DISASTER RECOVERY PLANNING WITH VIRTUALIZATION TECHNOLOGIES

Business Value Whitepaper
Double-Take Software, Inc.
Published: October 2009

Virtualization and Recoverability

Business continuity is not just a good business practice - it can mean success or failure if data and applications on a production server are lost. Disaster recovery planning ensures organizations have the capability to continue essential functions across a wide range of situations that could disrupt normal operations. High availability is the cornerstone for most business continuity plans and is one of the most compelling reasons for evaluating and deploying data protection solutions. However, traditional data protection strategies focus on just the data and not the application.

IT departments design the organization's infrastructure with continuity of business operations in mind. However, most organizations aren't doing enough to protect mission-critical data, applications and systems from unexpected disruption and potential loss -- volatilities, such as viruses, power outages, natural disasters, corruption, human error and media failures can't always be prevented. Environments today are characterized by rapid data growth, complexity, stringent business requirements and the increasing government regulations, making it difficult for organizations to get their arms around their data protection strategies. In many cases, the focus is on just protecting data - not necessarily on recovering it. And when there is a focus on recovery, it usually involves just making data available to an application.

Double-Take Software promotes different approach - one that goes beyond recovery. Our concept of "recoverability" involves layers of protection that not only mitigate the risk of data loss, but, importantly, maintain the health and uptime of systems and applications. Although the distinction is slight, the results are dramatic. By preparing for, preventing and minimizing the impact of a catastrophic event, such as a natural disaster, hardware failure, system corruption or operational error, companies gain in productivity (spend significantly less time recovering from a negative event), customer satisfaction (through its ability to meet recovery objectives), recovery cost savings and more.

We know we know we can't prevent disasters, but we can help end-users recover quickly and overcome the effects of the disaster. With no downtime, no loss of important components and, importantly, no hassle. We distinguish this level of "recoverability" from the basic "recovery" most data protection solutions provide. Double-Take Software has a broad spectrum of solutions that are proof points to this concept of "Recoverability".

The Trouble with Tape

Tape backup can provide for the long-term archival needs of the virtual servers; however tape cannot provide the level of recoverability required for critical business applications. Rebuilding one application from tape can be a difficult and lengthy process. Recovering four or more applications at the same time from tape to rebuild one physical server will result in an excessive period of downtime, likely more than the business can afford.

Organizations may not understand how vulnerable their data and business remain to disaster - even after they've made a huge up-front and ongoing investment in tape-based disaster recovery. It's estimated that as

many as 20% of routine nightly backups fail to capture all data and 40% of recoveries from a tape fail. This is a significant concern for corporations that are regulated as they can face the risk of being out of compliance if they cannot produce required data when they need it.

Tape backup also places limits on your recovery point objective (RPO), the point in time to which you can recover your systems should disaster strike. Periodic tape backup guarantees hours of lost data in the event of a disaster. Suppose, for example, that a critical system fails anytime today; the best you can do is recover to yesterday's data, which will be at least twelve hours old. The later in the day disaster strikes, the older the data from which you'll recover. In addition, recovering from a disaster, any data not backed up is lost for good - unless you recreate it.

The cost of permanently lost data is high and includes the cost of the revenue that the data represents, the business value you can extract from it, and the cost to recreate it. Consider:

- How much money would your business lose if you lost all your transaction data for the last twelve hours, or even the last ten minutes?

- What is the value of the knowledge contained in your company's last twelve hours worth of e-mails and e-mail attachments? What would it cost to have your engineers recreate the last twelve hours worth of original or edited CAD/CAM drawings?

In *The Cost of Lost Data*, a Pepperdine University report, before the advent of Sarbanes-Oxley - Dr. David Smith estimates the average cost of irrecoverably lost data at more than $10,000 per megabyte lost.[1] But if the data lost is business transaction data or data that's especially expensive to reproduce and key to your company's disaster recovery plan, your costs could be much, much higher.

The Cost of Downtime

When a large-scale disaster strikes, with tape backup you're out of business until you can restore your systems and your data from your tapes. This kind of restoration takes a minimum of several hours, and can easily take days or even weeks.

Gartner Group estimates that the average cost of network downtime for larger corporations is $42,000 per hour; Contingency Planning Research pegs the average hourly downtime costs for many businesses at roughly $18,000. The key to a successful disaster recovery plan is to focus not just on the data (RPO) but also on the applications that end users run to gain access to that data. Recovery Time Objective (RTO) is generally defined as the amount of time it takes to regain access to business-critical data. Solutions like tape backup, which have an RTO of hours or days, don't provide the level of recoverability that most companies require.

