Failover control

No cloud solution would be workable without a viable disaster recovery solution. Virtualized workloads owned by business units in large enterprises or by customers of cloud hosting providers must be backed up regularly to prevent loss of continuity should a disaster occur on the provider's infrastructure. This chapter ends with a look at Hyper-V Replica, a new feature of Hyper-V in Windows Server 2012 that helps ensure that your cloud solutions can be recovered in the event of a disaster.

Hyper-V Replica

While many third-party backup solutions can be used for backing up and recovering VMs running on Hyper-V hosts, the Hyper-V Replica feature in Windows Server 2012 provides an in-box business continuity solution for cloud environments that can efficiently, periodically, and asynchronously replicate VMs over IP-based networks, including slow WAN links and across different types of storage subsystems. The Hyper-V Replica feature does not require any shared storage or expensive storage array hardware, so it represents a low-cost solution for organizations looking to increase the availability of their virtualized workloads and ensure that these workloads can be recovered quickly in the event of a disaster.

Hyper-V, together with Failover Clustering, allows VMs to maintain service availability by moving them between nodes within the datacenter. By contrast, Hyper-V Replica allows VMs to maintain availability across a datacenter where the node hosting the replica is located at a physically separate site. Hyper-V Replica provides host-based replication that allows for failover to a secondary datacenter in the event of a disaster. It's an application-agnostic solution because it operates at a VM level regardless of what guest operating system or applications are installed in the VM. It's a storage-agnostic solution because you can use any combination of SAN, direct attached storage (DAS), or SMB storage for storing your VMs. It also works in both clustered and nonclustered environments, and you can even replicate from a host on a shared cluster to a remote, stand-alone replica host. And it works with Live Migration and Live Storage Migration.

Typical cases for using Hyper-V Replica might include:

- Replicating VMs from head office to branch office or vice versa in large and mid-sized business environments
- Replication between two datacenters owned by a hosting provider to provide disaster recovery services for customers
- Replication from the premises of small and mid-sized businesses to their hosting provider's datacenter

Implementing Hyper-V Replica

Hyper-V Replica can be enabled, configured, and managed from either the GUI or by using Windows PowerShell. Let's briefly look at how to enable replication of a VM by using Hyper-V Manager. Begin by selecting the Replication Configuration section in Hyper-V Settings on the hosts that you plan on replicating VMs to or from. Select the Enable This Computer As A Replica Server check box to enable the host as a replica server and configure the authentication, authorization, and storage settings that control the replication process:

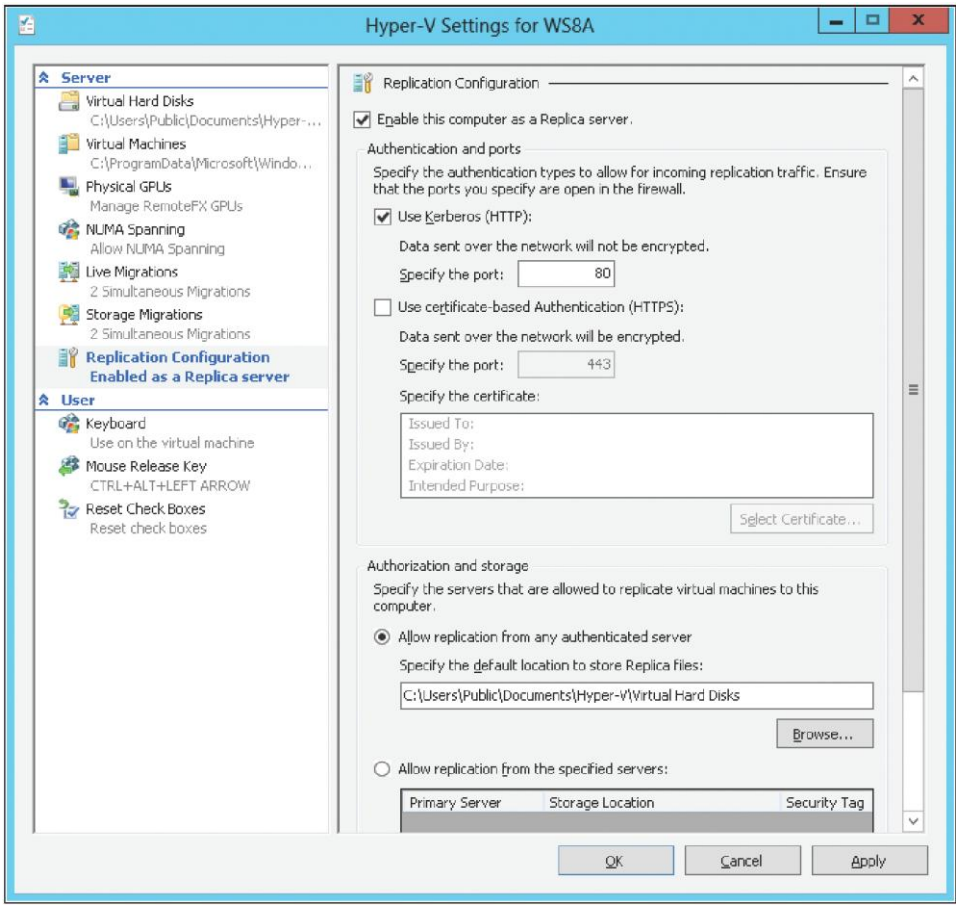

Once you've performed this step on both the primary and replica servers (the primary server hosts the virtualized production workloads, whereas the replica server hosts the replica VMs for the primary server), you then can enable replication on a per-VM basis. To do this, right-click a VM in Hyper-V Manager and select Enable Replication.

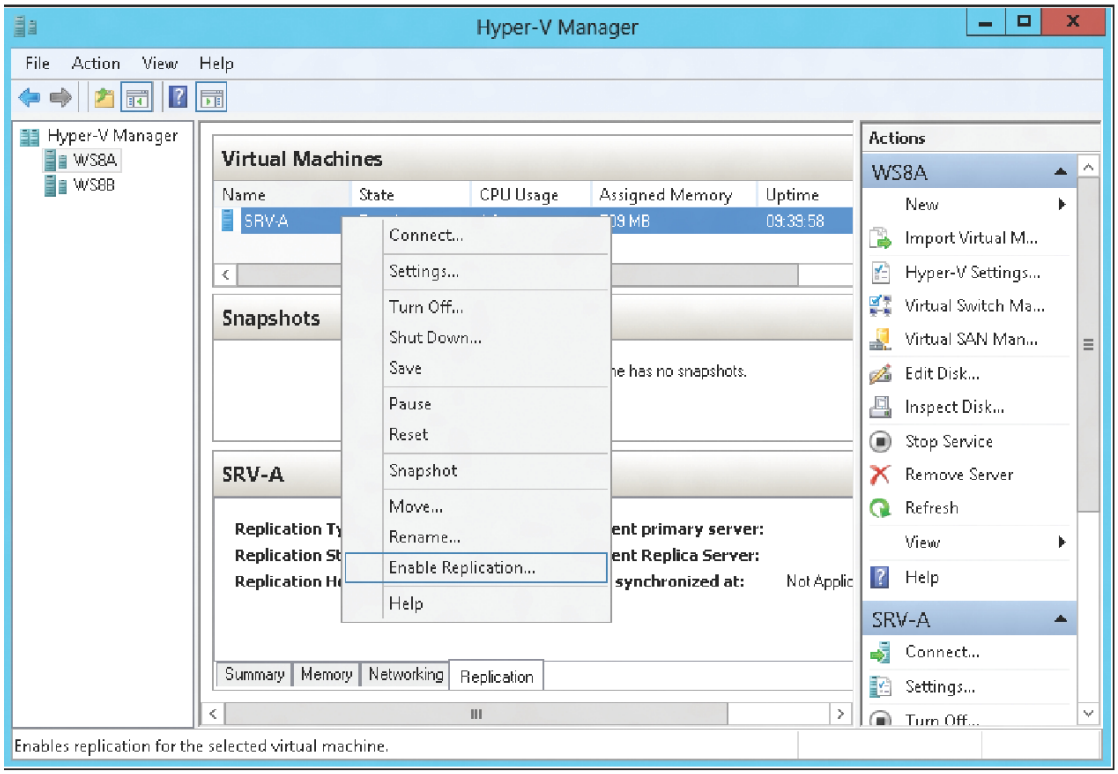

When the Enable Replication wizard launches, specify the name of the replica server that you want to replicate the selected production VM to:

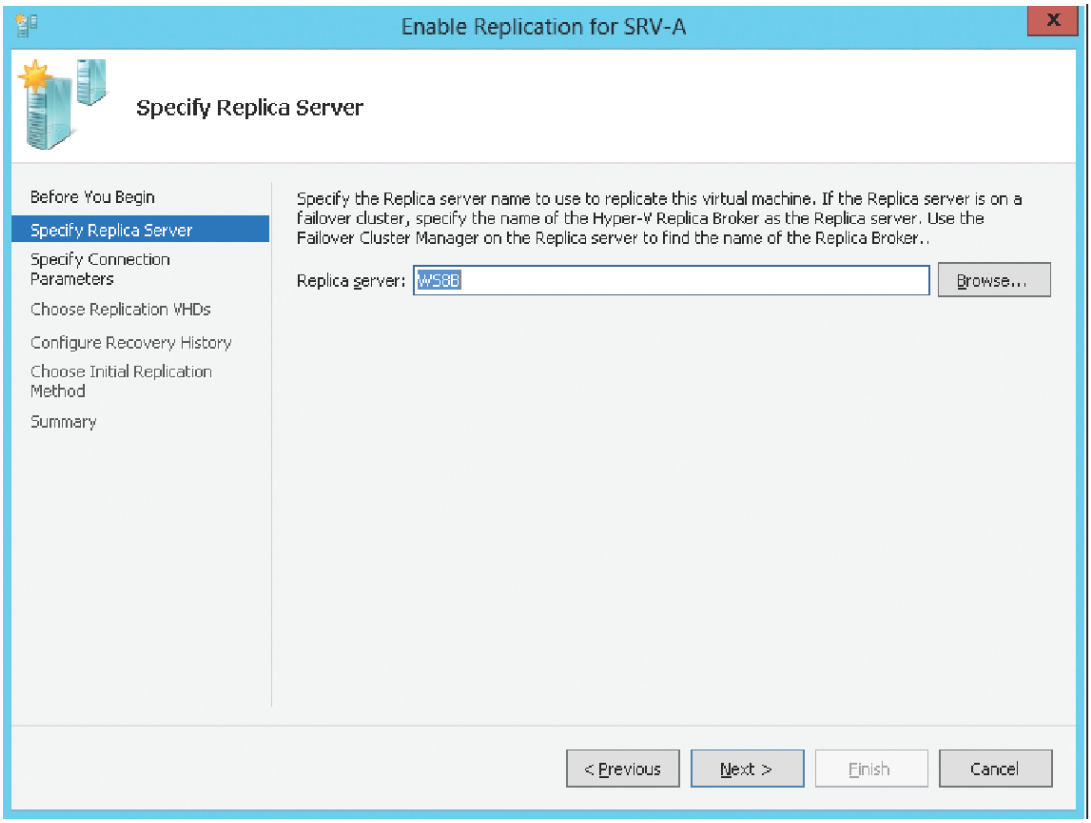

Specify connection parameters that define the port and authentication method used for performing replication:

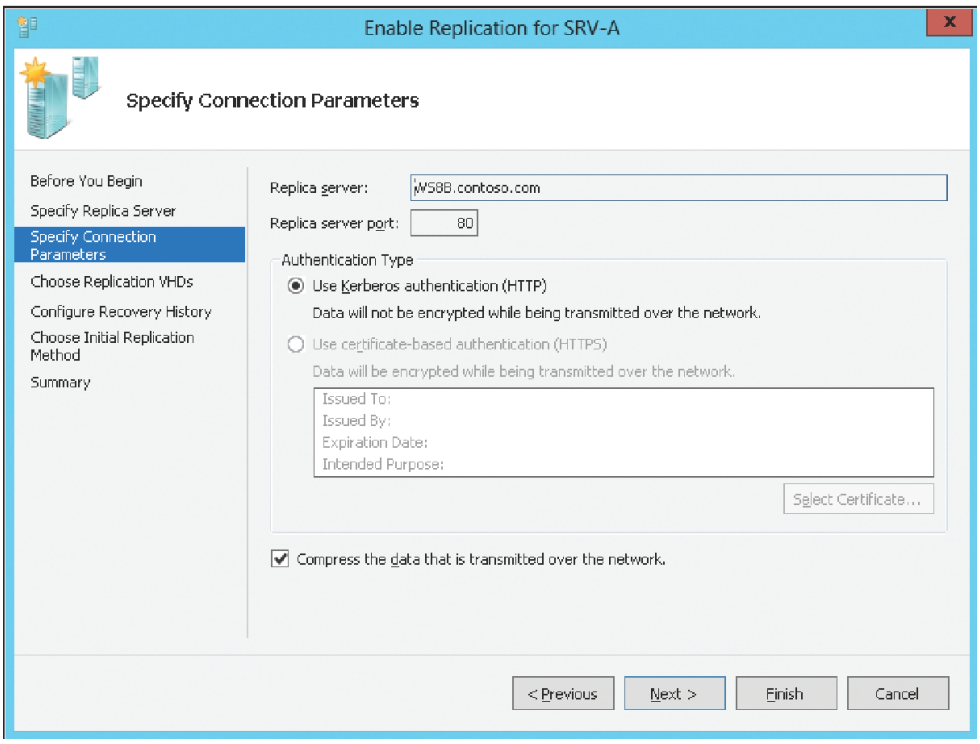

Continue through the wizard until you reach the Choose Initial Replication Method page, where you specify how and when the VM first will be copied over to the replica server:

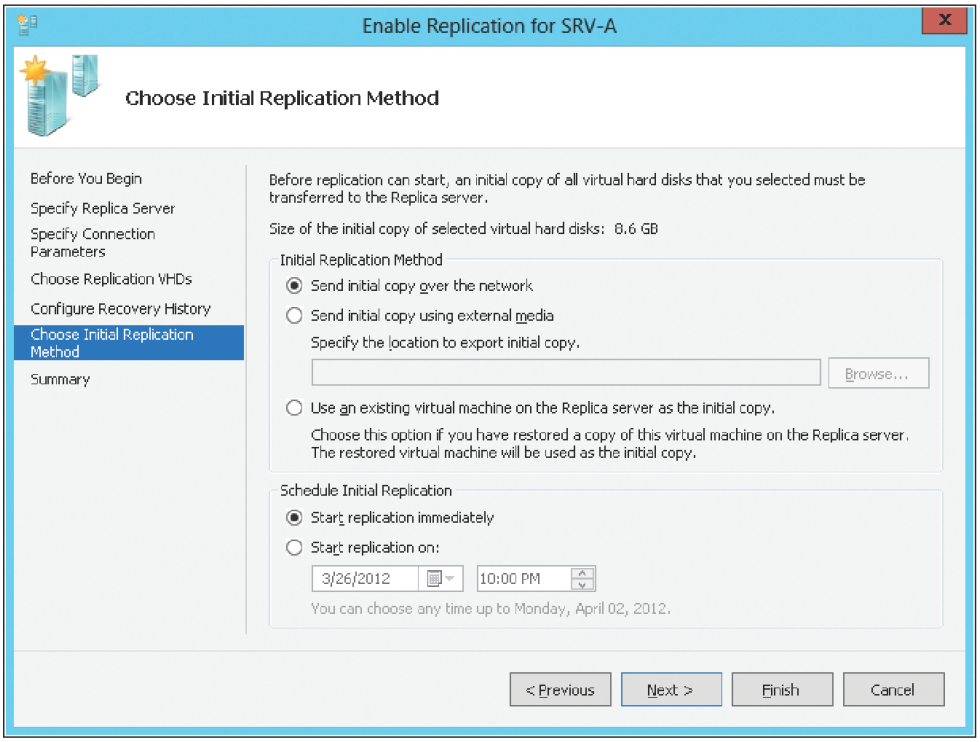

Once you've completed the wizard and clicked Finish, replication will begin. You can view the replication process as it takes place by selecting the Replication tab in the bottom-central pane of Hyper-V Manager:

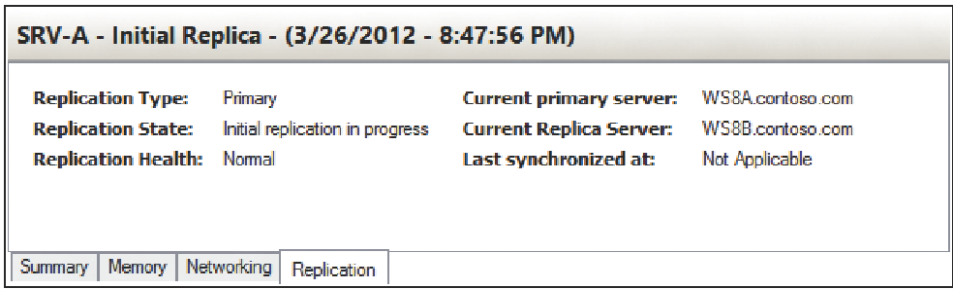

You also can use the Measure-VMReplication cmdlet in Windows PowerShell to view the success or failure of the replication process:

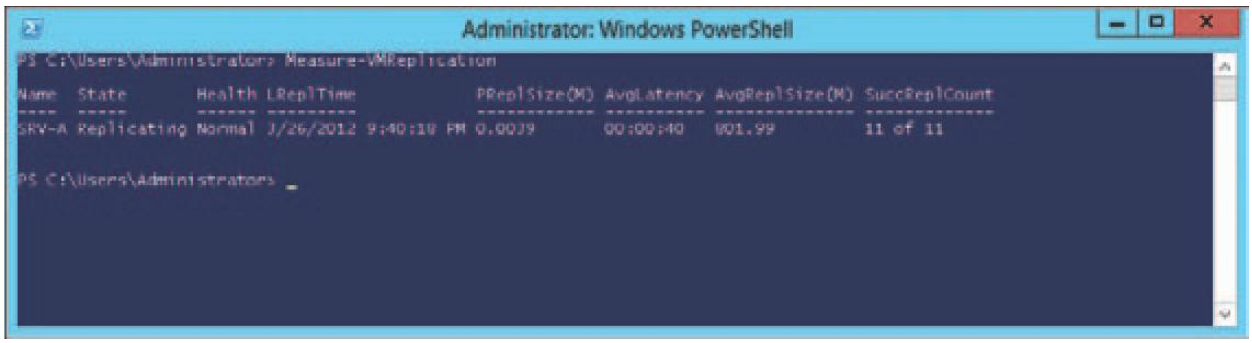

To view all the Windows PowerShell cmdlets for managing the Hyper-V Replica feature, use the Get-Command cmdlet, as shown here:

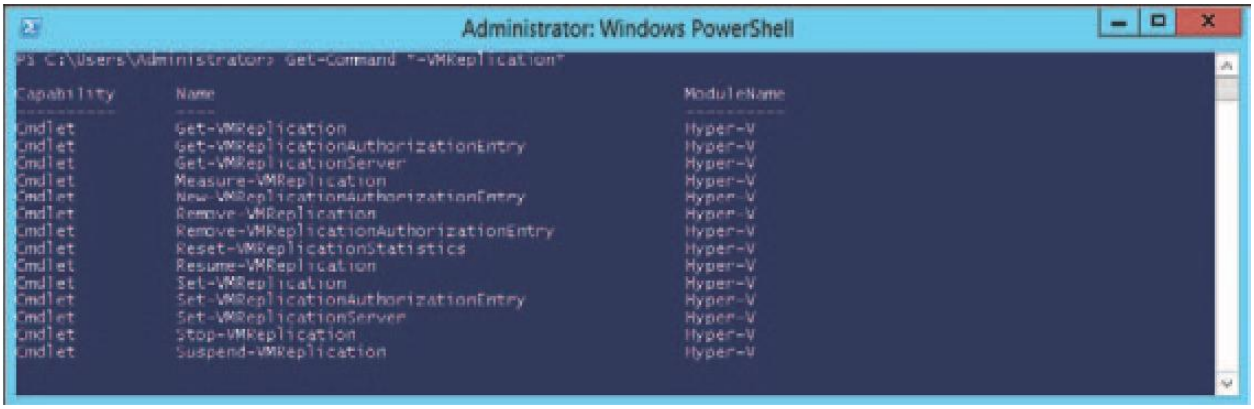

Continuous availability

Guaranteeing continuous availability of applications and services is essential in today's business world. If users can't use the applications they need, the productivity of your business will be affected. And if customers can't access the services your organization provides, you'll lose their business. Although previous versions of Windows Server have included features like Failover Clustering and NLB that help you ensure the availability of business-critical applications and services, Windows Server 2012 adds a number of improvements that can greatly help ensure application uptime and minimize service disruptions.

Key availability improvements include enhancements to Failover Clustering such as greater scalability, simplified updating of cluster nodes, and improved support for guest clustering. The new SMB 3.0 Transparent Failover capability lets you perform maintenance on your cluster nodes without interrupting access to file shares on your cluster. Storage Migration now allows you to transfer the virtual disks and configuration of VMs to new locations while the VMs are still running. Windows NIC Teaming now provides an in-box solution for implementing

fault tolerance for the network adapters of your servers. Improvements to Chkdsk greatly reduce potential downtime caused by file system corruption on mission-critical servers. Easy conversion between installation options provides increased flexibility for how you configure servers in your environment, whereas Features On Demand lets you install Server Core features from a remote repository instead of the local disk. And DHCP failover improves resiliency by allowing you to ensure continuous availability of Dynamic Host Configuration Protocol (DHCP) services to clients on your network.

Failover Clustering enhancements

Failover Clustering is a feature of Windows Server that provides high availability for server workloads. File servers, database servers, and application servers are often deployed in failover clusters so that when one node of the cluster fails, the other nodes can continue to provide services. Failover Clustering also helps ensure workloads can be scaled up and out to meet the demands of your business.

Although the Failover Clustering feature of previous versions of Windows Server provided a robust solution for implementing high-availability solutions, this feature has been significantly enhanced in Windows Server 2012 to provide even greater scalability, faster failover, more flexibility in how it can be implemented, and easier management. The sections that follow describe some the key improvements to Failover Clustering found in Windows Server 2012. Note that some other cluster-aware features, such as concurrent Live Migrations and Hyper-V Replica, were discussed previously in Chapter 2, "Foundation for building your private cloud."

Increased scalability

Failover Clustering in Windows Server 2012 now provides significantly greater scalability compared to Windows Server 2008 R2 by enabling you to do the following:

- Scale out your environment by creating clusters with up to a maximum of 64 nodes, compared to only 16 nodes in the previous version.
- Scale up your infrastructure by running up to 4,000 VMs per cluster and up to 1,024 VMs per node.

These scalability enhancements make Windows Server 2012 the platform of choice for meeting the most demanding business needs for high availability.

CSV2 and scale-out file servers

Version 1 of Cluster Shared Volumes (CSV) was introduced in Windows Server 2008 R2 to allow multiple cluster nodes to access the same NTFS-formatted volume simultaneously. A number of improvements have been made to this feature in Windows Server 2012 to make it easier to configure and use a CSV and to provide increased security and performance.

For example, a CSV now appears as a single consistent file namespace called the CSV File System (CSVFS), although the underlying file system technology being used remains NTFS. CSVFS also allows direct I/O for file data access and supports sparse files, which enhances performance when creating and copying VMs. From the security standpoint, a significant enhancement is the ability to use BitLocker Drive Encryption to encrypt both traditional failover disks and CSVs. And it's also easier now to back up and restore a CSV with in-box support for CSV backups provided by Windows Server Backup. Backups of CSV volumes no longer require redirected I/O in version 2. The volume snapshots can be taken on the host that currently owns the volume, unlike version 1, where they were taken on the node requesting the backup. Configuring a CSV can now be performed with a single right-click in the Storage pane of Failover Cluster Manager.

CSV2 also supports the SMB 3.0 features described in the previous chapter, making possible scale-out file servers that can host continuously available and scalable storage. Scale-out file servers are built on top of the Failover Clustering feature of Windows Server 2012 and the SMB 3.0 protocol enhancements. Scale-out file servers allow you to scale the capacity of your file servers upward or downward dynamically as the needs of your business change. This means you can start with a low-cost solution such as a two-node file server, and then later add additional nodes (to a maximum of four) without affecting the operation of your file server.

Scale-out file servers can be configured by starting the High Availability Wizard from Failover Cluster Manager. Begin by selecting File Server from the list of cluster roles (formerly called clustered services and applications):

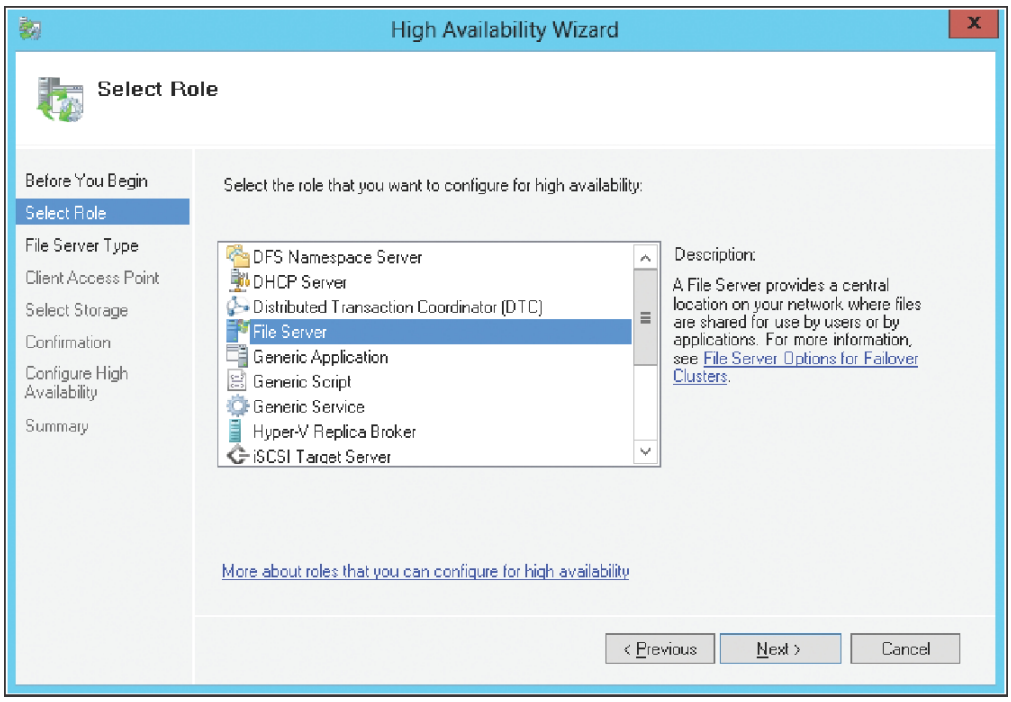

Then, on the next page of the wizard, select the File Server For Scale-Out Application Data option, as shown here, and continue through the wizard:

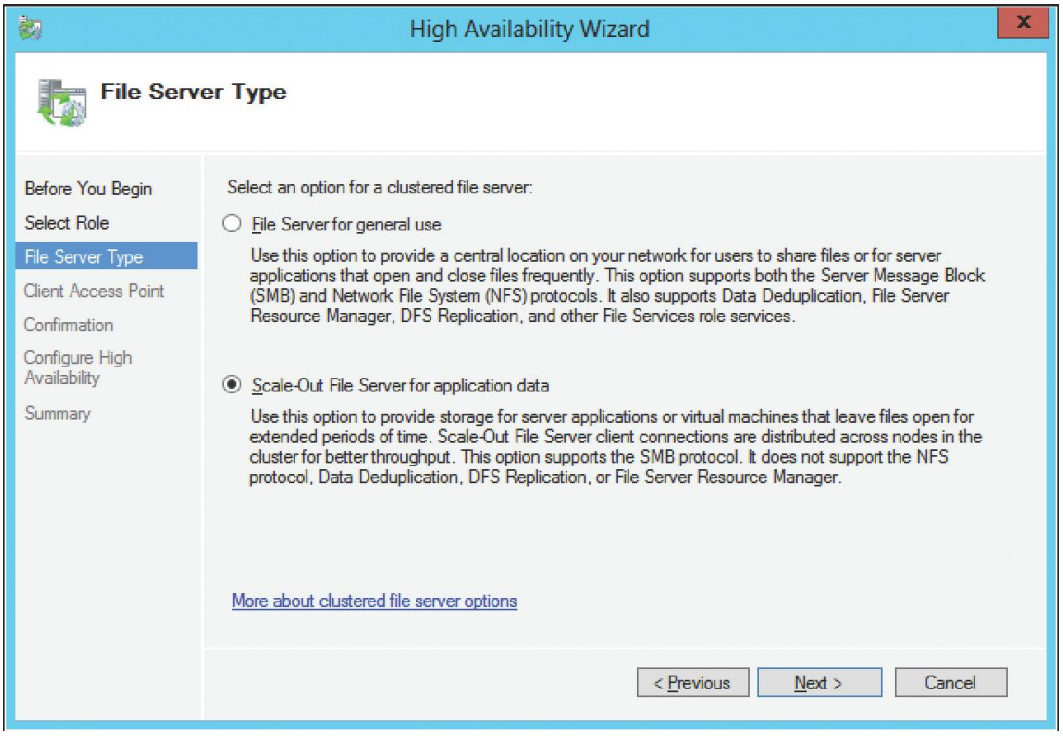

When the wizard executes, a series of steps is performed to create the scale-out file server. These steps are summarized in a report that the wizard generates:

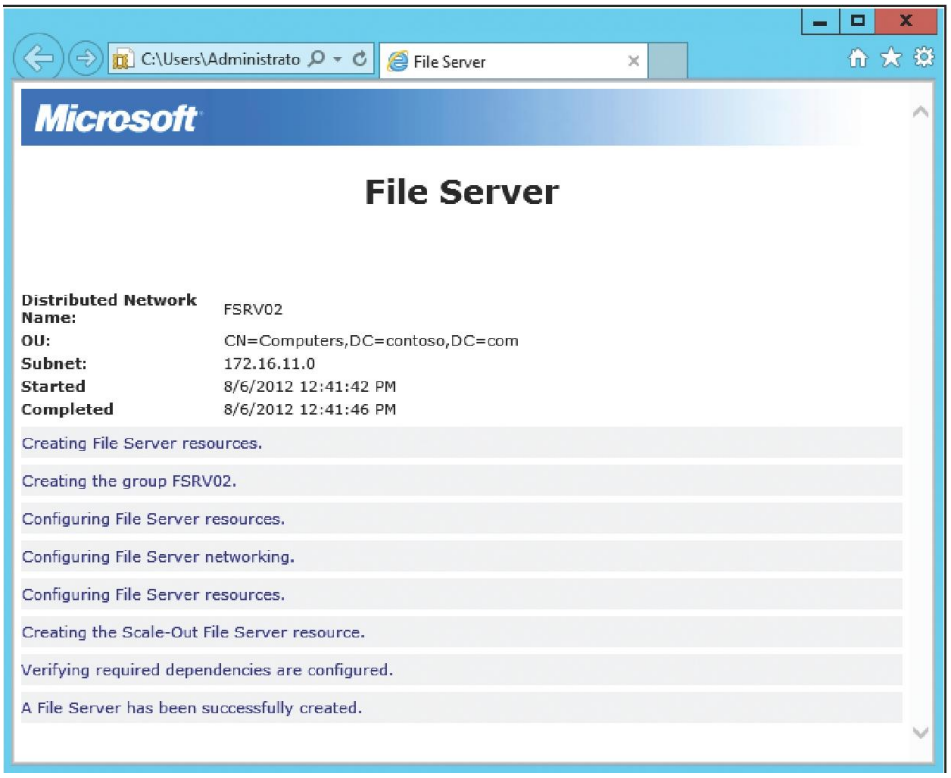

Scale-out file servers have a few limitations that general-use file servers don't have. Specifically, scale-out file servers don't support:

- File Server Resource Management (FSRM) features like Folder Quotas, File Screening, and File Classification.
- Distributed File Services Replication (DFS-R).
- NFS.
- Data deduplication.

