

# Clustering Options for Windows 2000

*By Mike Kosacek and Edward Yardumian*

**Windows 2000 clustering components—Cluster Service, Network Load Balancing, and Component Load Balancing—can be used to create highly available, scalable environments for many types of applications. These components can be combined with Dell PowerEdge servers, Dell PowerVault storage, a high-speed, low-latency interconnect such as the Giganet cLAN adapter, and a management solution such as Dell OpenManage Cluster Assistant with ClusterX to build cluster solutions that meet the scalability and availability requirements of high-end, enterprise-class applications.**

With Windows 2000, Microsoft expands its cluster options to achieve both availability and scalability. This article discusses four different clustering and load-balancing mechanisms within the Windows 2000 server product line (see Figure 1).

On the surface, it may appear that these are different options for doing the same thing, but that is far from true. These complementary solutions can be used separately to provide high availability and load balancing for specific types of applications, or they can be combined to provide true application solutions. The foundation of the Dell® Scalable Enterprise Computing initiative is to build solutions by interconnecting and managing multitier and multicluster configurations using industry-standard components that together provide scalability and high availability for an application.

## Windows 2000 Advanced Server Cluster Service: Two Node

Windows 2000 Advanced Server includes Windows 2000 Cluster service. Cluster service implements a two-node, failover cluster, as illustrated in Figure 2, that delivers high availability for applications and services. Both nodes connect to a common SCSI or Fibre Channel storage system.

As with Windows NT Server 4.0 Enterprise Edition clusters, both nodes can be configured to host resources (Active/Active), or one node can be simply a backup, always waiting for the first node to fail (Active/Passive). Windows 2000 Cluster service now supports plug-and-play for networking and disks and has greater support for core Windows NT applications such as dynamic host configuration protocol (DHCP), Windows Internet Naming System (WINS), and Distributed File System (DFS) root.

Windows 2000 Cluster service is positioned for dynamic, transactional applications:

- Databases such as Microsoft SQL Server, Oracle®, and IBM® DB2®
- Groupware such as Microsoft Exchange and Lotus® Domino™
- Enterprise resource planning (ERP) applications such as SAP™ and Baan™
- File sharing, print queues, and Web servers

| WINDOWS 2000 FAMILY CLUSTER OPTIONS | | | | | |
|---|---|---|---|---|---|
| Server Version | Processor/Memory Support | Cluster Service Two-Node | Cluster Service Four-Node | Network Load Balancing | Component Load Balancing and Web Farm Management |
| Windows 2000 Server | 1 to 4 Processors 4 GB RAM | No | No | No | No |
| Windows 2000 Advanced Server | 1 to 8 Processors 8 GB RAM | Yes | No | Yes | No |
| Windows 2000 AppCenter Server | Can be used with Server or Advanced Server | No | No | Optional | Yes |
| Windows 2000 Datacenter Server | 1 to 32 Processors | No | Yes | Yes | No |

Figure 1. Windows 2000 Server Family

Exchange 2000, SQL Server 2000, and Internet Information Services 5.0 are among the new cluster-aware applications written specifically for Windows 2000. Figure 3 shows a comparison of application uses for Windows Cluster services and Network Load Balancing (NLB).

Like Microsoft Cluster Server (MSCS) 1.0, the Windows 2000 Cluster service is somewhat hardware dependent (the storage system must be able to support multiple systems attached to it) and has a separate Microsoft certification specifically for Windows 2000. Each server-storage configuration must be thoroughly tested on Windows 2000 and pass a vigorous certification process before Microsoft will list it on the Hardware Compatibility List (see www.microsoft.com/hwtest).

Each configuration—a specific server with a specific RAID Controller or Fibre host bus adapter (HBA) and storage-RAID array—must be certified independently; all must be recertified when key components, such as a server

