

Configuring Microsoft Cluster Service in a VMware vSphere 4 Environment Using Infortrend Storage Application Note

Abstract

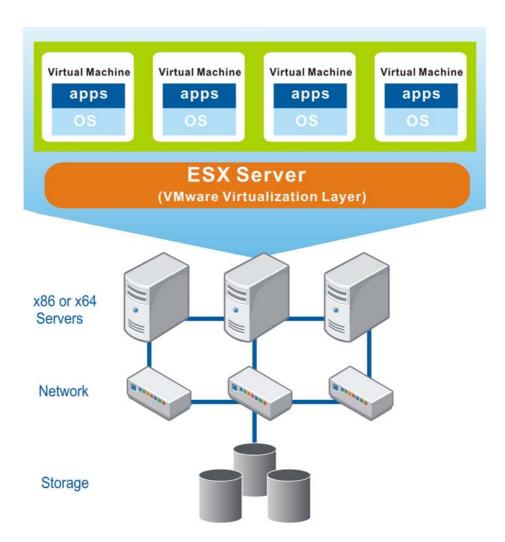
This application note explains the configuration details of using Microsoft Cluster Service to cluster virtual machines in VMware vSphere 4 to enhance system availability.



VMware Virtualization

The concept of virtualization originated in 1960s but was not applied to the x86 architecture until 1990s. Since 1980s, x86 servers have been widely adopted in IT environment because they are much cheaper than mainframe computers. This distributed system of computing reduces TCO but gives birth to other challenges, such as low infrastructure utilization, increasing physical infrastructure costs, increasing IT management costs, insufficient failover and disaster protection, and etc. Virtualization is found an effective way to deal with these challenges.

In VMware's virtualization technology, ESX Server is the foundation of virtualized environments.



Installed on an x86 or x64-based server, VMware ESX Server provides a virtualization layer on the host OS to consolidate all hardware resources, including processors, memories, storage and networking, and apply them to virtual machines (i.e. virtual servers) running on the physical server. Each virtual machine can have its own OS and



applications. By sharing hardware resources across multiple virtual machines, users can improve resource utilization and in turn greatly reduce the costs spent on building datacenter's physical infrastructure.

Besides the optimized resource utilization, VMware virtualization also reduces deployment efforts and simplifies management tasks. Free from the physical considerations and requirements, deploying virtual machines can be done in minutes or hours. After the deployment, managers can easily monitor the entire virtual datacenter through a unified management pane. When any of the physical device fails, the high availability features of VMware can ensure continuous system operation.

EonStor Storage Systems in VMware Virtualized

Environments

Storage Area Network (SAN) is an architecture brining all storage resources into a pool and making them shared among multiple hosts. It is the most suitable storage architecture for virtualized environments because when the operating systems (OS) and applications are moved around virtual machines, the data do not have to be copied and moved accordingly. Infortrend provides both FC-host RAID arrays and iSCSI RAID arrays to support SAN in VMware virtualized environments. EonStor arrays' features of centralized management and configuration platform, redundancy designs, and smart scalability can bring many benefits to the environments.

In a data center implemented with numerous storage systems, a powerful storage management tool is absolutely necessary. Infortrend's proprietary management suite, SANWatch, allows system managers to configure, administer and monitor multiple EonStor arrays locally or remotely through a user-friendly graphic interface. Integrating Infortrend's storage management tool into a VMware virtualized environment can make the infrastructure more robust and easier to maintain.

Besides the software complement, the fault-tolerant hardware modules and RAID functionalities of EonStor arrays further improve system availability. When controllers, power supplies, cooling modules or hard drives fail, they can all be hot-swapped without causing downtime or data loss.

Along with data growth, the capacity of a single EonStor subsystem can be scaled by connecting to expansion enclosures, and the scaling will not downgrade their industry-leading performance. By allowing users to start with the capacity they need



now and expand for more when future needs arise, the flexible way of scaling minimizes waste of investment and in turn contributes to enhancing virtualization's core value of optimized resource utilization.

To ensure seamless integration of EonStor Storage Area Network (SAN) storage into VMware virtualized environments, Infortrend has performed comprehensive testing to verify its full compatibility with VMware Infrastructure 3 (VI3) and vSphere 4. VI3 is VMware's flagship server and datacenter product and vSphere 4 is its renamed next major version. These powerful software suites can optimize and manage IT environments through virtualization. They both greatly reduce operating costs and increase IT service availability, security and scalability while providing the flexibility to choose any OS, application and hardware. Building on the proven power of VI3 platform, vSphere 4 delivers many enhanced features in simplified management, application services, infrastructure services and compatibility and third-party extensibility. For more details on the new features of vSphere 4, please refer to http://www.vmware.com/support/vsphere4/doc/vsp 40 new feat.html.