Full system rebuilds and tape restores are unacceptable recovery methods for meeting the RPO/RTO of mission critical applications and leave organizations vulnerable to lengthy recovery times and potential data loss. Architecting for maximum availability throughout various types of outages presents a challenge that can be solved through a combination of real-time replication, application availability and virtualization technology. Using virtualization along with high availability storage solutions and data protection software like Double-Take Availability can help businesses economize on equipment, bandwidth, and budget dollars, while allowing them to architect to a RTO of minutes with low risk of data loss (RPO).

How Virtualization Can Help

Virtual server technologies provide businesses and IT departments with the ability to do more with less, enabling the consolidation of data and applications onto a single server. The result is reduced costs, simplified IT management, and minimized space requirements. As projects for server consolidation and server rationalization are realized, the need to protect these virtualized systems is paramount as they are running multiple business-critical applications,

[1] Smith, David M. "The Cost of Lost Data." Graziadio Business Report Vol. 6 No. 3

requiring a higher level of protection. In fact, consider the impact of a failure of a virtualized system. Instead of incurring downtime of just one application, you will incur downtime of multiple applications. Virtual machine technologies require advanced disaster recovery and availability solutions that provide protection against data loss and downtime for the entire environment.

Even for companies which are not implementing virtualization technologies to support their production infrastructure, the same benefits of virtualization can assist their disaster recovery efforts. 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 like Double-Take Availability which are based on replication and failover usually 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 either is not recommended or not possible to failover multiple physical workloads to a single OS instance running on standard server hardware. This usually requires organizations to 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 guest operating system (Guest OS) 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 one another 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. Though many factors affect server consolidation savings, enterprises have reported up to 70% cost savings with virtual machines instead of physical servers.

The success of business operations at an alternate facility is dependent upon the availability and redundancy of critical communications systems to support connectivity to internal organizations, other businesses, critical customers and the public. Double-Take Availability provides remote availability to essential applications and services running on any Windows® Server operating system and ensures failover to a standby server in the event of a local failure whether either system is running on physical or virtualized hardware.

Protection Strategies for Virtualized Environments

Businesses are becoming increasingly dependent on continuous access to stored data and as a result, storage usage is growing at an unprecedented rate. As the number of mission-critical servers and storage resources grow, so does the importance of protecting against service interruptions that can threaten an organization's ability to provide access to key data.

There are a number of strategies that can be employed to protect important data, and each has strengths and weaknesses. The most common method of storage protection is also the oldest: backing up to and restoring from magnetic tape. This method has been around for almost forty years and is still the bedrock of most recovery strategies. The cost per megabyte for tape storage is low; it's easy to move tapes to secure offsite storage, and the technology continues to scale well for many applications. However, tape backups have limitations, such as the amount of time required to back up and restore large volumes of data, the accompanying latency between when the data was protected and when the loss occurs, and the security involved in moving tapes to offsite storage. Accordingly, much attention is being focused on replication-based technologies.

Replication-Based Technologies - An Overview

Replication-based technologies offer the promise of capturing a data set at a particular point in time with minimal overhead required to capture the data or to restore it later. There are four main methods of interest in today's storage environments:

- **Whole-file replication copies files in their entirety.** This is normally done as part of a scheduled or batch process since files copied while their owning applications are open will not be copied properly. The most prevalent use of this technology is for login scripts or other files that don't change frequently.

- **Application replication copies a specific application's data.** The implementation method (and general usefulness) of this method varies dramatically based on the feature set of the application, the demands of the application and the way in which replication is implemented. This model is almost exclusively implemented for database-type applications

- **Hardware replication copies data from one logical volume to another and copying is typically done by the storage unit controller.** Normally, replication occurs when data is written to the original volume. The controller writes the same data to the original volume and the replication target at the same time. This replication is usually synchronous, meaning that the I/O operation isn't considered complete until the data has been written to all destination volumes. Hardware replication is most often performed between storage devices attached to a single storage controller, making it poorly suited to replicating data over long distances. Most hardware replication is built out of SAN-type storage or proprietary NAS filers.

- **Software replication integrates with the Windows operating system to copy data by capturing file changes as they pass to the file system.** The copied changes are queued and sent to a second server while the original file operation is processed normally without impact to application performance. Protected volumes may be on the same server, separate servers on a LAN, connected via storage-area network (SAN), or across a wide-area network. As long as the network infrastructure being used can accommodate the rate of data change, there is no restriction on the distance between source and target. The result is cost-effective data protection.

