WHITE PAPER

Disaster Recovery Virtualization

Protecting Production Systems Using VMware Virtual Infrastructure and Double-Take®

vmware®

Double-Take Software™

Contents

Introduction

In today's businesses, the need to protect critical IT system data is no longer in question. There are many types of protection solutions available, and choosing the right solution for your data can be challenging. One of the biggest obstacles to choosing an effective protection solution is cost. Server virtualization and real-time data replication and failover technology are capabilities that provide a high level of data protection while keeping costs at a minimum. VMware Infrastructure and Double-Take are software solutions that provide these cost-saving technologies and can be combined to provide affordable data protection for businesses of any size.

What is VMware Infrastructure?

VMware Infrastructure is the most widely deployed industry-standard virtualization software suite. The VMware Infrastructure software suite is fully optimized, rigorously tested, and certified for the widest range of hardware, operating systems, and software applications. VMware Infrastructure provides built-in management, resource optimization, application availability, and operational automation capabilities — very useful in disaster recovery environments. Using these capabilities, you can achieve cost savings — saving capital as well as operational costs. In many cases, using virtualization allows customers to implement disaster recovery solutions that would be cost prohibitive in nonvirtual environments.

VMware Infrastructure has two key components:

VMware ESX Server is datacenter-class virtual infrastructure software for partitioning, consolidating, and managing systems. ESX Server, included in VMware Infrastructure, provides a highly scalable virtual machine platform with advanced resource management capabilities that can be managed by VMware VirtualCenter. ESX Server runs directly on x86 hardware, providing high performance and complete hardware resource control. This component frees you of the need for having identical hardware at the primary site and the disaster recovery site. It also eliminates the need for complex server recovery processes.

VMware VirtualCenter is the centralized management component that provides operational automation, resource optimization, and high availability to IT environments. Virtualization-based distributed services equip the datacenter with unprecedented levels of responsiveness, serviceability, efficiency, and reliability. VirtualCenter delivers the highest levels of simplicity, efficiency, security, and reliability required to manage virtualized IT environments of any size.

What is Double-Take?

Double-Take protects data stored on Windows file servers and application data used by mission-critical applications such as Microsoft SQL Server and Microsoft Exchange Server. It combines continuous real-time backup and automatic failover capabilities for disaster recovery, high availability, and centralized backup. The real-time protection of Double-Take provides a recovery point objective (RPO) of seconds rather than minutes or hours as other solutions provide. By providing application monitoring and automated failover, Double-Take can provide a recovery time objective (RTO) of minutes rather than hours or days by providing immediate availability of protected services on a secondary standby system.

The Double-Take Server Recovery Option (SRO), an add-on component for Double-Take, is a full server data protection solution that, when combined with Double-Take, simplifies server protection and restoration. The entire server, including operating system, applications, and data, is protected with continuous real-time replication and can be recovered to the same or a different server in the event of a failure.

Double-Take is Microsoft Windows 2000 and 2003 certified at all levels, one of the few replication products to have achieved this level of certification. It delivers better protection than many hardware-based solutions, and it costs tens of thousands of dollars less.

Double-Take benefits:

- Real-time data protection — Replicates continuously at the byte level over any shared or private IP network, ensuring that changed data is protected and can be quickly restored at all times
- Application agnostic - —Works with existing hardware to protect applications such as Microsoft Exchange, Microsoft SQL Server, and Microsoft SharePoint.
- Continuous data protection — Guarantees business continuity and high availability by restoring access to data in minutes with failover capabilities to maintain a seamless working environment.
- Easily installed and maintained — Allows companies of any size looking for data protection solutions to install and maintain Double-Take.
- Cost-effective — Provides the best possible protection at the lowest cost with an accelerated return on investment, paying for itself within months.

Data Protection Criteria and Objectives

Virtualization and real-time data replication and failover are just a few of the many tools available on the market today for protecting critical data and systems. To understand when these tools are appropriate and how to implement them, it is important to know what systems are to be protected and the level of protection they require, based on some key metrics.

The first key metric is the recovery point objective. For a data protection solution, RPO defines how much data could be lost in the recovery procedure. If a single disk drive in a RAID 5 array fails, there is no associated data loss. This is an RPO of zero. If an entire disk array fails and data must be restored from the last good tape backup, all data that has changed since that backup will be lost. In most cases, this is an RPO of 12–24 hours.

The second critical metric is recovery time objective. Used hand-in-hand with RPO, RTO measures or defines how long the entire recovery process should take before users can reconnect to their applications and continue working. If a virus corrupts a database or a single file, a simple tape restoration might resolve the issue. Tape backup, in this scenario, would provide an RTO of hours. If an entire production server fails and must be rebuilt using the same tape backup solution, the RTO could quickly grow to several hours or even days, depending on how long it takes to repair the existing hardware or locate new hardware, reinstall the operating systems and applications, retrieve tape media, and restore the data.

Once such recovery objectives as RPO and RTO have been quantified, most organizations discover that traditional tape backup alone is not good enough to achieve their goals for many critical applications and servers. For this reason, many organizations deploy such technologies as data replication and application availability coupled with operating system virtualization. Deployed properly, a solution based on these technologies can provide protection superior to tape backup at a fraction of the cost of hardware-based or clustering-based solutions.

