Examining Web Serving

on the PowerEdge 1655MC

The Dell™ PowerEdge™ 1655MC server blade system running Intel® Pentium® III processors provides a cost-effective and flexible platform for horizontally scaled, midtier applications. Dell tested the basic functionality of the various PowerEdge 1655MC components as well as the management software that ships with each system. This article also explains how the Dell team installed and optimized a Web-serving application on the PowerEdge 1655MC to examine its scalability and availability as a Web server.

BY DAVE JAFFE, PH.D.

The Dell™ PowerEdge™ 1655MC server blade system provides a flexible, scalable Web platform that is cost- and space-efficient, well suited for horizontally scaled, midtier applications. The modular construction of the 1655MC shares power and cooling among multiple servers, greatly reducing the cabling required for power; Ethernet; and keyboard, video, mouse (KVM) switches. To demonstrate the flexibility and scalability of the PowerEdge 1655MC as a Web platform, the Dell Product Group tested the PowerEdge 1655MC in the Dell Technology Showcase lab. This article describes the test environment and results, and details the process of creating and optimizing a Web-serving application on the PowerEdge 1655MC.

Setting up the PowerEdge 1655MC environment

To test the PowerEdge 1655MC server blade system in the Technology Showcase lab, the Dell team built a network around the server (see Figure 1). The PowerEdge 1655MC enclosure was equipped with three server blades (labeled blade11, blade12, and blade13); two integrated managed Layer 2 Gigabit Ethernet[1] switches; a local KVM switch; and an Embedded Remote Access/MC (ERA/MC) controller. Each blade contained two Intel® Pentium® III processors at 1.26 GHz, 2 GB of memory, two 36 GB disks, and two Gigabit Ethernet embedded network interface cards (NICs).

A Dell PowerEdge 1650 server, called Blademaster, served as a management console for various systems management functions, as well as an image server for capturing and deploying blade operating system images. A Dell PowerVault™ 775N network attached storage (NAS) server with 205 GB of internal disk capacity was added to store Web content. The Dell team also added two PowerEdge 1650 servers, Bladework and Bladework2, as workload generators.

The network was built around the Dell PowerConnect™ 5224 managed Layer 2 Gigabit Ethernet switch. Ports 1 through 8 of the switch were assigned to virtual LAN (VLAN) 50, and ports 9 through 16 to VLAN 60. Thus the switch was virtually configured as two separate switches, each handling a separate network segment. VLAN 50 handled the 192.168.50.x subnet and VLAN 60 handled the 192.168.60.x subnet. As shown in Figure 1, the first

[1] Gigabit Ethernet indicates compliance with IEEE® 802.3ab and does not connote speeds of 1 Gbps.

embedded NIC of each blade connected through the PowerEdge 1655MC switch 1 to the 50 subnet, while the second embedded NIC connected through the PowerEdge 1655MC switch 2 to the 60 subnet. Each support server was similarly multihomed. The ERA/MC was attached directly to port 11 of the PowerEdge 1655MC integrated switch 1 inside the enclosure.

A Domain Name System (DNS) server ran on the Blademaster system. The PowerEdge 1655MC server blades and auxiliary systems were assigned names within the dellblades.com domain.

Installing the operating system and applications

The PowerEdge 1655MC server blades may be ordered with the Microsoft® Windows® 2000 or Red Hat® Linux® operating system pre-installed. Alternatively, administrators can install the operating system using Dell OpenManage™ Server Assistant. Installation of either Windows 2000 or Red Hat Linux is very similar with Server Assistant. In both cases, Server Assistant loads the necessary drivers, gathers the required configuration information (such as host name and IP addresses), and prompts for the operating system CD when needed. Currently, Windows 2000 Server, Windows 2000 Advanced Server, and Red Hat Linux 7.3 are supported.

Mirroring disks in a RAID-1 configuration

If the blade has two disks, they can be mirrored in a RAID-1 configuration for fault tolerance (the recommended configuration) using the PowerEdge Expandable RAID Controller (PERC 4/im) embedded on each blade. Four methods exist for mirroring the disks:

- Specify the RAID-1 configuration when purchasing the blade
- Create the mirror using the Array Manager feature of Dell OpenManage Server Assistant
- Create the mirror using the PERC software directly (press Ctrl + M during blade boot)
- Deploy the RAID configuration from the Dell OpenManage Remote Install image server (see the “Capturing and deploying the blade system image” section of this article)

Administrators should perform any disk mirroring before installing the operating system and applications.

Installing the applications and Web content

Upon completion of the operating system installation, Server Assistant prompts the user to insert the Dell OpenManage Systems Management CD to install the OpenManage Server Administrator systems management software and the Remote Install agent. The system is now ready for installing and configuring the Web server.