To best understand how to protect data, it's important to consider what the data is being protected *from*. Evaluating the usefulness of replication for particular conditions requires us to examine four separate scenarios in which replication might lead to better business continuity:

- **Loss of a single resource** - In this scenario, a single important resource fails or is interrupted. For example, losing the web server that end-users use for product ordering would cripple any business that depends on electronic procurement. Likewise, many businesses would be seriously affected by the loss of one of their primary e-mail servers. For these cases, some organizations will investigate fault-tolerant architectures, don't invest in fault-tolerance technology for file and print servers-even though the failure of a single file server may simultaneously prevent several departments' employees from accessing their data. Planning for this case usually revolves around providing improved availability and failover for the production resources.

- **Loss of an entire facility** - In this scenario, entire facilities, and all of their resources, are unavailable. This can happen as the result of natural disasters, extended power outages, failure of the facility's environmental conditioning systems, and persistent loss of communications or terrorist acts. For many businesses, the normal response to the loss of a facility is to initiate a disaster recovery plan and resume operations at another physical site.

- **Loss of user data files** - This unfortunately common scenario involves the accidental or intentional loss of important data files. The most common mitigation is to restore the lost data from a backup, but this normally involves going back to the previous RPO - often with data loss.

- **Planned outages for maintenance or migration** - The goal of planned maintenance or migrations is usually to restore or repair service in a way that's transparent to the end users

Double-Take Availability Data Replication and Application Failover

Double-Take Availability is a real-time data replication and failover software product. It augments existing data protection strategies and more traditional backup technologies like disk-based backup by reducing downtime and data loss, and it provides these services with minimal impact on existing network resources. It allows selection of data sets that must be protected and then replicates, in real-time, that data from a primary machine to a secondary system.

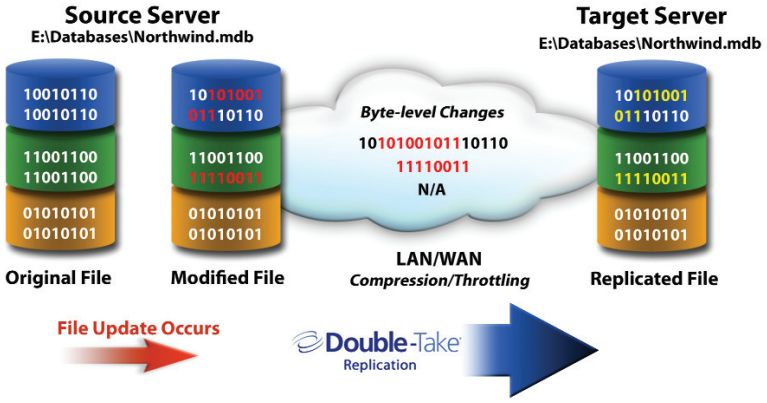

Figure 1 - Double-Take Availability Byte-Level Replication

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

Double-Take Availability

Protecting the data on a virtual machine is crucial and implementing a disaster recovery solution can be simplified using virtual technologies and result in a much lower overall cost. Double-Take Availability delivers enterprise-class disaster recovery and high availability solutions for cutting-edge virtual environments. Double-Take Availability is the industry-leading replication and availability solution that offers local and long-distance protection against most any outage, from a failed power supply to a regional disaster. Double-Take Availability is the industry's first disaster recovery and availability solution designed for virtualized server environments.

Figure 2 - VM to VM Protection Using Double-Take Availability

Double-Take Availability deploys easily within each virtual machine to protect individual workloads independently. Options include replicating protected data either to a single physical server or to a virtualized server. Double-Take Availability can maintain availability of the virtual servers by providing failover monitoring and client redirection to a standby virtual server to minimize lost productivity (Figure 2). The patented Double-Take Availability architecture provides the flexibility you need to design and implement the solution that's right for your business.

With physical production systems, most applications typically require their own dedicated server, providing high availability or failover protection to these applications requires a dedicated target server. With Double-Take Availability, each workload running on a physical server can replicate and failover to a dedicated virtual machine on a single physical server, eliminating the need to deploy and manage multiple physical servers at the recovery site (Figure 3). In a disaster or outage, Double-Take Availability will automatically start the appropriate application services within the virtual machines and seamlessly redirect via automatic updates to DNS or Microsoft Active Directory.

The result is a very high level of application protection and availability with a much lower TCO as virtual machine solutions offer improved efficiency and lower IT costs. Double-Take Availability was developed specifically for these virtualized environments to offer businesses advanced solutions for protecting these highly valuable servers.

Protecting vSphere

For virtual machines running in production on VMware ESX Servers, Double-Take Availability provides the ability to recover those virtual machines, in their entirety, on another ESX Server locally or at a disaster recovery site. Working at the ESX host level, it leverages VMware's snapshot capabilities to regularly replicate changes to protected VMs. In the event of an outage, the replicated virtual machine can be started on a second ESX server with the most recent data.