As shown in Figure 1, tape delivers poor RPO and RTO, measured in hours or days. Double-take reduces RPO and RTO to seconds or minutes.

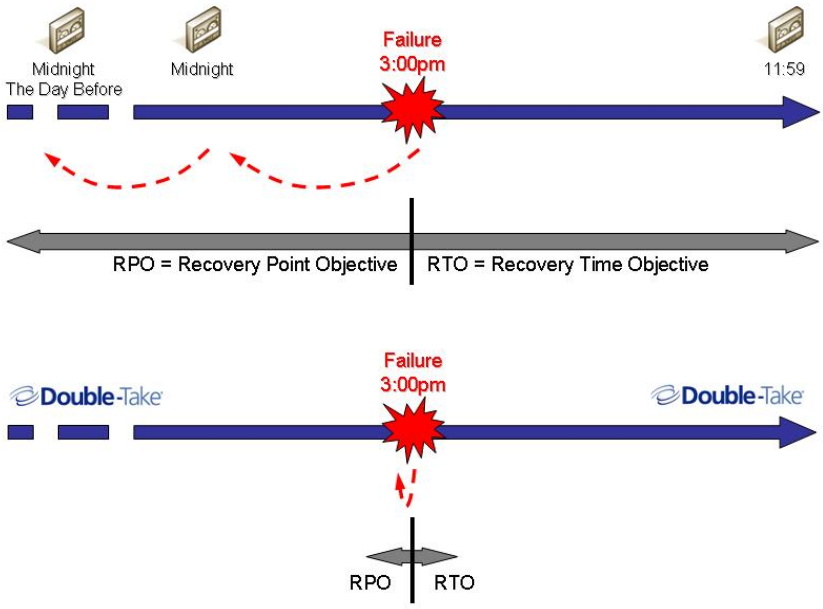

Figure 1 – Tape backup (top) compared to Double-Take RPO and RTO (bottom)

Using Virtual Machines and Data Replication for Disaster Recovery

There are many benefits to virtualization technologies. Among them is additional flexibility and cost savings in the deployment of a disaster recovery solution. Simply put, virtualization can reduce the amount of hardware required at a disaster recovery site and simplify recovery operations.

Solutions such as Double-Take that are based on replication and failover often require a one-to-one pairing of production systems with disaster recovery systems. Due to interoperability issues with some server-based applications and the complexity of managing such a configuration, it is often either not recommended or not possible to fail over multiple physical workloads to a single operating system instance running on standard server hardware. As a result, organizations usually must purchase enough hardware for the disaster recovery site to handle production capacity or make sacrifices by choosing not to protect certain systems.

By leveraging virtual machines as secondary servers in a standard replication and failover scenario, each virtual machine is its own self-contained, unmodified server image. Many of these virtual machines can be run simultaneously on a single piece of hardware, allowing many physical production servers to be protected by a single piece of hardware in a disaster recovery facility. Because each virtual machine is independent of the others and workloads do not need to be consolidated, managing applications and services during the recovery process is no more difficult than managing them in production.