Easier cluster migration

The Migrate A Cluster Wizard makes it easy to migrate services and applications from a cluster running Windows Server 2008, Windows Server 2008 R2, or Windows Server 2012. The wizard helps you migrate the configuration settings for clustered roles, but it doesn't migrate settings of the cluster, network, or storage, so you need to make sure that your new cluster is configured before you use the wizard to initiate the migration process. In addition, if you want to use new storage for the clustered roles you're migrating, you need to make sure that this storage is available to the destination cluster before running the wizard. Cluster migration also now supports Hyper-V and allows you to export and re-import VMs as part of the migration process.

Now support is also included for copying the configuration information of multiple VMs from one failover cluster to another, making it easier to migrating settings between clusters. And you can migrate configuration information for applications and services on clusters running Windows Server 2008, Windows Server 2008 R2, and Windows Server 2012.

Improved Cluster Validation

Cluster validation has been improved in Windows Server 2012 and is much faster than in the previous version of Failover Clustering. The Validate A Configuration Wizard, shown in Figure 3-1, simplifies the process of validating hardware and software for the servers that you want to use in your failover cluster. New validation tests have been added to this wizard for the Hyper-V role and VMs (when the Hyper-V role is installed) and for verification of CSV requirements. And more detailed control is now provided so that you can validate an explicitly targeted logical unit number (LUN).

Simplified cluster management

The Failover Clustering feature is now fully integrated with the new Server Manager of Windows Server 2012, making it easier to discover and manage the nodes of a cluster. For example, you can update a cluster by right-clicking the cluster name, which in Figure 3-2 has been added to the server group named Group 1.

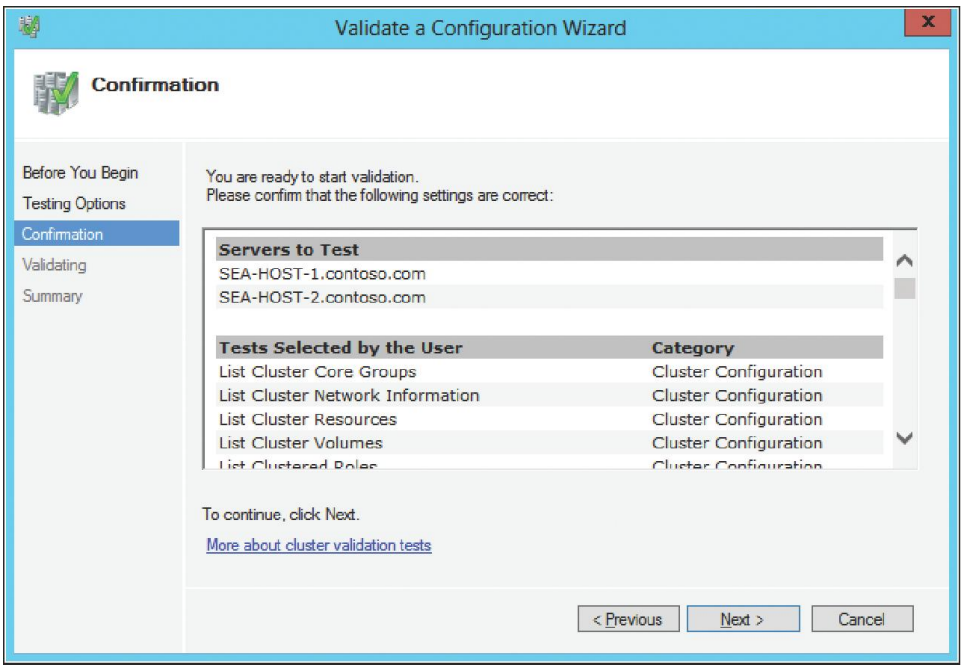

FIGURE 3-1 – Validating a failover cluster using the Validate A Configuration Wizard.

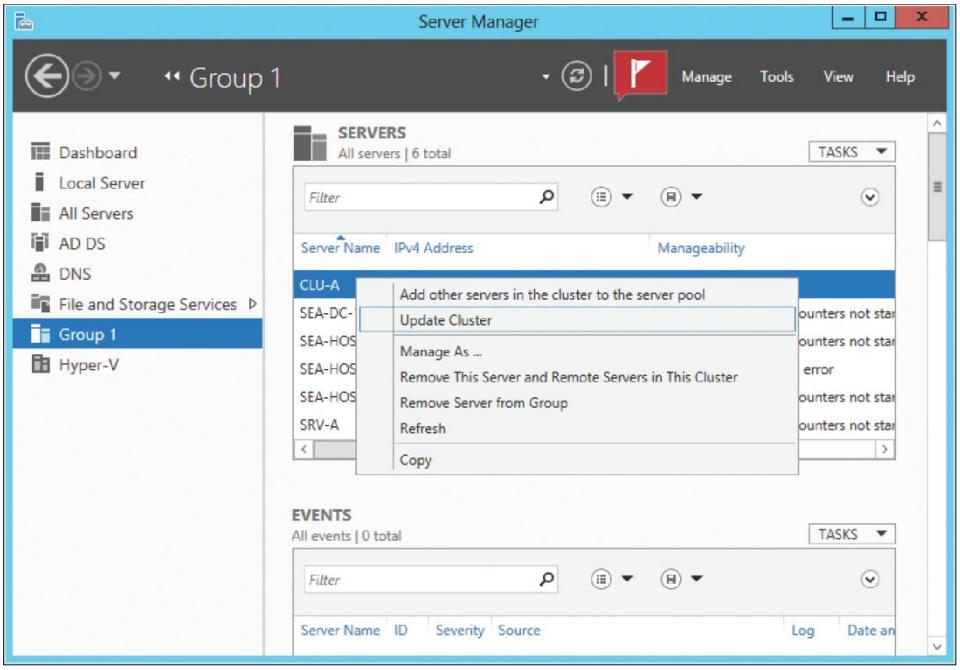

FIGURE 3-2 – You can perform cluster-related tasks from the new Server Manager

Server groups simplify the job of managing sets of machines such as the nodes in a cluster. A single-click action can add all the nodes in a cluster to a server group to facilitate remote multi-server management.

The capabilities of the new Server Manager of Windows Server 2012 are described in more detail later in this chapter.

Active Directory integration

Failover Clustering in Windows Server 2012 is more integrated with Active Directory than in previous versions. For example, support for delegated domain administration is now provided to enable intelligent placement of cluster computer objects in Active Directory. This means, for example, that you can now create cluster computer objects in targeted organizational units (OUs) by specifying the distinguished name (DN) of the target OU. And as a second example, you could create cluster computer objects by default in the same OUs as the cluster nodes.

Task Scheduler integration

Failover Clustering in Windows Server 2012 is also integrated into the Task Scheduler, which allows you to configure tasks you want to run on clusters in three ways:

- ClusterWide tasks are scheduled to run on all nodes in the cluster.
- AnyNode tasks are scheduled to run on a single, randomly selected cluster node.
- ResourceSpecific tasks are scheduled to run only on the cluster node that currently owns the specified resource.

You can configure clustered tasks by using Windows PowerShell. Table 3-1 lists the cmdlets available for this purpose. For more information on any of these cmdlets, use Get-Help <cmdlet>

TABLE 3-1 – Windows PowerShell Cmdlets for Configuring Clustered Tasks

Windows PowerShell Cmdlet	Description
Register-ClusteredScheduledTask	Creating a new clustered scheduled task
Unregister-ClusteredScheduledTask	Delete a clustered scheduled task
Set-ClusteredScheduledTask	Update existing clustered task
Get-ClusteredScheduledTask	Enumerating existing clustered tasks

VM priority

Efficient automatic management of clustered VMs and other clustered roles is now possible in Windows Server 2012 by assigning a relative priority to each VM in the cluster. Once this has been configured, the cluster will then automatically manage the VM or other clustered role based on its assigned priority.

Four possible priorities can be assigned to a clustered VM or clustered role:

- High
- Medium (the default)
- Low
- No Auto Start

Assigning priorities to clustered VMs or other clustered roles lets you control both the start order and placement of each VM or other role in the cluster. For example, VMs that have higher priority are started before those having lower priority. The benefit of this is to allow you to ensure that the most important VMs are started first and are running before other VMs are started. In addition, support for preemption is included so that low-priority VMs can be automatically shut down in order to free up resources so that higher-priority VMs can successfully start. And although Hyper-V in Windows Server 2012 now supports concurrent Live Migrations, the order in which VMs queued for Live Migration but not yet migrated can also be determined on the basis of priority.

VMs that have higher priority are also placed on appropriate nodes before VMs with lower priority. This means, for example, that VMs can be placed on the nodes that have the best available memory resources, with memory requirements being evaluated on a per-VM basis. The result is enhanced failover placement, and this capability is also Non-Uniform Memory Access (NUMA)-aware.

Figure 3-3 shows Failover Cluster Manager being used to manage a two-node cluster that has two cluster roles running on it: a scale-out file server and a VM. Right-clicking the clustered VM and selecting Change Startup Priority allows you to change the priority of the VM from its default Medium setting to High.

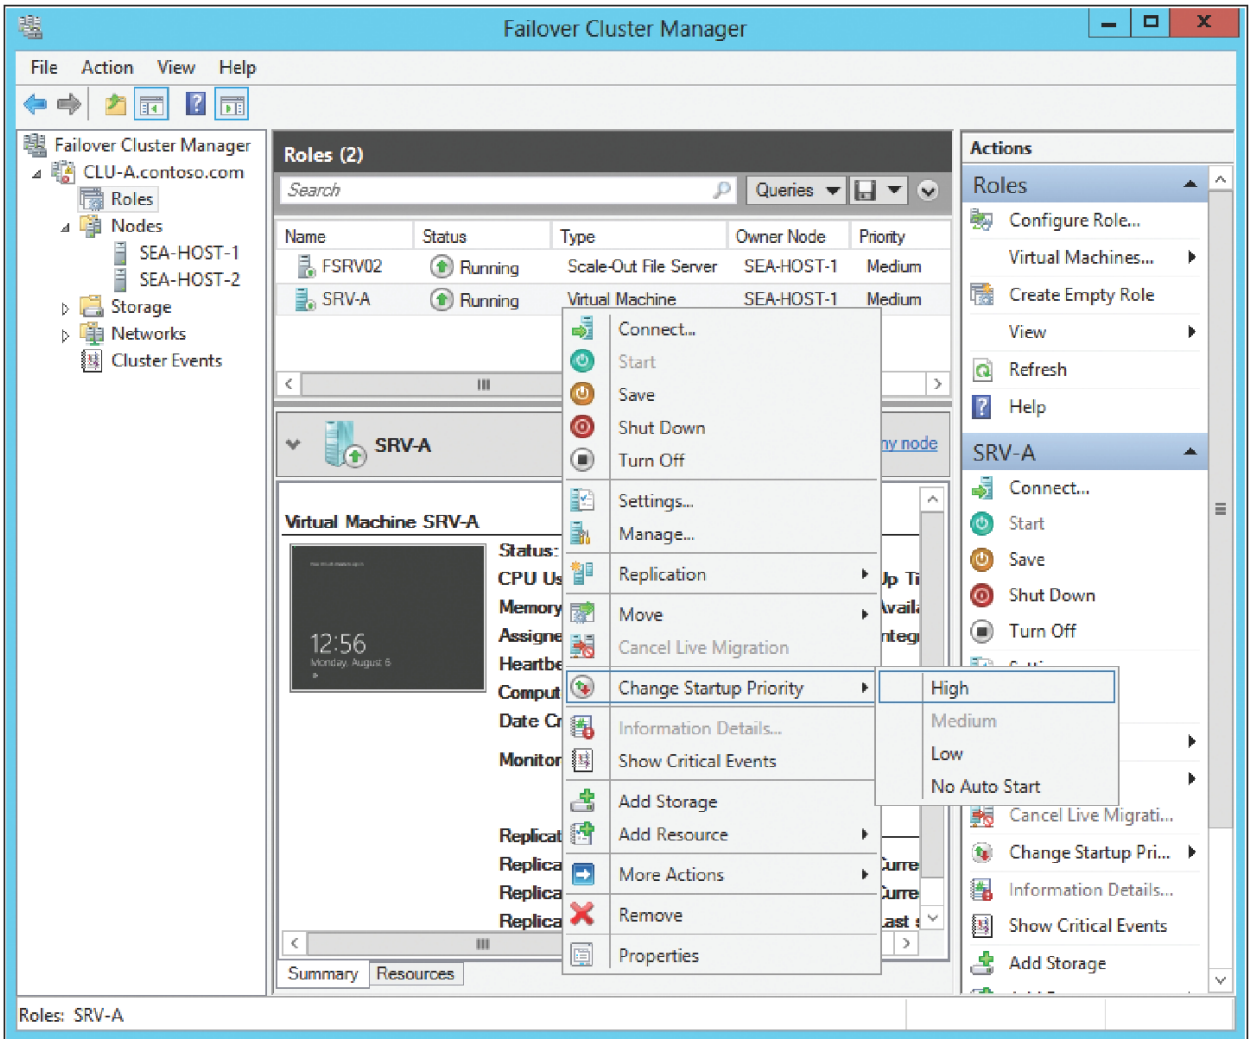

FIGURE 3-3 – Using Failover Cluster Manager to configure the priority of a clustered VM

Failover Clustering placement policies for Hyper-V

Windows Server Failover Clustering provides a critical piece of Hyper-V infrastructure not just for high availability, but also for mobility. A key concept of a virtualized or private cloud environment is to abstract workloads from their underlying physical resources, and Failover Clustering enables this by allowing the movement and placement of VMs between different physical hosts using live migration with no perceived downtime. There are a few placement best practices that can allow you to optimize the cluster for different Hyper-V scenarios.

Default failover policy. When there is a failure of a node, VMs are distributed across the remaining cluster nodes. In previous versions of Windows Server, any resource would be distributed to the nodes hosting the fewest number of VMs. In Windows Server 2012, enhancements in this logic have been made to redistribute

the VMs based on the most commonly constrained resource, host memory. Each VM is placed on the node with the freest memory resources, and the memory requirements are evaluated on a per-VM basis, including checks to see if the VM is NUMA-aware.

If a cluster node hosting several VMs crashes, the Cluster Service will find the highest-priority VM, then look across the remaining nodes to determine which node currently has the freest memory. The VM is then started on that node. This process repeats for all the VMs, from the highest priority to the lowest priority, until all VMs are placed.

VM Priority. In Windows Server 2012, each VM running on a cluster can be assigned a priority: High, Medium, or Low. This can be used to ensure that the high-priority VMs are given preferential treatment for cluster operations. This could be used to ensure that the organization's most critical services or key infrastructure roles can come online before less important workloads.

If a cluster node hosting several VMs crashes, the high-priority VMs will start first, then the medium-priority VMs, then finally the low-priority ones. This same logic will be applied for other cluster operations, such as multiple live migrations or Node Maintenance Mode, where the high-priority VMs will always be moved first.

Preferred Owners. From earlier versions of Windows Server, it has been possible to configure the preference for node failover order for each VM. This can be helpful in an environment where it is important for certain VMs to stay on certain nodes, such as if there is a primary datacenter where the VMs should usually run (the Preferred Owners), and a backup datacenter available for a disaster recovery for the VMs if the primary site is unavailable.

If a cluster node hosting several VMs crashes, a high-priority VMs will attempt to move to the first node in the list of Preferred Owners. If that node is not available, then the VM will attempt to move to the second node in the Preferred Owners list. If none of those Preferred Owners are available, then it will move to the first node that is on the Possible Owners list.

Possible Owners. The Possible Owners setting for each VM also existed in earlier versions of Windows Server. It enables VMs to move to and start on a cluster node when none of the Preferred Owners are available. This can be used in an environment when VMs should still run on a host, even when none of the Preferred Owners are available. In a multisite cluster, the nodes at the backup site would be assigned as a Possible Owner, but not as a Preferred Owner. In this

scenario, the VMs would fail over to the secondary site only when none of the nodes at the primary site (Preferred Owners) are available.

