



White Paper
Hyper-V* powered by
Intel® Xeon® processors

High Availability with Windows Server* 2008 Hyper-V* Powered by Intel® Xeon® Processors



Table of Contents

Why Plan for High Availability?	3
High Availability with Windows Server 2008 Hyper-V	3
Reduce the possible points of failure	3
Minimize downtime during maintenance	3
Different Ways to Achieve High Availability with Hyper-V	4
Hyper-V Quick Migration	4
Failover clustering	4
Network load balancing across Hyper-V	4
High Availability Design Considerations	4
Processor considerations	4
Hardware requirements	5
Software requirements	5
Planning utilization and failover setup	5
Disaster recovery requirements and geographically dispersed clustering	6
Conclusion	6

Why Plan for High Availability?

High availability is a concept and a terminology that can have different meanings to different people. For the purpose of this document, high availability will be defined as making IT-related services available with minimal downtime, as well as making sure there are enough resources provided to the service to ensure that it can respond to the needs of the projected user population using the services.

Business processes are heavily dependent on IT services. As a result, more and more IT services are considered critical to the everyday functioning of the business. An example of this is the ubiquitous reliance on e-mail. At first, e-mail was a convenience. Then it became a standard productivity tool. Now, to most organizations, it is a mission-critical tool handling the vast majority of business communications. That same can be said for corporate intranets: first introduced as a means to communicate information back to the company, they have evolved to host group workspaces and to allow teams to span the globe.

Because of the increased reliance on these IT services, as well as many others, when one of these services experiences downtime, an organization incurs very real costs. The cost of downtime can be measured in the form of idle worker time, underserved customers, and jeopardized partner relationships. To sum up, all organizations should plan for high availability because in today's world downtime costs money.

Avoiding downtime, therefore, has very real financial rewards. Downtime can come in one of two forms: planned or unplanned downtime; a well-prepared high-availability plan prepares an organization to respond to either type of downtime.

High Availability with Windows Server* 2008 Hyper-V*

Virtualization is a powerful technology that provides inherent high availability benefits. Virtual machine portability allows a level of flexibility that allows for rapid deployment and re-deployment of critical workloads, which greatly contributes to a highly available environment. Furthermore, virtual machine isolation provides the ability to consolidate multiple virtual machines onto a single physical host, reducing the possible points of failure.

Windows Server* 2008 Hyper-V* is the virtualization technology from Microsoft that provides a platform from which to create and manage virtual machines.

Reduce the possible points of failure

With virtual machine consolidation, a single physical server has the potential to host many virtualized servers, in the form of virtual machines. This increases the utilization of the hardware resources, which in turn increases the organization's return on investment. It also reduces the possible points of hardware failure. On the other hand, it now creates a situation where a single hardware failure can affect many critical workloads at once.

To turn this risk into an advantage, organizations should re-invest some of the costs saved by consolidating servers into deploying redundant components for the reduced number of physical servers that are left. Thus, when calculating the ROI of a server consolidation project, the cost of making these physical host servers highly available should always be part of the calculation.

Moreover, it will be necessary to plan for how these physical host servers will be maintained, and what procedures will be followed in the event of physical server failure. Maintenance strategies need to be planned not just for the virtual workloads, but also the physical server.

Minimize downtime during maintenance

When physical servers are brought down for maintenance, there is usually a need for a service interruption. With the 24-hour nature of many organizations, even planned outages have an associated cost; increasingly organizations are looking to eliminate these costs as well.

A well-planned high availability solution can allow for maintenance on both the physical server and the virtual servers it hosts, while minimizing planned downtime. This is critical to achieving the goal of minimal downtime maintenance. It is a good idea to plan for availability while performing maintenance, even if it is not a requirement. It is also important to create a rollback strategy that may include elements of high availability planning. This allows the organization to roll a system back to a point just before the maintenance was performed.

Different Ways to Achieve High Availability with Hyper-V

Hyper-V Quick Migration

One of the key advantages of server virtualization is the ability to deploy servers that are not tied to hardware. This translates into the ability to re-deploy a virtual server, or copy from one physical host to another without having to worry about reconfiguring the operating system for different chips sets, network cards, or other hardware constraints. With Hyper-V, Microsoft has optimized the transition process to quickly migrate a virtual server from one physical host to another in a matter of seconds. The ability to quickly migrate from one physical host to another is very beneficial when performing planned maintenance for on physical servers. This can be hardware maintenance or patch management. In order to use Quick Migration it is necessary to host the virtual machines on a shared storage device.

