

Module 12: VSS Backups and Cluster Environments

Upon completion of this module, you should be able to:

- Describe:
 - Microsoft Volume Shadow Copy Service and its components
 - How VSS works
- Describe how to:
 - Back up and recover VSS SYSTEM save sets
 - Configure a NetWorker client in a cluster
 - Configure cluster nodes as NetWorker storage nodes
- Explain how:
 - To recover a clustered client's data
 - A NetWorker server works in a cluster



This module focuses on NetWorker with VSS and backing up and recovering with NetWorker in a cluster environment.

Module 12: VSS Backups and Cluster Environments

Lesson 1: VSS Backup and Recovery

During this lesson the following topics are covered:

- Volume Shadow Copy Service (VSS)
- VSS SYSTEM save sets
- Windows system recovery



This lesson covers VSS backup and recovery, including the use of VSS with backup and recovery, VSS SYSTEM save sets, and an overview of Windows system disaster recovery.

Volume Shadow Copy Service

- Volume Shadow Copy Service (VSS) is a Microsoft framework for creating, managing, and restoring consistent point-in-time copies of NTFS volumes on running systems.
- VSS coordinates the activities of the following to create point-in-time copies of Windows volumes:
 - ▶ Backup applications (VSS requestors).
 - ▶ Business applications and services running on the system (VSS writers).
 - ▶ Storage subsystems (VSS providers).
- The point-in-time volume copies are called “shadow copies”.



VSS provides the backup infrastructure for Microsoft Windows operating systems for Windows XP, 2003, and later, as a built-in component and works at a low level with storage and file systems.

VSS was designed to create a Windows service that is able to coordinate the actions required to create a consistent shadow copy (also known as a snapshot or a point-in-time copy) of the data that you want to backup. The actions to coordinate are between the backup application, the service or application that contains the data, and the I/O subsystem to get a consistent snapshot of the data. If your operating system, applications, backup software, and SAN manufacturer all support VSS, you can create flexible storage solutions that can easily be protected, without the need to stop servicing clients.

Produces clean (uncorrupted) snapshots of a volume by enabling applications to flush partially committed data from memory in a coordinated fashion with the snapshot requestor and the hardware. This ability prevents “torn writes” – the occurrence of the system not being able to complete the write of a block of data to disk.

VSS provides multiple volume-level snapshots, application-level snapshots, and management of shadow copies.

Shadow copies can be used for other purposes but backup is a major use.

Need for VSS

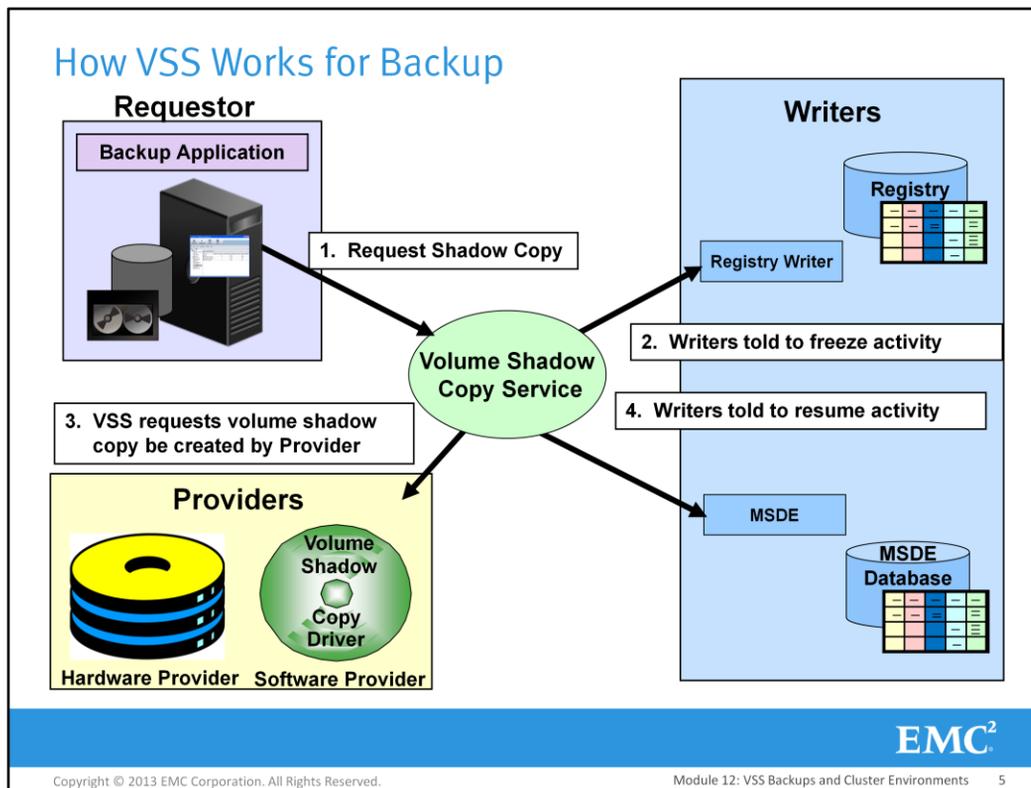
- Designed to address common backup issues with Windows XP, Server 2003, and later versions.
- Common Windows Volume Backup issues:
 - ▶ Inaccessible files during a backup – Applications keep files open in exclusive mode.
 - ▶ Inconsistent file state – Application files opened/read by backup may not all reflect the same application state.
 - ▶ Service interruptions – Applications run 24/7 and hence, cannot be stopped for obtaining consistent backups.
- Components include:
 - ▶ Requester – request the snapshot (NetWorker)
 - ▶ Writer – writes the actual data
 - ▶ Provider – creates and maintains the shadow copy



VSS is designed to address common backup issues for Windows XP, Server 2003 (and later) versions. Earlier versions of Windows still require Open File Manager.

- Applications that are running need to keep files open in exclusive mode. This prevents backup programs from copying them.
- There is a finite time requirement to open, backup, and close a file. If an application keeps multiple files open for updates, it is possible that files copied to the backup media may not all reflect the same application state.
- To overcome these problems, a backup program could require the suspension or termination of all running programs to ensure file accessibility and integrity of data. However, service interruption is not tolerated in many environments.

The basic VSS components include: the requester which requests the snapshot, the writer which writes the actual data, and the provider which is responsible for creating and maintaining the shadow copy.



Volume Shadow Copy Service (VSS) manages the creation of a point-in-time shadow copy, or snapshot, of a disk volume or set of data. This slide describes the process of creating a snapshot.

1. The requestor, such as a NetWorker, requests VSS to create a snapshot of a particular volume or data set.
2. VSS notifies the application-specific writer to prepare its data for making a shadow copy. The writer creates an XML description of the backup components and defines the restore method. The writer prepares the data by completing all open transactions, rolling transaction logs and flushing caches. VSS then directs the writer to temporarily freeze requestor I/O write requests for the time required to create the shadow copy. VSS flushes the file system buffer and then freezes the file system, which ensures that file system metadata is written and that the data is written in a consistent order.
3. VSS notifies a provider to create and maintain the shadow copy until it is no longer needed. A point-in-time copy of the complete volume mapping is created using XML. The XML file is a bitmap of the current state of the volume but does not contain any file data.
4. Once the snapshot has been successfully created, VSS thaws the file system and instructs the writers to resume normal activities, or thaw. VSS provides the location information for the shadow copy back to the requestor.

The requestor uses the snapshot to create the backup. The backup is as of the point in time that the snapshot is taken. The snapshot can also be backed up to secondary storage through a proxy client by either using an existing snapshot or with the use of temporary snapshots.

How NetWorker Works with VSS (1 of 2)

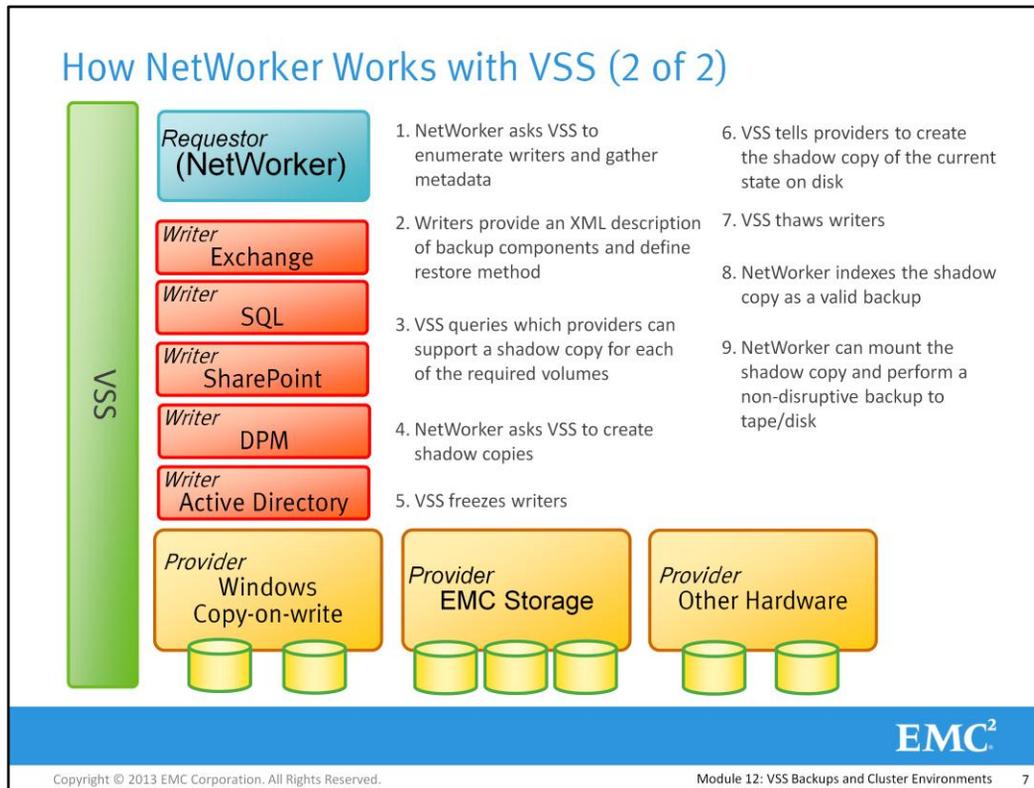
- NetWorker utilizes the VSS framework to safely and completely back up sets of data that may contain “open files”, including:
 - ▶ SYSTEM Save Sets
 - ▶ Databases
 - ▶ Entire file systems
- VSS technology is used to backup Windows clients by default
 - VSS is the only backup method for Windows 7, Microsoft Server 2008 and later versions.
 - VSS is the only backup method for most Microsoft applications.



