

# Enhancements to Microsoft Cluster Server for Windows 2000, Advanced Server

By Mike Kosacek and Edward Yardumian

**As mission-critical applications are developed for Windows 2000 and its new technologies such as Active Directory, PowerEdge clusters running Windows 2000 Cluster service should be deployed when high availability is a requirement. Windows 2000 Cluster service offers tighter integration and implements the new infrastructure of networking, storage, plug-and-play, and hotplug underlying Windows 2000 Server.**

A cluster can be loosely defined as two or more systems working together to provide a single-system image to clients and administrators. The usual goals of a cluster include availability, scalability, manageability, or any combination of these.

A high-availability (HA) cluster, also referred to as a failover cluster, is one common implementation of clustering. From the client perspective, the application or resource should always be available to service requests, no matter which node (or nodes) the application is running on. An HA cluster allows an application to run on multiple nodes, but multiple nodes usually do not run the application simultaneously. For example, if a database is running on a two-node HA cluster, one of the nodes will service requests (database queries). If the node actively running the database services fails, then the other node will take over and restart the database.

## Microsoft Cluster Server for Windows NT Server 4.0

The first Microsoft cluster product, Microsoft Cluster Server (MSCS) 1.0, was delivered with Windows NT Server 4.0

Enterprise Edition. MSCS (code named “wolfpack”) implemented a two-node HA cluster. Out of the box, the product enabled the core services in Windows NT—including file, print, and Web—to restart or failover resources from one node in the cluster to the other.

Microsoft also created a cluster application programming interface (API) to allow third parties to make their applications MSCS cluster aware. Microsoft followed up closely with several cluster-aware applications, such as SQL Server and Exchange Server. Some additional applications, such as Oracle Fail Safe, also fully supported the MSCS environment.

## Windows 2000 Cluster Service

With the Windows 2000 Server family, Microsoft has enhanced its cluster offering to include IP load balancing and applications based on Microsoft Component Object Model (COM) in addition to its HA cluster software. MSCS, now called Windows 2000 Cluster service, is an option included in both Windows 2000 Advanced Server and Windows 2000 Datacenter Server, and can be classified as an HA cluster.

Windows 2000 Advanced Server supports two-node failover, and Windows 2000 Datacenter Server is expected to support failover for up to four nodes. Both versions of Windows 2000 Cluster service will offer many enhancements over MSCS 1.0.

### Installation and Setup During Initial OS Setup

For installation with Windows NT Server 4.0 Enterprise Edition, MSCS 1.0 resembles a separate piece of software more than an integrated feature of the operating system (OS). Windows 2000 Cluster service, however, can be installed during the initial OS installation process. If Cluster service is selected during the initial installation, Windows 2000 copies the necessary files onto the system.

Because some aspects of the Cluster service configuration, such as the shared disk configuration, must be defined after the initial OS configuration, Cluster service is not configured during the OS setup. The Cluster Service Configuration Wizard is launched manually after Windows 2000 setup. Windows 2000 reminds the administrator to finish the setup by following the instructions in the Configure Your Server dialog box that automatically pops up after installation if any installed service needs further configuration.

Cluster service can also be installed following the initial installation by using the Add/Remove Programs icon in the Control Panel and the Add/Remove Windows Components utility (see Figure 1).

Selecting Cluster Service from the Windows Components Wizard adds the software to the computer and automatically launches the wizard to configure Cluster service. Windows 2000 also eliminates the need to reboot after Cluster service is configured.

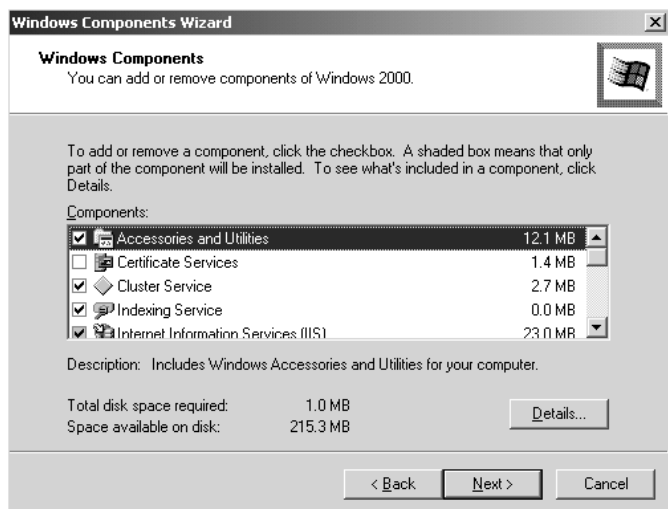


Figure 1. The Microsoft Add/Remove Windows Components Dialog Box

## Windows 2000

*eliminates  
the need to  
reboot after  
Cluster service  
is configured.*

Windows 2000 Advanced Server supports three methods of installation:

- A new installation
- An upgrade from Windows NT Server 4.0 Enterprise Edition
- A rolling upgrade from Windows NT Server 4.0 Enterprise Edition Microsoft Cluster Server

A standard upgrade involves stopping the cluster and upgrading both nodes simultaneously. This requires the cluster to be off-line during the upgrade. A rolling upgrade is performed by moving all of the cluster's virtual

servers to Node B, stopping Cluster service on Node A, and upgrading Node A to Windows 2000 Advanced Server. Once the installation is complete on Node A, resources are moved back to Node A, and the process is repeated on Node B. Cluster services are only off-line for the time necessary to move the groups from node to node.

### Plug-and-Play Networking

MSCS 1.0 cannot detect a failure in the public network used to communicate with clients. For example, if a cable connecting a network interface card (NIC) used to process client requests failed, MSCS would not detect a failure as long as the other server was still seeing the heartbeat through the (private) cluster interconnect. Network cards that support adapter teaming, such as the Intel® Pro 100 series, can help to minimize the possibility of failure.

With Windows 2000 plug-and-play networking, the Windows 2000 Cluster service recognizes this type of failure in any network segment—public or private. If a public network segment fails, Cluster service looks for any IP address resources that use the failed segment, and it moves the appropriate corresponding groups to the other node.

Adapter teaming can still be used with Windows 2000 to add another layer of protection for the public network. Even greater system protection can be achieved by plugging the nodes and redundant NICs into separate network switches (on the same subnet). Under that arrangement, if a network switch or cable fails, the redundant NIC can continue operations or the applications and groups will automatically failover to the other node, where they can continue to service client requests through another route.

Windows 2000 has also improved network adapter identification and integration. In MSCS 1.0, each network segment receives a name, such as Public or Private, during cluster configuration. This name, however, is exclusive to MSCS.

With Windows 2000, the functionality for naming a network segment was integrated into the core operating system. The Windows 2000 network plug-and-play service keeps the

naming of the networks synchronized with all nodes of the cluster. When Cluster service is configured on the first node, it will use the names in the Network Connections folder for the network segments. When the second node joins the cluster, it applies the network names used on the first node. If the networks are renamed from the Cluster Service Configuration or the Network Connections folder on any node, all nodes will reflect the change.

Dell supports Hotplug Peripheral Component Interconnect (HPPCI) infrastructure for specific hardware and software configurations only. Using HPPCI, a failed NIC can be replaced without rebooting the server (a hot replace). If virtual servers running on the cluster node with the failed NIC have been moved to another cluster node, they can be moved back automatically (if failback is enabled) or manually after initializing the replacement NIC.

In addition, through HPPCI, a NIC for a new network segment can be added without rebooting (a hot add), and the Cluster Service with virtual servers can use the new segment instantly. Many of the Dell departmental and enterprise-class servers, including the PowerEdge 4400 and PowerEdge 8450, support HPPCI. (Contact your Dell representative for more information.)

### Plug-and-Play Storage

Windows 2000 also adds plug-and-play storage capability. Windows NT Server 4.0 required a reboot to recognize new disks added to the disk subsystem, but Windows 2000 has

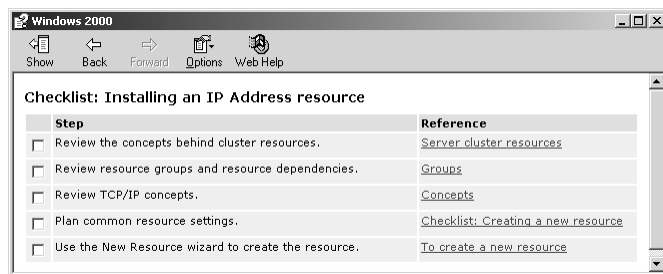


Figure 2. Checklist for Creating an IP Address Resource



Figure 3. Configure Cluster Application Wizard

facilities to scan for and recognize new disks. Users can name and partition new disks, give them drive letters, and use them without rebooting. If new disks are configured on each node, Cluster service can immediately use these disks as disk resources.

Note that if a disk is removed from the shared storage system, Cluster service does not assume it was done intentionally. It treats the loss of the disk as a failure, and therefore does not remove the corresponding disk resource and dependent resources. If the removal was intentional, the user must manually remove the disk and dependent resources.

