

# Deploying Microsoft Windows Compute Cluster Server 2003

## on Dell PowerEdge Servers

Microsoft® Windows® Compute Cluster Server 2003 (CCS) can help provide a simple, cost-effective way to deploy and manage clusters. This article discusses CCS installation and configuration on Dell™ PowerEdge™ 1950 servers.

BY RON PEPPER AND VICTOR MASHAYEKHI, PH.D.

### Related Categories:

Dell PowerEdge servers

High-performance computing (HPC)

Microsoft Windows Compute Cluster Server 2003

Microsoft Windows Server 2003

System deployment

Visit [www.dell.com/powersolutions](http://www.dell.com/powersolutions) for the complete category index.

In September 2006, Dell began bundling Microsoft Windows Compute Cluster Server 2003 (CCS) with its PowerEdge 1950 servers. This product enables administrators to create high-performance computing (HPC) clusters with servers running Microsoft Windows Server® 2003 x64 operating systems using a standard MPICH-based Message Passing Interface (MPI) library, and can allow easy code porting from UNIX® OS-based parallel applications to Windows.

CCS includes two components: the Windows Server 2003 Compute Cluster Edition OS and the Compute Cluster Pack (CCP). The Compute Cluster Edition is a limited version of Windows Server 2003 that does not allow some server services to function.<sup>1</sup> The CCP contains the necessary components to create server clusters, along with a job

scheduler, MPICH MPI library, and Microsoft Windows Remote Installation Services (RIS) extensions.

### Configuring cluster hardware

CCS is currently supported for Dell PowerEdge 1950 servers using embedded Ethernet interconnects as the compute fabric. Administrators can provide additional storage for the head node by adding a Dell PowerVault™ MD1000 disk expansion enclosure or network attached storage.

Some parallel applications do not benefit from Intel® Hyper-Threading Technology, so administrators may want to disable it on both the head node and compute nodes. Because the head node OS is installed manually or at the factory, administrators should also disable the Pre-boot Execution Environment (PXE) on the head node.

<sup>1</sup> For more information about these limitations, see the Windows Server 2003 Compute Cluster Edition end-user license agreement.

They should enable PXE on the compute nodes and place the first embedded network interface card (NIC) before the local hard drive in the system boot order.

Figure 1 illustrates a CCS-based HPC cluster configuration. This configuration uses both head node NICs, with NIC1 connecting to the compute nodes and NIC2 connecting to the public network; the compute nodes use only NIC1. If the compute nodes require public network access, administrators can enable Internet Connection Sharing (ICS) on the head node or use the secondary network connection (NIC2) on the compute nodes.

The appropriate head node configuration is typically determined by the environment. If a domain controller already exists in the environment and administrators want to set up network access between the cluster and this environment, they can configure the head node as a member server in that Microsoft Active Directory® directory domain. However, if they are building a stand-alone cluster, then the head node must be its own domain controller.<sup>2</sup>

If administrators plan to reinstall the head node OS and software, they should typically use the Dell OpenManage™ Server Assistant CD provided with PowerEdge 1950 servers. This CD can help streamline the installation process and automatically install the network or storage drivers needed for embedded controllers. If administrators plan to automate the compute node installations using RIS, they should leave some storage space un-partitioned or use secondary disks, because RIS requires an independent drive (different from the system drive) where a copy of the OS image can be stored.

### Preparing the head node for Compute Cluster Pack installation

Before installing the CCP on the head node, administrators should configure this node as an Active Directory member server or domain controller. Stand-alone clusters also require a Domain Name System (DNS) server; when promoting the head node or another server to domain controller, administrators are prompted to set up a DNS server if one is not already present.

Using RIS requires a Dynamic Host Configuration Protocol (DHCP) service. Administrators should run this service on the cluster interconnect (NIC1 in the Figure 1 example), not the primary or public network (NIC2). If administrators plan to use RIS, they should also leave space for a second partition, or have additional disk(s) available.