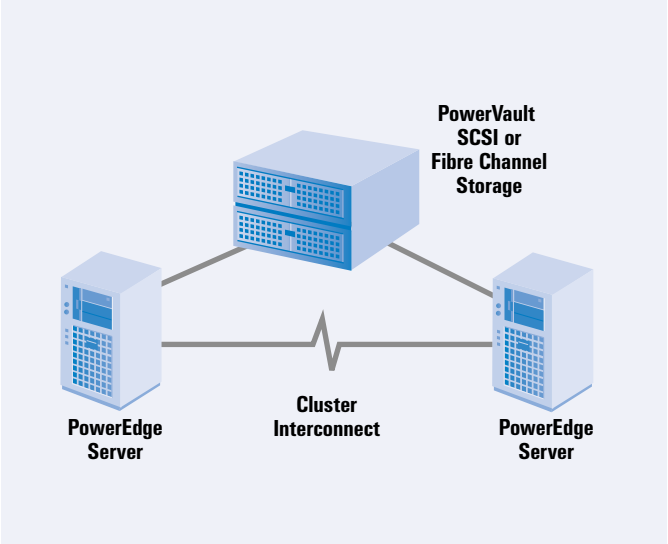


Figure 2. Cluster Service Two-Node Cluster

| MICROSOFT RECOMMENDED CLUSTER TECHNOLOGY APPLICATIONS | | | |
|---|---|---|---|
| | Technology | | |
| Scenario | Cluster Service | Network Load Balancing | Benefits |
| Web Server Farm | | ✔ | ■ Quickly expand capacity<br>■ Minimize site downtime |
| Terminal Servers | | ✔ | ■ Quickly expand capacity<br>■ Minimize effects of server failures |
| File/Print Servers | ✔ | | ■ Minimize service downtime<br>■ Ensure data consistency after failover |
| Database/Messaging | ✔ | | ■ Minimize application downtime<br>■ Ensure data consistency after failover |
| E-Commerce Sites | ✔ | ✔ | ■ Quickly expand capacity<br>■ Minimize the effects of server and application downtime |

Source: *Introducing Windows 2000 Clustering Technologies,* Microsoft (January 25, 2000)

Figure 3. Application Uses for Windows Cluster Service and Network Load Balancing

BIOS, change. With the release of Windows 2000 Advanced Server, Dell intends to support two-node, high-availability clusters with SCSI or Fibre Channel storage (as illustrated in Figure 2). Various combinations of servers, storage, and cluster interconnects will be thoroughly tested and certified.

## Windows 2000 Datacenter Server Multinode Cluster Service: Four Node

Windows 2000 Datacenter Server will include a four-node Cluster service, the Windows Multinode Cluster service similar to the two-node Cluster service in Windows 2000 Advanced Server. A failover cluster providing high availability for applications and data, the Multinode Cluster service will enable an application or service to failover from its server to any of three other servers.

To build a Multinode Cluster, four nodes connect to a common Fibre Channel, shared storage array (SCSI-attached storage is not supported), as shown in Figure 4. For application configurations, Multinode Clusters offer numerous options. Example configurations include all nodes running one or more applications (for example, Active/Active/Active/Active) and N-1 nodes running one or more applications and one node serving as the failover site for the N nodes (for example, Active/Active/Active/Passive).

### Applications

Windows 2000 Datacenter Server's Cluster service supports the same applications as the two-node service. Applications such as file sharing and Web serving can failover to any of the other cluster nodes.

For applications written to specifically address a primary node and a secondary node, such as SQL Server 7.0 Enterprise Edition and Exchange 5.5 Enterprise Edition, the application will not failover to the third and fourth server (if they are installed). The server (or servers) running SQL Server, however, can be the first, second, or third failover location for another server's application. SQL Server 2000 and Exchange 2000 will address this requirement by permitting up to four nodes to be active and available for the failover of another node.