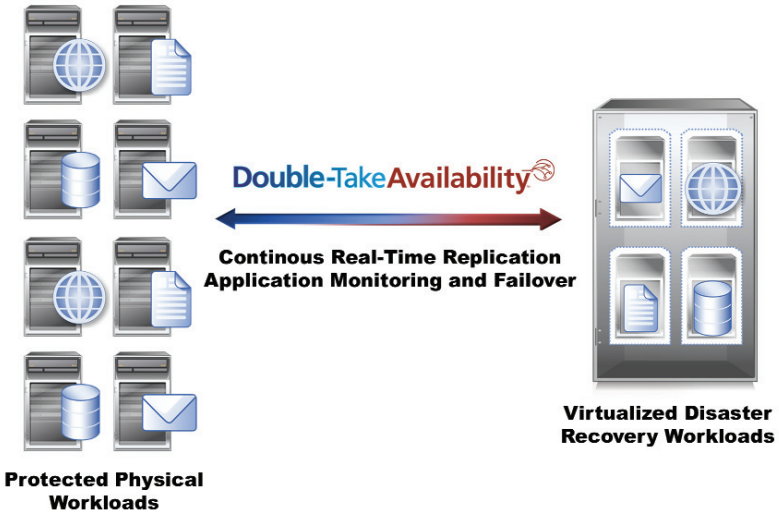

Figure 3 - P2V Protection Using Double-Take Availability

Double-Take Availability greatly reduces downtime and data loss in virtual environments, making it an ideal component for customers implementing cost-effective disaster recovery and business continuity plans. It is a comprehensive solution that protects ESX-based virtual machines and provides administrators centralized management capabilities for ease of use and lower total cost of ownership. Double-Take Availability replicates the entire virtual machine: the OS, applications and data, allowing you to replicate everything to any location and resume operations with minimal downtime.

Using VMware APIs for virtual machine snapshot functionality, Double-Take Availability captures changes regularly, keeping secondary virtual disk(s) up-to-date and ready for recovery at any time. In an outage, the replicated virtual machine can be started on a second ESX server with the most recent data. By leveraging VMware-supported APIs, protected data is crash-consistent and time-coherent across all virtual disks on a protected virtual machine. Further, by replicating entire VMs, any Guest OS supported by VMware (including Linux, UNIX and other non- Microsoft® Windows® OS's) can be protected on a secondary ESX Server with no reconfiguration at recovery time (Figure 4).

Double-Take Availability captures changes regularly, keeping the target virtual disk(s) up to date and ready for recovery at any time. In a disaster or outage, the replicated virtual machine can be started on a second ESX server with the most recent data. Double-Take Availability has a flexible licensing model - there is no per-host charge; instead, licenses protect groups of virtual machines, regardless of the host or configuration.

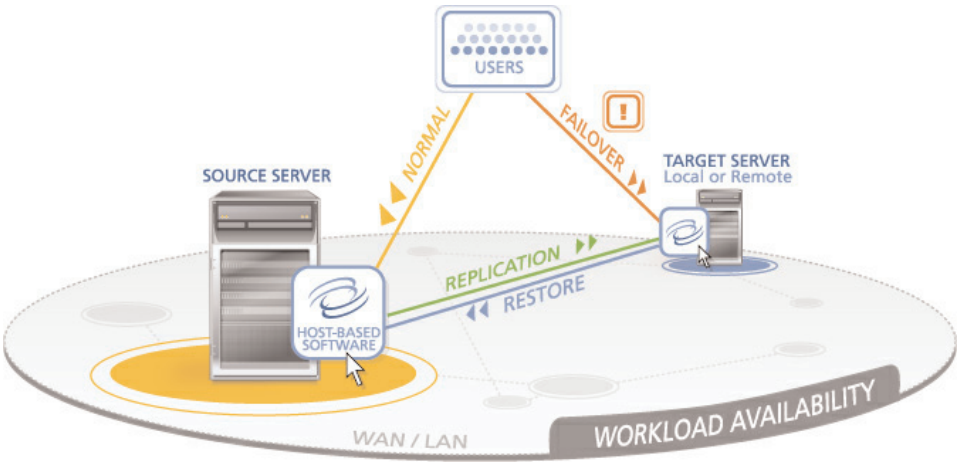

Figure 4 - Overview of Double-Take Availability

Protecting Hyper-V

Double-Take Availability has been protecting host-based servers and applications for well over a decade. This provides robust and time-tested features that allow you to schedule and shape bandwidth traffic, which lets you match business requirements to your recovery point objectives. In addition, you can choose your recovery time objectives using automatic or manual failover methods that bring your applications online and let you continue working seamlessly. The Double-Take Availability replication engine has been proven thousands of times in the largest environments to provide a scalable solution that can match your production performance requirements.

