DELL POWER SOLUTIONS

TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

MAY 2007 • $4.95

Access **Everywhere.** Business **Anywhere.**

How Microsoft Exchange Server 2007 Makes It Happen

Planning for Exchange Server 2007, Windows Vista, and 2007 Microsoft Office Deployments

Exploring Exchange Server 2007 Architectural Alternatives

Drafting a Blueprint for Exchange Server 2007 Storage

Access Everywhere. Business Anywhere. How Microsoft Exchange Server 2007 Makes It Happen


intel
Leap ahead™
Acceleration Ace.
It could be you.
Ready for faster data delivery?
Intel® I/O Acceleration Technology (Intel® I/OAT) moves network data more efficiently through Dual-Core and Quad-Core Intel® Xeon® processor-based servers.
Scalability across multiple Ethernet ports? Yes.
Integration with popular operating systems? Absolutely.
Sluggish throughput? No way.
Speed your server performance.
With Intel I/OAT, data flows free. And easy.
Learn more: intel.com/go/ioat
©2007, Intel Corporation. All rights reserved. Intel, the Intel logo, Intel. Leap ahead. and the Intel. Leap ahead. logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries. *Other names and brands may be claimed as the property of others.

DELL POWER SOLUTIONS

May 2007

TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

FEATURE SECTION: MICROSOFT EXCHANGE SERVER 2007

COVER STORY • page 10

Access **Everywhere.** Business **Anywhere.**

How Microsoft Exchange Server 2007 Makes It Happen

By Jeanne Feldkamp, Deb McDonald, and Tom Kolnowski

With intellectual capital flowing through messaging systems and an increasingly mobile workforce, enterprises face significant challenges surrounding data access, operational efficiency, security, and data protection. Microsoft Exchange Server 2007 addresses these concerns by providing a comprehensive, integrated, end-to-end messaging solution designed to reduce costs and complexity, enhance performance, and simplify compliance.

The Business of IT • page 18

Planning a Strategic Framework for Anywhere Data Access

By Chris Auger

As communication technologies converge in the workplace, IT decision makers must find new approaches to messaging infrastructure design, security, and compliance. By keeping the needs of highly mobile clients in mind during the Microsoft Exchange Server 2007 planning process, organizations can simplify operations, optimize solutions, and sustain value across the overall IT infrastructure.

SPECIAL TEAR-OUT POSTER • page 48a

Dell PowerEdge High-Availability Clusters

To help protect access to business-critical applications, Dell cluster solutions are designed to eliminate single points of failure in a wide array of environments. Integrating Dell PowerEdge servers, Dell PowerVault and Dell/EMC storage, and other key components, Dell high-availability cluster configurations can scale flexibly and cost-effectively to meet the management and performance needs of any size enterprise, offering a host of benefits:

- **Manageability:** Centralize administration and resource management
- **Scalability:** Expand to address resource requirements throughout the enterprise
- **Consolidation:** Host multiple applications to maximize utilization
- **Customization:** Choose the industry-standard components and configurations that best suit your needs

Now Playing: Dell Datacenter of the Future Podcast Series

How do companies like Dell manage their desktop environment and maintain standards? "Best Practices: Systems Management with SMS" explores how enterprises can use Microsoft Systems Management Server to inventory hardware and software, distribute software updates and patches, and meter use—all using an open API. For more information, visit www.dell.com/podcast.

EDITOR'S COMMENTS

FEATURE SECTION: MICROSOFT EXCHANGE SERVER 2007

SEE IT HERE FIRST

more Online

www.dell.com/**powersolutions**

Exploring the Dell PowerVault MD3000 Storage System

By Manjusha Gopakumar, Vishal Kadam, and Paul Winston

The Dell PowerVault MD3000 allows multiple servers to access a shared pool of expandable storage, helping avoid the costs and complexities of Fibre Channel or Ethernet storage infrastructures.

Installing Dell OpenManage Server Administrator on Linux Platforms

By O. R. Senthil Kumaran and Rajkumar V. J.

Administrators can choose from an array of local and remote deployment methods when installing Dell OpenManage Server Administrator 5.2 on supported Linux operating systems.

Using Dell OpenManage IT Assistant for Effective Systems Management

By Maheswarry Ramakrishnan and Rammya Krishnamurthy

Following best practices can help administrators efficiently utilize Dell OpenManage IT Assistant for simple, effective remote systems management.

Discovering Dell/EMC Storage Devices with Dell OpenManage IT Assistant 8.0

By Eranna Talur and Maheswarry Ramakrishnan

Dell OpenManage IT Assistant 8.0 helps administrators discover, inventory, monitor, and manage Dell/EMC storage devices.

Integrating Dell Remote Console Switches with Microsoft Active Directory

By Rayan Ghosal

The multiuser Dell 2161DS-2 and 4161DS remote console switches are designed to authorize and authenticate Active Directory users for switch access.

Advanced Configuration for the Dell PowerConnect 5316M Blade Server Chassis Switch

By Surendra Bhat and Saurabh Mallik

Enterprises can take advantage of advanced configuration options for the Dell PowerConnect 5316M Gigabit Ethernet switch to help optimize network access for blade servers.

Seamless Updates Using Dell OpenManage Server Update Utility 1.5

By Gagan Shrestha, Deepak Panambur, and Panna Hegde

Dell OpenManage Server Update Utility 1.5 includes several enhanced features designed to increase flexibility and simplify hardware and software updates for Dell PowerEdge servers.

Simplifying OS Deployment with the Dell OpenManage Server Assistant Driver Extraction Utility

By Pritesh Prabhu and Joseph Razo

The Dell OpenManage Server Assistant Driver Extraction Utility enables administrators to easily integrate drivers for Dell PowerEdge servers into custom OS deployment processes.

Your Information, Your Way

The *Dell Power Solutions* articles on this page are available exclusively online at **www.dell.com/powersolutions**. Plus: Check the *Dell Power Solutions* Web site for our early-release articles, how-to's, implementation studies, and expert tips you won't find anywhere else.

Need specific subject matter at your fingertips? Search our **Related Categories** index and then visit full-text articles from current and back issues online. Can't wait? For real-time access to content as soon as it's posted, subscribe your favorite RSS (really simple syndication) reader to the feed at the *Dell Power Solutions* Web site or **www.dell.com/rss**.

DEPARTMENTS

SCALABLE ENTERPRISE SERVICES SHOWCASE • page 40

Unlocking the Potential of Microsoft Exchange Server 2007

With turnkey solutions and proven expertise, the Dell Services team helps ease migration to Microsoft Exchange Server 2007—enabling businesses to gain a competitive edge by taking advantage of new features like Unified Messaging from data center to desktop and beyond.

It's time for you to lay off the juice.

The energy-efficient AMD Opteron™ processor.

AMD
Smarter Choice

www.amd.com/lessenergy

ADVERTISER INDEX

STOCK NUMBER: DPS 2007-05 (MAY 07)

DELL POWER SOLUTIONS TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

WWW.DELL.COM/**POWERSOLUTIONS**

TALK BACK

We welcome your questions, comments, and suggestions. Please send your feedback to the *Dell Power Solutions* editorial team at us_power_solutions@dell.com.

Remember when attacks occurred on the outside?
Today, more and more attacks are coming from within your network.

Do you have the right security solution?

The Viking V 2-Gigabit LAN Firewall

The Viking V 2-Gigabit LAN Firewall offers premium performance at a fraction of the price. It's the only firewall that performs at full Gigabit speeds in its price range – even when fully configured.

And independent tests confirm it!*

Traditional firewall use at the network's perimeter protects from outside attacks, but today more and more attacks are occurring from the inside. Is your LAN protected?

Learn more at: www.vikinginterworks.com/security

VIKING INTERWORKS™

Tel: +1 408 964 3947
Tel: 800 529 7710 (US and Canada only)
Fax: +1 408 964 3437
Email: security@vikinginterworks.com

*Up to Spec is The Tolly Group's flagship custom testing program for certification of Information Technology products, services and components. The Tolly Group offers uncompromised, independent evaluations.

Brave New Worlds

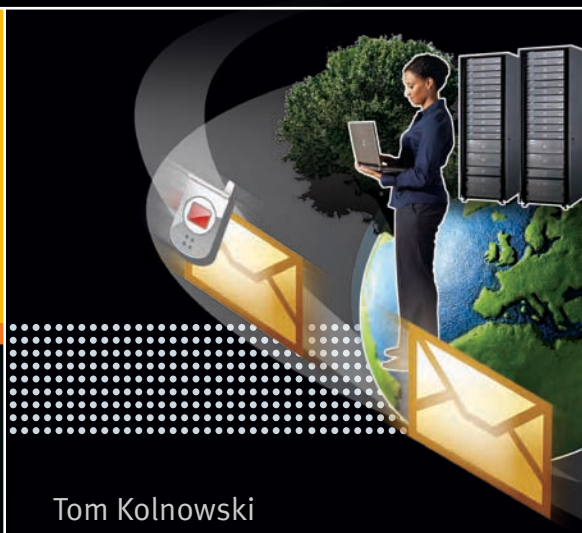

Tom Kolnowski
Editor-in-Chief and Publisher
tom_kolnowski@dell.com
www.dell.com/**powersolutions**
www.dell.com/rss

The concept of a server role—a logical grouping of features and components—has been introduced in the recently launched Microsoft® Exchange Server 2007 messaging system. In addition to the server roles that IT architects and administrators might normally expect to find in an enterprise messaging infrastructure—such as Client Access, Edge Transport, Hub Transport, and Mailbox—the Redmond development team has seamlessly integrated the new Unified Messaging role into the Exchange Server 2007 management console. And the implications of unified messaging are blurring the lines between corporate e-mail and traditionally separate, proprietary telephony systems, which may have a significant impact on businesses.

Is unified messaging another technological paradigm shift in the making? We don't have a crystal ball, but you can read about the world of Exchange Server 2007, unified messaging, and anywhere access for an increasingly mobile workforce starting on page 10. Our cover story, "Access Everywhere. Business Anywhere. How Microsoft Exchange Server 2007 Makes It Happen," leads our extensive lineup—30 pages in all of late-breaking business and technical information about this significant Exchange Server upgrade, plus insight into why you should also be considering upgrades to the Microsoft Windows Vista™ OS and 2007 Microsoft Office.

On another development front, the *Dell Power Solutions* editorial and art teams have been very busy the past six months reinventing our presentation layer—the design of the magazine. In an extensive design effort led by our new art director, David Chan, every element of the magazine—ranging from the front cover to the back and every pixel in between—has been rethought, re-sketched, refined, and redesigned to visually heighten the viewing experience while improving readability. We hope you enjoy our use of a fresh, eye-easy font family and text artfully flowed into new two- and three-column layouts that are enveloped by bold colors and high-tech dot-matrix creative touches throughout.

For content beyond the borders of the print edition, look for the new "more online" logo. This indicates there is additional value-added content at our home base of www.dell.com/powersolutions or elsewhere on the Web. While you are online, we hope you will try our new granular RSS feeds at www.dell.com/rss, where you can pick from 20 different categories of lab-fresh *Dell Power Solutions* content.

The blood pressure of the IT Director is connected to the corrupted Exchange database
is connected to the CEO's email
is connected to Symantec Backup Exec 11d which
recovered the data in a heartbeat.

Prepare for the worst with the best. Introducing Backup Exec™ 11d from Symantec. The worldwide leader in backup and recovery. Now with continuous protection for Exchange. Mailbox backups are history. Messages, folders, and emails can be recovered in seconds. No wonder Backup Exec is still the gold standard in Windows® data recovery. So for the optimal health of your IT infrastructure (and your IT Director) look into Backup Exec 11d. Visit symantec.com/BE11d

Confidence in a connected world.

symantec™

Access **Everywhere.** Business **Anywhere.**

How Microsoft Exchange Server 2007 Makes It Happen

With intellectual capital flowing through messaging systems and an increasingly mobile workforce, enterprises face significant challenges surrounding data access, operational efficiency, security, and data protection. Microsoft® Exchange Server 2007 addresses these concerns by providing a comprehensive, integrated, end-to-end messaging solution designed to reduce costs and complexity, enhance performance, and simplify compliance.

Related Categories:

2007 Microsoft Office
Dell PowerEdge servers
Dell/EMC storage
Enterprise management
Microsoft Exchange Server 2007
Microsoft Office Outlook 2007
Microsoft Windows Server 2003
Microsoft Windows Vista

Visit www.dell.com/powersolutions for the complete category index.

BY JEANNE FELDKAMP
DEB MCDONALD
TOM KOLNOWSKI

If corporations are complex, living bodies, e-mail is the blood pulsing through their veins. It circulates vital information where needed and helps keep operations running smoothly. Without e-mail, business can grind to a halt.

E-mail is now the primary means of communication within the enterprise—and it has become much more than a simple messaging tool. Employee e-mail messages contain an enormous volume of critical data that must be managed, protected, and accessed. As the volume of messaging-related data grows exponentially, enterprises are facing significant data storage and security challenges. Spam, viruses, and malware present major threats to 24/7 availability. Compliance has also been added to the mix—enterprises now must be able to pull information out of archives easily and securely to help satisfy governmental regulations.

For these reasons, organizations must protect themselves effectively against threats, manage the increasing volume and cost of e-mail storage, meet compliance and risk-management requirements, and build a resilient e-mail infrastructure. Microsoft Exchange Server 2007 helps enterprises accomplish these imperatives with anywhere access, built-in protection, and outstanding operational efficiency.

Expanding the definition of messaging

Ask enterprise users today how important e-mail is, and they are likely to say that being without e-mail would present a greater hardship than losing telephone service. But the reality is that enterprise users need *both* the telephone and e-mail to do their jobs—and increasingly, they need other forms of messaging as well. Employees need rapid, easy access not only to e-mail, but also to voice mail, instant messaging, and fax communications. In addition, they need to access these messages using a variety of tools, including mobile devices, computers connected through wireless networks, and cell phones as well as conventional telephones.

Understanding the challenges of a mobile workforce

As the definition of messaging within the enterprise has expanded, so too have the demands on IT infrastructures. An increasingly mobile workforce creates challenges in terms of access control, security, and bandwidth. Operational efficiency can suffer as current architectures fail to keep pace with the increasing volume of users and data. The threat of viruses, worms, and spyware looms large. And because significant portions of enterprise intellectual capital reside within e-mail, those messages must be backed up and protected against tampering in the same manner as other critical files.

"Enterprises need more than a quick fix to support heterogeneous messaging systems. They need a comprehensive, end-to-end solution—with all the elements tested and validated to work together."

Traditional messaging architectures lack the level of control needed to handle exponential growth in e-mail volume and proliferation of message sources. It is not uncommon for users to require 10 MB of storage per day. This explosion in usage and volume is testing the performance limits of many enterprise IT environments.

Limitations of 32-bit processors put a lid on the number of users that e-mail systems can handle as well as the size of individual mailboxes. IT administrators typically try to add 32-bit servers to increase capacity, but this tactic also adds complexity to the IT infrastructure. Eventually, the system approaches a tipping point where administrators simply cannot squeeze additional performance out of the existing architecture. The organization cannot take advantage of the increased memory, processing power, and I/O speeds enabled by 64-bit systems without replacing its entire messaging infrastructure.

Besides creating storage challenges, messaging growth has produced serious IT management issues. Historically, e-mail management has been performed manually. Increased volumes lead to rising costs and additional administrators—and many IT teams lack the tools and best practices needed to make messaging management easy and cost-effective.

Security is also a major issue in traditional messaging systems. Data is under constant threat of attack from both outside and inside an organization, which means that protection has become a 24/7 necessity. In recent years, there has been a huge increase in spam, viruses, and malware that can harm or even eliminate critical data. According to a recent computer crime and security publication, 65 percent of surveyed organizations reported virus attacks, and 52 percent of those respondents also reported unauthorized use of computer systems.[1] Spam, which is a frequent carrier of viruses and malware, accounts for a substantial amount of e-mail volume and is still trending upward. In addition, many viruses typically enter an organization through the e-mail gateway—so as the number of applications and devices accessing the e-mail infrastructure increases, security considerations become more pressing than ever.

Compliance regulations such as the Sarbanes-Oxley Act (SOX) and the Health Insurance Portability and Accountability Act (HIPAA) further complicate e-mail management within the enterprise. These laws require quick and accurate retrieval of e-mail messages to meet standards, yet many archived messages are still locked in vaults on tape. The shift to easily searchable near-line storage systems can add cost and complexity in managing the e-mail life cycle.

Despite the potential costs, data protection is another reason enterprises are rethinking traditional messaging infrastructures. Invaluable intellectual property is commonly contained within e-mail messages, and organizations involved in an electronic data discovery request are expected to produce messages as part of legal or regulatory proceedings. Data recovery from crashes and disasters is often difficult and must be performed manually to help ensure that everything from e-mail messages to attachments to shared files is restored.

Rethinking the messaging paradigm

Unfortunately, no easy fix exists to address the problems associated with heterogeneous messaging infrastructures. Many IT managers believe they can cobble together solutions that will help them keep up with performance demands while using the same IT infrastructure. For example, some enterprises try to decouple servers from storage or try to focus on reducing costs on the storage side. However, this approach does not solve the root problem—it only adds additional storage to a flawed infrastructure.

Enterprises need more than a quick fix to support heterogeneous messaging systems. They need a comprehensive, integrated, end-to-end messaging solution. In an integrated solution, all the elements (including servers, storage, tools, services, and support) have been tested and validated to work together. Messaging infrastructures must not only integrate messages from multiple inputs, but also provide a platform that helps reduce costs and complexity, increase performance, improve scalability, and simplify compliance. For information about end-to-end messaging solutions from Dell, see the "Comprehensive Dell Solution for Microsoft Exchange Server 2007" sidebar in this article.

By implementing a cohesive messaging system, enterprises help ensure that all of the components work and will be supported. Standards-based components enable the messaging infrastructure to interoperate with existing systems, helping ensure that the messaging environment can be expanded as needed. The result: a solution that addresses the e-mail issues, yet is ready for an end-to-end communications infrastructure that also includes voice mail, instant messaging, fax, Microsoft Office applications, and the Microsoft Windows Vista™ OS (see Figure 1).

Microsoft Exchange Server 2007 introduces a wide array of new features and capabilities specifically designed to help enterprises address the challenges of unified messaging. It provides administrators with powerful new tools to help them do their jobs and gives end

[1] "2006 CSI/FBI Computer Crime and Security Survey," by the Computer Security Institute, i.cmpnet.com/gocsi/db_area/pdfs/fbi/FBI2006.pdf.

Guided Tour: Enabling Microsoft Exchange Server 2007 for Anywhere Access

Exchange Server 2007 introduces unified messaging capabilities that broaden the scope of business as usual while dramatically enriching the user experience. From the Microsoft Office Outlook 2007 desktop client, a MAPI (Messaging Application Programming Interface)–enabled corporate mainstay, to remote and mobile user end points through Outlook Web Access, an AJAX (Asynchronous JavaScript and XML)–based Web application—Exchange Server 2007 provides an integrated, end-to-end messaging solution.

In addition, Exchange Server 2007 makes it easy for administrators to enable and manage these unified messaging capabilities throughout their organizations. High-level views of some of the new administration tools for enabling these capabilities are shown here.

- **Outlook Anywhere:** New in Exchange Server 2007, Outlook Anywhere (see Figure A) was formerly known as Remote Procedure Call (RPC) over HTTP in Exchange Server 2003. It enables Internet access to Exchange mailboxes through the full Outlook 2007 or Outlook 2003 desktop client—without the need for a virtual private network connection.
- **Unified Messaging:** When Unified Messaging is enabled in Exchange Server 2007 (see Figure B), administrators can allow client access to the new telephony capabilities, including integrated voice mail, faxes, and more.
- **Outlook Web Access:** Web-based messaging through Outlook Web Access (see Figure C) has been enhanced to provide full access to Unified Messaging capabilities, while Outlook Web Access Light has been added for efficient access over low-bandwidth connections.

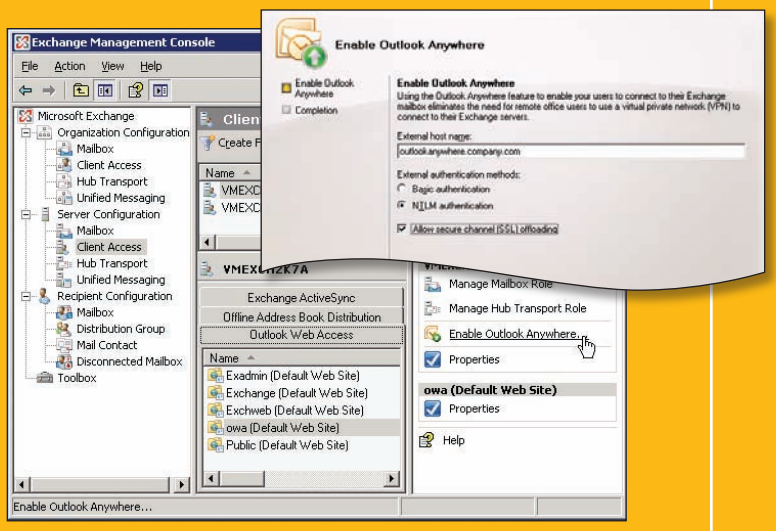

Figure A. *Enabling Outlook Anywhere capabilities through the Exchange Management Console*

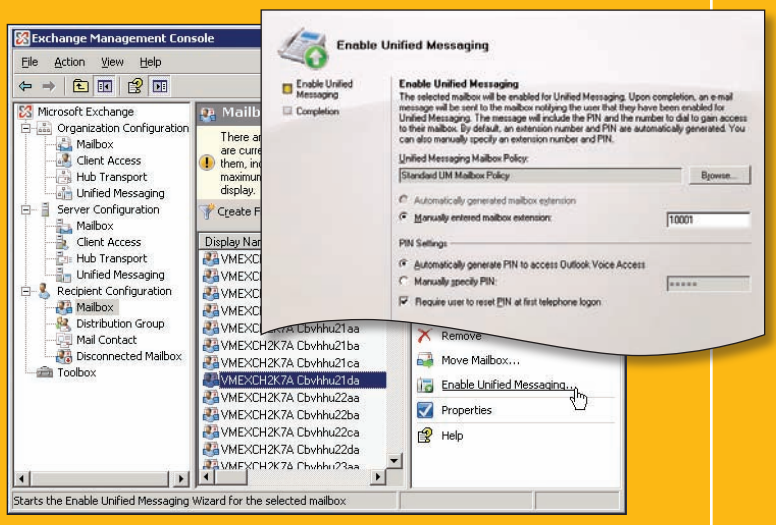

Figure B. *Enabling Unified Messaging capabilities for an individual mailbox*

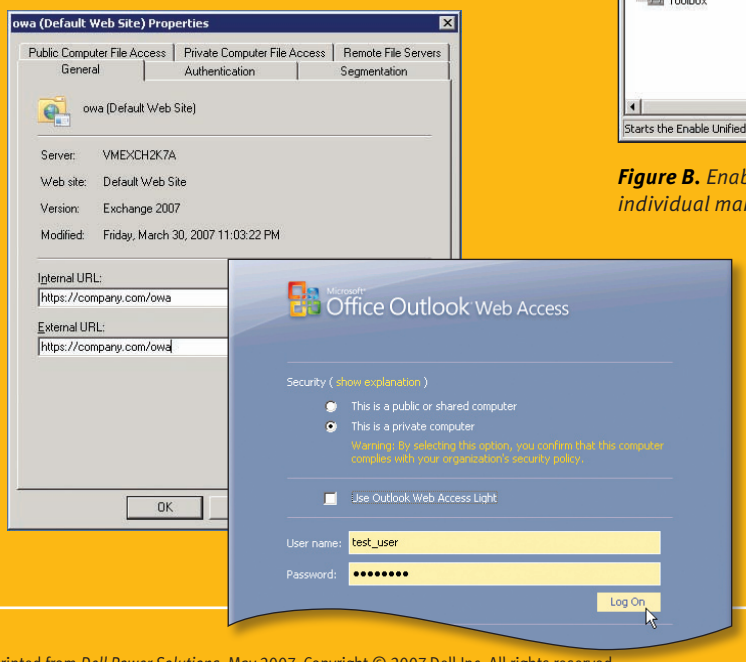

Figure C. *Establishing Outlook Web Access URL properties and logging in through the Web-based interface*

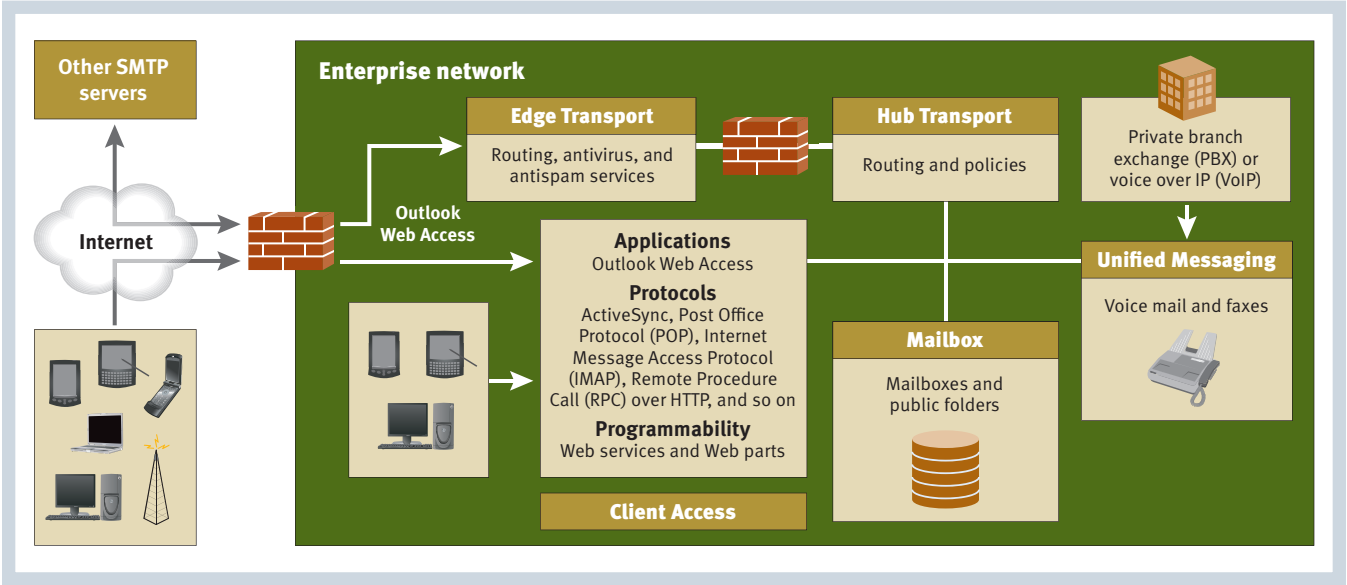

Figure 1. *Microsoft Exchange Server 2007 integrated messaging architecture*

users additional ways to connect to the information they need wherever they happen to be throughout the day.

Expanding the boundaries of the workplace with anywhere access

Exchange Server 2007 helps simplify connectivity for end users by integrating the various types of communication used throughout the enterprise. The experience of using Microsoft Office Outlook® software remains consistent from desktop to mobile devices, allowing employees to remain productive without having to adapt to different interfaces (see Figure 2). The Outlook Anywhere feature allows users to easily access their mailbox from Internet-connected computers running the full Outlook client without connecting to a virtual private network, while Outlook Web Access gives users access to their mailbox and network resources through a Web browser–based interface (see the sidebar on the preceding page). Mobile devices enabled for Microsoft Exchange ActiveSync® technology can provide mobile access to Exchange data. When Unified Messaging capabilities are enabled, Outlook Voice Access gives users access to their mailbox, calendar, personal contacts, and enterprise directories through their phones by using speech recognition or touch tones.

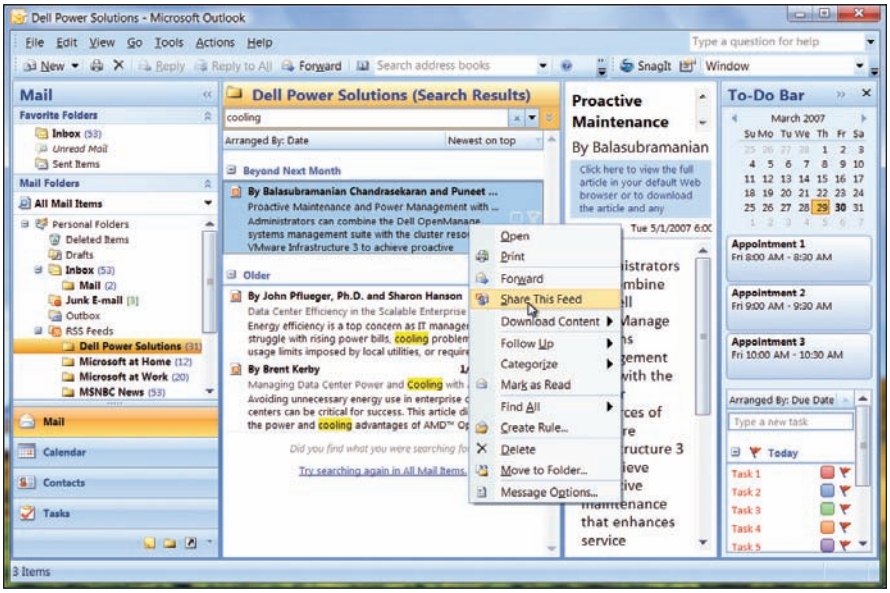

Figure 2. *Microsoft Office Outlook 2007 on Windows Vista, showing enhanced search capabilities and sharing of RSS feeds*

Exchange Server 2007 also helps facilitate efficient collaboration, which is critical for meetings and document sharing. Unified Messaging gives users anywhere access to their vital enterprise communications—including e-mail, voice mail, instant messaging, and faxes—through desktop, mobile, Web, or phone clients. The Calendar Concierge offers a set of features that help simplify and automate scheduling people and resources, and the Scheduling Assistant feature analyzes attendee and resource schedules and suggests meeting times using a color-coded user interface. The Calendar Attendant marks meeting requests as tentative on recipient calendars until users act on the request. The Resource Booking Attendant manages resource availability and allows for resource policies such as available hours and scheduling permissions.

Anywhere access extends to file and attachment viewing as well. The WebReady Document Viewing feature allows users to view a variety of document types (including Microsoft Word,

Microsoft Excel®, and Microsoft PowerPoint® files) in a Web browser, even if the application that created the document is not installed on the client.

Guarding against threats with built-in protection

Exchange Server 2007 is designed with several types of built-in protection, helping provide cost-effective, enterprise-class availability through two similar replication technologies. First, cluster continuous replication (CCR) provides virtual servers and failover capabilities. Because each node has its own copy of the Exchange Information Store, shared storage is not required. CCR allows organizations to implement a variety of storage options such as direct attach storage, Serial Attached SCSI (SAS)–based storage, and storage area networks (SANs).

Local continuous replication (LCR) helps promote high availability in another way by using the database replication technology from CCR and applying it to a stand-alone Exchange server. Instead of replicating data to virtual servers, databases are replicated to another location on the local server. The Exchange Information Store can be pointed to the local copy, and service can continue if the database becomes corrupt or a disk fails. This method of replication can be an affordable way for enterprises to enhance availability for key applications.

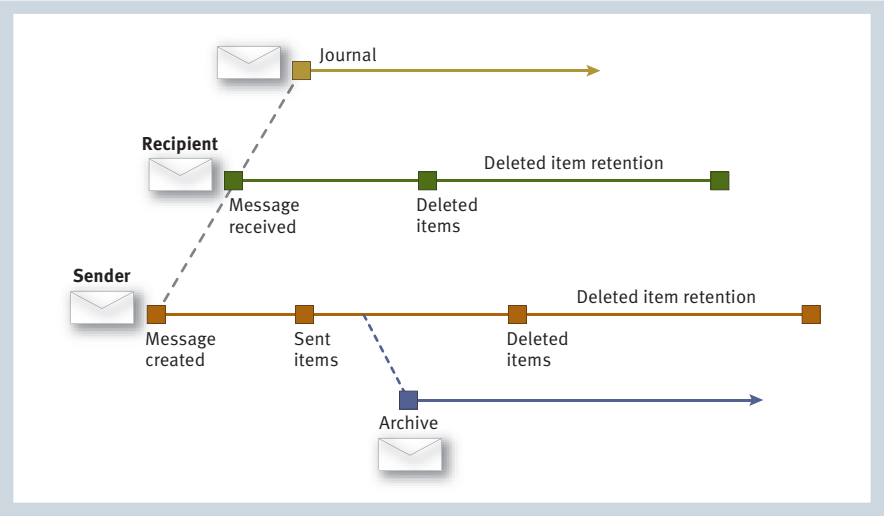

Figure 3. *Simple life cycle of an e-mail message*

Exchange Server 2007 also helps guard against external threats through expanded protection features that block spam in several ways. Safe-sender lists created by Outlook users help reduce false positives, allowing messages that users deem safe into the organization's network regardless of their spam confidence level. Message-specific puzzles and solutions, known as postmarks, also can be attached to outgoing messages. When Exchange receives an incoming message with an attached postmark, it verifies the puzzle and solution. Highly complex postmarks help reduce the likelihood that the message is spam. Potential spam can be quarantined for review by an e-mail administrator, and Exchange can dynamically analyze and update sender reputation. In addition, as spammers develop new methods for evading antispam safeguards, Exchange can automatically update the spam content filter to stay ahead of the game. Exchange also can update virus signatures, IP reputation services, and antispam filters several times each day.

To help simplify compliance, Exchange Server 2007 includes messaging records management features designed to allow administrators to define the e-mail life cycle (see Figure 3). These

COMPREHENSIVE DELL SOLUTION FOR MICROSOFT EXCHANGE SERVER 2007

Dell offers an integrated, end-to-end messaging implementation for Microsoft Exchange Server 2007 from the end user to the data center. Designed to help IT departments deploy every aspect of a unified messaging infrastructure with confidence, Dell's approach ties together multiple aspects of a messaging environment—servers, storage, networking, tools, services, and support—in a cost-effective, easy-to-use way. Rigorously tested and validated, the Dell solution helps dramatically reduce the cost and complexity of managing messaging environments. It includes the following components:

- **Servers:** Ninth-generation Dell™ PowerEdge™ servers (which are 64-bit ready)
- **Storage:** Dell PowerVault™ and Dell/EMC storage
- **Custom tools:** Dell Exchange 2007 Advisor
- **OS:** Microsoft Windows Server® 2003 Release 2 (R2)
- **E-mail:** Microsoft Exchange Server 2007
- **Security:** Symantec Mail Security for Microsoft Exchange
- **Backup and recovery:** Symantec Backup Exec
- **Compliance:** Symantec Enterprise Vault
- **Dell Services offerings:** Exchange Migration, Backup and Recovery, and Enterprise Support

To learn more about how to get started, see "Design Considerations for Deploying Microsoft Exchange Server 2007," by Suman Kumar Singh and Bharath Vasudevan, in *Dell Power Solutions,* May 2007, www.dell.com/downloads/global/power/ps2q07-20070403-Singh.pdf; and "Unlocking the Potential of Microsoft Exchange Server 2007," in *Dell Power Solutions,* May 2007, www.dell.com/downloads/global/power/ps2q07-50070342-DellSvcs.pdf.

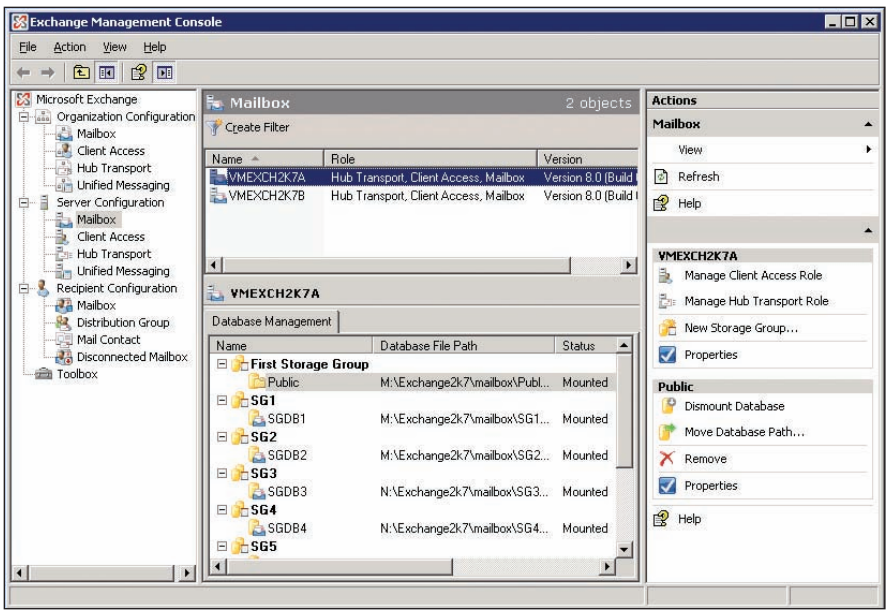

Figure 4. *Exchange Management Console in Microsoft Exchange Server 2007*

features help satisfy legal discovery and compliance regulations by making it easy to perform fast, accurate retrieval of e-mail messages.

Enhancing operational efficiency to help improve performance and cut costs

Unlike previous versions of Exchange, Exchange Server 2007 has been optimized for performance and scalability with a 64-bit architecture. Whereas the previous 32-bit architecture allowed up to 4 GB of addressable memory (split between the kernel and applications), the 64-bit design allows more of the application to reside on the processor than was possible in the 32-bit architecture.

Exchange Server 2007 offers expanded role-based deployment, allowing administrators to assign predefined roles to specific servers easily and flexibly. These roles, which can be predefined and chosen during installation, allow administrators to control e-mail flow, increase security, and distribute services as follows:

- **Mailbox:** Hosts user mailboxes stored in databases that can be replicated or clustered
- **Client Access:** Proxies Internet client traffic to the correct Mailbox server
- **Hub Transport:** Enforces messaging policies and provides internal routing of messages from Edge Transport servers and Unified Messaging servers, and between two users on the same Mailbox server database
- **Edge Transport:** Provides on-premise e-mail security, antivirus, and antispam services for Exchange
- **Unified Messaging:** Enables voice mail and fax messages to be delivered to Exchange mailboxes and provides voice dial-in capabilities to Exchange

The Exchange Management Console (see Figure 4) is designed to help IT departments increase administrator productivity. The console is divided into four independent work centers, which facilitate effective management of Exchange roles and permissions without requiring administrators to drill down several layers to get to the object to be managed.

Exchange Server 2007 also promotes operational efficiency by giving enterprises advanced control over e-mail retention and discovery. An enhanced search feature allows administrators to search across many mailboxes at once for efficient discovery. Journaling features allow administrators to audit e-mail sent to and received by a group of users—a capability that is required by several regulations and is useful to organizations for internal policies or audits. Messages journaled by Exchange Server 2007 can be stored in an Exchange database on a Microsoft Office SharePoint® Server site, or sent to an external Simple Mail Transfer Protocol (SMTP) address used by third-party journaling companies.

Embracing the future of messaging

As mobile access and a global workforce become the norm, enterprises must stay abreast of constantly emerging challenges surrounding data access, operational efficiency, security, and data protection. A new messaging paradigm is emerging after years of stopgap measures: enterprises are implementing comprehensive, integrated, end-to-end messaging solutions that incorporate not only e-mail, but also voice mail, instant messaging, and fax communications. With a wide array of new features and capabilities specifically designed to help enterprises address these imperatives, Microsoft Exchange Server 2007 delivers on this promise of unified messaging.

Jeanne Feldkamp *is a business and technology writer based in San Francisco. She has worked on several publications for leading high-tech corporations.*

Deb McDonald *is the managing editor of* Dell Power Solutions. *Prior to that, she served in editorial management positions for numerous computer and technology publications at Ziff Davis Media and Reed Elsevier.*

Tom Kolnowski *is the editor-in-chief and publisher of* Dell Power Solutions.

more Online
www.dell.com/powersolutions

QUICK LINKS

Microsoft Exchange solutions on Dell and Dell Exchange 2007 Advisor tool:
www.dell.com/exchange2007

Dell Services:
www.dell.com/services

Network Nirvana

Connectivity, Convergence and Compliance—**all Secure**

Enterasys security-enabled network infrastructure products give you granular visibility and control over who's using your network and your voice, video and data communications—wired and wireless; switching and routing; LAN and WAN.

Our advanced security applications produce compliance reports and proactively prevent threats against your IT assets with our intrusion prevention and network access control solutions.

Let us show you how we can automatically protect you and your information without sacrificing performance.

Leading companies in more than 70 countries ensure the integrity and performance of their IT services with Enterasys Secure Networks.™

We Secure *Any* Network

Set up a time to see how our unique approach can secure *any* network from *any* vendor while leveraging your existing investments. Call **+1 877-801-7082** or visit **enterasys.com/securenetworks.**

enterasys® Secure Networks™

There is nothing more important than our customers.