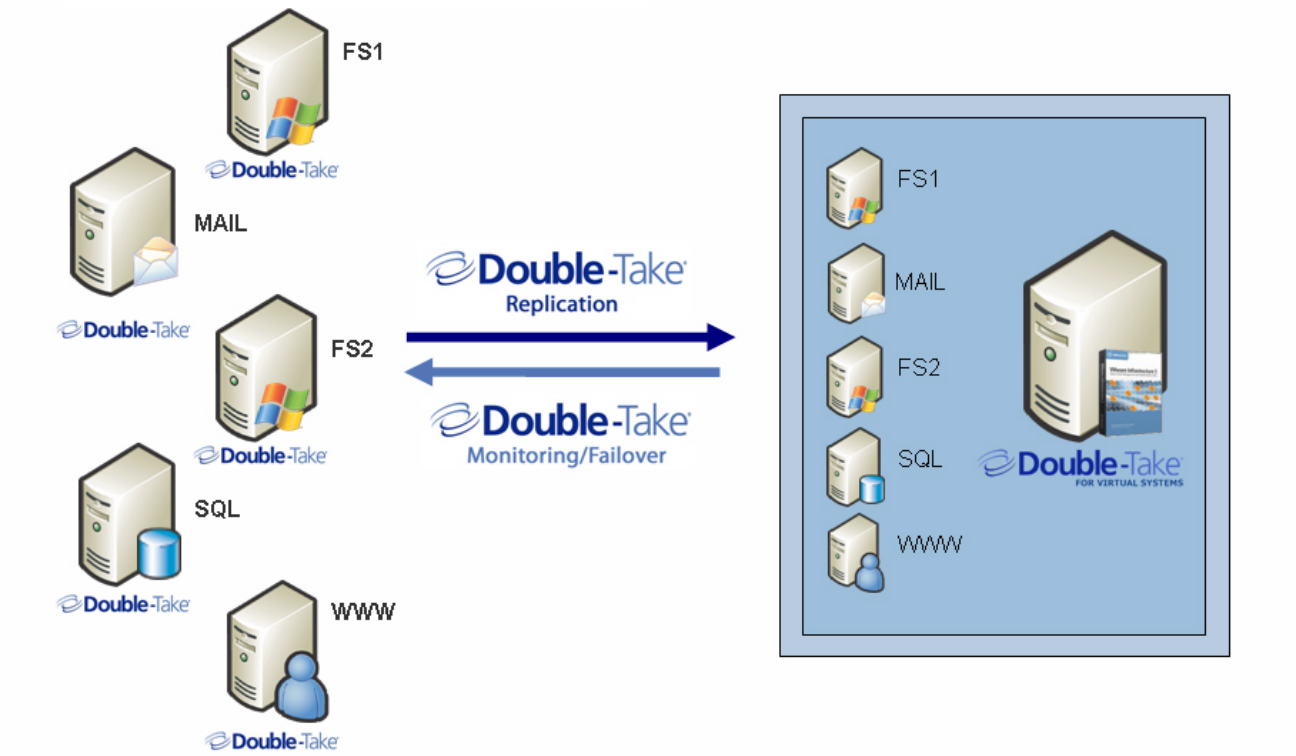

Figure 2 – Example of physical servers protected by virtual machines

VMware Infrastructure helps customers realize huge benefits when implementing their disaster recovery plans. With the built-in VMware High Availability feature, customers can reduce the unexpected downtime completely. Resource pools in VMware Infrastructure offer customers the advantage of aggregating their CPU, memory, and I/O resources. They can then deploy the applications on the resource pools, rather than worrying about individual hosts. Many tasks can be monitored easily using VirtualCenter. Using this technology increases the efficiency and the effectiveness of the disaster recovery process. It also offers cost-savings opportunities for customers.

Because virtual machines are hardware independent, customers can implement disaster recovery with hardware that is not identical to that at the primary site. This approach can offer customers many opportunities to save on hardware costs.

More than 55 percent of VMware customers virtualize in order to implement an effective disaster recovery solution using virtualization technology. According to an industry analyst, more than 70 percent of VMware customers deploying disaster recovery are doing so for the first time. The main driver behind this move is the cost savings achieved by customers using virtualization for disaster recovery.

Virtualization also simplifies disaster recovery operations. Virtual machines, by their nature, are independent of the physical hardware they are running on and are completely self-contained. Moving a virtual machine from one physical server to another is as simple as copying all of the associated configuration and virtual disk files to the other server and importing their configuration within ESX Server. It is also much easier to recover a virtual machine than it is to recover a physical machine, using either tape-based or disk-based backups.

From the RTO example earlier in this white paper, consider the restoration of an entire server if it is lost or damaged beyond repair. First, new hardware must be acquired. Next, the operating

system must be installed and configured, along with any needed service packs and drivers. Next, the application running on the server (Microsoft Exchange or Microsoft SQL Server, for example) must be installed and configured properly to mimic the previous installation. After all these steps are finished, the tape backup software must be installed, the tapes must be retrieved, and the actual data must be restored from the last good backup. With virtual machines, restoring a server is as simple as retrieving a copy of the virtual machine's files, restoring them to a new ESX Server host, then importing the virtual machine configuration. This alone could reduce recovery time from hours or days to minutes.

How ESX Server Works

VMware ESX Server is a robust, production-proven virtualization layer that abstracts processor, memory, storage, and networking resources into multiple virtual machines. VMware ESX Server is the bare-metal virtualization software that runs directly on the x86 hardware platform. By eliminating the host operating system, customers can achieve acquisition and management cost benefits while improving performance at the same time.

VMware ESX Server supports virtualization for networking by allowing customers to create Virtual switches inside the server. These virtual switches can use multiple Ethernet adapters on the system to provide network redundancy and load balancing. By programming these switches, customers can change the networking environment without any hardware configuration changes. This virtual switch reconfiguration technique can be effectively used in testing the disaster recovery site. Simply reprogramming the virtual switch isolates the virtual machines running on the secondary site from the production network. Once testing is done, the virtual switch can be reconfigured to connect back to the production network.

The VMFS file system built into ESX Server also offers similar benefits for DISASTER RECOVERY testing. Customers can use snapshot feature of the VMFS file system to create copies of the VMs for testing. Once testing is done, the copies can be destroyed by deleting the snapshots. Apart from snapshots, VMFS is a logical volume manager that can manage the IO going from the server to the storage devices. VMFS is also a clustered file system. This clustering allows customers to move virtual machines freely between servers, thus load-balancing the environment,

VMware ESX Server management is done through VirtualCenter GUI. The agents are pre-installed on the ESX server. Of course, every ESX server also has a service console attached that can be used for running administrator commands or for troubleshooting. VMware VirtualCenter stores information about the virtual infrastructure data center in a database. VMware ESX server does not store any configuration information for the virtual machines. It has a very small amount of networking configuration information stored. Thus its very easy to re-build the ESX server in case of a loss or a failure. This property of the ESX server makes it useful in a disaster recovery environment. Because there is never any real need for backing up and recovering the ESX Server host itself, customers can rebuild the server on the fly, if needed.