NetWorker will leverage VSS on Windows servers to allow for the backup of open and active files by default. Additionally, VSS is the only option for backup of most Microsoft applications.

If VSS is unavailable or if the snapshot process fails for some reason, the backup reverts to the legacy method if available on the operating system or application. For applications or operating systems that do not support non-VSS backups, the backup will simply fail.

How NetWorker Works with VSS (2 of 2)



Looking at NetWorker in the context of the VSS framework and its supported providers, the process of capturing a consistent copy of application data is shown. One of the significant advantages with NetWorker in this process is that once a shadow copy or snapshot is created, NetWorker indexes the shadow copy as a valid backup from which it can recover. From this snapshot backup, NetWorker is able to perform a secondary backup onto tape, disk or virtual tape.

VSS SYSTEM Save Sets

- VSS SYSTEM save sets are automatically generated during VSS backups when a client's save set list contains **All**.
- VSS SYSTEM save sets include data components which are dynamically determined at the time of backup.
- The save sets for VSS SYSTEM backups have the following save set names:
 - ▶ VSS SYSTEM BOOT:
 - ▶ VSS SYSTEM FILESET:
 - ▶ VSS SYSTEM SERVICES:
 - ▶ VSS USER DATA: (Windows Server 2003 only)
 - ▶ VSS OTHER: (Windows Server 2003 only)



VSS SYSTEM save sets are automatically generated during VSS backups instead of the standard SYSTEM save sets when a client's save set list contains **All**. If the save set is set to anything other than **All** and you want any of the SYSTEM save sets to be backed up, you must explicitly specify them in the save set list of the client resource.

Save sets for VSS SYSTEM backups include data components which are dynamically generated. This may differ each time NetWorker software runs. Some components are backed up only if they are installed and have active writers.

Notes:

Please refer to the *EMC NetWorker Administration Guide* for a detailed description of the contents of the VSS SYSTEM save sets.

Beginning with NetWorker v8 clients, ASR backup is no longer supported. However, ASR recovery using pre-NetWorker v8 client ASR save sets is still supported.

Disaster_Recovery Save Set

- Provides support for Windows Disaster Recovery on Windows 2008 and Windows 7 and newer platforms.
- Encapsulates all critical volumes required for complete Windows DR.
- Created during full backups when either the save set All or the DISASTER_RECOVERY:\ save set is specified.
- Backup of the DISASTER_RECOVERY:\ save set includes:
 - ▶ All critical volumes.
 - ▶ VSS System State save sets: VSS SYSTEM BOOT, VSS SYSTEM SERVICES, VSS SYSTEM FILESET.
- Save sets included in a DISASTER_RECOVERY:\ save set can be displayed in the Save Set tab of the Media window.



Beginning with NetWorker 7.6 SP2, a disaster recovery save set is introduced for Windows 2008, Windows 7, and newer client platforms.

The DISASTER_RECOVERY : \ save set is included in full backups when the save set list is All or DISASTER_RECOVERY:\. This save set encapsulates all of the critical volumes required to provide complete Windows DR capabilities. The following VSS System State save sets are also included: VSS SYSTEM BOOT, VSS SYSTEM SERVICES, and VSS SYSTEM FILESET.

Notes:

You can revert the definition of the All save set to be equivalent to that of a Windows 2003 host with VSS enabled by adding the **VSS:DISASTER_RECOVERY=off** keyword to the **Save operations** attribute of the NetWorker client resource.

Files that are associated with application VSS writers are not backed up as part of the DISASTER_RECOVERY:\ save set and cannot be recovered unless they are backed up by an application backup program, such as NMM.

Critical volumes are: any volume that contains OS files or files needed by an installed service, a non-critical volume with a critical volume mounted on it, a non-critical volume serving as a parent to a critical volume, all volumes on all dynamic disks if at least one of those volumes is critical. Please refer to the *EMC NetWorker Administration Guide* for a complete discussion of the components of the DISASTER_RECOVERY:\ save set.

savefs and Save Set Level for VSS

- `savefs` (client) returns the save set worklist:
 - ▶ Local drives: `C:\`, `D:\`, `E:\`, etc.
 - ▶ `VSS SYSTEM` save sets / `DISASTER_RECOVERY:\` save set.
- The **Save operations** attribute in the NetWorker client resource can be used to disable VSS for individual drives and writers.
- `VSS SYSTEM` save sets are always backed up at a Full level:
 - ▶ However, if **VSS SYSTEM FILESET:** has not changed since the last backup, no save is performed.
- `DISASTER_RECOVERY:\` save set is backed up only during Full backups:
 - ▶ During incremental backups, all other specified save sets are backed up.



When backing up using VSS, if the client has All in its save set attribute, `savefs` is executed on the client to determine which save sets to back up. The resulting work list is returned to the `savegrp` command. The work list includes local drives such as `C:\`, `D:\`, `E:\`, etc. and `VSS SYSTEM` / `DISASTER_RECOVERY:\` save sets.

In the NetWorker client resource, there is a **Save operations** attribute that you can use to modify the default VSS behavior.

`VSS SYSTEM` save sets are always backed up at a full level, even if another level is specified. However, the `VSS SYSTEM FILESET:` save set is unique in that if none of its components have changed since the last backup, that save set is not saved.

The `DISASTER_RECOVERY:\` save set is backed up only during full backups. During incremental backups, all other specified save sets are backed up.

VSS Recovery Workflow

The steps performed by NetWorker to recover VSS save sets are:

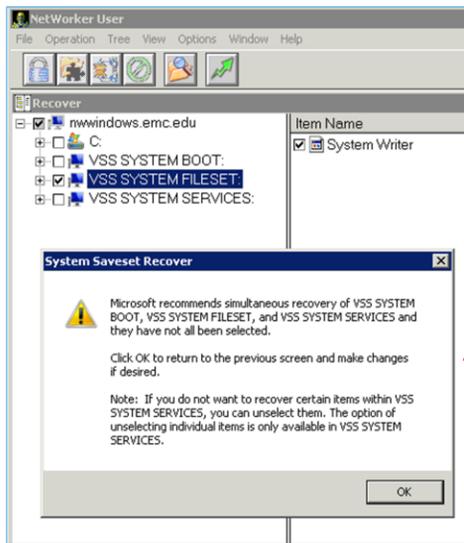
1. The writer's metadata document is recovered and used by NetWorker to determine the writer's restore method.
2. Active writers are queried to determine if an alternate recover location will be used; this step is dependent on the type of VSS save set.
3. NetWorker asks VSS to mark each component that will be recovered.
4. A pre-restore warning is sent to all affected writers.
5. The files are recovered to the correct locations as defined by the recovery method for each writer.
6. Once files have been recovered, a post-restore message is sent to all affected writers.
7. Any final steps required to complete the recovery are performed, as detailed in the writer's metadata document.



The steps performed by NetWorker to recover a VSS save set are listed on the slide.

Note: VSS SYSTEM save sets may not be relocated during recovery. For most VSS SYSTEM save sets, the final step involves displaying a message that the system needs to be rebooted to complete the recovery.

Recovering VSS Save Sets (1 of 2)



Microsoft recommends simultaneous recovery of VSS SYSTEM BOOT, VSS SYSTEM FILESET, VSS SYSTEM SERVICES.

For a full recovery, always recover the operating system drives, including the system and boot drives.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 12

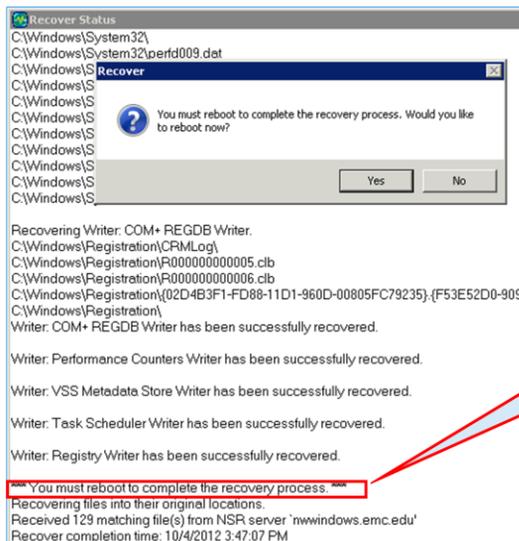
The NetWorker User interface provides a selection of VSS save sets for recovery.

If recovering the Windows system state, recover the VSS SYSTEM save sets in the following order:

- VSS SYSTEM SERVICES
- VSS SYSTEM FILESET
- VSS USER DATA
- VSS OTHER
- VSS SYSTEM BOOT

Microsoft recommends recovering VSS SYSTEM BOOT, VSS SYSTEM FILESET, and VSS SYSTEM SERVICES at the save time. VSS SYSTEM BOOT is backed up as a single, indivisible unit. VSS SYSTEM FILESET components cannot be backed up or restored individually. The components of the VSS SYSTEM SERVICES, VSS USER DATA and VSS OTHER save sets mentioned cannot be backed up individually, but can be restored individually.

Recovering VSS Save Sets (2 of 2)



VSS SYSTEM save sets can only be recovered to their original location. Directed recoveries are not supported.

After recovering a VSS SYSTEM save set, you must reboot the computer.

A list of recovered files and errors is written to ...\\nsr\\logs\\networkr.raw

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 13

VSS SYSTEM save sets can be recovered only to their original location. Directed recoveries are not supported because of differences in hardware information among machines.

After recovering a VSS SYSTEM save set, you must reboot the computer. The content of the **Recover Status** window is also written to \\nsr\\logs\\networkr.raw file for future viewing. This file is overwritten each time a client-initiated save or recovery is performed.

Windows Disaster Recovery Overview

- NetWorker 8.1 supports Bare Metal Recovery (BMR)
 - ▶ Full and Incremental backup and recovery support
 - ▶ Incremental support provides reduced backup time / size
- Windows disaster recovery tasks:
 - ▶ Create a NetWorker Windows Disaster Recovery bootable media.
 - ▶ The disaster recovery boot image can be downloaded from Powerlink at Home → Support → Software Downloads and Licensing → Downloads J-O → NetWorker
 - ▶ Ensure prerequisites are met.
 - ▶ Boot from the bootable CD or network location and run the recovery wizard.
 - ▶ Lastly, recover data that was not recovered as part of disaster recovery.