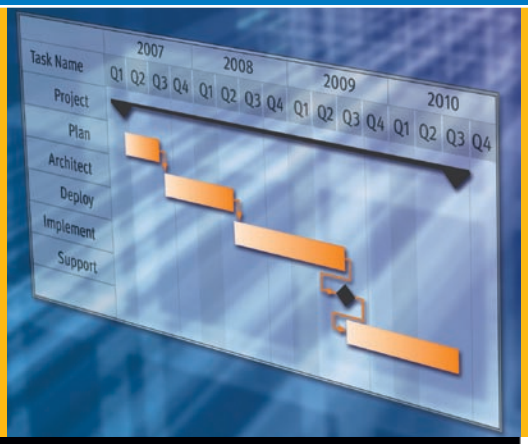

The Business of IT: Planning a Strategic Framework for Anywhere Data Access

BY CHRIS AUGER

As communication technologies converge in the workplace, IT decision makers must find new approaches to messaging infrastructure design, security, and compliance. By keeping the needs of highly mobile clients in mind during the Microsoft® Exchange Server 2007 planning process, organizations can simplify operations, optimize solutions, and sustain value across the overall IT infrastructure.

Related Categories:

2007 Microsoft Office

Data center technology

Microsoft Exchange Server 2007

Microsoft Windows Vista

Storage

Unified communications

Visit www.dell.com/powersolutions for the complete category index.

Nearly every enterprise faces difficulties when its technology fails, even momentarily. No cell coverage? Phone lines down? These situations can be bothersome and inconvenient—but nothing interrupts business as usual quite like an e-mail outage.

Unlike the telephone, which is simply a tool for communication, e-mail typically contains a significant amount of intellectual capital—and according to the Radicati Group, corporate users send and receive an average of 133 e-mail messages per day.[1] Business usually grinds to a halt without access to this data, so e-mail systems and content must be protected just like other mission-critical infrastructures.

In addition, the proliferation of mobile devices such as wireless computers, PDAs, and cell phones is catalyzing a major change in the way IT administrators think about enterprise communication infrastructures. Because e-mail messages can be created and stored in a wide variety of locations both inside and outside the enterprise, the IT department must coordinate with telephony administrators to anticipate the needs of mobile workers. Security and compliance must be approached differently. Most importantly, organizations must begin to look at the messaging environment as a cohesive whole.

Intertwined client and data center decisions

By looking at the messaging infrastructure as a single system, administrators are quickly realizing that client and data center decisions and deployments cannot be made in a vacuum. Desktop systems such as the Microsoft Windows Vista™ OS, Microsoft Office Outlook® and other Microsoft Office software, and thin clients affect data center elements such as Microsoft Exchange Server, Microsoft Office SharePoint® Server, systems management, storage, and file and print services. Conversely, data center decisions determine which client systems the organization can use. Because these decisions are inextricably linked, it is crucial that IT decision makers consider client systems and data center elements together.

For example, the method employees use to access their e-mail can significantly affect the number of servers and the amount of storage required to support the e-mail application. Cached Outlook access demands far fewer database spindles and servers than remote PDA access. Figure 1 provides a basic example designed to highlight the impact client choices can have on Exchange Server 2007 enterprise designs—specifically, differences in the number of database spindles and servers required to support 5,000 users and a 500 MB mailbox size limit with different types of clients.

[1] "Taming the Growth of Email—An ROI Analysis," by the Radicati Group, March 2005, www.radicati.com/publication.asp?id=302.

As the workforce becomes increasingly mobile, IT departments no longer have the option of offering only one type of e-mail access. Growing e-mail volumes and new types of media attachments require a data center infrastructure that can provide not only anywhere access, but also the scalability and control needed to help ensure outstanding performance as well as built-in protection to mitigate risk.

Microsoft Exchange, Windows Vista, and Office: Better together

Microsoft Exchange Server 2007, Microsoft Windows Vista, and 2007 Microsoft Office are designed to build on each other's strengths and provide a platform that is optimized to meet the needs of a mobile workforce. Implemented together, Exchange Server 2007, Windows Vista, and 2007 Microsoft Office can enhance messaging from the user to the data center.

For example, the Exchange Server 2007 Calendar Concierge feature and Outlook 2007 streamline scheduling meetings at the client level by automatically recommending the best days and times for a meeting. The Out of Office Assistant in Outlook can send customized internal and external out-of-office messages. Windows Vista Instant Search allows users to search their entire PC for files and e-mail messages.

In the data center, Exchange Server 2007 supports Windows Vista and 2007 Microsoft Office with its Unified Messaging capabilities. Exchange Server 2007 is designed to deliver all types of Outlook items—e-mail, voice mail, instant messages, and faxes—to user in-boxes. Users can tag and organize messages with customized classifications (such as "Do Not Forward" or "Attorney-Client Privileged"), while built-in e-mail retention policies through managed folders or SharePoint Server help ensure that compliance requirements are met.

Anywhere access for enhanced productivity and collaboration

The basic requirement for a mobile workforce is the ability to access multiple types of messages and content from a variety of devices. Without restrictions on where and when they can access key data, employees can accomplish tasks in the manner that best suits them—helping increase productivity and providing opportunities for collaboration.

However, anywhere access is not as simple as delivering e-mail to several types of devices. Truly efficient messaging systems provide not only mobile access but also a consistent interface across devices. They integrate messaging as well as calendaring, document management, and search functions.

Used together, Exchange Server 2007, Windows Vista, and 2007 Microsoft Office offer a range of complementary features designed to enhance productivity and encourage collaboration through anywhere access. Public folders are de-emphasized in Exchange Server 2007 to enhance focus on Microsoft Windows® SharePoint Services. A single mailbox for e-mail, voice mail, instant messages, and faxes provides a streamlined center for communication, and the Calendar Concierge feature helps improve scheduling, resource booking, and assistance.

Windows Vista extends the benefits of Exchange Server 2007 to client systems by enabling users to easily organize and search large volumes of information. The Instant Search feature provides integrated search functionality for the Internet, multimedia, and e-mail, helping individual users access important content quickly. Desktop search capabilities also make it easy for users to find files regardless of location by providing quick access to Microsoft Office documents, e-mail messages, programs, and media files throughout the computer. Microsoft Enterprise Search facilitates finding the right information regardless of where it resides—from redirected folders, offline folders, and SharePoint Server documents to removable hard drives and other PCs running Windows Vista.

Windows Vista helps users organize files with metadata tags indicating that information pertains to a particular project or category. Users can then search, filter, and organize files based on certain criteria, no matter where the files reside on the hard drive. Search folders allow users to save frequent search queries. Saved searches can be executed when clicked, automatically organizing information for easy and fast access.

In addition, roaming capabilities enable the files of mobile Windows Vista users to sync with enterprise servers once the end point reconnects to the network—helping minimize data loss, maintain organization of key intellectual property, and maximize productivity of remote workers.

Scalability and control for operational efficiency

To support anywhere access, organizations must build messaging infrastructures that provide outstanding performance, scalability, and control at every level. Exchange Server 2007, Windows Vista, and 2007 Microsoft Office offer complementary hardware and software functionality designed not only to boost performance and capacity but also to simplify messaging management.

Access mode	Database spindle requirements		Server requirements
	Exchange Server 2003	Exchange Server 2007	
Microsoft Outlook cached	38	22	One Exchange Mailbox server
Microsoft Outlook online	52	42	One Exchange Mailbox server
Microsoft Outlook Web Access	52	42	• One Exchange Mailbox server • Two Exchange Client Access servers
BlackBerry mobile device	71	71	• Two Exchange Mailbox servers • Two Exchange Client Access servers • Two BlackBerry Enterprise Server systems

Figure 1. *Example enterprise requirements for Microsoft Exchange Server supporting 5,000 heavy users and a 500 MB mailbox size limit*

In the data center, Exchange Server 2007 enables several improvements over Exchange Server 2003 that help lay the foundation for optimal messaging performance. By supporting only 64-bit servers in production, such as Dell™ PowerEdge™ 2950 and PowerEdge 2970 servers, Exchange Server 2007 allows increased memory and caching capabilities compared with 32-bit servers, which can significantly reduce storage performance requirements. The maximum number of storage groups or databases has been increased to 50, which means administrators can provide employees with larger mailboxes than in Exchange Server 2003 while retaining the same backup and restore sizes. In addition, an increase to an I/O size of 8 KB means fewer I/Os per second and more flexible storage options in Exchange Server 2007 compared with Exchange Server 2003.

Exchange Server 2007 can also help simplify messaging management. Expanded role-based deployment allows administrators to assign predefined roles to specific servers easily and flexibly. These roles, which can be predefined and chosen during installation, enable administrators to control e-mail flow, increase security, and distribute services. The Exchange Management Console also helps administrators increase productivity through four independent work centers, which facilitate effective management of Exchange roles and permissions without having to drill down several layers to get to the managed object.

Windows Vista builds on the management capabilities of Exchange Server 2007 with Group Policy Objects, which provide hundreds of policy settings for flexibly managing new and existing features. XML-based policies support multilingual environments and use a common syntax across policy settings to help simplify messaging management.

Furthermore, significant improvements in the event logging infrastructure in Windows Vista Event Viewer make it easy for administrators to perform troubleshooting tasks. New power management features promote power conservation and maximize battery life for mobile users, and enhanced system diagnostics and remote assistance tools empower employees to troubleshoot while helping increase the efficiency of IT staff. In addition, image management features enable simplified OS migrations and image updates across multiple hardware systems.

Built-in protection for risk mitigation

Security and compliance requirements are major concerns for enterprise IT departments. As electronic transaction processing and real-time communication become fixtures of global commerce, governments around the world are introducing strict, detailed regulations designed to ensure that critical messaging systems stay up and running and that enterprise data is traceable throughout its life cycle.

Exchange Server 2007 is designed with several types of built-in protection, helping mitigate risk with two related replication technologies. Cluster continuous replication (CCR) provides virtual servers and failover capabilities, and allows organizations to implement a variety of storage options such as direct attach storage, Serial Attached SCSI (SAS)–based storage, and storage area networks (SANs). Local continuous replication (LCR) helps promote high availability by using the same database replication technology as CCR but within a stand-alone Exchange Server 2007 server. This method of replication can provide a cost-effective way for enterprises to increase availability for key applications.

Other Exchange Server 2007 features designed for recovery and high availability include single-mailbox restores using the Restore-Mailbox feature and a multi-mailbox search function. This enhanced search functionality across multiple file types and thumbnails of documents makes it easy to find important files.

A range of features designed for security and compliance help protect enterprises against risk. In the data center, Exchange Server 2007 provides antivirus and antispam protection through the Edge Transport server role. Administrators can use built-in e-mail retention policies through managed folders or SharePoint Server, and Exchange Hosted Services can deliver outsourced antivirus and antispam protection.

At the client level, Windows Vista helps mitigate risk through User Account Control (UAC) features. By default, UAC limits privileges to install applications or make configuration changes that could compromise the integrity of the client system. When a task requires administrator-level access, UAC asks the user to accept the elevation and then bonds the administrator token to the standard user account to provide enhanced protection against malware. Built-in features such as an enhanced firewall and spyware detection are designed to protect organizations by blocking harmful Web sites and hackers from accessing the system.

The 2007 Microsoft Office system also helps counter risk through enhancements to the Microsoft Internet Explorer® 7 Web browser. Protected Mode options help guard against malware, while the built-in Windows Firewall supports bidirectional filters and includes rules designed to protect services and ports. Windows Resource Protection features block write access to system files and registry keys. In addition, Microsoft BitLocker™ Drive Encryption in Windows Vista provides a native disk encryption tool to mitigate risk when hardware is lost or stolen.

Effective game plan for an enterprise-wide messaging infrastructure

As employees become increasingly mobile, IT decision makers must consider new approaches to messaging access, infrastructure performance, security, and compliance. This shifting paradigm demands that businesses make executive decisions with both the front-end user interface and the data center in mind. By facilitating anywhere access, operational efficiency, and risk mitigation, the combination of Microsoft Exchange Server 2007, Microsoft Windows Vista, and 2007 Microsoft Office can help optimize enterprise messaging and sustain the value of the overall IT infrastructure from the user to the data center.

Chris Auger *is an enterprise technologist for Dell-focused Microsoft products. He has a B.S. in Mathematics from the University of Texas at Austin.*

BROADCOM®
Connecting everything®

Efficient server performance was never easier – and never more affordable…

Dell™ 9th Generation PowerEdge™ Servers With TCP/IP Offload Engine Technology!

Dell's innovative 9th generation PowerEdge Servers with Broadcom® NetXtreme II™ TOE enabled Gigabit Ethernet deliver enterprise level performance, whether you are an enterprise or small business. Tower, rack and blade models provide flexibility and expandability for growing environments with a wide variety of configuration options available. The next generation of servers offering power and performance that you have come to expect from Dell.

Dell PowerEdge 2950

Designed to Reduce Complexity

- Featuring Intel® Xeon® processors
- Factory installed OS available to choose from Microsoft, Red Hat and Novell
- Dual embedded Broadcom® NetXtreme II™ 5708 Gigabit Ethernet NIC with fail-over and load balancing
- TOE (TCP/IP Offload Engine) supported on Microsoft® Windows Server® 2003*
- 12 FBD DIMM sockets for supporting up to 48GB of main memory
- 3.5 inch SAS or SATA drives supporting up to 6.0TB of storage**
- 6 expansion slots:
  1 x8 PCI Express - x8 lane with x8 connector
  3 x4 PCI Express - x4 lane with x8 connector
  2 x 64-bit/133MHz PCI-X – supports full-height, full-length 3.3v PCI or PCI-X cards
- Optional 3 year Gold Enterprise Support: 7x24 HW/SW, 4hr 7x24 onsite***

TCP/IP Offload Engine Technology is also available in Dell PowerEdge models 1900, 1950, 2900 and 6950

* SP1 or higher with Scalable Networking Pack
**For hard drives, GB means 1 billion bytes and TB equals 1 trillion bytes; actual capacity varies with preloaded material and operating environment and will be less.
***Service may be provided by a third party. Technician will be dispatched if necessary in parallel with phone-based troubleshooting for severity level 1 issues on systems with Gold or Platinum 2-hour or 4-hour on-site response service. Technician will be dispatched if necessary following phone-based troubleshooting for all other severity levels and on-site response service levels. 6-hour repair time is measured from dispatch time of the on-site response service. Subject to parts availability, geographical restrictions and terms of service contract. Service timing dependent upon time of day call placed to Dell. U.S. only.

DELL®

Understanding the Architecture and Features of **Microsoft Exchange Server 2007**

BY ANANDA SANKARAN

SUMAN KUMAR SINGH

Microsoft® Exchange Server 2007 includes significant changes from Exchange Server 2003. This article describes key technical aspects of these changes—including its core architecture, mailbox data access, and secure deployment features—to enable administrators to effectively prepare for migrations to Exchange Server 2007.

Related Categories:

Dell PowerEdge servers

Dell/EMC storage

Microsoft Exchange Server 2007

Microsoft Windows Server 2003

Outlook Web Access

Regulatory compliance

Security

Unified communications

Visit www.dell.com/powersolutions for the complete category index.

Messaging systems are an integral component of many enterprises. Over the years, they have evolved from providing basic functionality such as e-mail to encompassing rich collaboration features as well. Users' messaging requirements have also changed significantly, from basic desktop e-mail access to remote collaboration available from multiple devices. In such an evolving environment, administering and maintaining messaging systems has become increasingly complex for IT administrators—and enterprise requirements for security, regulatory compliance, and availability only add to that complexity. Enterprises also face an enormous challenge when implementing a suitable messaging system that meets their targets for total cost of ownership and return on investment.

Microsoft Exchange Server 2007 incorporates features that enable enterprises to effectively meet these challenges without fundamentally altering the traditional methods used by end users to access e-mail. This article highlights some of these features, discussing the core architecture, mailbox data access, and secure deployment components of Exchange Server 2007.

Core architecture of Exchange Server 2007

Exchange Server 2007 introduces several architectural changes and other features designed to enhance its scalability and performance compared with Exchange Server 2003. Figure 1 summarizes the major differences between these two versions, which include key changes to server roles, memory use, and storage groups.

Server roles

Exchange Server 2007 includes five defined server roles—sets of functionality that administrators can deploy individually on servers or combined with other roles, with certain restrictions:

- **Mailbox:** Hosts user mailboxes, public folders, and calendar data
- **Client Access:** Provides functionality for Microsoft Office Outlook® Web Access (OWA), Microsoft Exchange ActiveSync® client access, Post Office Protocol 3/Internet Message Access Protocol 4 (POP3/IMAP4) client access, and Outlook Anywhere access—referred to in Exchange Server 2003 as Remote Procedure Call (RPC) over HTTP
- **Hub Transport:** Handles internal e-mail flow across Exchange components and user message delivery, along with functionality such as journaling, server-side rules, and an additional layer of antivirus and antispam security; this role is mandatory in Exchange Server 2007 deployments
- **Edge Transport:** Resides in the perimeter network and routes external Simple Mail Transport Protocol (SMTP) e-mail flow to and from the Internet, and provides

antivirus and antispam security, filtering, and rule-based protection

- **Unified Messaging:** Integrates e-mail, voice mail, and faxes into user mailboxes and provides Outlook Voice Access functionality, which allows users to access e-mail, voice mail, faxes, calendars, contacts, and directory entries from a telephone

These five server roles refine and add features to the traditional roles available with Exchange Server 2003. The mailbox and public folder functionality of Exchange Server 2003 back-end servers is provided by the Mailbox role in Exchange Server 2007. Some of the functionality of Exchange Server 2003 bridgehead servers is provided and enhanced by the Hub Transport role. Figure 2 illustrates the different server roles as part of an example Exchange Server 2007 infrastructure.

With the exception of the Edge Transport server role, all Exchange Server 2007 server roles are members of Microsoft Active Directory® domains. The Edge Transport role resides in the perimeter network, outside of Active Directory frameworks, using Active Directory Application Mode and EdgeSync to obtain the required Active Directory information. Because of its location, the Edge Transport role cannot be consolidated and deployed on a single-server system with other server roles. The other restriction on role consolidation is that when the Mailbox role is deployed in a highly available clustered configuration using single-copy clustering (SCC)

	Exchange Server 2003	Exchange Server 2007
Platform	32-bit application based on x86 platforms	64-bit application based on x86-64 platforms
Memory scalability	Limited to 4 GB of system memory	Scales beyond 4 GB of system memory
Number of supported storage groups	• Up to 4 storage groups and 1 recovery storage group • Up to 5 mailbox databases per storage group	• Up to 50 storage groups • Up to 50 mailbox databases total
Server roles	Back-end, front-end, and bridgehead	Mailbox, Client Access, Hub Transport, Edge Transport, and Unified Messaging
Mailbox high-availability options	Microsoft Cluster Service–based shared storage clustering	• Single-copy clustering (Microsoft Cluster Service–based shared storage clustering) • Cluster continuous replication (Microsoft Cluster Service–based majority node set clustering)
Supported backup methods	Legacy and Volume Shadow Copy Service (VSS) backups on active databases	Legacy and VSS backups on both active databases and database copies using local continuous replication and cluster continuous replication
Business continuance volumes	Provided by storage hardware or third-party software mechanisms	Provided through local continuous replication, in addition to storage hardware or third-party mechanisms
Unified Messaging features	None	• Unified mailbox for e-mail, voice mail, and faxes • Voice access to e-mail, voice mail, faxes, calendars, contacts, and directory entries

Figure 1. *Major architectural differences between Microsoft Exchange Server 2003 and Exchange Server 2007*

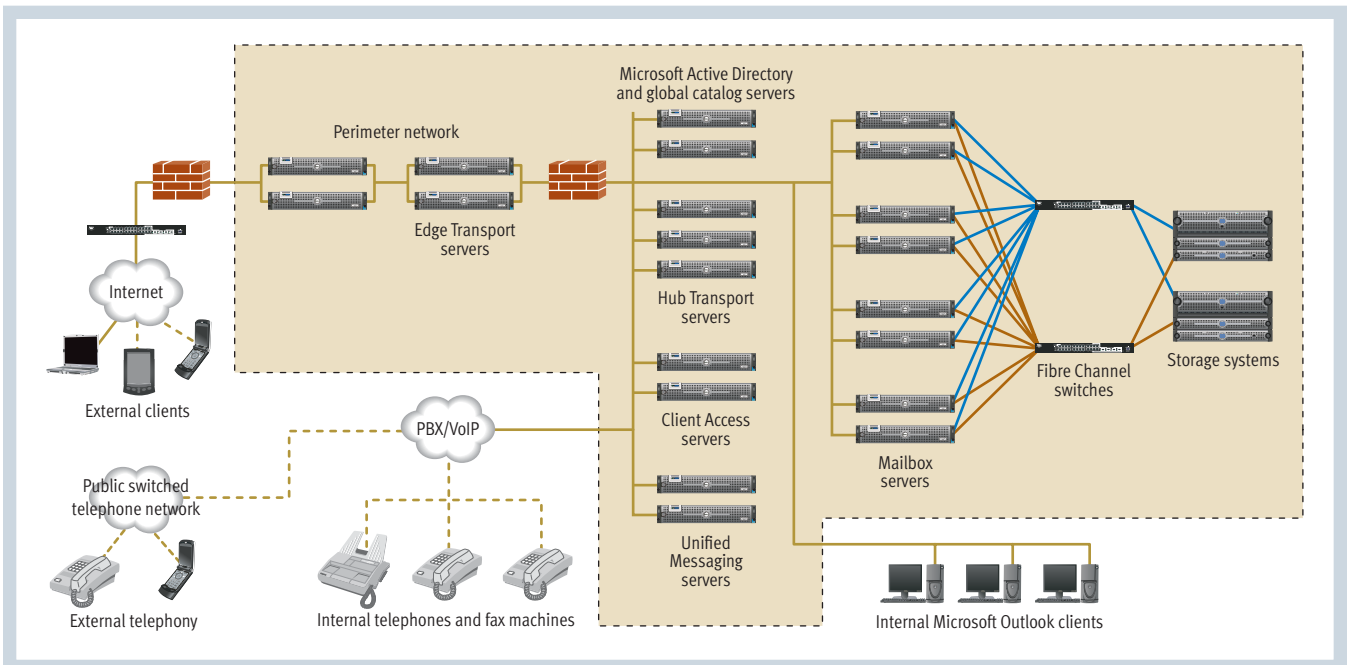

Figure 2. *Server roles as part of an example Microsoft Exchange Server 2007 infrastructure*

or cluster continuous replication (CCR) based on Microsoft Cluster Services, no other role can be consolidated with the Mailbox role on the same server.

Administrators should be sure to appropriately size the hardware for each server role for capacity and performance. High availability and scalability for server roles other than the Mailbox role can be achieved by deploying multiple server nodes hosting each role and using network load balancing. This flexibility to either consolidate roles or split them to different servers makes Exchange Server 2007 a highly scalable application: as an organization grows, administrators can choose to add servers to a particular role or split roles among additional servers.

Memory use

Exchange Server 2003 is a 32-bit application, and is limited to 4 GB of addressable memory: it can utilize roughly 3 GB of user-mode virtual memory, with the remaining 1 GB required for the kernel. Because Exchange Server 2003 does not support using Address Windowing Extensions (AWE), it cannot access physical memory above 4 GB made available through Physical Address Extension (PAE) on supported Microsoft Windows® OS–based systems, requiring the application to rely highly on disk subsystems and perform a large number of I/O operations. Beyond the physical memory limitation, the kernel running Exchange Server 2003 is also under constant stress because of the limited kernel-mode memory available for user connections and other processing.

Exchange Server 2007 overcomes these memory limitations by providing support as a 64-bit application capable of running on supported Microsoft Windows Server® 2003 x64 Editions platforms, which make approximately 8 TB of addressable memory available for user-mode and kernel-mode applications. Windows Server 2003 Enterprise x64 Edition supports up to 1 TB of physical RAM. Both the application and kernel can have sufficient memory for operations, allowing the Exchange Server 2007 Extensible Storage Engine to utilize additional memory to buffer data pages and thereby help reduce the number of required I/Os (specifically read operations) to the disk subsystem.

Storage groups

Exchange Server 2007 provides support for up to 50 storage groups to host mailbox stores or databases, a significant increase from the 4 storage groups supported in Exchange Server 2003. This support enables administrators to split mailboxes across multiple storage groups and helps simplify administrative operations such as backup and restore. Splitting mailboxes across multiple storage groups also increases the checkpoint depth available for user data operations. In certain situations, dirty data pages—those that have not yet been written to disk—can be optimized to reside in memory for additional time to help reduce the number of required I/O write operations to the disk subsystem. Exchange Server 2007 also increases the data page size from 4 KB to 8 KB, which allows it to optimize I/O in some situations by containing large messages and internal data structures within a single page.

Mailbox data access in Exchange Server 2007

Exchange Server 2007 enables users to access their mailbox data in multiple ways using computers, telephones, and handheld devices. The main capabilities of this access are enabled by back-end functions, including mobile messaging, Web-based messaging, and the Unified Messaging feature. In addition to these three key capabilities, Exchange Server 2007 works closely with Outlook 2007 to provide rich features such as the scheduling assistant, which can automatically find appropriate meeting times based on free and busy data for each participant. Outlook 2007 also provides features for organizing different message types, such as e-mail, voice mail, and faxes.

Mobile messaging

Mobile messaging is enabled by Exchange ActiveSync, which allows users to access data on different supported mobile devices through low-latency wireless data networks. The Direct Push feature helps keep Outlook mobile client devices up-to-date by providing new item notifications, and many mobile devices provide the ability to read and edit attachments that use common Microsoft Office file formats. Exchange Server 2007 ActiveSync also supports HTML messages, enhanced message flagging, server-side searching for items not stored locally, and access to Microsoft Windows SharePoint® Services and Windows file shares. The Exchange Server 2007 Client Access server role hosts the services required for ActiveSync functionality, allowing Client Access servers to communicate with mobile clients that can connect to the Internet through high-speed mobile data networks.

Web-based messaging

Web-based messaging is enabled by OWA and the Outlook Anywhere feature. OWA provides Web browser–based access to mailbox data, and Exchange Server 2007 provides enhanced browser-based functionality that includes access to Unified Messaging data, HTML data conversion, and reduced local client data storage to enhance security. OWA also provides access to Windows SharePoint Services and Windows file shares. This feature requires Internet connectivity through a Web browser, but not an Outlook client.

The Outlook Anywhere feature provides Outlook client connectivity to Exchange Server 2007 through the Internet, without requiring a connection to an internal enterprise network; in Exchange Server 2003, similar functionality was referred to as RPC over HTTP. The Client Access server role hosts the services and functionality required for OWA and Outlook Anywhere support, and facilitates the access of OWA Web browser clients and remote Outlook clients to data hosted on Exchange Server 2007 Mailbox servers.

Unified Messaging

Traditionally, e-mail, voice mail, and fax data have been maintained in separate systems. Exchange Server 2007 Unified Messaging enables these three types of data to be routed

to Outlook in-boxes, allowing the in-boxes to serve as consolidated repositories for messaging data. Users can access this data using Outlook clients on a computer or mobile device, or through OWA using a Web browser. Voice mail messages and incoming faxes appear as e-mail attachments with unique identification flags; voice mail can be played on computer or device speakers. Unified Messaging also provides telephone access to the messaging data through Outlook Voice Access: users can dial in from their office extension or an external telephone and access their e-mail, voice mail, faxes, calendar, contacts, and directory entries.

The Exchange Server 2007 Unified Messaging server role hosts the services and functionality required to implement Unified Messaging. Unified Messaging servers use voice over IP (VoIP) protocols, enabling them to receive voice mail and fax messages from existing private branch exchange (PBX) telephony systems and store these messages in user mailboxes hosted by Exchange Server 2007 Mailbox servers. Through their connection with the telephony system, Unified Messaging servers also provide the interface for voice access to Outlook data. A VoIP gateway device may be needed to translate between Unified Messaging servers and the PBX system; it is critical that this gateway meet the requirements for Unified Messaging server integration.

Secure enterprise deployments for Exchange Server 2007

Exchange Server 2007 introduces several features designed to enhance security, compliance and archiving, and backup and recovery functionality.

Security

Exchange Server 2007 provides enhanced built-in antivirus and antispam security features, including filtering based on IP address, content and attachments, and sender ID, as well as message stamping to identify scanned messages in the system. Edge Transport servers provide these features in perimeter networks and can filter potential threats before they enter internal networks. Administrators can also configure Hub Transport servers with antivirus and antispam functionality when these servers have been set up to relay external messages directly without utilizing Edge Transport servers.

Compliance and archiving

Exchange Server 2007 includes features designed to help organizations effectively meet their regulatory compliance and archiving requirements. Administrators can configure transport rules for Hub Transport servers that enforce restrictions on internal and external communications based on the users involved or the message content. Messaging records management features provide managed folders—Outlook user folders that can be centrally managed by administrators. Journaling features have been enhanced to include journaling at the mailbox, distribution list, or mailbox database level, and messages included in journaling can be sent to a third-party SMTP server, an Exchange mailbox, or a Microsoft Office SharePoint Server folder for archiving. User-wide mailbox scan and search functionalities are included to help meet compliance and retention requirements.

Backup and recovery

Exchange Server 2007 supports backups of database copies when using local continuous replication (LCR) or CCR. Backing up copies, rather than active production databases, helps increase performance by reducing the load on production servers and e-mail clients and enables administrators to remove time restrictions on backup and maintenance windows. The database portability feature allows other servers to host a failed server's user mailboxes, enabling users to continue sending and receiving messages while the backup mailbox data is being recovered.

Enhanced Exchange Server 2007 architecture and features

Microsoft Exchange Server 2007 introduces multiple features and enhancements designed to meet ever-growing enterprise messaging needs. Dell™ PowerEdge™ servers and Dell PowerVault™ and Dell/EMC storage provide a standard hardware platform for deploying Exchange Server 2007 messaging solutions, and Dell Services can provide assessments, designs, and implementations tailored for those messaging deployments. Dell also offers end-to-end Exchange messaging solutions that include partner offerings for security, archiving, and backup and recovery. Enterprises can take advantage of these services when planning to upgrade to or deploy Exchange Server 2007 in their data centers.

Ananda Sankaran *is a lead engineer in the Dell End-to-End Solutions team. His interests include databases, application consulting, performance engineering, and sizing, and he led the release of the Dell SQL Server 2005 Advisor and Dell Exchange 2007 Advisor tools. Ananda has a master's degree in Computer Science from Texas A&M University.*

Suman Kumar Singh *is a lead systems engineer on the Dell End-to-End Solutions team. He specializes in messaging systems architecture and sizing, and led the release of the Dell Exchange 2003 Advisor tool. His other interests include storage area networks, virtualization, and security, and he has published and presented several papers at industry conferences. Suman has a master's degree from the University of Texas at El Paso.*

more Online

www.dell.com/**powersolutions**

QUICK LINKS

Microsoft Exchange Server 2007 on Dell:
www.dell.com/exchange2007

Dell Secure Exchange solutions:
www.dell.com/secureexchange

A Blueprint for Implementing **Microsoft Exchange Server 2007 Storage Infrastructures**

BY DERRICK BAXTER

SURESH JASRASARIA

Designing a consolidated storage infrastructure is the first step toward successfully implementing Microsoft® Exchange Server 2007 to help manage ever-increasing e-mail growth. EMC has created an integrated, cost-effective storage infrastructure blueprint for Exchange Server 2007 deployments using Dell/EMC CX3 UltraScale™ series storage arrays and EMC® software that can deliver high availability and performance while scaling effectively for future growth.

Related Categories:

Dell/EMC storage

Microsoft Exchange Server 2007

Microsoft Windows Server 2003

Visit www.dell.com/powersolutions for the complete category index.

Many organizations rely on Microsoft Exchange for e-mail and collaboration services, and for these organizations, Exchange can be one of their most critical communication tools—often even more critical than telephone and fax systems. This reliance can increase the number of e-mails as well as the variety and richness of attachments users send and receive, leading to explosive growth in both e-mail message volume and mailbox size. Exchange Server 2007 introduces many features designed to accommodate this growth, including a native 64-bit architecture, an increased number of message stores, local and cluster continuous replication, flexible journaling, and policy-based message records management.

Designing a consolidated storage infrastructure to support Exchange Server 2007 is the first step toward an efficient and cost-effective deployment. After extensive testing of Exchange Server 2007 workloads running on Dell/EMC CX3 UltraScale series storage arrays with an EMC software infrastructure, EMC has created a blueprint for Exchange Server 2007 deployments designed to provide an integrated, cost-effective solution that delivers high availability and performance while scaling effectively for future growth.

Although changes to the database engine in Exchange Server 2007 due to its 64-bit memory support, increased I/O size (8 KB in Exchange Server 2007, compared with 4 KB in Exchange Server 2003), and increased Extensible Storage Engine cache size have lowered I/O requirements compared with Exchange Server 2003, I/O subsystem availability and performance remain critical for Exchange Server 2007 deployments. Dell/EMC CX3 storage arrays include high-performance Fibre Channel drives to store production data as well as cost-effective, high-capacity Serial ATA (SATA) drives to store backup and archive data. These storage arrays support native Internet SCSI (iSCSI) and Fibre Channel connectivity—including a minimum of four Gigabit[1] Ethernet ports and a minimum of four 4 Gbps Fibre Channel ports—and provide dual redundant active storage processors with battery-backed caches. They are designed to provide data-in-place upgrades to allow scaling from small- to very-large-capacity configurations.

Storage-based EMC software such as the EMC Navisphere® and SnapView™ applications (with SnapView providing snapshot and clone functionality), and host-based EMC software such as the EMC Replication Manager and RepliStor® applications, are tightly integrated with Microsoft technologies like Volume Shadow Copy Service (VSS). This integration is designed to provide rapid e-mail, mailbox, database, and entire server recovery and restore, helping maximize operations efficiency.

[1]This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

The global architecture of the EMC blueprint for Exchange Server 2007 storage infrastructures includes four major modules:

- **Store:** Uses a building block approach to provide a solid foundation for the infrastructure; this module is at the core of the EMC blueprint, with the other three modules being added as needed
- **Backup:** Adds backup and rapid recovery functionality at the server, storage group, database, and mailbox levels
- **Archiving:** Supports unrestricted mailbox quotas while keeping databases manageable, and does so transparently to end users of Microsoft Outlook® software
- **Protection:** Adds disaster recovery functionality

This article outlines the physical and network architectures of this blueprint and discusses the Store and Backup modules. An article addressing the Archiving and Protection modules is planned for a future issue of *Dell Power Solutions*.

Physical and network architectures

Figure 1 shows the physical architecture of the Store and Backup modules in the EMC blueprint for Exchange Server 2007 deployments using a Dell/EMC CX3-20 storage array. To help simplify testing, the Exchange Server 2007 support infrastructure, including the domain controller and global catalog, is consolidated on VMware® ESX Server 3.0.1 virtual machines (VMs) running on a Dell™ PowerEdge™ 6850 server with two dual-core Intel® Xeon® processors at 3.0 GHz and 32 GB of RAM. The Exchange Server 2007 Client Access and Hub Transport roles are assigned to a separate PowerEdge 1850 server with two dual-core Intel Xeon processors at 2.8 GHz with Intel Hyper-Threading Technology and 8 GB of RAM.

In this blueprint, Exchange Server 2007 runs on a PowerEdge 2950 server with two dual-core Intel Xeon processors at 2.4 GHz, 16 GB of RAM, and the Microsoft Windows Server® 2003 Enterprise x64 Edition OS with Service Pack 1 (SP1). This server also runs EMC PowerPath® 4.6

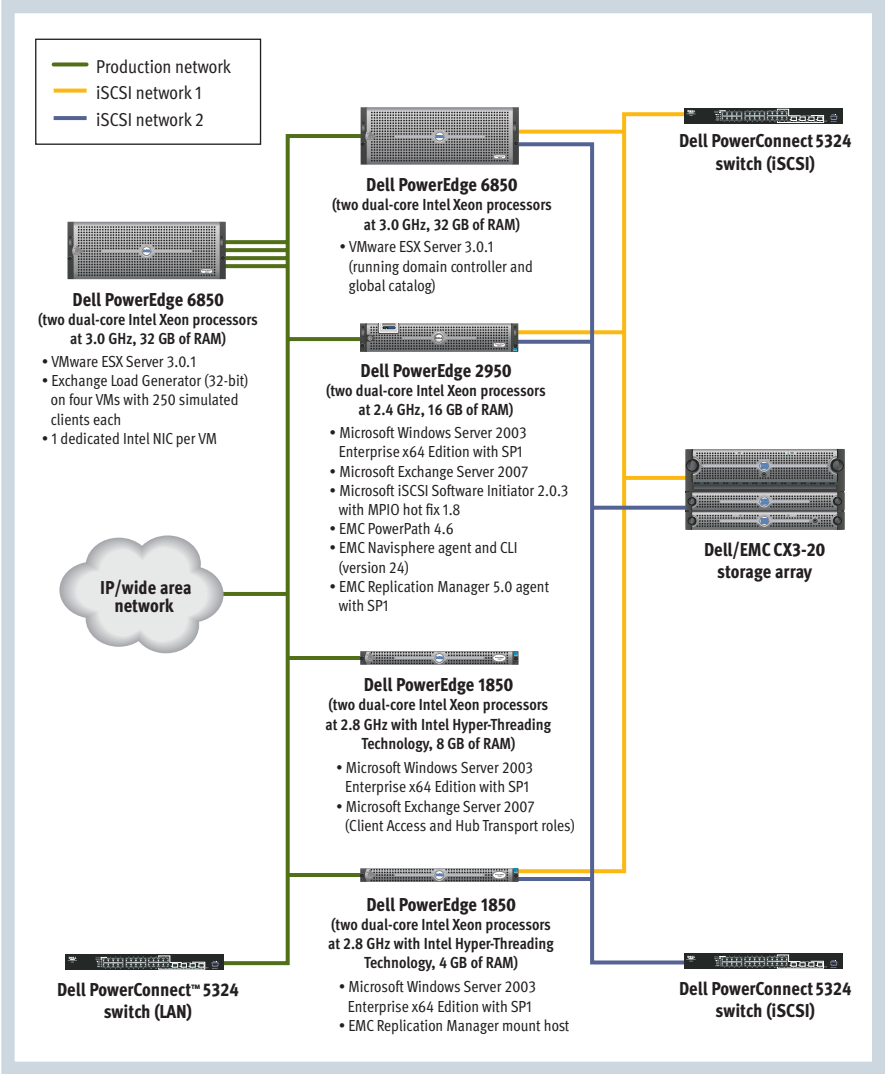

Figure 1. *Physical architecture of the Store and Backup modules in the EMC blueprint for Microsoft Exchange Server 2007 deployments*

software to help provide resiliency against a storage processor failure in the Dell/EMC CX3-20 array, as well as the EMC Replication Manager 5.0 agent with SP1 to help orchestrate rapid backup and recovery of Exchange Server 2007 data. In addition, this PowerEdge 2950 server includes Microsoft iSCSI Software Initiator 2.0.3 with Microsoft Multipath I/O (MPIO) hot fix 1.8 and version 24 of the EMC Navisphere agent and command-line interface (CLI) to allow iSCSI access to the Dell/EMC CX3-20 array. EMC PowerPath uses the MPIO framework to provide path failover functionality.

EMC Replication Manager 5.0 with SP1 is designed to simplify replication and recovery of Exchange Server 2007 data by coordinating and automating the steps required at the application, host, and array levels. It is a simple, easy-to-use application that utilizes the Microsoft VSS and Virtual Device Interface (VDI) frameworks; in addition to supporting Exchange Server 2007, it can provide rapid backup and recovery of Exchange Server 2003 and the Microsoft SQL Server™ 2000 and SQL Server 2005 database platforms. In the EMC blueprint, the EMC SnapView clone application running on Dell/EMC CX3 arrays helps create backup copies of Exchange Server 2007 data in coordination with the Replication Manager agent running on the Exchange Server 2007 production servers. Although SnapView snapshots use less

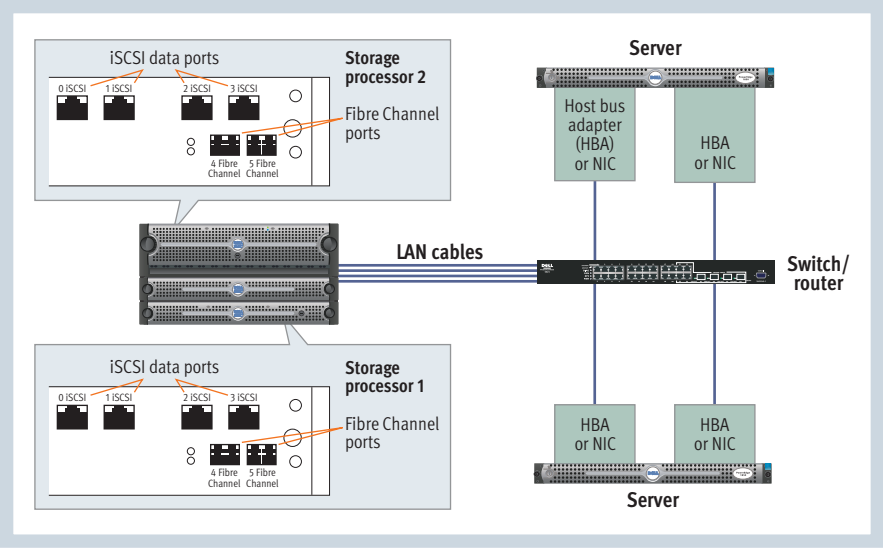

Figure 2. *iSCSI data ports for the Dell/EMC CX3-20 storage array*

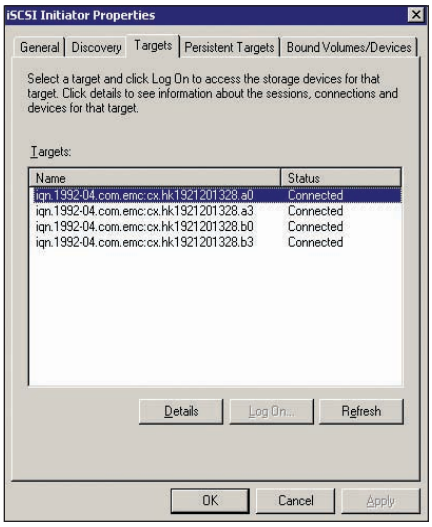

Figure 3. *iSCSI connections in Microsoft iSCSI Software Initiator*

space than SnapView clones—snapshots store only changes to production data, while clones make full copies—the appropriate choice for creating backup replicas is highly dependent on the production data change rate. Because of online maintenance and online defrag (OLM/OLD) operations, many Exchange databases can change significantly. As a result, SnapView clones can provide a higher-performance configuration than SnapView snapshots in Exchange environments. In the EMC blueprint, the EMC Replication Manager mount host mounts Exchange Server 2007 instant clones for backup and recovery of e-mail, mailbox, and database data.