Quick Migration can transfer a virtual machine between physical servers with minimal downtime and no data loss, usually only last several seconds.

Failover clustering

When planning a highly available Hyper-V deployment, Windows Server 2008 Failover Clustering is one of the primary tools used to achieve a fault tolerant deployment. Windows Failover Clustering is the next version of Windows clustering. There are significant improvements that make configuring cluster environments more streamlined. The Create Cluster Wizard drastically reduces the steps required to configure server clustering. Additionally, the Cluster Validation Wizard validates the cluster configuration. It should be noted that Microsoft will only support cluster configurations that pass the Cluster Validation Wizard. A best practice is to ensure that any production environment passed this wizard.

Failover clustering provides child virtual machines with a fault-tolerant infrastructure. By combining Windows Failover Clustering with Hyper-V, Microsoft allows organizations to leverage a mature, well-tested technology for clustering virtual machines. Moreover, many enterprise administrators are already familiar with Windows Failover Clustering technology, and thus current skill sets can be leveraged to deploy such a solution.

In the event that the primary cluster node fails, Windows Failover Clustering will restart the virtual machines from that node on another physical cluster node. It is important to note that due to this restart, data loss may occur as anything that was in the virtual machine memory when the physical server experienced a failure will be lost. The advantage is that if a server fails during off hours the failover to a new physical server will be automatic.

Network load balancing across Hyper-V

Another high availability technology that can be deployed using Hyper-V is network load balancing. Hyper-V can be used to add additional nodes to a load balance cluster. Using systems monitoring solutions such as System Center Operations Manager, this can be done so that resources are dynamically brought online, or physical machine resource allocation can be changed during periods of peak demand.

High Availability Design Considerations

Designing Hyper-V clusters requires special considerations in order to ensure desired functionality and performance. The most important of these considerations are the following:

Processor considerations

One of the primary components that will impact all aspects of the Hyper-V solution is the processor. Since reliability is of utmost importance, you should consider using a processor that offers the latest in reliability features.

Intel® Xeon® processors are designed to provide leading reliability availability and serviceability. The Intel Xeon processor family includes many features that work to prevent hardware errors, as well as features that work to automatically detect and correct hardware errors when they do occur. The features combine to provide a robust and reliable processor platform that is ideal for powering a high availability solution based on Hyper-V. These features include:

- Fully buffered DIMM technology: Enables a more robust memory system, preventing memory errors from causing system disruption.
- Memory Error Correcting Code (ECC): Helps avoid memory errors.
- System bus enhancements: provide for more reliable data flow into and out of memory.
- Memory mirroring: Protects data that is loaded into RAM.
- Hot-Add Memory support: Allows ranges of memory to be added without rebooting or downtime.
- I/O hot-plug support: Enhances system reliability and prevents downtime due to I/O failure.
- Support for Windows Hardware Error Architecture (WHEA): WHEA provides a common infrastructure for handling hardware errors on Windows platforms. WHEA is intended to reduce mean-time-to-recovery for fatal hardware errors through richer error reporting and to reduce system crashes related to hardware errors through effective operating system hardware error recovery and health monitoring.

These all combine to allow the Intel Xeon processor family to provide a foundation of reliable components that minimize the possibility of hardware failure, allowing organizations to deploy a highly available solution with confidence.