NetWorker 8.1 now supports backup and recoveries at all levels except for synthetic full operations. This means that NetWorker can perform Bare Metal Recoveries using Incrementals. Prior to NetWorker 8.1 only BMR full backups were supported.

For a complete discussion of Windows server disaster recovery operations with NetWorker, please refer to the *EMC NetWorker Administration Guide*. As with all recovery operations, it is recommended that the process and procedures for Windows server disaster recovery be tested without completing the entire recovery process (exit before formatting the drives and performing the actual recovery) to ensure successful recovery when needed. Be aware that running the wizard to completion will format the disks chosen to restore which will erase any existing data.

Module 12: VSS Backups and Cluster Environments

Lesson 1: Summary

During this lesson the following topics were covered:

- Volume Shadow Copy Service (VSS)
- VSS SYSTEM save sets
- Windows system recovery



This lesson covered VSS backup and recovery, including the use of VSS with backup and recovery, VSS SYSTEM save sets, and an overview of Windows system disaster recovery.

Module 12: VSS Backups and Cluster Environments

Lesson 2: Backup and Recovery of Clusters

During this lesson the following topics are covered:

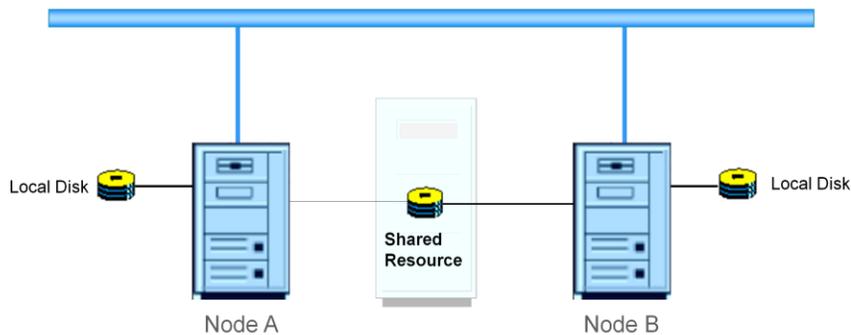
- Cluster components and characteristics
- Configuring cluster clients
- Managing path ownership issues
- Configuring storage nodes and NetWorker servers in a cluster



This lesson covers backup and recovery of clusters as well as the configuration of storage nodes and NetWorker servers in a cluster. Topics include cluster components and characteristics; the procedures for configuring cluster clients, storage nodes, and NetWorker servers; and the management of path ownership with clusters.

Understanding Cluster Basics

- A *node* is a physical host.
- Nodes are able to run applications installed on their own physical disks, while common data is stored on a shared resource (disk).



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 17

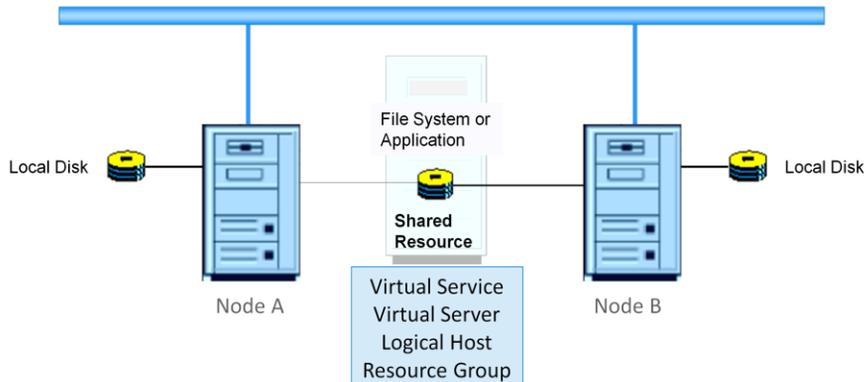
Clustering is a common practice that can help ensure that data or applications are continuously available to clients on a network. The basic premise of clustering is simple: two or more nodes (physical hosts) are connected and appear to network users as a single, highly available system.

When using a clustering application, all nodes in a cluster share one or more disk resources. In an active/passive cluster, only one of the nodes in the cluster is active at any given time. The active node is responsible for managing the shared resources. All other nodes in the cluster are passive nodes. If the active node fails for any reason, one of the passive nodes will take control of the shared resources.

Clustering can involve more than two nodes and may also involve load balancing. Clustering can also be configured in active/active arrangements where there are multiple shared resources and each of the nodes is the active node for one or more resources. This module covers a basic cluster environment of two nodes in an active/passive configuration.

Creating a Virtual Client from a Shared Resource

- A shared resource houses the application data which can be accessed from either of the two nodes, but not from both at the same time.
- A shared resource is referred to by many different names depending on the clustering software being used.
- It has its own hostname and IP address.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 18

A shared resource may be either a set of files or an application. There may be many shared resources within a cluster. A shared resource within a cluster is referred to by any of several different names, depending on the clustering software being used. For the remainder of this course, a shared resource will be referred to as a *virtual service*. A virtual service is always managed by the active node.

A virtual service is not a physical host, but rather a shared resource that each node of the cluster can access. Each shared resource may be comprised of multiple components, such as files, processes, data, and so on, and is assigned its own hostname and IP address. It is seen by hosts outside the cluster as a normal physical host.

During normal operation, the active node manages all communication between the virtual services and other hosts on the network. If a planned shutdown or failure of the active node occurs, control of the virtual services is transferred to the other node in the cluster, which changes from the passive to the active node.

When the failed node is returned to a functional condition, it becomes the passive node and is available for failover in the event of a failure of the current active node.

NetWorker Cluster Configuration Types

NetWorker Host Type	Description
Client	Install NetWorker client software on each node of the cluster.
Storage Node	Install NetWorker client and storage node software on each node of the cluster and run the appropriate NetWorker clustering script on each node. Configure device resources to be managed by each physical cluster node.
Server	Install client, storage node, and server software on each node. Run the appropriate NetWorker clustering script to cluster the NetWorker server processes and data, thereby guaranteeing high availability of the NetWorker server.



NetWorker supports three types of cluster configurations: a clustered NetWorker client, a clustered NetWorker client that is also configured as a NetWorker storage node, and a clustered NetWorker server.

Clustering a NetWorker client involves installing NetWorker client software on each node in the cluster. In addition to creating NetWorker client resources for each of the nodes, one or more NetWorker client resources (referred to as **virtual clients**) are created for each virtual service. In this configuration, the virtual client is backed up through the active node.

Clustering a NetWorker client which is also configured as a NetWorker storage node is essentially the same as configuring a clustered NetWorker client, but involves configuring each node in the cluster as a NetWorker storage node. NetWorker client and storage node software are installed on each node, and each node controls one or more backup devices. The virtual client is backed up to a device managed by the active node. When using this configuration all devices within the cluster are created as remote devices, and the storage node keyword **curphyhost** is used to tell NetWorker to backup to the devices attached to the current physical host.

When clustering the NetWorker server, the NetWorker client, storage node, and server software are installed on each node in the cluster. Using a NetWorker clustering script, NetWorker itself is configured as a virtual service. NetWorker binaries reside on the nodes but the remaining NetWorker files, such as resources, client file indexes, the media database, and logs, are moved to a shared resource. Thus, it is the NetWorker server application which is clustered, resulting in increased availability.

This course covers the generic steps for configuring NetWorker in a clustered environment. For more information on installing NetWorker server in a cluster configuration refer to the *NetWorker Cluster Installation Guide*.

Configuring a Clustered NetWorker Client

To configure a clustered NetWorker client:

1. Install the NetWorker client software on each node in the cluster.
2. Configure NetWorker client software as highly available
3. Create NetWorker client resources for each node and virtual service in the cluster.
4. Configure the **Save set** attribute of each client.
5. Configure the **Remote access** attribute of each virtual client to include all nodes within the cluster.

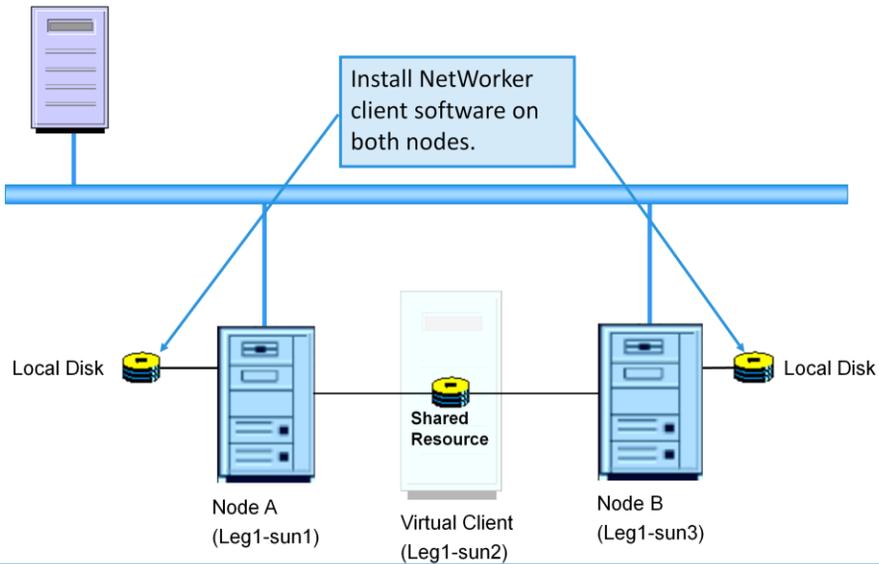


A number of steps are required to configure a clustered NetWorker client. These steps include installing the NetWorker client software on each node in the cluster, creating NetWorker client resources for each node and virtual service in the cluster, configuring the **Save set** attribute of each client resource, and configuring the **Remote access** attribute of each virtual client to include all nodes within the cluster.

These steps are covered in more detail on the subsequent slides.

Installing NetWorker Client Software

NetWorker Server(Leg1-sun10)



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 21

NetWorker client software is installed on every node in the cluster enabling each node to function as the active node if necessary. NetWorker software is not required to be installed on the shared resource.

Creating a Cluster Client Resource

The image displays three sequential screenshots of the 'Create Client' dialog box in EMC NetWorker, illustrating the configuration for different client types in a cluster environment. Each screenshot shows the 'Identity' and 'Backup' sections.