### Ease of Use and Operation

The only documentation or configuration information provided with MSCS 1.0 was the *MSCS Administration Guide*, included on the CD-ROM. With Windows 2000, online help for clusters is readily available. A new help feature using checklists is also included in the help for clustering. Windows 2000 provides checklists for installing all the built-in cluster applications and for many other cluster procedures.

Since all the help files are linked, users can find help immediately. If users do not know how to perform a step during the installation, they need only a single mouse click to find information on the procedure, and possibly another checklist.

Figure 2 shows an example of the checklist for creating an IP address resource. The checklist shows several required steps in the left column, and concepts or procedures associated with each step in the right column.

Within the Cluster Administrator, Windows 2000 Cluster service adds wizards for setting up resource groups for Windows 2000 core applications. A wizard is launched from the Cluster Administrator (see Figure 3) by choosing the Configure Application item from the File menu or by right-clicking over the Cluster Administrator. The wizard helps administrators to create virtual server groups for the built-in resources listed in Figure 4.

Visually, the Windows 2000 Cluster Administrator differs little from the one in Windows NT Server 4.0. Microsoft simplified the interface, as shown in Figure 5, by organizing the folders for resource types (cluster resource .DLL) files and network adapters under a separate folder called Cluster Configuration.

### Clustering Windows 2000 Core Services

With Windows 2000, Microsoft adds support for dynamic host configuration protocol (DHCP) and other core services, including Windows Internet Naming System (WINS) and stand-alone Distributed File System (DFS) roots. Internet Information Services (IIS) 5.0, included in all versions of Windows 2000 Server, can be made highly

## BUILT-IN RESOURCES FOR VIRTUAL SERVER GROUPS

- Dynamic Host Configuration Protocol (DHCP) Service
- Distributed Transaction Coordinator
- File Share
- Generic Application
- Generic Service
- Internet Information Services Server Instance
- IP Address
- Message Queuing
- Microsoft Message Queue Server
- Network Name
- Network News Transfer Protocol (NNTP) Server Instance
- Physical Disk
- Print Spooler
- Time Service
- Simple Message Transfer Protocol (SMTP) Server Instance
- Windows Internet Naming Service (WINS)

Figure 4. Built-in Resources for Virtual Server Groups

available through the IIS Server Instance resource included with Cluster service.

### Clustering Applications

Applications that were written to be cluster aware for MSCS 1.0 do not need to be revised to support Windows 2000 Cluster service. If the application supports MSCS 1.0 and Windows 2000, it should work with Windows 2000 Cluster service. Microsoft's upcoming SQL Server 2000 and Exchange Server 2000 will be cluster aware.

Many applications not designed to be cluster aware can still run on a Windows 2000 cluster, using a virtual server and failover from node to node in case of a hardware failure. However, these applications may not be able to take full advantage of some of Cluster service's more advanced features. Any application not designed to be cluster aware should be thoroughly tested before production deployment in a Windows 2000 Cluster service environment.

### Cluster-Aware Tape Backup

The most common way to back up MSCS 1.0 configurations is across one of the following networks: the client (public) network, a dedicated network using a high-speed interconnect such as Gigaset's cLAN, or a storage area network (SAN) using Dell's PowerVault line of Fibre Channel storage and SANs. Using internal tape backup devices in cluster nodes is not optimal because clustered applications can run on any cluster node. For example, suppose Node 1 backs up Application A, and Node 1 fails. Application A will then move to Node 2, where it will continue to run and be updated while its backup server, Node 1, is off or in a failed state.

In backing up cluster configurations over a network, an administrator can point the backup software at the virtual servers running on the cluster. The backup software will not care which server the virtual server runs on, only that it is available and ready for backup.

Cluster-aware tape backup software has not become widely available for MSCS 1.0 because the API lacks

facilities for tape backup software; that is, there is no interface for creating the software. The Windows 2000 Cluster API adds these extensions, which enables backup software vendors to release new versions of their backup software with agents that fully support the Microsoft Cluster service environment.

With MSCS 1.0, it was difficult to back up the cluster's configuration (virtual servers, IP addresses, groups, and so on). Most common backup utilities, including Microsoft Backup, cannot back up or restore the metadata that describes the cluster configuration. Dell OpenManage Cluster Assistant with ClusterX does offer this capability for MSCS 1.0. The Windows 2000 Backup utility also adds the ability to back up both the configuration database and the quorum resource for the cluster.