Finally, administrators should apply any CCP and Microsoft Management Console (MMC) updates, including the following:

- ICS update for Windows Server 2003 x64 (available at [go.microsoft.com/fwlink/?linkid=55166](http://go.microsoft.com/fwlink/?linkid=55166))




Figure 1. Example Microsoft Windows Compute Cluster Server 2003–based cluster configuration

- RIS update for Windows Server 2003 x64 (available at [go.microsoft.com/fwlink/?linkid=55167](http://go.microsoft.com/fwlink/?linkid=55167))
- MMC 3.0 for Windows Server 2003 x64 (available at [go.microsoft.com/fwlink/?linkid=62400](http://go.microsoft.com/fwlink/?linkid=62400))

After installing these files, administrators must reboot the head node before installing the CCP.

### Installing the Compute Cluster Pack

Administrators can begin the CCP installation by launching the `setup.exe` file on the CCP CD. The CCP installer then helps ensure that the proper updates have been installed on the system. If the head node is connected to the Internet, the installer can download and begin installation of these patches as necessary.

During installation, administrators must select whether the head node will also be a compute node; if not, they should select the “Create a new compute cluster with this server as the head node” option without selecting the sub-option to include compute node installation. The installer, after providing several destination directory prompts, then installs Microsoft .NET Framework 2.0 and Microsoft SQL Server™ Desktop Engine—which are included on the CCP CD—and completes the CCP installation.

### Configuring the Compute Cluster Pack

Following CCP installation, a To Do List screen appears that includes four task sections: Networking, RIS, Node Management, and User Management (see Figure 2).

<sup>2</sup> For more information about installing a domain controller, visit [www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/stepbystep/dmctrnl.mspx](http://www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/stepbystep/dmctrnl.mspx).

## Networking task section

The Networking section includes the Configure Cluster Network Topology and Manage Windows Firewall Settings wizards to help simplify configuration. The Configure Cluster Network Topology wizard displays various possible network configurations—“Compute nodes isolated on private network,” for example, places only the head node on a public network such as the Internet or a corporate intranet and connects the compute nodes only to the head node. If administrators select this option, the installer prompts them to select a network connector for each network—in the Figure 1 configuration, the private (MPI) network uses NIC1 and the public network uses NIC2. If administrators want to set up compute node access to the public network, they can also enable ICS at this point.

The Manage Firewall Settings wizard enables or disables public network firewall settings, which should typically be

enabled. Administrators can later provide firewall access to individual services as needed.

## RIS task section

The wizards in this section enable administrators to install and uninstall RIS and manage OS images. Installing and configuring RIS can help administrators save time by automating compute node OS installation. Even if the compute node operating systems were factory installed, administrators must still add the nodes to Active Directory and install the CCP, which can be time-consuming to perform manually even for small clusters.

Administrators can use the Install RIS wizard to install the necessary OS components; this process may require the head node OS CD. After RIS is installed, the Manage Images wizard becomes available, which administrators can use to install or remove OS images and manage OS product keys. Following initial deployment of a head node and RIS, administrators can launch this wizard and select “Add new image,” then follow a series of prompts to copy an OS image to the previously prepared RIS partition. This process requires the compute node OS CD; administrators should keep in mind that copying files from this CD to the system can be time-consuming.

After creating an image on the head node, administrators should run the Manage Images wizard again and select “Modify image configuration,” which allows them to change the image description and the product key used for installation. They can provide the key manually or have the wizard search the installation CD for one. At this point they should also add other necessary device drivers, as described in the “Adding specific drivers for Dell PowerEdge 1950 servers to the Remote Installation Services OS image” sidebar in this article.

