure of one of the platforms fails, the workload is distributed over the remaining platforms. This approach is attractive because it allows companies to maintain business operations even after the disaster has occurred. Business operation is continuous and uninterrupted.
- **On-line and near-line hot sites**
  This strategy assumes availability of a failover site, which is equipped with power, cooling, network connections, physical security, and any other critical requirements. Sufficient equipment is available to resume business operations at the failover site in case a disaster is declared. Such a site can either be owned by the business or be provided by a third party. The data, application, and system information is replicated to the failover site either by using data replication methods or by shipping media with backed up data.

  This approach is attractive because it does not require a complete system overhaul like the continuous availability approach does; however, it still permits recovery from a regional disaster such as an earthquake or a hurricane.
- **Backup to tape**
  Finally, there is the well-understood method of backing up the data to tape using one of the popular backup management software packages. With this methodology, the backup is performed on a file-by-file basis. The backup may be full (all files are backed up), incremental (only files that have been modified since the last backup are saved to tape), or differential (all files modified after the last full backup are saved to tape). The tapes can then be stored in an off-site location for disaster mitigation.

  This approach is the least expensive and allows the use of the same methodology for discreet data problems (i.e. accidentally deleted files) and recovery in the case of a major disaster. However, a simple backup schedule does not amount to a comprehensive business continuity plan. It also doesn't encompass a plan to recover potential data loss which makes it difficult to predict the time it will take to resume the business operations.

Commercially feasible business continuity strategy represents a fine balance between the cost of implementing the business continuity plan and the impact and likelihood of the potential outage. The more cost effective methods allow IT professionals to put in place continuity plans with wider application coverage, allowing near uninterrupted operations.

Virtual Infrastructure makes business continuity commercially feasible for mid-size enterprises and expands the plans to cover applications beyond the top 5% of the most critical application set.

Virtual Infrastructure in the Enterprise

Virtual infrastructure provides a layer of abstraction between the computing, storage, and networking hardware, and the software that runs on it. Virtual infrastructure simplifies IT computing architecture, so companies can leverage their storage, network, and computing resources to control costs and respond faster. In a virtual infrastructure, users see resources as if they were dedicated to them. The administrator manages and optimizes resources globally across the enterprise.

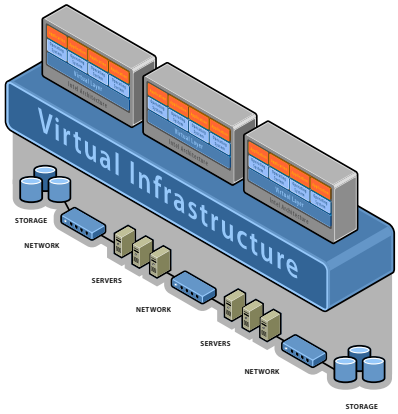

VMware's virtual infrastructure architecture enables businesses to lower IT costs through increased efficiency, flexibility, and responsiveness. Managing a virtual infrastructure enables IT to connect resources to business needs quickly. IT organizations can provision new services and change the amount of resources dedicated to a software service. The datacenter can be treated as a single pool of processing, storage and networking power.

Adopting virtual infrastructure lets IT be responsive to business needs, including:

- 60-80% utilization rates for Intel servers – up from today's 5-15%
- Provisioning times for new applications measured in tens of seconds, not days
- Response times for change requests measured in minutes
- Zero-downtime hardware maintenance without waiting for maintenance windows

The VMware virtual hardware platform implemented by VMware virtual machines makes virtual infrastructure possible. It creates a uniform hardware image – implemented in software – on which operating systems and applications run. On top of this platform, the VMware family of products provides management and provisioning of virtual machines, continuous workload consolidation across physical servers, and VMotion™ technology for virtual machine mobility.

With virtual infrastructure, IT organizations can provision new services and change the amount of resources dedicated to a software service. Hardware management is completely separated from software management, and hardware equipment can be treated as a single pool of processing, storage, and networking power to be allocated and de-allocated on the fly to various software services.