The Dell test used two different Web servers to serve static content: Microsoft Internet Information Services (IIS) 5.0 for blade11, which ran Windows 2000 Advanced Server; and Apache HTTP

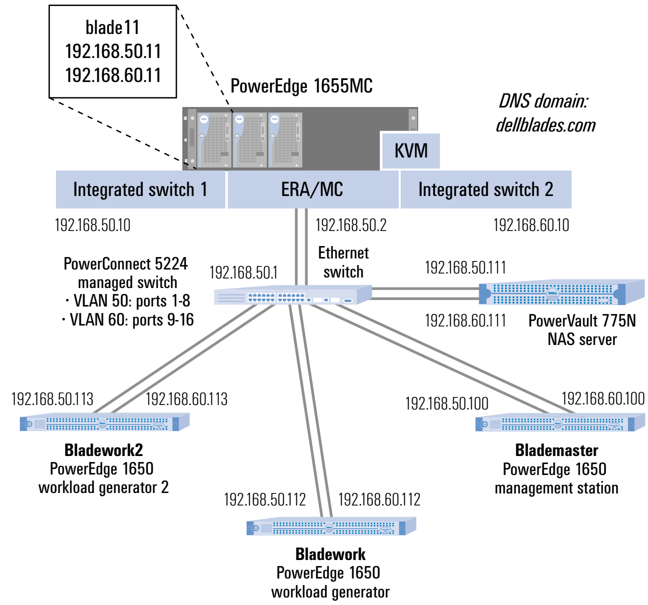

Figure 1. Technology Showcase modular computing test environment

server, version 1.3.23, for blade13, which ran Red Hat Linux 7.3. The Dell team copied approximately 60 MB of HTML pages, Microsoft Word documents, PDF files, and images from http://www.dell.com/servers to C:\inetpub\wwwroot\dell on the Windows-based system. The IIS remote administration tool, Internet Services Manager, was used to create a new virtual directory with the alias *dell* and the C:\inetpub\wwwroot\dell pathname. On the Linux-based system, the team copied the same content to /var/www/html/dell, the default Apache document root. Thus, users browsing http://blade11.dellblades.com/dell (the Windows-based blade) or http://blade13.dellblades.com/dell (the Linux-based blade) saw the same Web pages.

Storing content on the NAS server. To avoid the need to replicate Web content between server blades, the content should be kept on a NAS device. The PowerVault 775N is well suited for use in a mixed-platform environment because it can serve files to a Windows-based server through the Common Internet File System (CIFS) or to a UNIX® or Linux platform through Network File System (NFS).

The Dell PowerEdge 1655MC server blade system provides a flexible, scalable Web platform that is cost- and space-efficient, well suited for horizontally scaled, midtier applications.

To demonstrate this alternative approach, the Dell team created a share on the PowerVault 775N called *data* that contained the same *dell* directory of server Web site files as the one created on the

```
LABEL=/               /                          ext3       defaults                  1 1
LABEL=/boot           /boot                      ext3       defaults                  1 2
none                  /dev/pts                   devpts     gid=5,mode=620            0 0
LABEL=/home           /home                      ext3       defaults                  1 2
none                  /proc                      proc       defaults                  0 0
none                  /dev/shm                   tmpfs      defaults                  0 0
LABEL=/tmp            /tmp                       ext3       defaults                  1 2
LABEL=/usr            /usr                       ext3       defaults                  1 2
LABEL=/var            /var                       ext3       defaults                  1 2
/dev/sda9             swap                       swap       defaults                  0 0
PV775N1:/data/dell    /var/www/html/data/dell    nfs        defaults                  0 0
/dev/cdrom            /mnt/cdrom                 iso9660    noauto,owner,kudzu,ro     0 0
```

Figure 2. Sample Red Hat Linux /etc/fstab file

blades. To access the NAS files under IIS, the testing team used Internet Services Manager to create a new virtual directory, *data*, and selected the "A share located on another computer" option. Under Red Hat Linux, testers first created a directory stub, /var/www/html/data. Then the *data* share was NFS-mounted over the stub. After these steps, the content accessed at http://blade11.dellblades.com/data/dell and http://blade13.dellblades.com/data/dell was exactly the same content. To make the NFS mount permanent, administrators can add a line to the /etc/fstab file (see Figure 2).

Capturing and deploying the blade system image

The PowerEdge 1655MC server blades are well suited for applications that scale horizontally, such as Web serving. Horizontal scaling lets administrators accommodate a bigger workload by adding more servers of the same type, usually in a load-balanced cluster with failover capability. To facilitate the addition of new blades to run the same application, the PowerEdge 1655MC comes with Dell OpenManage Remote Install software, which enables the capture of blade system images and deployment to new blades.