Following Exchange Server 2007 installation, administrators typically should run the Microsoft Exchange Server Best Practices Analyzer (ExBPA) and follow its recommendations for performance tuning and hot-fix installation.

Network architecture

The network architecture in the EMC blueprint focuses on iSCSI technology, which allows servers to connect to storage arrays using existing IP network infrastructures. iSCSI can be cost-effective and easy to implement while still providing good performance for applications like Exchange that require low bandwidth and a high number of I/Os per second (IOPS). The small I/O size in Exchange Server 2007 also allows Exchange Server 2007 implementations using iSCSI to scale to large numbers of users by utilizing Gigabit Ethernet network interface cards (NICs) and Microsoft iSCSI Software Initiator.

Figure 2 shows the iSCSI data ports for the Dell/EMC CX3-20 array. Exchange Server 2007 and the EMC Replication Manager mount host connect to the external storage through two Intel PRO/1000 network adapters using four MPIO-based iSCSI connections. As shown in Figure 3, these connections are created using Microsoft iSCSI Software Initiator 2.0.3, with NIC0 logging in using iSCSI connections a0 and b3 and NIC1 logging in using iSCSI connections a3 and b0.

Store module

The architecture of the Store module in the EMC blueprint is based on a building block design approach, which can scale well as enterprises grow and require additional space and fault tolerance. This design helps provide scalable performance while minimizing configuration and operations complexity by dividing users into groups of 500 per building block, with a maximum mailbox size of 250 MB. Each building block logical unit (LUN) is created in a four-disk RAID-10 group to help provide fault tolerance and high performance while minimizing the number of disks required.

Figure 4 shows a building block for 500 users in two storage groups utilizing backup to disk, and Figure 5 shows a building block for 500 users in two storage groups utilizing two EMC Replication Manager clones. The databases reside on the corresponding LUNs (drive or mount point) created on the RAID-10 disk group. Log files reside on the corresponding LUNs (drive or mount point) created on the RAID-10 disk group.

In February and March 2007, EMC engineers used the Microsoft Exchange Server Jetstress tool to test one building block (four database disk spindles and four log disk spindles in shelf 0_0 of a Dell/EMC CX3-20 storage array containing 146 GB, 15,000 rpm disks). Figure 6 shows the results, which indicate that one building block can

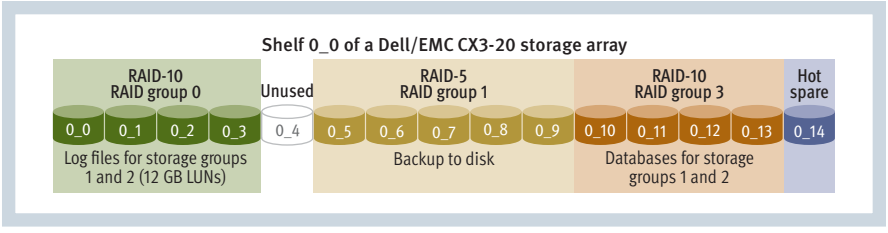

Figure 4. *Building block for 500 users in two Microsoft Exchange Server 2007 storage groups utilizing backup to disk*

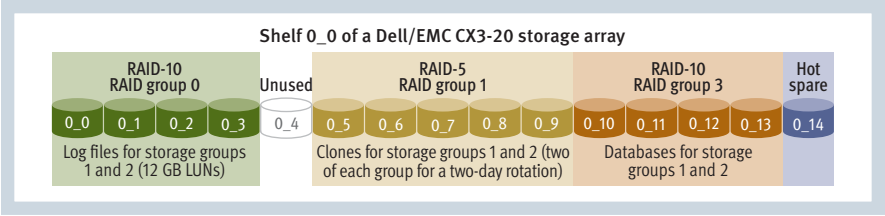

Figure 5. *Building block for 500 users in two Microsoft Exchange Server 2007 storage groups utilizing two EMC Replication Manager clones*

achieve up to 755 I/Os per second (IOPS) while staying within the Microsoft recommended database read latency limit of 20 milliseconds. The test team was able to create 725 mailboxes of 250 MB each with a user profile of 0.5 IOPS while maintaining a database read latency of 9 milliseconds and a database write latency of 5 milliseconds.

In the EMC blueprint, the Exchange Server 2007 message stores and logs on the Dell/EMC CX3-20 array are accessed through the assigned ports logged in to using Microsoft iSCSI Software Initiator. EMC PowerPath 4.6 provides LUN storage processor failover.

Backup module

The combination of Microsoft Windows Server 2003 Enterprise Edition and Microsoft Exchange Server 2007 supports Exchange data backup using VSS, a Microsoft service that enables administrators to create restartable copies of Exchange databases. Both the EMC NetWorker™ and EMC Replication Manager applications support VSS. Administrators can perform backups directly using the EMC NetWorker Module for Microsoft Exchange Server, which performs a VSS backup using either tape or SATA disk drives as the target location.

EMC Replication Manager 5.0 with SP1 provides local replication capabilities for Exchange Server 2007. The SnapView enabler software is required on the Dell/EMC CX3 storage arrays to support clone and snapshot capabilities. The EMC blueprint utilizes EMC SnapView clones to create instant backup copies of the Exchange Server 2007 Information Store. The Replication Manager agent resides on the production Exchange node and the mount host server. After the clones are mounted to the mount host server, administrators can run the Microsoft Eseutil utility using the `eseutil /k` command to validate the Exchange clone and help ensure database integrity.[2] They can then back up the clones as flat files to another medium, such as tape, SATA disk, or a disk library using EMC NetWorker, Symantec NetBackup, CommVault Galaxy, or other industry-leading backup tools. In this blueprint, the use of a mount host helps reduce performance impact on production Exchange servers when carrying out database checksum or backup operations.

The Replication Manager agent on Exchange servers takes advantage of array replication technology to create clones of Exchange Information Stores. Administrators should create these clones during non-production hours and when the process of creating clones does not conflict with other database maintenance procedures, such as OLM/OLD.

Administrators should not mount replicas back to Exchange production servers to run Eseutil. When the Exchange Mailbox server is not using a single-copy cluster (SCC) configuration, the replicas can technically be mounted back to this server without explicitly requiring a mount host; however, doing so does not offload the Eseutil workload from the Mailbox server. In SCC configurations, a mount host is explicitly required, because the LUN replica cannot be mounted back to the SCC server. Loading the Exchange administrative tools also loads Eseutil and the necessary dynamic-link library (DLL) files required for the checksum operation. Production Exchange servers must have the same service pack and hotfix revisions of the Exchange DLL files installed.

After the message stores have been checked for consistency, administrators can use EMC NetWorker to copy the clones to tape or other media. Array-based clones require clone and production LUNs to have equal block sizes, but do not require them to use the same RAID level; the EMC blueprint, for example, uses RAID-10 for the production databases and RAID-5 for clones, which helps reduce the number of disks. The number of required clone LUNs is directly proportional to the number of Exchange database LUNs: if administrators are using only one Exchange database LUN and require two clones,

Storage IOPS	Database read latency (milliseconds)	Database write latency (milliseconds)	Log write latency (milliseconds)
228	8	4	1
368	9	5	1
453	11	5	1
507	12	5	1
549	14	6	2
575	16	6	2
604	17	6	2
628	18	6	2
755	18	6	2

Figure 6. *Test results using the Microsoft Exchange Server Jetstress tool on one building block*

[2] For more information about Microsoft Exchange best practices and the Eseutil utility, visit support.microsoft.com/kb/822896.

they must create two clone LUNs that are the same size as the Exchange database LUN.

Basic backup

The basic backup architecture in the EMC blueprint uses Exchange-aware backup software such as EMC NetWorker and Microsoft NTBackup. With a 500 GB LUN, administrators can quickly perform a single nightly backup to disk and then present this backup copy to offline backup tape devices, helping avoid the tight backup windows Exchange administrators typically struggle with. Nightly backup-to-disk jobs are dependent on the tape backup completing successfully, so administrators must test the data integrity of each night's backup to disk and ensure that no backup errors (such as -1018, -1019, or -1022 errors) have occurred. These processes allow administrators to maintain the last viable full backup on disk to facilitate rapid restore times.

In February and March 2007, EMC engineers tested streaming backup-to-disk performance by using the Eseback.dll file provided with the Jetstress tool to back up an Exchange Server 2007 Information Store to the backup-to-disk LUN on shelf 0_0 of a Dell/EMC CX3-20 storage array containing 146 GB, 15,000 rpm disks, as shown in Figure 4. Figure 7 shows the results, which indicate that an Exchange Server 2007 Information Store can be backed up to disk at an approximate throughput of 50 GB/hour.

The cluster continuous replication (CCR) feature in Exchange Server 2007 helps provide redundancy for Exchange Information Stores to facilitate rapid recovery without using Exchange backups. An article addressing Exchange Server 2007 backup processes using CCR is planned for a future issue of *Dell Power Solutions*.

Advanced backup

The advanced backup architecture in the EMC blueprint uses EMC Replication Manager 5.0 software with SP1 to create local clones (within the same array) of Exchange databases and log files. Replication Manager allows administrators to create multiple clones of Exchange Information Stores, enabling them to quickly restore e-mail, mailboxes, databases, and servers following database or log file corruption. Administrators should typically maintain two clones of each storage group, which Replication Manager can automatically rotate to keep two days of backup copies available on a local array.

Database instance	Database size	Total backup time (hours:minutes:seconds)	Transfer speed
Instance1072.1	105,997 MB	2:08:41	13.73 MB/sec
Instance1072.2	106,045 MB	2:15:50	13.01 MB/sec

Figure 7. *Test results for streaming backup-to-disk performance*

EMC recommends using EMC NetWorker to back up clones to tape for off-site backup maintenance. Using two clone backup copies on a local array and tape for off-site backup enables administrators to provide additional time for Exchange OLM/OLD operations to complete. Replication Manager can also facilitate Eseutil integrity checks of Exchange databases on remote hosts, helping minimize performance impact on Exchange production servers during these checks.

EMC Assessment for Microsoft Exchange Server

Enterprises can face many challenges when upgrading to Exchange Server 2007 or implementing a new Exchange environment. EMC, which has more than a decade of experience with Exchange deployments, offers its expertise through EMC Assessment for Microsoft Exchange Server. When performing these assessments, EMC employs industry standards, best practices, and tools such as Microsoft ExBPA, Quest MessageStats, and EMC VisualSRM™ software to gather, analyze, and report on Exchange e-mail use.

EMC Assessment for Microsoft Exchange Server can help enterprises understand their existing environment before moving forward with a migration or upgrade to Exchange Server 2007 by providing an analysis of the current environment and a recommendation for the future design—with the ultimate goal of helping organizations meet their e-mail and collaboration service levels and minimize their total cost of ownership for this environment.

Blueprint for Exchange Server 2007 storage infrastructures

The EMC blueprint for Microsoft Exchange Server 2007 storage infrastructures is designed to help organizations plan for Exchange Server 2007 upgrades or new deployments. Built on Dell/EMC CX3 UltraScale series storage arrays and EMC software, this blueprint can help administrators cost-effectively manage and scale Exchange environments while providing robust backup and recovery functionality. Following the guidelines and best practices described in this article, and taking advantage of services such as EMC Assessment for Microsoft Exchange Server, can help enterprises significantly simplify the deployment of Exchange Server 2007 storage infrastructures in their data centers.

Derrick Baxter *is an EMC practice manager and Microsoft Exchange expert responsible for midsize enterprise solutions validation for Exchange Server 2003 and Exchange Server 2007.*

Suresh Jasrasaria *is an EMC senior manager for midsize enterprise solutions responsible for the solutions validation program.*

more Online
www.dell.com/powersolutions

QUICK LINKS

Dell/EMC storage:
www.dell.com/emc

EMC solutions for midsize enterprises:
www.emc.com/solutions/smb/medium

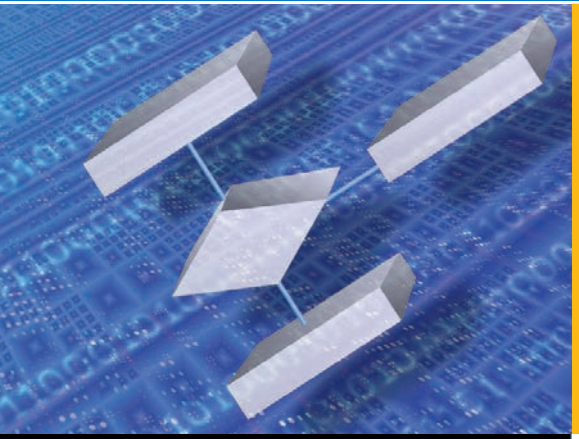

Design Considerations for **Deploying Microsoft Exchange Server 2007**

BY SUMAN KUMAR SINGH

BHARATH VASUDEVAN

Creating appropriate Microsoft® Exchange Server 2007 deployments can be challenging. The Dell Exchange 2007 Advisor online tool is designed to reduce the complexity of such deployments, recommending Exchange configurations for enterprises of all sizes based on their specific goals, mailbox design, and requirements for application and data availability.

Related Categories:

Application server sizing

Dell PowerEdge servers

E-mail technology

Microsoft Exchange Server 2007

Microsoft Windows Server 2003

Performance characterization

Visit www.dell.com/powersolutions for the complete category index.

E-mail has become an essential productivity tool, effectively replacing the telephone in many enterprises. This increased reliance on e-mail creates the need for a comprehensive e-mail system for enterprises of all sizes, not just the largest corporations. However, designing appropriate messaging systems can be a challenging, complex task that must take multiple enterprise requirements and other criteria into account.

Microsoft Exchange Server 2007 introduces several features and other variables that enterprises should understand before deployment. For example, advancements in server architecture now allow applications to be written using an expanded instruction set, enabling them to use available hardware resources more efficiently than they could previously. In legacy 32-bit operating systems, memory configurations larger than 4 GB could be installed in a system, but the management and utilization of this memory often carried a performance penalty. Some applications, such as Exchange Server 2003, were designed to execute in a 32-bit environment and did not support these large memory configurations. The latest-generation processors and chipsets, however, can implement a 64-bit instruction set, which offers direct access to system memory greater than 4 GB without a performance penalty. Because Exchange Server 2007 is designed to utilize this architecture, enterprises should take the increase in available memory into account when planning Exchange Server 2007 deployments.

To help reduce the complexity of planning Exchange Server 2007 deployments for enterprises of all sizes, Dell has performed extensive testing and characterization of Exchange Server 2007 on Dell™ PowerEdge™ servers and Dell PowerVault™ and Dell/EMC storage, and incorporated the information collected from these tests into the Dell Exchange 2007 Advisor tool, available at www.dell.com/exchange2007. This tool is designed to take basic information about an enterprise's requirements, goals, and current Exchange deployment (see Figure 1) and design a solution using standard Dell components that can meet those goals (see Figure 2). Enterprises can try out different requirements in the tool to help them understand different possible deployments before contacting their account team.

This article outlines some of the criteria used by the Dell Exchange 2007 Advisor tool, including their purpose and how they can affect the resulting recommendation. The criteria cover three major categories: enterprise goals, mailbox design, and application and data availability.

Enterprise goals

IT managers may have many reasons for upgrading to or deploying Exchange Server 2007, including a need to consolidate servers and storage or to address security, archiving and compliance, or disaster recovery requirements. To help understand these needs, the Dell Exchange 2007 Advisor tool asks enterprises to select from a list of

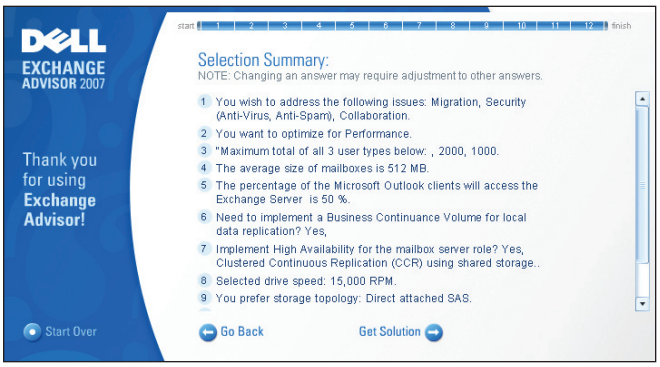

Figure 1. *Summary of sample information provided in the Dell Exchange 2007 Advisor tool*

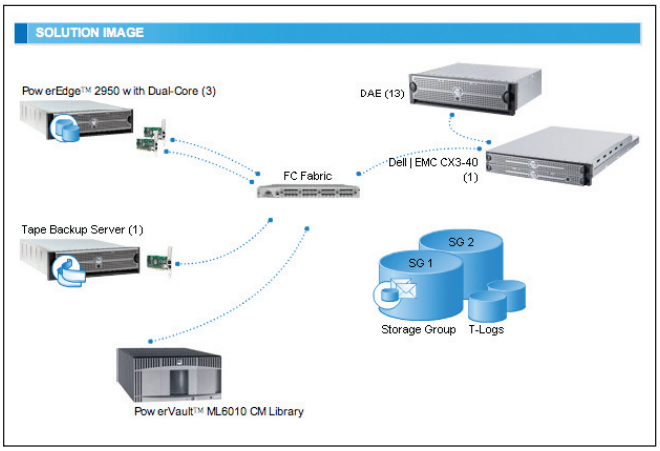

Figure 2. *Sample solution recommended by the Dell Exchange 2007 Advisor tool*

possible requirements for their Exchange Server 2007 systems. Dell has created a suite of professional services offerings for each of these requirements to help enterprises deploy, upgrade, and maintain their Exchange Server 2007 systems.

Multiple configurations may be able to satisfy a given set of criteria, but sometimes performance requirements can outweigh price considerations and vice versa. Because some high-performance components can increase the cost of a given Exchange Server 2007 system and may not be necessary for all enterprises, the Dell Exchange 2007 Advisor tool allows enterprises to select whether they want the recommended solution optimized for price or performance.

Mailbox design

Two significant factors that must be considered when planning Exchange Server 2007 deployments are the number of mailboxes and the average mailbox size, both of which enterprises can specify in the Dell Exchange 2007 Advisor tool. The tool can size up to 15,000 mailboxes with an average size of up to 2 GB.

Just as important as the number and size of the mailboxes, however, are the different types of users, who can have different impacts on memory, processor, and disk resource requirements. Proper capacity planning must take into account not only the disk resources used for mailbox data, but also those used for the transaction logs—records of messages sent and received, which are used primarily to maintain consistency and help ensure the recoverability of mailbox data. To help estimate the necessary resources, the Dell Exchange 2007 Advisor tool allows enterprises to classify their Exchange e-mail users into three categories:

- **Light users:** Each light user typically sends 5 messages and receives 20 messages in an eight-hour period.
- **Medium users:** Each medium user typically sends 10 messages and receives 40 messages in an eight-hour period.
- **Heavy users:** Each heavy user typically sends 20 messages and receives 80 messages in an eight-hour period.

RAID volume types and performance response times also help determine disk drive count and configuration. Best practices determined through Dell characterization studies have determined that a RAID-1 or RAID-10 volume is typically appropriate for transaction log hosting, while the appropriate RAID level for mailbox hosting depends on the nature of the users and the load on mailboxes. For light loads, RAID-5 can help provide the necessary performance with fewer drives than would be required for heavy loads, where the demands placed on the physical hard drives can require RAID-10 volumes.

Application and data availability

Exchange Server 2003 provides multiple options to achieve high availability for applications and data—for example, administrators can employ single-copy clustering (SCC) using shared clustered storage based on Microsoft Cluster Service to help provide both application and data availability, and many third-party applications and utilities can help provide data replication systems. Exchange Server 2007 continues to offer application availability using SCC with Microsoft Cluster Service, but also provides cluster continuous replication (CCR), which is based on the majority node set (MNS) framework within Microsoft Cluster Service. For data availability, administrators can continue to use third-party tools, but Exchange Server 2007 also allows

"The Dell Exchange 2007 Advisor tool takes advantage of extensive Dell testing to help recommend solutions based on specific requirements."

them to create business continuance volumes using local continuous replication (LCR).

The Dell Exchange 2007 Advisor tool allows enterprises to select whether they require high-availability clustering, business continuance volumes, or both. If enterprises require high availability at the application level, the tool provides options using both SCC and CCR. It also provides a framework for deploying business continuance volumes using third-party tools or LCR; this option requires additional disks, and the actual implementation choice is left to individual administrators.

Customized Microsoft Exchange Server 2007 solutions

Many different criteria affect how servers and storage should be sized and deployed in Microsoft Exchange Server 2007 messaging deployments, including enterprise goals, mailbox design, and requirements for application and data availability. The Dell Exchange 2007 Advisor tool takes advantage of extensive Dell testing and characterization of Exchange Server 2007 on Dell servers and storage to help recommend solutions based on specific requirements and existing environments, which enterprises can take into account when planning an upgrade to or new deployment of Exchange Server 2007.

Suman Kumar Singh *is a lead systems engineer on the Dell End-to-End Solutions team. He specializes in messaging systems architecture and sizing, and led the release of the Dell Exchange 2003 Advisor tool. His other interests include storage area networks, virtualization, and security, and he has published and presented several papers at industry conferences. Suman has a master's degree from the University of Texas at El Paso.*

Bharath Vasudevan *manages the Dell End-to-End Solutions team. His current interests include application performance characterization and storage technologies, and he has previously designed server hardware and led multiple clustering releases. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.*

more Online

www.dell.com/**powersolutions**

QUICK LINK

Microsoft Exchange solutions on Dell and Dell Exchange 2007 Advisor tool:
www.dell.com/exchange2007

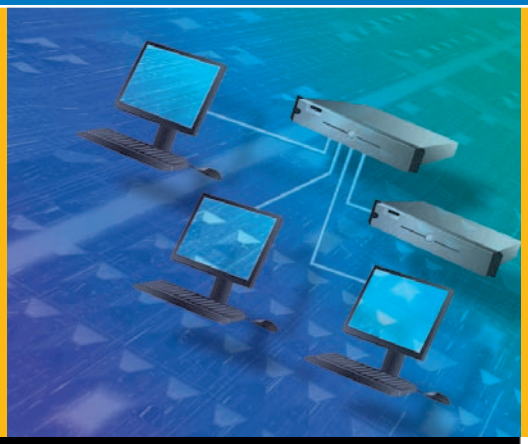

Protecting Critical Enterprise Applications with Symantec Backup Exec 11d

BY CHARLES BUTLER

Symantec® Backup Exec™ 11d for Windows Servers introduces significant backup and recovery capabilities for Microsoft® Exchange environments, including Granular Recovery Technology and Continuous Protection Server, as well as robust capabilities for systems running Microsoft SQL Server™, Microsoft Office SharePoint® Portal Server, and Oracle® software.

Related Categories:

Backup, recovery, and archiving (BURA)

Microsoft Exchange

Microsoft Office SharePoint Portal Server

Oracle

Storage software

Symantec

Visit www.dell.com/powersolutions for the complete category index.

Symantec Backup Exec 11d for Windows Servers is a leading data protection and recovery solution for the Microsoft Windows® OS. Designed for use with disk and tape media and with both 32-bit and 64-bit systems, it enables fast, comprehensive, cost-effective backup and recovery, including continuous data protection for Microsoft Exchange and Microsoft SQL Server, file servers, and desktop and notebook workstations. Its robust capabilities are designed to eliminate backup windows, allow extremely rapid data recovery, and enhance data security with strong Advanced Encryption Standard (AES) encryption.

The centralized administration of Backup Exec 11d provides scalable management for distributed backup and remote servers, and its intuitive interface and wizards help simplify data protection and recovery procedures for both small and large networks while providing the flexibility to administer storage area networks (SANs), remote offices, and departmental workgroups (see Figure 1).

Backup Exec 11d includes a comprehensive set of high-performance agents and options to help protect data on servers running Windows, Linux®, or UNIX® operating systems as well as data on desktops and notebooks. Significant enhancements to these agents and options in Backup Exec 11d help increase backup and recovery performance for Microsoft Exchange, Microsoft SQL Server, Microsoft Office SharePoint Portal Server, and Microsoft Active Directory® environments. In addition, innovative Granular Recovery Technology (GRT) enables individual document recovery from these key Microsoft applications, and the Continuous Protection Server (CPS) feature can help provide continuous data protection for Exchange and SQL Server systems. Backup Exec 11d can also provide protection for IBM® DB2, SAP®, and Oracle applications (including Oracle Recovery Manager and Real Application Clusters). And by using the Web-based Backup Exec Retrieve file recovery solution, end users can even retrieve their own files—without IT staff intervention.

This article outlines the features enabling Backup Exec 11d to protect Microsoft Exchange, SQL Server, and SharePoint Portal Server software and Oracle databases, as well as its powerful, flexible data encryption capabilities. Backup Exec 11d also supports an agent for Microsoft Active Directory and several add-on options designed to simplify management, provide additional protection for network devices, and help protect end-user desktops and notebooks (see Figure 2).

Backup Exec 11d Agent for Microsoft Exchange Server

Backup Exec has supported both Microsoft Exchange and Microsoft Windows Server® operating systems since their inception. Backup Exec 11d and the Backup Exec 11d Agent for Microsoft Exchange Server introduce GRT and CPS for Exchange environments, two key features that help address problems with traditional Exchange backups. Together, these

features help eliminate not only individual mailbox backups but also daily backups, significantly simplifying Exchange data protection.

Granular Recovery Technology

GRT can help enterprises do the following:

- Eliminate slow individual mailbox backups
- Reduce backup time and storage by performing fast, single-pass Exchange database backups
- Recover individual messages, folders, and mailboxes, or complete Exchange databases
- Reduce storage and media costs

Because this feature uses single-pass full or incremental Exchange backups, it can help dramatically reduce both the time required to back up mailboxes and storage requirements, while enabling rapid granular data recovery (see Figure 3). GRT is designed to work with or without recovery storage groups.

Continuous Protection Server

CPS can help enterprises do the following:

- Eliminate daily backups by continuously protecting Exchange data
- Automatically truncate Exchange transaction logs for automated log growth control
- Provide comprehensive disaster recovery of Exchange databases up to the latest transaction log

Administrators can use GRT-enabled backups to help protect Exchange data at the storage group and mailbox store database level while also providing granular recovery of individual messages, folders, and mailboxes from single-pass backups. Enterprises might run these traditional full or incremental backups of Exchange databases nightly using Backup Exec 11d. However, as Exchange has become increasingly critical to many enterprises, the need for data recovery beyond daily backups has also increased—and recovering data from a previous night's backup may no longer be acceptable. The CPS feature uses the same GRT-enabled technology for granular or full database recovery, but extends it by enabling these backups to occur more frequently than simple daily backups to help increase the flexibility and availability of data recovery. CPS enables Backup Exec 11d to create GRT-enabled *recovery points* of Exchange data at intervals administrators can specify in the Backup Exec 11d console.

When using CPS, administrators perform full backups once a week or once a month. Meanwhile, Backup Exec 11d continuously protects Exchange transaction logs and automatically consolidates them into easily managed recovery points to help ensure the Exchange databases are protected up to the latest complete transaction log. When administrators have

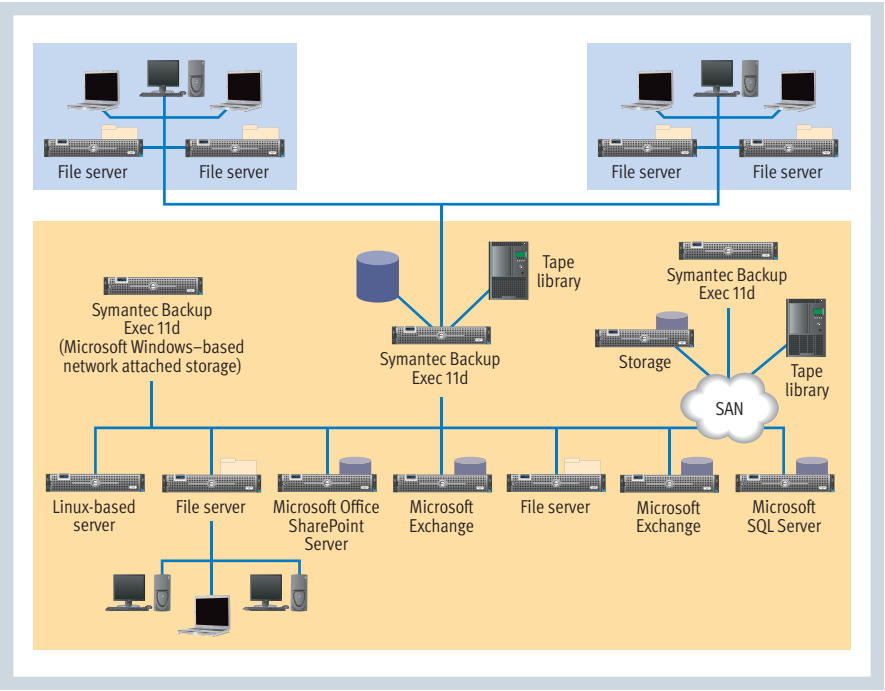

Figure 1. *Network data protection using Symantec Backup Exec 11d centralized management*

Feature	Description
Agent for Active Directory	Licensed as a separate add-on component on a per-domain-controller basis, this agent allows administrators to quickly recover from accidental deletions or changes to Microsoft Active Directory data, enabling them to granularly restore objects or attributes without performing authoritative or non-authoritative full restores and without rebooting or taking Active Directory domain controllers offline.
Central Admin Server Option	This option enables simple, centralized management of multiple Backup Exec 11d media servers. By allowing administrators to centrally manage operations, load balancing, fault tolerance, monitoring, and reporting, it helps increase both administrative efficiency and visibility into the environment.
NDMP Option	This option allows Backup Exec 11d to use the Network Data Management Protocol (NDMP) to manage filer-attached tape backup and restore operations. It is designed to provide an additional layer of protection for enterprises using NetApp and EMC® filers, and supports both filer-to-tape and filer-to-filer-to-tape backup topologies.
Desktop and Laptop Option	Designed to protect end-user workstations, this option supports continuous backups, periodic scheduled backups, and manual backups of desktops and notebooks. It is designed to allow users to easily retrieve their own data through an intuitive interface, and because it does not require a dedicated application server, it can fit easily into existing infrastructures.

Figure 2. *Add-on components for Symantec Backup Exec 11d*

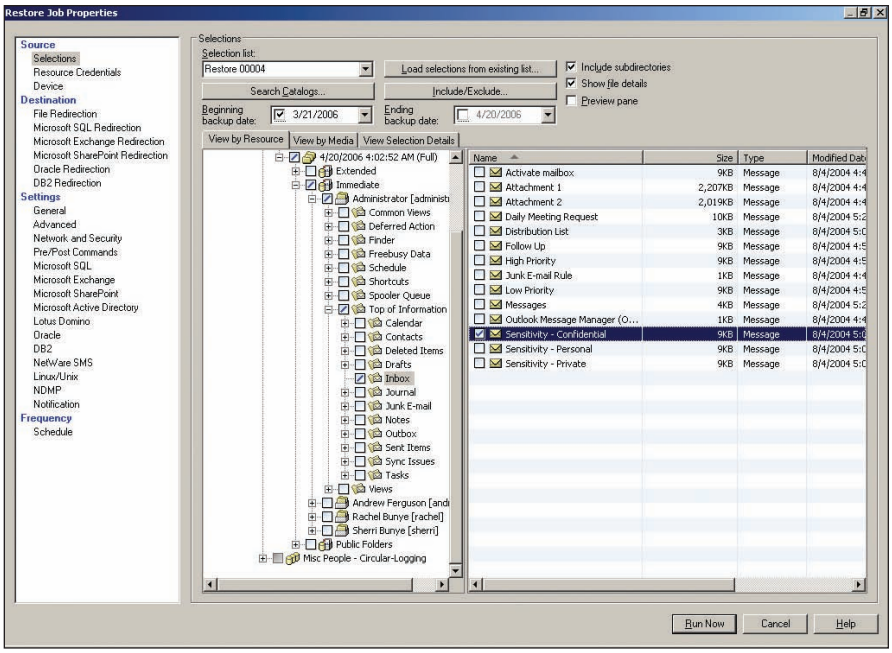

Figure 3. *Recovery of individual Microsoft Exchange messages using Backup Exec 11d Granular Recovery Technology*

enabled recovery points at intervals between the weekly or monthly full backups, they can restore individual messages, folders, and mailboxes for all Exchange components, including embedded objects and attributes, to specific times when the recovery points were created.

Backup Exec 11d Agent for Microsoft SQL Server

The Backup Exec 11d Agent for Microsoft SQL Server helps protect critical online transaction processing, online analytical processing, and e-business data to enable data recovery following application- or hardware-based corruption or failure. Designed to be flexible and easy to use, this agent provides comprehensive and customizable protection for SQL Server 7.0, SQL Server 2000, and SQL Server 2005 down to the level of individual file groups. Enterprises with backup windows too small for full backups can use this agent to perform differential as well as transaction log backups with automatic truncation.

This agent also helps simplify restoring data to another SQL Server system, supports rollback and single-pass recovery so administrators can restore databases based on a point in time rather than a specific backup job, and takes advantage of the Microsoft Virtual Device Interface to help protect SQL Server databases quickly and easily. To help resolve backup window problems, administrators can perform off-host backups or use CPS.

The Backup Exec 11d Agent for Microsoft SQL Server introduces the following features:

- **Database snapshots:** Backup Exec 11d can quickly create and maintain full database snapshots with minimal impact on SQL Server performance.
- **Copy-only backups:** Also known as out-of-band backups, copy-only backups operate as full backups but avoid disrupting future full or differential backup rotations.
- **Simplified point-in-time restores:** Administrators simply select a database and a point in time to recover, and Backup Exec 11d can assemble the necessary full, differential, and log-sequenced backups. It also can identify selection or redirection conflicts, and determine whether the database could be restored to a later date than the one specified.
- **Full-text catalog support:** Backup Exec 11d helps seamlessly protect and recover SQL Server full-text catalogs as part of standard backup and recovery jobs.
- **Concurrent instance protection for i386 and x86-64:** Backup Exec 11d can protect i386 and x86-64 SQL Server instances running simultaneously on the same server.
- **Continuous data protection:** The CPS feature enables Backup Exec 11d to continuously back up and recover SQL Server databases and file groups to any recovery point.

Backup Exec 11d Agent for Microsoft SharePoint Portal Server

The Backup Exec 11d Agent for Microsoft SharePoint Portal Server automates the many steps required to comprehensively protect a SharePoint Portal Server environment. It enables administrators to restore individual documents from full backups of SharePoint Portal Server 2003, and helps protect and restore Windows SharePoint Services, including recovery of individual documents. Protecting SharePoint Portal Server 2003 requires a Backup Exec media server as well as this agent, which includes the necessary remote agent for single-server or small server farm configurations. The agent also enables enterprises to back up a SQL Server system serving as the back end to SharePoint Portal Server without requiring a separate SQL Server agent, although protecting multiple SQL Server instances may require additional agents if those instances reside on a separate server from SharePoint Portal Server.

The Backup Exec 11d Agent for Microsoft SharePoint Portal Server supports backup and restore for SQL Server databases, document libraries, index databases, and additional metadata, and enables administrators to scale backup and recovery operations from single-server environments to large, distributed server farms. The use of server farms allows administrators to distribute the various SharePoint Portal Server components over multiple servers. The agent also allows administrators to configure SharePoint Portal Server systems independently of other servers to help customize data protection strategies.

"Symantec Backup Exec 11d for Windows Servers offers comprehensive data protection functionality designed to eliminate backup windows and dramatically simplify backup and recovery."

Backup Exec 11d Agent for Oracle on Windows and Linux Servers

Oracle databases are a foundation of many enterprise applications, so even short periods of downtime and small reductions in performance can cause significant problems. Oracle database backup and recovery is an inherently challenging process that becomes increasingly difficult as databases grow and availability demands increase, limiting the time available for backups. Many organizations use custom scripts and manual backups to protect this data; however, these methods typically do not provide the high levels of reliability and availability required for critical enterprise applications. A reliable backup and recovery solution should help eliminate downtime and provide the performance and efficiency to support these environments.

The Backup Exec 11d Agent for Oracle on Windows and Linux Servers offers nondisruptive data protection for critical Oracle9*i* and Oracle Database 10*g* databases. Advanced features such as granular protection of individual tablespaces, comprehensive application and database backup, and protection of archived redo logs and control files help enterprises protect Oracle applications and databases without taking them offline.

Backup Exec 11d data encryption capabilities

Backup Exec 11d includes 128-bit and 256-bit AES encryption capabilities that provide an additional layer of protection for sensitive data without hindering the backup or restore processes critical to safeguarding enterprise assets. This high level of encryption typically meets strict U.S. government and enterprise standards, and is designed to support both files and databases.

The Backup Exec 11d encryption capabilities help address areas of concern traditionally associated with backup encryption:

- **Cost-effectiveness:** Backup Exec 11d includes encryption at no additional cost, helping ensure that enterprises using Backup Exec 11d have access to easily encrypted backups to help safeguard important data.
- **Manageability:** Backup Exec 11d encryption is designed to be not only strong, but also easy to manage, allowing administrators to use the encrypted key management system from within the familiar Backup Exec 11d console.
- **Flexibility:** Backup Exec 11d encryption enables administrators to encrypt only the data they want—all backup jobs or only specific ones, such as those determined by automated policies or the type of storage. Using software rather than media hardware as the encryption controller allows administrators to encrypt and decrypt data regardless of the hardware platform used for backup and recovery.
- **Performance:** The flexibility of Backup Exec 11d encryption allows administrators to use it only during specific backup stages. For example, enterprises using disk-to-disk-to-tape backups can enable encryption only on the tape backup. Alternatively, those concerned about performance impact on production systems can perform fast, un-encrypted backups to secure disk locations using the included backup-to-disk technology, then configure a duplication job to run immediately after the backup or after a scheduled interval to create an encrypted backup for off-site storage, where encryption can be most critical. Using Backup Exec 11d encryption in this manner can help keep initial backup processes quick without affecting performance on production systems.

Comprehensive, easy-to-manage backup and recovery

Symantec Backup Exec 11d for Windows Servers offers comprehensive data protection functionality designed to eliminate backup windows and dramatically simplify backup and recovery. It can be used to protect data residing on Microsoft Windows–based file servers; application servers for Microsoft Exchange, SQL Server, and SharePoint Portal Server; and Oracle databases. Advanced features such as Granular Recovery Technology and Continuous Protection Server allow the recovery of individual Microsoft Exchange e-mail messages, folders, and mailboxes while helping ensure continuous data protection, and end users can even recover their own data without IT staff intervention. Deploying Backup Exec 11d can help enterprises implement a comprehensive strategy for protecting critical enterprise data in Microsoft Windows environments.

***Charles Butler** is a technical product manager in the Data and Systems Management Group at Symantec. He has a B.S. in Electrical and Computer Engineering from the University of Colorado at Boulder and an M.B.A. from St. Edward's University.*

more Online

www.dell.com/**powersolutions**

QUICK LINKS

Symantec:
www.symantec.com

Symantec Backup Exec 11d:
www.backupexec.com

Advancing Microsoft Exchange Server 2007 with 64-bit AMD Opteron™ Processors

Microsoft® Exchange Server 2007 takes a big step forward in managing large e-mail databases and offers breakthrough features such as Unified Messaging, which allows e-mail, voice mail, attachments, and faxes to be combined into a common Microsoft Outlook® in-box. Energy-efficient 64-bit AMD Opteron™ processors can accelerate Microsoft Exchange Server 2007 performance while minimizing power consumption to help keep operating costs low.

Microsoft Exchange Server 2007, released in January, is the first upgrade to Microsoft's business-critical messaging system in three years. By leveraging 64-bit x86 processors, Microsoft Exchange Server 2007 is designed to improve security and increase availability while simplifying access to a broad range of communication tools for an increasingly complex user base.

