

Hyper-V with Failover Clustering

Enhancing Computing Efficiency

Cabinet Secretariat Informatics Division's Data Centre has multiple physical hosts configured as a single cluster, running Hyper-V services in failover clustering mode and housing different virtual machines for various servers. Virtual hard disks are connected to these machines as shared resources via Cluster Shared Volume, which is configured through Storage Area Network.



DR. SHUBHAG CHAND

Technical Director
shubhag@nic.in



VINOD KUMAR J
Scientist 'B'
nj.vinod@nic.in

Edited by
MOHAN DAS VISWAM

Hyper-V with Failover Clustering is a combination of three technologies: Hyper-V, Failover Clustering and Cluster Shared Volume. Combination of these ensures optimal usage of the available resources, high availability and live-migration of running machines. Implemented at the Cabinet Secretariat Data Centre (CSDC), this combination possesses several advantages that ensure high efficiency in performances. Hyper-V is one of the roles present in the Server® 2008 R2 or higher that enables creation of a virtualized server computing environment improving the efficiency of computing resources.

FAILOVER CLUSTERING

Failover Clustering is one of the features of Windows Server Operating System that facilitate to create and manage failover clusters. A failover cluster is a group of independent computers that work together to increase the availability of applications and services. If one of the machines in the cluster fails, another machine begins to provide service (a process known as failover). This minimises disruption of services.

CLUSTER SHARED VOLUMES

Cluster Shared Volumes are volumes present in the shared storage area and are shared by multiple machines pres-

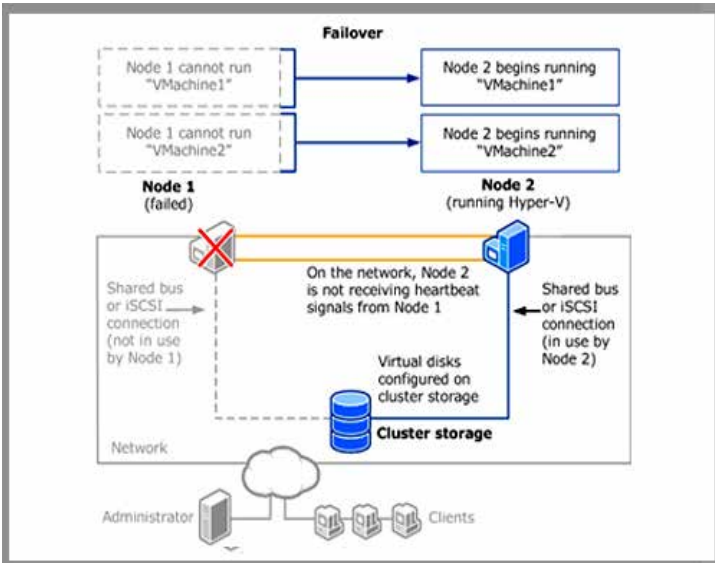
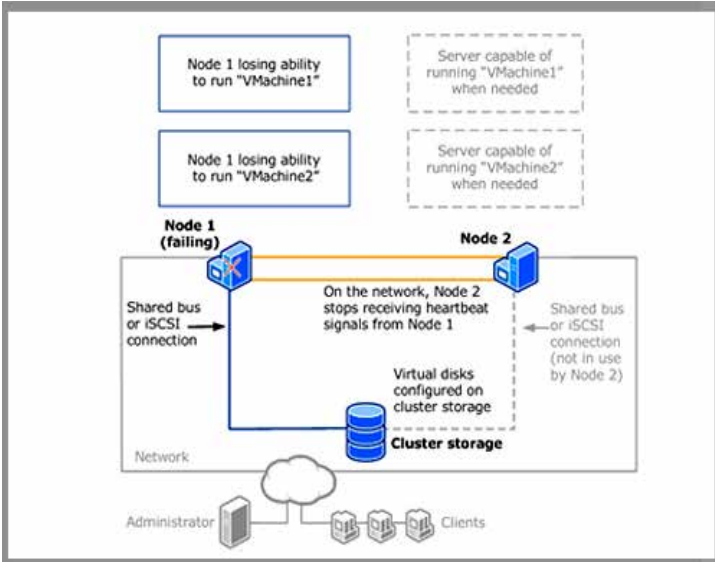
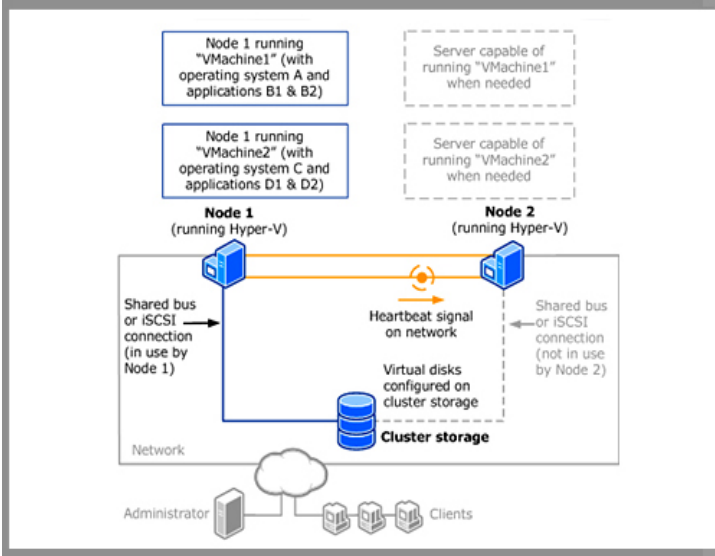
ent in the same cluster, meaning they can read from and write to on the same Volume. The machines coordinate the reading and writing activity so that the disk is not corrupted. Whereas, disks with LUN (Logical Unit Number) in Cluster Storage that are not Cluster Shared Volumes are always owned by a single machine at a time.

