

# Real-time Protection for Hyper-V

## Real-Time Protection for Hyper-V

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. Many maintain the false belief that protecting virtual machines is as simple as backing up the virtual disk files. This is true - and special care must be taken when performing backup of virtual machines.

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 as shown in the diagram.

## Hyper-V Systems Architecture

Hyper-V™ is a role of Windows 2008 and an integral part of the operating system platform that works on 64-bit computers. Hyper-V 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.

# Real-time Protection for Hyper-V

## Performing Hyper-V Backup

While Hyper-V builds on the mature foundation of the Windows 2008 operating system, new tools and technologies are required to properly backup and restore Hyper-V virtual machines. Performing backup from within the virtual machine is accomplished the same as the backup processes performed on their physical machine counterparts. 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 Components.

However, many operating systems do not yet have access to Hyper-V integration component software 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 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, also called Quick Migration, and unplanned recovery.

When performing planned failover 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.

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.


# Real-time Protection for Hyper-V

Windows 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. 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 days and sometimes hours. Even performing snapshots at frequent intervals doesn't help because state must be preserved, and thus the impact is only slightly less than a backup. 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 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.

# Real-time Protection for Hyper-V

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.

## 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 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. By definition, high availability solutions are protection architectures that reside on the same local area network infrastructure that avoids name to IP address and other network mapping changes. These factors provide high availability solutions RTO measured in seconds.

Double-Take Availability is installed on your existing Hyper-V systems 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 server. You can configure automatic failover when a failure is detected to immediately transfer virtual machine services to the target server and continue processing. When you're ready to grow your Hyper-V infrastructure, you can add nodes to your existing Double-Take Availability protection architecture. This provides maximum return on your high availability investment without requiring a significant investment in hardware and training.

## Disaster Recovery for Hyper-V

Using Double-Take Availability to provide off-site disaster recovery of your virtual machines is a process similar to the high availability architecture. The primary differentiation resides in the network architecture and the change in expectations for recoverability. Disaster recovery solutions typically span geographic regions and thus wide area networks that almost always require traversing IP subnets. While there are methods for stretching or re-configuring your network after a failover, it is usually a complex and error-prone process. While Double-Take Availability can certainly work within those network environments, it provides other methods that are less complicated and less error prone.

Expectations for disaster recovery solutions should always be different from high availability even though they may look similar on paper. The recovery time and recovery point objectives for disaster recovery solutions are typically measured in longer intervals (minutes) and it's not uncommon for less critical workload RPO/RTO measurement to be hours or days. Using Double-Take Availability gives your company numerous recovery options that you can match to business requirements. You can even share resources between more than two sites to create hained or meshed distributed recovery architectures to gain maximum value.

When a failure is detected, you are notified through standard network monitoring tools such as Performance Monitor, Event Log, WMI, SNMP and email. This lets the IT staff determine what the problem is and whether a failover is necessary. This avoids false failure detections and thus split brain (having the same application on-line in two locations) which can occur if you were using automatic fail-over and had taken a router off-line for maintenance. Once you decide that a failover is necessary, the Double-Take Availability manual failover process requires a single click. You can also choose to failover individual machines for failure and migration scenarios when you don't want to recover the entire Hyper-V host.

# Real-time Protection for Hyper-V

## Remote Office Recoverability for Hyper-V

Another popular architecture provides remote office recoverability for Hyper-V virtual machines from a centralized location. This solutions architecture works fundamentally like the Double-Take Availability disaster recovery architecture since it spans wide area networks. The primary difference is in the business expectation for failure rates. This lets you over-subscribe your disaster recovery plans for remote sites since most geographically distributed remote office networks are far enough apart that a geographic disturbance won't affect them all simultaneously. Thus, you can further reduce your total cost of ownership by centrally provisioning enough hardware capacity to handle failover of one or two of your remote sites.

## Summary

Double-Take Availability builds upon and extends the platform provided by Microsoft Windows 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.

QUESTIONS?

1-888-674-9495

info@doubletake.com

www.doubletake.com



© Double-Take Software, Inc. All rights reserved. Double-Take, Balance Double-Take Cargo, Double-Take Flex, Double-Take for Hyper-V, Double-Take for Linux, Double-Take Move, Double-Take ShadowCaster, Double-Take for Virtual Systems, GeoCluster, Livewire, netBoot/i, NSI, sanFly, TimeData, TimeSpring, winBoot/i and associated logos are registered trademarks or trademarks of Double-Take Software, Inc. and/or its affiliates and subsidiaries in the United States and/or other countries. Microsoft, Hyper-V, Windows, and the Windows logo are trademarks or registered trademarks of Microsoft Corporation in the United States and/or other countries. Linux is a registered trademark of Linus Torvalds. Red Hat is a registered trademark of Red Hat, Inc. Novell, the Novell logo, the N logo, SUSE are registered trademarks of Novell, Inc. in the United States and other countries. All other trademarks are the property of their respective companies.