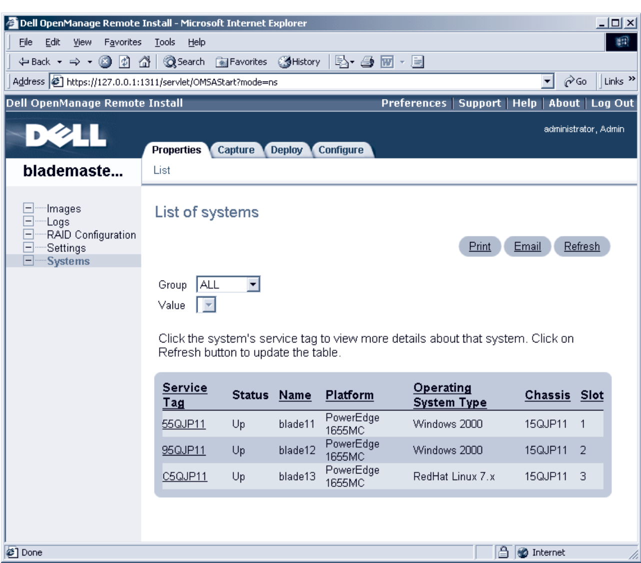

Figure 3. The Systems Properties tab of OpenManage Remote Install

The Remote Install software, which the Dell team installed on the Blademaster system, provides a central repository of system images (an image server) as well as a management station for all blades accessible on the network. The main screen of the Web-based Remote Install user interface shows the service tag (an ID number unique to each blade), name, status, operating system, and slot of each blade under the software's management (see Figure 3). To use OpenManage Remote Install to manage a blade, at least one Broadcom® NIC must be ahead of the hard drive in the blade's boot order (see sidebar, "Changing the boot order").

After configuring blade11, the testing team used the Remote Install user interface to capture the blade's Windows-based Web server image and to deploy the image to blade12. To deploy an image to a blade, administrators must specify which blade and image to use, and they must customize a configuration file with the correct host name, IP addresses, and other information. In addition, the configuration file enables users to enter a command line (which may call a script) that will be executed once the deployed blade boots. Administrators can use the command line to add configuration details beyond those allowed by the configuration file. A useful command line to add is:

```
net time \\192.168.50.100 /set /yes
```

This command line synchronizes the clock in the blade with that of Blademaster, the management station.

Administrators can predeploy a blade by specifying in Remote Install the blade's service tag, which Dell will supply before shipping, along with the image and configuration information. Then, when the blade arrives, administrators simply slide it into the enclosure and power it up. The blade will boot using the Preboot Execution Environment (PXE) from the image server, which will check the service tag and deploy and configure the waiting image in approximately 15 minutes.

In addition to the system images, Remote Install also manages a set of RAID configuration packages that can be deployed much like operating system images. Thus, if the image administrators wish to deploy is based on a mirrored or non-mirrored disk, they can apply the CreateMirror or DeleteMirror configuration as appropriate.

Clustering blades using NLB

To provide redundancy and load balancing, administrators should cluster the blades. The team chose Network Load Balancing (NLB), a component of Microsoft Windows 2000 Advanced Server, as the clustering software. NLB provides load balancing among the cluster nodes (with either an equal or weighted distribution), failover from a failing blade to the other blades, and remote control of the cluster.

In the lab network, the team configured NLB on the second NIC of blade11 using the Network and Dial-up Connections applet. The configuration parameters fall into three categories: cluster parameters, host parameters, and port rules. Figure 4 lists the parameters that were changed from the defaults. The primary IP address for each node in the cluster (192.168.60.200) was specified under NLB cluster parameters as well as in the IP properties for the second NIC; an entry was added to the DNS server to map the host name blades.dellblades.com to this IP address. Of these, only the Priority and the Dedicated IP address parameters are host-specific.

CHANGING THE BOOT ORDER

The default boot order of the blade puts the hard disk ahead of the two embedded NICs, so that the system will boot from the hard disk once the operating system is installed. To use OpenManage Remote Install, administrators must move the hard disk after the NICs in the boot order. This way, the blade can use the Preboot Execution Environment (PXE) mechanism to communicate with the Remote Install image server. The boot order can be changed by running the Setup program, accessible by pressing F2 during boot. The use of both NICs for PXE provides fault tolerance but can lead to extra delays, because both NICs must communicate with the PXE server. The second NIC should be removed from the boot order for development systems.

To facilitate the addition of new blades to run the same application, the PowerEdge 1655MC comes with Dell OpenManage Remote Install software.