- Node A:** The 'Name' field is 'Leg1-sun1', the 'Comment' is 'Physical Node A', and the 'Backup type' is empty. The 'Scheduled backup' checkbox is checked, the 'Directive' is empty, and the 'Save set' is 'All'.
- Virtual Client:** The 'Name' field is 'Leg1-sun2', the 'Comment' is 'Virtual Application Server', and the 'Backup type' is empty. The 'Scheduled backup' checkbox is checked, the 'Directive' is empty, and the 'Save set' is '/sharedapps'.
- Node B:** The 'Name' field is 'Leg1-sun3', the 'Comment' is 'Physical Node B', and the 'Backup type' is empty. The 'Scheduled backup' checkbox is checked, the 'Directive' is empty, and the 'Save set' is 'All'.

Each 'Identity' section is highlighted with a red box. The 'Backup' section is also visible in each screenshot.

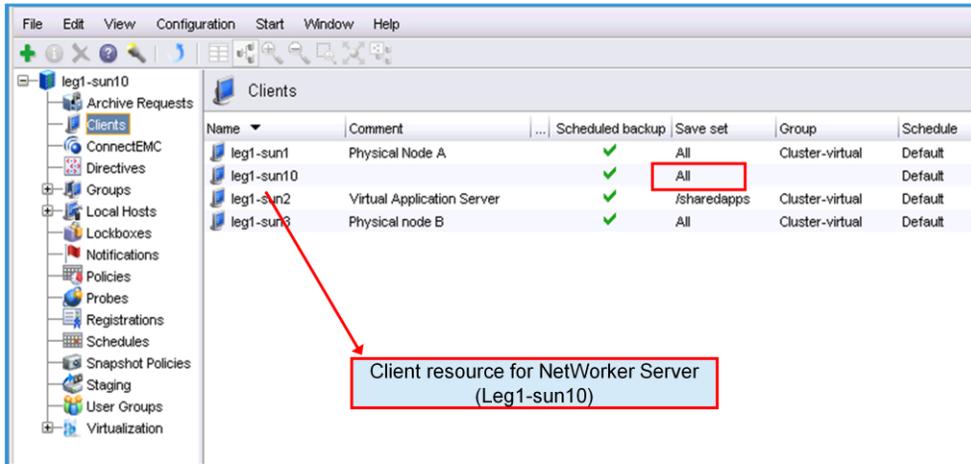
EMC²
Copyright © 2013 EMC Corporation. All Rights Reserved. Module 12: VSS Backups and Cluster Environments 22

NetWorker client resources are created for each node in the cluster as well as for each virtual service. In a cluster environment with two nodes and one virtual service, you would configure at least three NetWorker client resources.

Each node backs up data residing on its own local disks. A virtual client backs up the shared clustered data. If the cluster has multiple virtual services which require multiple hostnames and IP addresses, it is necessary to create at least one NetWorker client resource for each virtual service. As with any NetWorker client, multiple client resources may be configured for each node and virtual service.

Remember that each virtual client will have its own hostname and IP address and that all hosts must be listed in the appropriate name service database, such as local hosts file, DNS, NIS or NIS+. It is important that reverse lookups behave correctly.

Updating the Save Set Attribute



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 23

When creating the client resources, it is necessary to properly configure the **Save set** attribute. Ensure that the virtual client is backing up all shared data and that the NetWorker client resource of each node includes the local data on that host.

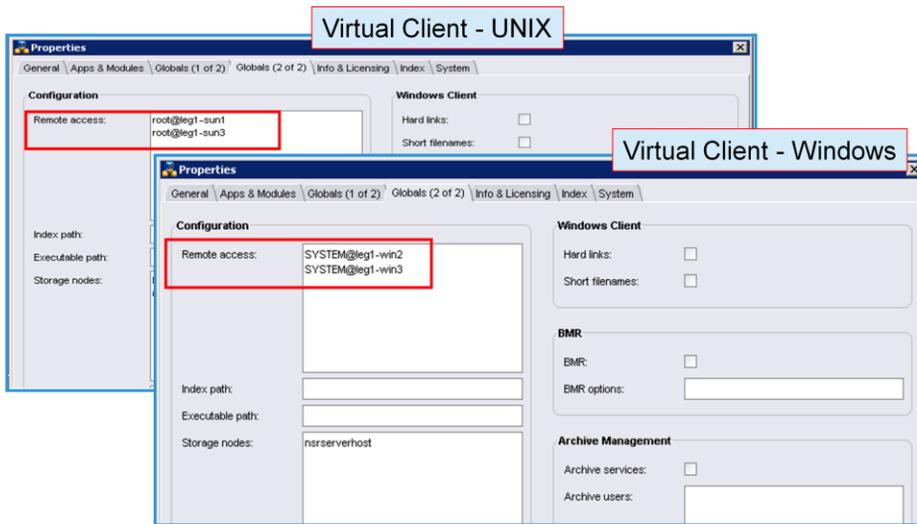
Make sure that the **Save set** attribute of the virtual client(s) and the nodes account for all data, shared and non-shared, on the systems.

Although use of the **All** save set is supported for a virtual client, it is recommended that you use the **All** save set only for the nodes. When **All** is specified for a node, it does not include the shared data.

Ensure that you perform test backups and monitor the save sets that are backed up. It is important to watch which client file index (CFI) is updated when a client is backed up. If a backup of a node results in the virtual client's CFI being updated or, conversely, a backup of a virtual client results in the active node's CFI being updated, difficulties may result when browsing for files during a recovery.

Later, we discuss the steps that can be taken if an incorrect CFI is being updated during backup of any of the cluster hosts.

Configuring Remote Access

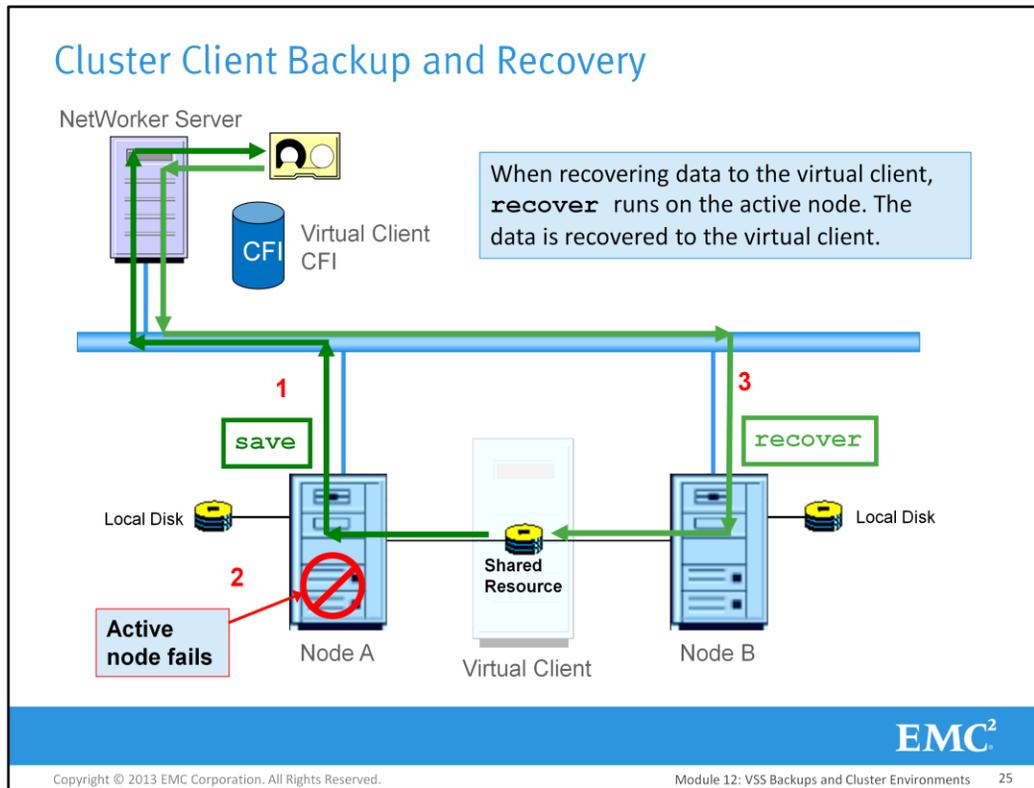


Configure the **Remote access** attribute of the virtual client to include all nodes within the cluster.



Configure the **Remote access** attribute of each virtual client to include the host names of both nodes in the cluster. This allows recoveries of the virtual client to be performed by the active node, regardless of which node is currently active. Thus, if the virtual client is normally backed up by Node A, and Node A becomes unavailable due to a failover to Node B, Node B can use the virtual client's CFI to perform a recovery to the virtual client.

Cluster Client Backup and Recovery



The virtual client is configured as a normal NetWorker client resource. The clustered data is backed up as though it belongs to the virtual client.

When the virtual client backs up, its CFI is updated, regardless of which node is active. If a recovery of the clustered data is required, whichever node is active can perform the recovery. This assumes that both nodes are listed in the virtual client's **Remote access** attribute.

In a UNIX cluster, the virtual client's shared data is mounted on the active node. To recover data belonging to the virtual client, a normal browsable or save set recovery is performed from the active node. However, the virtual client is selected as the "source" client and the data must be relocated to the directory on the active node where the shared data is mounted.

To recover data to the virtual client in a Windows environment, the active node is the "administering" client in the recovery while the virtual client is both the "source" and "destination" clients.

Path Ownership Rules (1 of 2)

- Path ownership rules are the criteria used by the clustering software to determine which host owns each save set being backed up.
- If there are conflicts with path ownership rules, NetWorker might have trouble assigning a save set to the correct client.
- A potential problem can be detected by monitoring the savegroup completion report and by checking the virtual client's CFI after a backup of the virtual client.

```
# /usr/sbin/savegrp -pu -c leg1-sun2 Cluster-virtual
leg1-sun2:/sharedapps level=incr
06/06/10 14:53:13 savegrp: Group will not limit job parallelism
06/06/10 14:53:13 savegrp: leg1-sun2:probe started
savefs -s leg1-sun10 -c leg1-sun2-g Cluster-virtual -p -l full -R -v -F /sharedapps
savefs leg1-sun2: succeeded.
06/06/10 14:53:14 savegrp: leg1-sun2:probe succeeded.
--- Probe Summary ---

leg1-sun2:Probe level=full, dn=-1, mx=0, vers=ssbrowse, p=4
leg1-sun2:Probe level=full, pool=Default, save as of Tue Jun 6 14:53:14 GMT-0400 2010
leg1-sun2:/sharedapps level=full, dn=0, mx=1, vers=ssbrowse, p=4
leg1-sun2:/sharedapps level=full, pool=Default, save as of Tue Jun 6 14:53:14 GMT-0400 2010
leg1-sun2:index level=full, dn=-1, mx=0, vers=ssbrowse, p=4
leg1-sun2:index level=full, pool=Default, save as of Tue Jun 6 14:53:14 GMT-0400 2010
06/06/10 14:53:14 savegrp: nsrim run recently, skipping
```

Verify proper client and group resource configuration

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 26

In a clustered environment, NetWorker must determine which save sets are owned by the nodes and which save sets are owned by the virtual client(s). The criteria used to determine save set ownership are called *path ownership rules*. These rules determine which CFI save set tracking information is written to.