Whereas Microsoft Exchange Server 2003 was limited to 2 GB of RAM for the OS and 2 GB for users, Microsoft Exchange Server 2007 soars to a maximum 1 TB of physical memory and also supports 64 GB of virtual memory. This increased memory enhances data center efficiency by supporting more Microsoft Outlook users per server as well as larger user mailboxes than Microsoft Exchange Server 2003. Getting peak performance from Microsoft Exchange Server 2007 calls for 2 GB of RAM for the OS and 10 MB per user, so the move from Microsoft Exchange Server 2003 to Microsoft Exchange Server 2007 may require a significant hardware upgrade.

Upgrading can provide many benefits beyond increased capacity. For example, Microsoft Exchange Server 2007 includes built-in support for receiving telephone voice messages and routing those messages to Outlook in-boxes. Outlook Web Access (OWA) is an almost complete clone of the Outlook 2007 desktop client. Beefed-up security integrates antivirus, antispam, and anti-phishing technology. To help with legal and regulatory compliance, new tools set retention rules, scan and act on messages in transport, journal, and perform rich text searches across mailboxes in an organization.

AMD Opteron™ processors enable improved business value

Part of the AMD family of 64-bit processors, AMD Opteron™ processor models are designed for high performance while enabling organizations to run Microsoft Exchange Server 2007 at a low cost in power consumption and thermal load. The result can help lower total cost of ownership and advance data center efficiency by allowing additional mailboxes per server and easy expandability. These benefits derive from a long list of outstanding technologies developed by AMD, including Direct Connect Architecture with HyperTransport™ technology, AMD PowerNow!™ technology with Optimized Power Management, and Enhanced Virus Protection.[1]

Direct Connect Architecture with HyperTransport™ technology. AMD Opteron processors make a clean break from the aging x86 frontside bus architecture, a legacy design that can restrict and interrupt the flow of data, thereby creating bottlenecks. Slow data flow means slow system performance, while interrupted data flow limits system scalability. With Direct Connect Architecture, the processors, memory controller, and I/O are directly connected to the CPU. The architecture also utilizes an integrated, on-die memory controller, helping optimize memory performance and bandwidth per CPU. The controller enables

AMD Smarter Choice

AMD Opteron™ processor design for Socket F (1207)

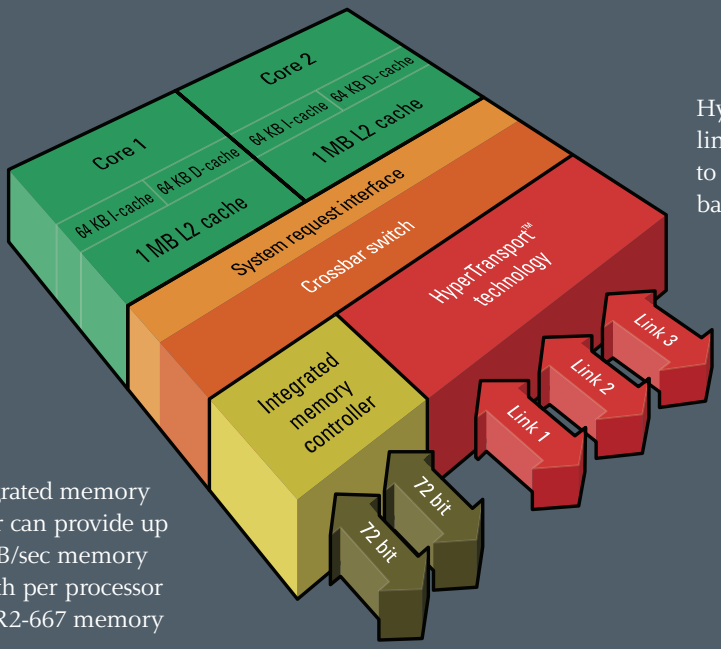

HyperTransport technology links can provide up to 24 GB/sec peak bandwidth per processor

The integrated memory controller can provide up to 10.7 GB/sec memory bandwidth per processor with DDR2-667 memory

Find out more about the advantages of running Microsoft Exchange Server 2007 on energy-efficient AMD Opteron™ processor–based Dell PowerEdge 2970 servers now!

AMD Opteron™ processors:
www.amd.com/opteron

Dell PowerEdge 2970 servers:
www.dell.com/servers

memory bandwidth to scale with the number of processors, unlike legacy designs that typically scale poorly because access to the main memory is limited by external Northbridge chips.

Direct Connect Architecture is supported by HyperTransport technology—a high-speed, bidirectional, low-latency, point-to-point communication link providing a scalable bandwidth interconnect between computing cores, I/O subsystems, and other chipsets. AMD Opteron processors support up to three coherent HyperTransport links, providing up to 24 GB/sec peak bandwidth per processor.

AMD PowerNow!™ technology with Optimized Power Management. AMD Opteron processors are built from the ground up to offer industry-leading performance per watt, helping hold down power consumption and thermal load. AMD has a track record of conserving processor power consumption, which earned the company special Energy Star recognition in 2005 from the U.S. Department of Energy and the U.S. Environmental Protection Agency.[2]

AMD PowerNow! technology works with operating systems such as the Microsoft Windows Server® 2003 OS to dynamically adjust power states, voltage, and frequencies depending on workload. AMD tests have shown that AMD PowerNow! technology can reduce CPU power consumption by up to 65 percent at 60 percent processor utilization, and by up to 75 percent during processor idle.[3] The high level of integration achieved with Direct Connect Architecture also helps reduce power consumption by eliminating external chips such as memory controllers.

Enhanced Virus Protection. The old saying still goes: Security is a process, not a product. So the strong security in Microsoft Exchange Server 2007 is not the only answer to protecting IT infrastructures. Enhanced Virus Protection,[1] working together with Windows Server, sets portions of system memory aside as "data only"—meaning code residing in these areas cannot be executed, only read from or written to. This helps prevent the spread of certain malicious viruses and worms and improves the integrity of office networks.

Dell PowerEdge 2970 servers bring it all together

Microsoft Exchange Server 2007 and AMD Opteron processors both come together in the new Dell™ PowerEdge™ 2970 server, a champion performer. As with other PowerEdge systems, the PowerEdge 2970 server offers easy expandability, robust remote server management, and high availability.

Expandability is especially important for systems running Microsoft Exchange Server 2007, because the number of Outlook users and the quantity of e-mail, calendar items, contacts, notes, and tasks they accumulate seldom goes down. IT departments can start with a single Dual-Core AMD Opteron processor with double data rate 2 (DDR2) memory in each PowerEdge 2970 server, then fill the other sockets over time. Dual-core processors can even be replaced with Quad-Core AMD Opteron processors for a maximum of eight cores. Quad-Core AMD Opteron processors, planned for introduction later this year, will be designed to maintain electrical, thermal, and socket compatibility with existing Second-Generation AMD Opteron processors, enabling a seamless upgrade path without altering the data center infrastructure.

Even IT departments not immediately shifting to Microsoft Exchange Server 2007 should consider acquiring 64-bit servers now to help prepare their hardware for a future transition. Now is also the time to start proof-of-concept evaluations. With AMD Opteron processor–based Dell PowerEdge 2970 servers, those evaluations just might show that Microsoft Exchange Server 2007 offers such gains in productivity and efficiency that there is no reason to wait.

[1] Enhanced Virus Protection (EVP) is enabled only by certain operating systems, including the current versions of the Microsoft Windows®, Linux®, Sun Solaris, and BSD UNIX® operating systems. After properly installing the appropriate OS release, users must enable the protection of their applications and associated files from buffer overrun attacks. They should consult the OS documentation for information on enabling EVP, and contact application software vendors for information regarding use of applications in conjunction with EVP. AMD strongly recommends that users continue to use third-party antivirus software as part of their security strategy.

[2] Special Recognition—Advancement of Energy-Efficient Computer Technologies, www.energystar.gov/ia/partners/pt_awards/2005_winners.pdf.

[3] These results are based on a four-socket internal AMD test platform with four AMD Opteron™ processors Model 8220 SE; four 1 GB, 667 MHz double data rate 2 (DDR2) dual in-line memory modules (DIMMs) per socket (for a total of 16 GB of memory); a 250 GB Serial ATA (SATA) hard drive running Windows Server® 2003 in 64-bit mode; and an AMD internal CPU utilization utility. For more information about AMD PowerNow! technology and processor utilization, see "Managing Data Center Power and Cooling with AMD Opteron Processors and AMD PowerNow! Technology," by Brent Kerby, in *Dell Power Solutions*, February 2007, www.dell.com/downloads/global/power/ps1q07-20070204-AMD.pdf.

Unlocking the Potential of Microsoft Exchange Server 2007

With turnkey solutions and proven expertise, the Dell Services team helps ease migration to Microsoft® Exchange Server 2007—enabling businesses to gain a competitive edge by taking advantage of new features like Unified Messaging from data center to desktop.

Brainstorming ways to translate unique messaging requirements into strategic business solutions

Although deploying a new or updated messaging platform can be complex and time-consuming, many organizations see the migration to Microsoft Exchange Server 2007 as essential for continuing business success. That's because advanced platforms like Exchange Server 2007 help provide the enhanced functionality, availability, and security required to maintain a stable communications environment. Furthermore, Exchange Server 2007 helps organizations ramp up operational efficiency through standards-based technologies that support valuable new features such as Unified Messaging. So the question has become not whether to migrate, but how to tackle the migration process to reap enhanced business benefits as quickly and efficiently as possible with minimal disruption to the existing environment.

To help businesses achieve this goal and take full advantage of emerging features, Dell Services provides an end-to-end solution that facilitates a complete migration to Exchange Server 2007, either from existing Exchange platforms or competitive applications. Each phase of the solution is designed to address a specific stage in the migration process and then guide businesses to the next step, from understanding requirements to implementing the new platform and realizing its full potential.

Understanding how the migration works

Businesses can opt to begin the migration by performing a preliminary self-assessment that results in a fast, high-level view of requirements. For organizations with large or complex environments, Dell Services can quickly engage to provide an accurate determination of a project's size and scope.

Next, working diligently alongside customers, Dell Services experts dig deep to understand the conditions imposed by an organization's unique infrastructure, and then translate those conditions into solutions-oriented recommendations designed to maximize the value of the messaging environment. These recommendations help businesses make informed decisions about the migration. Dell Services then acts on those decisions by providing detailed, customized messaging architecture and implementation plans. Finally, the Dell Services team implements the complete hardware and software solution so that organizations can reap the benefits without delay.

Businesses already running Exchange can migrate rapidly by selecting a bundled upgrade option that includes assessment, design, and delivery of a comprehensive Exchange Server 2007 solution. Because this approach is pre-scoped, the bundled option minimizes the length of the Dell Services engagement—enabling organizations to achieve an efficient upgrade while keeping costs low.

Testing the complete hardware and software implementation so organizations can benefit right away

Getting up to speed

Following a successful migration, the Dell Services team can help organizations unlock the benefits of significant new capabilities such as Unified Messaging, which enables the consolidation of key messaging applications and technologies. With Unified Messaging, critical communication avenues—including e-mail, voice mail, and fax—converge into the Microsoft Outlook in-box, giving businesses the potential to increase mobility, streamline workflow, improve employee productivity, and cut costs. (For Unified Messaging solutions from Dell Services, see the sidebar in this article, "Tying It All Together.")

In addition, Exchange Server 2007 leverages enhanced data replication capabilities to increase system availability, so organizations can keep vital communications running smoothly while minimizing tape backups. Exchange Server 2007 also enriches security through enhanced antivirus and antispam functionality, giving businesses heightened protection against messaging-related threats along with advanced e-mail filtering and scanning.

Delivering exceptional value

Given the complicated nature of a messaging platform migration, working with an experienced, savvy implementation partner helps ensure success. The Dell Services team has completed more than 1,000 Exchange migrations. During these projects, Dell Services also migrated more than 5 million Exchange and Microsoft Active Directory® directory service users.

Dell Services experts draw on a comprehensive background in deploying Microsoft products, as well as in supporting standards-based components and technologies. As a result, the Dell Services team offers organizations an elite level of engineering expertise that helps drive the migration process completely, from the data center to the desktop. This is a core advantage particularly when dealing with emerging features like Unified Messaging, which affects servers, storage, and PCs. Organizations need to understand and capitalize on strategic capabilities that can help achieve a competitive business advantage.

Finally, Dell Services stands out in terms of project accountability and management, offering an optimized service delivery process that encompasses a comprehensive migration solution with a single point of contact for all issues. As a result, organizations can enjoy a rapid, hassle-free migration that puts them on the fast track to realizing the full value of the Exchange Server 2007 messaging platform.

Tying It All Together

Without question, a successful migration to Microsoft Exchange Server 2007—one that unlocks the full potential of Unified Messaging functionality—demands specific services and expertise. Organizations must meet certain hardware requirements and will likely need detailed guidance to determine the optimal way to build, deploy, or integrate existing data center components to get the most value from the new messaging environment.

Dell Services helps simplify the process with a turnkey Unified Messaging solution that extends the Exchange Server 2007 migration to include a move to voice over IP (VoIP) for messaging. Leveraging the experience and proficiency of Dell Services, organizations can seamlessly integrate the Exchange Server 2007 platform with a VoIP network to help ensure a robust and scalable foundation for unified communications, even as the infrastructure and technology continue to evolve.

more Online

www.dell.com/**powersolutions**

QUICK LINK

To find out how Dell Services can assist your business in a migration or upgrade to Microsoft Exchange Server 2007, visit **www.dell.com/services/exchange2007**.

Comparing Virtualization Performance: Dell PowerEdge and HP ProLiant Servers

BY TODD MUIRHEAD
DAVE JAFFE, PH.D.
TERRY SCHROEDER

To demonstrate the advantages of two-socket servers over four-socket servers for virtualization, Dell engineers tested a two-socket Dell™ PowerEdge™ 2950 server with quad-core Intel® Xeon® processors alongside a four-socket HP ProLiant DL585 G2 server with dual-core AMD Opteron™ processors in a virtualized environment based on VMware® Infrastructure 3.

Related Categories:

Dell PowerEdge servers

Dell/EMC storage

Microsoft SQL Server 2005

NetBench

Novell SUSE Linux

Performance characterization

Power and cooling

Scalable enterprise

Visit www.dell.com/powersolutions for the complete category index.

Taking advantage of server virtualization has often meant deploying farms of four-socket servers, which had the necessary processing power to handle a large number of virtual machines (VMs) and allowed enterprises to efficiently consolidate their physical servers. However, this type of environment also typically required sophisticated load balancing and management to achieve appropriate levels of redundancy and performance across the VMs, and acquiring and licensing software for additional servers could be expensive.

The recent introduction of quad-core processors has enabled two-socket servers to achieve similar performance to four-socket servers, but at a lower cost. Both a two-socket server with quad-core processors and a four-socket server with dual-core processors have eight total processing cores available, but two-socket servers typically cost less than four-socket servers and can require less-expensive software licensing. VMware virtualization software, for example, is currently licensed per socket; running this software on a two-socket server costs approximately half as much as running it on a four-socket server, regardless of the number of available processing cores. In addition, because of the multi-process nature of running multiple VMs on a single server, additional cores allow virtualized environments to scale well at the node level.

To compare the virtualization performance of two-socket and four-socket servers, in December 2006 Dell engineers tested a two-socket Dell PowerEdge 2950 server with quad-core Intel Xeon X5355 processors alongside a four-socket HP ProLiant DL585 G2 server with dual-core AMD Opteron 8220 SE processors, both running VMware Infrastructure 3 software. They then extrapolated these results to compare a farm of three PowerEdge 2950 servers with a farm of two ProLiant DL585 G2 servers. The results showed that the PowerEdge 2950 farm can provide an average of 44 percent higher performance, an average of 58 percent more performance per watt, and an average of 95 percent better price/performance compared with the ProLiant DL585 G2 farm.

Hardware test configuration

The Dell PowerEdge 2950 is a dual-socket server that supports Intel Xeon 5000, 5100, and 5300 series processors. The Dell test team configured a PowerEdge 2950 with two quad-core Intel Xeon X5355 processors at 2.66 GHz. Because the Intel Xeon X5355 is essentially two dual-core Intel Xeon 5160 processors combined, it has an 8 MB level 2 (L2) cache, with 4 MB shared by each set of two cores, along with a 1.333 GHz front-side bus. This server was configured with 16 GB of RAM.

The HP ProLiant DL585 G2 is a four-socket server that supports AMD Opteron 8000 series processors. The test team configured a ProLiant DL585 G2 with four dual-core AMD Opteron 8220 SE processors at 2.8 GHz, each with a 1 MB L2 cache per core. In this server, each processor is coupled with RAM and connected to the others through a 1 GHz

HyperTransport bus; this direct connection keeps the processors close to the memory, and in many cases can increase access speeds. The server used in the test was configured with 32 GB of RAM.

Figure 1 summarizes the server configuration used in the tests along with the pricing of each configuration. The hardware prices include the cost of a QLogic QLE2462 Fibre Channel host bus adapter (HBA) for both the PowerEdge 2950 and ProLiant DL585 G2. Because this HBA is not available as an option with the ProLiant DL585 G2 on the HP Web site, the US$1,219 price for the QLogic adapter listed on the Dell Web site was used for both systems.

Both the PowerEdge 2950 and ProLiant DL585 G2 were connected to a storage area network (SAN) with dual-port QLogic QLE2462 PCI Express HBAs, utilizing storage on a Dell/EMC CX3-80 array with twenty 146 GB, 15,000 rpm disks. The three types of VMs—each running a different workload, as described in the "Test workloads: Microsoft SQL Server 2005, SUSE LAMP, and NetBench" section in this article—were spread across the 20 disks on each storage array. These disks were divided into four 5-disk (4+1) RAID-5 logical units (LUNs). The three types of VMs were evenly divided across the LUNs so that a quarter of each type were on each LUN. Figure 2 summarizes the storage configuration used in the test environment.

	HP ProLiant DL585 G2	Dell PowerEdge 2950
Processors	Four dual-core AMD Opteron 8220 SE processors at 2.8 GHz with one 1 MB L2 cache per core	Two quad-core Intel Xeon X5355 processors at 2.66 GHz with one shared 4 MB cache for each set of two cores
HyperTransport or frontside bus speed	1 GHz	1.333 GHz
Memory	32 GB (sixteen 2 GB, 667 MHz PC2-5300 double data rate 2 dual in-line memory modules [DIMMs])	16 GB (eight 2 GB, 667 MHz fully buffered DIMMs)
Internal disks	Two Serial Attached SCSI (SAS) 73 GB, 15,000 rpm drives	Two SAS 146 GB, 15,000 rpm drives
NICs	Two internal 10/100/1,000 Mbps NICs	Two internal 10/100/1,000 Mbps NICs
Disk controller	HP Smart Array P400	PowerEdge Expandable RAID Controller (PERC) 5/i
HBA	QLogic QLE2462	QLogic QLE2462
Virtualization software	VMware ESX Server 3.0.1	VMware ESX Server 3.0.1
Hardware price as configured	US$28,088	US$13,498
VMware software license price	US$13,914	US$7,188
Total price (hardware plus VMware software)	US$42,002	US$20,686

Note: All prices are given as listed at www.hp.com and www.dell.com on December 18, 2006.

Figure 1. *Configurations and pricing for the HP ProLiant DL585 G2 and Dell PowerEdge 2950 servers used in the test environment*

Two-socket Dell PowerEdge servers

The test team chose the Dell PowerEdge 2950 as the two-socket Dell server because it provides a good balance of density and number of PCI slots for most virtualized workloads. Two other two-socket Dell servers—the PowerEdge 1950 and PowerEdge 2900—provide performance similar to the PowerEdge 2950, but with differences in form factor, number of PCI slots, maximum memory configuration, and amount of rack space consumed. Figure 3 summarizes the different hardware features of these three servers.

Some enterprises may prefer one of these other servers if they require more network interface cards (NICs), higher server density, or more RAM than the PowerEdge 2950 provides. For example, I/O diversity requirements, which are most common at the NIC level, might make the six PCI slots in the PowerEdge 2900 an essential feature. If density is the goal, then the 1U PowerEdge 1950 might be the most appropriate choice. All three platforms share a common system image and can deliver similar performance.

Because VMs are essentially given a processor core as a virtual processor, the same number and type of multi-virtual-processor VMs can be deployed in a similar way on any of these two-socket Dell servers with quad-core processors or on the four-socket HP ProLiant DL585 G2 server with dual-core processors.

Virtualization test platform

The performance tests used VMware Infrastructure 3 as the virtualization platform;

TALK BACK

Tell us how the Dell Enterprise Technology Center can help your organization better simplify, utilize, and scale enterprise solutions and platforms. Send your feedback and ideas to enterprise_techcenter@dell.com.

Controller cache	10,384 MB (3,072 MB write, 7,312 MB read)
Fibre Channel speed	Fibre Channel 4 (FC4)
Disk enclosures	Four DAE3P disk array enclosures
Disks	Twenty 146 GB, 15,000 rpm disks
LUNs	Four 5-disk RAID-5 LUNs
Software	EMC® Navisphere® Manager and Access Logix™ software

Figure 2. *Configuration of the Dell/EMC CX3-80 storage used in the test environment*

this package includes ESX Server 3 and VirtualCenter 2 as well as features such as load balancing and VMware High Availability (VMware HA). ESX Server allows multiple VMs to run simultaneously on a single physical server. Each VM runs its own OS, which in turn has its own set of applications and services. Because ESX Server isolates each VM from other VMs on the same physical server just as physical systems are isolated from one another, administrators have flexibility in using ESX Server to run different types of applications and operating systems at the same time. VirtualCenter 2 enables administrators to consolidate control and configuration of ESX Server systems and VMs, which can improve management efficiency in large environments.

Both the Dell PowerEdge 2950 and the HP ProLiant DL585 G2 used ESX Server 3.0.1 and were managed by a VirtualCenter 2.0.1 console. All VMs were first tested on the PowerEdge 2950, and then migrated to the ProLiant DL585 G2 and retested.

Test workloads: Microsoft SQL Server 2005, SUSE LAMP, and NetBench

To compare the relative performance of the Dell PowerEdge 2950 and HP ProLiant DL585 G2 servers, the test team ran three workloads on each server: the Microsoft® SQL Server™ 2005 database platform with an online transaction processing (OLTP) workload, the Novell® SUSE® Linux® Enterprise Server OS with a LAMP (Linux, Apache, MySQL, PHP) stack, and the Microsoft Windows Server® 2003 OS with NetBench 7.03. To simulate how enterprises typically run applications on VMs using ESX Server in a production environment, the test team increased the number of VMs until processor utilization for the entire physical server was as close to 85 percent as possible, with all tests within a range of 84 to 86 percent—a reasonably high level of usage that still allows for workload spikes. The test team calculated utilization levels by averaging the values from the esxtop utility run on the ESX Server service console during each test.

Each workload ran simultaneously on multiple VMs under the same load. By keeping all settings on the VM and driver systems identical and then observing how many VMs could be run simultaneously, the test team was able to measure how many VMs each physical server could support as well as the total throughput for that workload. Figure 4 shows the configuration for each type of VM in the test environment.

Microsoft SQL Server 2005. On the SQL Server 2005 VMs, the test team installed 32-bit versions of Microsoft Windows Server 2003 Release 2 (R2) Enterprise Edition and SQL Server 2005 with Service Pack 1 (SP1).[1] The SQL Server version of the Dell DVD Store database was loaded into SQL Server 2005 using the scripts provided with the DVD Store download to create the medium-size database. The complete DVD Store application code, including SQL Server and MySQL versions, is freely available for public use under the GNU General Public License (GPL) at linux.dell.com/dvdstore. The DVD Store database simulates the database back end of a simple Web-based storefront. The database size is small (approximately 1 GB) and representative of a database used for development or testing.

To simulate a load on the VMs, the test team used the DVD Store driver program, which is included in the DVD Store download. Each SQL Server 2005 VM was driven by four threads of the driver application with a 20-millisecond delay.

	PowerEdge 1950	PowerEdge 2900	PowerEdge 2950
Rack size	1U	5U	2U
PCI slots	2	6	3
Maximum memory	32 GB	48 GB	32 GB

Figure 3. *Hardware comparison of three models of two-socket Dell PowerEdge servers*

Workload	Memory	Disk	Number of virtual NICs	Number of virtual processors
Microsoft SQL Server 2005	512 MB	10 GB	1	1
SUSE LAMP	1,024 MB	10 GB	1	1
NetBench	512 MB	10 GB	1	1

Figure 4. *Configurations for the virtual machines used in the test environment*

[1] The use of Microsoft SQL Server 2005 in these tests does not indicate that Dell or Microsoft has tested or certified SQL Server with VMware virtualization software. As described at support.microsoft.com/?kbid=897615, Microsoft typically does not support problems with Microsoft operating systems or applications that run on VMs using non-Microsoft virtualization software unless the same problem can be reproduced outside the VM environment.

Sometimes, having the best cards has nothing to do with luck.

Don't gamble with your graphics. Optimized and certified for CAD and DCC applications, the ATI FireGL™ 3D graphics accelerator from AMD is the essential workstation graphics card — delivering increased performance, reliability, and value.

AMD
Smarter Choice

Featured in Dell Precision™ Workstations. Visit ati.amd.com/FireGL for more information.

SUSE LAMP. For the LAMP workload, the test team installed 32-bit versions of Novell SUSE Linux Enterprise Server 9, Apache 2, and MySQL 5 on a VM. The MySQL version of the DVD Store application was loaded into MySQL 5, and the PHP version of the DVD Store Web tier application was set up on Apache. In this setup, the Web tier and the database tier ran on the same VM to create a complete LAMP stack.[2]

The driver for the LAMP stack differs from the driver used in the SQL Server testing in that it sends HTTP requests and receives HTML code returned from the Apache/PHP layer, whereas the SQL Server driver communicates directly with the database. However, the LAMP workload measures the same parameters: total orders per minute (OPM) handled by the application, and average response time experienced by the simulated customers. Each SUSE LAMP VM was driven by a single thread of the driver program with a 20-millisecond delay.

NetBench. NetBench 7.03, developed by *PC Magazine,* is a benchmark tool designed to simulate a file server workload. The program creates and accesses a set of files according to predefined scripts. NetBench is typically run with an increasing number of client engines running against a single server to measure how much throughput (in megabits per second) can be achieved with a given number of connections.

The NetBench VMs were installed with the 32-bit version of Microsoft Windows Server 2003 R2 Enterprise Edition. To determine how many VMs could run on an ESX Server host, the test team increased the number of VMs and the number of client engines at the same rate until the processor utilization on the ESX Server host reached 85 percent. NetBench 7.03, with the included standard DiskMix script, was used with a 0.6-second think time to connect two client engines to each VM.[3] This simulates multiple file servers on the same ESX Server host, similar to a file server consolidation scenario. The driver systems on which the client engines ran had mapped drives to the test VMs. In NetBench the test directories path file was modified so that as successive client engines were added, they would use the next drive letter, which corresponded to the next VM.

Test results

The test team first ran the VMs on the two-socket Dell PowerEdge 2950 server in successive tests, adding VMs in each round as described in the "Test workloads: Microsoft SQL Server 2005, SUSE LAMP, and NetBench" section in this article, to measure two elements of virtualization performance: the number of VMs each server could support and the aggregate performance those VMs could achieve. Next, they cold migrated the VMs to the four-socket HP ProLiant DL585 G2 server and repeated the tests. A power meter attached to the servers measured the actual power consumption (in watts) during these tests. The test team then combined the performance and power consumption results with the system costs to calculate performance per watt and price/performance.

	HP ProLiant DL585 G2			Dell PowerEdge 2950		
Workload	Number of VMs	Performance	Power consumption	Number of VMs	Performance	Power consumption
Microsoft SQL Server 2005	36	31,729 OPM	719 W	32	29,346 OPM	449 W
SUSE LAMP	46	10,093 OPM	743 W	44	9,852 OPM	447 W
NetBench	52	1,028 Mbps	735 W	42	1,001 Mbps	444 W

Figure 5. *Workload performance and power consumption results for each server in the test environment*

Performance and power consumption

The differences in the number of VMs and the associated performance metrics—OPM for SQL Server 2005 and SUSE LAMP and megabits per second for NetBench—indicated the relative differences in performance. The test team calculated the performance results for the SQL Server 2005 and SUSE LAMP VMs by totaling the OPM from all the VMs running in the test environment; NetBench provides the megabits-per-second metric as part of the results displayed at the end of a test. Figure 5 summarizes the performance and power consumption results for the three workloads on each server.

The test team arrived at the power consumption results by logging the power readings on the meter while the test was running and then averaging those measurements. Because power consumption can be heavily dependent on system load and configuration options such as the amount of RAM, number of PCI adapters, and number of internal disks, the test team kept these elements as equal as possible between the systems based on the number of sockets.

Scaling out a virtualization server farm is different from scaling out a server cluster. Some types of database clusters, for example, coordinate individual cluster nodes to maintain database integrity, and the clustering software that manages this coordination requires additional performance overhead. The servers in a VMware Infrastructure 3–based farm, in contrast, are each running multiple VMs that are not coordinated or interdependent on one another; each server operates its VMs independently of the other servers in the farm. Because of how this type of virtualized farm operates, performance and power consumption figures for a

[2] The LAMP stack has been fully documented in "MySQL Network and the Dell PowerEdge 2800: Capacity Sizing and Performance Tuning Guide for Transactional Applications," by Todd Muirhead, Dave Jaffe, and Nicolas Pujol, Dell Enterprise Product Group, April 2005, www.dell.com/downloads/global/solutions/mysql_network_2800.pdf.

[3] The NetBench client driver systems were two PowerEdge 6650 servers with four Intel Xeon processors at 2.8 GHz, nine PowerEdge 1855 servers with two dual-core Intel Xeon processors at 2.8 GHz, and two PowerEdge 1950 servers with two dual-core Intel Xeon 5160 processors at 3.0 GHz. All client driver systems ran Windows Server 2003 and had 8 GB of RAM and Intel Gigabit Ethernet adapters. The NetBench client driver systems and ESX Server hosts were connected to a Dell PowerConnect™ 5224 Gigabit Ethernet switch. The NetBench controller ran Windows Server 2003 Enterprise Edition and used an Intel Gigabit Ethernet adapter.

		Two HP ProLiant DL585 G2 servers	Three Dell PowerEdge 2950 servers	Dell PowerEdge advantage
SQL Server 2005	Number of VMs	72	96	33%
	Performance	63,458 OPM	88,038 OPM	39%
SUSE LAMP	Number of VMs	92	132	43%
	Performance	20,186 OPM	29,556 OPM	46%
NetBench	Number of VMs	104	126	21%
	Performance	2,056 Mbps	3,003 Mbps	46%
Average power consumption for all three workloads		1,465 W	1,340 W	9%
Hardware price as configured		US$56,176	US$40,494	39%
VMware software license price		US$27,828	US$21,564	29%
Total price (hardware plus VMware software)		US$84,004	US$62,058	35%

Figure 6. *Performance, power consumption, and price extrapolation for two HP ProLiant DL585 G2 servers and three Dell PowerEdge 2950 servers*

single server can be extrapolated to multiple servers simply by multiplying the results.

The test team performed this type of extrapolation to extend the results shown in Figure 5 to multiple servers. Figure 6 shows how a farm of three PowerEdge 2950 servers would compare to a farm of two ProLiant DL585 G2 servers in terms of performance, power consumption, and price. The three PowerEdge 2950 servers outperform the two ProLiant DL585 G2 servers in OPM by an average margin of 44 percent, while the ProLiant DL585 G2 servers cost 35 percent more as configured.

Figure 7 shows SQL Server 2005 OPM performance and power consumption results if an entire industry-standard 42U rack were fully populated by PowerEdge 2950 or ProLiant DL585 G2 servers. Up to 21 PowerEdge 2950 servers can fit into this type of rack, but because the ProLiant DL585 G2 takes up about twice as much space as the PowerEdge 2950, only 10 of these servers can fit into such a rack. The comparison, which maintains the 2:3 server ratio in Figure 6, shows that in addition to allowing increased server density, the PowerEdge 2950 servers can consistently outperform the ProLiant DL585 G2 servers while using less power.

Performance per watt and price/performance

Performance-per-watt calculations help demonstrate server efficiency. Figure 8 shows performance per watt for all three workloads on individual Dell PowerEdge 2950 and HP ProLiant DL585 G2 servers based on the results shown in Figure 5. On average, the PowerEdge 2950 delivered 58 percent more work per watt than the ProLiant DL585 G2.

In addition to providing better virtualization performance while using less power than the ProLiant DL585 G2, the PowerEdge 2950 in this test also cost less: a single ProLiant DL585 G2 server typically costs more than twice as much as a PowerEdge 2950 as configured in these tests, and two ProLiant DL585 G2 servers typically would cost 35 percent more than three PowerEdge 2950 servers. Software licensing is an important factor in these costs: because Virtual Infrastructure 3 is licensed by socket, licensing a four-socket ProLiant DL585 G2 would cost almost twice as much as licensing a two-socket PowerEdge 2950, and licensing eight sockets across the two ProLiant DL585 G2 servers would cost 29 percent more than licensing six sockets across the three PowerEdge 2950 servers. As shown in Figure 9, on average, the ProLiant DL585 G2 cost almost twice as much per unit of work as the PowerEdge 2950.

Price in a production environment

An additional factor enterprises should keep in mind is that the three-server Dell PowerEdge 2950 and two-server HP ProLiant DL585 G2 environments illustrated by these tests would typically not be suitable for a production environment because of their availability constraints: with each server operating with a processor utilization of around 85 percent, if a node failed, the

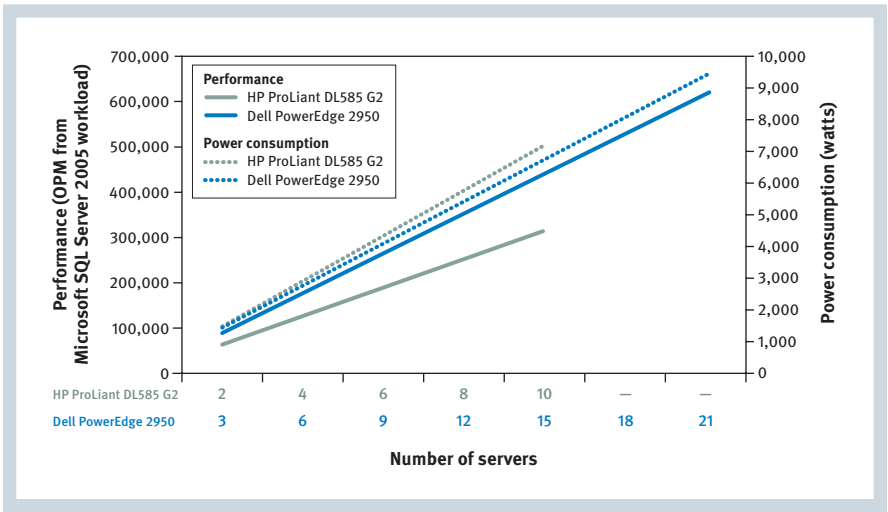

Figure 7. *Performance and power consumption extrapolation for HP ProLiant DL585 G2 servers and Dell PowerEdge 2950 servers on a single 42U rack*

remaining node or nodes would not be able to take over its workload. Appropriate load balancing would require one additional node in each setup. When this requirement is taken into account, the three ProLiant DL585 G2 servers would cost 52 percent more than the four PowerEdge 2950 servers, rather than the previous 35 percent, as shown in Figure 10.

Efficient, scalable virtualization

One barrier to the adoption of two-socket servers over four-socket servers in virtualized environments is the perceived complexity of managing additional nodes. But the VMware Infrastructure 3 management stack, and in particular the Distributed Resource Scheduler feature, helps provide simplified, policy-based management of pooled resources, which helps greatly reduce administrative overhead. Using this management system enables enterprises to take advantage of efficient, cost-effective two-socket nodes rather than four-socket nodes.

And as the tests and calculations in this article demonstrate, an environment of two-socket Dell PowerEdge 2950 servers with quad-core processors can provide increased performance with lower power consumption and lower costs than one consisting of four-socket HP ProLiant DL585 G2 servers with dual-core processors. Three PowerEdge 2950 servers provided an average of 44 percent higher performance, 58 percent higher performance per watt, and 95 percent better price/performance than two ProLiant DL585 G2 servers. These results illustrate the level of performance, scalability, and cost-effectiveness offered by two-socket Dell servers over comparable four-socket servers in virtualized data center environments.

	HP ProLiant DL585 G2	Dell PowerEdge 2950	Dell PowerEdge advantage
Microsoft SQL Server 2005	44.1 OPM/watt	65.4 OPM/watt	48%
SUSE LAMP	13.6 OPM/watt	22.0 OPM/watt	62%
NetBench	1.4 Mbps/watt	2.3 Mbps/watt	64%
Average			58%

Figure 8. *Performance-per-watt test results*

	HP ProLiant DL585 G2	Dell PowerEdge 2950	Dell PowerEdge advantage
Microsoft SQL Server 2005	US$1.32/OPM	US$0.70/OPM	89%
SUSE LAMP	US$4.16/OPM	US$2.10/OPM	98%
NetBench	US$40.86/Mbps	US$20.67/Mbps	98%
Average			95%

Figure 9. *Price/performance test results*

	Three HP ProLiant DL585 G2 servers	Four Dell PowerEdge 2950 servers	Dell PowerEdge advantage
Hardware price as configured	US$84,264	US$53,992	56%
VMware software license price	US$41,742	US$28,752	45%
Total price (hardware plus VMware software)	US$126,006	US$82,744	52%

Figure 10. *Price comparison when including an additional server to meet production availability requirements*

Todd Muirhead *is a senior engineering consultant on the Dell Enterprise Technology Center team. Todd has a B.A. in Computer Science from the University of North Texas and is Microsoft Certified Systems Engineer + Internet (MCSE+I) certified.*

Dave Jaffe, Ph.D., *is a senior consultant on the Dell Enterprise Technology Center team who specializes in cross-platform solutions. He has a B.S. in Chemistry from Yale University and a Ph.D. in Chemistry from the University of California, San Diego.*

Terry Schroeder *is an enterprise technologist in the Advanced Systems Group at Dell, supporting Dell field system consultants and engineers by communicating Dell systems management products and initiatives to customers. Terry has a B.S. in Social Sciences and an M.S. in Library Science and Information Management from Emporia State University.*

more Online

www.dell.com/powersolutions

QUICK LINKS

Dell Enterprise Technology Center:
www.dell.com/techcenter

Dell and VMware:
www.dell.com/vmware

What could your enterprise do with unparalleled flexibility and scalability?

DELL | EMC²

DELL/EMC DELIVERS A FULL PORTFOLIO OF MANAGEABLE, AFFORDABLE AND HIGHLY CAPABLE IP STORAGE SOLUTIONS

The unparalleled flexibility and scalability you've been waiting for is here. Dell/EMC iSCSI solutions offer an ideal choice for a remote server environment, virtualization or limited budgets.

By using standard IP switches to link the storage array to servers, Dell/EMC arrays can provide extreme versatility for your business. The CX-3 dual-mode arrays create integration between your Fibre-Channel SANs and IP networks, enabling the use of Fibre-Channel in areas where performance needs are most critical and iSCSI for less critical or remote servers. Offering both Fibre-Channel and iSCSI connections on the same array can help streamline costs and IT efficiency by consolidating data management and security within one environment. You can use each separately, or combine iSCSI and Fibre-Channel technologies for the ultimate power and flexibility.

BUSINESS SOLUTIONS FOR YOUR ENTERPRISE

Simple and Affordable SAN Solution

Dell/EMC AX150 and AX150i Storage Platform:

- Either Fibre-Channel or iSCSI host connectivity for flexible connectivity options
- Excellent for storage consolidation for up to 10 servers and 6TB of data
- Integrated functionality and ease-of-use at an entry-level price
- Includes the easy-to-use software you need including Navisphere® Express, SnapView™ Express and PowerPath®
- High availability via redundant architecture, integrated snapshot software and other safeguards

Flexible, Highly Scalable Enterprise SAN

Dell/EMC CX3-10, CX3-20, CX3-40 Storage Platforms:

- Both Fibre-Channel and iSCSI host connectivity for maximum deployment flexibility and highly efficient storage consolidation
- End-to-end 4 Gb/s capability for outstanding bandwidth and application performance
- Easy to deploy, manage and upgrade for excellent investment protection
- Comprehensive software packages including EMC's robust suite of array management software
- High availability and scalability for even your most mission critical data

Please contact your Dell sales representative about a professional file assessment. CALL 877.671.3355 toll free www.dell.com/SAN

Implementation Study: Dell IT Scales Supply Chain Management with Oracle RAC 10*g*

BY DAVE JAFFE, PH.D.
TODD MUIRHEAD
TIONG TEY
RAVEENDRA AVUTU

When the expensive proprietary servers running the Dell supply chain management systems had reached their limits, the Dell IT group migrated to cost-effective, standards-based Dell™ PowerEdge™ servers running Oracle® Real Application Clusters 10*g*. This architecture helped enhance database performance while providing scalability for future growth.

Related Categories:

Database

Dell Enterprise Technology Center

Implementation study

Oracle

Scalable enterprise

Visit www.dell.com/powersolutions for the complete category index.