If a cluster node hosting several VMs crashes, a high-priority VMs will attempt to move to the first node in the list of Preferred Owners. If none of those Preferred Owners are available, then it will move to the first node that is on the Possible Owners list. If the first node in the Possible Owners list is not available, then it will move to the next node on the list. If none of the nodes in either the Preferred Owners nor Possible Owners lists are available, then the VM will move to any other node, but remain offline. Depending on Failback policies, the VM can move back to a Preferred Owner or Possible Owner and start as soon as one of those nodes becomes available.

Failback. Another setting for each VM that continues to be important in Windows Server 2012 is the option to move the VM back to Preferred Owners or Possible Owners, starting from the most Preferred Owner. This feature is helpful if you wish to keep certain VMs on the same hosts, and return those VMs to the host once it recovers from a crash.

If a cluster node recovers from a crash and rejoins cluster membership, any VMs that are not running on a Preferred Owner will be notified that this node is now available for placement. Starting with the high-priority VMs that are running on a Possible Owner (or are offline on another node), each VM will determine if this node is a better host, then live-migrate (or start) the VM on that Preferred Owner.

Persistent Mode. One problem that is often seen in highly virtualized environments is a "boot storm," which happens when simultaneously starting a large number of VMs. Starting a VM requires more host resources than standard running operations, so starting a lot of VMs can sometimes overload the host, affecting its performance, or even causing it to crash (if certain host reserves are not set). As a safety precaution, during failover or when a node is restarted, the number of VMs that will start simultaneously is limited (High priority first), and the rest will be queued up to start on that node. Even when these VMs are simultaneously starting, they are slightly staggered to help spread out the demands on the host. There are still some settings that can be configured to avoid these "boot storms."

Persistent Mode was introduced in Windows Server 2008 R2 and provides the ability to keep a VM on the last host it was deliberately placed on (either by an administrator or a System Center Virtual Machine Manager placement policy). If an entire cluster crashes, each VM will wait for the node is was previously

hosted on to come online before starting up, still honoring high-priority VMs first. This prevents all of the VMs across the cluster from trying to start up on the first node(s) that come online, helping to avoid a "boot storm." There is a default amount of time the cluster service will wait for the original node to rejoin the cluster. If the node does not join within this period, the VM will be placed on the most Preferred Owner, ensuring that the VM will still come online, while having given that new host an opportunity to start its own VMs.

Auto-Start. There may be cases when there are unimportant VMs that should not be started after a cluster failover or a crash, giving the other VMs an opportunity to fail over and come online quickly. The Auto-Start property has also existed in previous versions of Windows Server, and if it is disabled, the VM will not be automatically started when it is placed on a node.

This can be useful in highly virtualized environments when it is important to keep hosts and critical infrastructure VMs running, while not worrying about constraining resources or "boot storms" caused by VMs that do not need to be continually available, yet are still hosted on the cluster. These VMs can be started later by the administrator or automatically using a script.

Anti-Affinity. The final placement policy has also existed before Windows Server 2012, but looks at other VMs, rather than the hosts. The cluster property, AntiAffinityClassName (AACN), enables custom tagging of a VM so that different VMs may share or have different AACNs. VMs that share the same AACN will distribute themselves across different hosts automatically. This can be useful to separate tenets or VMs with the same infrastructure roles across different nodes in the cluster. For example, having all the virtualized DNS servers or guest cluster nodes on the same host would be a single point of failure if that node crashes, so spreading these VMs out across different hosts helps maintain continual service availability.

If there is a cluster with four nodes and four VMs that have the AntiAffinityClassName of "blue," then by default, each node would host one of the "blue" VMs. If there are more "blue" VMs with the same AACN than there are nodes in the cluster, then there will be more than 1 "blue" VM on each node, but they will still distribute themselves as evenly as possible.

Virtual machine monitoring

Ensuring high availability of services running in clustered VMs is important because service interruptions can lead to loss of user productivity and customer

dissatisfaction. A new capability of Failover Cluster Manager in Windows Server 2012 is the ability to monitor the health of clustered VMs by determining whether business-critical services are running within VMs running in clustered environments. By enabling the host to recover from service failures in the guest, the cluster service in the host can take remedial action when necessary in order to ensure greater uptime for services your users or customers need.

You enable this functionality by right-clicking the clustered VM and selecting Configure Monitoring from the More Actions menu item, as shown here:

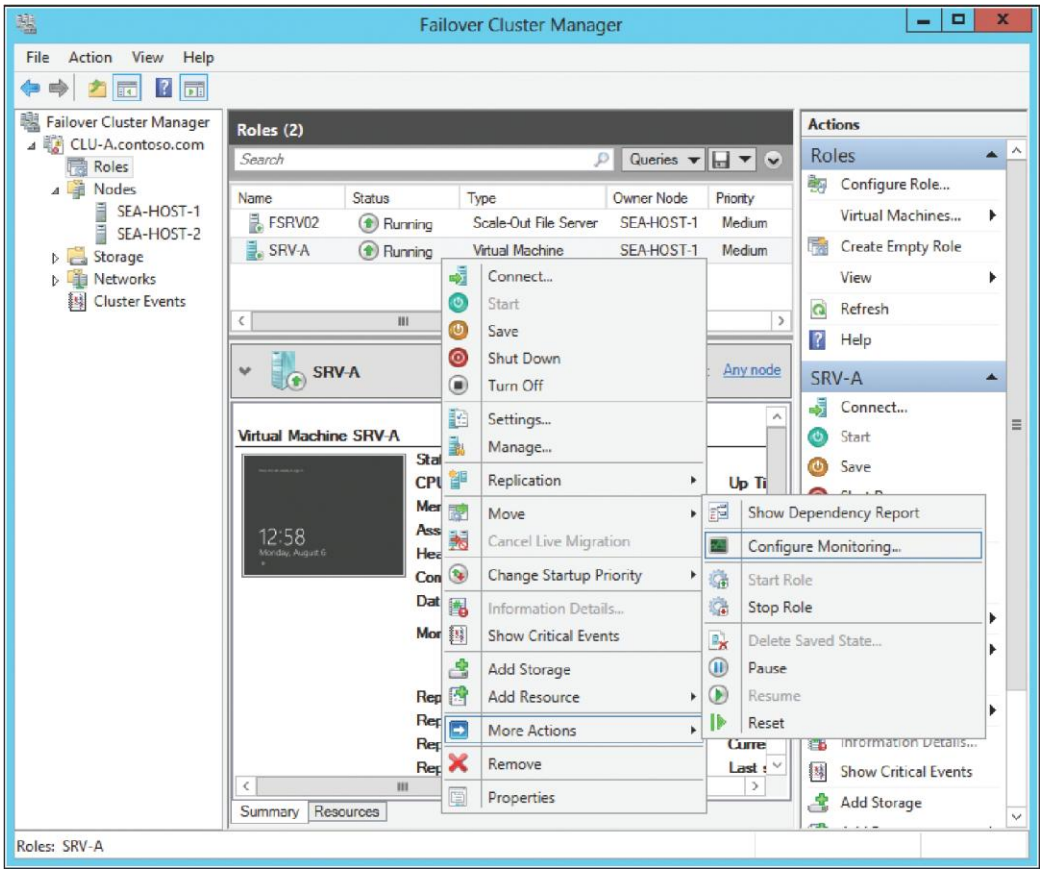

You then select the service or services you want to monitor on the VM, and if the selected service fails, the VM can either be restarted or moved to a different cluster node, depending on how the service restart settings and cluster failover settings have been configured:

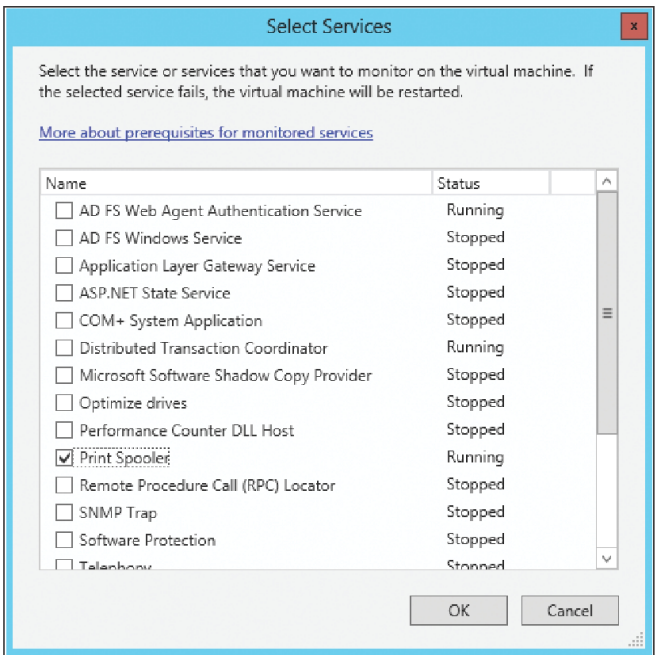

You can also use Windows PowerShell to configure VM monitoring. For example to configure VM monitoring. For example, to configure monitoring of the Print Spooler service on the VM named SRV-A, you could use this command:

```
Add-ClusterVMMonitoredItem -vm SRV-A -service spooler
```

For VM monitoring to work, the guest and host must belong to the same domain or to domains that have a trust relationship. In addition, you need to enable the Virtual Machine Monitoring exception in Windows Firewall on the guest:

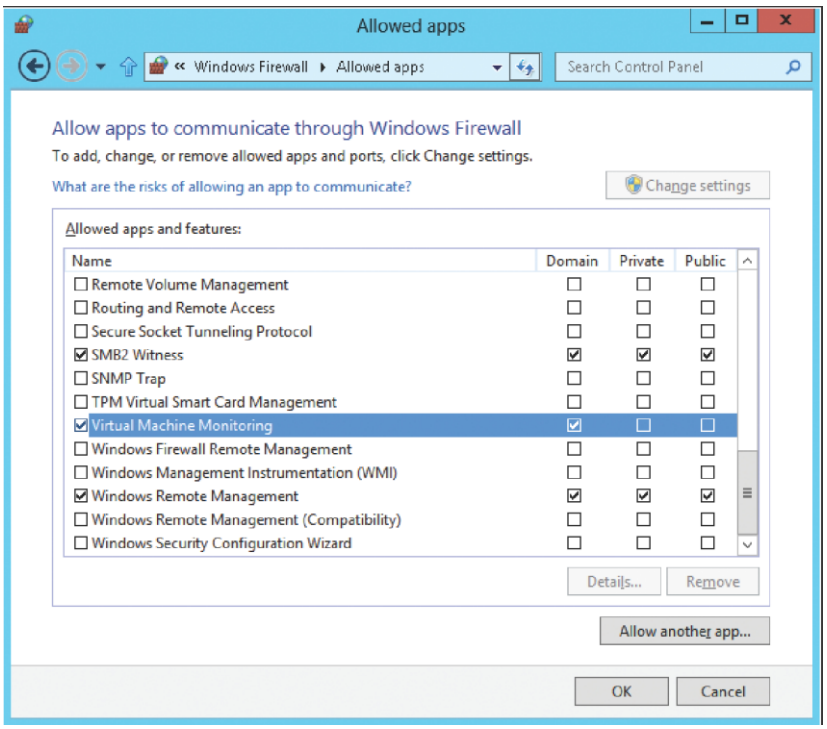

If Windows PowerShell Remoting is enabled in the guest, then you don't need to enable the Virtual Machine Monitoring exception in Windows Firewall when you configure VM monitoring using Windows PowerShell. You can enable Windows PowerShell Remoting by connecting to the guest, opening the Windows PowerShell console, and running this command:

```
Enable-PSRemoting
```

Then, to configure monitoring of the Print Spooler service on the guest, you would open the Windows PowerShell console on the host and run these commands:

```
Enter-PSSession
Add-ClusterVMMonitoredItem -service spooler
Exit-PSSession
```

VM monitoring can monitor the health of any NT Service such as the Print Spooler, IIS, or even a server application like SQL Server. VM monitoring also requires the use of Windows Server 2012 for both the host and guest operating systems.

Node vote weights

The quorum for a failover cluster is the number of elements that need to be online in order for the cluster to be running. Each element has a "vote," and the votes of all elements determine whether the cluster should run or cease operations. In the previous version of Failover Clustering in Windows Server 2008 R2, the quorum could include nodes, but each node was treated equally and assigned one vote. In Windows Server 2012, however, the quorum settings can be configured so that some nodes in the cluster have votes (their vote has a weight of 1, which is the default), whereas others do not have votes (their vote has a weight of 0).

Node vote weights provide flexibility that is particularly useful in multisite clustering scenarios. By appropriately assigning a weight of 1 or 0 as the vote for each node, you can ensure that the primary site has the majority of votes at all times.

Dynamic quorum

Another new feature of Failover Clustering in Windows Server 2012 is the ability to change the quorum dynamically based on the number of nodes currently in active membership in the cluster. This means that as nodes in a cluster are shut down, the number of votes needed to reach quorum changes instead of remaining the same, as in previous versions of Failover Clustering.

Dynamic quorum allows a failover cluster to remain running even when more than half of the nodes in the cluster fail. The feature works with the following quorum models:

- Node Majority.
- Node and Disk Majority.
- Node and File Share Majority.

It does not work, however, with the Disk Only quorum model.

Node drain

When a failover cluster node needs to be taken down for maintenance, the clustered roles hosted on that node first need to be moved to another node in the cluster. Some examples of the kind of maintenance you might need to perform on a cluster node might be upgrading the hardware on the node or applying a service pack.

In the previous version of Failover Clustering in Windows Server 2008 R2, taking down a node for maintenance was a manual process that required placing the node into a Paused state and then manually moving the applications and services running on the node to another node on the cluster.

However, Failover Clustering in Windows Server 2012 now makes performing maintenance on cluster nodes much easier. A new feature called node drain now lets you automate the moving of clustered roles off from the node scheduled for maintenance onto other nodes running on the cluster.

Draining a node can be done either manually by a single click in the Failover Cluster Manager console (as shown in Figure 3-4), or you can script it with Windows PowerShell for automation purposes by using the Suspend-ClusterNode cmdlet.

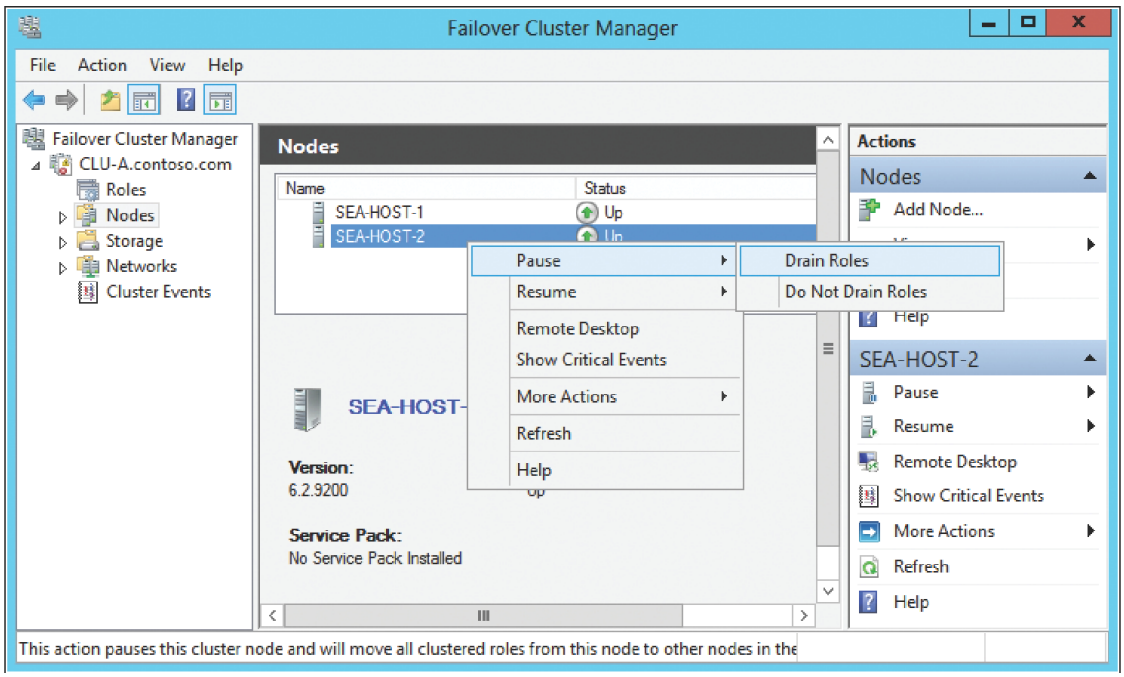

FIGURE 3-4 – Initiating a node drain to take down a node for maintenance

Initiating the node drain process does the following:

1) Puts the node into the Paused state to prevent roles hosted on other nodes from being moved to this node
2) Sorts the roles on the node according to the priority you've assigned them (assigning priorities to roles is another new feature of Failover Clustering in Windows Server 2012)
3) Moves the roles from the node to other nodes in the cluster in order of priority (VMs are live-migrated to other hosts)

Once the process is completed, the node is down and is ready for maintenance.