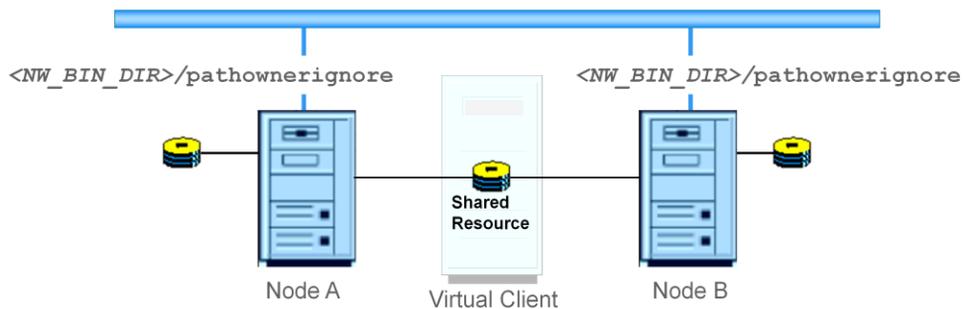
If NetWorker determines that a save set defined in a client resource is not owned by that client, NetWorker might not back up the save set during a server-initiated backup. This prevents a clustered host from writing to multiple client file indexes. Writing to multiple indexes can cause recovery problems.

The following conditions mean a save set will **not be backed up** during a scheduled backup: if a save set owned by a virtual client is defined in the **Save set** attribute of a node's client resource, and, if a save set owned by a node is defined in the **Save set** attribute of a virtual client resource.

To determine if an incorrect CFI will be used, preview a server-initiated backup of each node and virtual client after the cluster is configured. Run a test probe with the `savegrp` command as shown on the slide, once for each node and virtual client.

Path Ownership Rules (2 of 2)

- To configure NetWorker to ignore path ownership rules, create an empty **pathownerignore** file in the directory containing the NetWorker binaries on each node.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

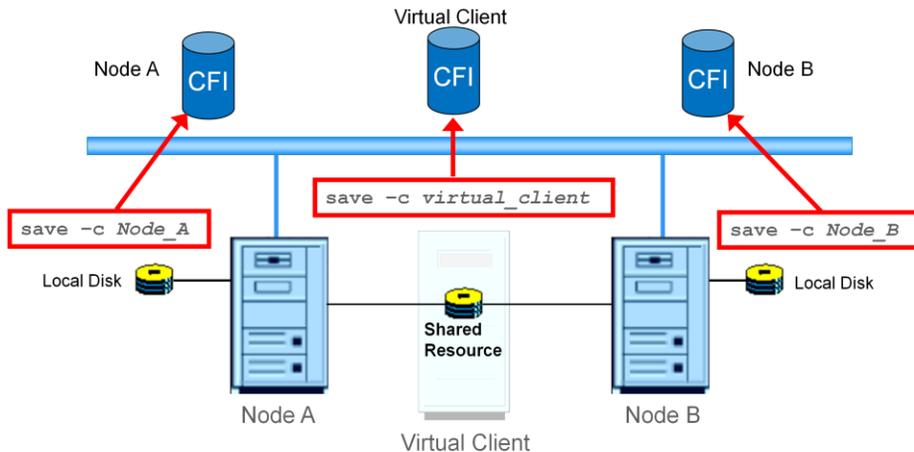
Module 12: VSS Backups and Cluster Environments 27

To ignore path ownership rules, you can create an empty **pathownerignore** file in the directory containing the NetWorker binaries. This file should be created on each node. Its existence forces NetWorker to back up all specified save sets regardless of ownership conflicts.

It is important to realize that creating the **pathownerignore** file is not recommended, but may be necessary if the cluster resources are incorrectly configured. Remember that this file does not override the path ownership rules, it simply ignores them. This may result in tracking information being sent to an incorrect CFI, possibly causing problems when performing browsable recoveries.

Forcing Save Sets to the Correct CFI

- In the NetWorker Client resource, the **Backup command** attribute can be configured to force NetWorker to write to the proper CFI.
- This may be necessary if using the **pathownerignore** file.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

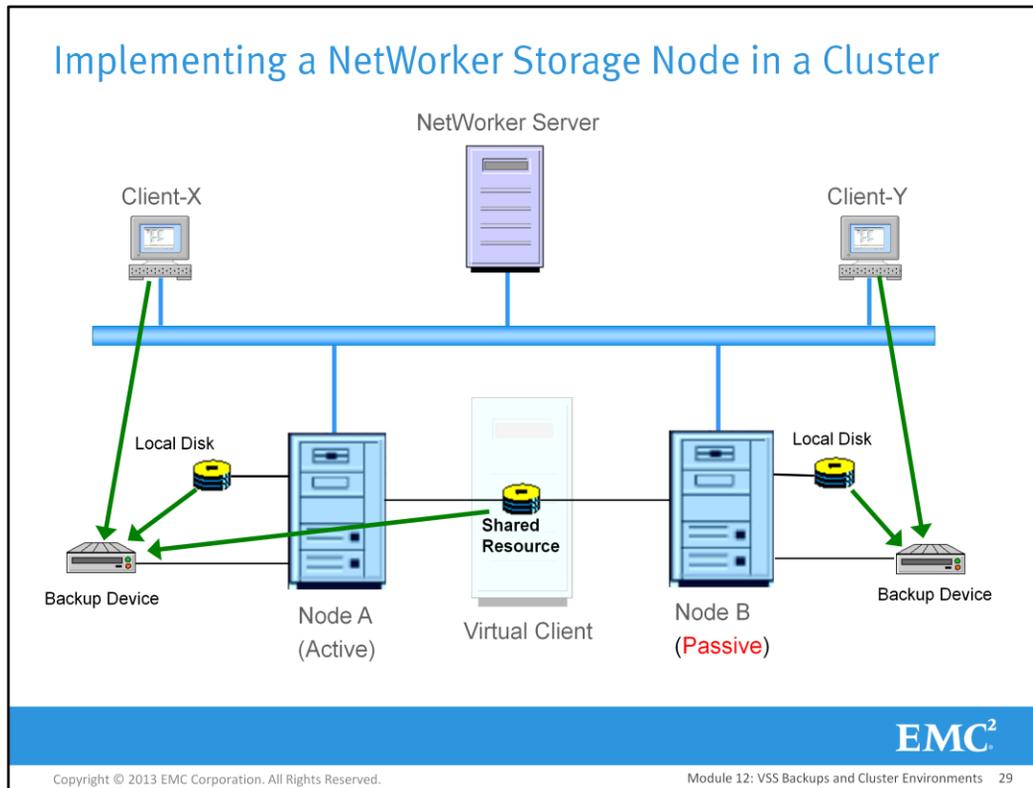
Module 12: VSS Backups and Cluster Environments 28

If you create a **pathownerignore** file, check whether the save set tracking information is written to the correct client file index. If it goes to the wrong CFI, you can force the tracking information to go to a specific client's index.

To force save sets to be written to a specific CFI, it is necessary to modify the **Backup command** attribute of the client whose data is being sent to the incorrect CFI. The following command should be placed in this attribute: `save -c client_name` where *client_name* is the hostname of the client being backed up.

If you are backing up an application server using a NetWorker module, make sure that you are using the `-c client_name` arguments (or similar arguments) required by the NetWorker module. Refer to the applicable module documentation for details on options for the backup command used by each NetWorker module.

Implementing a NetWorker Storage Node in a Cluster



It is often desirable to back up clustered data to devices managed by the cluster nodes, thereby avoiding TCP/IP traffic. NetWorker supports the environment where each node in a cluster is configured as a NetWorker storage node. Backing up a virtual client to a device on the active node requires a simple modification of its client resource.

The important thing to realize in this configuration is that both cluster nodes are functional storage nodes. As shown on the slide, the active node (Node A), backs up its local save sets to its own backup device, and the passive node (Node B), backs up its local save sets to its own backup device. Save sets belonging to the virtual client are backed up by the active node (Node A), to a device controlled by the active node.

Additionally, clients outside the cluster can be configured to direct their save sets to any NetWorker storage node residing within the cluster.

It is also important to recognize that since the storage node is not a shared resource, if either Node A or Node B fails, the *storage nodes* list of each physical or virtual client backing up to the failed node will be consulted to determine where to redirect the backup.

Although some clustering products have the ability to fail over backup devices between nodes, it is beyond the scope of this course.