Configuring Microsoft Cluster Service in VMware vSphere 4

For high system availability, VMware vSphere 4 supports using Microsoft Cluster Service (MSCS) to cluster virtual machines. MSCS here refers to Microsoft Cluster Service with Windows 2000 Server and Windows Server 2003 and Failover Clustering with Windows Server 2008. There are three clustering configurations: clustering virtual machines on a single host, clustering virtual machines across physical hosts and clustering physical machines with virtual machines. No matter which configuration users adopt, each clustering setup includes shared disks and private network heartbeat between clustering nodes. Below we will illustrate the procedures of setting up the three clustering configurations. For detailed configuration steps, please refer to http://kb.vmware.com/selfservice/microsites/search.do?language=en_US&cmd=displayKC&externalId=1004617.

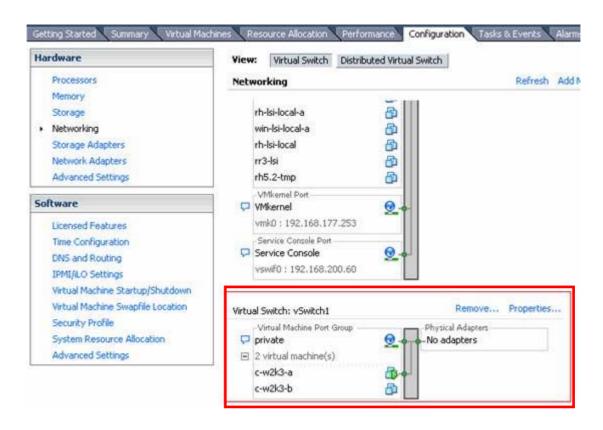
Cluster Virtual Machines on One Physical Host Step 1. Create the First Node for Clusters

In the vSphere Client inventory, right-lick the ESX host and select *New Virtual Machine*. Then go through the wizard to create the virtual machine. When creating disks, select *Support clustering features such as Fault Tolerance* under Disk Provisioning. Then create two network adapters for the virtual machine, labeling one public and the other private. Lastly, install a Windows Server operating system on the virtual machine.



Step 2. Create the Second Node for Clusters

In the vSphere Client inventory, right-click the virtual machine created as the first node and select *Clone*. Then go through the wizard to clone the virtual machine. For Disk Format, select *Same format as source*; for Customization, select *Do not customize*.



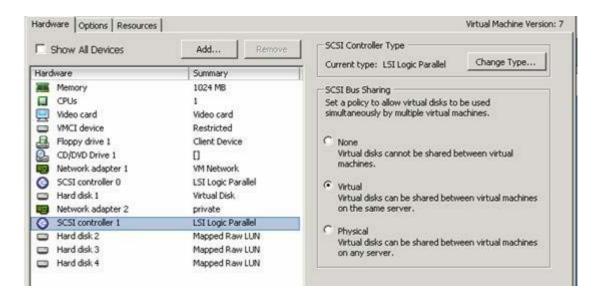
Step 3. Add a Shared Quorum Disk to the First Node

In an MSCS cluster, you have to set up a shared quorum disk to store meta-data that allows all nodes in a cluster to understand how they are related to each other. The quorum disk must be separate and isolated from all other data disks.

In the vSphere inventory, select the virtual machine created as the first node and select *Edit Settings...* to add hard disks to it. The disk types can be either a new virtual disk or a RDM set to virtual compatibility mode. Under Disk Provisioning, select *Support clustering features such as Fault Tolerance*. In the *Virtual Device Node*, select a new SCSI controller, such as *SCSI* (1:0), instead of *SCSI* (0:1). After the new hard disk and the new SCSI controller are created, select the appropriate type of controller depending on your operating system¹ in the Virtual Machine Properties dialog box. And the SCSI Bus Sharing should be set to *Virtual*.

If you use Windows 2000 Server and Windows Server 2003, the type of controller should be LSI Logic Parallel. If you use Windows Server 2008, the type should be LSI Logic SAS.