Cluster-Aware Updating

Cluster-Aware Updating (CAU) is a new feature of Windows Server 2012 that lets you automatically apply software updates to the host operating system in clustered servers with little or no downtime. CAU thus both simplifies update management of cluster nodes and helps ensure your cluster remains available at all times.

CAU functionality works seamlessly with your Windows Server Update Services (WSUS) infrastructure and is installed automatically on each cluster node. CAU can be managed from any server that has the Failover Cluster feature installed but does not belong to the cluster whose nodes you wish to update.

As shown previously in Figure 3-2, you can use Server Manager to initiate the process of updating a cluster. Selecting the Update Cluster menu item opens the Cluster-Aware Updating dialog box and connects to the cluster you selected in Server Manager:

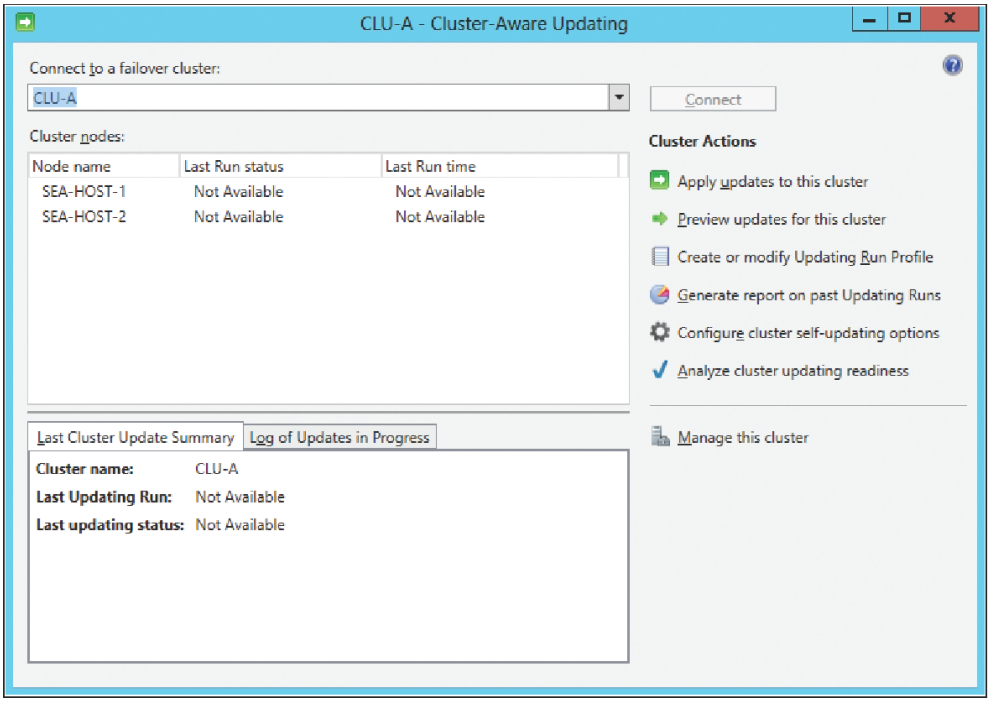

You can also open the Cluster-Aware Updating dialog box from Failover Cluster Manager.

Clicking the Preview Updates For This Cluster option opens the Preview Updates dialog box, and clicking Generate Update Preview List in this dialog box downloads a list of the updates available for nodes in the cluster:

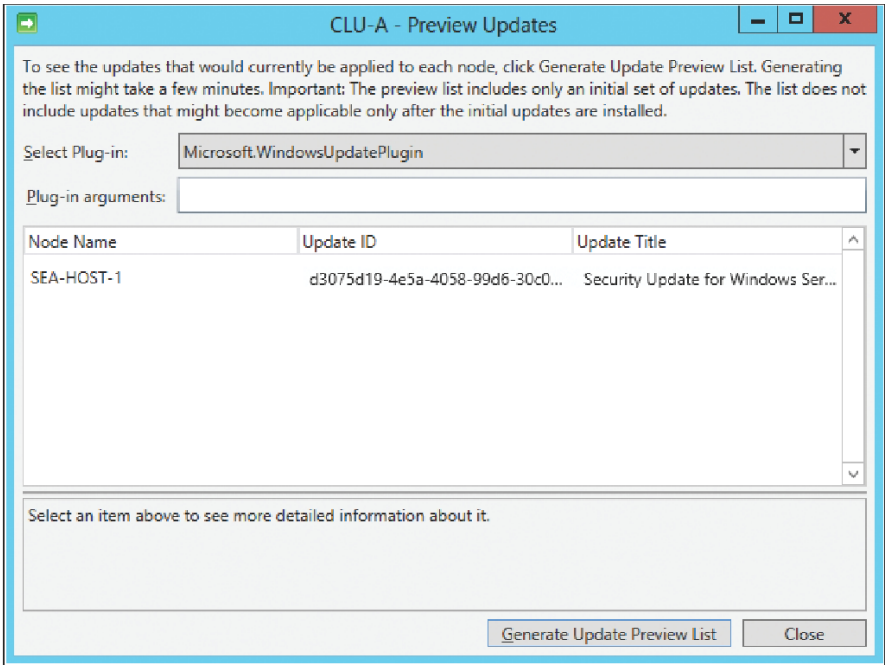

Closing the Preview Updates dialog box returns you to the Cluster-Aware Updating dialog box where clicking the Apply Updates To This Cluster option starts the Cluster-Aware Updating Wizard:

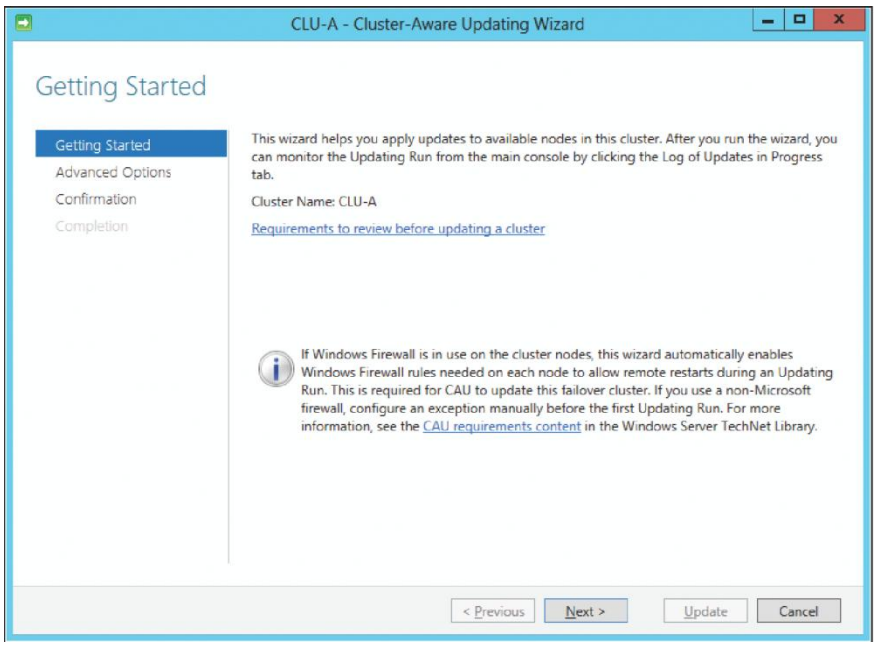

Once you've walked through the steps of this wizard and clicked Next, the update process begins. Cluster nodes are then canned to determine which updates they require in the following way:

1) Nodes are prioritized according to the number of workloads they have running on them.
2) The node with the fewest workloads is then drained to place it into maintenance mode. This causes the workloads running on the node to be moved automatically to other active nodes in the cluster.
3) The Windows Update Agent on this node downloads the necessary updates from either Windows Update or from your WSUS server if you have one deployed in your environment.
4) Once the node has been successfully updated, the node is resumed and becomes an active node in the cluster again.
5) The process is then repeated on each remaining node in the cluster in turn, according to priority.

CAU employs an updating run profile to store the settings for how exceptions are handled, time boundaries for the update process, and other aspects of the node updating process. You can configure these settings by clicking the Create Or Modify Updating Run Profile option in the Cluster-Aware Updating dialog box shown previously. Doing this opens the Updating Run Profile Editor, as shown here:

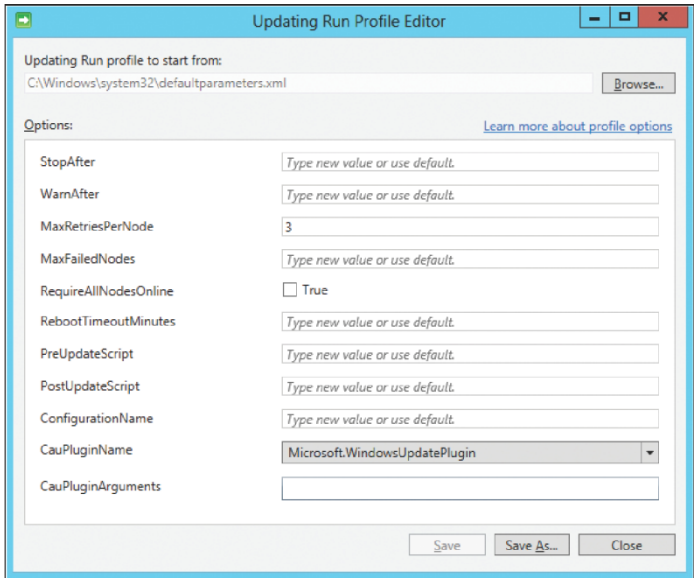

Guest clustering

Failover Clustering of Hyper-V can be implemented in two ways:

1) **Host clustering**, in which the Failover Clustering feature runs in the parent partition of the Hyper-V host machines. In this scenario, the VMs running on the hosts are managed as cluster resources and they can be moved from one host to another to ensure availability of the applications and services provided by the VMs.
2) **Guest clustering**, in which the Failover Clustering feature runs in the guest operating system within VMs. Guest clustering provides high availability for applications and services hosted within VMs, and it can be implemented either on a single physical server (Hyper-V host machine) or across multiple physical servers.

Host clustering helps ensure continued availability in the case of hardware failure or when you need to apply software updates to the parent partition. Guest clustering, by contrast, helps maintain availability when a VM needs to be taken down for maintenance. Implementing guest clustering on top of host clustering can provide the best of both worlds.

Guest clustering requires that the guest operating systems running in VMs have direct access to common shared storage. In previous versions of Windows Server, the only way to provision such shared storage in a guest clustering scenario was to have iSCSI initiators running in the guest operating systems so they could connect directly with iSCSI-based storage. Guest clustering in previous versions of Windows Server did not support using Fibre Channel SANs for shared storage. VMs running Windows Server 2008 R2 in a guest clustering scenario can use Microsoft iSCSI Software Target 3.3, which can be downloaded from the Microsoft Download Center. Figure 3-5 illustrates the typical way guest clustering was implemented in Windows Server 2008 R2.

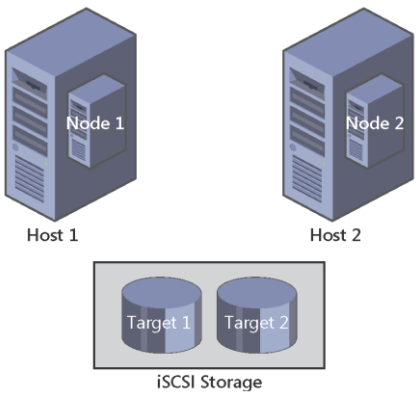

FIGURE 3-5 – Implementing guest clustering with Failover Clustering in Windows Server 2008 R2 using iSCSI Software Target

In Windows Server 2012, iSCSI Software Target is now an in-box feature integrated into Failover Clustering, making it easier to implement guest clustering using shared iSCSI storage. And by starting the High Availability Wizard from the Failover Clustering Manager console, you can add the iSCSI Target Server as a role to your cluster quickly. You can also do this with Windows PowerShell by using the Add-ClusteriSCSITargetServerRole cmdlet.

But iSCSI is now no longer your only option as far as shared storage for guest clustering goes. That's because Windows Server 2012 now includes an in-box Hyper-V Virtual Fibre Channel adapter that allows you to connect directly from within the guest operating system of a VM to LUNs on your Fibre Channel SAN (see Figure 3-6). The new virtual Fibre Channel adapter supports up to four virtual HBAs assigned to each guest with separate worldwide names (WWNs) assigned to each virtual HBA and N_Port ID Virtualization (NPIV) used to register guest ports on the host.

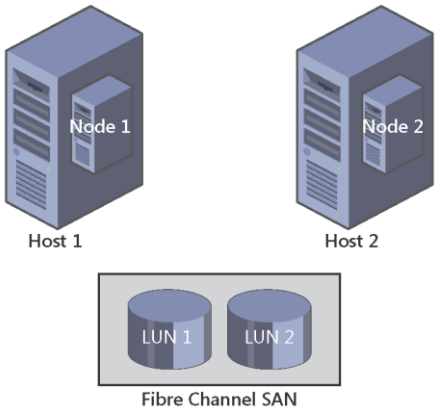

FIGURE 3-6 – Failover Clustering in Windows Server 2012 now allows VMs to connect directly to a Fibre Channel SAN

Configuring Fibre Channel from the guest

Before you configure Fibre Channel as the shared storage for VMs in a guest cluster, make sure that you have HBAs installed in your host machines and connected to your SAN. Then, open the Virtual SAN Manager from the Hyper-V Manager console and click Create to add a new virtual Fibre Channel SAN to each host:

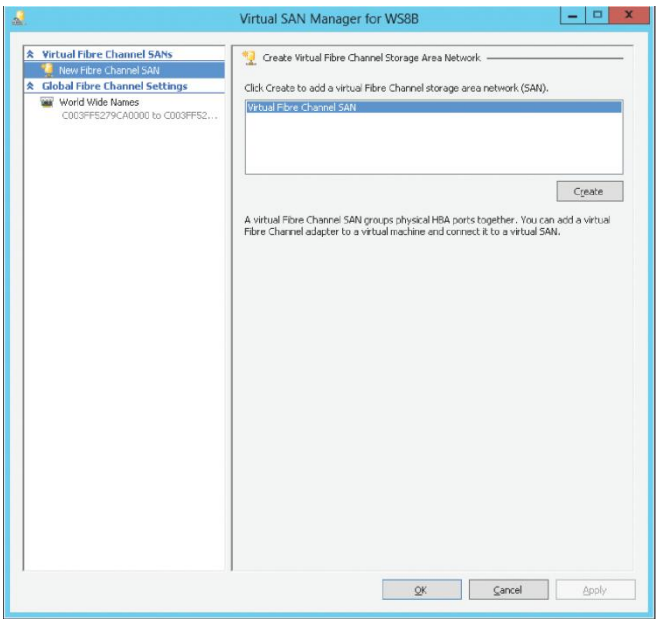

Provide a name for your new virtual Fibre Channel and configure it as needed. Then open the settings for each VM in your guest cluster and select the Add Hardware option to add the virtual Fibre Channel adapter to the guest operating system of the VM:

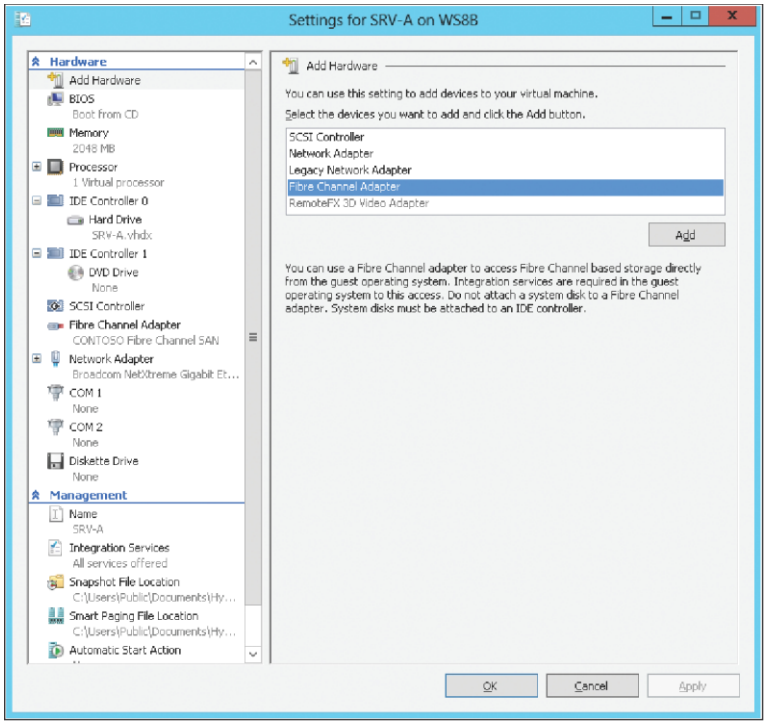

Then simply select the virtual SAN you created earlier, and once you're done, each VM in your guest cluster can use your SAN for shared storage:

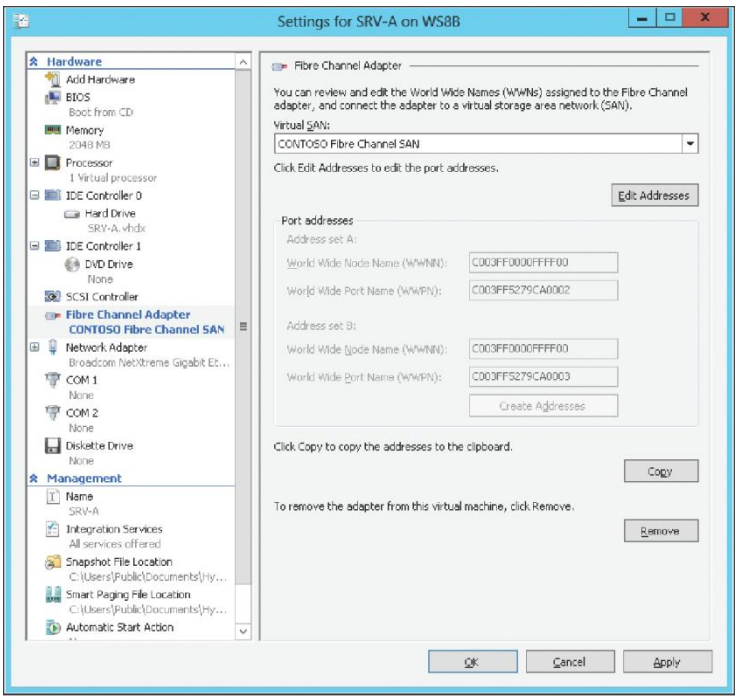

Guest clustering in Windows Server 2012 also supports other new Failover Cluster features, such as CAU, node drain, Storage Live Migration, and much more.

Guest clustering vs. VM monitoring

Guest clustering in Windows Server 2012 is intended for server applications that you currently have clustered on physical servers. For example, if you currently have Exchange Server or SQL Server deployed on host clusters, you will have the additional option of deploying them on guest clusters (which can themselves be deployed on host clusters) for enhanced availability when you migrate your infrastructure to Windows Server 2012.

VM monitoring by contrast can enhance availability for other server roles in your environment, such as your print servers. You can also combine VM monitoring with guest clustering for even greater availability.

Enhanced Windows PowerShell support

Failover Clustering in Windows Server 2012 also includes enhanced Windows PowerShell support with the introduction of a number of new cmdlets for managing cluster registry checkpoints, creating scale-out file servers,

monitoring health of services running in VMs, and other capabilities. Table 3-2 lists some of the new Windows PowerShell cmdlets for Failover Clustering.

Table 3-2 – New Windows PowerShell Cmdlets for Failover Clustering

Windows PowerShell cmdlet	Purpose
Add-ClusterCheckpoint	Manages cluster registry checkpoints, including cryptographic checkpoints
Get-ClusterCheckpoint	
Remove-ClusterCheckpoint	
Add-ClusterScaleOutFileServerRole	Creates a file server for scale-out application data
Add-ClusterVMMonitoredItem	Monitors the health of services running inside a VM
Get-ClusterVMMonitoredItem	
Remove-ClusterVMMonitoredItem	
Reset-ClusterVMMonitoredState	
Update-ClusterNetworkNameResource	Updates the private properties of a Network Name resource and sends DNS updates
Test-ClusterResourceFailure	Replaces the Fail-ClusterResource cmdlet

Storage migration

Storage migration is a new feature of Hyper-V in Windows Server 2012 that lets you move all of the files for a VM to a different location while the VM continues running. This means that with Hyper-V hosts running Windows Server 2012, it's no longer necessary to take a VM offline when you need to upgrade or replace the underlying physical storage.

When you initiate a storage migration for a VM, the following takes place:

1) A new VHD or VHDX file is created in the specified destination location (storage Migration works with both VHD and VHDX).
2) The VM continues to both read and write to the source VHD, but new write operations are now mirrored to the destination disk.
3) All data is copied from the source disk to the destination disk in a single-pass copy operation. Writes continue to be mirrored to both disks during this copy operation, and uncopied blocks on the source disk that have been updated through a mirrored write are not recopied.

4) When the copy operation is finished, the VM switches to using the destination disk.
5) Once the VM is successfully using the destination disk, the source disk is deleted and the storage migration is finished. If any errors occur, the VM can fail back to using the source disk.

Storage migration of unclustered VMs can be initiated from the Hyper-V Manager console by selecting the VM and clicking the Move option. Storage migration of clustered VMs cannot be initiated from the Hyper-V Manager console; the Failover Clustering Manager console must be used instead. You can also perform storage migrations with Windows PowerShell by using the Move-VMStorage cmdlet.

Windows NIC Teaming

Windows NIC Teaming is the name for the new network adapter teaming functionality included in Windows Server 2012. Network adapter teaming is also known as **load balancing and failover** (LBFO) and enables multiple network adapters on a server to be grouped together into a team. This has two purposes:

- to help ensure availability by providing traffic failover in the event of a network component failure;
- to enable aggregation of network bandwidth across multiple network adapters.

Previously, implementing network adapter teaming required using third-party solutions from independent hardware vendors (IHVs). Beginning with Windows Server 2012, however, network adapter teaming is now an in-box solution that works across different NIC hardware types and manufacturers.

Windows NIC Teaming supports up to 32 network adapters in a team in three modes:

- **Static Teaming.** Also called **Generic Teaming** and based on IEEE 802.3ad draft v1, this mode is typically supported by server-class Ethernet switches and requires manual configuration of the switch and the server to identify which links form the team.
- **Switch Independent.** This mode doesn't require that the team members connect to different switches; it merely make it possible.
- **LACP.** Also called **dynamic teaming** and based on IEEE 802.1ax, this mode is Supported by most enterprise-class switches and allows automatic creation of a team using the Link Aggregation Control Protocol (LACP), which dynamically

identifies links between the server and a specific switch. To use this mode, you generally need to enable LACP manually on the port of the switch.

Configuring NIC teaming

NIC teaming can be enabled from Server Manager or using Windows PowerShell. For example, to use Server Manager to enable NIC teaming, you can begin by right-clicking the server you want to configure and selecting Configure NIC Teaming:

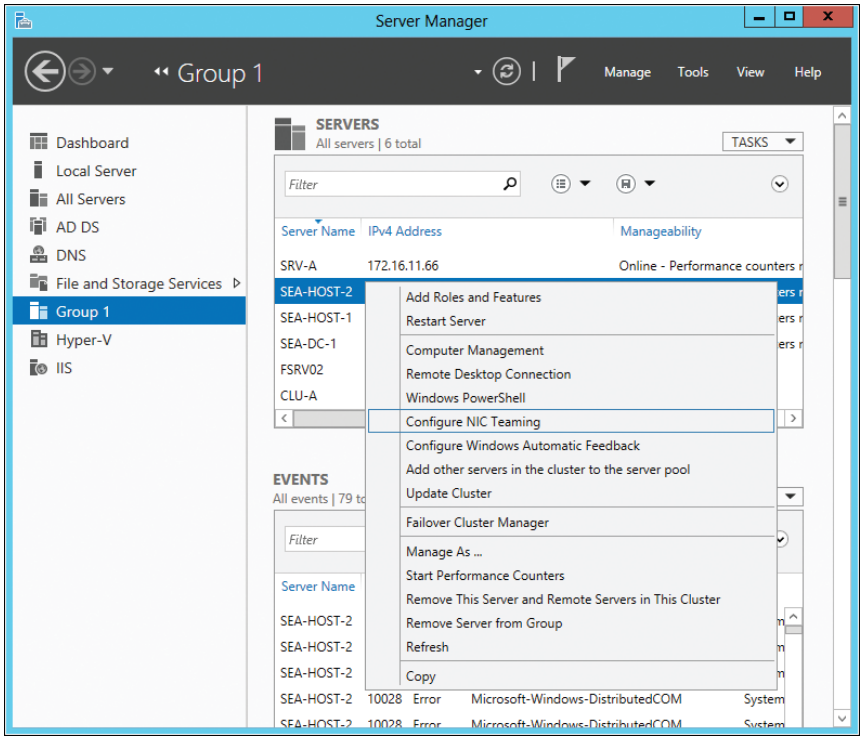

In the NIC Teaming dialog box that opens, select the network adapters you want to team. Then right-click and select Add To New Team:

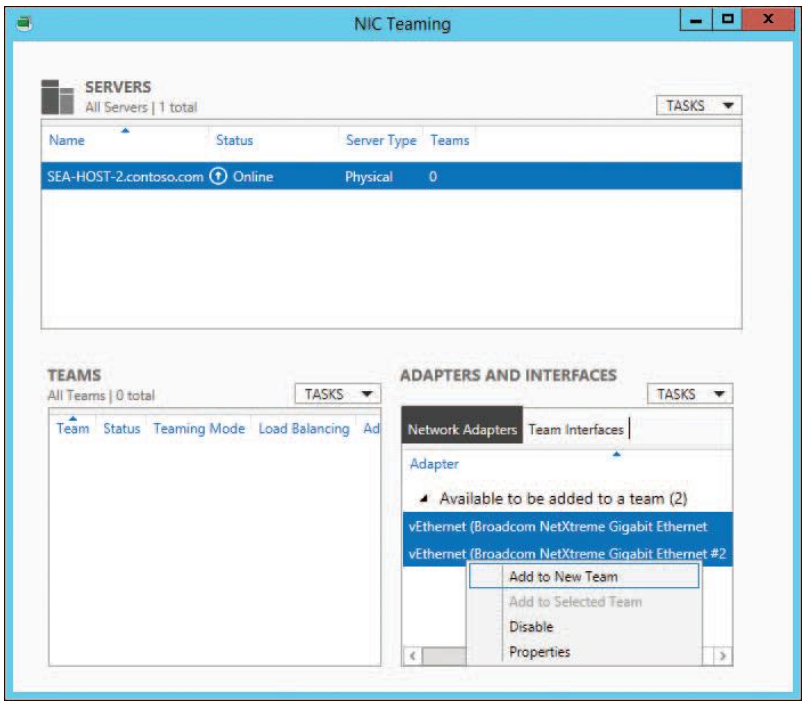

In the New Team dialog box, configure the teaming mode and other settings as desired:

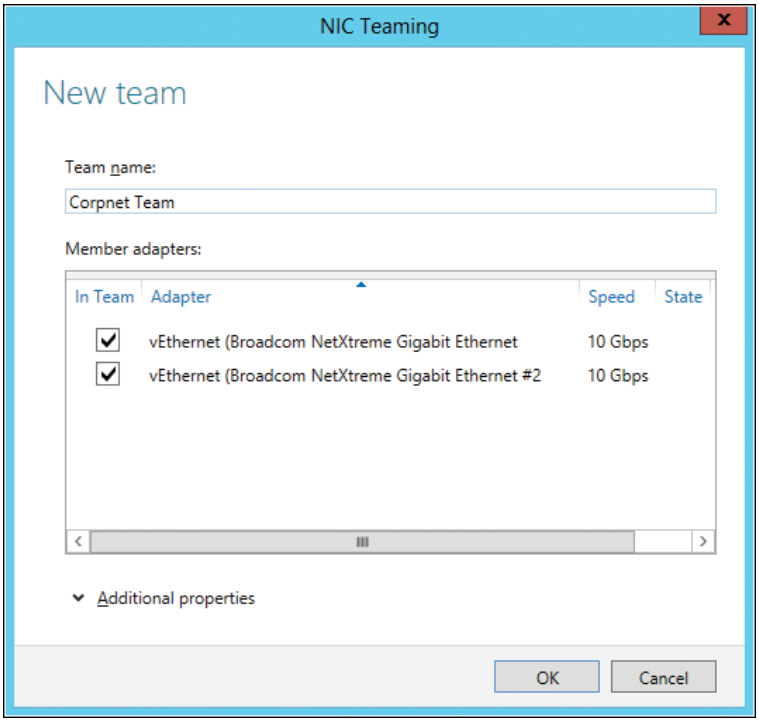

Clicking OK completes the process and, if successful, the new team will be displayed in the Teams tile of the NIC Teaming dialog box:

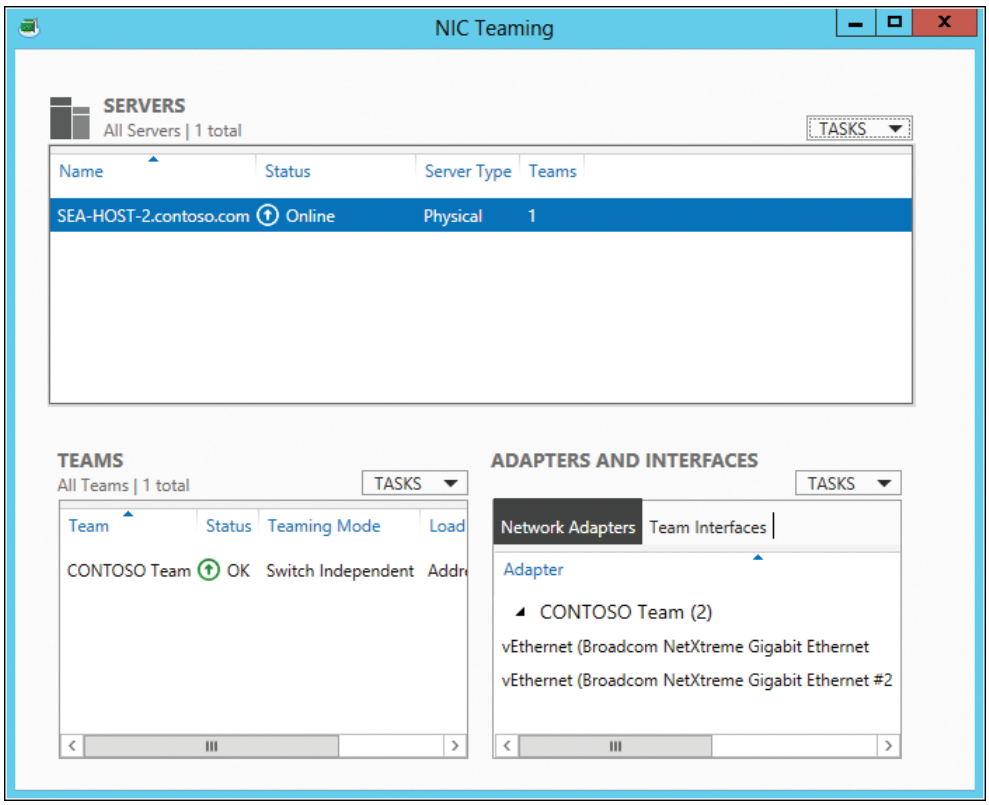

To configure and manage NIC teaming using Windows PowerShell, use cmdlets such as New-NetLbfoTeam to add a new team or Get-NetLbfoTeam to display the properties of a team. The cmdlets for managing NIC teaming are defined in the Windows PowerShell module named NetLbfo, and as Figure 3-7 shows, you can use the Get-Command cmdlet to display all the cmdlets defined in this module.