Configuring a NetWorker Storage Node in a Cluster

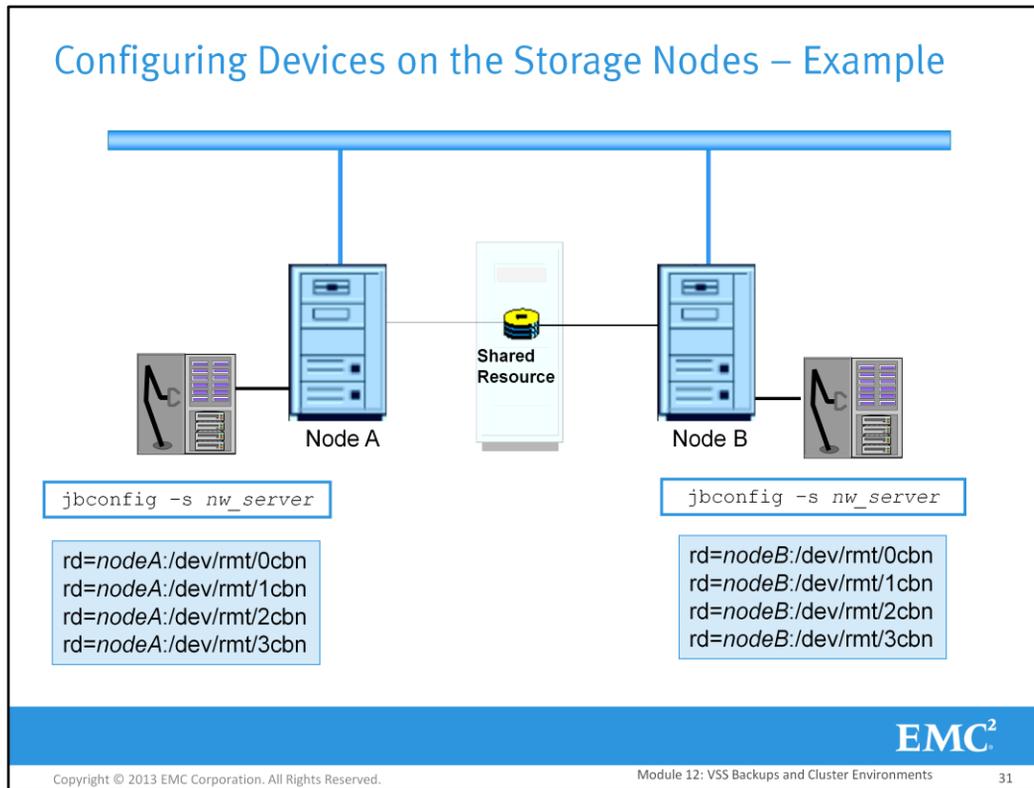
To configure a NetWorker storage node in a cluster:

1. Install NetWorker client and storage node software on each node in the cluster.
2. Configure NetWorker client software as highly available.
3. Create NetWorker client resources for each node and virtual service in the cluster.
4. Configure the **Save set** attribute for each client.
5. Configure the **Remote access** attribute of the virtual client to include all nodes in the cluster.
6. Configure the devices attached to the nodes as remote devices.
7. Enter **curphyhost** as the only value in the storage nodes attribute of the virtual client.



Configuring a NetWorker storage node in a clustered environment is identical to configuring a clustered NetWorker client, except for the additional steps of configuring the remote devices and adding a special hostname, **curphyhost**, to the virtual client's **storage nodes** attribute. "curphyhost" stands for Current Physical Host. On this slide, Steps 2 through 5 are identical to those described earlier for configuring a clustered NetWorker client.

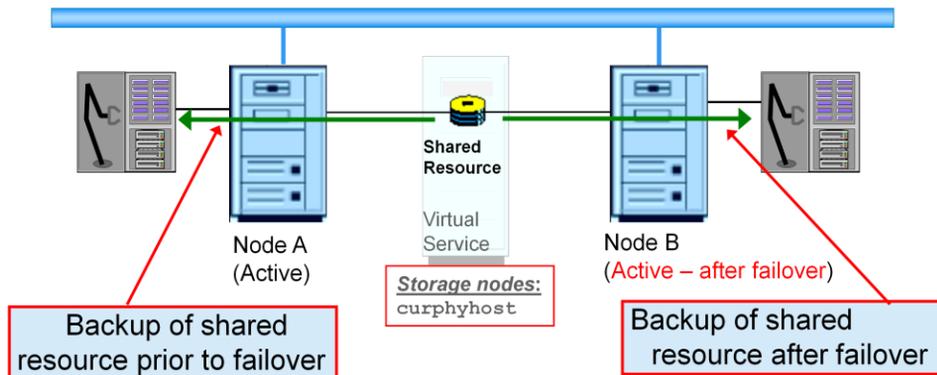
Configuring Devices on the Storage Nodes – Example



Once the storage node software has been loaded on all nodes of the cluster, devices can be configured to the same way as a standalone storage node. The only difference with clustered storage devices is that they are all created as remote devices even though they are local to the node. This means that each device will be created with the **rd=storage_node:pathname** syntax.

Configuring the Storage Nodes Attribute: curphyhost

- A value of **curphyhost** in the **storage nodes** attribute of the virtual client resource resolves to the current physical host (node) where the virtual service is running.
- It causes data backed up from the virtual client to always be sent to a device managed by the active node, regardless of which node is active.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 32

In the case of a dedicated storage node, it is important to ensure the data is passed to the active node of the cluster. This ensures the data is sent directly to the storage device and not out over the TCP/IP network to the passive node.

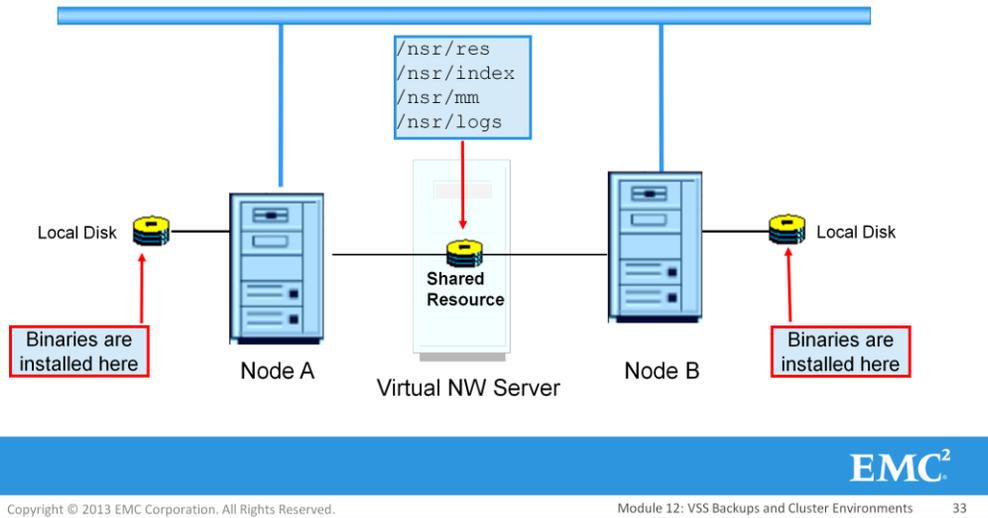
To ensure that data is backed up to the active node of a clustered dedicated storage node NetWorker provides a special keyword called **curphyhost**. This keyword is placed in the **storage nodes** attribute of the virtual storage node's client resource.

If the **curphyhost** keyword is not listed in the virtual client's **storage nodes** attribute, the save stream is directed to a device on Node A. If the virtual client fails over to Node B and a backup of the virtual client is performed, the save stream is still saved to a device on Node A because it is still the first entry in the **storage nodes** list. As a result, the save stream is sent over the TCP/IP network.

Placing **curphyhost** first in the virtual client's **storage nodes** attribute modifies NetWorker operations so that if a failover to Node B occurs and a backup of the virtual client is performed, the save stream is directed to a device on Node B. This action takes place because, after the failover, the virtual service is being managed by Node B, the "current physical host."

Implementing a NetWorker Server in a Cluster

- NetWorker client, storage node, and server software are installed on both nodes.
- NetWorker files, with the exception of binaries, are configured as a shared resource (virtual service).

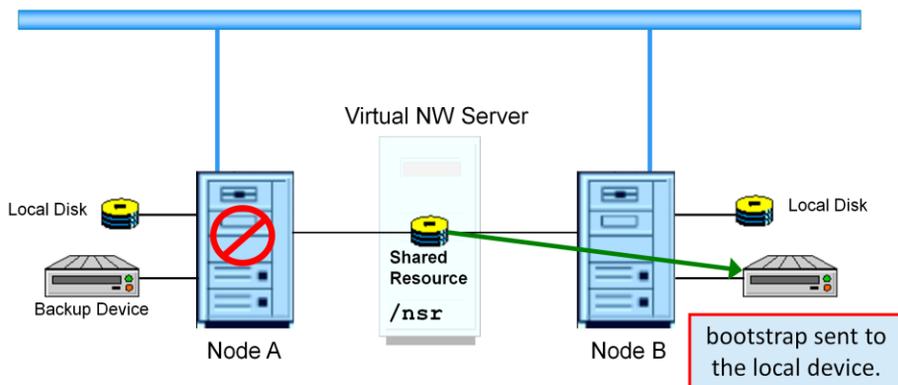


When the NetWorker server is configured in a clustered environment, all NetWorker files, with the exception of the binaries, are located on the shared disk resource.

With these directories being part of the shared resource, if the NetWorker server virtual service fails over to Node B, Node B will have access to all the NetWorker resource files, logs, databases, and licensing information.

Failover of a NetWorker Server in a Cluster

- If Node A fails, Node B starts running the NetWorker server virtual service.
- The bootstrap save set contains the resource files and media database that reside on the virtual NetWorker server.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments

34

Even if the NetWorker server is clustered failover of the cluster is still an interruptive event; meaning that if a group backup was in process when the failover occurred, the same group would be restarted after the failover, potentially directing save streams to different devices than before the failover.

Just as with any NetWorker server, each node in the cluster must have a locally-managed backup device to which the bootstrap save set is saved.

When a group is restarted, it behaves like a manually restarted group. Save sets backed up successfully will not be saved again. Clients that have already failed will not be retried if the **client retries** value has already been met. Save sets being backed up when the failover occurred will be restarted from the beginning. Save sets that have not started will be saved.

If `savegrp` is in the process of saving the bootstrap save set, the group will not be restarted. Additionally if NDMP file history processing is in progress when the failover occurs it will not restart.

If a recovery is in progress, you must manually restart the recovery.

In some cases the requested tape may reside in a device controlled by the failed node. If this is the case, you will need to manually eject the tape and move it to another device.

It is important to note two requirements for automatically restarting a group backup: **Autorestart** must be enabled in the group resource and the **Options** attribute in the group resource must not have manual restart set.

Configuring a NetWorker Server in a Cluster

To configure a clustered NetWorker server:

1. Install the NetWorker client, storage node, and server software on each node in the cluster.
2. Configure NetWorker client and server software as highly available.
3. Create a NetWorker client resource for each node.
4. Configure the **Save set** attribute for each client.
5. Set the **Remote access** field of the virtual client to include each node within the cluster.
6. Edit the NetWorker server **administrators** list to include access for both nodes as well as for the virtual client.
7. Update the **servers** file on all NetWorker clients in the data zone to include all hostnames and aliases of the NetWorker server (virtual client and physical nodes).
8. Create necessary jukebox and device resources.
9. Modify the **storage nodes** attribute of the NetWorker server's virtual client to include both nodes and the value **nsrserverhost**.