ESX Server

VMware® ESX Server™ transforms physical systems into a pool of logical computing resources. Operating systems and applications are isolated in multiple virtual machines that reside on a single piece of hardware. System resources are dynamically allocated to any operating system based on need, providing mainframe-class capacity utilization and control of server resources. VMware ESX Server simplifies server infrastructure by partitioning and isolating server resources in secure and portable virtual machines. VMware ESX Server enables these server resources to be remotely managed, automatically provisioned, and standardized on a uniform platform. Advanced resource management controls allow IT administrators to guarantee service levels across the enterprise. VMware ESX Server runs directly on the system hardware to provide a secure, uniform platform for deploying, managing, and remotely controlling multiple virtual machines.

With VMware ESX Server:

- Applications running on dedicated systems can be moved into separate virtual machines on a single, more reliable, and scalable system.
- Servers can be remotely managed from any location, simplifying server maintenance.
- Service levels can be guaranteed with advanced resource management.

GSX Server

VMware GSX Server™ is virtual infrastructure for enterprise IT administrators who want to consolidate servers and streamline development and testing operations. VMware GSX Server is easy to install and manage and provides rapid Return On Investment (ROI). Unlike other virtualization products, VMware GSX Server is enterprise-proven, preserves freedom of choice, and offers an upgrade path to datacenter-class virtualization.

With many years of proven success, thousands of customers trust VMware GSX Server as their virtualization solution. Easily installed on Microsoft® Windows or Linux platforms, VMware GSX Server provides advanced capabilities that make it the most flexible server virtualization product on the market. VMware GSX Server is part of the widely deployed VMware virtual infrastructure solution. Virtual machines are compatible across all VMware products and unified management and provisioning is provided by VMware VirtualCenter.

VMware GSX Server simplifies computing infrastructure by partitioning and isolating servers in secure and transportable virtual machines, each of which can run standard Windows, Linux, or Novell® NetWare® operating systems and applications. VMware GSX Server allows these virtual machines to be remotely managed, automatically provisioned, and standardized on a secure, uniform platform.

Thousands of enterprise customers rely on VMware GSX Server to deliver server scalability, reliability, and high availability and to maximize return on IT investments. VMware GSX Server is used across the enterprise to:

- Streamline software development and testing operations with easily provisioned and managed server-based virtual machines.
- Implement server consolidation for new and legacy departmental server applications.
- Provision servers rapidly to local or remote locations.
- Streamline OS and application patch management.

VirtualCenter and VMotion

VMware VirtualCenter is virtual infrastructure management software that provides a central point of control for your datacenter's virtual computing resources.

VirtualCenter allows IT managers to create a more responsive datacenter through:

- Instant provisioning
- Zero-downtime maintenance
- Intelligent workload management

These features allow faster reconfiguration and reallocation of applications and services. Servers are instantly provisioned. User-downtime is decreased. The datacenter becomes optimized.

VirtualCenter is a powerful way to connect IT services to business needs. With VirtualCenter, IT infrastructure becomes more flexible, efficient, and responsive. VirtualCenter uniquely leverages virtual computing, storage, and networking to improve datacenter management and reduce costs.

VirtualCenter centrally manages an enterprise's ESX Servers and virtual machines as a single, logical pool of resources. Key management functions built into VirtualCenter include:

- Dashboard of virtual machines
- Monitoring of system availability and performance
- Automated notifications with email alerting.

Using robust access control integrated with Windows authentication, VirtualCenter runs as a service on Windows 2000, Windows XP Professional, and Windows 2003. VirtualCenter works with the hardware and software infrastructure enterprises already have. VirtualCenter manages Intel-architecture based computing. It works with blades, 1-U rack-mounts – any server form factor from 2-16 CPUs per server.

How Virtualization Helps Business Continuity

Pain Points in Current Implementations of Business Continuity

There are many varied tools available for business continuity, however, due to the specifics of the Windows operating system design, even the most advanced tools can only provide a seamless restoration when the target and source physical platforms are identical. Maintaining identical physical platforms at the failover site means lock-step hardware upgrades in the primary and fail-over locations, which is prohibitively expensive. If no failover site is available, it can be impossible to locate identical hardware. Even hardware of the same series that is purchased from the same manufacturer is likely to have different firmware revisions, stepping levels, BIOS settings, or support lifecycles. Restoration to different platforms is often unreliable and includes complex manual steps. These manual elements and the need to troubleshoot problems cause long recovery times and lack of repeatability.

