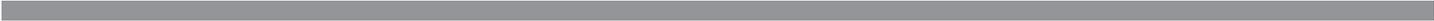

Real-time Protection for Microsoft Hyper-V



Introduction

Computer virtualization has come a long way in a very short time, triggered primarily by the rapid rate of customer adoption. Moving resources to virtual platforms have allowed IT departments to consolidate and better utilize computer resources. However, many methods for system protection that were tried and true in the physical world haven't lived up to their expectations when tasked with protecting virtual machines. Protecting virtual machines is not always as simple as backing up the virtual disk files and special care must be taken.

Furthermore, consolidating systems to the same physical hardware has increased the amount of exposure to system failure. The concern for system outages typically entails factors such as application failure, operating system failure, hardware or storage failure and site failure. Using virtualization technology adds additional failure points, exacerbated by physical failures since they now directly affect multiple systems at the same time. Thus, having a regularly tested availability strategy is drastically growing in importance.

Developing a strategy for virtual machine recovery begins with defining the criticality of your virtual machines. The two primary methods of measuring criticality relate how much data you can afford to lose, called the Recovery Point Objective (RPO), and how quickly the application must be recovered, the Recovery Time Objective (RTO). Using these two primary measures will help you understand your cost of downtime, help define a budget and determine the technology that meets your needs within your budget.

Defining your company's RPO typically begins with examining the current backup schedule of how frequently backup takes place. Since backup is an intrusive process to systems, they are not typically performed more frequently than several hours apart. This means that your backup RPO is measured in hours of system state and data loss that is typically acceptable for very few applications in the modern data center.

Even if backups are performed regularly, it may take much longer to actually restore the virtual machine backups when they're needed the most – and you must take that into consideration when determining your company's RTO. Next, you should determine how long it takes to provision servers, storage, networking resources and virtual machine configurations. These are all major factors that need to take place before your users have access to their applications and data.

Navigating the numerous solutions available on the market may seem daunting at first, but finding the right balance of features and price to meet your RPO and RTO is one of the most critical things that an IT department can do to protect the business. The three primary solution categories are backup, high availability and disaster recovery solutions. Each solution has its own RPO and RTO expectations.

Hyper-V Systems Architecture

Hyper-V™ is a role of Windows Server 2008 which provides a virtualization hypervisor that hosts virtual machines in isolated containers that have their own virtualized CPU, memory, disk and networking resources.

Hyper-V virtual machines are also called 'partitions'. Each Hyper-V server has a single Parent partition that has direct access to the machine's hardware resources and any number of Child partitions. Child partitions do not have direct access to the physical machine resources, but instead receive virtualized views of the resources allocated by the hypervisor to each Child.

Virtualized disks presented to Child partitions reside as files on the host machine's file system in Virtual Hard Disk (VHD) format. These VHD files can physically reside on DAS, NAS, Fiber Channel or iSCSI SAN storage devices mounted to the host machine. However, protecting Hyper-V virtual machines requires additional features beyond performing backup of the VHD files themselves in order to properly protect the machines in their running state.

Performing Hyper-V Backup

While Hyper-V builds on the mature foundation of the Windows Server 2008 operating system, new tools and technologies are required to properly backup and restore Hyper-V virtual machines. A backup from within the virtual machine is performed just as it would be from a physical server. However, performing backup of the virtual machine disk files from the host presents several challenges that do not allow a simple file-based backup process to preserve the data consistency of the virtual system. Microsoft addressed the concern for host-based virtual machine backup by creating guidelines for how to build and backup a Hyper-V environment to ensure system state and data consistency.

Point-in-time, host-based backup of running virtual machines and their applications can be accomplished with the Microsoft VSS (Volume Shadow Services) feature and a VSS-aware backup application. When the backup application wants to perform a backup it notifies VSS; which notifies any VSS registered applications that a backup is about to be performed so those application can put their data into a consistent state. Hyper-V virtual machines running supported operating systems with VSS that also have Hyper-V integration components can preserve their state just like other VSS-aware applications when a backup is performed since VSS notifications are passed into the virtual machine via Integration Services.

However, many operating systems do not yet have access to Hyper-V Integration Services and cannot participate in the native backup process, so alternative solutions are required. For those that can participate, native backup may not prove good enough to protect production systems because of the long RPO and RTO inherent with backup technologies. Real-time protection solutions are required to provide instant protection and recoverability for Hyper-V virtual machines.

Hyper-V Availability and Disaster Recovery

Microsoft® Windows Server 2008 Failover Cluster features can be employed to create high-availability protection for Hyper-V implementations. This enables failover of Child partitions between nodes of a cluster for both planned and unplanned recovery. There are two types of planned failovers in Failover Clustering, Quick Migration and Live Migration.

When performing a Quick Migration of Child partitions the failover clustering software triggers the affected virtual machines to perform a Save State, which puts the machine into a saved state and copies its memory contents to disk resources shared by the cluster. Once Save State completes, then another node gains access to the virtual machine's shared disk and resume the Child partition by copying the saved memory back into RAM and continuing where it left off. This process typically takes one to two minutes to complete and provides a graceful method of redistributing virtual machines throughout the available computing resources when performing maintenance on a hosting node.

A Live Migration can be performed on a cluster utilizing Clustered Shared Volumes (CSVs). The process is similar to a Quick Migration but happens much faster as there is not a delay in the second node gaining control of the disks in order to obtain the copy of the saved state. A Live Migration typically occurs in seconds rather than minutes.

Unplanned failover occurs when a catastrophic event happens without warning - such as a host node crash. This type of sudden failure doesn't provide the luxury of saving the machine memory state to disk before restarting on a different node. The failover process works the same as during a planned failover, but the virtual machine must boot from a cold state and it should be checked to ensure the virtual machine and its applications recover properly since their state was in flux during the crash.