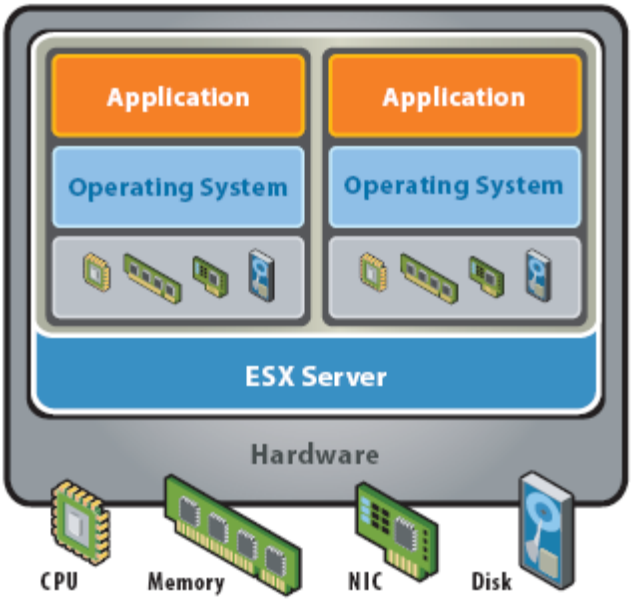

Figure 3 – VMware ESX Server architecture

How Double-Take Works

Double-Take is a real-time data replication and failover software product. Double-Take augments your existing data protection strategy by reducing downtime and data loss, and it provides these services with minimal impact on existing network and communication resources. Double-Take allows you to specify mission-critical data that must be protected and replicates that data, in real-time, from a production machine known as the source to a backup machine known as the target. The target machine, on a local network or at a remote site, stores the copy of the critical data from the source.

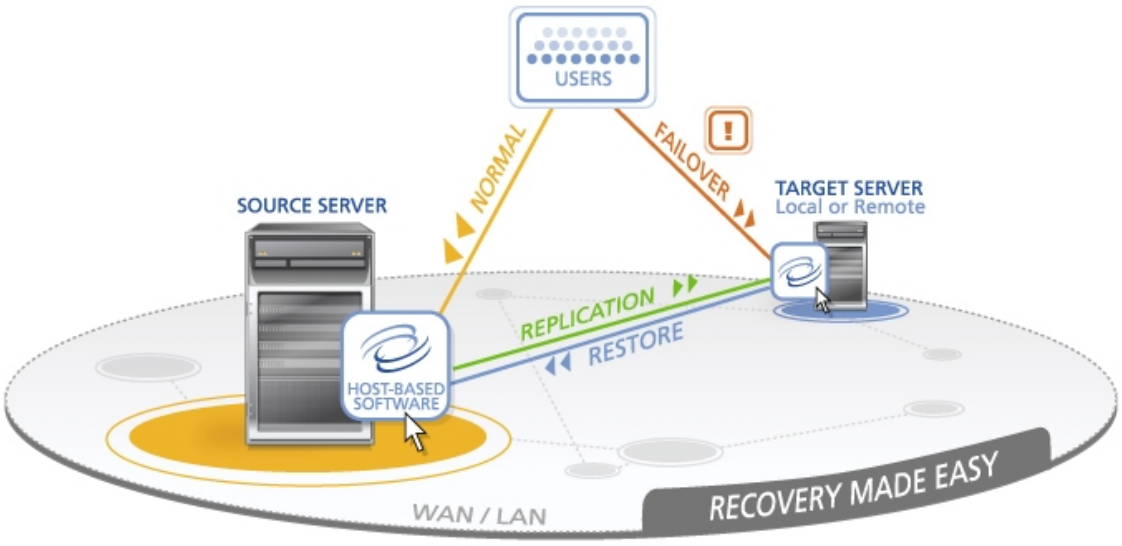

Figure 4 – Simple Double-Take Configuration

Double-Take monitors any changes to the protected data and sends the changes to the target machine. By replicating only the byte-level changes rather than copying an entire file, Double-Take allows you to more efficiently use resources.

To eliminate downtime during an outage, the real-time copy of the data on the Double-Take target can be used to resume running protected applications like email or database services on the secondary server. The Double-Take service running on the secondary server can monitor the production server and in the event of an outage automatically start the appropriate application services on the secondary server, then seamlessly redirect end-users' requests. This combination of real-time replication and application availability enables the implementation of various data protection solutions including:

- Off-site disaster recovery services
- Local and remote high-availability services
- Enhanced centralized backup using third-party backup systems

How the Double-Take Server Recovery Option Works

Double-Take, when combined with the Double-Take Server Recovery Option (SRO), provides a single solution to continuously protect and recover an entire server. Protection is provided by the industry-proven real-time replication of Double-Take. Recovery is performed by the Server Recovery Option. The Recovery Manager, provided as part of the Double-Take Server Recovery Option, presents the recovery task of a server as a series of easy-to=understand steps. Because Double-Take replication protects the entire production server — its operating system, applications, and data — restoring the server requires as few steps as possible and provides a significantly better RTO than existing solutions, such as tape backup.

Double-Take Server Recovery Option replicates entire servers to a central image server where the images can be restored in only a few simple steps. Recovery requires only that the operating system be installed on the recovery server. SRO takes care of the rest, pushing the installation of Double-Take to the server and restoring all of the data and system state in one recovery process,

unlike the multiple tape restores required by tape backup. Both the backup and recovery processes can be performed over any distance on standard IP networks.