To assist enterprises in disaster recovery planning, operating system vendors, applications vendors, and backup management software vendors have developed specialized APIs, tools, and best practices. As a rule, these practices involve separate processes for backing up and restoring system information, OS information, applications, and data. Some applications and data, for example Microsoft Exchange, have modules that exhibit significantly different behavior, requiring each module to have a different disaster recovery strategy. Furthermore, each mission critical application has a different backup and recovery API. The differences are especially significant if the application needs to remain accessible for the duration of the backup. With a plurality of tools, enterprise IT managers have to learn new tools and design new strategy for each of the applications covered by the disaster recovery strategy. To complicate the situation even more, the methodologies and the APIs can change completely from one version of the application to the next. For example, a disaster recovery strategy for Exchange 2003 based on native Exchange APIs is entirely different than a disaster recovery strategy for Exchange 2000.

Because of the differences in applications and tools, disaster recovery strategy often includes several application-specific plans. Each plan has variations which have to be tested. In addition, if the strategy does not include a failover site, locating hardware to test the recovery can be a challenge.

Method	Target Recovery Time	Pain Points
Recovery from tape using backup agent native functionality	over 24 hours	Failover hardware not available for testing. Not suitable for mission critical applications.
Application specific backup agent deployed in conjunction with shadow storage volumes	4 – 24 hours	Many complex diverse application specific processes, manual processes to address dissimilarity of recovery hardware reduce reliability, separate processes for system, application, and data recovery.
Failover software with storage & server mirroring	Under 4 hours	High ongoing capital costs to maintain identical servers at the failover site. High management and maintenance expenses.

Built-in Continuous Availability

Virtual Infrastructure changes the way information systems are architected. Features such as migration of virtual machines between any virtualization platforms, snapshotting, and VMotion create environments where outages are limited to having to restart a virtual machine at most. Customers who wish to implement continuous availability solutions based on virtual infrastructure have several options.

For a continuous availability solution to guard against application or hardware failure, customers can use VirtualCenter to monitor a warning site and then migrate the failed application to another platform before catastrophic failure. For a more robust continuous availability solution without application interruption, N+1 clustering between virtual machines hosted on different physical hardware platforms can be implemented. Finally, if regional disasters are a concern, Virtual Infrastructure in conjunction with SAN and data replication technology offers the highest degree of protection. Customers can use data replication between primary and failover storage arrays and bring up virtual machines at the consolidated failover site.

Hardware Independence

One of the main benefits of virtualization for business continuity is independence of the recovery process from the recovery hardware. Because virtual machines encapsulate the complete environment, including data, application, operating system, BIOS, and virtualized hardware, applications can be restored to any hardware with a virtualization platform without concern for the differences in underlying hardware. The physical world limitation of having to restore to an identical platform does not apply.

Not only does hardware independence allow IT managers to eliminate manual processes associated with adjusting drivers and BIOS versions to reflect the change in platform, it also eliminates Windows registry issues and plug-and-play issues.

Hardware Consolidation

VMware enterprise customers actively take advantage of VMware consolidation benefits for their production and staging servers. The consolidation benefits are even greater for the failover hardware. Because it is extremely unlikely that all the workloads will fail at once and it is often acceptable to provide somewhat lower application performance in the failover facility on a temporary basis, customers experience a consolidation ratio of failover equipment that often reaches twice the consolidation ratio of the primary datacenter. The unexpected outcome of workload mobility and high hardware consolidation is that enterprises are able to oversubscribe hardware to multiple workloads with very little performance impact, which in turn makes in-sourcing a disaster recovery model much more economically attractive.

BC Solutions with Virtualization

Cost Effective High Availability

With virtualization, IT managers are able to create a cluster between a physical machine running mission critical workloads and a similarly configured virtual machine. The virtual machines do not consume computing resources in stand-by mode and can be consolidated to one or a few physical platforms at a very high consolidation ratio. As a result, the enterprise is able to realize high availability (HA) benefits without having to invest in twice the amount of hardware or having to manage and patch sprawling servers. Redundancy is reduced from 2N to N+1.