The Dell supply chain management (SCM) database systems handle key business functions that support worldwide manufacturing operations, including the efficient Dell inventory management model and fast, direct delivery of computers, accessories, parts, and supplies. These systems must be designed for reliability and cost-effective scalability: a failure can cost thousands of dollars per minute in factory downtime, and the SCM systems must be able to handle increasing workloads as the company grows.

When Dell was a smaller company than it is now, before the development of powerful, industry-standard servers of the type that Dell manufactures, the Dell IT group ran its SCM database applications on large, expensive, proprietary servers based on the UNIX® OS. However, as the company grew, servers lacking the necessary capacity had to be replaced with even larger, more powerful servers. And because the servers were not redundant, updating a single server often required shutting down entire systems.

The increased performance of industry-standard Dell PowerEdge servers, however, has enabled Dell IT to create cost-effective, highly scalable systems using Oracle Real Application Clusters (RAC) 10*g*. By sharing a large database across multiple PowerEdge servers, Dell IT can easily deploy additional low-cost servers when necessary to handle increased workloads rather than buying additional large, expensive, proprietary UNIX-based servers. Dell IT has implemented this type of system for Dell operations at multiple locations around the world. Using the same processes for disaster recovery, backup, and monitoring across all Dell operations enables Dell IT to take advantage of a cost-effective and readily supported deployment model. This article describes how Dell IT migrated its critical SCM applications from proprietary UNIX-based servers to industry-standard Dell hardware.

Dell supply chain management

The Dell SCM system must handle an enormous number of transactions and pieces of information, and includes multiple core components necessary to keep operations running smoothly:

- **Configuration management:** The configuration management component manages over 1 million Dell part numbers per year across approximately 200 product families, and over 2 million bills of materials (BOMs) per year. BOMs listing component part numbers are created for manufacturing facilities to build assemblies and subassemblies for Dell products.
- **Procurement:** The procurement component manages nearly 1.8 million purchase order lines per year from more than 5,000 suppliers worldwide. To streamline the procurement process, Dell uses an automated application that includes

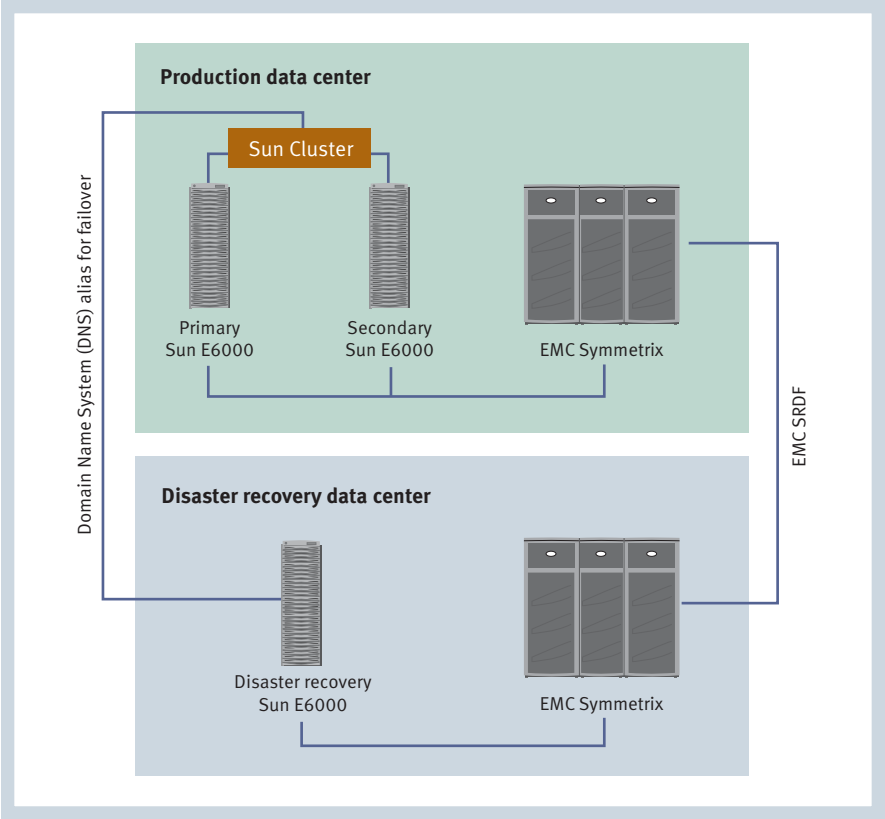

Figure 1. *Previous Dell supply chain management system based on proprietary UNIX-based servers*

workflow approvals and vendor communication and enables services such as defective part replacement.

- **Cost:** The cost component runs mostly in batch mode to calculate the costs to Dell for all BOMs. These batch jobs run weekly, monthly, and quarterly, with each job aggregating total material costs.
- **Inventory:** The inventory component manages more than 3 million inventory movements daily from stock rooms to factory floors across all Dell sites, along with the corresponding 3 million messages transmitted to different systems for reporting, analysis, and factory scheduling.
- **Accounts payable:** The accounts payable component handles approximately 15,000 items per day, including payments to Dell suppliers, invoices, and receipts. Vendor information includes vendor ID number, location, negotiated terms, and contact information.

In addition to these order-related transactions, the SCM system also runs several other batch process jobs to aggregate data each week, month, or quarter.

In the North America region, the Oracle Database application for SCM consists of approximately 3,000 database objects (functions, packages, procedures, triggers, tables, and views). This SCM system is supported by six PowerEdge 2650 application servers, five internally developed Web-based applications, more than 50 system-to-system integrations, approximately 125 batch jobs, and approximately 500 user interfaces.

Previous system: Proprietary UNIX-based servers

The previous Dell SCM system used Sun E6000-class UNIX-based servers running the Sun Solaris 8 OS and Sun Cluster 2.2 (see Figure 1). The primary and secondary servers, each with 16 processors at 336 MHz and 11 GB of memory, managed the production database on Oracle Database 8.0.6. The disaster recovery server had 12 processors at 336 MHz and 6 GB of memory. Arrays running the EMC® Symmetrix® platform provided disk storage.

To make the system highly available, Dell IT used Sun Cluster to cluster the primary and secondary servers in an active/passive configuration connected to shared disk storage, enabling the database to fail over to the passive node if the active one failed. The failover was initiated when the passive node no longer detected a heartbeat signal from the active node. The failover disconnected all users, who would then have to reconnect to the newly active server.

Current system: Oracle RAC on Dell PowerEdge servers

By 2005, the performance of the proprietary system was becoming a concern, with many of the batch processes taking a long time to complete—the end-of-quarter batch process job, for example, could take up to 31 hours. In addition, the Oracle Database version was outdated and unsupported, and because it did not support Oracle RAC, Dell IT could not add capacity by horizontal scaling. To continue using this system would have required a large investment in upgrading these large, expensive servers.

Instead, Dell decided to migrate to a cost-effective, industry standards–based platform,

TALK BACK

Tell us how the Dell Enterprise Technology Center can help your organization better simplify, utilize, and scale enterprise solutions and platforms. Send your feedback and ideas to enterprise_techcenter@dell.com.

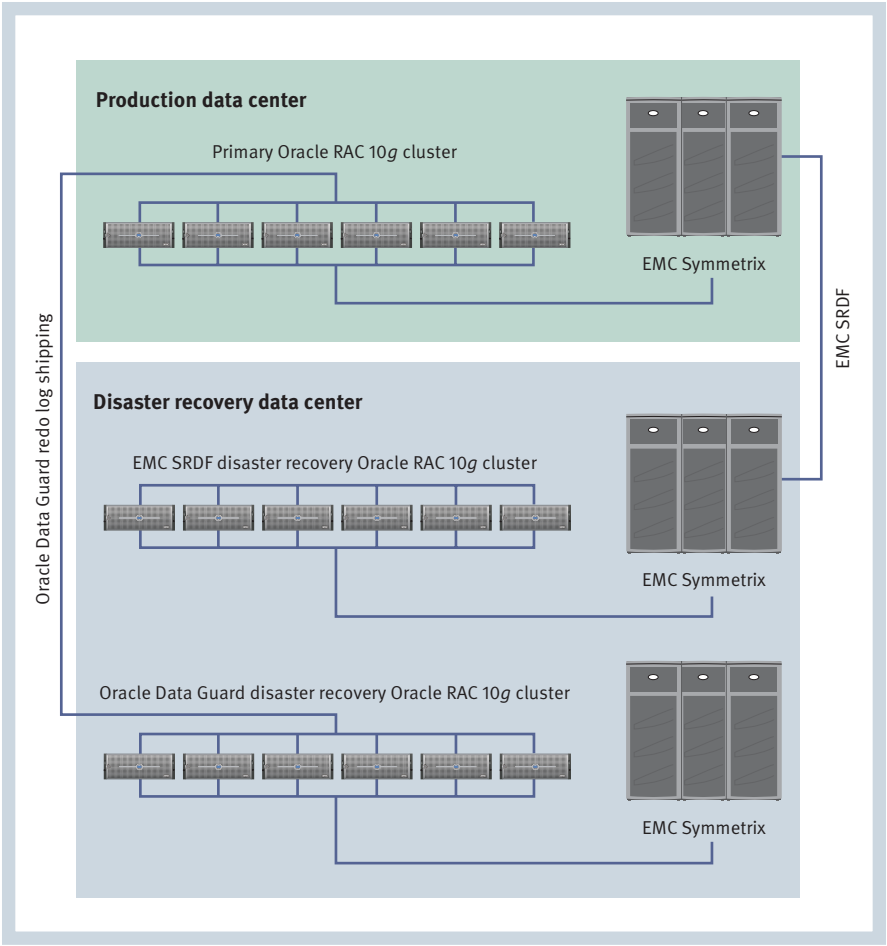

Figure 2. *Current Dell supply chain management system based on Oracle RAC 10g and Dell PowerEdge servers*

replacing the Sun servers with Oracle RAC 10*g* clusters of Dell PowerEdge 6650 servers, each with four Intel® Xeon® processors (see Figure 2). The environment includes a primary cluster and two disaster recovery clusters. While the previous system could not scale beyond one physical server, the current system allows Dell IT to easily add servers to the clusters to handle increased workloads. The multi-node Oracle RAC architecture can also provide higher I/O throughput than the previous system because of the increased number of available interfaces.

The Oracle RAC 10*g* architecture includes built-in high-availability features. All Oracle RAC nodes and instances share the same physical database; if any of these experience problems, the users and connections can fail over to the other nodes or instances. Because this failover is seamless and transparent to users, users do not have to reconnect, as they did in the previous system.

One of the disaster recovery clusters uses Oracle Data Guard software, and the other uses EMC Symmetrix Remote Data Facility (SRDF®) software. The first cluster uses the Data Guard Broker component to ship logs in real time from the primary site to the disaster recovery site and apply them, which allows the database to stay in sync and helps prevent physical corruptions from being copied to the disaster recovery database. The second cluster uses SRDF to ship changed blocks at the storage level in real time from the primary site to the disaster recovery site. This method is fast but does not guard against data corruption. The combination of Data Guard and SRDF provides both fast and secure data replication.

Dell IT also uses several other programs to provide key backup, load balancing, and administration capabilities:

- **Oracle Recovery Manager:** Dell IT uses this program to back up the primary database and archive logs. Two full (hot) backups are made to tape weekly through the EMC NetWorker™ interface; archive logs are backed up hourly to enable Dell IT to recover the database to a specific point in time if necessary. Dell IT can also back up the database from the disaster recovery site rather than taxing the primary nodes.
- **Oracle Database 10*g* Services:** This software provides load balancing for each cluster—for example, Dell IT can have three nodes handle online users and three nodes handle batch processing.
- **Oracle Grid Control:** This software provides monitoring and management capabilities and helps simplify many day-to-day database administration tasks; it also allows administrators to set thresholds for different events to create trouble tickets through the problem management system.

Dell IT has implemented these clusters at multiple locations around the world (see Figure 3).

"Standardizing on cost-effective Oracle RAC 10*g* clusters of PowerEdge servers with Dell/EMC storage helped Dell IT deploy the systems globally in just eight months."

Standardizing on cost-effective Oracle RAC 10*g* clusters of PowerEdge servers with Dell/EMC storage, with similar processes for deployment, disaster recovery, and backup, helped Dell IT deploy the systems globally in just eight months.

Performance increases

Dell engineers measured database performance for both the previous and current SCM systems. Figure 4 shows the 10 longest end-of-month and end-of-quarter transactions. The most significant time savings occurred in the data extraction for all material movements transactions, which dropped from almost 5 hours to just 35 minutes for an 88 percent improvement, while the time for the entire end-of-quarter jobs processing decreased from 31 hours to 23 hours.

Cost-effective, scalable supply chain management

Supply chain management is essential to Dell operations around the world, with both factory operations and internal systems dependent on SCM systems to provide real-time information about key business functions. Running these systems on Oracle RAC 10*g* clusters of industry-standard Dell PowerEdge servers enables Dell IT to scale them efficiently and cost-effectively to handle increased workloads. By moving the systems to Dell servers when it did, Dell IT avoided significant additional expenditures for proprietary UNIX-based servers, enhanced performance, and provided a clear path for future growth.

Region	Primary site	Disaster recovery site
North America	Six-node Oracle RAC 10*g* cluster	Two sets of six-node Oracle RAC 10*g* clusters
China	Three-node Oracle RAC 10*g* cluster	Three-node Oracle RAC 10*g* cluster
Brazil	Two-node Oracle RAC 10*g* cluster	Two-node Oracle RAC 10*g* cluster
Europe	Two-node Oracle RAC 10*g* cluster	Two-node Oracle RAC 10*g* cluster
Malaysia	Two sets of two-node Oracle RAC 10*g* clusters	Two sets of two-node Oracle RAC 10*g* clusters

Figure 3. *Global deployment of Dell supply chain management systems*

Transaction	Time on previous system (hours:minutes)	Time on current system (hours:minutes)	Performance gain
Entire end-of-quarter jobs processing	31:00	23:00	26%
Cost-per-order calculation	8:52	5:40	36%
Quarterly BOM calculation	7:30	1:56	74%
Entire end-of-month jobs processing	6:00	4:00	33%
New product material costs calculation	5:56	2:17	62%
Data summary for all levels of order details	5:33	1:58	65%
Data extraction for all material movements transactions	4:49	0:35	88%
Costs roll-up (materials, royalties, and transportation)	3:18	3:12	3%
Material management transactions creation based on inventory transactions	3:00	0:43	76%
Financial journal entries creation	2:37	0:40	75%

Figure 4. *Performance gains for 10 longest database transactions in Dell supply chain management systems*

Dave Jaffe, Ph.D., *is a senior consultant on the Dell Enterprise Technology Center team. He has a B.S. in Chemistry from Yale University and a Ph.D. in Chemistry from the University of California, San Diego.*

Todd Muirhead *is a senior engineering consultant on the Dell Enterprise Technology Center team. Todd has a B.A. in Computer Science from the University of North Texas and is Microsoft Certified Systems Engineer + Internet (MCSE+I) certified.*

Tiong Tey *is an IT strategist on the Dell World Wide Procurement IT team. Tiong has a B.S. and a master's degree in Computer Science from Wichita State University.*

Raveendra Avutu *is an IT manager in the Global Data Management group supporting databases for Dell IT. Raveendra has a master's degree in Computer Science from Oklahoma State University.*

more online

www.dell.com/**powersolutions**

QUICK LINKS

Dell Enterprise Technology Center:
www.dell.com/techcenter

Dell and Oracle Database 10*g*:
www.dell.com/oracle10g

Oracle and Dell:
www.dell.com/oracle

Monitoring Dell PowerEdge Servers with Oracle Enterprise Manager 10*g* Grid Control

BY DAVE JAFFE, PH.D.

The integration of Oracle® Enterprise Manager 10*g* Grid Control with Dell OpenManage™ software allows administrators to monitor Oracle database metrics and ninth-generation Dell™ PowerEdge™ servers from a single console. This article discusses best practices for installing, configuring, and using this software to help simplify systems management.

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Oracle

Scalable enterprise

Systems management

Visit www.dell.com/powersolutions for the complete category index.

Oracle Enterprise Manager (OEM) 10*g* Grid Control enables Oracle system and database administrators to manage large Oracle grids. Dell and Oracle have worked together to integrate OEM 10*g* Grid Control with Dell OpenManage Server Administrator (OMSA) to provide Dell hardware–specific data within OEM 10*g* Grid Control. This integration allows administrators to monitor Oracle database metrics and ninth-generation Dell PowerEdge servers from a single console, helping eliminate the need to learn and monitor multiple interfaces. Because this integration uses the industry-standard Simple Network Management Protocol (SNMP), no additional Dell hardware or software components are required. Administrators can use the OEM 10*g* Grid Control console with Dell OpenManage to carry out discovery and reporting of Dell servers, manage hardware event notifications, execute event-driven policies, and maintain server assets and inventory.

This article describes best practices for installing, configuring, and using OEM 10*g* Grid Control and OMSA to help simplify systems management. To demonstrate these processes, Dell engineers installed the OEM 10*g* Grid Control console on a Dell PowerEdge 1950 server, then installed OMSA and emagent—the OEM 10*g* agent—on this PowerEdge 1950 as well as on two PowerEdge 2950 nodes of an Oracle Real Application Clusters (RAC) system (see Figure 1). OMSA uses SNMP to communicate with emagent, which in turn uses HTTP over Secure Sockets Layer (HTTPS) to communicate with the central OEM 10*g* Grid Control console.

The servers ran the 32-bit version of Red Hat® Enterprise Linux® AS 4, Update 3, but the integration also supports other Linux distributions and the Microsoft® Windows Server® OS. When using an OS other than Red Hat Enterprise Linux, administrators must modify the commands described in this article appropriately.

Installing and configuring Oracle Enterprise Manager 10*g* Grid Control

The first step in setting up a system integrating OEM 10*g* Grid Control and OMSA is to install, patch, and launch OEM 10*g* Grid Control. This section outlines the steps necessary to carry out these tasks; for complete details, visit www.oracle.com/technology/documentation.

Installing Oracle Enterprise Manager 10*g* Grid Control

OEM 10*g* Grid Control is a separate component from Oracle Database that requires its own licenses. The current version is based on the 32-bit version of Oracle Database 10*g* Release 2 (R2), which must be installed on a 32-bit OS. The OS must be configured for a standard 32-bit Oracle Database 10*g* R2 installation, including the proper packages, kernel parameters, and users, groups, and directories. These requirements are described in the *Oracle*

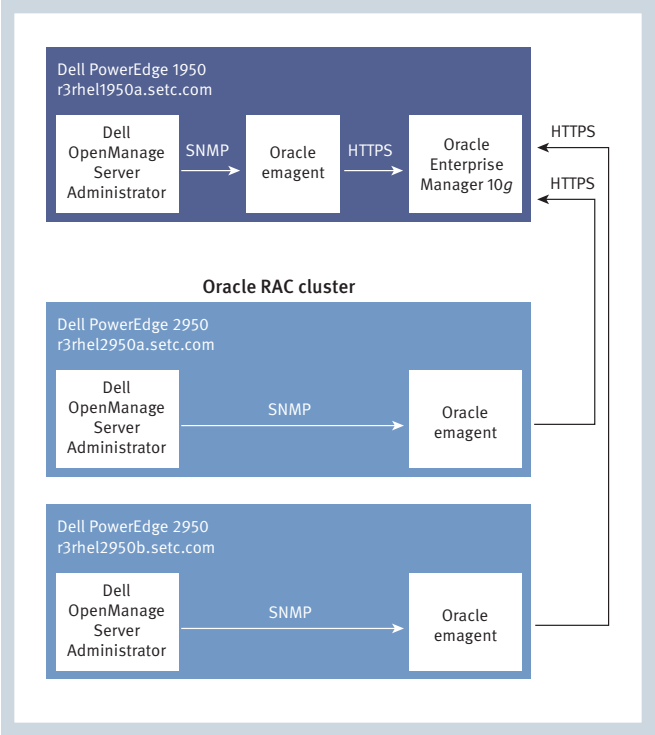

Figure 1. *Example configuration integrating Oracle Enterprise Manager 10g Grid Control and Dell OpenManage Server Administrator*

Enterprise Manager Grid Control Quick Installation Guide for Linux x86, 10g Release 2 (10.2).[1] The Dell engineers followed this guide when preparing the Red Hat Enterprise Linux 4 OS for OEM 10*g* Grid Control installation, including defining the user "oracle" as a member of the oinstall and dba groups and creating the /opt/oracle directory with oracle.oinstall as the owner.

The OEM 10*g* Grid Control files, including agent software for 64-bit systems and patch 3731593, are available from the Oracle Web site in the following five zip files:

- Linux_Grid_Control_full_102010_disk1.zip
- Linux_Grid_Control_full_102010_disk2.zip
- Linux_Grid_Control_full_102010_disk3.zip
- Linux_x86_64_Grid_Control_agent_download_10_2_0_2_0.zip
- p3731593_10202_LINUX.zip

The Dell team installed these files by running the OEM 10*g* Grid Control installer in an X Window shell as user "oracle" in the directory where they unzipped the OEM Grid Control files. First, they executed the following Linux command:

```
export ORACLE_BASE=/opt/oracle
```

[1] This installation guide is available at download-east.oracle.com/docs/cd/B19306_01/install.102/b28091.pdf.

Parameter	Value
Installation Type	Enterprise Manager 10*g* Grid Control Using a New Database
Installation Location	*Parent Directory:* /opt/oracle/OracleHomes (default)
Inventory directory and credentials	• *Path*: /opt/oracle/oraInventory (default) • *OS group name*: oinstall (default)

Figure 2. *Oracle Enterprise Manager 10g Grid Control installation parameters*

Parameter	Value
Repository Database Name	emrep (default)
Repository Database File Location	/opt/oracle/oradata (default)
Database Administrator (OSDBA) group	dba
Database Operator (OSOPER) group	oinstall (default)

Figure 3. *Oracle Enterprise Manager 10g Grid Control configuration parameters*

Next, they executed the command `./runInstaller` and specified the installation parameters shown in Figure 2. After the product-specific prerequisite checks, the Dell team entered the configuration parameters shown in Figure 3. They skipped the Specify Optional Configuration step and set the passwords in the Specify Security Options step. For demonstration purposes, the Dell engineers used "oracle1" for all passwords.

After completing the installation, the Dell team ran the configuration scripts as instructed and accepted the default answer to all questions. When the configuration scripts finished, the main installation screen returned and various configuration assistants ran. Because the Agent Configuration Assistant failed despite repeated attempts during the setup,

TALK **BACK**

Tell us how the Dell Enterprise Technology Center can help your organization better simplify, utilize, and scale enterprise solutions and platforms. Send your feedback and ideas to **enterprise_techcenter@dell.com**.

the Dell team skipped that step and installed the agent later (see the "Installing Oracle Enterprise Manager 10*g* Grid Control agents on target systems" section in this article).

After the configuration assistants had run, the Dell team returned to the main installation screen and selected "End" and then "Exit." A message provided the following information: "Use the following URL to access the Enterprise Manager Grid Control: http://r3rhel1950a:4889/em."

Patching Oracle Enterprise Manager 10*g* Grid Control

Before applying patch 3731593 to OEM 10*g* Grid Control, the Dell team checked that the repository database (emrep) was running by executing the following commands as the user "oracle":

```
export ORACLE_SID=emrep
/opt/oracle/OracleHomes/db10g/bin/lsnrctl
   status
/opt/oracle/OracleHomes/db10g/bin/lsnrctl start
  # if necessary
/opt/oracle/OracleHomes/db10g/bin/sqlplus "/ as
   sysdba"
  # if connected to idle instance: startup
```

They next stopped the management server by executing the following command as the root user:

```
/opt/oracle/OracleHomes/oms10g/opmn/bin/opmnctl
   stopall
```

Finally, the Dell team stopped the OEM 10*g* agent running on the server by executing the following command as the user "oracle":

```
/opt/oracle/OracleHomes/agent10g/bin/emctl stop
   agent
```

The patch must be installed four times—once each for Oracle Configuration Manager (OCM), the Oracle Management Server (OMS), the database repository, and the OEM agent. Because OCM was unnecessary for the demonstration, the Dell team declined the OCM license agreement, which resulted in the software being installed but not started. Before installing the patch, the Dell team ran the Linux command `unset ORACLE_HOME`. They then installed the patch by running the command `./runInstaller` four times from the directory where the patch was unzipped, specifying Oracle Home in turn as OCM, oms10g, db10g, and agent10g.

Launching Oracle Enterprise Manager 10*g* Grid Control

The Dell team launched OEM 10*g* Grid Control by executing the following command as the root user:

```
/opt/oracle/OracleHomes/oms10g/opmn/bin/opmnctl
   startall
```

They could then access OEM 10*g* Grid Control from the URL shown after installation (http://r3rhel1950a:4889/em) with the username "sysman" and the password "oracle1." They could start OEM 10*g* Grid Control after a reboot by first executing the following commands as the user "oracle":

```
export ORACLE_SID=emrep
/opt/oracle/OracleHomes/db10g/bin/lsnrctl start
/opt/oracle/OracleHomes/db10g/bin/sqlplus "/ as
   sysdba"
  SQL> startup
  SQL> quit
```

Next, the Dell team executed the following command as the root user:

```
/opt/oracle/OracleHomes/oms10g/opmn/bin/opmnctl
   startall
```

Finally, they executed the following command as the user "oracle":

```
/opt/oracle/OracleHomes/agent10g/bin/emctl
   start agent
```

Installing Oracle Enterprise Manager 10*g* Grid Control agents on target systems

Managing remote Oracle systems, known as target systems in OEM 10*g* Grid Control, requires installing emagent, the OEM 10*g* agent, on the target systems. To demonstrate this process, the Dell team installed emagent on a two-node Oracle RAC cluster consisting of two PowerEdge 2950 servers (r3rhel2950a and r3rhel2950b) running the 64-bit version of Red Hat Enterprise Linux 4. The RAC database instance was r3rac1.

"Administrators can use OEM 10*g* Grid Control to monitor Dell hardware data, configure alert notifications and corrective actions, and configure target groups to view multiple servers in a single display."

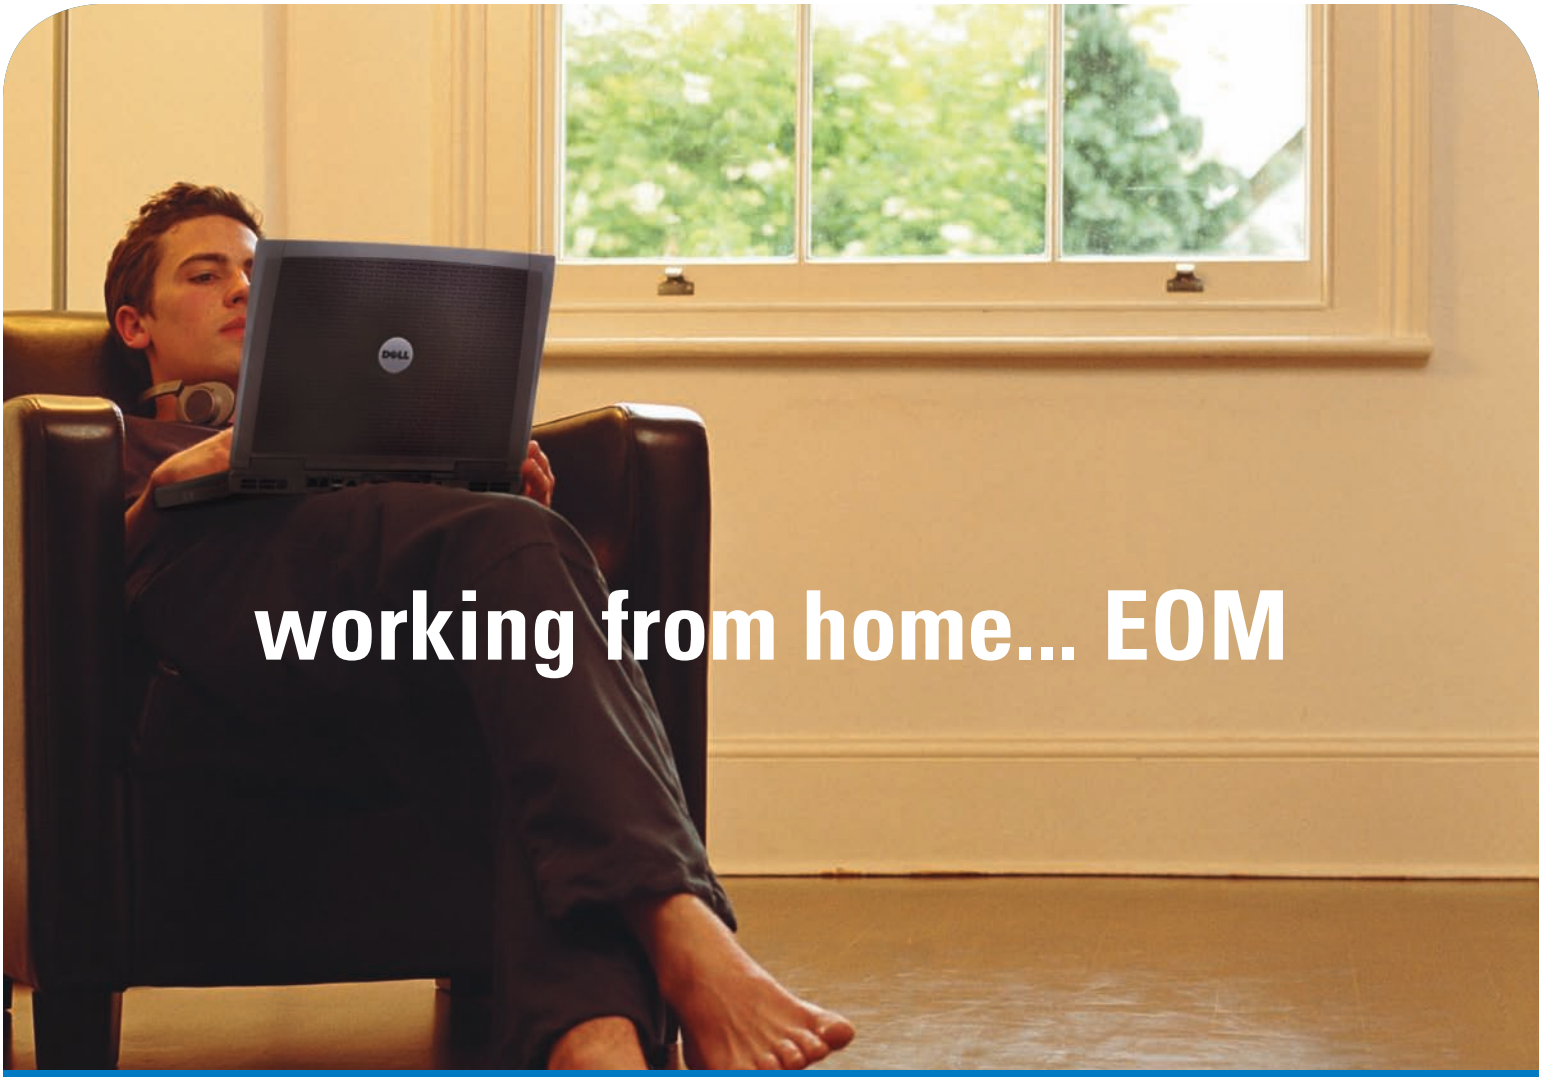

Access the data center from the comfort of anywhere.

The Dell™ 2161DS-2/4161DS remote console switches give today's server administrators the power to monitor and control servers from virtually anywhere. For maximum flexibility, the 4161DS provides **access for up to four simultaneous users and by using Dell's remote console software, simply point-and-click to take control of your servers. Whether you're at the rack, in your office, across the globe, or in your living room, your data center can always be within your reach.**

Digital Availability. Easy as DELL™

www.dell.com

Preparing for installation

After OEM 10*g* Grid Control has been installed on the OEM 10*g* Grid Control host (the PowerEdge 1950), the directory /opt/oracle/OracleHomes/oms10g/sysman/agent_download/10.2.0.2.0 contains linux and win32 subdirectories. To avoid an installation error, the Dell team edited the already-installed agent_download.rsp file in the 10.2.0.2.0 directory by changing the line `s_OMSHost="r3rhel1950a.setc.com\:4889__Management_Service"` to `s_OMSHost="r3rhel1950a.setc.com"`.

Because the PowerEdge 2950 nodes were running a 64-bit Linux OS and the 64-bit Linux agent is not included in OEM 10*g* Grid Control by default, the Dell team added this agent using the following command from the 10.2.0.2.0 directory (where *path* is the location of the 64-bit Linux agent zip file):

```
unzip path/Linux_x86_64_Grid_Control_agent_
   download_10_2_0_2_0.zip
```

These commands create a new linux_x64 subdirectory in the 10.2.0.2.0 directory. The Dell team next edited two lines in the agentDownload .linux_x64 file in this new subdirectory to read as follows:

```
OMSHost="r3rhel1950a.setc.com"
httpPort="4889"
```

They then copied this file to a directory owned by the user "oracle" (the Dell team created the /home/oracle/10g/agent directory for this purpose) on each target. Only one cluster node needs a copy of this file.

On each target, administrators should verify that the user "oracle" has the wget and jar (Java Archive Tool) programs in its path. If wget is not present, they should add it to the path. If jar is not present, they should download the Java SE Development Kit for the Linux x64 platform from java.sun.com (the Dell team used jdk-1_5_0_06-linux-amd64-rpm.bin) and install it on each target.

Installing the agent

To install the agent, administrators can run the agent installation script on each target to be managed. For the two-node Oracle RAC cluster, the Dell team ran the following command on one node only:

```
./agentDownload.linux_x64 -b /opt/oracle/
   product/10.2.0 -c "r3rhel2950a,r3rhel2950b"
   -n r3rac1
```

When prompted, the Dell team entered the "oracle1" password; then, as the root user on each node, they ran the following script and accepted all defaults:

```
/opt/oracle/product/10.2.0/agent/agent10g/
   root.sh
```

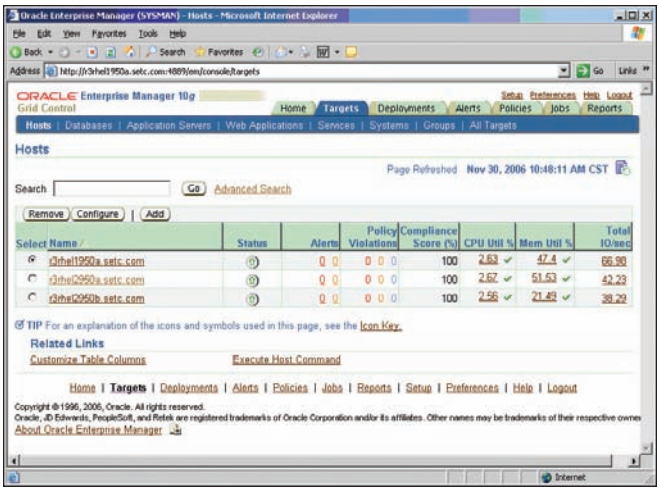

Figure 4. *Hosts screen in Oracle Enterprise Manager 10g Grid Control following emagent installation on a two-node Dell PowerEdge 2950 Oracle RAC cluster*

This script creates a new agent home at /opt/oracle/product/10.2.0/agent10g. To verify that the agent was running, the Dell team ran the following command as the user "oracle" on each node:

```
/opt/oracle/product/10.2.0/agent10g/bin/emctl
   status agent
```

Figure 4 shows the OEM 10*g* Grid Control Hosts screen following successful emagent installation.

Installing and configuring Dell OpenManage Server Administrator

The next step is to install OMSA and configure SNMP to allow OMSA to interface with emagent on each system. These instructions describe installing OMSA on the system that is also the OEM 10*g* Grid Control host (r3rhel1950a in the Dell test scenario), but they can apply to all systems being managed by OEM 10*g* Grid Control. For more information, see the *Dell OpenManage Server Administrator User's Guide*.

Installing Dell OpenManage Server Administrator

Before installing OMSA, administrators should use the following command to verify that the snmpd agent is installed:

```
rpm -qa | grep -i snmp
```

This command should produce output similar to the following:

```
net-snmp-libs-5.1.2-11.EL4.6
net-snmp-5.1.2-11.EL4.6
```

Note: Because OMSA installation modifies the /etc/snmp/snmpd.conf SNMP configuration file, administrators should make a copy of this file with a different name before proceeding with installation.

To install OMSA, the Dell team first downloaded the OMSA Managed Node software for Linux (OM_5.1_ManNode_LIN_A00.tar.gz) from support.dell.com by selecting Drivers and Downloads › PowerEdge 1950 › Enterprise Linux 4 › Systems Management › Dell OpenManage Server Administrator Managed Node. Next, they created a /root/srvadmin directory, copied the OMSA Managed Node software to it, and executed the following commands from that directory:

```
gunzip OM_5.1_ManNode_LIN_A00.tar.gz
tar xvf OM_5.1_ManNode_LIN_A00.tar
cd linux/supportscripts/
./srvadmin-install.sh --express
./srvadmin-services.sh start
```

These commands installed OMSA to /opt/dell/srvadmin. If this process results in dependency errors, administrators can install the necessary Red Hat Package Manager files from the appropriate OS subdirectory in /root/srvadmin/linux/RPMS/supportRPMS.

Administrators can access OMSA over the Web using the host name with port 1311—for example, https://r3rhel1950a:1311. This Web access is not required for OMSA to integrate with OEM 10*g* Grid Control, and the Web access option can be omitted during OMSA installation for security or performance reasons. Administrators can also access OMSA from a Linux shell using `omreport` commands such as `omreport server summary`, `omreport chassis fans`, and `omreport chassis temp`. They can view the available commands with `omreport -?`.

Configuring SNMP

Installing OMSA adds the following lines to the snmpd.conf file to allow OMSA to communicate with other services using SNMP:

```
# Allow Systems Management Data Engine SNMP to
    connect to snmpd using SMUX
smuxpeer .1.3.6.1.4.1.674.10892.1
```

To enable OMSA to communicate with the emagent software running on a particular server, administrators can add the localhost IP address to snmpd.conf. The Dell team added the following line:

```
rocommunity public 127.0.0.1
```

They then restarted the snmpd service with the command `service snmpd restart`. (If this command does not work, administrators should reboot the server.) To make the snmpd process start up on boot, the Dell team executed the following command as the root user:

```
chkconfig --level 35 snmpd on
```

Using the localhost IP address restricts the services that can access OMSA to those running on that server—most importantly, emagent. This restriction helps prevent external access to the critical hardware-level OMSA features, although administrators can still configure the system to pass SNMP traps to other destinations without compromising security. Administrators should also keep in mind that although OMSA uses SNMP to support remote inventory, it does not support remote configuration using SNMP set or write actions: HTTPS is used for all remote configuration tasks as well as for the interface between emagent and the OEM 10*g* Grid Control console.

Verifying the integration

Administrators can verify the integration of OEM 10*g* Grid Control and OMSA by accessing the URL http://*oem_hostname*:4889/em (for example, http://r3rhel1950a:4889/em), logging in, and selecting the Targets tab, then selecting the host and clicking the All Metrics link (see Figure 5). The Fans, Memory Devices, Power Supplies, Processors, Remote Access Card, System BIOS, and Temperature metrics are Dell-specific and come from OMSA.

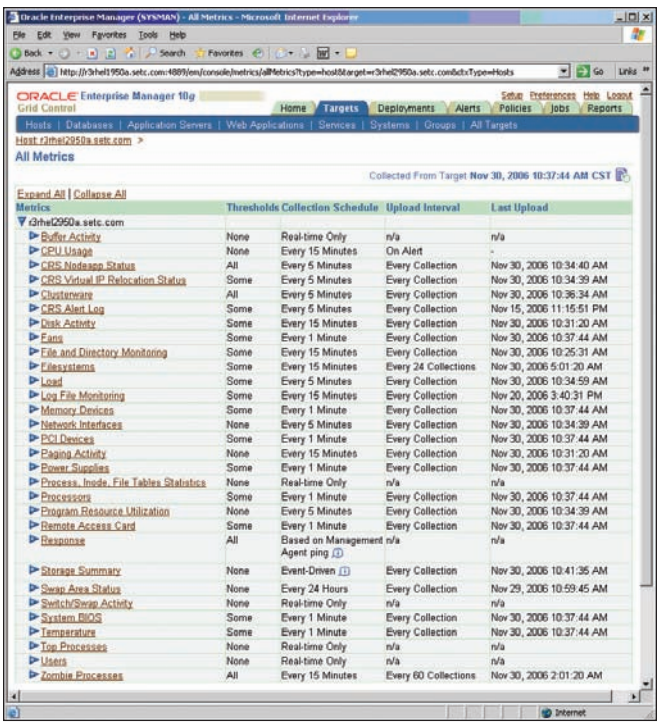

Figure 5. *All Metrics screen in Oracle Enterprise Manager 10g Grid Control displaying Dell-specific metrics from Dell OpenManage Server Administrator*

```
Temperature Probes Information
  -----------------------------------
  Main System Chassis Temperatures: Ok
  -----------------------------------
  Index                      : 0
  Status                     : Ok
  Probe Name                 : System Board
                               Ambient Temp
  Reading                    : 30.0 C
  Minimum Warning Threshold  : 8.0 C
  Maximum Warning Threshold  : 42.0 C
  Minimum Failure Threshold  : 3.0 C
  Maximum Failure Threshold  : 47.0 C
```

Figure 6. *Temperature information for a Dell PowerEdge 2950 server displayed using the omreport command*

Using Oracle Enterprise Manager 10*g* Grid Control to monitor and manage Dell servers

Administrators can use OEM 10*g* Grid Control to monitor Dell hardware data, configure alert notifications and corrective actions, and configure target groups to view multiple servers in a single display.