Testing Web workloads

Once NLB was configured on blade11, the image was captured and stored on the image server. The testing team then ran a Web workload against the single blade running as a single-instance NLB cluster until the blade's CPU utilization exceeded 80 percent. At this point, the team added another Web server. They installed a blade into the second slot and deployed the NLB image to it, using the Remote Install configuration file to set its host name to blade12 and its IP address to 192.168.50.12. Additionally, the team pasted a Microsoft Visual Basic® script (see Figure 5) into the configuration file's script box to set the NLB Priority (Unique host ID) to 2 and its dedicated IP address to 192.168.50.12.

After booting the second blade, the testing team remotely started NLB on it. Without stopping the first blade, the second blade immediately ramped up and handled half the Web requests. After a few minutes, the CPU utilization of both blades was approximately 40 percent.

Aggregating switch ports to increase bandwidth

The integrated PowerEdge 1655MC switches are fully managed Layer 2 switches.[2] In this study, the Dell team performed an experiment to demonstrate the improved throughput available by trunking two or more of the switch uplink ports together.

Cluster parameters	
Primary IP address	192.168.60.200
Remote control	Enabled
Host parameters	
Dedicated IP address	192.168.50.11
Priority (Unique host ID)	1
Initial cluster state	Unselected (cluster will not start at operating system boot)
Port rules	
Affinity	None (equal host weighting)

Figure 4. Network Load Balancing parameters

[2] For information about the switches used in different network configurations, see "Examining Network Performance of the PowerEdge 1655MC Blade Server" by Yinglin Yang and Mike J. Roberts in Dell *Power Solutions*, November 2002.

```
' setnlb2.vbs - Set NLB Info
' Dave Jaffe Dell 11/30/02
' based on code from MS TechNet Script Center

Set objWMIService = GetObject("winmgmts:\\.\root\MicrosoftNLB")
Set objNodes = objWMIService.InstancesOf("MicrosoftNLB_NodeSetting")

For Each objNode in objNodes
 Exit For
Next

objNode.DedicatedIPAddress = "192.168.50.12"
objNode.HostPriority = 2
objNode.Put_
```

Figure 5. Sample NLB configuration script

A Web workload was applied to the NLB cluster described in the previous section. As before, in each server the NLB cluster NIC was the second NIC, which connected to PowerEdge 1655MC integrated switch 2. A single uplink port of integrated switch 2 was connected through the PowerConnect 5224 switch to the simulated Web users running on the two workload generator servers, Bladework and Bladework2.

Many Web and other applications can be efficiently and robustly implemented on a set of horizontally scaled servers.

The workload—1,700 users requesting five static files—was ramped up until the 1 Gbps single port was saturated at 974 Mbps. The two Web servers delivered 47,009 Web pages per minute, with an average response time of 0.172 seconds. Then a second uplink port of the PowerEdge 1655MC integrated switch 2 was cabled to the PowerConnect 5224, and port aggregation (static trunking) was enabled on both sides using their respective Web browser-based interfaces. The two-port trunk eliminated the network bottleneck, permitting a total of 1031 Mbps throughput, which improved the Web performance by nearly 8 percent, up to 50,597 pages per minute, with a reduced response time of 0.021 seconds. Up to four ports of each integrated switch may be aggregated in this manner, enabling a maximum full-duplex throughput of 8 Gbps per switch.

Providing a flexible and scalable Web platform

Testing in the Dell Technology Showcase lab focused on examining the basic functionality of the various PowerEdge 1655MC components as well as the management software that ships with each system. Using a Web-serving application, the testing team was able to demonstrate the use of network attached storage for storing Web content, the use of the Microsoft NLB cluster software for scalability and failover, and the use of port aggregation to increase the bandwidth of the PowerEdge 1655MC.

Many Web and other applications can be efficiently and robustly implemented on a set of horizontally scaled servers. The PowerEdge 1655MC server, the initial entry from Dell into the modular computing arena, is well suited for such applications.

Acknowledgments

The author would like to thank Mike J. Roberts, the PowerEdge 1655MC marketing manager; members of the PowerEdge 1655MC development team (Scott Ramsey, Mike Kolb, Brian Peil, Jeremey Pionke, Jin Masuyama, Ron Salmon, and Julie Silver); the Remote Install team (Danny Bertram, Weijia Zhang, Syam Poluri, Hai Phung, and David Schmidt); Yinglin Yang and Mike Morrin of the networking team; software guru William Collins; testers Brian Doty and Matthew Moore; and Fred Gonce and Roger Foreman from the Technology Showcase Group.

Dave Jaffe, Ph.D. (dave_jaffe@dell.com) is a senior consultant in the Dell Technology Showcase Group, specializing in cross-platform solutions. Dave has a B.S. in Chemistry from Yale University and a Ph.D. in Chemistry from the University of California, San Diego.

FOR MORE INFORMATION

Dell PowerEdge 1655MC: http://www.dell.com/blades