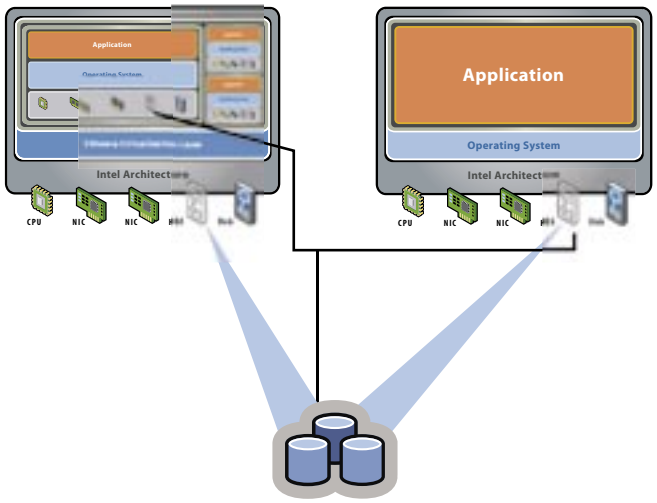

Physical to virtual machine clustering supports the same clustering software as physical to physical machine clustering. In fact, the same clustering software is supported for virtual machines as for their physical equivalent including Microsoft clustering, Veritas clustering, and Legato AAM, so no IT ramp-up is required. At the same time, reduced cost allows implementation of HA and Service Level Agreements (SLAs) for more workloads.

Continuity with Virtual Machines on Storage Area Networks

Virtual Infrastructure deployed in conjunction with a Storage Area Network (SAN) has an additional built-in level of robustness. Any virtual machine that resides on a SAN can survive a crash of server hardware, which runs this Virtual Machine (VM), and can be restarted on another ESX Server.

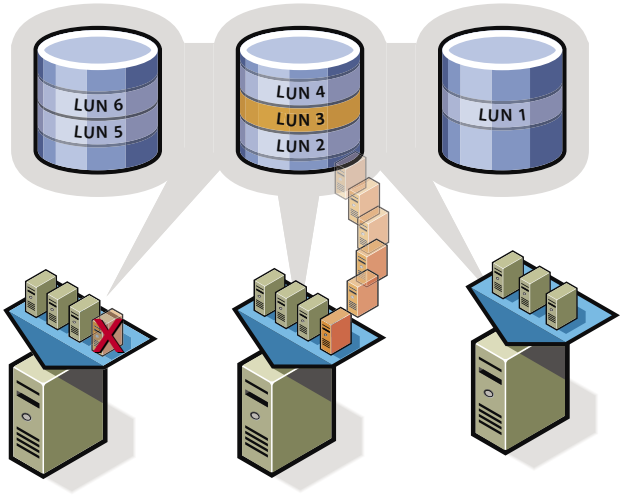

Better yet, VMware VMotion technology allows migration of a workload off the machine which is about to experience outage with no user downtime. Furthermore, virtualizing IT infrastructure improves the business case for increased SAN deployment. Because each server and associated host bus adapters (HBAs) are shared between multiple workloads, the per workload cost of SAN attachment is reduced dramatically. In addition, using multiple HBAs for failover and multi-pathing becomes more affordable, thereby improving availability and eliminating single points of failure.

Inter-ESX Server Clustering

This solution is based on a more traditional view of high availability which minimizes IT management cycles and reduces the need to develop complex application and OS specific clustering strategies. This approach clusters together two ESX Servers at the virtualization layer.

If a virtual machine fails on ESX Server A, it will be immediately restarted on the second server of the clustered pair, ESX Server B. The advantage of this approach, in addition to drastically lower HA management costs, is that it enables failover even for the applications that are not inherently cluster aware. The disadvantage is that the failover is crash consistent (i.e., the state of the application is not preserved at failover).

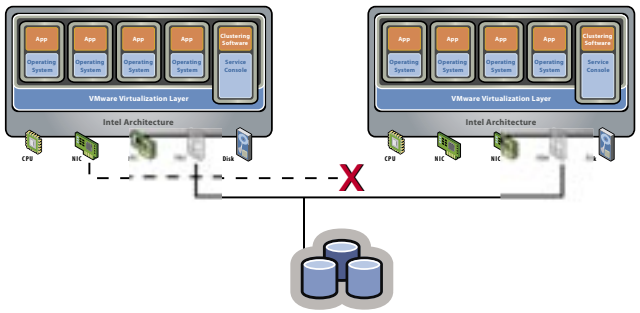

Storage Array Based Replication

For the most critical applications, many enterprises turn to storage array based data replication to a failover site. This approach provides the most up-to-date copy of the data and applications at a remote location, thereby protecting data from a regional disaster as well as from a hardware failure.