Monitoring Dell hardware metrics

OEM 10*g* Grid Control manages Dell hardware data through a set of critical and warning alerts. The Dell team first compared the temperature readings provided in OEM 10*g* Grid Control with those provided by the `omreport` command for one of the PowerEdge 2950 servers, r3rhel2950a. In the OEM 10*g* Grid Control console, they selected Targets › r3rhel2950a › All Metrics › Temperature › Current Temperature; OEM 10*g* Grid Control displayed the last known value as 30. Next, from a Linux shell on the PowerEdge 2950, they ran the command `omreport chassis temps`, which provided the information shown in Figure 6.

Setting	Value
Name	Temp_notification_rule
Description	Notification rule for monitoring host chassis temperature
Make Public	Selected
Target Type	Host (apply rule to all host targets)
Availability	All options unselected
Metrics	Temperature Probe Status (both critical and warning)
Methods	Send Me E-mail

Figure 7. *Example notification rule settings for a temperature alert e-mail notification in Oracle Enterprise Manager 10g Grid Control*

"Integrating Dell-specific server monitoring and management with Oracle Enterprise Manager 10*g* Grid Control enables Oracle system and database administrators to manage their hardware and software from a single console."

Both methods reported the same value, as expected. OMSA uses the maximum and minimum warning and failure thresholds shown in Figure 6 to send out-of-bounds parameter alerts through OEM 10*g* Grid Control. To demonstrate these alerts, the Dell team first lowered the parameter collection schedule from 15 minutes to 1 minute in OEM 10*g* Grid Control by selecting Targets › r3rhel2950a › Metric and Policy Settings and selecting the Collection Schedule entry for Temperature Probe Status. They changed the Collection Frequency setting to repeat every minute, the Use of Metric Data setting to "Alerting and Historical Trending," and the Upload Interval setting to one collection. (Administrators should keep in mind that modifying collection schedules affects groups of metrics, not individual metrics.)

Next, the Dell team forced an alert by setting the temperature Maximum Warning Threshold value to 25 (lower than the current temperature). To do so, they executed the following command as the root user in a Linux shell:

```
omconfig chassis temps index=0 maxwarnthresh=25
```

This configuration forced a warning alert in the OEM 10*g* Grid Control console within a minute, with OEM 10*g* Grid Control changing the r3rhel2950a yellow (warning) number under "Alerts" from 0 to 1. The Dell team could then display the details of the alert metric and show a graph of the parameter over time. After this type of problem is fixed, the alert disappears.

Configuring alert notifications and corrective actions

Administrators can configure OEM 10*g* Grid Control to send alert notifications as e-mail or text messages and to trigger a corrective action based on an alert. Sending alert notifications requires adding the e-mail or text message addresses, then creating a new notification rule that

Setting	Value
Name	fix_temp
Description	Reset temperature warning threshold
Command Type	Single Operation
Command	/usr/bin/omconfig chassis temps index=0 warnthresh=default
Credentials	Selected "Override Preferred Credentials" and entered "root" and password

Figure 8. *Example corrective action settings to reset the temperature warning threshold in Oracle Enterprise Manager 10g Grid Control*

uses those addresses. To demonstrate this, the Dell team selected Preferences › General › E-mail addresses › Add Another Row, then checked the "Select" check box, entered an e-mail address, selected "Short Format" from the Message Format drop-down menu, and clicked the Apply button.

Next, they selected Preferences › Notification Rules › Host Availability and Critical States › Create Like and entered the settings shown in Figure 7. After this configuration is complete, setting off a temperature alert as described in the preceding section results in an e-mail being sent to the specified address.

OEM 10*g* Grid Control can also respond to an alert with a configurable corrective action. To demonstrate this, the Dell team set up a corrective action that simply resets the temperature warning threshold. First, they selected Setup › Corrective Action Library › Create Library Corrective Action, selected the OS Command option from the drop-down menu, and clicked the Go button. They then entered the settings shown in Figure 8.

After saving this action to the library, they selected Targets › r3rhel2950a › Metric and Policy Settings › Edit › Temperature Probe Status, then Monitored Objects › Edit › Edit Advanced Settings, and finally

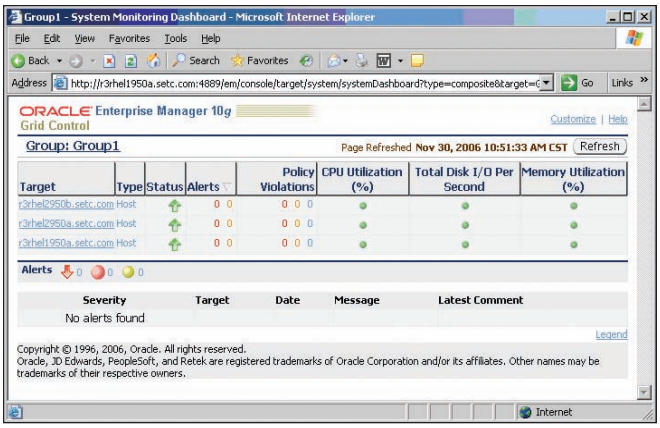

Figure 9. *System Monitoring Dashboard for an example group in Oracle Enterprise Manager 10g Grid Control*

Edit › Corrective Action (Warning) › Add › Add Corrective Action. They then selected the From Library option in the drop-down menu and fix_temp as the action, and clicked through the rest of the screens to complete the configuration. After this configuration is complete, setting off a temperature alert as described in the preceding section triggers the correction to reset the temperature warning threshold to the default value. The alert itself appears briefly, then disappears about a minute later.

Configuring target groups

To create a group to display multiple targets in a single screen, the Dell team selected Targets › Groups › Add, then created a group called Group1 that included r3rhel1950a, r3rhel2950a, and r3rhel2950b. They then selected Group1, clicked Configure › Dashboard, and set the Refresh Frequency setting to 1 minute. Finally, to view the group, they selected Groups › Group1 › Launch Dashboard. Figure 9 shows the System Monitoring Dashboard for this group.

Simplifying Dell server management with Oracle Enterprise Manager 10*g* Grid Control

Integrating Dell-specific server monitoring and management with Oracle Enterprise Manager 10*g* Grid Control enables Oracle system and database administrators to manage their hardware and software from a single console. By combining these two tools and using features such as alert notifications, corrective actions, and target groups, administrators can help simplify their server management duties in environments running Oracle databases on Dell PowerEdge servers.

Dave Jaffe, Ph.D., *is a senior consultant on the Dell Enterprise Technology Center team who specializes in cross-platform solutions. He has a B.S. in Chemistry from Yale University and a Ph.D. in Chemistry from the University of California, San Diego.*

more Online
www.dell.com/**powersolutions**

QUICK LINKS

Dell Enterprise Technology Center:
www.dell.com/techcenter

Oracle documentation:
www.oracle.com/technology/documentation

Dell OpenManage Server Administrator User's Guide:
support.dell.com/support/edocs/software/svradmin

Understanding the Challenges of Delivering Cost-Effective, High-Efficiency Power Supplies

BY STUART BERKE
DAVID MOSS
RANDY RANDALL

It's no small matter. Power supplies can account for as much as one-fifth of the wasted power in data center servers. By emphasizing energy-efficient power supply design, Dell is addressing a key factor that can help reduce overall power and cooling requirements and contribute to lower total cost of ownership.

Related Categories:

Energy efficiency

Power and cooling

Power management

Power supplies

Total cost of ownership (TCO)

Visit www.dell.com/powersolutions for the complete category index.

Rising power and cooling costs have made energy efficiency a strategic concern for data center managers. To help organizations lower total cost of ownership (TCO) and create a strategy for continuing growth, Dell is addressing energy efficiency at every level—from component and server design to the overall data center infrastructure. This article explains how and why Dell is setting its sights on the redesign of the power supply as a key contributing factor to energy efficiency in the data center.

Although data center managers have the flexibility to choose many different system components to strike the most advantageous balance of performance and energy efficiency, power supply options have been limited. Common industry practice is to provide a single power supply design to cover the full range of system configurations and loads for a given server model. For years, cost-effectiveness has driven power supply design considerations, and energy efficiency has been a secondary concern.

In an ongoing effort to optimize energy usage in the data center, Dell has identified power supplies along with fans, processors, and memory as the components that typically consume the most energy within a server. A Dell power usage study in summer 2006 revealed that 20 percent of power loss in a typical data center server can be attributed to inefficient power supplies.[1] Clearly, improving power supply efficiency offers a significant opportunity to help lower energy costs for data center servers.

Dell has made it a priority to reduce power loss and improve energy efficiency in power supplies. Dell is working with key power supply manufacturers such as Artesyn Technologies, Delta Electronics, and Astec Power to help improve power supply designs and move power supply technology forward.

Understanding power supply specifications

When systems engineers specify the power supply requirements for an enterprise-class server, their approach is similar to that of a civil engineer designing a tunnel. Even though most traffic passing through the tunnel may be cars, the tunnel must be able to accommodate the largest trucks on the road. Likewise, common industry practice has been to specify a single power supply for each server based on maximum system configuration and load requirements.

Theoretically, systems engineers could utilize a different power supply for each server configuration, but designing to multiple power supply specifications could be cost-prohibitive. Instead, manufacturers typically purchase high volumes of a single power supply design to reduce cost and pass that component savings on to the customer.

When it comes to power supply design, balancing cost, performance, and efficiency is an ongoing challenge. Dell

[1] For more information, see "Data Center Efficiency in the Scalable Enterprise," by John Pflueger, Ph.D., and Sharon Hanson, in *Dell Power Solutions*, February 2007, www.dell.com/downloads/global/power/ps1q07-20070210-CoverStory.pdf.

continually weighs the need for improved energy efficiency against increasing cost. In addition to cost and energy efficiency considerations, Dell is committed to offering the most flexible server configurations possible.

Given current power supply design guidelines, the goal for systems engineers is to specify a single maximum power supply output rating that can handle a wide variety of configuration options and load requirements. For example, a Dell™ PowerEdge™ 2950 server can be configured in such a way that processor, memory, and hard disk power needs vary by as much as 480 W (see Figure 1). (Other configuration options—such as I/O adapters, internal storage controllers, and tape backup—can further widen the load requirements.) Although the minimum server component configuration in this example may draw as little as 60 W and the maximum server configuration may draw as much as 540 W, today's industry practice is to specify one power supply for both configurations.

	Minimum configuration		Maximum configuration	
	PowerEdge 2950 server component	Power draw	PowerEdge 2950 server component	Power draw
Processor	One dual-core Intel® Xeon® 5148 processor at 2.33 GHz	40 W	Two dual-core Intel Xeon 5080 processors at 3.73 GHz	260 W
Memory	Two 256 MB dual in-line memory modules (DIMMs) at 533 MHz	20 W	Eight 8 GB DIMMs at 667 MHz	120 W
Hard disk drives	None	0	Eight Serial Attached SCSI (SAS) drives at 15,000 rpm	160 W
Total		**60 W**		**540 W**

Figure 1. *Illustrative budgetary power requirements for Dell PowerEdge 2950 server component minimum and maximum configurations*

Comparing power supply output ratings

Even taking into account different component suppliers, system configurations, and design principles, published output data specifications indicate that maximum power supply ratings are fairly uniform across the industry. Figure 2 lists the published power supply output ratings of comparable enterprise-class servers—noting processor, memory type, chassis size, and whether they support redundant power supplies.

Each server manufacturer must support a wide range of configuration options, making the improvement of power supply efficiency an industry-wide challenge. However, careful choice of configuration flexibility and overall power supply output can drive significant distinctions in energy efficiency even though different manufacturers may incorporate similar power supply components.

Aligning design specifications with energy efficiency

Meeting the power requirements of a flexible, configurable enterprise-class server while improving energy efficiency requires a fundamental shift in power supply design specifications. Tackling this challenge—a challenge shared by all vendors in the industry—necessitates examining the causes of inefficiency in today's power supply designs.

When it comes to efficiency, a power supply may be likened to an internal combustion engine. Like an idling car burning fuel at a stoplight, a power supply uses some energy just to operate. And similar to automobiles cruising at highway speeds, power supplies typically exhibit top efficiency as they do more work. Still, this does not indicate that a power supply becomes more efficient at higher loads. Rather, the energy

Server model	Processor	Memory type	Chassis size	Support for redundant power supplies	Power supply output rating
Dell PowerEdge 1950	Intel Xeon	667 MHz fully buffered DIMM (FBD)	1U	Yes	670 W
HP ProLiant DL360 G5[a]	Intel Xeon	667 MHz FBD	1U	Yes	700 W
Intel Server Chassis SR1550[b]	Intel Xeon	667 MHz FBD	1U	Yes	650 W
IBM System x3455[c]	AMD Opteron 2000 series	667 MHz double data rate 2 (DDR2)	1U	No	650 W
IBM System x3550[d]	Intel Xeon	667 MHz FBD	1U	Yes	670 W
Dell PowerEdge 2950	Intel Xeon	667 MHz FBD	2U	Yes	750 W
HP ProLiant DL380 G5[e]	Intel Xeon	667 MHz FBD	2U	Yes	800 W
IBM System x3650[f]	Intel Xeon	667 MHz FBD	2U	Yes	835 W
IBM System x3655[g]	AMD Opteron 2000 series	667 MHz DDR2	2U	Yes	835 W

[a] h10010.www1.hp.com/wwpc/pscmisc/vac/us/en/ss/proliant/dl360g5-models.html
[b] www.intel.com/design/servers/chassis/sr1550/index.htm
[c] www-03.ibm.com/systems/x/rack/x3455/specs.html
[d] www-03.ibm.com/systems/x/rack/x3550/specs.html
[e] h10010.www1.hp.com/wwpc/pscmisc/vac/us/en/ss/proliant/dl380g5-models.html
[f] www-03.ibm.com/systems/x/rack/x3650/specs.html
[g] www-03.ibm.com/systems/x/rack/x3655/specs.html

Figure 2. *Maximum power supply output ratings for enterprise-class servers comparable to Dell PowerEdge 1950 and PowerEdge 2950 servers (as of January 1, 2007)*

losses become a smaller percentage of the total power being consumed, and it is this change in the energy consumption–to–energy loss ratio that results in higher efficiency. In other words, when a power supply is operating below peak loads—a state that might describe a minimally configured server or a server at idle—energy losses account for a greater percentage of the total energy being consumed, which results in a corresponding drop in efficiency. The "Redundancy Penalty: Availability Versus Energy Usage" sidebar in this article illustrates a common power supply efficiency curve, and explains how energy losses can be compounded by the use of redundant power supplies.

Hybrid vehicle technology may have solved the idling inefficiency dilemma in automobiles, but a different approach is needed to improve power supply efficiency. Because peak loads are not typical for every server deployed, design engineers are working toward improvements in energy efficiency throughout the power supply output range. To that end, Dell is working with power supply manufacturers to take the necessary steps to improve power supply efficiency. For example, changes in component materials and circuit design can help control power losses, as can low-resistance transistors or transistors operating in parallel. Alternatively, power supply designers can adopt new technologies that inherently offer higher efficiency than current practices.

Dell is working toward a goal of achieving over 80 percent efficiency throughout the power supply output range to help minimize energy loss regardless of server configuration or workload. Technology is available to accomplish this goal, but it is rare that advances in efficiency come without additional cost. More expensive circuit materials increase up-front power supply expenses. But with energy costs climbing, a growing number of organizations may view the investment in energy efficiency to be a worthwhile one. In many cases, long-term energy cost savings may offset the incremental up-front cost of energy-efficient components.

Benefiting from power efficiency in the data center and beyond

Optimizing power supply efficiencies throughout load ranges with effective designs and materials can help minimize power loss in data center servers and contribute to reduced power and cooling requirements in the data center as a whole, particularly in server farms and other dense computing environments. By minimizing power loss, energy-efficient power

REDUNDANCY PENALTY: AVAILABILITY VERSUS ENERGY USAGE

Configuring servers with redundant power supplies is common practice throughout data centers. Utilizing both power supplies helps ensure that a server can continue running if one power supply fails, or if the AC power to one of the supplies is lost (many large data centers provide dual redundant AC power delivery grids). Redundant power supplies are a characteristic element of highly available data center systems in particular.

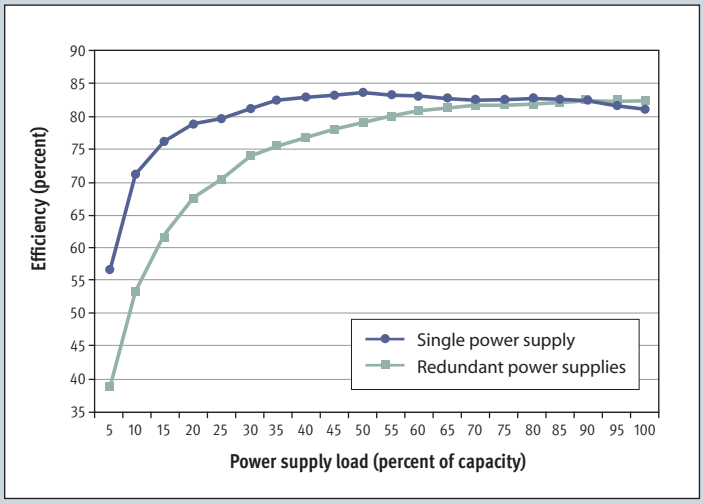

Figure A. *Energy efficiency of power supplies in single and redundant configurations*

However, this enhanced availability comes at the expense of an energy usage penalty associated with redundant power supplies. Because power supplies are inherently less efficient at lower loads, sharing loads across two power supplies reduces overall energy efficiency. For example, in a 200 W configuration, a single power supply rated at 800 W might carry the entire 200 W load, or 25 percent of its capacity. If a second power supply is added, the original 200 W load is distributed to two 800 W power supplies providing 100 W each. That means the power supplies in the redundant configuration are running at 12.5 percent of capacity. As shown in Figure A, that drops the efficiency from 83 percent for a 25 percent load to less than 65 percent for a 12.5 percent load.

The efficiency curve in Figure A is based on power supply evaluation testing performed by Dell engineers in May 2006, and is representative of current-generation redundant power supply design characteristics. Actual power supply efficiencies will vary, but the penalty for redundancy is relatively consistent across current designs.

Although high availability may trump efficiency for many data center managers, organizations that are interested in reducing energy costs should carefully consider availability requirements, or employ alternative high-availability methods, before uniformly configuring servers with redundant power supplies.

supplies can help reduce energy costs and contribute to lower TCO in the data center.

Today, many Dell products and services are available to help organizations improve efficiency throughout the data center life cycle. For example, the Dell Datacenter Capacity Planner is available online at www.dell.com/calc. This tool can help data center managers configure servers, storage, and peripherals in a drag-and-drop manner to produce reports that estimate power and cooling needs for various data center configurations.

While the Capacity Planner helps data center managers get a view of the big power and cooling picture, Dell is concerning itself with optimizing energy efficiency at every level of the data center. Energy efficiency in power supply design is one of the many facets of building an energy-conscious data center that is reliable, scalable, and highly available.

Stuart Berke *is an engineer strategist for the Dell Enterprise Product Group and currently focuses on server power efficiency and memory technologies. He was a lead systems engineer for the Dell PowerEdge 1950, PowerEdge 2900, and PowerEdge 2950 servers, and has over 20 years of experience in design and development of computing systems and servers. Stuart has a B.S. in Computer and Systems Engineering from Rensselaer Polytechnic Institute and an M.S. in Electrical Engineering from Worcester Polytechnic Institute.*

David Moss *is a senior thermal/mechanical architect with Dell and has over 20 years of experience in electronics packaging design. In his role at Dell, he holds 18 U.S. patents. David has a B.S. in Mechanical Engineering from Texas Tech University.*

Randy Randall *is a senior engineer and consultant currently working in Dell Enterprise Server AC-DC development. He has over 25 years of experience in power supply design and development. Randy has a B.S. in Electrical Engineering from Oklahoma State University and an M.S. in Electrical Engineering from the University of Texas at Austin.*

more Online

www.dell.com/**powersolutions**

QUICK LINKS

Dell energy efficiency:
www.dell.com/energy

Dell Datacenter Capacity Planner:
www.dell.com/calc

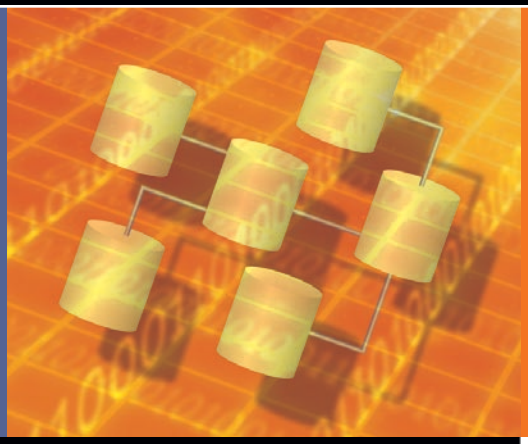

iSCSI: Changing the Economics of Storage

Part 1—Understanding iSCSI in Enterprise Environments

BY TRAVIS VIGIL

Internet SCSI (iSCSI) technology is changing the face of networked storage, providing a cost-effective and easy-to-manage alternative to Fibre Channel networks. The first article in this ongoing series explores the basics of iSCSI architecture; clarifies several common misperceptions about its performance, manageability, and security; and discusses how enterprises can implement iSCSI using Dell™ PowerVault™ and Dell/EMC storage.

Related Categories:

Dell PowerVault storage

Dell/EMC storage

Internet SCSI (iSCSI)

Storage

Storage architecture

Storage area network (SAN)

Visit www.dell.com/powersolutions for the complete category index.

Storage is a critical component of many enterprise infrastructures, and the need to support steadily growing amounts of digital content, provide 24/7 access, and comply with regulatory requirements has only increased the importance of controlling costs and maximizing the efficiency of storage resources. The most common solution to these issues is the deployment of a Fibre Channel–based storage area network (SAN), which is typically easier to manage and back up, provides higher utilization, and offers lower total cost of ownership than internal storage or direct attach storage (DAS) in large environments. However, Fibre Channel SANs also tend to have a higher cost of entry than other types of storage, and these costs combined with a steep management learning curve can place Fibre Channel SANs outside the reach of many small and medium-size enterprises. For similar reasons, these factors often prevent large enterprises from extending their Fibre Channel SANs to departments and workgroups.

Internet SCSI (iSCSI) can provide the advantages in scalability, availability, and manageability of Fibre Channel but at a lower acquisition cost, helping eliminate this barrier to entry for organizations that want to deploy a SAN without investing in a Fibre Channel infrastructure. Importantly, iSCSI also allows IT staff to work with familiar TCP/IP networks and standard Ethernet components, avoiding the need to learn how to deploy and manage specialized Fibre Channel equipment. iSCSI storage can be implemented in several ways, including as an alternative to Fibre Channel for enterprises that have not yet implemented a SAN, as a complement to an existing Fibre Channel SAN, and as a complement to network attached storage (NAS) for enterprises seeking integrated file- and block-level storage access.

This ongoing series of articles explores how iSCSI is changing the economics of storage across enterprises of all sizes—making SANs a cost-effective option for small and medium-size enterprises and offering a complement to existing Fibre Channel environments for large enterprises. This introductory article explores the basics of the technology; clarifies several common misperceptions about its performance, manageability, and security; and discusses how enterprises can implement iSCSI using Dell PowerVault and Dell/EMC storage. For more information on Dell iSCSI storage solutions, visit www.dell.com/iscsi.

Comparing traditional storage approaches

To understand the advantages offered by iSCSI technology, enterprises should understand the strengths and weaknesses of traditional approaches: internal storage, DAS, and networked storage such as NAS and Fibre Channel SANs (see Figure 1).

Internal storage. Internal storage is typically simple and inexpensive, but does not scale well. If the server's

> "iSCSI enables small and medium-size enterprises to deploy SAN technology that may have been previously out of reach, while enabling large enterprises to create a complementary storage tier for secondary servers."

maximum storage capacity or processing power is reached, administrators must add servers to meet the new requirements, potentially underutilizing either their storage or processing resources. Internal storage also offers limited high-availability options because a server failure renders that server's storage unavailable. Because each server's storage must be managed and protected separately, this type of storage can be difficult and time-consuming to administer in large environments.

Direct attach storage. DAS is a logical extension of internal storage, and typically consists of a rack of external hard drives utilized by a single server to expand capacity—although innovations such as the Dell PowerVault MD3000 have recently provided the flexibility for several servers to access the same DAS array. Other improvements—such as multiple drive options like Serial ATA (SATA) and Serial Attached SCSI (SAS), enhanced scalability and data availability, and the introduction of applications with built-in replication capabilities like Microsoft Exchange Server 2007—have further increased the usefulness of DAS. However, because the number of attached hosts is limited, DAS can still create the same sort of management difficulties as internal storage in large environments.

Networked storage. NAS and SANs enable many hosts to share storage resources, helping enhance storage utilization, simplify management, and reduce total cost of ownership compared with internal storage and DAS. NAS is primarily used for file sharing, while SANs are typically used for application data. Given that the initial entry cost for SANs is typically much higher than internal storage or DAS, deployment has primarily been restricted to large enterprises with the resources to invest in this type of infrastructure. Fibre Channel SANs can also have a steep learning curve for administrators unfamiliar with their deployment and management.

Understanding iSCSI

This is where iSCSI fills a longtime gap in enterprise storage—a networked storage environment that can offer the benefits of a SAN while avoiding the high entry costs and specialized knowledge typically required for Fibre Channel infrastructures. iSCSI enables small and medium-size enterprises to deploy SAN technology that may have been previously out of reach, while enabling large enterprises to create a complementary storage tier for secondary servers or expand their existing networked storage to include departments and workgroups, helping consolidate storage for increased capacity utilization.

The iSCSI protocol allows SCSI commands to be sent over an Ethernet network, enabling enterprises to build SANs with standard equipment such as Gigabit Ethernet switches and network interface cards (NICs). An iSCSI SAN accesses and transfers data in blocks, with each connected server seeing the remote storage array as a local hard drive. (This is a different implementation from NAS, which abstracts and organizes block data so that connected systems see the storage as files.)

Much like Fibre Channel, iSCSI networks use two types of nodes: initiators (the hosts requesting data from storage) and targets (the storage holding that data). Unlike Fibre Channel, administrators can install free initiator software, widely available from sources such as Microsoft and standard Linux® OS distributions. iSCSI also offers a variety of connectivity options. Servers can connect to iSCSI SANs using standard Ethernet NICs or, if extremely low CPU utilization is required, iSCSI host bus adapters (HBAs). Another key technology that benefits iSCSI is the TCP/IP Offload Engine (TOE); this dedicated hardware is responsible for TCP/IP processing and helps avoid using CPU cycles

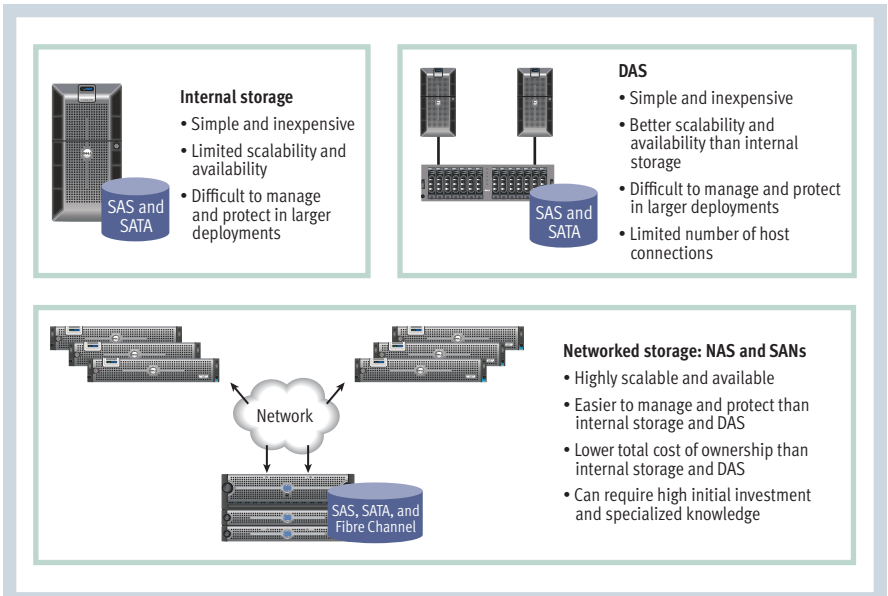

Figure 1. *Characteristics of internal, direct attach, and networked storage environments*

for network traffic. Like iSCSI HBAs, TOEs also help reduce the CPU utilization of hosts connected to an iSCSI SAN, leaving the CPU available for application processing. TOEs are a standard feature of ninth-generation Dell PowerEdge™ servers.[1]

Clarifying common misperceptions

Several common misperceptions have persisted regarding how iSCSI compares to traditional Fibre Channel deployments. With iSCSI becoming an increasingly mainstream technology and large vendors like Dell and EMC offering iSCSI products, these notions are no longer necessarily true.

Performance. The first misperception is that iSCSI cannot provide the performance necessary for enterprise applications. Many enterprises compare Fibre Channel's 4 Gbps bandwidth to Ethernet's 1 Gbps bandwidth and assume that Fibre Channel is four times faster than iSCSI. However, for many key applications, the "size of the pipe" is not the bottleneck for performance. Many key applications have random I/O data patterns, and the performance bottleneck ends up being the time it takes to write and read data from hard disk drives, not the network bandwidth.

With an appropriately configured Ethernet network, iSCSI can provide performance levels comparable to Fibre Channel for key applications such as databases and messaging—and typically at a much lower cost. Fibre Channel's bandwidth does benefit applications that use sequential I/O, such as backups to disk or streaming media.

Manageability. Another misperception is that iSCSI is more difficult to manage than Fibre Channel. Storage device management—for example, creating RAID groups and virtual disks—is typically network independent, and thus iSCSI and Fibre Channel exhibit equivalent manageability in this respect. Dell/EMC arrays even use the same management application, EMC® Navisphere® Manager, for both iSCSI- and Fibre Channel–based storage. However, because iSCSI utilizes standard Ethernet equipment and because many IT staffs are more familiar with Ethernet than Fibre Channel, iSCSI networks may be easier to manage.

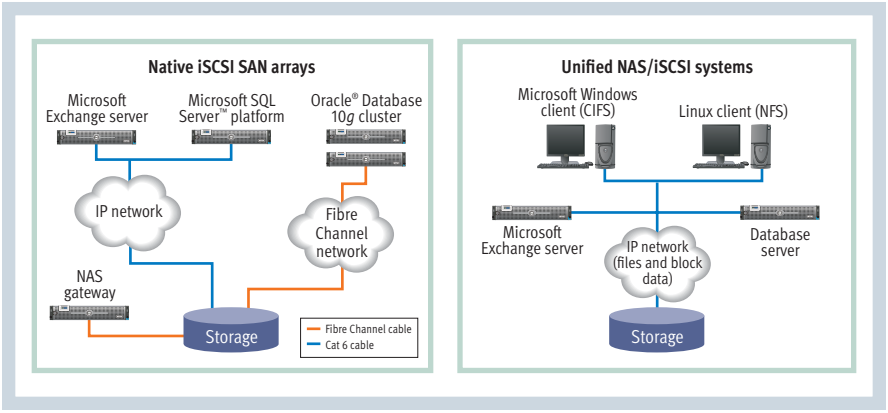

Figure 2. *Example environments using native iSCSI SAN arrays and unified NAS/iSCSI systems*

Network security. A very common misperception is that iSCSI SANs are not as secure as Fibre Channel SANs. In fact, when logically or physically separated, iSCSI networks are just as secure as Fibre Channel. An easy way to secure iSCSI traffic is to implement a dedicated iSCSI network using an industry-standard Gigabit Ethernet switch such as a Dell PowerConnect™ 2700, PowerConnect 5300, or PowerConnect 6200 series switch. If iSCSI is used over a shared corporate data network, implementing virtual LANs (VLANs) typically should be sufficient to secure iSCSI data from unwanted access.

Because iSCSI offers comparable SAN benefits at a lower cost of entry than Fibre Channel—while allowing administrators to work with familiar, industry-standard Ethernet components—many enterprises are exploring ways to implement iSCSI in their own data centers.

Implementing iSCSI with Dell PowerVault and Dell/EMC storage

A convergence of network storage technologies has allowed iSCSI to be integrated into several types of storage—SAN arrays can support both iSCSI and Fibre Channel, and NAS systems can include iSCSI functionality. Enterprises have two primary options when implementing iSCSI-based networked storage: native iSCSI SAN arrays (such as the Dell/EMC AX150i and CX3 series storage) and unified NAS/iSCSI systems (such as the Dell PowerVault NX1950).

Figure 2 illustrates these two implementations in example network environments. Native iSCSI SAN arrays are optimized for block-level storage, and are well suited both for small and medium-size enterprises that want to consolidate storage and for large enterprises that want to complement a Fibre Channel SAN with a second storage tier using iSCSI. In this type of implementation, integrated storage controller redundancy helps provide high data availability, and attaching a NAS gateway or file server to the SAN can provide file-level support.

Unified NAS/iSCSI systems are optimized for file sharing, but still provide block-level capabilities that allow administrators to connect them to application servers. In this type of implementation, clustering with an additional head unit helps provide the equivalent of storage controller redundancy for high data availability. This implementation is well suited for organizations with file-sharing and application data requirements and for those that prefer the simplicity of file- and block-level access in a single device.

Native iSCSI SAN arrays

The Dell/EMC AX150i and CX3 series are designed to meet different needs within different types of storage environments. Figure 3 summarizes some characteristic differences.

[1] For more information about TOEs in PowerEdge servers, see "Boosting Data Transfer with TCP Offload Engine Technology on Ninth-Generation Dell PowerEdge Servers," by Pankaj Gupta, Allen Light, and Ian Hameroff, in *Dell Power Solutions*, August 2006, www.dell.com/downloads/global/power/ps3q06-20060132-Broadcom.pdf.

Dell/EMC AX150i. The Dell/EMC AX150i, a cost-effective entry-level array, offers only iSCSI support. It is designed for easy installation and management, and includes several streamlined management applications: the EMC Navisphere Express management tool, the EMC SnapView™ Express snapshot application, and the EMC PowerPath® application for automatic load balancing and failover. This array is typically used for production deployments where low to moderate performance is acceptable, such as light to moderate application data (for example, a Microsoft Exchange system with fewer than 300 users) and light database, Web server, and e-mail environments.

Dell/EMC CX3 series. The iSCSI-capable arrays in the Dell/EMC CX3 series are designed to provide enhanced performance, scalability, and mirroring capabilities, and add the flexibility to support both iSCSI and Fibre Channel connections simultaneously. This dual support allows enterprises to create tiered configurations to help meet the needs of specific application workloads, helping streamline data center costs and management efficiency.

The entry-level iSCSI array of this series is the Dell/EMC CX3-10c, which is designed to be cost-effective while providing extensive functionality and high-availability features. It includes EMC PowerPath for automatic load balancing and failover, and optional EMC Navisphere Manager, Navisphere Quality of Service Manager, SnapView, SAN Copy™, and MirrorView™ software provide additional management and data protection capabilities. This array is typically used for moderate database and e-mail environments, high-speed large block data transfers for operations such as replication and disk backup, and department-level software such as multimedia and collaboration applications.

In this series, the Dell/EMC CX3-20c and CX3-40c are midrange and high-end native iSCSI arrays, respectively. They provide significantly higher capacities than the AX150i and CX3-10c to offer plenty of room for growth. Optional EMC software allows administrators to easily discover, configure, manage, and monitor these arrays and provide efficient backup and disaster recovery capabilities. Both the CX3-20c and CX3-40c are suitable for the same types of environments as the CX3-10c, but are designed to accommodate high performance and capacity requirements.

Unified NAS/iSCSI systems

The Dell PowerVault NX1950 is a networked storage solution designed to simplify storage deployment, management, and integration into existing storage environments. It can store both file and application data and is designed to work with multiple operating environments and communication protocols, including iSCSI. Its support for Network File System (NFS) and Common Internet File System (CIFS) enables the PowerVault NX1950 to work seamlessly with Microsoft Windows®, Linux, UNIX®, and Mac OS environments, and the included Microsoft Windows Unified Data Storage Server 2003 OS provides a single console to help administrators easily create and manage file shares, iSCSI targets, point-in-time snapshots, performance logs and metrics, and more.

For more information about the PowerVault NX1950, see "Deploying the Dell PowerVault NX1950 for High-Availability Storage," by Kevin Guinn and Ananda Sankaran, in *Dell Power Solutions,* February 2007, www.dell.com/downloads/global/power/ps1q07-20070138-Guinn.pdf.

Changing the economics of storage

SANs offer numerous advantages over alternatives such as internal storage and DAS, including streamlined management, enhanced scalability, and low total cost of ownership. While the cost and specialized expertise required to deploy and manage Fibre Channel SANs has primarily restricted them to large enterprises, iSCSI—with its use of standard Ethernet components and familiar management tools—has placed SANs within the reach of small and medium-size enterprises. At the same time, iSCSI SANs enable large enterprises to create cost-effective, complementary tiered storage for existing Fibre Channel infrastructures. With 10 Gigabit Ethernet on the horizon, iSCSI has the potential to surpass Fibre Channel in bandwidth over the next few years, which would make iSCSI a viable option for even the highest-performance enterprise environments.

	Supported connections/ iSCSI support	Available connectivity per array	Maximum number of SAN-connected hosts	Maximum number of drives	Maximum capacity
Dell/EMC AX150i	iSCSI/native	Four Gigabit* Ethernet ports	10	12	6 TB
Dell PowerVault NX1950	NAS (CIFS and NFS) and iSCSI/ unified	Four Gigabit Ethernet ports	16	45	13.5 TB
Dell/EMC CX3-10c	iSCSI and Fibre Channel/ native	Four iSCSI ports and four Fibre Channel ports	64	60	30 TB
Dell/EMC CX3-20c	iSCSI and Fibre Channel/ native	Eight iSCSI ports and four Fibre Channel ports	128	120	57 TB
Dell/EMC CX3-40c	iSCSI and Fibre Channel/ native	Eight iSCSI ports and four Fibre Channel ports	128	240	117 TB

*This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

Figure 3. *Dell iSCSI storage solutions*

Travis Vigil *is a senior product marketing consultant for Dell iSCSI and Dell PowerConnect solutions. He has a B.S. from Stanford University and an M.B.A. from Northwestern University's Kellogg School of Management.*

Building Highly Available Systems: The Dell PowerEdge Cluster SE600W and PowerVault MD3000

BY NAM NGUYEN
DANIEL MOGES
SHABANA M.
TRAVIS VIGIL

The Dell™ PowerEdge™ Cluster SE600W is a Serial Attached SCSI (SAS)–based cluster solution comprising Dell PowerEdge servers, Microsoft® Windows® operating systems, and the Dell PowerVault™ MD3000 storage system. Deploying this solution enables enterprises to achieve highly available service levels in a cost-effective way.

Related Categories:

Clustering

Dell PowerEdge servers

Dell PowerVault storage

Direct attach storage (DAS)

High availability (HA)

Microsoft Windows Server 2003

Serial Attached SCSI (SAS)

Storage architecture

Visit www.dell.com/powersolutions for the complete category index.

Global commerce and the demand for 24/7 access have made information availability vital to enterprise success. A lack of data access during maintenance, hardware or software failure, or any other period of downtime can lead to decreased productivity, customer dissatisfaction, and lost revenue. As high-availability technology continues to advance, both the cost and complexity associated with clustering for availability and failover have been reduced significantly, making these configurations feasible for small businesses as well as large enterprise-class data centers.

For data to be continuously available, IT systems must maintain uninterrupted data access during both planned and unplanned downtime. Host-based clustering, in conjunction with high availability at the storage level, can greatly enhance IT service levels, and represents a crucial step toward ensuring data availability and minimizing downtime problems. Allowing multiple servers (for example, a two-node cluster) shared access to data on a single storage system helps provide application availability even if one of the servers becomes unavailable.

The Dell PowerEdge Cluster SE600W solution and Dell PowerVault MD3000 storage array are designed to allow enterprises of all sizes to build highly available systems in a cost-effective way. Deploying this solution can help maintain access to critical application data even following a hardware or software failure.

Dell PowerEdge Cluster SE600W

Dell PowerEdge clusters are based on Microsoft Cluster Service software and designed to keep applications and services available during any single failure within the cluster. When Microsoft Cluster Service detects a failure, it automatically moves cluster resources from the failed cluster node to a healthy one and restarts the applications. PowerEdge clusters support cluster-aware applications such as Microsoft SQL Server™, Microsoft Exchange Server, and Oracle® Database with Oracle Fail Safe software.

In addition to application-level availability, enterprises should consider redundancy in the server-to-storage I/O path. Because a failure of any component along this path (such as a server, adapter, controller, cable, or disk drive) jeopardizes system availability, the storage system's ability to maintain data access during such a failure is a key part of a highly available storage deployment.