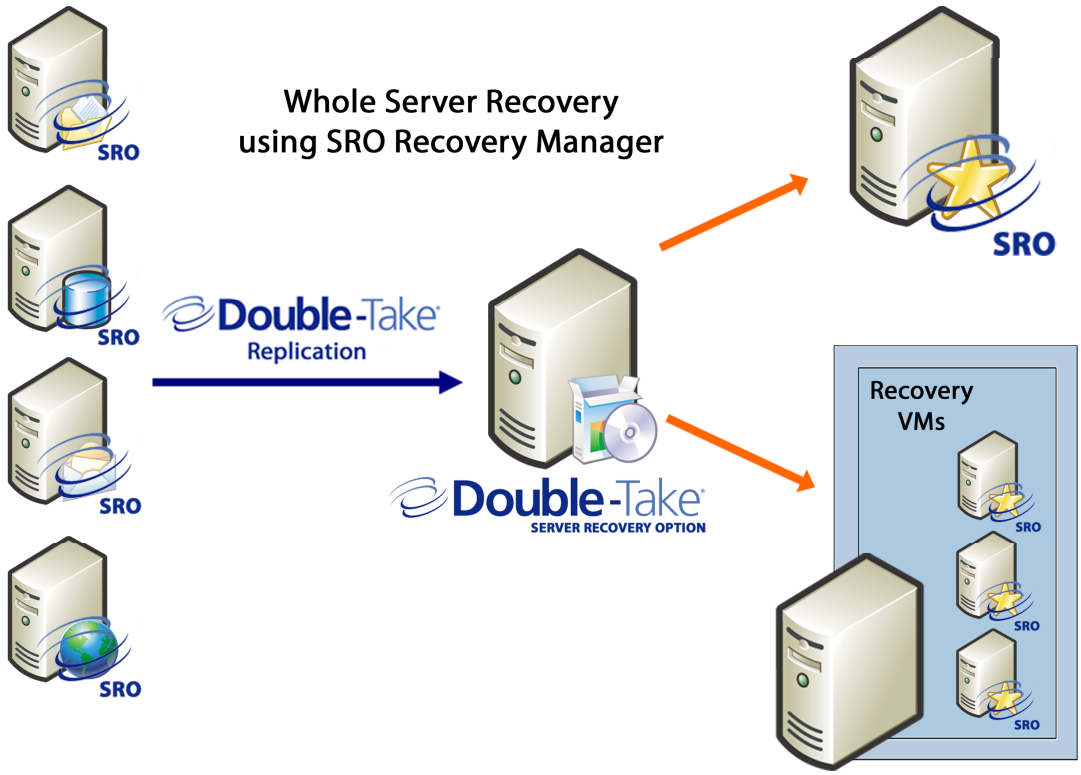

Figure 5 – SRO Replication and Recovery to Physical or Virtual Machines

Physical-to-Virtual Disaster Recovery — Implementation Overview

The first step in protecting physical servers and workloads with virtual machines and Double-Take data protection and failover is to select which production servers will be protected. For each production server, Double-Take needs to be installed and a virtual machine must be created running on an ESX Server host to serve as the Double-Take target or as a recovery server for Double-Take Server Recovery Option.

Double-Take System Requirements

There are two editions of Double-Take applicable to protecting Windows servers. If the production server is running Windows 2000 Server or Windows Server 2003 Standard Edition, Double-Take for Windows Server Edition or Double-Take for Windows Advanced Edition can be used. If the production server is running on Windows 2000 Server Advanced Edition or Windows Server 2003 Enterprise Edition, Double-Take for Windows Advanced Edition is required.

If using Double-Take Server Recovery Option, the Server Recovery Option must be installed on the source servers and the image server. The source server and image server do not require the same operating system version or Double-Take license version. However, the recovery server operating system version must match the operating system version of the source image being recovered to it.

There are also various hardware requirements for Double-Take, though many of these requirements are already satisfied by the minimum requirements for most production applications.

- Network — TCP/IP LAN connection with one network adapter and one IP address
- System memory — 64MB minimum
- Disk space on source and target (application) — approximately100MB free space

Secondary (Failover) Server — Virtual Machine Creation

The first step in configuring the secondary server to be used for real-time replication and failover is creating the virtual machines that will be used. This can be done in several ways, including manually creating a new virtual machine, then installing a copy of Windows Server in the newly created virtual machine. But that approach involves constantly managing the virtual machine to keep it in synch with the physical server of which the virtual machine is the copy. To save time and reduce the possibility for errors, it is recommended that tools such as VMware Converter be utilized to automate the creation of a virtual machine based on the physical server running in production.

VMware Converter is the next-generation conversion tool designed to create virtual machines easily from sources such as physical machines, other VMware virtual machine formats (VMware Server or VMware Workstation, for example) or third-party formats. VMware Converter provides centralized management plus fast and reliable conversions. The goal of VMware Converter is to shrink the amount of time that IT professionals spend converting physical machines to virtual machines by reducing the amount of manual effort required for each conversion. VMware Converter provides the ability to do remote conversion with zero disruption to the source server.

VMware Converter is managed through a simple, tas-based user interface that enables users to convert physical servers or various disk image formats to VMware virtual machines in three easy steps:

1. Specify the source physical server, virtual machine ,or third-party format to convert.
2. Specify the destination format, virtual machine name, and location for the new virtual machine to be created.
3. Create or convert to destination virtual machine and configure the virtual machine.

VMware Converter achieves faster conversion speeds through the use of sector-based copying (rather than the file-level copying in other products).