However, the question of how quickly the operations can be restored at the secondary site remains to be answered. If only storage arrays are maintained at the secondary site, the servers needed to run the applications must be procured if a disaster is declared. This configuration results in recovery time ranging from days to weeks. Also, the recovery to dissimilar hardware is a very risky manual process and in some cases is not possible.

For guaranteed recovery within hours, server hardware at the secondary site must be upgraded in lock-step with the primary datacenter equipment. Even this approach may not meet recovery time objectives for the more demanding workloads.

Virtual Infrastructure combined with array based replication allows enterprises to replicate the encapsulated virtual machine to the secondary site and bring it up at the secondary site in a programmatic way, without human intervention, on any available ESX Server. The ESX Server hardware at the secondary datacenter does not need to match the ESX Server hardware configuration at the primary datacenter. Furthermore, a higher ratio of server consolidation can be maintained at the secondary site.

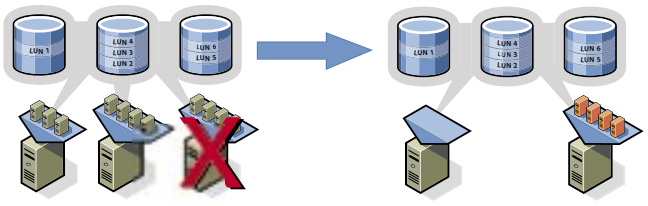

Virtualized Failover Site

To test data restoration, IT administrators need to locate a test failover server for each of the backed up machines, install the OS and backup agent, and then try to adjust the Windows registry and other system configurations on the test failover server. If system adjustments are successful, the backup server and the backup agent can be used to test data restoration.

The two obvious drawbacks of this approach to enabling data recovery are having to provision several new servers and the fact that it is not always possible (or at best a long manual process) to adjust Windows registry and other system characteristics of a dissimilar failover server.

All these issues are resolved by using virtualized failover hardware. Moreover, OS installation, backup agent installation, and Windows registry adjustment only need to be done once. Thereafter, a fully configured VM template is stored in a VM template library. For all the subsequent recovery tests the step by step process would be:

1. Source ONE physical machine, regardless of the number of servers that are slated for recovery testing.
2. Install ESX Server.
3. Copy (from a library) a VM with the appropriate Windows version and backup agent pre-installed.
4. Edit IP addresses and register the VM.
5. Start the VM.
6. Restore (from tape) into the VM using the backup agent.

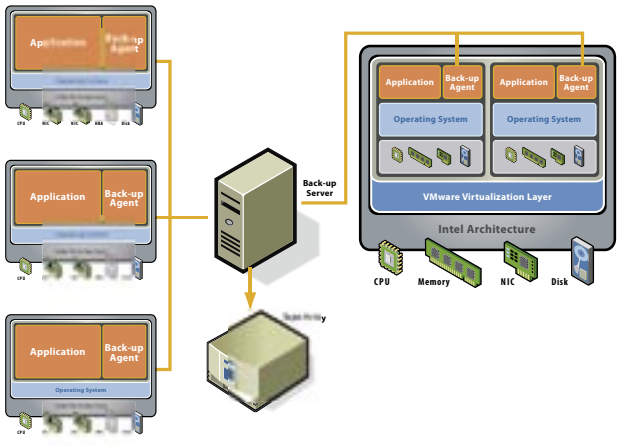

Integration with Backup Software

With Virtual Infrastructure, IT managers have an option to continue using the existing backup processes even with virtualized hardware. VMware software supports a wide variety of backup agents operating inside VMs, which allows the backup server to control the backup and file restoration process in a consistent manner for physical and virtual servers. Another alternative is to use a backup agent in the ESX Server Management Console or on the GSX Server Host OS. This method allows IT managers to take advantage of the encapsulation property of virtual machines (i.e., to take a consolidated backup of the entire VM including the system configuration, the application, and the data). For applications with very strict RTO requirements, the snapshot feature can be used to cut out booting time. Some additional options include: using disk log/redo feature to enable "live" application backup without performance impact.