This slide shows the steps necessary to configure a NetWorker server in a cluster. These steps are covered on the following slides.

Configuring a NetWorker Server as Highly Available

Operating System	Setup Script
EMC FullTime AutoStart	<code><networker_bin_dir>/networker.cluster</code>
IBM HACMP	<code>/usr/bin/networker.cluster</code>
HP MC/ServiceGuard or LockManager Cluster	Configure the <code>.nsr_cluster</code> file. <code>/opt/networker/bin/networker.cluster</code>
Sun Cluster	<code>/usr/sbin/networker.cluster</code>
Veritas Cluster	<code>/usr/sbin/networker.cluster</code>
TruCluster	<code>/usr/opt/networker/bin/networker.cluster</code>
Microsoft Windows MSCS	<code>regcnsrd [-c -r]</code> Configure manually using Microsoft Cluster Administrator.

Note: Make sure cluster services are running on both nodes.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments

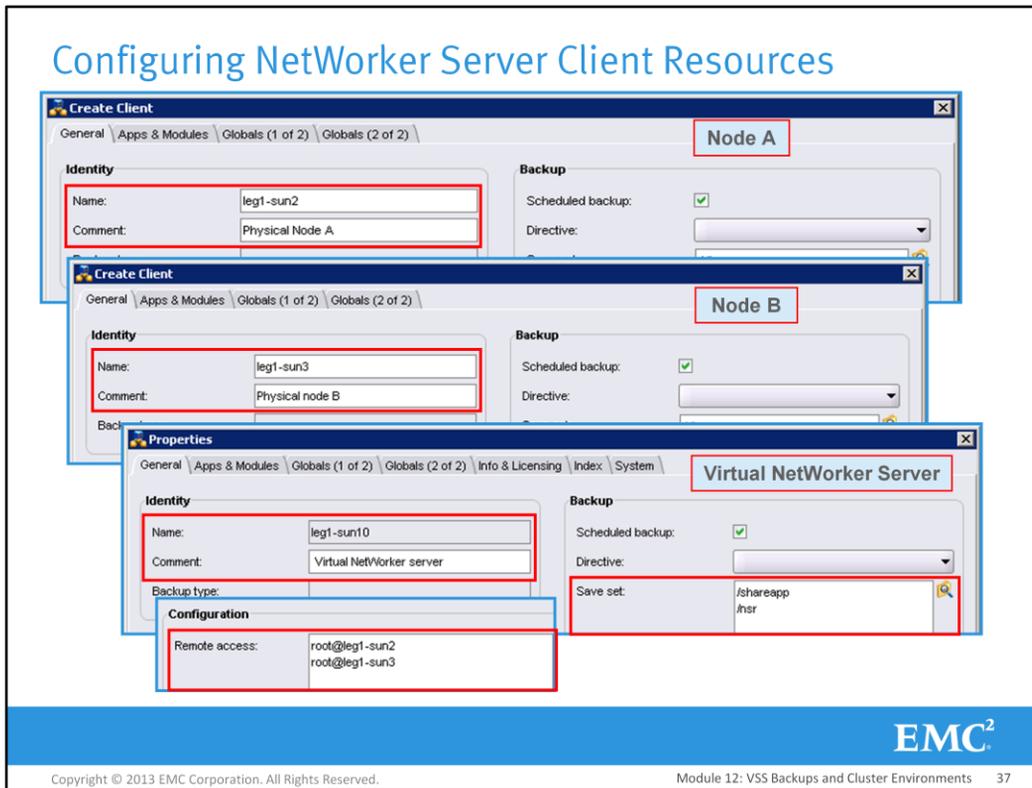
36

To make a NetWorker server highly available, the cluster software needs to be aware of NetWorker. To accomplish this, the NetWorker software includes a script called `networker.cluster` which must be run to configure NetWorker within the clustered environment. This slide lists the location of the script for each supported clustering product. This is the same script used to configure a clustered NetWorker client.

In a Microsoft MSCS environment, there is no `networker.cluster` script. `regcnsrd` is used to create the NetWorker server cluster resource type and the remaining configuration is done using the Microsoft Management Console (MMC) Cluster Administrator snap-in program.

As the `networker.cluster` script runs, you are asked to specify a shared directory to which all the NetWorker resources, tracking data, and logs are moved. During configuration of the shared `/nsr` directory, a symbolic link named `/nsr.NetWorker.local` is created and points to the NetWorker files on the node, while `/nsr` points to the shared directory on the virtual client.

Configuring NetWorker Server Client Resources



The virtual client for the NetWorker server, is the NetWorker server for the data zone while the nodes are NetWorker clients.

Ensure that the entire cluster is backed up by configuring the following items. First, create at least one NetWorker client resource for each node and virtual client in the cluster.

Secondly, configure the **Save set** attribute of each client. Ensure that the NetWorker servers virtual client is backing up all shared data. Since there are files under `/nsr` that are not backed up by the bootstrap and CFI save sets, make sure that the **Save set** attribute of the virtual client includes the `/nsr` save set as well as any shared data not related to NetWorker.

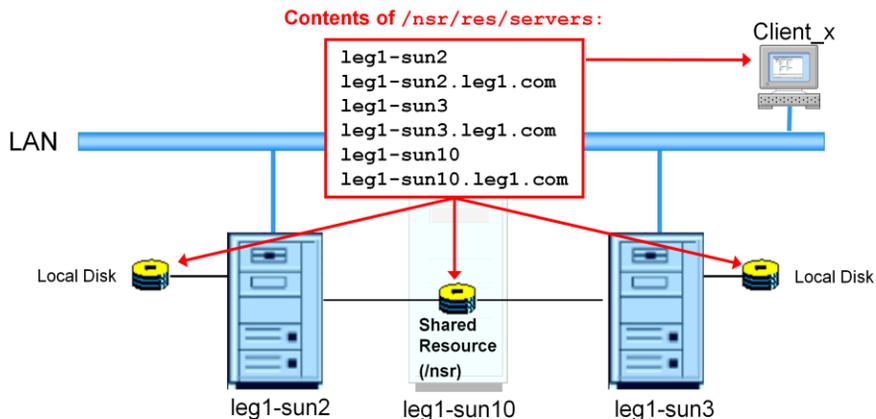
Finally, configure the **Remote Access** attribute of the virtual client to include the hostname of each node. This allows the nodes to access the virtual client's CFI during recoveries.

As with clustered clients you must modify the **Administrator** attribute of the virtual NetWorker server resource to include administrative users on the nodes.

For UNIX systems, include the root user from each node. For Windows systems, include the **SYSTEM** account from each node.

Updating the NetWorker Server File

- The **server** file on every NetWorker client in the data zone must contain the hostname and aliases of the virtual NetWorker server and both nodes.
- The server file on the shared `/nsr` directory must also be updated.



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 38

When configuring a NetWorker server in a cluster, three hostnames are associated with the cluster: the hostname of the virtual NetWorker server and the hostnames of both nodes. Additionally, each of these hosts may have aliases, such as fully qualified domain names.

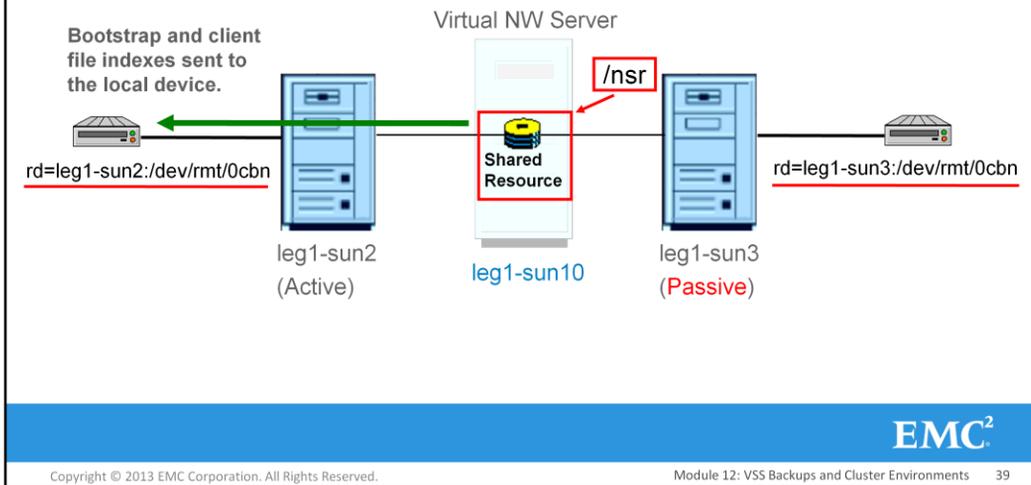
You must add all of these hostnames and aliases, one hostname per line, to the **server** file on every NetWorker client within the data zone. The **server** file should reside in the NetWorker resource directory on each client.

After modifying the **server** file, it is necessary to restart all NetWorker daemons or services on each machine, thereby causing **nsrexecd** to read the file.

To ensure that the **server** file is used, check the NetWorker boot-time startup mechanism for **nsrexecd** to verify that **the nsrexecd** command does NOT include any “-s” options; if there is a “-s” option, the **server** file is ignored. In Microsoft Windows, look in the **NetWorker Remote Exec** service properties tab. In most UNIX environments, you can look in the NetWorker `run-control` script.

NetWorker Server Local Device Support

The virtual NetWorker server treats devices attached to the active node as local devices, even though they are configured in NetWorker as remote devices.



Just like with clustered storage nodes all devices on the virtual NetWorker server must be created as remote devices on each of the physical nodes within the cluster. This is also true for any jukeboxes that will be attached to the NetWorker server, these must also be created as remote jukeboxes

The virtual NetWorker server always treats a remote device managed by the active node as a local device. The device managed by the passive node is perceived as a remote device.

Configuring Storage Nodes Attribute for NetWorker Server

Properties

General | Apps & Modules | Globals (1 of 2) | Globals (2 of 2) | Info & Licensing | Index | System

Configuration

Remote access: root@leg1-sun2
root@leg1-sun3

Index path:

Executable path:

Storage nodes: leg1-sun2
leg1-sun3
nsrserverhost

Windows Client

Hard links:

Short filenames:

BMR

BMR:

BMR options:

Archive Management

Archive services:

Archive users:

For proper backup of bootstrap and indexes, the **storage nodes** attribute for the virtual NetWorker server must include the hostnames of both nodes.