Step 4. Add Hard Disks to the Second Node

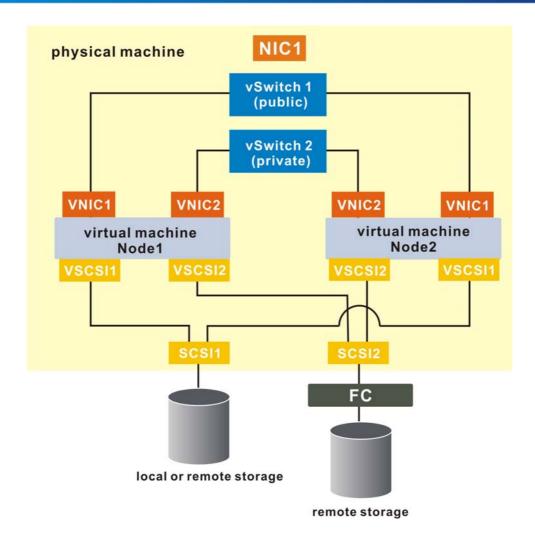
In the vSphere inventory, select the virtual machine created as the second node and select *Edit Settings...* to add a hard disk to it. The disk type should be *Use an existing virtual disk*. Select the same virtual device node you used for the first node's shared storage disks, such as *SCSI* (1:0). Then in Disk File Path, browse to the location of the quorum disk or data disks specified for the first node.

Step 5. Install Microsoft Cluster Service

For Windows Server 2003, after setting up the first and second nodes, you must configure Microsoft Cluster Service. Please see Microsoft documentation for details: http://support.microsoft.com/

Complete setup is as below:





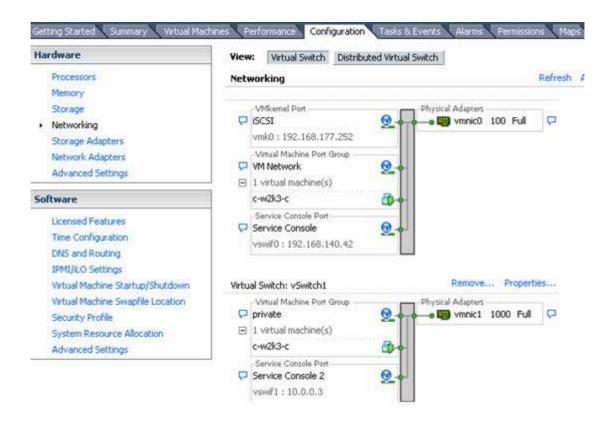
Cluster Virtual Machines Across Physical Hosts Step 1. Create the First Node for Clusters

In the vSphere Client inventory, right-lick the ESX host and select *New Virtual Machine*. Then go through the wizard to create the virtual machine. When creating disks, select *Support clustering features such as Fault Tolerance* under Disk Provisioning. Then create two network adapters for the virtual machine, labeling on public and the other private. Lastly, install a Windows Server operating system on the virtual machine.

Step 2. Create the Second Node for Clusters

In the vSphere Client inventory, right-click the virtual machine created as the first node and select *Clone*. Then go through the wizard to clone the virtual machine. For Disk Format, select *Same format as source*; for Customization, select *Do not customize*.





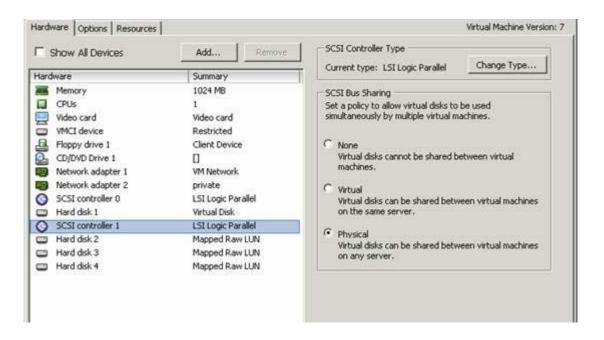
Step 3. Add a Shared Quorum Disk to the First Node

In an MSCS cluster, you have to set up a shared quorum disk and for meta-data that allows all nodes in a cluster to understand how they are related to each other. The quorum disk must be separate and isolated from all other data disks.

In the vSphere inventory, select the virtual machine created as the first node and select *Edit Settings...* to add hard disks to it. The disk type must be RDM and it is suggested to be set to physical compatibility mode². In the *Virtual Device Node*, select a new SCSI controller, such as *SCSI* (1:0), instead of *SCSI* (0:1). After the new hard disk and the new SCSI controller are created, select the appropriate type of controller depending on your operating system in the Virtual Machine Properties dialog box. And the SCSI Bus Sharing should be set to *Physical*.

The RDM volume must reside on a Fibre Channel (FC) SAN and only Windows 2000 Server and Windows Server 2003 support clustering across physical machines with RDM in virtual compatibility mode.