The PowerEdge Cluster SE600W is the latest addition to the Dell family of feature-rich clustering solutions, and the first based on Serial Attached SCSI (SAS) technology. It currently supports two-node clustering based on the components shown in Figure 1. The supported PowerEdge servers provide key high-availability features such as error-correcting code memory; software or hardware RAID for the internal drives; hot-swappable drives, power supplies, fans, and PCI slots; optional dual host bus adapters (HBAs); and optional redundant paths to the storage systems.

Operating systems	Microsoft Windows Server 2003 and Windows Server 2003 Release 2 (R2) Enterprise Edition and Enterprise x64 Edition
Servers	Dell PowerEdge models 1800, 1850, 2800, 2850, 6800, 6850, 1950, 2900, 2950, 2970, and 6950
Storage	Dell PowerVault MD3000

Figure 1. *Supported components for the Dell PowerEdge Cluster SE600W*

The PowerEdge Cluster SE600W offers three configurations:

- **Entry-level configuration:** This configuration offers only one path from each server to the storage (see Figure 2). If a component in one path fails, the server loses access to the storage; when this happens, Microsoft Cluster Service moves resources from the affected server to the other server and restarts the applications.
- **Redundant configuration with dual-port HBAs:** This configuration provides higher availability than the entry-level configuration, incorporating two ports on each SAS 5/E HBA to create two paths from each server to the storage (see Figure 3). If an HBA fails, Microsoft Cluster Service handles application migration the same way it does in the entry-level configuration. If another storage component fails, I/Os are transparently rerouted to the alternate path without affecting applications or clients.
- **Fully redundant configuration:** This configuration provides the highest availability of the three configurations, incorporating two SAS 5/E HBAs and two paths from each server to the storage (see Figure 4). Administrators can configure I/Os to use any path for load balancing. If a storage path fails, I/Os are transparently rerouted to the alternate path without affecting applications or clients.

Dell PowerVault MD3000

The Dell PowerVault MD3000 storage array, with two high-performance, active/active external RAID controllers and a mirrored cache, can be a key part of designing a cost-effective, highly available cluster system. It is based on SAS technology and offers high-availability features such as hot spares; hot-swappable drives, power supplies, cooling fans, and storage controllers; redundant storage controllers; redundant back-end paths with dual-port SAS drives; and optional premium features such as snapshots and virtual disk copy.

To appreciate the potential value of clustering with PowerVault MD3000 storage, enterprises should first understand the *host-based RAID* method of clustering. This method uses a direct connection from a storage device to each server, and the RAID controller itself resides on a PCI RAID card installed in the server. In an *external RAID* configuration, like that of the PowerVault MD3000, the RAID controller resides outside the server in the external storage system itself.

Among the issues associated with host-based RAID clustering, one of the most significant is related to the cache and performance. The cache is high-speed memory designed to reduce the time required to read or write data. In host-based RAID configurations, the cache resides on the RAID adapter card located in the server. To help ensure that in the event of a server failure the unwritten cache data is written to disk, vendors disable the use of the cache on RAID adapters to maintain cache coherency for the cluster. Unfortunately, doing

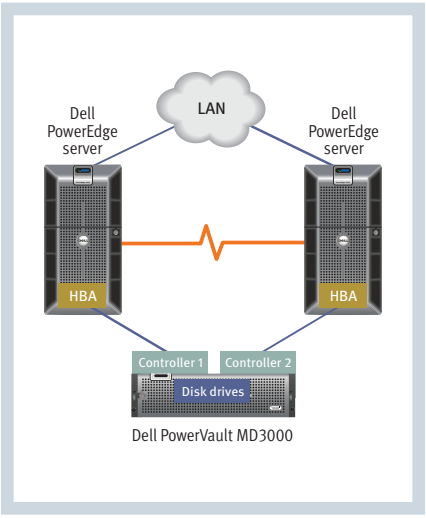

Figure 2. *Entry-level Dell PowerEdge Cluster SE600W configuration*

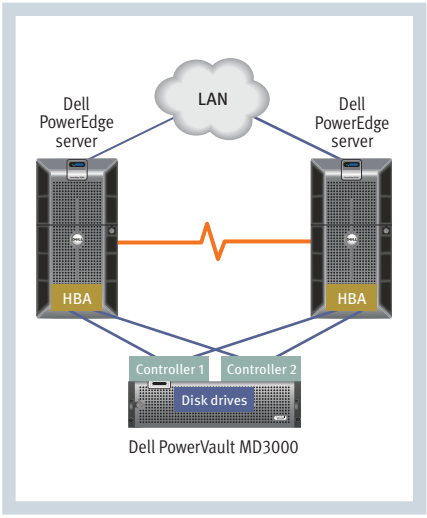

Figure 3. *Redundant Dell PowerEdge Cluster SE600W configuration with dual-port HBAs*

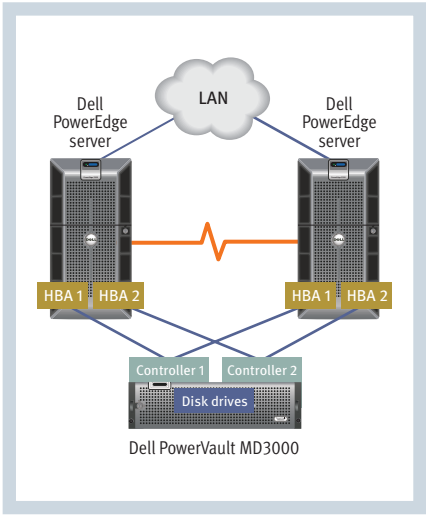

Figure 4. *Fully redundant Dell PowerEdge Cluster SE600W configuration*

so can cause a significant degradation in cluster performance.

When using external RAID, the cache resides on the external controller within the storage system, and cache coherency can be maintained without disabling the cache functionality. For this reason, external systems are preferable for a clustered configuration. Storing the cache separately from the server helps eliminate the risks associated with a server failure in a host-based RAID cluster. As an external storage system, the PowerVault MD3000 is ideal for highly available two-node clustering applications on PowerEdge servers when performance and access to data are critical.

Cluster features and advantages

The PowerEdge Cluster SE600W and PowerVault MD3000 are designed to build upon the reliability of parallel SCSI while addressing that technology's performance, reliability, and scalability limitations. Figure 5 summarizes the cluster features of the SAS-based PowerEdge Cluster SE600W compared with the parallel SCSI–based PowerEdge Cluster SE500W.

"The PowerEdge Cluster SE600W and PowerVault MD3000 are designed to build upon the reliability of parallel SCSI while addressing that technology's performance, reliability, and scalability limitations."

Performance

The PowerVault MD3000 can provide better performance than PowerVault 22*x*S storage because of its increased data transfer rate over the SAS link and its enabled write cache. While the maximum data transfer rate for Ultra320 SCSI is 320 Mbps (and the complications of the shared bus architecture make increased speeds unlikely), SAS currently offers speeds of up to 3 Gbps, and the SCSI Trade Association plans for it to eventually support up to 12 Gbps.

	Dell PowerEdge Cluster SE500W	Dell PowerEdge Cluster SE600W
HBAs	Up to two PowerEdge Expandable RAID Controller (PERC) 4/DC or PERC 4e/DC adapters per node	Up to two SAS 5/E HBAs per node
Storage	Up to four PowerVault 22*x*S enclosures	One PowerVault MD3000 enclosure connected to up to two PowerVault MD1000 enclosures
Cables	68-pin shielded P-type SCSI cable for connection to the host, up to 12 meters long	Compact cable and connector, up to 6 meters long
Hard drives	13 SCSI hard drives per enclosure	15 SAS hard drives per enclosure (with expansion to up to 45 drives by adding PowerVault MD1000 enclosures)
Redundant path to the storage system	No	Yes, including redundant storage controllers with failover capability, dual-port SAS drives, and a multipath failover driver
RAID	RAID-1, RAID-5, RAID-10, and RAID-50	RAID-0, RAID-1, RAID-5, and RAID-10
Write cache	Disabled in cluster mode	Enabled in cluster mode
Storage manageability	Dell OpenManage Storage Services and Dell OpenManage Array Manager, installed on a host directly connected to the storage	Dell Modular Disk Storage Manager with advanced features such as snapshots and virtual disk copy, installed on a management station (either in-band through the SAS link or out-of-band over the network)

Figure 5. *Feature comparison of Dell PowerEdge Cluster SE500W and PowerEdge Cluster SE600W*

In a PowerEdge Cluster SE500W, the write cache is disabled in the RAID controller to maintain cache consistency, because the RAID intelligence is part of the controller residing on the host. In a PowerEdge Cluster SE600W, the write cache is enabled in the RAID controller, because the RAID intelligence is part of the controller residing on the external storage device. Each storage controller has 512 MB of cache, which helps significantly improve overall system performance.

Reliability

When deployed in an environment using parallel SCSI and a shared bus architecture, a misbehaving device can bring down the entire system. The PowerVault MD3000 offers a point-to-point architecture in which such devices can be easily isolated without affecting the entire system. In addition, the fully redundant PowerEdge Cluster SE600W configuration provides two paths from the servers to the storage, and includes dual-port SAS drives to allow data access from either port from one of the two storage controllers. As a result, if a component in a storage path fails, the multipath software can automatically reroute the I/Os to the alternate path without affecting the applications or clients.

Scalability

Parallel SCSI can support up to 16 devices on the shared bus, while a SAS domain can support up to 16,384 devices without performance degradation. The PowerEdge Cluster SE600W can support up to 45 SAS hard drives by adding PowerVault MD1000 disk expansion enclosures to the PowerVault MD3000 using a daisy-chain topology (see Figure 6).

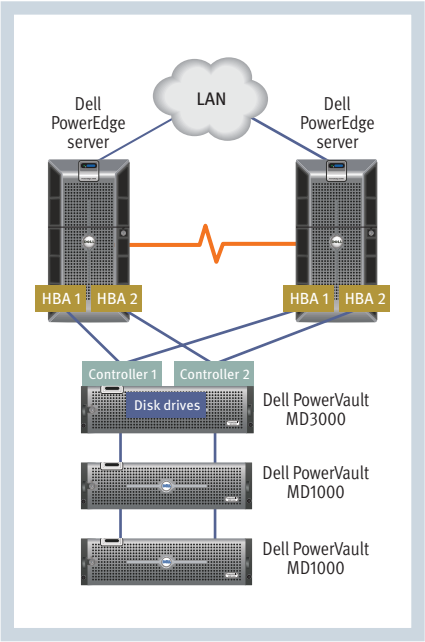

Figure 6. *Fully redundant Dell PowerEdge Cluster SE600W configuration with two daisy-chained Dell PowerVault MD1000 enclosures*

Highly available cluster solution

The Dell PowerEdge Cluster SE600W and PowerVault MD3000 take advantage of SAS technology to offer enhanced performance, reliability, and scalability, and include additional features not found in parallel SCSI–based PowerEdge Cluster SE500W configurations. This cost-effective clustering solution can benefit enterprises of all sizes that require high availability for critical applications such as databases, messaging systems, and file, print, and Web servers.

Nam Nguyen *is a senior consultant in the High-Availability Cluster Development Group at Dell, and the lead engineer for SAS, Internet SCSI (iSCSI), and Fibre Channel Dell PowerEdge Cluster products. His current interests include business continuity, clustering, and storage technologies. He has a B.S. and an M.S. in Electrical Engineering from the University of Texas at Austin.*

Daniel Moges *is a systems engineer adviser in the High-Availability Cluster Development Group at Dell. His responsibilities include developing SCSI-, iSCSI-, and Fibre Channel–based high-availability clustering products; his current interests related to high-availability clustering include enterprise storage technologies and database systems. Daniel has a B.S. in Electrical Engineering from the University of Illinois at Urbana-Champaign.*

Shabana M. *is an engineering analyst in the High-Availability Cluster Development Group at Dell. Her interests include the development of SCSI- and Fibre Channel–based clustering solutions and application performance. She has a B.E. in Computer Science and Engineering.*

Travis Vigil *is a senior product marketing consultant for Dell PowerVault storage. He has a B.S. from Stanford University and an M.B.A. from Northwestern University's Kellogg School of Management.*

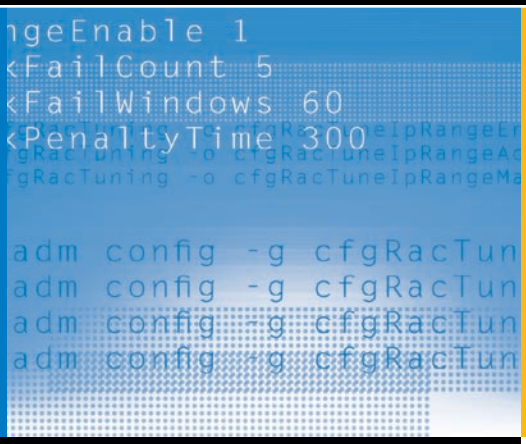

Using Reliable Datagram Sockets Over InfiniBand for Oracle Database 10*g* Clusters

BY ZAFAR MAHMOOD
ANTHONY FERNANDEZ
GUNNAR K. GUNNARSSON

Reliable Datagram Sockets (RDS) Over InfiniBand can provide a horizontally scalable, high-performance alternative to traditional vertical scaling for enterprises using Oracle® Database 10*g* and Oracle Real Application Clusters (RAC). This article discusses the advantages of using RDS Over InfiniBand to build scalable, high-performance Oracle RAC clusters with cost-effective, industry-standard Dell™ and QLogic components.

Related Categories:

Clustering

Database

Dell/EMC storage

InfiniBand

Oracle

Reliable Datagram Sockets (RDS)

Visit www.dell.com/powersolutions for the complete category index.

In the past, large database systems were often synonymous with costly mainframes. Today, however, grid computing with commodity servers can provide many advantages for large databases, including cost-effectiveness, scalability, high performance and availability, network consolidation, and simple installation and management. Two key technologies enabling this type of system for large databases are Oracle Real Application Clusters (RAC) and industry-standard grid components with efficient interconnects that provide high throughput and low latency for data traffic between components. Oracle Database 10*g* and Oracle RAC enable the creation of scalable, high-performance, highly available shared databases on clusters of cost-effective industry-standard servers, helping increase return on investment and reduce total cost of ownership.

The cluster interconnect can have a major effect on Oracle RAC performance. In Oracle RAC systems with interconnect-intensive workloads, Gigabit Ethernet can often become a bottleneck for high-volume cluster messaging and Oracle Cache Fusion traffic between nodes for many applications. InfiniBand, in contrast, can provide significant advantages in both raw bandwidth and reduced latency compared with Gigabit Ethernet, and typically can provide higher performance than Gigabit Ethernet for Oracle RAC systems.

Although support problems and the lack of a standard protocol has historically made implementing InfiniBand for clusters of this size a challenge, Oracle Database 10*g* Release 2 (R2) and the 10.2.0.3 patch set support a cluster interconnect protocol developed by Oracle and QLogic specifically for Oracle RAC called Reliable Datagram Sockets (RDS), which is agnostic to underlying Remote Direct Memory Access (RDMA)–capable devices and simplifies implementation. This protocol can work over either an RDMA-capable Ethernet network interface card (NIC) or an InfiniBand host channel adapter (HCA).

The RDS protocol uses InfiniBand delivery capabilities that offload end-to-end error checking to the InfiniBand fabric, which frees processor cycles for application processing and enables significant increases in processor scaling compared to Gigabit Ethernet implementations. The key advantages of using this method with Oracle Database 10*g* and Oracle RAC can include high bandwidth and availability, low latency and processor utilization, reliable packet delivery with no discards or retransmissions, ease of use, and a simplified infrastructure compared with Gigabit Ethernet. Figure 1 illustrates the architecture differences between traditional interconnect protocols and RDS.

Oracle RAC configurations: Gigabit Ethernet and InfiniBand

Typical Oracle RAC 10*g* configurations, such as those using Gigabit Ethernet, require three distinct networks:

- **Dedicated, secure Oracle RAC cluster interconnect:** Provides data coherency between multiple servers to scale out the database; interconnect bandwidth, latency, and processor overhead are critical factors in database scaling and performance
- **Storage area network (SAN):** Provides access to shared storage resources supporting the database cluster; these connections are sensitive to latency and available I/Os per second
- **Public network for client and application tier:** Provides network communications between the database tier and the client and application tiers that require the data; these connections are sensitive to processor overhead from transmission protocols

These three separate networks, with redundant components and connections, can require four to six NICs and host bus adapters (HBAs) in each server in the database tier, increasing network complexity and cost. An additional challenge in building Oracle RAC clusters is that rack-mount and blade servers include a limited number of PCI slots, typically making it difficult or impossible to construct clusters with the necessary I/O connectivity, throughput, and availability. And because these same servers now include increased component density and multi-core technology, their processing capacity creates high demands on cluster interconnects and I/O resources.

An InfiniBand infrastructure can be implemented in two ways:

- RDS Over InfiniBand only for the interconnect traffic, which can reduce latency and increase throughput for cluster messaging and Oracle Cache Fusion traffic between Oracle RAC nodes
- RDS Over InfiniBand for the interconnect traffic, and SCSI RDMA Protocol (SRP) for the SAN running over the InfiniBand infrastructure

In the second implementation, a Fibre Channel gateway allows SRP to translate InfiniBand traffic into Fibre Channel Protocol (FCP) frames to utilize the Fibre Channel SAN. This architecture helps greatly increase throughput while reducing cabling, component complexity, and latency. In addition, applications can take advantage of the InfiniBand infrastructure's inherent RDMA capability by utilizing Socket Direct Protocol (SDP) for zero-copy data transfers between application servers and database servers.

Oracle RAC scalability in a test environment

In August 2006, the Dell Database and Application Solutions engineering team evaluated the scalability of different Oracle RAC configurations by building an eight-node Oracle RAC 10*g* cluster on Dell PowerEdge™ 1850 servers, then running the Transaction Processing Performance Council TPC-H workload with both Gigabit Ethernet and RDS Over InfiniBand to evaluate total runtime and average response time. The RDS Over InfiniBand tests used the second implementation described in the preceding section, deploying RDS Over InfiniBand for the interconnect traffic and SRP for the SAN running over the InfiniBand infrastructure. In both tests the storage consisted of two Dell/EMC CX700 arrays with 60 spindles each. Figure 2 shows the test configuration; Gigabit Ethernet was used as the interconnect in the first round of tests, and was then replaced by RDS Over InfiniBand. Figure 3 summarizes the hardware and software used in the test environment.

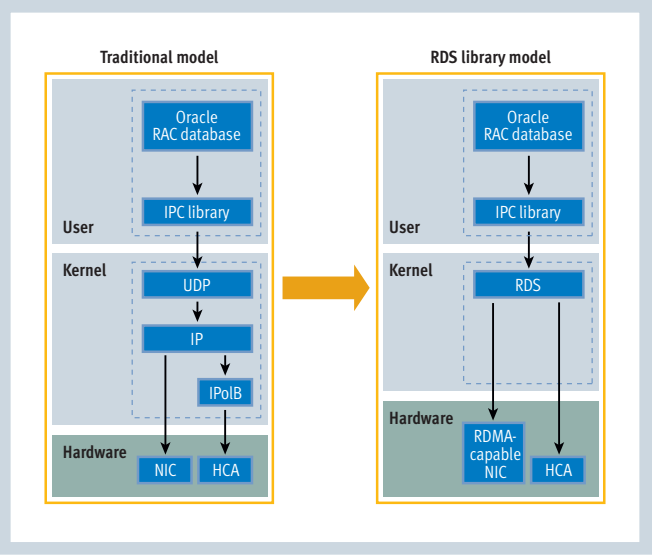

Figure 1. *Architectures for traditional interconnect and Reliable Datagram Sockets protocols*

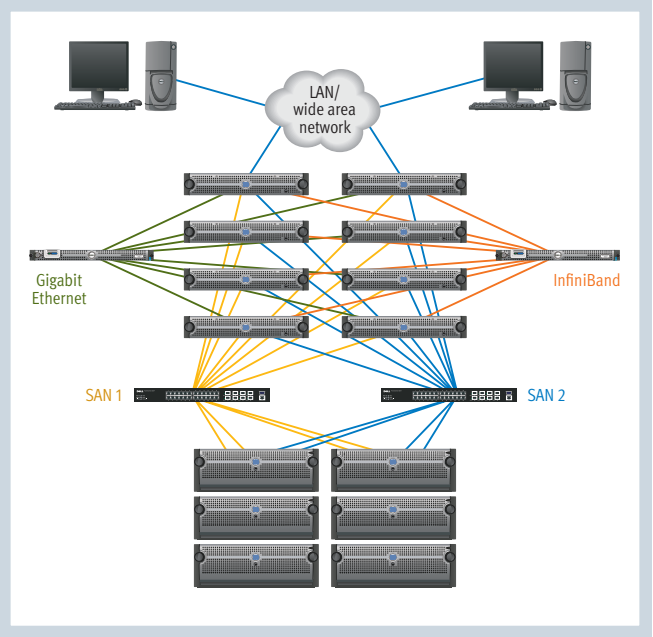

Figure 2. *Test configuration*

IMPLEMENTING RDS OVER INFINIBAND

The Dell Database and Application Solutions engineering team conducted five tests to compare the scalability of Oracle Database 10*g* R2 clusters using various components. For test 5, the Gigabit Ethernet interconnect was replaced by an InfiniBand infrastructure. The process of replacing a Gigabit Ethernet interconnect with InfiniBand is simple and does not require reinstalling Oracle binaries to make changes to the data. Enabling the RDS Over InfiniBand fabric can be accomplished with the following steps:

1. After cabling the InfiniBand HCAs and switches, install the required InfiniBand drivers and RDS support libraries on each host. These include the InfiniBand network stack, fast fabric, IP Over InfiniBand (IPoIB), and RDS drivers. Use the `ifconfig` command to ensure that the InfiniBand interface appears in the host network interface list, as shown in the `ib1` section in Figure A. Oracle Database 10*g* should initially be configured to use IPoIB to validate Oracle RAC functionality over the InfiniBand fabric.
2. Define the set of host names and associated IP addresses that will be used for the Oracle RAC private Interprocess Communication (IPC) network and specify these in /etc/hosts along with IPoIB host names, as shown in Figure B. *Note:* Figure B assumes a single InfiniBand connection to each host; however, redundant InfiniBand connections, if configured in IPoIB, are used automatically as part of RDS failover.
3. Define the set of host names that make up the Oracle RAC cluster and start Oracle Database.
4. Once Oracle RAC cluster operation is validated over Gigabit Ethernet, shut down Oracle Database and modify the /etc/hosts files on each node to reference the IPoIB interface IP addresses for the Oracle RAC private interconnect, as shown in Figure C.
5. Through the Oracle Interface Configuration Tool (oifcfg), change the cluster interconnect interface on each node with the following three commands (note that oifcfg requires Oracle Cluster Ready Services to be running):

```
oifcfg getif -global
oifcfg delif -global <ifname, ex.  bond0>
oifcfg setif -global ib1/192.168.1.0:cluster_interconnect
```

Then restart Oracle Database. Oracle RAC should now utilize InfiniBand, with User Datagram Protocol (UDP) IPC traffic passed using IPoIB. Confirm IPoIB use through IPoIB interface statistics or by checking switch port statistics through the switch graphical user interface.

6. After validating the Oracle RAC cluster operation over IPoIB, shut down Oracle Database and build the Oracle RAC IPC library (libskgxp.so) by performing the following commands (as user "oracle") on each Oracle RAC node:

```
[root@node3 ~]# cd $ORACLE_HOME/rdbms/lib
```

```
bond0   Link encap:Ethernet  HWaddr 00:00:00:00:00:00
        inet addr:192.168.0.34  Bcast:192.168.0.255  Mask:255.255.255.0
        inet6 addr: fe80::200:ff:fe00:0/64 Scope:Link
        UP BROADCAST RUNNING MASTER MULTICAST  MTU:1500  Metric:1
        RX packets:0 errors:0 dropped:0 overruns:0 frame:0
        TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:0
        RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)

eth0    Link encap:Ethernet  HWaddr 00:14:22:21:E7:F1
        inet addr:155.16.5.34  Bcast:155.16.255.255  Mask:255.255.0.0
        inet6 addr: fe80::214:22ff:fe21:e7f1/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        RX packets:38093292 errors:0 dropped:0 overruns:0 frame:0
        TX packets:3080668 errors:0 dropped:0 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:2979697141 (2.7 GiB)  TX bytes:318248063 (303.5 MiB)
        Base address:0xdcc0 Memory:fe4e0000-fe500000

eth0:2  Link encap:Ethernet  HWaddr 00:14:22:21:E7:F1
        inet addr:155.16.5.134  Bcast:155.16.255.255  Mask:255.255.0.0
        UP BROADCAST RUNNING MULTICAST  MTU:1500  Metric:1
        Base address:0xdcc0 Memory:fe4e0000-fe500000

ib1     Link encap:Ethernet  HWaddr 06:06:6A:00:6E:6D
        inet addr:192.168.1.34  Bcast:192.168.1.255  Mask:255.255.255.0
        inet6 addr: fe80::406:6aff:fe00:6e6d/64 Scope:Link
        UP BROADCAST RUNNING MULTICAST  MTU:2044  Metric:1
        RX packets:21764931 errors:0 dropped:0 overruns:0 frame:0
        TX packets:21672564 errors:0 dropped:2 overruns:0 carrier:0
        collisions:0 txqueuelen:1000
        RX bytes:4786497450 (4.4 GiB)  TX bytes:2947031664 (2.7 GiB)
```

Figure A. *Host network interface list from the ifconfig command*

```
# GE RAC private network
# bond0 - Ethernet RAC private IPC

192.168.0.33    node3-ge        node3-priv
192.168.0.34    node4-ge        node4-priv
192.168.0.35    node5-ge        node5-priv
192.168.0.36    node6-ge        node6-priv
192.168.0.37    node7-ge        node7-priv
192.168.0.38    node8-ge        node8-priv
192.168.0.39    node9-ge        node9-priv
192.168.0.40    node10-ge       node10-priv
192.168.0.41    node11-ge       node11-priv
192.168.0.42    node12-ge       node12-priv

# InfiniBand
# ib1 - InfiniBand IPC (not in use for RAC)
192.168.1.100   sst9024

192.168.1.33    node3-ib
192.168.1.34    node4-ib
192.168.1.35    node5-ib
192.168.1.36    node6-ib
192.168.1.37    node7-ib
192.168.1.38    node8-ib
192.168.1.39    node9-ib
192.168.1.40    node10-ib
192.168.1.41    node11-ib
192.168.1.42    node12-ib
```

Figure B. *Set of host names, IPoIB host names, and their associated IP addresses used for the Oracle RAC private IPC network*

```
# GE RAC private network
# bond0 Ethernet RAC private IPC (no longer
  used for RAC)

192.168.0.33    node3-ge
192.168.0.34    node4-ge
192.168.0.35    node5-ge
192.168.0.36    node6-ge
192.168.0.37    node7-ge
192.168.0.38    node8-ge
192.168.0.39    node9-ge
192.168.0.40    node10-ge
192.168.0.41    node11-ge
192.168.0.42    node12-ge

# InfiniBand
# ib1 - InfiniBand RAC IPC (now used for RAC)

192.168.1.100   sst9024

192.168.1.33    node3-ib    node3-priv
192.168.1.34    node4-ib    node4-priv
192.168.1.35    node5-ib    node5-priv
192.168.1.36    node6-ib    node6-priv
192.168.1.37    node7-ib    node7-priv
192.168.1.38    node8-ib    node8-priv
192.168.1.39    node9-ib    node9-priv
192.168.1.40    node10-ib   node10-priv
192.168.1.41    node11-ib   node11-priv
192.168.1.42    node12-ib   node12-priv
```

Figure C. */etc/hosts file edited to reference IPoIB interface IP addresses for the Oracle RAC private interconnect*

```
[root@node3 ~]# make -f ins_rdbms.mk ipc_rds
```

Restart Oracle Database. Oracle RAC should now utilize InfiniBand and RDS.

7. To confirm that Oracle Database is utilizing RDS, look for the `IPC version` string in the Oracle Database alert logs, as follows:

```
[root@node4 ~]# vi /opt/oracle/admin/dwdb/bdump/
  alert_dwdb2.log
```

Here, `dwdb` is the name of the Oracle RAC 10*g* database. When RDS is enabled for Oracle RAC, the log contains an entry similar to the one shown in Figure D indicating the interconnect protocol and IP address associated with the InfiniBand HCA.

```
Cluster communication is configured to use the
  following interface(s) for this instance
  192.168.1.34
Sat Aug 14 12:46:40 2010
cluster interconnect IPC version:Oracle RDS/IP
  (generic)
IPC Vendor 1 proto 3
  Version 1.0
```

Figure D. *Oracle Database 10g alert log showing IPC version string*

Storage test configuration

The Dell/EMC CX700 storage was laid out so that the tablespaces spanned across the maximum number of disks utilizing both storage processors on both storage arrays, and both storage arrays were configured so that 100 percent of the cache was allocated to read ahead, to accommodate the sequential read-intensive nature of data warehousing queries. Each storage processor had 4 GB of cache, all of which was made available for read cache after data load.

Oracle Database 10*g* performs I/O in sizes specified by the value of the DB_BLOCK_SIZE and DB_FILE_MULTIBLOCK_READ_COUNT parameters. These parameters were set to 8 KB and 128, respectively, to achieve the I/O size of 1 MB (8 KB × 128). Each logical unit (LUN) was created using 16 spindles with a stripe size of 64 KB, matching the Oracle Database I/O size to optimize large sequential read performance.

Servers	Eight Dell PowerEdge 1850 servers
Processors	• Single-core Intel® Xeon® processors at 2.8 GHz with a 2 MB L2 cache and 800 MHz frontside bus • Dual-core Intel Xeon processors at 2.8 GHz with two 2 MB caches and an 800 MHz frontside bus
Memory	8 GB
I/O slots	Two PCI Extended (PCI-X) slots
LAN NIC	Intel Gigabit* Ethernet PCI-X
HBA	Dual-port QLogic QLA2342
Cluster interconnects	• *Gigabit Ethernet tests:* Two Intel Gigabit Ethernet interconnects • *RDS Over InfiniBand tests:* One dual-port QLogic 9000** double data rate (DDR) HCA
Switches	• *Gigabit Ethernet tests:* Dell PowerConnect™ 5224 switch • *RDS Over InfiniBand tests:* QLogic 9024** DDR switch
Storage	Two Dell/EMC CX700 arrays
OS	Red Hat® Enterprise Linux® AS 4 Update 3
Software	• Oracle Database 10*g* R2 Enterprise Edition with 10.2.0.2 patch set • EMC PowerPath 4.5 • QLogic QuickSilver** host access software version 3.3.0.5.2 (for RDS Over InfiniBand tests) • Quest Benchmark Factory, Spotlight on RAC, and Toad for Oracle • TPC-H 300 GB database using a scaling factor of 300

*This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

**At the time of the Dell tests, the QLogic 9000 HCA, 9024 switch, and QuickSilver software were products of SilverStorm Technologies. QLogic acquired SilverStorm in November 2006.

Figure 3. *Hardware and software used in the test environment*

Tests

The test team performed five tests, as shown in Figure 4, and compared total runtime and average response time for each test. Test 5 used RDS Over InfiniBand instead of Gigabit Ethernet for the cluster interconnect; for more information about how the test team reconfigured the environment for this test, see the "Implementing RDS Over InfiniBand" sidebar in this article.

The goal of the testing was to determine the relative scalability of Oracle Database 10*g* R2 clusters running a decision support system (DSS) workload for each test, which involved scaling the Oracle cluster from one to eight nodes using Oracle parallel execution and partitioning. The TPC-H database size was 300 GB, and the large tables were partitioned and sub-partitioned based on range and hash keys. Data was spread across the two Dell/EMC CX700 storage arrays as described in the preceding section.

Test results

Figures 5 and 6 show the total runtimes and average response times when using Gigabit Ethernet as the cluster interconnect in the test environment. This data has been normalized (by dividing all of the data by a constant) to show relative performance gains and does not represent actual test results.

While test 4 was running, the cluster interconnect began showing large cluster latency and Oracle consistent-read global-cache transfer times. Test 5 repeated this test, but used RDS Over InfiniBand as the cluster interconnect in place of Gigabit Ethernet. Figure 7 shows the average response times for test 4 (with Gigabit Ethernet) and test 5 (with RDS Over InfiniBand) for three TPC-H queries that seemed to be interconnect intensive. This data has also been normalized (by dividing all of the data by a constant) to show relative performance gains and does not represent actual test results. Adding the times for these three queries together for each interconnect shows that RDS Over InfiniBand provided an average performance gain of 33 percent over Gigabit Ethernet for these queries.

Test	Parallelism	Streams	Nodes	Processors	Cluster interconnect
1	Node-level	4	1	Single-core Intel Xeon	Gigabit Ethernet
2	Node- and cluster-level	4	4	Single-core Intel Xeon	Gigabit Ethernet
3	Node- and cluster-level	4	4	Dual-core Intel Xeon	Gigabit Ethernet
4	Node- and cluster-level	4	8	Dual-core Intel Xeon	Gigabit Ethernet
5	Node- and cluster-level	4	8	Dual-core Intel Xeon	RDS Over InfiniBand

Figure 4. *Tests used to evaluate Oracle RAC cluster scalability for different configurations*

Oracle RAC best practices

Several best practices can help administrators configure and scale Oracle RAC clusters to increase performance and availability in enterprise environments:

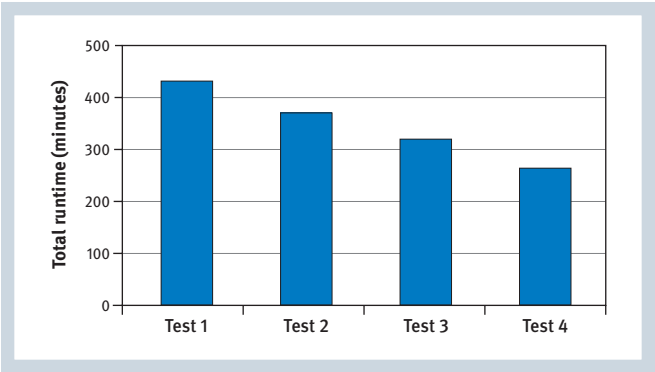

Figure 5. *Normalized total runtimes in tests using Gigabit Ethernet interconnects*

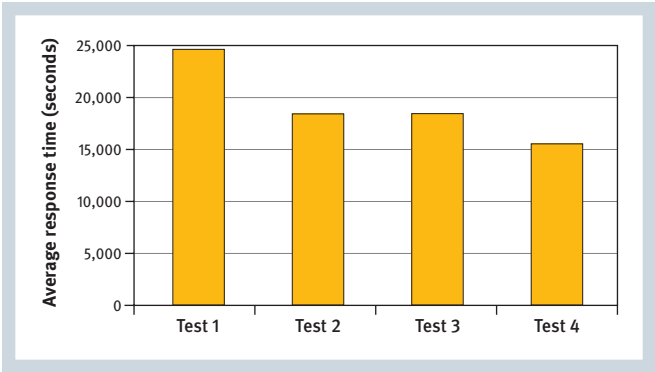

Figure 6. *Normalized average response times in tests using Gigabit Ethernet interconnects*

- Establish a baseline on a single node to provide a point of comparison when measuring performance increases.
- Use Oracle parallel execution to optimize I/O subsystems for parallelism. Adding new I/O paths (HBAs) as needed can mitigate high I/O waits, and administrators can use multipath software such as the EMC® PowerPath® and Microsoft® Multipath I/O applications for I/O balancing and failover across I/O paths.
- Use Oracle parallel execution to enable node-level parallelism with two parallel threads per processor, and to add nodes and enable cluster-level parallelism.
- For processor-intensive queries, scale up by adding processing power to existing nodes (for example, by upgrading to multi-core processors), which can be more efficient than scaling out by adding nodes. In contrast, I/O throughput bottlenecks (which are typical in data warehouses) require additional I/O channels—meaning that scaling out can be more efficient than scaling up.
- Use NIC bonding to provide interconnect scalability.
- If using Gigabit Ethernet, use jumbo frames for the cluster interconnect.
- Use a high-bandwidth, low-latency interconnect such as InfiniBand for Oracle RAC configurations with more than eight nodes or for database applications with high interconnect demand (for example, applications for which Oracle Enterprise Manager reflects significant cluster waits).

Scalable Oracle RAC clusters

Horizontally scalable InfiniBand-enabled Oracle RAC 10*g* clusters can provide a viable alternative to vertical scaling through traditional symmetric multiprocessing platforms. Deploying such clusters enables enterprises to achieve the processing power they need using cost-effective, industry-standard commodity servers, and to add capacity incrementally as application demands increase by adding cluster nodes. Oracle RAC clusters with InfiniBand also include effective high-availability features: unlike mainframes, Oracle RAC 10*g* clusters with *n*+1 or more database servers and a redundantly configured InfiniBand network can deliver a scalable, high-performance database platform with no single points of failure, enabling failed servers or switches to be replaced without application interruption or impact to users. Implementing these types of configurations helps provide scalable, high-performance, high-availability database clusters in enterprise data centers.

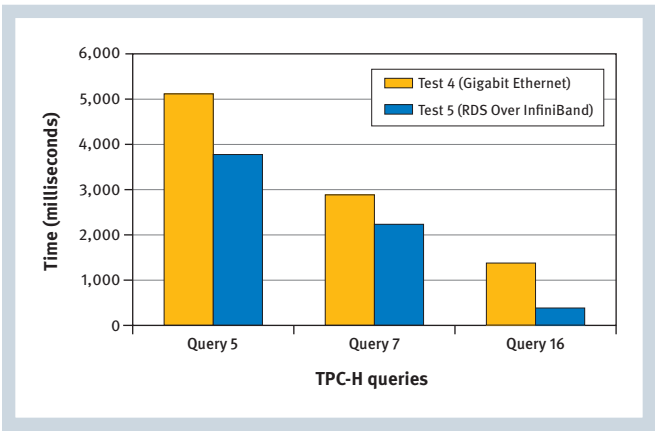

Figure 7. *Normalized average response times for three TPC-H queries using Gigabit Ethernet and RDS Over InfiniBand interconnects*

Zafar Mahmood *is a senior consultant in the Dell Database and Application Solutions team in Enterprise Solutions Engineering, Dell Product Group. Zafar has a B.S. and an M.S. in Electrical Engineering, with specialization in Computer Communications, from the City University of New York.*

Anthony Fernandez *is a senior analyst with the Dell Database and Application Solutions team in Enterprise Solutions Engineering, Dell Product Group. His focus is on database optimization and performance. Anthony has a bachelor's degree in Computer Science from Florida International University.*

Gunnar K. Gunnarsson *is a program manager and the Oracle global alliance manager at QLogic Corporation. Gunnar has a Bachelor of Electrical Engineering from the University of Delaware, a Master of Engineering with specialization in computer architecture from the Pennsylvania State University, and an M.B.A. from the Wharton School of the University of Pennsylvania.*

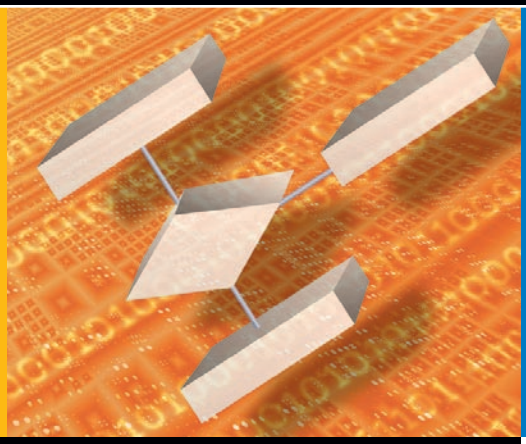

Deploying Dell PowerVault MD3000 Storage in Oracle Database 10*g* Cluster Environments

BY MAHMOUD AHMADIAN

CHETHAN KUMAR

The Dell™ PowerVault™ MD3000 storage array offers cost-effective enterprise storage that can scale to help meet future storage and application server requirements. This article discusses the benefits possible when deploying the PowerVault MD3000 along with Oracle® Database 10*g* with Oracle Real Application Clusters and Dell PowerEdge™ servers.

Related Categories:

Application servers

Characterization

Dell PowerEdge servers

Dell PowerVault storage

Microsoft Windows Server 2003 x64 Editions

Oracle

Performance

RAID

Serial Attached SCSI (SAS)

Visit www.dell.com/powersolutions for the complete category index.

Using cost-effective storage with enterprise-class performance, manageability, and scalability for databases can be important to many enterprises, particularly small and medium-size ones. The Dell PowerVault MD3000 offers such a storage array: easy to configure and manage, the PowerVault MD3000 is expandable to provide room for growth and includes features such as redundant storage paths, path failover, load balancing, and clustering support. This article discusses the potential benefits of using PowerVault MD3000 storage with Oracle Database 10*g* Release 2 (R2) and Oracle Real Application Clusters (RAC), and compares this storage array's performance with typical Fibre Channel–based storage.