HOW HYPER-V FAILOVER CLUSTER DIFFERS FROM NORMAL FAILOVER CLUSTER

In normal Failover Cluster, two or more computers are joined to form a cluster. An application is installed to each of them, which are then designated as a cluster resource. Other resources such as disk storage can also be designated as a cluster resource. From there, behaviour is dependent on the application or resource being clustered. In some cases, the application or resource is "owned" by one host at a time and can manually or automatically be transferred to any other host in the cluster. In other cases, the application runs on all hosts simultaneously and is balanced across the hosts. In Hyper-V failover cluster, the Hyper-V is the application and the Cluster Shared Volume is the shared resource.

BENEFITS

There are several advantages of using Hyper-V with Failover clustering. One among the key benefits of virtualization is 'High Availability', a



term which refer to a group of technologies that allows to keep a virtual machine online even when it's intended physical host is not available.

Normal clustering requires more time for reconfiguring whenever there is a change in the application. But in Hyper-V with failover Clustering, since every application runs on Virtual Server, any change on the Virtual Machine doesn't affect the Clustering and hence it is fast.

If all the machines in the Cluster fail, then the VHD (Virtual Hard Disk) can be attached to any other available physical server and can be made live within few minutes and with minimum effort.

One of the major advantages of Hyper-V Failover Clustering is live migration. It enables moving a running virtual machine from one physical server to another without interruption of the services. It moves running virtual machines from one node of the failover cluster to another node in the same cluster without a dropped network connection or perceived downtime.

REQUIREMENTS

Servers:

Though the best practice for clustering is to use identical computers, it is also possible to use mixed hardware within the Hyper-V Cluster. The Hyper-V cannot Live Migrate between CPUs from different manufacturers and hence building a cluster that combines Intel and AMD chips would prove to be a wasteful effort other than for pure failover purposes. Even then, there's no guarantee for virtual machines to perform well with the mixed environment. For any hardware used, its CPU should support native virtualization (VT on Intel and AMD-V on AMD) and that support must be enabled in the BIOS. It is required to have its hardware Data Execution Prevention enabled in the BIOS (XD on Intel, NX on AMD). In case the RemoteFX is used, the CPU will have to support Second-Level Address Translation (SLAT) and the computers would require graphics processors that can run DirectX 9 or higher. These CPUs must also have enough dedicated memory for all of the virtual machines which will be running on them. For optimal operations, each computer is required to have at least four network cards. A minimum of one more card would have to be added if it is to be connected to storage devices by Internet Small Computer System Interface (iSCSI).



1. Fiber Channel Switches (2 Nos Active – Passive Mode)

2. Multiple Physical host Configured as a Single Cluster and running Hyper-V as Services.

Storage Device(s):

A central requirement of clustering is shared storage. Depending on the amount of workload and budget allocation, shared storage can be a computer with Windows Server and iSCSI target software installed and made to work as Network Attached Storage (NAS) device or a powerful Storage Area Network (SAN). The device chosen must support iSCSI-3 persistent

reservations. Fiber channel SANs usually meet the requirement.

THE NETWORK EQUIPMENT

All the hardware are to be connected and the best way is through a switch. For this, an iSCSI system with a layer-2 managed switch with sufficient gigabit ports to handle all the network cards of the hardware can be used. It is considered a good practice to completely segregate the iSCSI traffic by placing it on its own physical switches that are not connected to any others. However, this is not an ideal solution when there are budgetary constraints. If this iSCSI network cannot physically be separated, at least setting up a separate Virtual Local Area Network (VLAN) may have to be ensured. Fiber channel installations will have their own requirements, but it is preferable to use fiber channel switches rather than to directly connect the hosts to the storage device.

MUST-DO WITH HYPER-V FAILOVER

Taking backup of the Virtual Hard Disk (VHD) file is the main requirement of the Hyper-V failover clustering to ensure reliability. Windows Server 2012 onwards support incremental backup of virtual hard disks while the virtual machine is running (with Hyper-V). During this process, only the differences are backed up. Manual backup of VHD file is also possible as the VSS (Volume Shadow Copy Service) writer allow reading the VHD file when it is mounted to any virtual machine.

LIMITATIONS

- A maximum limit of 64 nodes in the failover cluster

- An upper limit of 8000 virtual machines per cluster and with a maximum of 1024 virtual machines on one node
- The number of virtual machines allowed for each node does not change regardless of size of the cluster
- Servers running Hyper-V can have maximum of 320 Logical processors and 4TB physical memory only
- Each Virtual Machines created on the Hyper-V can have only a maximum of 64 virtual processors, 1TB of physical memory and hard disk of size 64TB, if VHDX format is used (Windows Server 2012 onwards)

IMPLEMENTATION AT CSDC

Hyper-V failover clustering satisfies all major requirements of the Cabinet Secretariat Informatics Division's Data Centre. The Data Centre demands high availability of servers, live migration during server maintenance, quick configuration in case of failure of all Nodes of the cluster. Thus the Hyper-V failover clustering was implemented with multiple physical hosts configured as a Single Cluster, running Hyper-V services in failover clustering mode and housing different Virtual Machines for various servers. VHDs are connected to these Virtual Machines as shared resources via Cluster Shared Volume configured through Storage Area Network. The SAN is connected to the Cluster via Fiber Channel Switches with iSCSI as standby.

For further information, please contact:

DR. SHUBHAG CHAND

HOD & Technical Director

Cabinet Secretariat Informatics Division

Rashtrapati Bhawan, New Delhi

E-mail: shubhag@nic.in