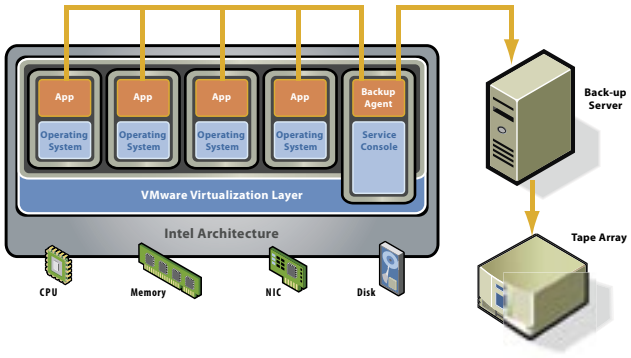

Before and After Examples

Online Hot Site

Let's consider a sample deployment by a VMware customer in the manufacturing industry.

The customer has implemented server consolidation with ESX Server connected to EMC Symmetrix SAN. All virtual machine images are complete with virtualized hardware, OSes, security patches, and applications residing on a SAN.

The deployment includes 78 physical dual-CPU Dell servers with ESX Server in the main data site, 38 HP servers in the secondary datacenter (17 miles away). The secondary datacenter is oversubscribed with the ratio of approximately 2:1. The data are synchronously replicated with SRDF (EMC remote data replication). The two sites are connected via fiber running a dedicated OC3 connection at 155Mb/s and IP over fiber.

The VMs in the secondary datacenter are powered off during normal operation. When a disaster is declared, a script is launched that uses EMC Control Center to break off the mirroring between the SANs in the primary and secondary datacenters, unlock Logical Unit Numbers (LUNs) in the secondary datacenter, and power on the VMs in the secondary datacenter.

The complete time to recover is now 16.5 minutes.

Before this system was implemented the disaster recovery was outsourced and tape restore method was used on mismatched hardware. The estimated time to recover was 6.5 days. Recovery tests were never successful.

The new default IT policy states that every new application installed in the datacenter will go to the VM and be covered under this disaster recovery strategy.

In this customer scenario, great improvements in TTR were realized, the process was made programmatic and more reliable, and per application protection costs were reduced. As a result, the customer is able to provide disaster recovery coverage to a wider range of applications.

Tape Backup and Recovery

Let's consider a pre-virtualization disaster recovery procedure with tapes.

A customer's datacenter applications are deployed on 15 physical servers. A backup server/agent solution is deployed.

A complete system backup-up is performed weekly, and an incremental file level backup is performed nightly.

If the server is not taken offline during a full backup, its performance is significantly impacted. Complete time to create tape backup of a 300GB datacenter can be estimated at 600 minutes (~500MBpm). For each server, there are three tapes (system, application, and data). Incremental backup is performed nightly. Assuming 10% of information has changed, the nightly backup window is on the order of 60 minutes.

For a disaster recovery plan that includes a secondary datacenter, tapes of the backup data need to be shipped to the secondary datacenter.

Restoration would involve locating the same number of available servers in the secondary datacenter (15 servers in this case). Subsequently, for each server, the IT administrator would have to assemble all the media pertaining to this server, make adjustments to the partition scheme and BIOS configuration, install the operating system, install the necessary patch level, configure networking and driver replacements, and then install the backup agent and restore the system, applications, and data from the last full backup. This work is likely to take about five hours per server. If the RPO requirement is not met, additional data can be recovered from file-based backup.

Disaster Recovery with Tapes for a Virtual Infrastructure

This is an example of how disaster recovery procedures can be improved if the datacenter infrastructure is virtualized. Other data and recovery architectures are possible depending on customer-specific configurations and requirements.

In this scenario, the original 15 physical servers in the customer's datacenter are consolidated into 15 VMs on three physical servers.

On a weekly basis, a full backup of VMFS is done using backup agent in the console OS. The information encapsulated in the files that are backed up includes (virtualized) hardware configuration, OS, system and application, and data for each virtual machine. IT managers must change the disk mode to append, which will unlock the virtual disk and make it available for backup without taking the virtual machine offline. At this point, all the disk changes resulting in server operation will go to a dynamic redo log.