### Enhancements in Windows 2000 Datacenter Server

Cluster service for Windows 2000 Datacenter Server is expected to include all the features of the two-node Cluster service for Windows 2000 Advanced Server. The major difference between the two services is that Cluster service for Windows 2000 Datacenter Server will support four-node

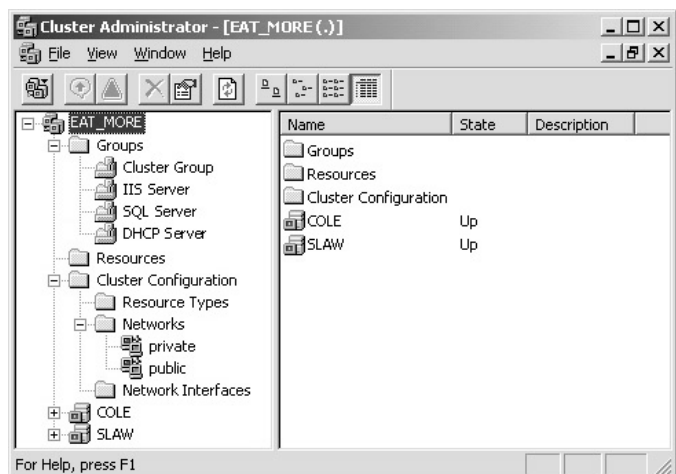


Figure 5. Windows 2000 Cluster Administrator

failover. Often called the Windows Multinode Cluster, the four-node Cluster service will enable one-to-many failovers of applications.

Another notable difference is that unlike the two-node Cluster service, which can use either SCSI or Fibre Channel for the shared storage array, the four-node Cluster service requires Fibre Channel shared storage.

### Developing Cluster-Aware Applications

Since the introduction of MSCS, major application vendors, including Oracle, IBM, Lotus, Tivoli®, and SAP, have developed applications that take full advantage of the monitoring and failover capabilities of MSCS. These applications communicate with the cluster through Microsoft's Cluster API. Windows 2000 Cluster service supports these applications, as well as new applications that use the Component Object Model to communicate with the cluster. The inclusion of COM support gives developers an additional option for having their applications communicate with the cluster.

Clustering is a core component of the Microsoft Application Specification for Windows 2000 for Servers. Before Microsoft will permit an application for Windows 2000 Advanced Server or Windows 2000 Datacenter Server to display the Certified for Windows logo, the application must meet either the core specification for desktop applications or the more comprehensive specification for server applications.

According to Microsoft Server Specification 1.2: "An application that meets this specification...is ready to run on a server cluster. By exploiting Cluster service features, it can minimize the service downtime caused by system failures or planned server maintenance and upgrades."

Therefore, application vendors that wish to qualify for a Windows Certification logo (see Figure 6) for the Windows 2000 Advanced Server, Windows 2000 Datacenter Server, or both, will, by default, deliver cluster-aware applications ready for the two-node or four-node Cluster service.

### PowerEdge Clusters Produce High Availability

PowerEdge clusters running Windows 2000 Cluster service should be deployed when high availability is a requirement

---

*Many applications not designed to be cluster aware can still run on a Windows 2000 cluster, using a virtual server and failover from node to node in case of a hardware failure.*

---

for mission-critical applications running under Windows 2000. Windows 2000 Cluster service more effectively implements and more tightly integrates the new infrastructure underlying Windows 2000 Server, including networking, storage, and plug-and-play.

Feature enhancements, additional application support, extensive online help, and expanded support for backup solutions will make Windows 2000 Cluster service easier to implement and manage than MSCS 1.0. In short, Windows 2000 Cluster service is an important feature of the Windows 2000 Server family and should be strongly considered for mission-critical applications such as databases, groupware, manufacturing software, and file systems. ♦

### Additional Resources

The following sources provide additional information about Windows 2000 Cluster service:

- **The Microsoft Application Specification:**  
<http://msdn.microsoft.com/certification/appspec.asp> (December 8, 1999)
- **Windows Clustering Technologies: Cluster Service Architecture, a Microsoft white paper:**  
<http://www.microsoft.com/windows2000/library/howitworks/cluster/clusterarch.asp> (October 26, 1999)

*Mike Kosacek (michael\_kosacek@dell.com) is in the Cluster Development group at Dell Computer Corporation. As the lead engineer on several Dell products, Mike's responsibilities include developing, testing, and certifying cluster solutions for the Dell server and storage product lines. Mike has an AAS in Electronics Technology and is a Microsoft Certified Systems Engineer (MCSE).*

*Edward Yardumian (edward\_yardumian@dell.com) is a member of the Internet Infrastructure Solutions team and Cluster Development group at Dell Computer Corporation. Edward's responsibilities include developing, testing, and certifying Dell PowerEdge clusters. He is currently the lead engineer for Windows 2000 cluster projects.*



Figure 6. Certified for Windows Logo

Test your applications before they go live.  
Visit a Dell Application Solution Center (ASC).  
For more information, see ASCs at [www.dell.com](http://www.dell.com).