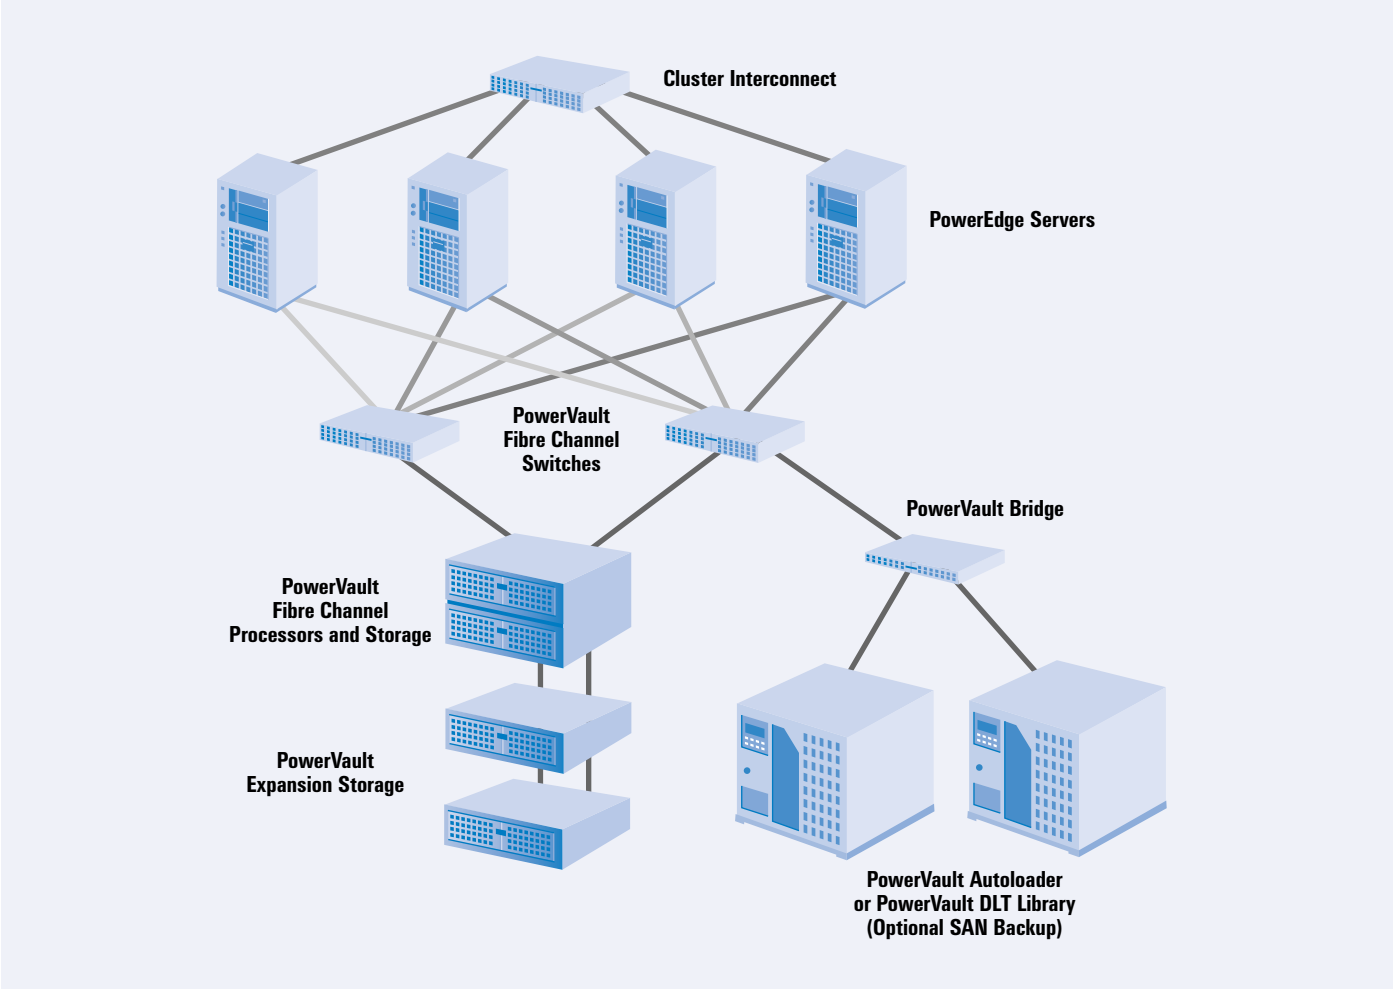


Figure 4. Possible Multinode Cluster Service Four-Node Configuration

### Implementation

Microsoft will require a separate, even more extensive certification process for Windows 2000 Datacenter Server four-node clusters (the details of this certification process were not yet final when this article went to press). With the release of Windows 2000 Datacenter Server, Dell intends to support four-node, high-availability clusters with Fibre Channel storage (Figure 4).

## Windows 2000 Network Load Balancing

In 1999, Microsoft purchased Valence Research and its Convoy Cluster product. A few months later, Microsoft repackaged the product and released it as a free download for licensed users of Windows NT Server 4.0 Enterprise Edition. The Windows NT Load Balancing Service (WLBS) is a load-balancing cluster capable of distributing IP traffic to a group of servers. It is a highly scalable cluster supporting up to 32 servers.