Step 4. Add Hard Disks to the Second Node

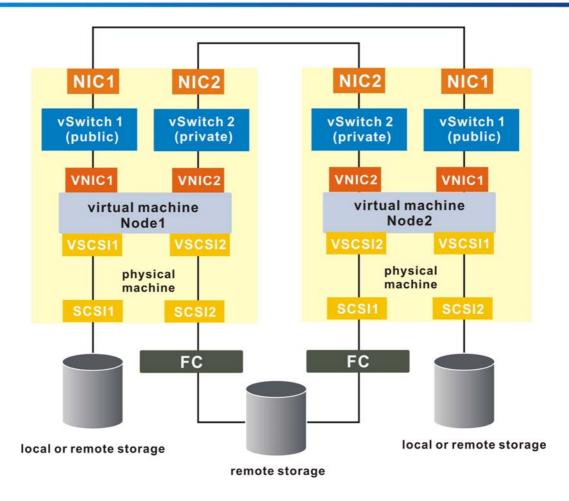
In the vSphere inventory, select the virtual machine created as the second node and select *Edit Settings...* to add a hard disk to it. The disk type should be *Use an existing virtual disk*. Then in Disk File Path, browse to the location of the quorum disk or data disks specified for the first node. Select *Physical* as the compatibility mode. In the *Virtual Device Node*, select the same virtual device node you used for the first node's shared storage disks, such as *SCSI (1:0)*. After the new hard disk and the new SCSI controller are created, select the appropriate type of controller depending on your operating system in the Virtual Machine Properties dialog box. And then set the SCSI Bus Sharing to *Physical*.

Step 5. Install Microsoft Cluster Service

For Windows Server 2003, after setting up the first and second nodes, you must configure Microsoft Cluster Service. Please see Microsoft documentation for details: http://support.microsoft.com/

Complete setup is as below:





Cluster Physical and Virtual Machines

Step 1. Create the First Node for Clusters

The first node is a physical machine. When deploying the machine, please note the following requirements.

- 1. If Windows 2003 is used, the Windows Cluster Administrator application of the physical machine should be set with *Advanced (minimum)* configuration.
- 2. The physical machine must have at least two network adapters
- 3. The physical machine must be able to access the same storage on a FC SAN as the ESX host on which the virtual machine as the second node will run.
- 4. Be sure to use STORport Miniport driver for the FC HBA (Qlogic or Emulex HBAs) installed in the physical machine.

For setup details, please refer to Microsoft documentation: http://support.microsoft.com/.

Step 2. Create the Second Node for Clusters

In the vSphere Client inventory, right-lick the ESX host and select *New Virtual Machine*. Then go through the wizard to create the virtual machine. When creating disks, select

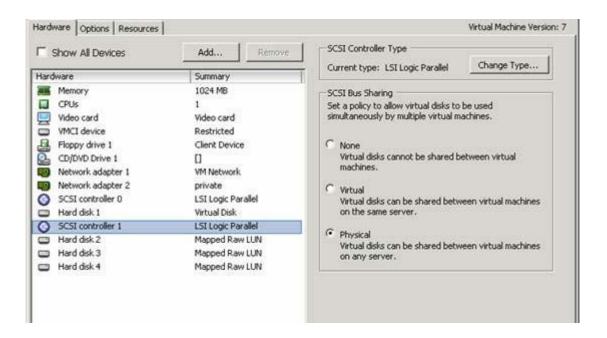


Support clustering features such as Fault Tolerance under Disk Provisioning. Then create two network adapters for the virtual machine, labeling on public and the other private. Lastly, install a Windows Server operating system on the virtual machine.

Step 3. Add Hard Disks to the Second Node

After creating the quorum disk and shared data disks (not required) for the physical machine, you should follow the procedure below to enable shared access to these disks for the virtual machine as second node.

In the vSphere inventory, select the virtual machine created as the second node and select *Edit Settings...* to add a hard disk to it. The disk type should be *Raw Device Mappings*. Select the LUN used by the physical machine and the datastore. Select *Physical* as the compatibility mode. In the *Virtual Device Node*, select a new SCSI controller, such as *SCSI* (1:0), instead of *SCSI* (0:1)). After the new hard disk and the new SCSI controller are created, select the appropriate type of controller depending on your operating system in the Virtual Machine Properties dialog box. And then set the SCSI Bus Sharing to *Physical*.



Step 4. Install Microsoft Cluster Service

For Windows Server 2003, after setting up the first and second nodes, you must configure Microsoft Cluster Service. Please see Microsoft documentation for details: http://support.microsoft.com/