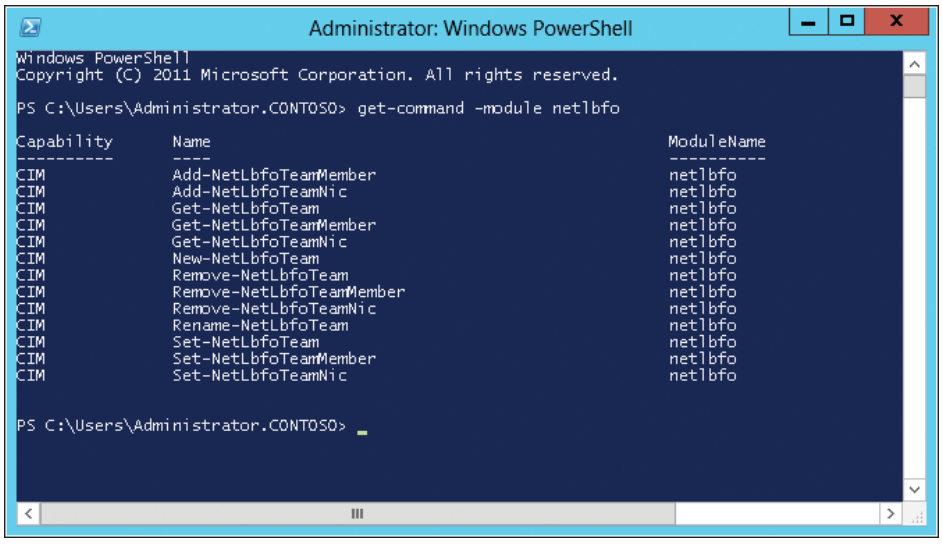

FIGURE 3-7 – Obtaining a list of cmdlets for configuring and managing NIC teaming.

Management efficiency

Provisioning and managing servers efficiently is an essential ingredient for cloud computing. Whether you are a mid-sized organization implementing a dedicated private cloud, a large enterprise deploying a shared private cloud, or a hoster managing a multitenant public cloud, Windows Server 2012 provides both the platform and the tools for managing your environment.

The new Server Manager of Windows Server 2012 can simplify the job of managing multiple remote servers across your organization. Enhancements to Active Directory can make your Active Directory environment much easier to deploy and manage than with previous versions of Windows Server. Domain controllers can now be safely cloned in order to save time when you need to deploy additional capacity, and restoring domain controller snapshots no longer disrupts your Active Directory environment. Foundational to successful cloud computing is automation, and version 3.0 of Windows PowerShell in Windows Server 2012 includes numerous enhancements that extend its capabilities and improve its usefulness in server administration.

The new Server Manager

Server Manager has been redesigned in Windows Server 2012 to facilitate managing multiple remote servers from a single administration console. Server Manager uses the remote management capabilities of Windows Management Instrumentation (WMI), Windows PowerShell, and the Distributed Component Object Model (DCOM) for connecting to remote servers to manage them. By default, servers running Windows Server 2012 are enabled for remote management, making it easy to provision and configure remote servers using Server Manager or Windows PowerShell. For example, in previous versions of Windows Server, you needed either physical access to a server or a Remote Desktop connection to the server if you wanted to add or remove a role or feature on the server. With Windows Server 2012, however, you can provision roles and features quickly and easily on remote servers from a central location by using Server Manager.

Server Manager is also included in the Remote Server Administration Tools (RSAT) for Windows 8, which enables administrators to manage their organization’s server infrastructure from a client workstation running Windows 8.

Server Manager can also be used to manage servers running Windows Server 2008 R2, Windows Server 2008, or Windows Server 2003, provided that remote management has been suitably configured on these systems.

Using Server Manager

The dashboard section of Server Manager shows you the state of your servers at a glance. The dashboard uses a 10-minute polling cycle so it's not a live monitoring solution like the System Center Operations Manager, but it does give you a general picture of what's happening with each server role in your environment. For example, in the following screenshot, the tile for the DNS role indicates an alert in the Best Practices Analyzer results for the DNS Server role:

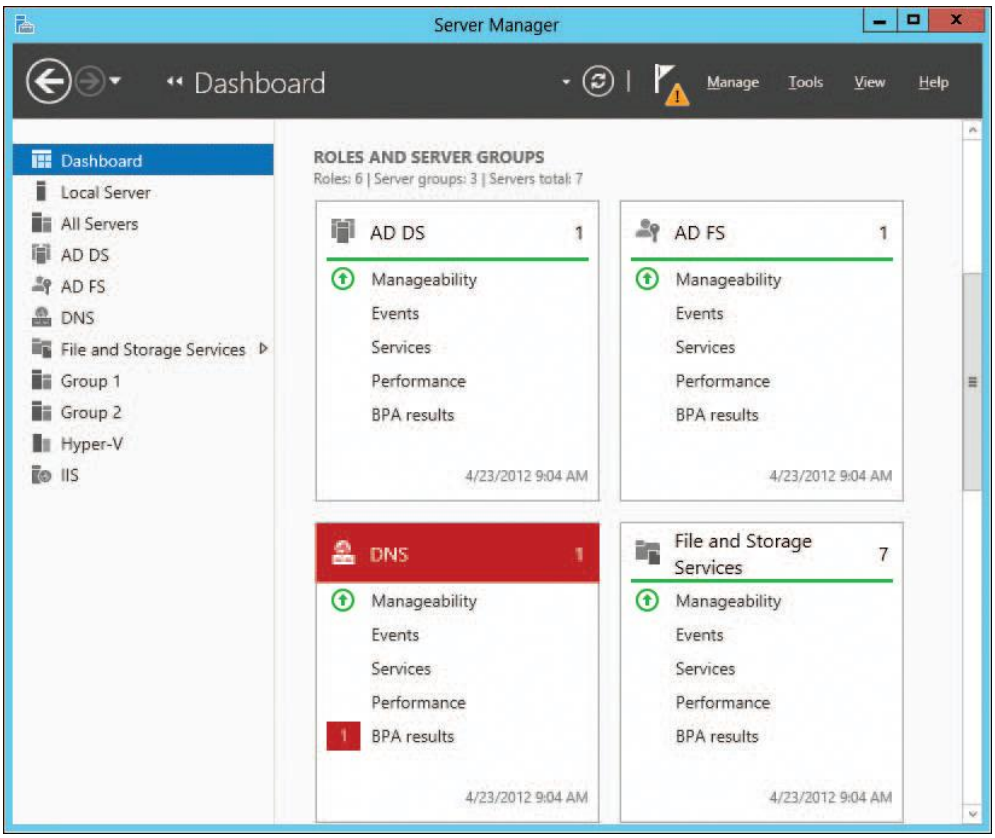

Clicking the alert brings up the details of the alert, indicating a possible problem with the configuration of one of the DNS servers in the environment:

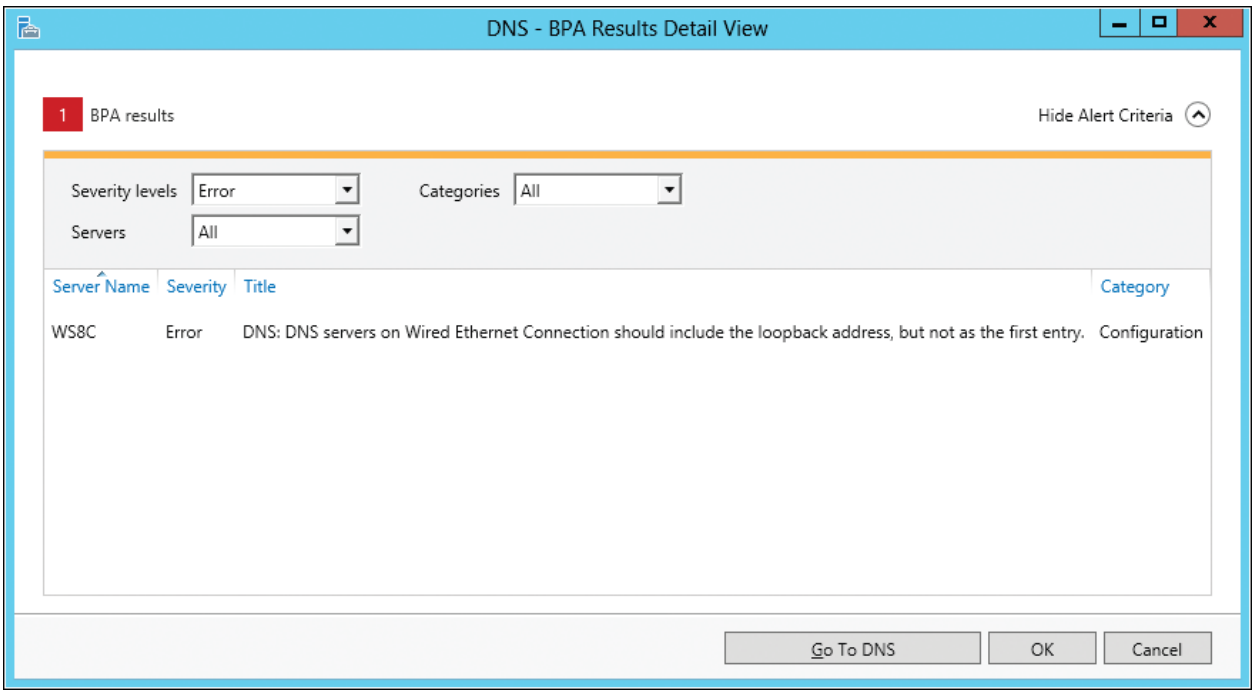

The Local Server section of Server Manager lets you view and configure various settings on your local server. You can also perform various actions on the local server, or on other servers in the available pool, by using the Manage and Tools menus. For example, you can add new roles or features to a server by selecting Add Roles And Features from the Manage menu:

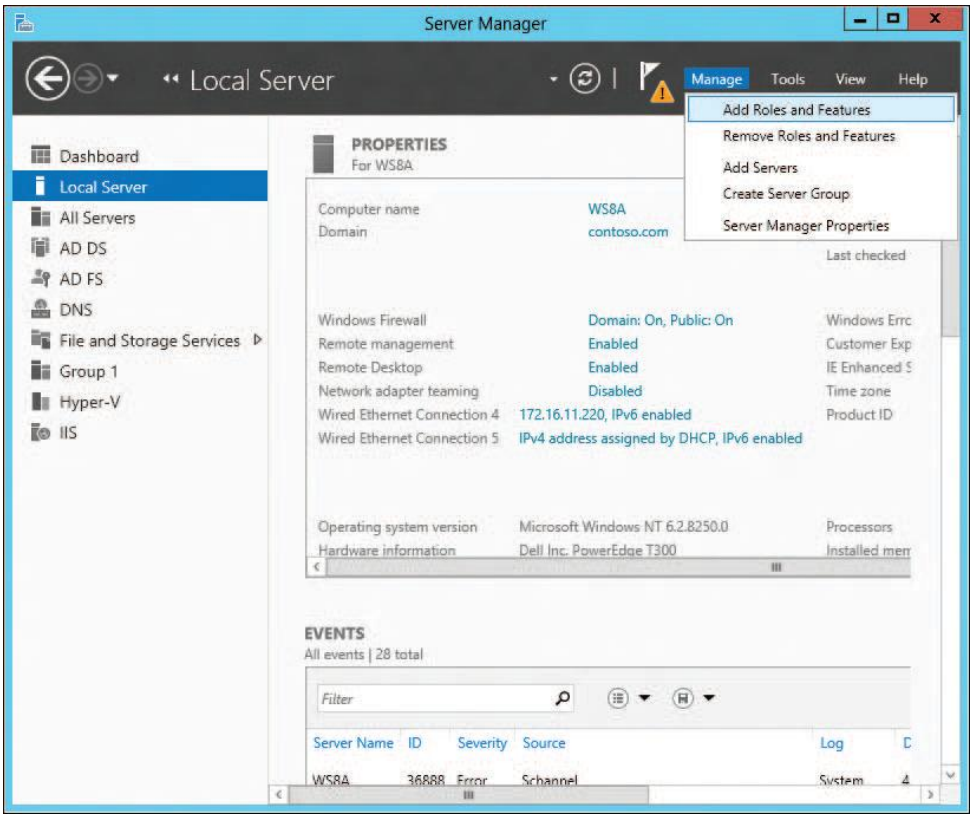

The Select Destination Server page of the new Add Roles And Features Wizard lets you select either a server from the server pool or an offline VHD as your destination server. The ability to provision roles and features directly to offline VHDs is a new feature of Windows Server 2012 that helps administrators deploy server workloads in virtualized data centers:

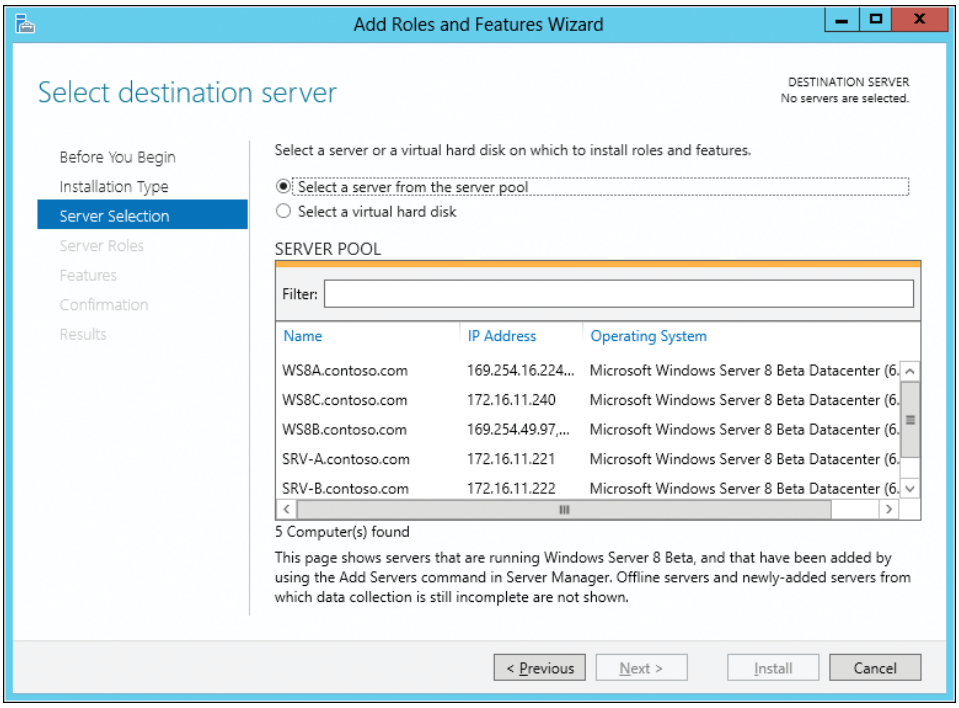

The All Servers section of Server Manager displays the pool of servers available for management. Right-clicking a server lets you perform different administrative tasks on that server:

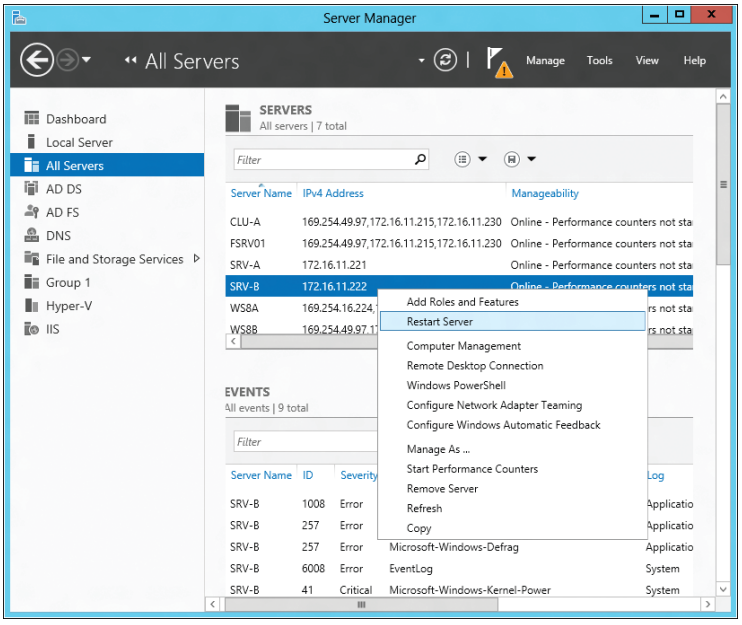

To populate the server pool, right-click All Servers in Server Management and select Add Server from the shortcut menu. Doing this opens the Add Servers dialog box, which lets you search for servers in Active Directory, either by computer name or IP address or by importing a text file containing a list of computer names or IP addresses. Once you've found the servers you want to add to the pool, you can double-click them to add them to the Selected list on the right:

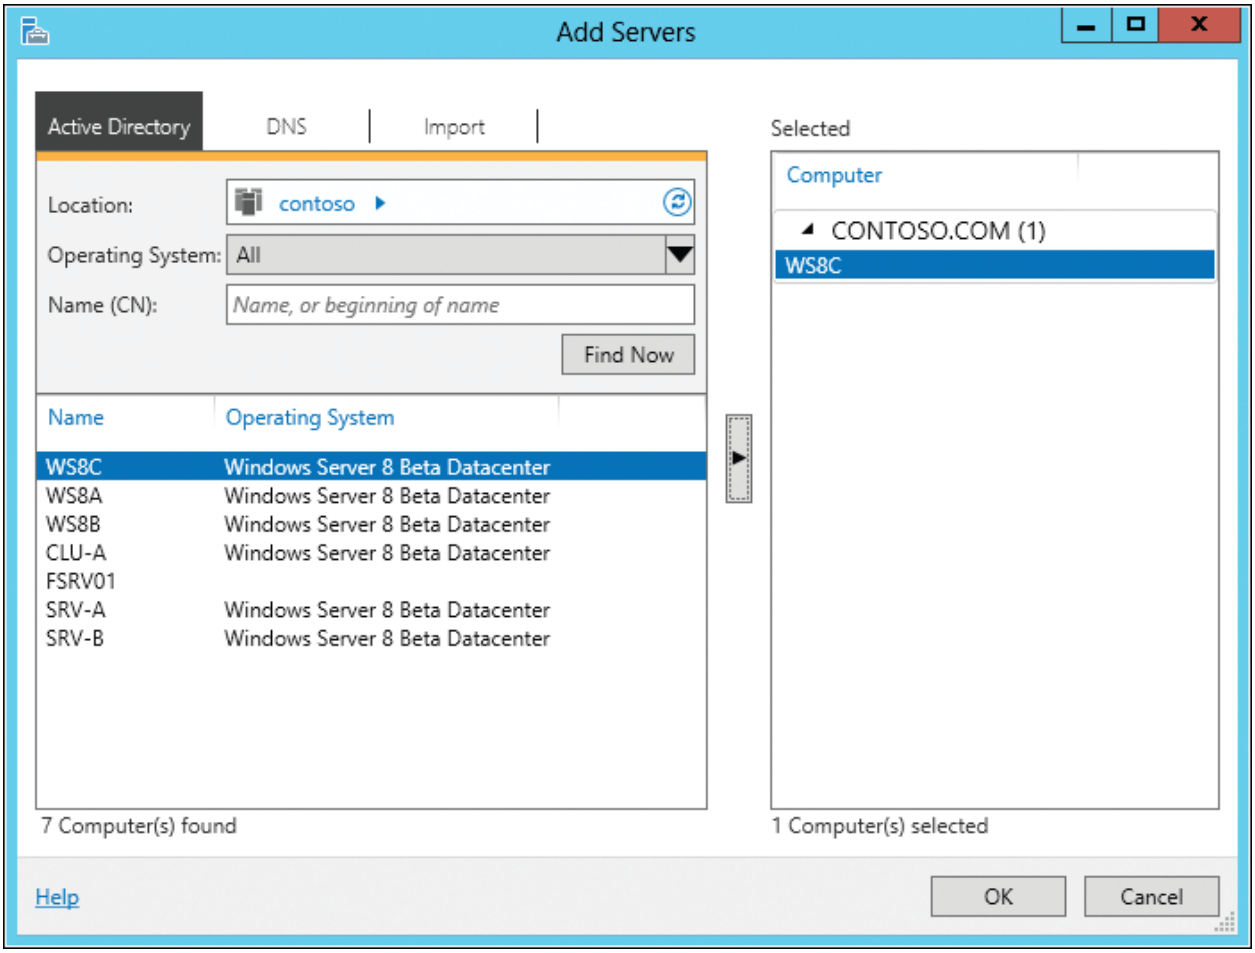

Servers are often better managed if they are grouped together according to their function, location, or other characteristics. Server Manager lets you create custom groups of servers from your server pool so that you can manage them as a group instead of individually. To do this, select Create Server Group from the Manage menu at the top of Server Manager. Doing this opens the Create Server Group dialog box, which lets you specify a name for the new server group and select multiple servers from your server pool to add to the group:

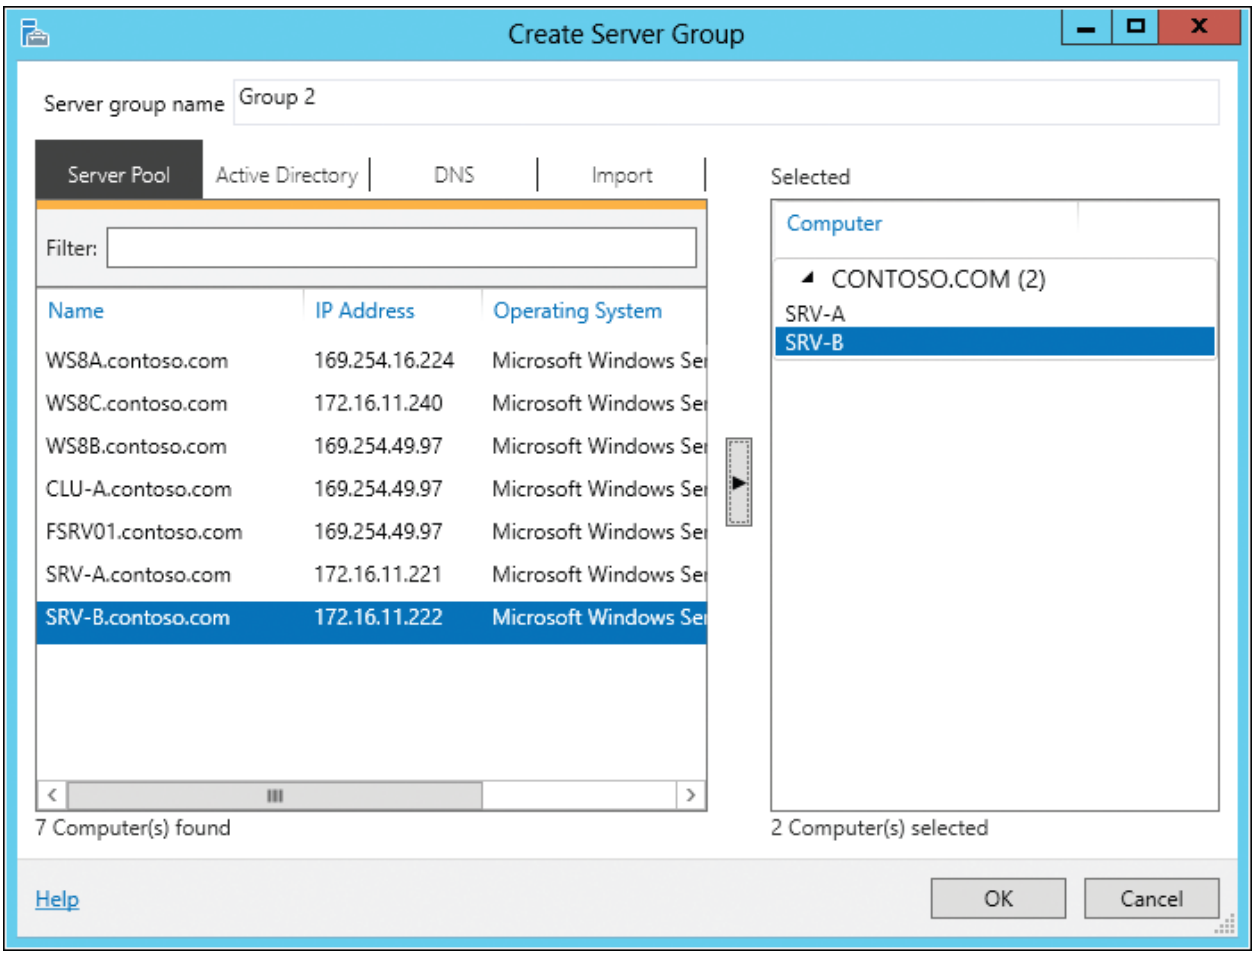

Once you've added servers to your new group, you can select multiple servers in your group and perform actions on them such as restarting:

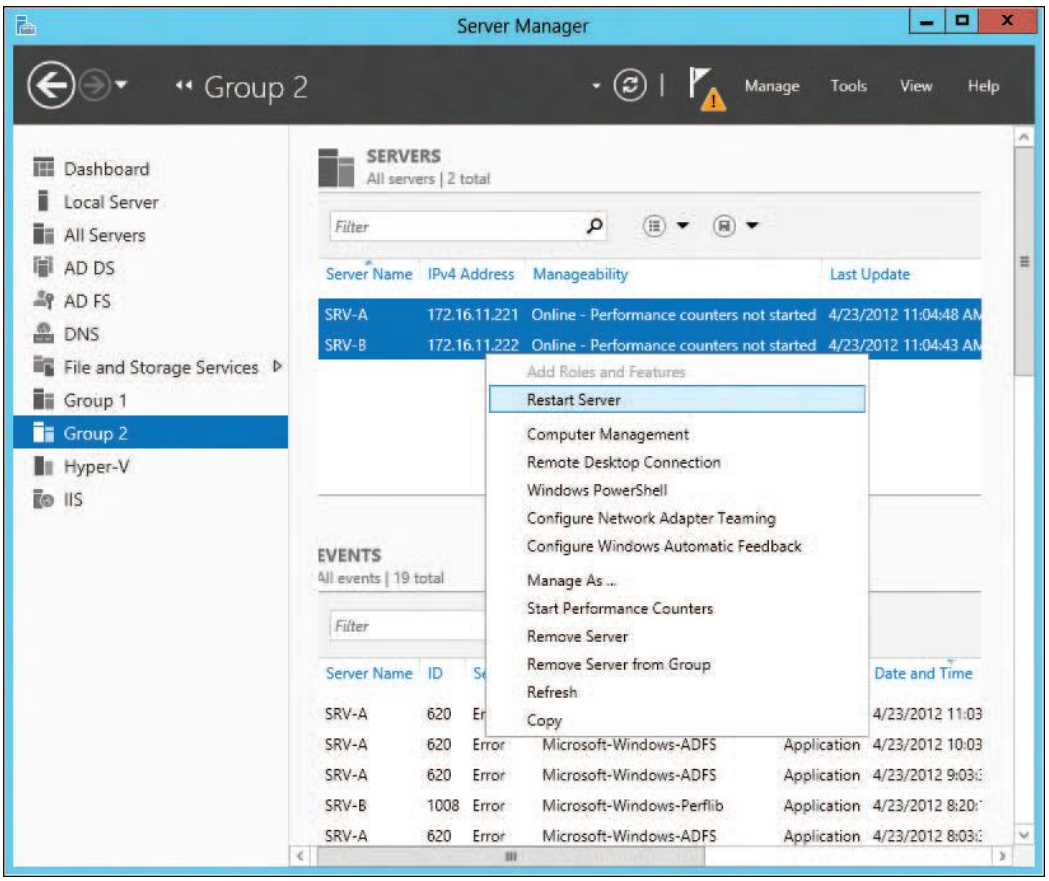

The Tools menu at the top of Server Manager can be used to start other management tools, such as MMC consoles. However, as the new Server Manager of the Windows Server platform evolves toward a true multi-server management experience, such single-server MMC consoles will likely become tightly integrated into Server Manager. With Windows Server 2012, such integration is already present for two roles: Remote Desktop Services and file and storage management. For example, by selecting File And Storage Services, you can manage the file servers, storage pools, volumes, shares, and iSCSI virtual disks in your environment:

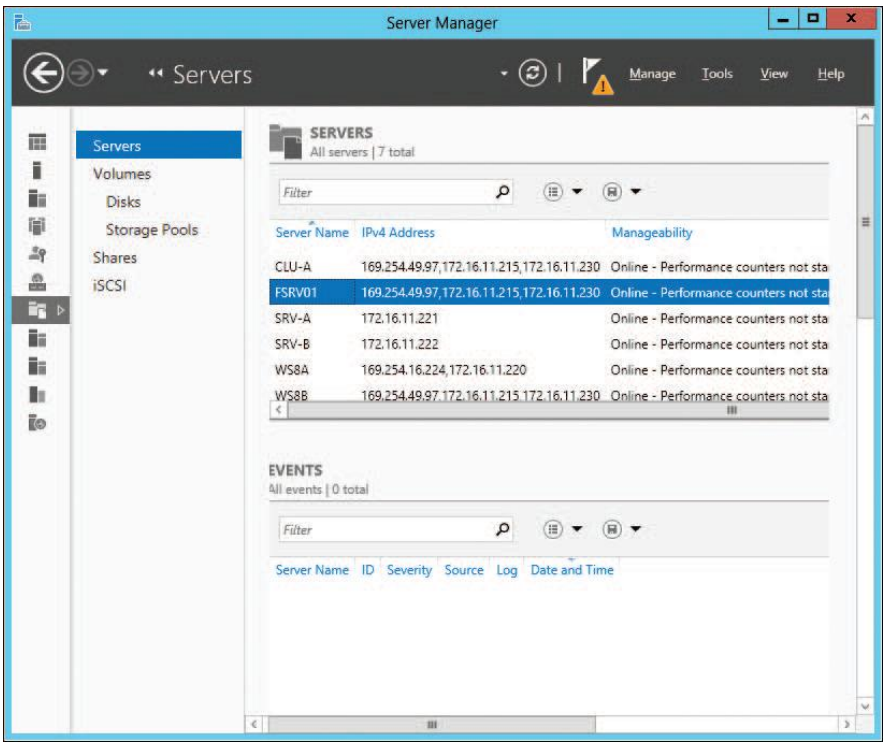
Server Manager
Servers
Manage
Tools
View
Help
Servers
Volumes
Disks
Storage Pools
Shares
iSCSI
SERVERS
All servers | 7 total
Filter
Server Name
IPv4 Address
Manageability
CLU-A
169.254.49.97,172.16.11.215,172.16.11.230
Online - Performance counters not sta
FSRV01
169.254.49.97,172.16.11.215,172.16.11.230
Online - Performance counters not sta
SRV-A
172.16.11.221
Online - Performance counters not sta
SRV-B
172.16.11.222
Online - Performance counters not sta
WS8A
169.254.16.224,172.16.11.220
Online - Performance counters not sta
WS8B
169.254.49.97,172.16.11.215,172.16.11.230
Online - Performance counters not sta
EVENTS
All events | 0 total
Filter
Server Name
ID
Severity
Source
Log
Date and Time