Renamed Network Load Balancing for Windows 2000, NLB acts as a front-end cluster for clients, distributing incoming IP traffic across all members. This fully distributed, software-based load balancing solution for IP is similar to hardware-based products from Cisco® and F5. It enables incremental scalability and outstanding availability for e-commerce Web sites. Clients see a single-system image, although their requests are distributed among up to 32 servers.

According to Microsoft, "When a computer fails or goes off-line for maintenance, Network Load Balancing

## CLUSTER BASICS

A cluster can be loosely defined as two or more systems working together to function as a single system. The goals of clustering include providing high availability, scalability, load balancing, and manageability for data, services, and applications. Different cluster offerings support different types of applications and offer various feature sets and capabilities.

**High Availability.** From a client's perspective, an application or resource should always be available to service requests. An application or system with a high percentage of uptime is said to achieve high availability.

High availability can be achieved with redundant hardware, load balancing, and software mechanisms that can stop, start, restart, or move applications if they fail. High-availability (HA) clusters may achieve some level of load balancing, but many times the load is distributed by manually moving resources from node to node.

**Failover.** One common implementation of high availability is in failover clusters. These are also called HA clusters because the primary objective of a failover cluster is to ensure high availability for an application or data while maintaining data integrity.

Failover clusters are not always synonymous with fault-tolerant servers. The term fault-tolerant usually implies that the client never observes even the slightest interruption in the application in the event of a component failure. With an HA cluster, if a component fails, services are restarted on another node. While the services are being started, the applications and data are unavailable to clients.

HA clusters usually boast uptime rates of 99 percent or more. Fault-tolerant systems often achieve uptime rates of 99.999 percent or higher.

**Scalable.** Scalability refers to a cluster's ability to accommodate expansion to achieve the desired level of performance. When a single server is overloaded, the problem can be resolved in two basic ways: use a larger server, or add servers and segment the load between them.

Within this article, scalability is defined as the ability to increase the performance capacity of the total cluster by adding resources (more nodes) to the cluster. Depending on the configuration, some nodes of the cluster may have a greater workload than others.

**Load Balancing.** Clusters that can distribute requests to the same application among multiple independent servers are load balancing. The term load balancing applies to clusters that include some number of nodes processing requests for the same type of application.

The load-balancing mechanism can vary greatly. In some cases, the clusters have no control over how requests are distributed. More intelligent load-balancing clusters can divide requests by response time, type of request, resource utilization on each server, and other defined metrics.

**Parallel.** Clusters that have parallel abilities are similar to load-balancing clusters because they do balance loads in a sense. The difference is that instead of working independently on the same type of application, as load-balancing clusters do, parallel cluster nodes work in tandem to complete a single request. For example, parallel computational clusters can take a request, divide the work among the cluster nodes, reassemble the results, and present them to the client as if a single supercomputer did the work.

**Single System Image.** To the client, a clustered resource appears as a single virtual server. End users do not care about connections as long as they can receive their e-mail, browse Web pages, run a database query, and perform other operations. Administrators, on the other hand, need to see and control each individual hardware and software component, even if they administer some aspects of the cluster as if it were a single entity.