Using Double-Take Availability lets you build high availability and disaster recovery solutions for your Hyper-V infrastructure without requiring you to wait hours to restore your backups before you can boot them. This dramatically improves your RTO to minutes or seconds from hours and days. Double-Take Availability doesn't require a SAN infrastructure, so you can use it with the hardware infrastructure that you already own, which further reduces your total cost of ownership while greatly improving your recoverability. You can also mix different storage solutions without regard to product line, vendor or even the underlying storage technology or configuration.

Double-Take Availability was designed to integrate with the Windows 2008 host operating system (Parent partition) and protect Hyper-V virtual machines in real time. This includes the virtual machine virtual hard drive (VHD) files and their associated configuration settings. Double-Take Availability replicates any changes to the virtual machine as soon as it occurs which provides you with a complete protection solution of your production machines. If a virtual machine fails for any reason, then it can be restarted on the target Hyper-V host and continue processing as normal without having to wait hours, as is required when restoring from a backup. Thus recovery time objectives of Double-Take Availability protected virtual machines are measured in minutes, or about as fast as the virtual machine can boot.

Double-Take Availability integrates with Hyper-V to provide discovery of each virtual machine and its associated system resources. Once virtual machines are discovered and cataloged, you can select each individual virtual machine that you would like to protect and the target location that you would like to replicate to. The replica location is another server running Hyper-V that can be physically located in the same data center or in an off-site location across country for geographic redundancy. Another key benefit of Double-Take Availability is that it doesn't require a SAN to provide virtual machine recoverability services. This eliminates storage as a single-point of failure in other Hyper-V protection solutions and lets you mix storage of different types of vendors to meet your recoverability objectives.

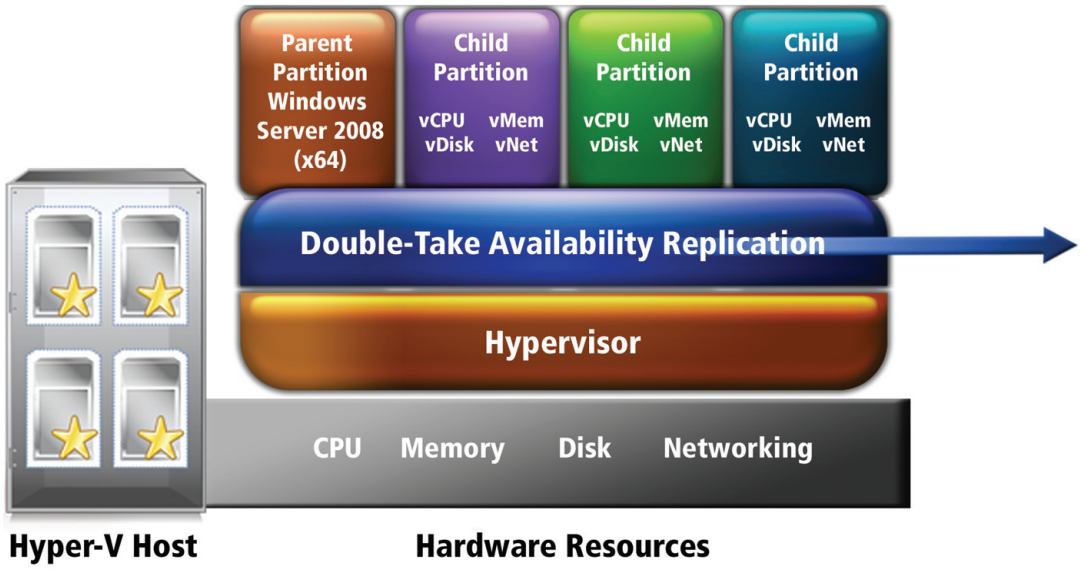

Summary

Double-Take Software and leading virtualization industry companies have combined efforts to offer some of the most advanced disaster recovery and application availability solutions while reducing the cost and complexity of the IT infrastructure. Businesses can have complete confidence in the unparalleled data protection, availability, and recoverability of these joint solutions provide while simplifying overall IT management.

As you have seen in this whitepaper, leveraging real-time replication and application availability for disaster recovery using solutions from Double-Take Software is a cost-effective way to ensure that a mission-critical application's RPO and RTO goals are adequately met. By combining this protection with virtualization technologies, organizations can achieve additional flexibility and cost efficiency while not sacrificing capabilities.