The configurations presented above are very basic in their implementation and management. There are numerous configurations supported by Microsoft for host-based failover clustering that should be consulted closely to avoid single points of failure and allow for future scalability. A challenge posed by Windows native failover clustering solutions is the lack of built-in protection for the Hyper-V virtual machine configuration files. One Hyper-V deployment architecture supports storing the virtual machine configuration files outside of the cluster entirely on a separate server's shared folder. However, it is still the responsibility of the system administrator to ensure adequate failover protection for the shared folder. Thus Hyper-V failover clustering has management challenges that require skilled and experienced system administrators to avoid self-inflicted downtime.

Windows Server 2008 Failover Clustering also introduces some new features that provide much better support for geographically distributed clustering. However, these solutions require third-party software and/or hardware support to facilitate the virtual machine configuration and disk replication process between sites. This further increases the management requirements for failover cluster protection of virtual machines using Quick Migration or Live Migration. You should also take care to fully evaluate the scalability claims made by replication technology vendors to ensure your systems are adequately protected under actual production performance loads.

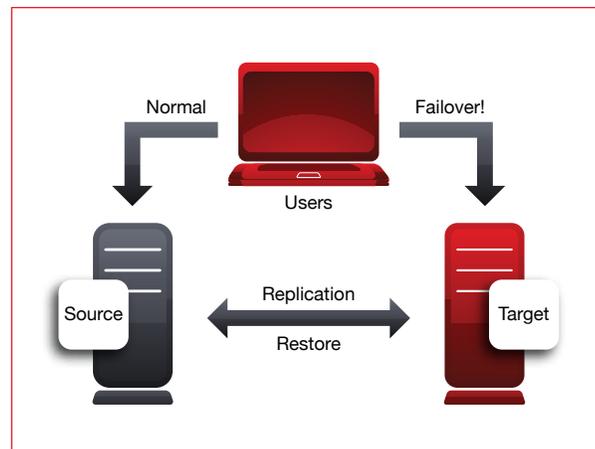
How Double-Take Availability Works

Performing a backup takes time and impacts applications negatively while they are putting themselves into a quiescent state, therefore backups are performed at intervals measured in hours and sometimes days.. To properly preserve data change that occurs between backups, a real-time protection solution is required.

Double-Take Availability is built on award-winning replication technology that 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, Double-Take Availability lets you 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 Infrastructure

Double-Take Availability was designed to integrate with the Windows Server 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 Hyper-V host that can be 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.

Use Cases

High Availability for Hyper-V

A first line of defense for protecting your virtual machines is to build high availability features into your Hyper-V infrastructure. If you don't have the ability to construct a failover cluster using Windows Server 2008 native features because you lack a SAN or enough cluster experience to feel comfortable, then Double-Take Availability provides a complete high availability solution for your environment. High availability solutions typically reside on the same local area network infrastructure for immediate resumption of failed services.

Double-Take Availability is installed on your existing Hyper-V hosts without any significant impact to their operation. After discovery, you can map the virtual machines that you want to protect with the target Hyper-V host. You can configure automatic failover when a failure is detected to immediately failover virtual machine services to the target host or you can manually failover virtual machines at any time, before or after a failure.

Remote Availability for Hyper-V

Server virtualization is often ideal for remote and branch offices although providing disaster recovery or high availability solutions at these sites can be complex at the least. Remote availability differs from high availability because there are different expectations for recovery times and accessibility because services are relocated across geographical distances. Double-Take Availability is optimized for WAN replication and can protect virtual machines on Hyper-V hosts from remote and branch offices directly to a central disaster recovery site or datacenter. When sites are geographically dispersed the overall expected failure rate is reduced because it is unlikely that a failure at one site will affect another site. Using a remote availability solution like Double-Take Availability allows you to fully subscribe or under-subscribe your DR infrastructure to failover just one or two sites at a time or all sites depending on your need.

Summary

Double-Take Availability builds upon and extends the platform provided by Microsoft Windows Server 2008 Hyper-V to provide numerous options for protection and recoverability of your virtual machine environment. It provides real-time data and system state protection that reduce the total cost of ownership for high availability, disaster recovery and remote office recoverability projects. Double-Take Availability integrates into your existing server and storage infrastructure to provide these recoverability solutions. This lets your IT staff spend more time solving business problems instead of struggling to overcome technical limitations all at an affordable price point.

Easy. Affordable. Innovative. Vision Solutions.

With over 20,000 customers globally, Vision Solutions is one of the industry's largest providers of high availability, disaster recovery and data management solutions for Windows, IBM i (i5/OS), AIX, Linux and Cloud environments. Vision's MIMIX, Double-Take and iTERA brands keep business-critical information continuously protected and available. With an emphasis on affordability and ease-of-use, Vision products and services help customers achieve their IT protection and recovery goals, which in-turn improves profitability, productivity, regulation compliance, customer satisfaction and quality of life.

Vision Solutions oversees a global partner network that includes IBM, HP, Microsoft, VMware, Dell and hundreds of resellers and system integrators. Privately held by Thoma Bravo, Inc., Vision Solutions is headquartered in Irvine, California with development, support and sales offices worldwide.

For more information call 801-799-0300 or toll free at 800-957-4511, or visit visionsolutions.com.



15300 Barranca Parkway
Irvine, CA 92618
800-957-4511
801-799-0300
visionsolutions.com

Double-Take[®] AVAILABILITY[™]
Double-Take[®] MOVE[™]

© Copyright 2010, Vision Solutions, Inc. All rights reserved. IBM and Power Systems are trademarks of International Business Machines Corporation. Windows is a registered trademark of Microsoft Corporation. Linux is a registered trademark of Linus Torvalds.