VMware Converter first takes a snapshot of the source machine before migrating the data, resulting in fewer failed conversions and no downtime on the source server. VMware Converter communicates directly with the operating system running on the source physical machine for hot-cloning these machines without any downtime and as a result has no direct hardware-level dependencies.

Secondary (Failover) Server — Double-Take Installation

Once the virtual machine has been created, Double-Take must be installed. If the virtual machine was created using VMware Converter, Double-Take installation may not be required. Complete Double-Take installation details are available in the *Double-Take Getting Started Guide*.

If the virtual machine is to be used as a recovery server for Double-Take Server Recovery Option recovery, you do not need to install Double-Take, because it will be installed during the recovery process.

Configuring Double-Take Replication

Once Double-Take is installed on the source machines and target virtual machines, you need to create a Double-Take replication set to protect the appropriate production data. Because the data that must be selected for replication depends on the application and services running on the server, the files that must be selected depend on the particular application. Double-Take application notes are available for most major applications that can be protected with Double-Take.

Once the data for replication has been selected the replication set needs to be connected to the target virtual machine. This can be done with either the Connection Wizard or the Connection Manager within the Double-Take Management Console.

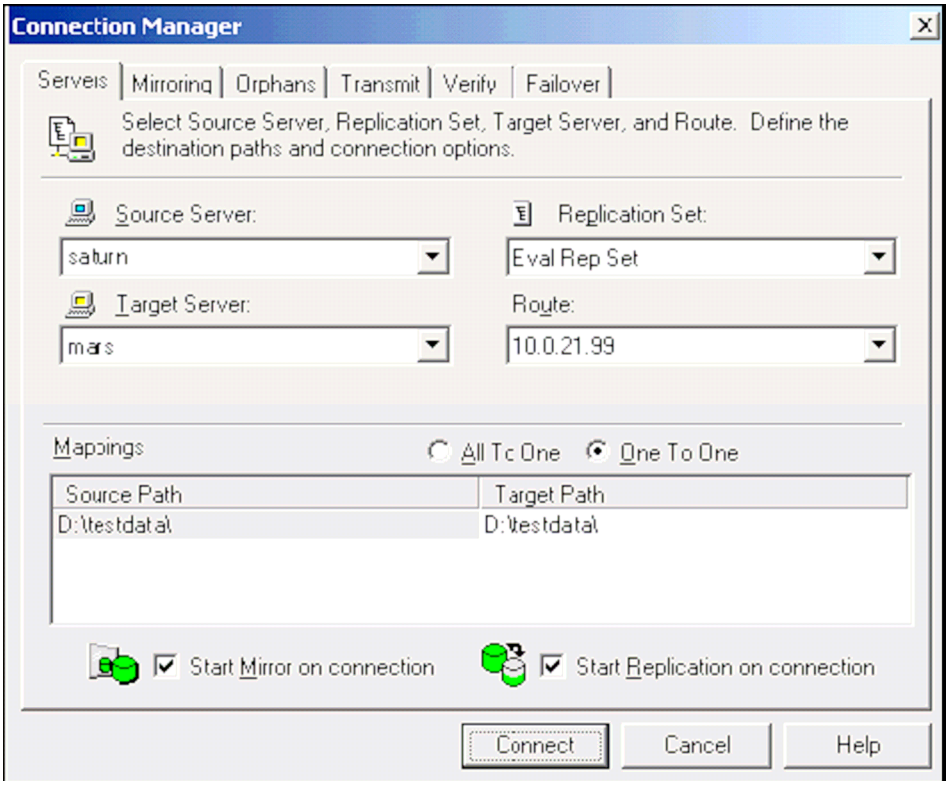

Figure 6 - Double-Take Connection Manager dialog box

If you are protecting servers using the Double-Take SRO for system state protection, use the Connection Wizard in the Double-Take Management Console. The Connection Wizard automatically selects all of the files required for full system protection, although you have the option of selecting which data volumes to protect.

For regular Double-Take replication, or for SRO protection, after the connection is made, replication begins immediately. A mirror (or synchronization) operation is also started after you make a new connection. The mirror process performs an initial synchronization of the data to ensure target data integrity. The target virtual machine data is not considered to be in a good state until the mirror has completed and mirror status is idle. Real-time replication occurs during and after the mirror to ensure all changes are replicated to the target virtual machine. The status of the connection can be viewed in the Double-Take Management Console.