## Node Management task section

The Node Management section consists of two wizards that allow administrators to add or remove cluster nodes. Administrators can add nodes manually or perform an automated deployment. When adding a node manually, they must ensure that the compute node is connected to the head node on the appropriate network and have local administrator access on that system. Administrators must also add the system to Active Directory if it is not already a member; they can then install the CCP and identify the head node, after which the CCP can add the node to the cluster.

Performing an automated deployment helps simplify the process of installing the OS, adding the system to Active Directory, and installing the CCP. Before performing this deployment, administrators must install RIS and prepare a proper OS image using the wizards in the RIS section. They must also provide a username and password for a user allowed to create Active

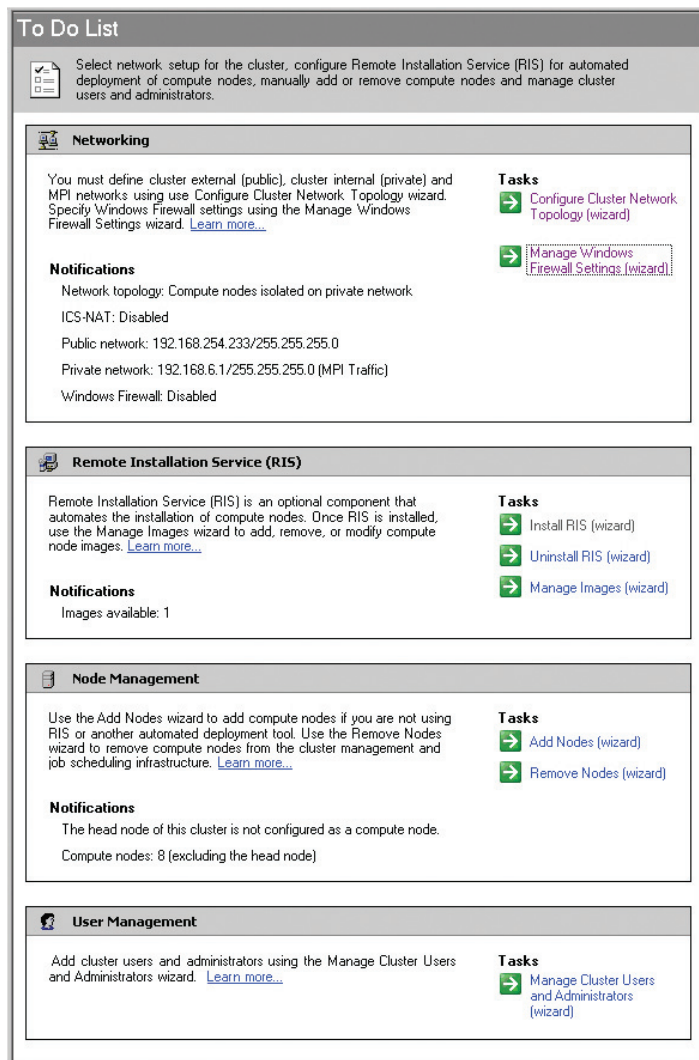


Figure 2. To Do List screen following Compute Cluster Pack installation

## ADDING SPECIFIC DRIVERS FOR DELL POWEREDGE 1950 SERVERS TO THE REMOTE INSTALLATION SERVICES OS IMAGE

To complete CCS configuration on Dell PowerEdge 1950 servers, administrators must install additional drivers for Dell PowerEdge Expandable RAID Controller (PERC) 5/i, SCSI/RAID, and Broadcom NetXtreme II devices. They can download these drivers from support.dell.com and integrate them into the RIS OS image by performing the following steps:

1. Open Windows Explorer and navigate to the image directory on the RIS image partition. Assuming that the D:\ drive is the RIS image partition and the default settings were used during the RIS OS image creation, this directory would be D:\RemotInstall\Setup\English\Images\WINDOWS.
2. Create an \$OEM\$ directory, then create two sub-directories in this directory—textmode and \$1\drivers\nic.
3. Run the Broadcom driver package and extract its files to C:\Broadcom\W2K364, assuming C:\ is the system boot directory.
4. Copy the files in C:\Broadcom\W2K364\RIS\_Drivers to the amd64 and \$OEM\$\\$1\drivers\nic sub-directories of D:\RemotInstall\Setup\English\Images\WINDOWS.
5. Execute the setup.exe program with the -a command-line option by going to Start > Run and entering C:\Broadcom\W2K364\setup.exe -a. This command extracts the additional required Plug and Play device drivers.
6. When prompted, enter C:\Broadcom as the network location.
7. Copy all the files from the Win2K3SNP\x64 and vbd\x64 sub-directories of C:\Broadcom\Program Files\Broadcom\Broadcom Driver and Management Applications\NetXtreme II to D:\RemotInstall\Setup\English\Images\WINDOWS\\$OEM\$\\$1\drivers\nic.
8. Copy the .inf and .sys files from the \$OEM\$\\$1\drivers\nic sub-directory of D:\RemotInstall\Setup\English\Images\WINDOWS to the amd64 sub-directory.
9. Extract the PERC 5/i drivers to D:\RemotInstall\Setup\English\Images\WINDOWS\\$OEM\$\textmode, which may require running an executable installer and then accessing the location of the installed files (for example, C:\Del\PERC5).
10. Copy the exact text in the SCSI section of the txtsetup.oem file—for example, DELL PERC 5 RAID Controller Driver (Windows Server 2003 x64)—and paste it into another file. This text can change between driver revisions.
11. Edit the ristndrd.sif file in D:\RemotInstall\Setup\English\Images\WINDOWS\amd64\templates. First, add a MassStorageDrivers section and add the SCSI section text copied in step 10. For example:
 

```
[MassStorageDrivers]
"DELL PERC 5 RAID Controller Driver (Windows
  Server 2003 x64)"="OEM"
```

Next, add an OEMBootFiles section and list the files in D:\RemotInstall\Setup\English\Images\WINDOWS\\$OEM\$\textmode, excluding .txt files:

```
[OEMBootFiles]
nodev.inf
oemsetup.inf
percsas.cat
percsas.pdb
percsas.sys
txtsetup.oem
```

Add the following line to the Unattended section:

```
OemPnpDriversPath="\drivers\nic"
```

Finally, save and close the file.
12. Restart RIS by opening a command prompt and entering net stop binlsvc and net start binlsvc.

Directory objects (typically a domain administrator). After providing this information, administrators can enter a node series name, which is used to provide consistent, sequential names for compute nodes—for example, if they provide “compute-” as the series name, the compute nodes would be named compute-001, compute-002, compute-003, and so on.

After administrators have accepted the end-user license agreement, they can click “Start RIS” on the Image Nodes screen to start RIS; they can then PXE boot the compute nodes to image them. RIS formats and completely re-images any system that is PXE booted on the private network at this time. If any compute nodes have previously been imaged, the wizard prompts administrators to press the F12 key when they are PXE booted to image the system again. After RIS has imaged the compute nodes, administrators must stop RIS before finishing the wizard. Figure 3 shows the Result screen, which lists the added nodes.

**User Management task section**

The Manage Cluster Users and Administrators wizard in the User Management section allows administrators to configure Active Directory users as either cluster users or cluster administrators. Cluster users can submit jobs to the cluster; cluster administrators can both submit jobs and cancel, pause, and rearrange jobs in the job scheduler.