Hardware requirements

- Hyper-V requires an x64-based processor, hardware-assisted virtualization, and hardware data execution protection. Intel Xeon processors provide state-of-the-art virtual machine acceleration with Intel® Virtualization Technology (Intel® VT). Enhancements such as FlexPriority allow virtual machines to have more direct access to hardware resources, thus improving virtual machine performance.
- Servers that are used in the cluster should use similar – and preferably identical – hardware. All hardware used in the cluster should be listed on the Windows Server 2008 Compatibility list. The Intel Xeon processor family, and related chipsets, provide parity across systems. By using Intel® chipsets with Intel Xeon processors, it is assured that compatibility between systems will be maintained, and further ensures that all Intel Xeon processor features will be available. The use of Intel VT and Intel® Core™ microarchitecture will also expedite the transition of virtual machines from one node to another.
- Network adapters should be dedicated to a single job. In the event that iSCSI storage is being used, a dedicated adapter should be used for communication with the iSCSI target, and another adapter should be used for network communication.
 - In addition, if using network adapters to communicate with iSCSI storage, it is recommended that more than one network interface be present to provide a redundant connection to the storage devices. If only one network interface exists, the Cluster Validation Wizard will throw a warning condition during the validation process.
 - Intel® dual- and quad-port network cards can provide a high density of network interfaces on a single card. This allows for dedicated iSCSI access to SAN for each virtual machine.
- If you are using Serial attached SCSI (SAS) or Fibre Channel to communicate with storage, the adapters on each cluster host should be identical. The firmware should also be the same on both cards.
- For iSCSI adapters and switching, equipment must be dedicated to storage communication. They cannot be used for normal network traffic. Furthermore, the network cards should be identical.
- Leverage multipath I/O wherever possible. This provides multiple connections to storage devices and allows for increased redundancy.

- The storage device must support a shared storage technology such as iSCSI, Fibre Channel or SAS. When creating the cluster, two volumes on the storage device will be required. The first volume is for the cluster's quorum information, also known as a witness disk. The second volume is where virtual machine files will be stored. As a best practice, each virtual machine should get its own volume on the storage device. Any server that is not part of the cluster should not have access to the virtual machine volume.

Software requirements

In order to configure Windows Failover Clustering, Microsoft Windows Server 2008 Enterprise Edition or Data Center edition must be used for the operating system on the physical servers. The servers can also be installed using a Server Core installation. This is often recommended for production environments, as it significantly reduces the attack surface of the host operating system. Note, however, that both cluster hosts must be configured with the same installation type.

The virtual servers may run any operating system that is necessary. Generally, it is best if Integration Services are installed to improve virtual machine performance.

Planning utilization and failover setup

When planning a two-node cluster server, it may be necessary to increase the utilization of server resources by having each node run several virtual machines. This allows organizations to get a better return on their hardware investment by allowing each node to act as the primary node for several complimentary virtual machine services. When planning for this kind of deployment, it is extremely important to ensure that the total peak load of one node does not exceed the available capacity on the fail over node.

If each node runs virtual machines that peak at 50 percent usage, then when one node fails, a single node should be able to accommodate the workload from the failed node. However, if the peak load exceeds 50 percent, you risk causing problems not only for the virtual machines that have failed over, but for the services that run on the node regularly.

With multiple nodes in a cluster, it is possible to redistribute one node's virtual machines to several other nodes in the cluster, therefore mitigating the chance that a single server will become over taxed for resources.

Disaster recovery requirements and geographically dispersed clustering

Virtualization is also becoming a key component in disaster recovery solutions. With the increasing options for high-speed long-haul networks, it is becoming increasingly affordable to create geographically dispersed clusters. Windows Server 2008 Failover Clustering has now been optimized to work across routed networks. If a location experiences a catastrophic event, it is now possible to have services fail over to a data center in another part of the world.

Conclusion

Designing highly available solutions is now a standard in many organizations, as they must ensure that their information systems remain operational in any event. Windows Server 2008 Hyper-V provides an extensive set of features ideal for providing high availability to virtual servers. In conjunction with failover clustering, and tools such as Virtual Machine Manager, those features provide a compelling list of benefits to help achieve the high availability goals of companies moving forward.

When designing a high availability solution with Hyper-V, it is important that the hardware platform includes reliability features to mitigate the risk of a single point of failure. Consider using Intel Xeon processor to provide the reliability and performance needed to power highly available solutions.

¹Intel® Virtualization Technology requires a computer system with an enabled Intel® processor, BIOS, virtual machine monitor (VMM) and, for some uses, certain platform software enabled for it. Functionality, performance or other benefits will vary depending on hardware and software configurations and may require a BIOS update. Software applications may not be compatible with all operating systems. Please check with your application vendor.

This document is for informational purposes only. INTEL MAKES NO WARRANTIES, EXPRESS OR IMPLIED, IN THIS DOCUMENT.

Copyright © 2008 Intel Corporation. All rights reserved. Intel, the Intel logo, Xeon, Core, and Xeon inside are trademarks of Intel Corporation in the U.S. and other countries.

*Other names and brands may be claimed as the property of others.