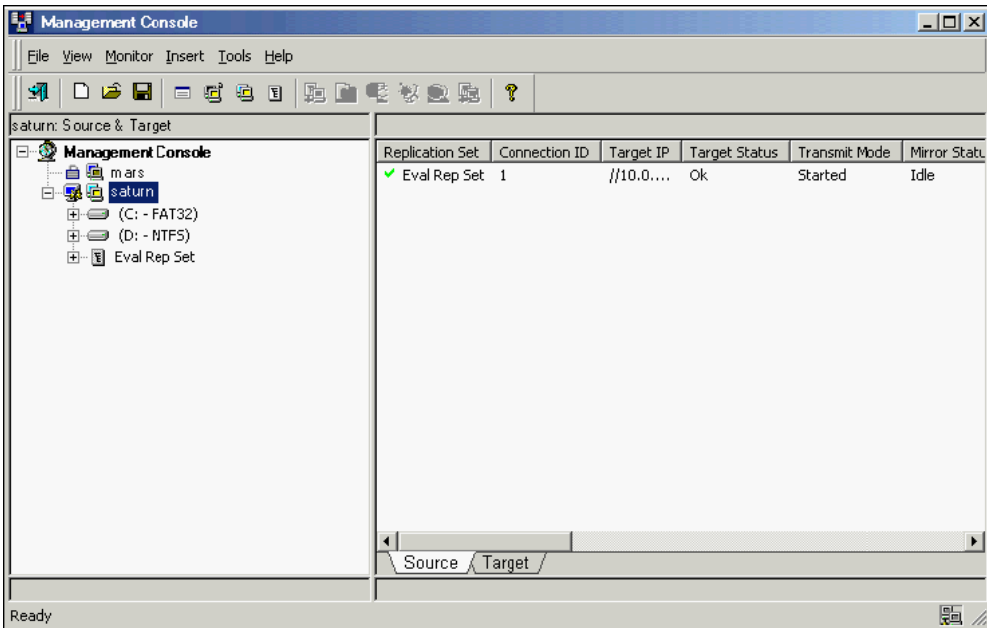

Figure 7 – Example of Double-Take mirror status at Idle

Configuring Double-Take Failover for the Production Server

Double-Take has the ability to monitor the Double-Take source machine and, in the event of a failure, assume the Double-Take source server's network name and IP address on the secondary system. By doing this, Double-Take can provide high availability for many different applications and types of data.

Double-Take Server Recovery Option is used for backup and recovery and does failover if not used with SRO protection.

Failover monitoring is configured using the Double-Take Failover Control Center. The target virtual machine can be configured to monitor one or more source IP addresses, and the failover monitor can be configured to do several types of failover including hostname and IP address failover. The failover monitor can also be configured to execute scripts during the failover and failback processes to perform additional operating system or application operations.

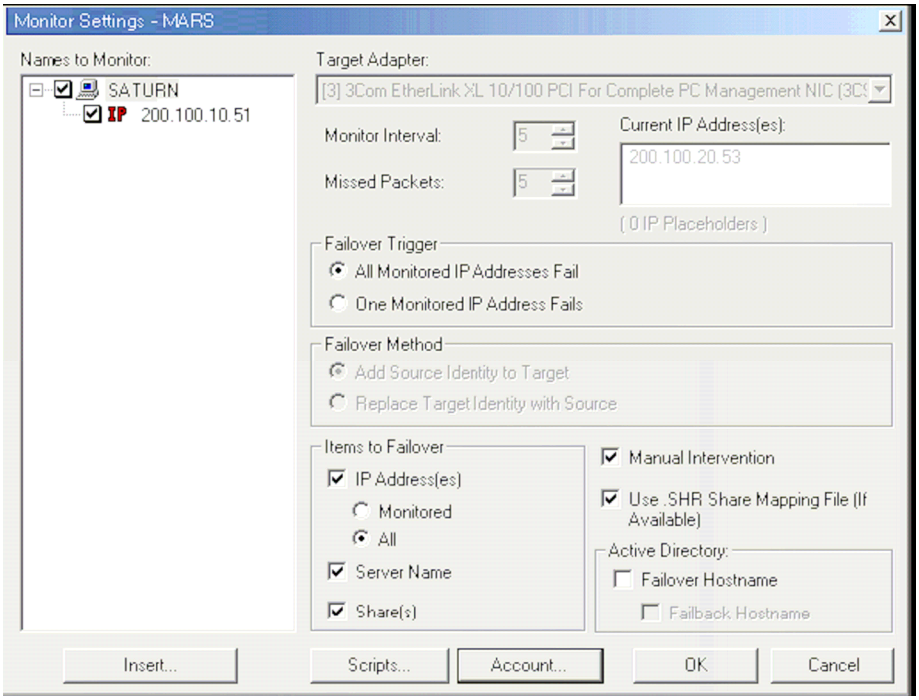

Figure 8 – Example of Double-Take Failover Control Center

In the event of a source machine failure, your target virtual machine will be ready to stand in for your source. For more information about configuring Double-Take failover monitoring for a production server's applications or services, refer to the *Double-Take User's Guide* or to one of the many available application-specific Double-Take application notes.

Restoring Data From a Double-Take Target

If your source experiences a failure, such as a power, network, or disk failure, your target virtual machine stands in for the source while you resolve the production machine's issues. During the source machine downtime, data is updated on the target virtual machine by end-users. When your source machine is ready to come back online, the data on the source machine is no longer current and must be updated with the new data from the target virtual machine. Verify that your source machine is not connected to the network. If it is, disconnect it. Restoration is performed using the Restoration Manager in the Double-Take Management Console. Restoration is normally done during a scheduled maintenance period when users are not connected to the server and no data is changing. Online restoration options are available for certain applications using the Double-Take Application Manager.

When you are prepared to restore the data to the source machine, perform failback on the target virtual machine using the Failover Control Center, then bring the source machine back onto the network. Once the source machine is back online, restoration is initiated from the Restoration Manager and the restore-connection status can be monitored from the Double-Take Management Console.

After your data is restored to your source machine, you can reconnect the replication set to the target virtual machine to resume replication. If you selected to continue failover monitoring during failback, the target is again available to stand in for the source in the event of a failure. If not, you may re-create your failover monitor in the Failover Control Center.