Note: The dynamic redo log is actually a write cache that holds the most current changed blocks of data.

The datacenter information is then backed up to tape. Although, the backup process will take approximately the same time as in the scenario without virtualization, the backup will result in little service performance impact.

Once all the information is on tape, the redo log is committed, and the virtual machines are converted to persistent mode. Next, the tapes with the information are sent to the secondary datacenter.

In the secondary datacenter, two physical servers are pre-configured with ESX Server. We are able to provision fewer physical servers in the secondary datacenter because VMs can be restored to any of the servers and moved between servers when utilization on one of the physical servers peaks, so a higher consolidation ratio for the secondary datacenter is acceptable.

The recovery will now take about five hours per physical server because of large amount of data that needs to be restored. However, because there are only two physical servers in the secondary datacenter now, we have reduced recovery time from 65 hours to 10 hours. We have additionally removed manual elements (hardware configuration and OS installation) from the restoration process and made it much more reliable.

What has improved:

- Cumulative number of servers in primary and secondary datacenter was reduced form 30 to 5 servers.
- Recovery time was reduced from 65 hours to 10 hours + tape shipment time.
- Impact of backup process on application performance was minimized.
- Manual elements were removed from the recovery process and it was made more reliable.

An additional performance gain for recovery can be achieved by using virtual tape devices, so server backups are not interleaved on the final backup tape.

Benefits of Virtualization for Business Continuity

If the recovery time objective for recovery to the state of the last full backup is under one hour, the only successful recovery strategy is to maintain a secondary datacenter equipped with the same model hardware as the primary datacenter, server for server.

A bare metal restore tool can be used in conjunction with backup or disk replication software to restore each of the mission critical servers – complete with the operating system and the mission critical application – to the corresponding server in the secondary datacenter.

The recovery procedure is specific to each server, so each server recovery has to be tested separately.

When the secondary datacenter is virtualized, there are three immediate benefits:

- There is no need to maintain the same model hardware because IT managers can restore applications encapsulated into VMs to any Intel architecture hardware, and they don't even need to license specialized bare metal restore tools.
- IT can pool together all the datacenter hardware and realize economy of scale benefits.
- IT managers only have to manage a single type of data -- encapsulated virtual machines — for capture and recovery. This approach drastically simplifies management complexity compared to the traditional approach of having to deal with disparate systems, applications, and data.

The secondary datacenter doesn't need the same model hardware, so upgrades of the secondary datacenter do not need to be done in lock-step with the primary datacenter. While servers in the primary datacenter are replaced on average once every three years, the servers in the secondary datacenter may have a life of six years. As servers in the primary datacenter are phased out, they can be redeployed to the secondary datacenter to augment capacity.

The ability to pool hardware resources together and balance all the mission critical workloads across the different servers in the datacenter results in an increase in the consolidation ratio in the secondary datacenter with minimal impact on the availability. For example, a twofold increase in consolidation ratio may only result in a decrease in availability from 99.98 to 99.95 in the restored applications, even in a case where all the DR protected applications failed at once.

Requirements for fewer servers immediately results in lower TCO because of lower power and cooling requirements, facilities requirements, wiring and networking elements, and savings on hardware maintenance.

More IT cycles are saved because there is less need for hardware upgrades (longer hardware life), simplified recovery testing (test once to recover all virtual machines), and shorter personnel training (uniform processes for all the applications)

Total cost savings given the same level of protection:

- Capital cost reduction: 50-70% a year
- Variable cost reduction: 60-80% a year
- IT resource requirement reduction: 70-90% a year

Insourcing vs. Outsourcing Trend

Outsourcing disaster recovery facilities, popular among medium size enterprises in the last several years, has been decreasing lately. The reason for the trend reversal is both the tendency of the leading outsourcers to over-extend their resources, decreasing the quality of service, and the first-come/first–served policies that create additional risk in regional disaster situations. Virtualizing failover facilities and consolidating failover servers to fewer physical boxes allows enterprises to insource, while maintaining the same or better cost structure as with an outsourcing model, and remain in control and avoid risk.