**Approving installed compute nodes**

As a final step before the cluster can run jobs, administrators must launch Compute Cluster Administrator and select “Node

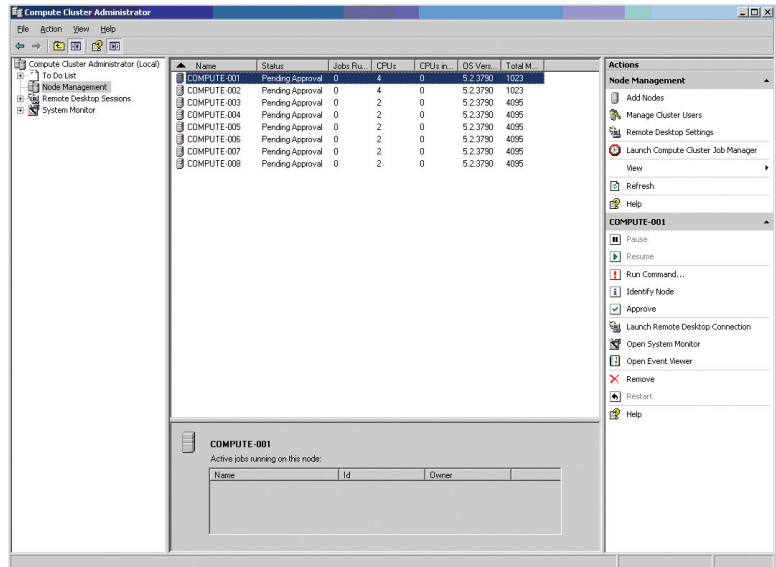


Figure 4. Compute Cluster Administrator Node Management screen showing newly installed compute nodes pending approval

Management.” They can then approve and un-pause the newly installed cluster compute nodes by selecting the compute nodes from the list and clicking “Approve” and “Resume” in the Actions pane (see Figure 4).

**Enabling simplified cluster installation and management**

Microsoft Windows Compute Cluster Server 2003 provides a comprehensive cluster deployment and management system for Dell PowerEdge 1950 servers running Windows Server 2003 x64 operating systems. Implementing Windows Compute Cluster Server 2003 can help administrators deploy and manage HPC clusters efficiently and cost-effectively.

**Ron Pepper** is a systems engineer and adviser in the Scalable Systems Group at Dell. He works on the Dell HPC Cluster team developing grid environments. Ron attended the University of Wisconsin at Madison, where he worked on a degree in Computer Science; he is continuing his degree at St. Edward’s University.

**Victor Mashayekhi, Ph.D.**, is the engineering manager for the Scalable Systems Group at Dell, and is responsible for product development of cluster offerings. His current research interests are HPC and high-availability clusters, virtualization, distributed systems, interconnect technologies, and computer-supported cooperative work. Victor has a B.A., M.S., and Ph.D. in Computer Science from the University of Minnesota.

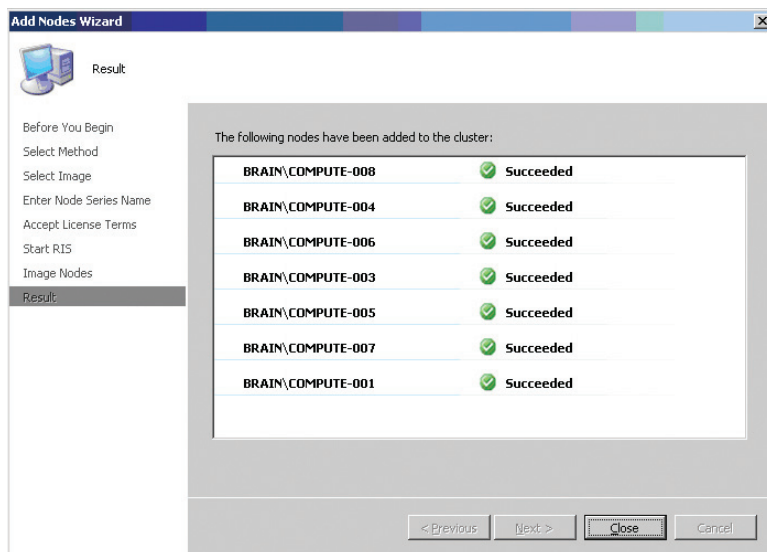


Figure 3. Result screen following completion of the Compute Cluster Pack Add Nodes wizard