Recovering Systems Using the Double-Take Server Recovery Option

If your source machine experiences a failure, you may recover the source system image from the SRO image server to the original source server (if recoverable) or to a virtual machine. Recovery is performed through the Double-Take Recovery Manager in only a few easy steps.

After the recovery server has been rebooted, the recovery server has the identity of the original source server, all of the applications from the original source server, and all of the data (unless you deselected any data during configuration). Users or processes that were connected to the original source server can now connect to the recovery server without any configuration changes. The recovery server has taken over all of the roles of the original source server.

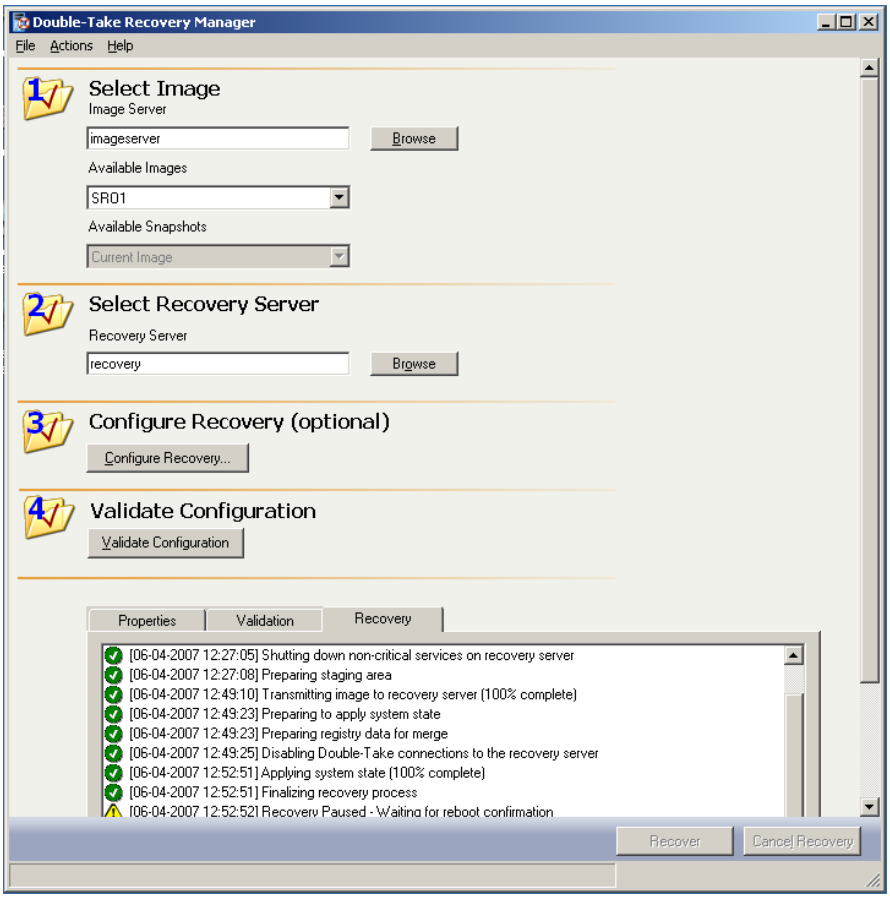

Figure 9 – Example of SRO Recovery Manager console

Summary

Double-Take Software and VMware have combined efforts to offer some of the industry's most advanced disaster recovery and application availability solutions while reducing the cost and complexity of the IT infrastructure. Two of the most trusted technology leaders in the industry, Double-Take and VMware provide cost-saving solutions to consolidate and reduce hardware requirements. Businesses can have complete confidence in the unparalleled data protection, availability, and recoverability these joint solutions provide while simplifying overall IT management.

As demonstrated in this white paper, leveraging real-time replication and application availability for disaster recovery using Double-Take is a cost-effective way to ensure that your RPO and RTO goals for a business-critical application are adequately met. By combining this protection with virtualization technology, such as VMware ESX Server, you can achieve additional flexibility and cost efficiency without sacrificing capabilities. Virtual machines powered by VMware ESX Server and used to protect production servers using Double-Take replication and failover provide an ideal solution for disaster recovery and local and remote high availability.

About Double-Take Software

Headquartered in Southborough, Massachusetts, Double-Take Software (Nasdaq: DBTK) is a leading provider of affordable software for recoverability, including continuous data replication, application availability and system state protection. Double-Take Software products and services enable customers to protect and recover business-critical data and applications such as Microsoft Exchange, Microsoft SQL Server, and Microsoft SharePoint in both physical and virtual environments. With its unparalleled partner programs, technical support, and professional services, Double-Take Software is the solution of choice for more than 10,000 customers worldwide, from SMEs to the Fortune 500. Information about Double-Take Software's products and services can be found at www.doubletake.com.

About VMware, Inc.

VMware is the global leader in virtual infrastructure software for industry-standard systems. The world's largest companies use VMware solutions to simplify their IT, fully leverage their existing computing investments, and respond faster to changing business demands. VMware is based in Palo Alto, California. For more information, visit www.vmware.com or call 650-475-5000.

Revision: 20070730 Item: WP-026-ISV-01-01

VMware, Inc. 3401 Hillview Ave. Palo Alto CA 94304 USA Tel 650-475-5000 Fax 650-475-5001 www.vmware.com

vmware®

Double-Take Software™