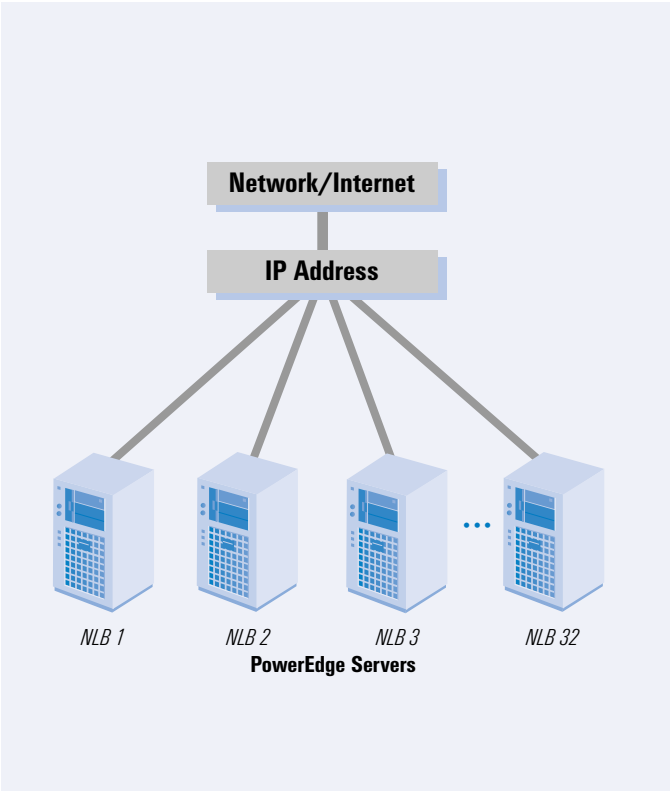


Figure 5. Possible NLB Load Balancing Cluster Configuration

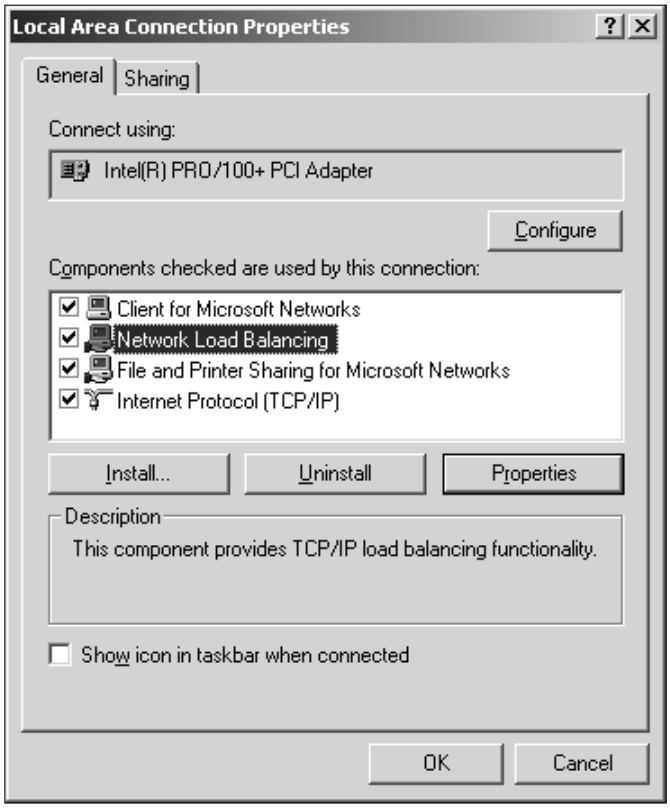


Figure 6. Windows 2000 Adapter Properties Page

automatically reconfigures the cluster to direct client requests to other servers, thereby maintaining continuous availability of network services." See Figure 5.

#### Enhancements over Windows NT

NLB for Windows 2000 is easier to set up and offers better integration with the operating system than WLBS (for Windows NT Server 4.0 Enterprise Edition). WLBS had to be downloaded and installed separately and then configured; only then could its network bindings be modified.

With Windows 2000 Advanced Server and Windows 2000 Datacenter Server, NLB software is always installed and appears as an option on the Network Properties tab. By default, the option is disabled, but NLB can be easily enabled from the network properties dialog (see Figure 6). The bindings configuration step has been removed, and the online help is now integrated with the help for the OS and other cluster offerings.

#### Application Support: NLB

NLB is most appropriate for applications in which the data is static and can be mirrored on other servers (static Web page servers, read-only file sharing, video streaming, and proxy servers) to load balance the incoming requests. NLB can also be used for multitier, multiple-cluster configurations in which the clients access a mission-critical application through a scalable Web front end that uses Cluster service on the back end to protect the dynamic, transactional data.

Microsoft intends to support Terminal Services (TS) in Windows 2000 using NLB for TS server load balancing under Windows 2000. Note that NLB is intended for IP load balancing and is not appropriate for applications like Exchange or SQL Server. NLB is complementary to Windows 2000 Cluster service and should not be run on nodes running Cluster service. See Figure 3 for a comparison of applications for Windows Cluster service and NLB.

#### Implementation

NLB does not require specialized hardware or separate Microsoft certification. Dell's departmental and workgroup servers (such as the PowerEdge® 1300, 2400, 2450, and 4400) make excellent choices for Web server farms. For applications that require more horsepower and can utilize NLB, such as Terminal Services, more powerful, rack-dense servers such as the PowerEdge 6350 are ideal.

### Windows 2000 COM+ Component Load Balancing

Component Load Balancing (CLB) is a dynamic load-balancing method for Component Object Model (COM)+ application components. This service, capable of scaling to

16 nodes, is integrated in and managed by Windows Application Center 2000.