A clustered NetWorker server does **not** support the **curphyhost** attribute value.

EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments

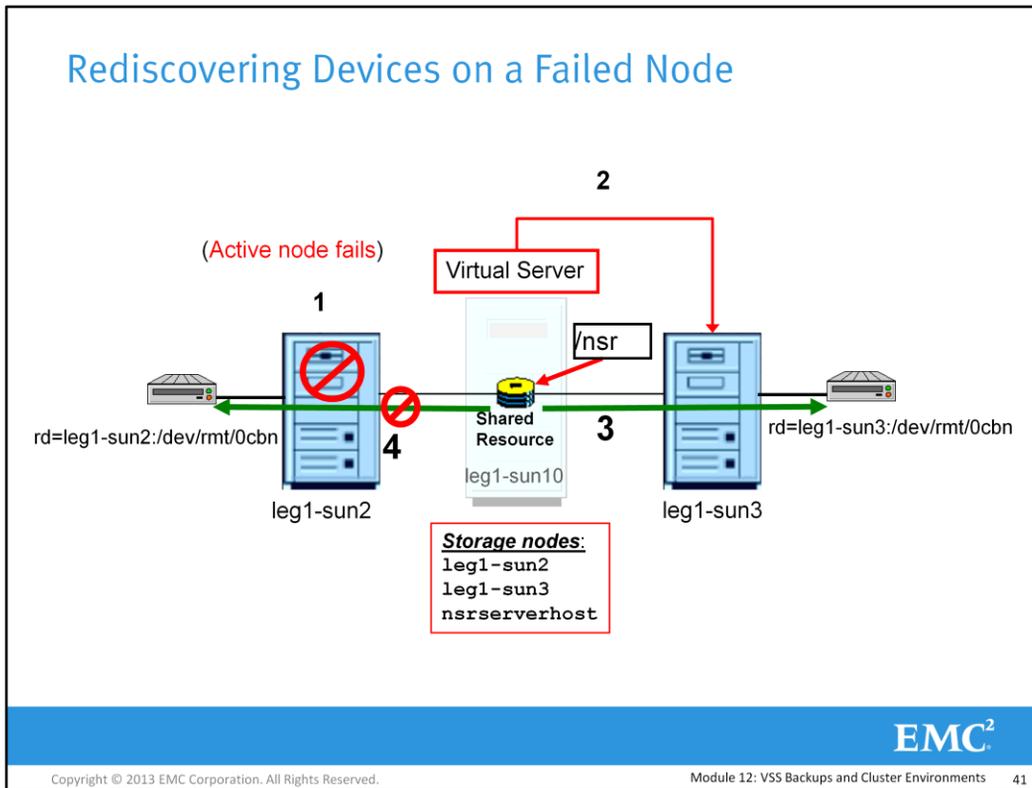
40

A virtual NetWorker server running in a two-node cluster requires that you configure each cluster node as a storage node for the virtual NetWorker server by creating a device or autochanger resource.

When configuring the virtual NetWorker server resource, ensure that the **Storage nodes** attribute in the virtual NetWorker server client resource lists the storage nodes in the following order: Node A which is the node that is normally active, Node B, and finally, nsrserverhost.

As a word of caution, do not use the **curphyhost** keyword in the virtual NetWorker server's client resource. This can cause undesirable behavior such as the bootstrap and index save sets being written to a device on a remote storage node, rather than to a local device on the NetWorker server.

Rediscovering Devices on a Failed Node



After failover occurs in a clustered NetWorker server environment, if the failed node becomes available as a passive node, a bootstrap backup may fail. This is because devices managed by a passive node are treated as remote devices by the NetWorker server. The following scenario is depicted in this slide:

In Step 1, the active node, leg1-sun2, fails. In step 2, the virtual NetWorker server service fails over to leg1-sun3. In step 3, as long as leg1-sun2 remains offline, leg1-sun3 is the first available storage node in the **Storage nodes** list. Since leg1-sun3 is the active node, and devices managed by the active node are recognized as local devices, the bootstrap backup succeeds.

In Step 4, leg1-sun2 comes back online as a passive node and a bootstrap backup occurs. The backup fails because devices managed by the passive node are treated as remote devices by the NetWorker server. Since leg1-sun2 is still the first storage node in the storage nodes list, the NetWorker server attempts to send the bootstrap save set to the device on leg1-sun2. However, since the bootstrap save set must be saved to a local device, the backup fails.

To correct this situation, it is crucial that you manually change the ordering of the **Storage nodes** list in the NetWorker server client resource so that the current active node is listed first.

NetWorker Cluster Licensing

- Clusters are licensed by physical node only.
 - ▶ Cluster client connection enabler is no longer required.
- Application module license is required for the client application the backup will be configured on.
 - ▶ If application is backed up from virtual node, module license is required for virtual node.
 - ▶ If application is backed up from physical node, module license is also required for physical node.
 - ▶ It is typically only backed up from virtual node so only one application license required.
- If using DDS, licensing is by drive to be shared.
 - ▶ One license per drive independent of storage node and library licensing.



Cluster client licensing is based solely on the client connection enablers, there is not specific cluster license required to enable backup of a clustered client.

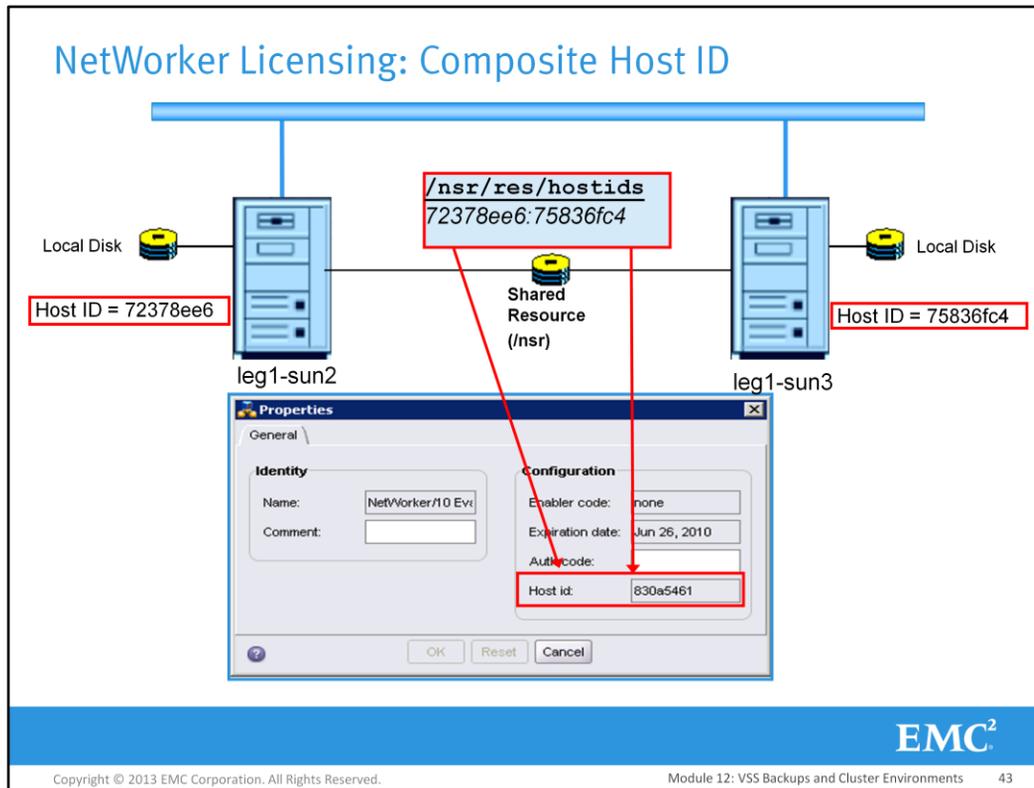
Virtual clients of a cluster do not require their own license. To backup with a NetWorker module from the virtual client, an application module license is required for the virtual client. Similarly, from the physical host, a separate application module license is required for the physical client computer.

To enable Dynamic Drive Sharing (DDS), a separate license is required for each drive that is to be shared. Once a drive is licensed, any number of storage nodes can be shared. DDS licensing is independent of library and storage node licensing.

Notes:

Refer to the *EMC NetWorker Licensing Guide* and the latest product release notes for current information about software licensing.

NetWorker Licensing: Composite Host ID



EMC²

Copyright © 2013 EMC Corporation. All Rights Reserved.

Module 12: VSS Backups and Cluster Environments 43

NetWorker licenses are tied to the NetWorker server's host ID. However, a virtual host does not have a host ID. Thus, when the NetWorker server is clustered, it is necessary to create a composite host ID. This is generated by combining the host IDs of all nodes in the cluster.

To create the composite host ID, determine the host IDs of all nodes in the cluster and enter them into the `hostids` file in the NetWorker resource directory. All of the host IDs should be placed on a single line, separated by colons. The line should not contain any spaces.

It is necessary to restart the NetWorker server daemons or services after creating or modifying the `hostids` file.

When creating a License resource, NetWorker checks for the existence of the `hostids` file and, if it exists, automatically generates the composite host ID from its contents.

The resulting composite host ID should be used when authorizing licenses on PowerLink.

Module 12: VSS Backups and Cluster Environments

Lesson 2: Summary

During this lesson the following topics were covered:

- Cluster components and characteristics
- Configuring cluster clients
- Managing path ownership issues
- Configuring storage nodes and NetWorker servers in a cluster



This lesson covered backup and recovery of clusters as well as the configuration of storage nodes and NetWorker servers in a cluster. Topics include cluster components and characteristics; the procedures for configuring cluster clients, storage nodes, and NetWorker servers; and the management of path ownership with clusters.

Module 12: Summary

Key points covered in this module include:

- Microsoft Volume Shadow Copy Service and its components
- How VSS works
- Backup and recovery of VSS SYSTEM save sets
- Configuring a NetWorker client in a cluster
- Recovering a clustered client's data
- Configuring cluster nodes as NetWorker storage nodes
- Using a NetWorker server in a cluster



This module covered an overview of Microsoft Volume Shadow Copy Service (VSS), backup and recovery of VSS SYSTEM save sets, and backing up and recovering with NetWorker in a cluster environment.

This slide is intentionally left blank.