Expanding Business Continuity Coverage

Incorporating virtualization with business continuity planning results in much lower fixed costs for implementing a hot failover site. In addition, per application business continuity TCO is lower because business continuity processes are standardized.

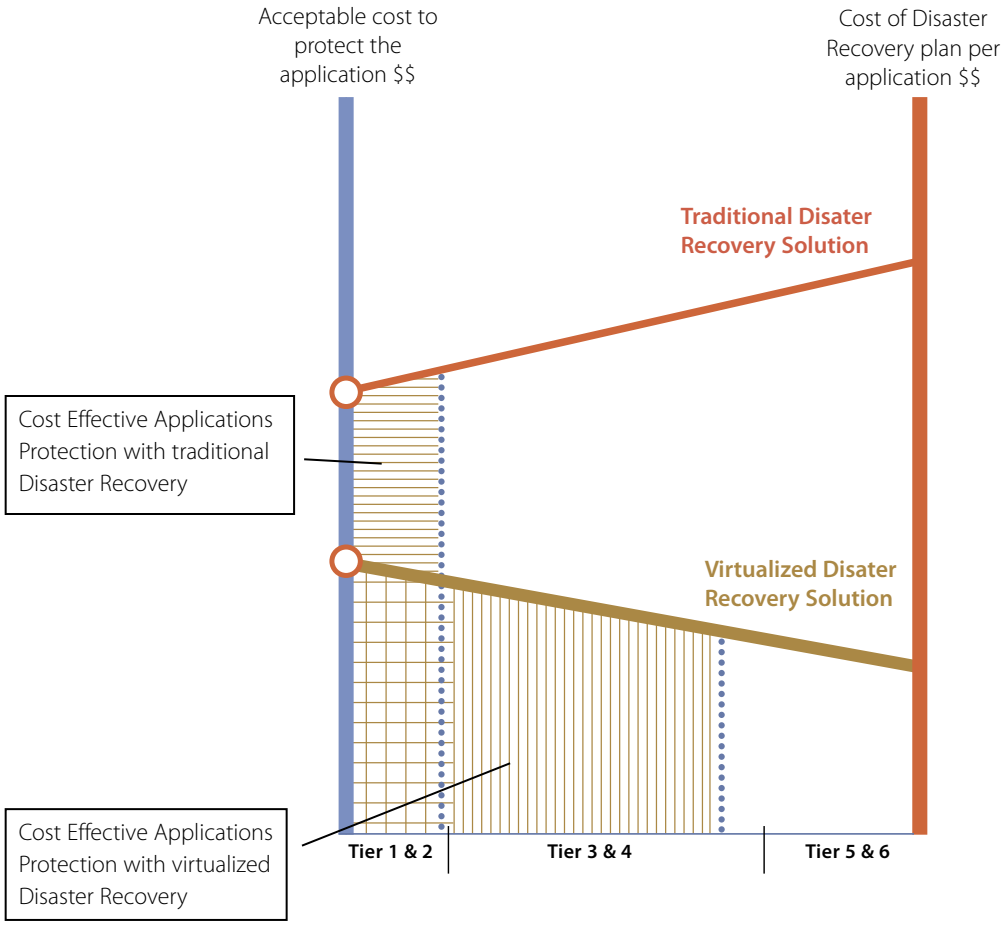

Conclusion

Using virtualization for disaster recovery allows companies to extend the disaster coverage to more applications while reducing the recovery time and making the process more reliable. To find out how to implement this solution in your environment e-mail sales@vmware.com or call 877-4VMWARE.

List of Acronyms and Abbreviations

API – Application Programming Interface

BC – Business Continuity

CRM – Customer Resource Management

DR – Disaster Recovery

ERP – Enterprise Resource Program

IT – Information Technology

IS – Information Systems

OS – Operating System

RTO – Recovery Time Objective

RPO – Recovery Point Objective

TTR – Time To Recover

SAN – Storage Area Network

SI – System Integrator

SLA – Service Level Agreement

VI – Virtual Infrastructure

VIN – Virtual Infrastructure Node

VM – Virtual Machine

VMFS – Proprietary file system optimized for VMware ESX Server Virtual Machines

(Footnotes)
1 The Definitive Handbook for Business Management

V00014-20001205

VMware, Inc. 3145 Porter Drive Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com

vmware®