In addition to enabling distributed load balancing for COM+ applications, AppCenter 2000 can help administer and deploy multitier Web configurations using IP load-balancing software, such as NLB, as shown in Figure 7 with CLB. Microsoft moved CLB, originally slated for all versions of the Windows 2000 Server family, to AppCenter to offer a consolidated management interface and better load-balancing capabilities.

CLB helps to load balance applications developed using Microsoft's COM technologies. A COM+ application component is installed on the cluster members through the management console, which replicates the components to the necessary servers.

### Application Support

CLB is intended for applications developed using the COM+ architecture. The COM+ architecture is an extension of the COM/DCOM model and technologies such as Distributed Transaction Coordinator, Transaction Server, and Message Queuing Server. CLB is not a load-balancing mechanism for applications like Microsoft SQL Server and Oracle.

For example, a COM+ application could be developed to perform order validation for an e-commerce Web site. When a user places a new order, the data collected from the Web site is passed to one of the load-balanced CLB servers to process the COM+ application. The COM+ application fulfills its responsibilities, possibly using data stored in a separate database server. It may either return an error message to the Web site (if an item the customer ordered is out of stock or the customer's credit card is over the limit) or pass the information to the order database and return a success notification to the customer.

### Implementation

CLB does not require specialized hardware. (Microsoft had not made a final decision about certification details as this article went to press.) Because each CLB configuration is application specific, the number of possible hardware configurations may be almost unlimited. In general, these systems will require substantial processing power because much of the business logic of the typical e-commerce applications and Web farms will be running on these systems.

## Cluster Solutions for Enterprise-Class Applications

With Windows 2000, Microsoft has expanded its cluster options to achieve enhanced availability and scalability. Windows 2000 clustering components—Cluster service, NLB, and CLB—can create highly available, scalable environments for many types of applications.

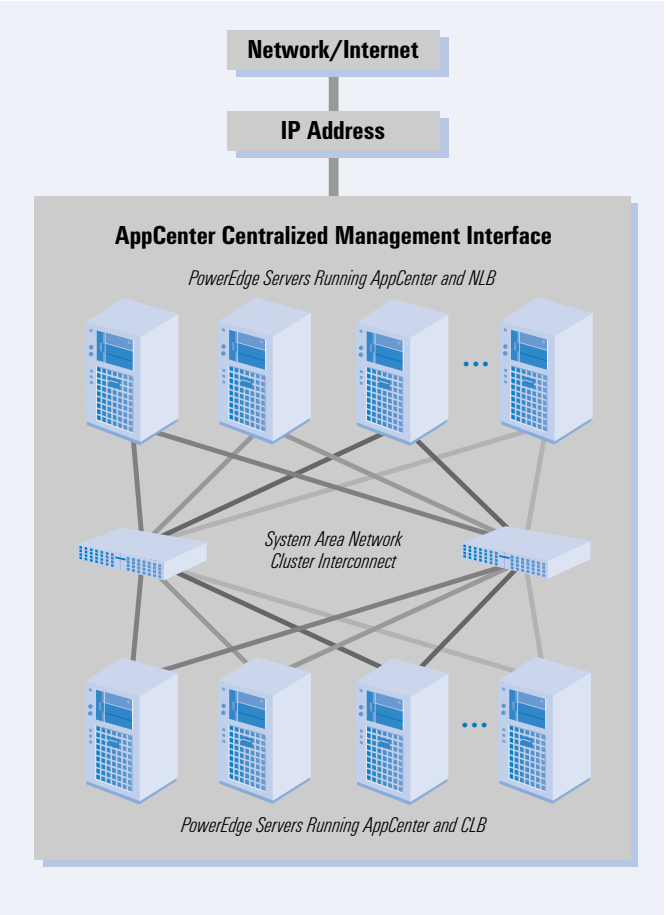


Figure 7. Web Farm Using NLB and CLB Clustering Managed, Monitored, and Replicated by AppCenter 2000

When used with Dell PowerEdge servers, Dell PowerVault™ storage, a high-speed, low-latency interconnect (such as the Giganet® cLAN™ adapter), and a management solution like Dell OpenManage™ Cluster Assistant with ClusterX™, these components can be used to build cluster solutions that meet the scalability and availability requirements of high-end, enterprise-class applications.◆

***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 A.A.S. 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.*