Understanding Dell PowerVault MD3000 features

The Dell PowerVault MD3000 uses Serial Attached SCSI (SAS) technology, which is well suited for critical applications and provides several features and benefits not available with ATA and SCSI. One example is the point-to-point topology of SAS, which enables dedicated and scalable throughput between devices. In addition, its full duplex communication of 300 MB/sec is almost twice that of Ultra320 SCSI because of the half-duplex nature of parallel SCSI. Because of their fast throughput, SAS disks can be used as external physical disks in mass storage devices requiring high I/O access. Serial connectivity also helps increase RAID volume creation and rebuild performance, and thin SAS cables help simplify cable and temperature management.

The PowerVault MD3000 has two RAID controllers with redundant power supplies and cooling fan modules, along with two input ports for host connections and one port for expansion. Other key features are designed to offer data protection, flexible storage partitioning, multiple management options, expansion scalability, and virtual disk snapshots and virtual disk copy.

Data protection

The PowerVault MD3000 offers RAID support to help provide fault tolerance and protect data. It supports the following RAID levels:

- **RAID-0:** This level provides the fastest performance of the supported RAID levels, but no redundancy. It is typically used for noncritical data.
- **RAID-1 and RAID-10:** These levels provide fast performance and the highest data availability of the supported RAID levels. They are typically used for accounting, payroll, and financial applications. In the Dell Modular Disk Storage Manager software, RAID-10 is automatically used when four or more physical disks are selected.
- **RAID-5:** This level is typically used in multiuser environments with a high proportion of read activity and a small typical I/O size, such as file, database, Web, e-mail, news, and intranet servers.

The PowerVault MD3000 also enables administrators to configure a separate physical disk as a hot spare to provide additional fault tolerance. This disk does not contain any data, but acts as a backup in case a physical disk fails in a RAID-1, RAID-10, or RAID-5 virtual disk.

Flexible storage partitioning

The PowerVault MD3000 provides storage partitioning capabilities for virtual disk management. A storage partition is a logical entity consisting of one or more virtual disks that can be accessed by a single host or shared among hosts that are part of a host group. Partitions can be useful when specific hosts must access specific virtual disks in the storage array or when hosts with different operating systems are attached to the same storage array. In the latter case, administrators must create a separate storage partition for each host type.

Multiple management options

The PowerVault MD3000 offers both in-band and out-of-band management, which differ in the type of connection they use for sending commands and receiving event updates. With in-band management, commands, events, and data all travel through host-to-controller SAS interface cables. With out-of-band management, data is separated from commands and events: data travels through host-to-controller SAS interface cables, while commands and events travel through Ethernet cables. Dell best practices recommend using both types of management.

An Ethernet port on each RAID controller in the PowerVault MD3000 provides the out-of-band management interface. When using out-of-band management, administrators must set the network configuration for each RAID controller module, including its IP address, subnet mask, and gateway. If administrators are using a Dynamic Host Configuration Protocol (DHCP) server, they can enable automatic network configuration; if not, they must configure network settings manually.

Expansion scalability

Administrators can add up to three Dell PowerVault MD1000 SAS enclosures to the PowerVault MD3000, allowing a total of up to 60 physical disk drives. Using an average SAS disk size of 146 GB, this setup provides up to 8.8 TB of raw disk space (or up to 4.4 TB of fault-tolerant disk space in a RAID-10 configuration), which exceeds the typical requirements of small and medium-size businesses. RAID-10 virtual disks configured on the PowerVault MD3000 can survive double disk failures, adding another layer of protection against downtime.

Virtual disk snapshots and virtual disk copy

The PowerVault MD3000 provides virtual disk snapshots and virtual disk copy as add-on premium features.

Virtual disk snapshots. A virtual disk snapshot is a point-in-time image of a virtual disk in a storage array. It is not an actual virtual disk containing data; rather, it is a reference to the data that was contained on a virtual disk at a specific time. A snapshot is the logical equivalent of a complete physical copy, but it can be created more quickly and uses less disk space than a physical copy.

The virtual disk on which the snapshot is based, called the source virtual disk, must be a standard virtual disk in the storage array. Typically, a snapshot is created so that an application can access the snapshot and read the data while the source virtual disk remains online and accessible (although no I/O requests are permitted on the source virtual disk while the snapshot is being created).

A virtual disk snapshot repository containing metadata and copy-on-write data is automatically created along with a snapshot. Because only data that has changed since the time of the snapshot is stored in this repository, it uses less disk space than a full physical copy. After the repository is created, I/O write requests to the source virtual disk can resume. Then, before a data block on the source virtual disk is modified, the contents of the block to be modified are copied to the repository for safekeeping. Because the repository stores copies of the original data in these data blocks, further changes to these data blocks write only to the source virtual disk.

Virtual disk copy. A virtual disk copy is a copy pair that creates a copy of a virtual disk, with the source and target virtual disks on the same storage array. The source virtual disk can be a standard virtual disk, a virtual disk snapshot, or the source virtual disk of a snapshot. The target virtual disk can be a standard virtual disk or the source virtual disk of a failed or disabled snapshot.

A source virtual disk accepts host I/O read activity and stores the data until it is copied to the target virtual disk. When a virtual disk copy is started, all data is copied to the target virtual disk, and the source virtual disk permissions are set to read-only until the virtual disk copy is complete. After the virtual disk copy is complete, the source virtual disk becomes available to host applications for write requests. To help prevent any data corruption, best practices recommend not accessing a source virtual disk that is participating in a virtual disk copy while the copy is in progress.

Virtual disk copy enables administrators to do the following:

- **Copy data to larger-capacity physical disks:** As the storage requirements for a virtual disk change, administrators can copy data to a virtual disk in a disk group in the same storage array that uses drives with larger capacity than the original drives.
- **Restore snapshot data to source virtual disks:** Administrators can restore the data from a virtual disk snapshot and then copy the restored data to the original source virtual disk.
- **Create backup copies:** Administrators can create a backup of a virtual disk by copying data from one virtual disk to another in the same storage array, minimizing the time that the source virtual disk is unavailable to host write activity. They can use the target virtual disk as a backup for the source virtual disk, as a resource for system testing, or as a means for copying data to another device, such as a tape drive.

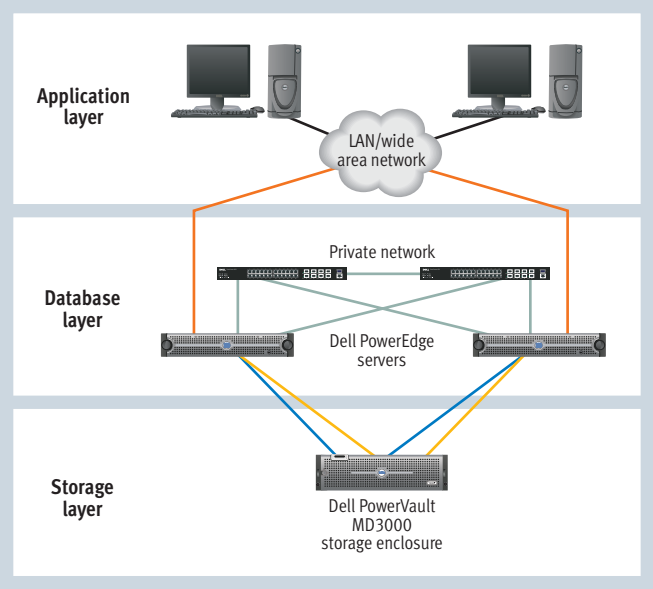

Figure 1. *Database cluster configuration using Dell PowerVault MD3000 storage*

- **Recover from backup copies:** Administrators can use the Edit Host-to-Virtual Disk Mappings feature in Dell Modular Disk Storage Manager to help recover data from a backup virtual disk. The Mappings option enables un-mapping the source virtual disk from its host and then mapping the backup virtual disk to the same host.

Using Dell PowerVault MD3000 storage with Oracle Database 10*g* and Oracle Real Application Clusters

Oracle Database 10*g* Standard Edition is well suited for medium-size enterprise environments, and it includes Oracle RAC capabilities to help protect against hardware failures and provide the flexibility to scale up hardware resources. It is also easy to install and configure, and includes its own cluster and storage management capabilities. Because Standard Edition is built from the same code base as Enterprise Edition, it can easily be upgraded to Enterprise Edition if necessary, thus providing scalability and helping reduce total cost of ownership.

Enterprises using Oracle Database 10*g* with Oracle RAC can take advantage of the PowerVault MD3000 for database storage as part of a validated two-node Dell PowerEdge server cluster configuration (see Figure 1). Each node in this configuration can be a PowerEdge 1850, PowerEdge 2800, PowerEdge 2850, or PowerEdge 6850 server with dual-core Intel® Xeon® processors; a PowerEdge 1950, PowerEdge 2900, or PowerEdge 2950 server with dual- or quad-core Intel Xeon processors; or a PowerEdge 6950 server with up to four AMD Opteron™ processors per node. Both Oracle RAC nodes share the PowerVault MD3000 as common storage connected through dual SAS storage paths. The PowerVault MD3000 includes dual storage processors, and each node is connected to both processors. This configuration is supported for systems running the Microsoft® Windows Server® 2003 Standard x64 Edition and Enterprise x64 Edition operating systems with Service Pack 2 (SP2).

High-availability clustering

In the validated Dell/Oracle configuration, the PowerVault MD3000 can connect to the cluster nodes in one of two ways, as shown in Figure 2: using two paths through a single Dell SAS 5/E host bus adapter (HBA) on each node, or using a single path through two SAS 5/E HBAs on each node. Both setups include a redundant path to the storage to enable path failover, but the additional HBA in the dual-HBA configuration enables higher availability than the single-HBA configuration can provide.

Both RAID controllers on the PowerVault MD3000 forward I/O requests to their respective virtual disks, and if one of the controllers fails, the requests are rerouted through the other RAID controller. If an HBA port on a cluster node fails, I/O paths fail over to the other HBA port (in a single-HBA configuration) or to the other HBA (in a dual-HBA configuration) in the node, and virtual disks fail over to the RAID controller in the storage. However, the advantage of a dual-HBA configuration is that administrators can quickly change the physical connection from the failed port to the other functional port on the HBA—allowing the cluster nodes to continue using all four paths until the SAS controller is replaced. Multipath drivers such as Microsoft Multipath I/O installed on host systems that access the storage array provide I/O path failover capabilities. Thus, in case an HBA, I/O path, or RAID controller fails, multipath drivers can automatically fail over the virtual disks to other available paths or RAID controllers, helping ensure disk accessibility for the cluster nodes.

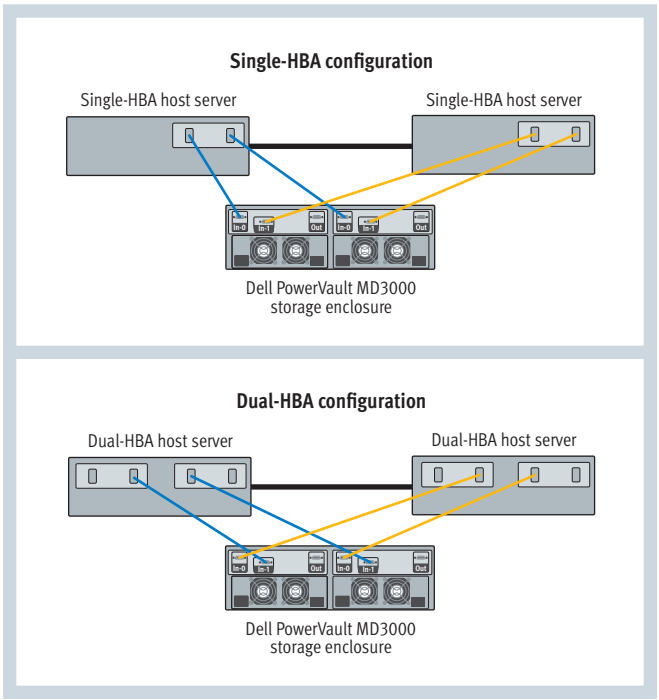

Figure 2. *Single- and dual-HBA cabling configurations for the Dell PowerVault MD3000 in a two-node Dell PowerEdge cluster*

Servers	Two Dell PowerEdge 2850 servers with Intel Xeon processors at 3.0 GHz
Memory	8 GB for each server
OS	Microsoft Windows Server 2003 R2 Standard x64 Edition
Software	Oracle Database 10*g* R2 (10.2.0.2) with Oracle RAC

Figure 3. *Cluster test configuration*

Each node runs a separate instance of the Oracle Database 10*g* Automatic Storage Management feature and Oracle database services, independently fulfilling client requests. Redundancy for the HBAs, network interface cards, and RAID controllers as well as software components help provide high availability. Multiple path management, load balancing on multiple paths, path failover software, and volume management are included on the Dell PowerVault MD3000 Resource CD.

Performance testing

In October 2006, Dell engineers set up a two-node Dell PowerEdge 2850 server cluster with PowerVault MD3000 external storage running Oracle Database 10*g* with Oracle RAC (see Figure 3). To evaluate cluster performance, they measured transactions per second under different loads by using Benchmark Factory from Quest, which provides an industry-standard TPC-C benchmark for testing online transaction processing (OLTP) databases. They then ran the same tests on this cluster configuration using 2 Gbps Fibre Channel–based storage and 4 Gbps Fibre Channel–based storage in place of the PowerVault MD3000.

Figure 4 shows the results, which have been normalized using the 2 Gbps Fibre Channel–based storage as a baseline for comparison. The results demonstrate that the PowerVault MD3000 performed similarly to Fibre Channel–based storage under low and moderate workloads, but exhibited better performance under heavy workloads, particularly when compared with the 2 Gbps Fibre Channel–based storage. Because the PowerVault MD3000 combines enterprise-class storage features with high performance and can typically cost less than comparable Fibre Channel implementations, it can offer better price/performance than Fibre Channel–based storage.

Deploying cost-effective database storage

The Dell PowerVault MD3000 offers enterprise-class performance, manageability, and scalability along with key features enabling data protection, storage partitioning, and virtual disk snapshots and copies. Deploying the PowerVault MD3000 as part of a Dell PowerEdge server cluster running Oracle Database 10*g* with Oracle RAC can help provide highly available storage clusters in a cost-effective way, particularly in small and medium-size enterprises.

Mahmoud Ahmadian *is an engineering consultant with the Database and Applications team of the Dell Product Group. He has an M.S. in Computer Science from the University of Houston, Clear Lake.*

Chethan Kumar *is a systems engineer and adviser in the Database and Applications team of the Dell Product Group. He has an M.S. in Computer Science and Engineering from the University of Texas at Arlington.*

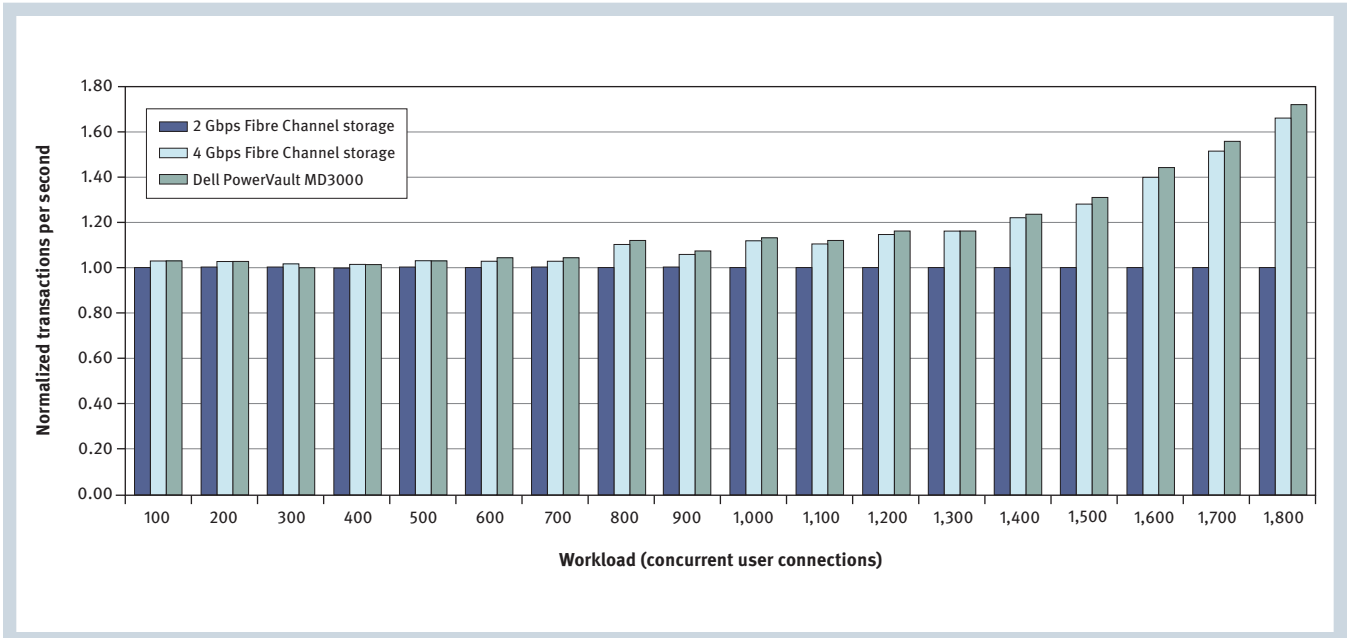

Figure 4. *Performance comparison of Dell PowerVault MD3000 storage against Fibre Channel–based storage*

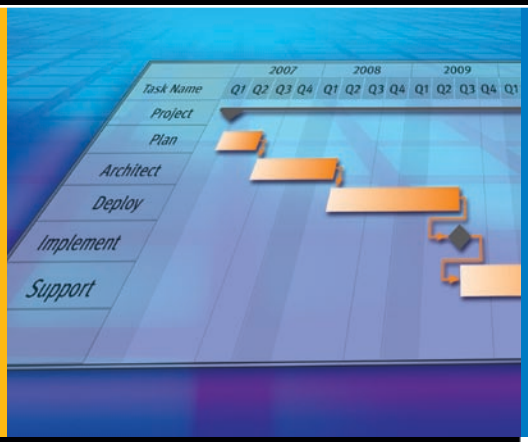

Using Oracle Streams Advanced Queuing in Application Messaging Infrastructures

BY RAMKUMAR RAJAGOPAL
MAHMOUD AHMADIAN

Oracle® Streams Advanced Queuing (SAQ) helps provide a flexible, scalable message queuing infrastructure through enqueue, propagation, and apply mechanisms. This article provides an overview of Oracle SAQ, describes an example messaging application created and tested by Dell engineers to demonstrate its functionality and performance, and discusses a number of available Oracle queue management tools.

Related Categories:

Application development

Application servers

Dell PowerEdge servers

Dell/EMC storage

Middleware

Oracle

Visit www.dell.com/powersolutions for the complete category index.

Scalability and flexibility are two important design considerations for building robust business integration architectures. Scalability is important for accommodating current applications, while flexibility is critical for seamless adaptation to future changes and requirements. Designing scalable and flexible business integration architectures requires strategic decision making.

Oracle Streams and Oracle Streams Advanced Queuing (SAQ) are designed to help provide such architectures. Oracle Streams can efficiently capture changes to database objects and replicate them to one or more databases; Oracle SAQ provides an application message queuing framework that integrates Oracle Advanced Queuing (OAQ) functionality with Oracle Streams through *enqueue, propagation,* and *apply* mechanisms.

Oracle SAQ provides a flexible, unified framework for building a messaging layer for the critical business integration software within an Oracle database. Using Oracle SAQ–based application programming interfaces (APIs), enterprises can efficiently build and deploy complex applications in an Oracle Database 10*g* environment.

Understanding Oracle message queuing

A *queue* in a database infrastructure is a holding place for messages, notifications, and events. Oracle Streams applications use the Oracle messaging infrastructure, comprising enqueue, propagation, and apply mechanisms. The enqueue mechanism puts messages in a queue, the propagation mechanism stages and forwards the messages to the remote destination, and the apply mechanism retrieves the messages from the queue and consumes them. The messages may be as small as a simple text string or as large as an entire table. The anydata type, which is an encapsulated presentation of all supported message types, allows a queue to handle multiple message types rather than just one type. The multi-consumer queue is often described as following the *publish/subscribe model,* the simplest form of which would have one publisher and one subscriber—in effect, a single-consumer queue.

Traditionally, messaging-based applications have used IBM® WebSphere MQ, Java Message Service (JMS), or OAQ. OAQ-based applications use a combination of stored procedures, database jobs, and triggers to run the application messaging process. The original OAQ APIs in Oracle8*i* and Oracle9*i* were built using DBMS_AQ packages; in Oracle SAQ, a set of APIs in DBMS_STREAMS packages either wrap these original APIs or extend their functionality. Strongly typed single-consumer queues are only supported with DBMS_AQ APIs. To take advantage of the ability of DBMS_STREAMS APIs to allow a queue to handle many types

of messages rather than just one, the queues must be anydata type and multi-consumer queues.

Publish/subscribe model

Oracle Streams allows applications to send messages of different types, propagate the messages to subscribing queues, and notify user applications that messages are ready for consumption. Oracle Streams interoperates with OAQ, which provides several standard message-queuing system features such as multi-consumer queues, publishing and subscribing, content-based routing, Internet propagation, transformations, and gateways to other messaging subsystems. A key advantage of the publish/subscribe model and the interoperability of Oracle Streams and OAQ is that Oracle SAQ processes messages automatically instead of requiring external APIs to process messages manually.

Message propagation

In Oracle SAQ, message routing or propagation is carried out by propagation jobs, which are created to dequeue messages from a local queue and enqueue them to a remote queue using database links and job queue processes. The default propagation latency is three seconds, meaning that messages not acknowledged within three seconds of submission will be resubmitted; this parameter is configurable for different application requirements. Propagation rules determine which remote queues received a given message. Each source queue can have one or more propagations set to different destination queues.

Enqueuing messages to the local queue, propagating them to multiple remote queues, and dequeuing them involve the following steps:

1. A queue is created on the local database, which stores messages in a predefined order—usually first in, first out (FIFO).
2. Queues that will receive and store the propagated messages are created on remote databases.
3. Database links are created to allow the propagation jobs to enqueue messages into the remote queues.
4. Propagation jobs are created.
5. Rules are created during the creation of the propagation jobs to determine which messages should be propagated to which remote queue. Rules are usually based on data contained within the message—for example, there may be a field named "destination" that identifies the remote queue.
6. User processes, most likely applications, create and enqueue messages into the local queue.
7. Propagation jobs dequeue the messages from the local queue and enqueue them into the remote queues depending on whether the rule for that propagation is satisfied.

If the queues and propagations are implemented using DBMS_STREAMS APIs, then subscribers can be user processes that manually dequeue from this queue using any of the supported APIs—including JMS, DBMS_AQ and DBMS_STREAMS, or apply processes that automate the dequeuing and processing.

Buffered queues

Oracle SAQ in Oracle Database 10*g* also introduces buffering for user enqueues, which includes many of the advantages of both persistent (disk-stored) and nonpersistent (memory-stored) queues. Buffered queues store messages in system memory as long as it is available; when this memory is unavailable, it stores them on disk. This type of queue combines the persistence advantages of disk storage with the speed advantages of memory storage.

Automated apply process

The Oracle Streams apply process is the automatic dequeuing and processing of events from an Oracle Streams queue. It can also apply business rules, handle conflicts in messages, and log the errors into queue tables. This apply process is carried out by one or more Oracle background processes. The apply engine consists of three components: a reader server, which dequeues messages and assembles them into transactions; a coordinator process, which distributes messages to apply servers in dependency order; and a number of parallel apply servers, which apply messages to the database or pass them to the appropriate handlers.

These processes can directly modify database objects or pass messages as parameters to an administrator-defined procedure called an apply handler. Apply handlers can take the form of message handlers, Data Manipulation Language (DML) handlers, Data Definition Language (DDL) handlers, pre-commit handlers, or error handlers. Like other database objects, apply processes can be created, altered, started, stopped, and dropped.

Errors or conflicts encountered during the apply process are dealt with by error handlers or conflict handlers. If an error arises during an apply process, the error handler for the apply process is invoked; if the error cannot be resolved by the apply server, it rolls back the transaction and places the transaction and all its messages in an error queue. The apply engine helps simplify the coding of message processing and automatically dequeues the message. Oracle SAQ can dramatically reduce the time required to build messaging applications, because most of the time and development resources can be spent on designing and developing the business functionalities, logic, or process and fine-tuning them through message handlers, rather than on writing code to poll, start, or send an alert to another application to process the messages.

Scalability

Oracle SAQ provides a framework for quickly and simply scaling out a messaging application. Oracle Real Application Clusters (RAC) 10*g* allows for load balancing and helps distribute messages across multiple nodes. It directs messages to nodes where processes are executing well and

Oracle software	Oracle Database 10*g*, patch set 10.1.0.5
OS	Red Hat® Enterprise Linux® AS 2.1 with the hugemem kernel
Servers	Dell PowerEdge 2850 servers with two Intel Xeon processors at 3.2 GHz with Intel Hyper-Threading Technology enabled
Memory	8 GB
Storage	Dell/EMC CX700 storage array
Network	Intel Gigabit Ethernet controller

Figure 1. *Software and hardware configuration for test environment*

resources are available, and automatically stops sending messages to slow, hung, or failed nodes. Adding adequate physical memory to an Oracle RAC 10*g* cluster can improve performance and scalability to meet application requirements. The publish/subscribe architecture enables adding new subscribers to expand capacity and future growth without causing downtime for critical applications.

Developing an Oracle Streams Advanced Queuing test application

To demonstrate Oracle SAQ functionality and performance, Dell engineers developed a test messaging application using the built-in Oracle SAQ packages. The test environment comprised five components: Web-based applications, a two-node clustered central system, a two-node Regional System 1 (Region1), a two-node Regional System 2 (Region2), and a single-node reporting system. The application used Oracle Database 10*g* on Dell™ PowerEdge™ 2850 servers with two Intel® Xeon® processors and a Dell/EMC CX700 Fibre Channel storage array. Figure 1 describes the software and hardware configuration for the test environment.

Propagation process stages

As shown in Figure 2, the following message propagations were configured:

- **Propagation1:** One-to-many propagation from the central system to the Region1 and Region2 systems
- **Propagation2:** Many-to-one propagation from the Region1 and Region2 systems to the central system
- **Propagation3:** One-to-one propagation from the central system to the reporting system

Apply process stages

As shown in Figure 3, the test application apply process included five stages:

- **Apply1:** The apply process dequeues messages from the central system, and the message handler propagates them to the remote queues.
- **Apply2:** Region1 dequeues the messages and calls the message handler to acknowledge the messages to the central system.
- **Apply3:** Region2 dequeues the messages and calls the message handler to acknowledge the messages to the central system.
- **Apply4:** The central system dequeues the messages, applies business rules using a message handler, and sends the messages to the reporting system.
- **Apply5:** The reporting system dequeues the messages and updates the reporting tables.

The complete message queuing process for the test application involved the following propagation and apply steps:

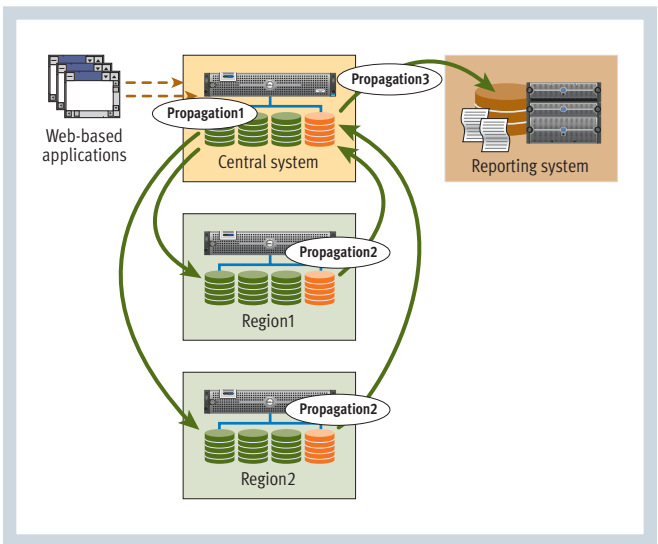

Figure 2. *Propagation process in the Oracle Streams Advanced Queuing test application*

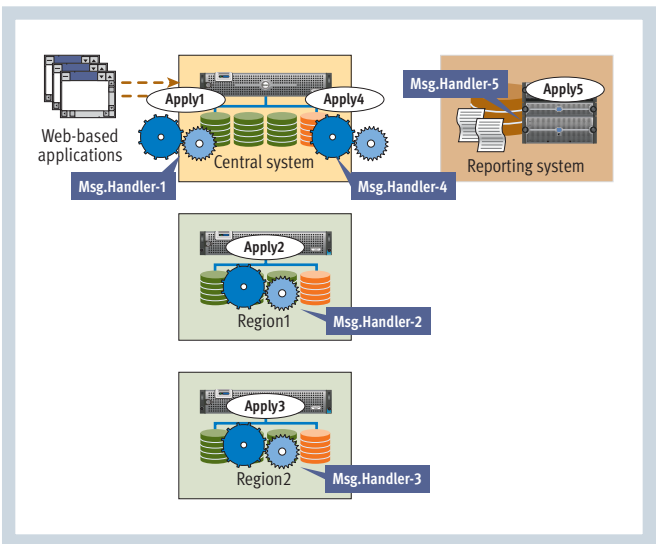

Figure 3. *Apply process in the Oracle Streams Advanced Queuing test application*

1. Web applications enqueue messages into the central system's local queue.
2. The central system converts the objects to the anydata type.
3. In the Apply1 stage, the apply process uses the Oracle Streams API to dequeue messages and set regional system location identification using the propagation rule.
4. The apply process message handlers call a stored procedure that enqueues the messages into the remote destination queues for the Region1 and Region2 systems.
5. The message propagation processes validate and send the messages to the particular regional system based on the location identification set in the apply process.
6. In the Apply2 and Apply3 stages at the regional destination local queues, the apply process dequeues the messages using its message rule and calls the message handler (a stored procedure), which builds an acknowledgment message and sends it back to the central system.
7. In the Apply4 stage at the central system, the apply process dequeues the messages, applies business rules using the message handler, and sends the messages to the reporting system.
8. In the Apply5 stage at the reporting system, the apply process calls the message handler to process the messages based on the apply rules, dequeues the messages, and updates the reporting tables.

Testing Oracle Streams Advanced Queuing performance

System architectures for critical applications are often not able to provide the fast, flexible workload balancing necessary to help maximize available processing and memory resources during system spikes. Enterprises can use Oracle SAQ to simplify messaging and database operations to help enhance resource utilization and scale out their systems in pragmatic ways. To demonstrate the potential benefits of Oracle SAQ, in July 2006 Dell engineers evaluated the performance of the test application described in the preceding section. The Dell team conducted the tests under the following conditions:

- All tests were repeated numerous times.
- The test results did not include network latencies for client connections.
- None of the systems exceeded 75 percent processor utilization or 75 percent memory utilization.

The dequeue rate and throughput of messages per second were measured using a customized Procedural Language/Structural Query Language (PL/SQL) procedure to log the start time and end time of each test in an Oracle table. This rate depended on the application logic and business rules specific to the test application. The tests measured the end-to-end processing of messages using Oracle SAQ–based APIs.

The test team ran a workload that generated messages ranging from 10 KB to 100 KB and measured the average dequeue rate for each message size. As shown in Figure 4, the average dequeue rates for the Apply1, Apply2, Apply3, and Apply5 process stages were nearly identical. In the Apply4 stage, however, the average dequeue rates were almost twice that of the other stages. This result was caused by the two-to-one subscription at the central system—both the Region1 and Region2 systems were publishing messages to it. Based on the workload, the apply process spawned multiple parallel slaves, using the available processing and memory resources to help automatically maximize performance.

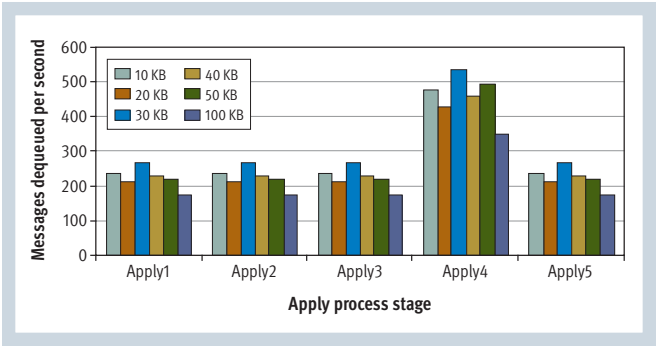

Figure 4. *Average apply-process dequeue rates for different-size messages in the Oracle Streams Advanced Queuing test application*

As shown in Figure 5, the system processed 100,000 messages from start to finish—a total of 500,000 dequeues for all five stages of the apply process—in approximately 47 minutes, resulting in an average end-to-end dequeue rate of about 177 messages per second. The aggregate processor utilization of the two central system nodes during this period was approximately 35 percent, as shown in Figure 6.

This test application scales well because many of the messaging functions traditionally performed by applications were abstracted to the database layer. As shown in Figure 6, the cluster had significant processor resources available before it would become a bottleneck. By using the load-balancing features of Oracle RAC, the test team was able to evenly distribute the workload between the two central system nodes to help ensure neither node inhibited performance. If aggregate processor utilization for the central system had become too high to run the application

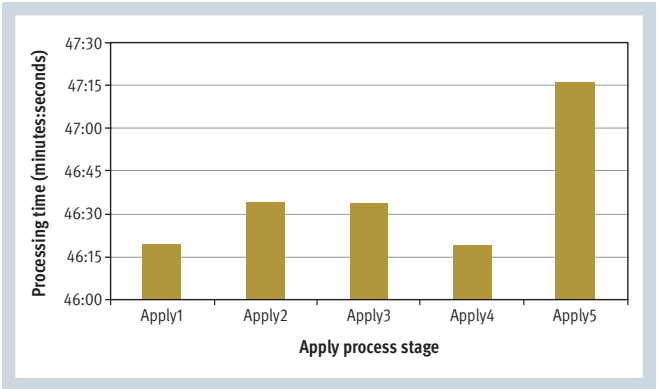

Figure 5. *Time to process 100,000 messages at each apply-process stage in the Oracle Streams Advanced Queuing test application*

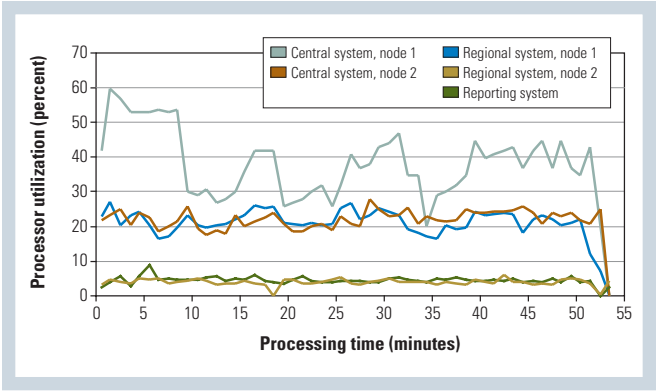

Figure 6. *Central system processor utilization during testing*

and performance degraded, nodes could have been added to the cluster to scale out and enable fast, flexible workload balancing to help maximize processing and memory resources.

Managing Oracle Streams Advanced Queuing

Several Oracle products and features help administrators manage Oracle SAQ, including Oracle Enterprise Manager 10*g* Release 2 (R2) Grid Control, Oracle Database 10*g* Automatic Optimizer Statistics Collection, the STRMMON Oracle Streams monitoring tool, and a health-check script.

Oracle Enterprise Manager 10*g* R2 Grid Control helps configure, administer, monitor, troubleshoot, and resolve problems quickly. The Oracle Database 10*g* Automatic Optimizer Statistics Collection feature can run regularly to gather updated table optimizer statistics. When using this feature, administrators should keep in mind that statistics collection for volatile tables (such as Oracle Streams queue tables) may occur when the tables do not include data representative of their full load period. To help avoid this problem, administrators can run the statistics-gathering job on these tables manually during an appropriate load period, then lock the statistics.

The STRMMON Oracle Streams monitoring tool, available in the rdbms/demo directory in $ORACLE_HOME, provides administrators with an overview of the Oracle Streams activity occurring within a database. Administrators can configure the reporting interval and the number of iterations to display.

The Oracle Streams health-check script is a collection of queries to determine the Oracle Streams environment configurations; administrators should run this script on each participating Oracle Streams database. The script also includes analysis of Oracle Streams rules and a guide to interpreting the output to help simplify troubleshooting.

In addition to using these tools, administrators can modify database instance initialization parameters to enhance performance. Figure 7 provides recommended values for these parameters.

Implementing a scalable application messaging infrastructure

Oracle Streams provides a flexible design framework for messaging applications by offloading or abstracting many functions traditionally implemented in the application layer into the database layer. Each application does not need to implement common messaging functionality; instead, these applications can simply publish or subscribe to Oracle Streams objects, and administrators can focus on improving performance and business logic. Oracle Streams Advanced Queuing also includes advanced tools for troubleshooting and manageability. By offering flexibility, transactional recoverability, high availability, and load balancing, the Oracle SAQ messaging database architecture can significantly improve performance through the use of buffered in-memory queues and parallelization of apply process. Oracle SAQ thus provides the capacity to scale out to meet future business needs within existing systems.

Ramkumar Rajagopal *is a senior database administrator for the Dell IT group.*

Mahmoud Ahmadian *is a senior development consultant with the Database and Applications team of the Dell Product Group.*

Parameter	Description	Default value	Recommended value
job_queue_processes	Specifies the maximum number of processes that can be created for the execution of jobs	0	Greater than 4
sga_max_size	Sets the initial size of the System Global Area (SGA) at startup, dependent on the sizes of different pools in the SGA, such as buffer cache, shared pool, large pool, and streams_pool_size	40% of physical memory	Large enough to support the streams_pool_size and shared_pool_size parameters
streams_pool_size	Specifies the size of the Oracle Streams pool; the Oracle Automatic Shared Memory Management feature manages this parameter when the SGA_TARGET initialization parameter is set to a nonzero value	0	Greater than 200 MB*
parallel_max_servers, processes, sessions, and open_links	Specifies the maximum number of parallel execution processes and parallel recovery processes for an instance	Derived from the values of CPU_COUNT, PARALLEL_THREADS_PER_CPU, and PGA_AGGREGATE_TARGET	Large enough to support the configuration

*The unit of measurement for this parameter is bytes; to specify kilobytes or megabytes, "K" or "M" should be added to the number (for example, 200 MB should be entered as "200M").

Figure 7. *Recommended values for Oracle database instance initialization parameters*

Intel I/O Acceleration Technology

Helps Eliminate Network Bottlenecks

Designed to move network data at optimal efficiency, Intel® I/O Acceleration Technology can help improve the performance, scalability, and reliability of Quad-Core and Dual-Core Intel Xeon® processor–based servers.

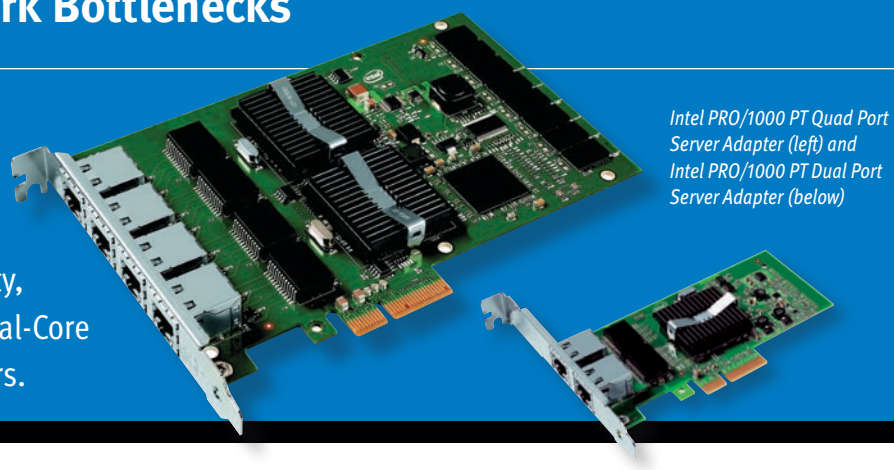

Intel PRO/1000 PT Quad Port Server Adapter (left) and Intel PRO/1000 PT Dual Port Server Adapter (below)

Today's servers are pushed to the limits by business-critical network applications. To help meet growing performance demands, IT managers typically increase server capacity and/or network bandwidth, only to discover that server performance does not increase as expected. The problem can often be traced back to three network I/O bottlenecks that can keep a server from achieving its full performance potential: system overhead, TCP/IP processing, and data copies within system memory.

Based on Intel QuickData Technology, Intel I/O Acceleration Technology (Intel I/OAT)—available in Quad-Core and Dual-Core Intel Xeon processor–based Dell™ PowerEdge™ servers equipped with Intel PRO/1000 Ethernet adapters for PCI Express—is designed to enhance performance by moving network data more efficiently through the server than was possible in previous-generation Intel Xeon processor–based servers. At the same time, Intel I/OAT scales seamlessly across up to six Gigabit[1] Ethernet ports—helping reduce total cost of ownership and boost server throughput.

Breaking through I/O bottlenecks

By taking advantage of architectural enhancements within the processor, chipset, network controller, and firmware of Quad-Core and Dual-Core Intel Xeon processor–based servers, Intel I/OAT implements a platform-oriented approach designed to speed up interactions between server applications and the network. This approach enables Intel I/OAT to avoid performance-limiting bottlenecks by using key server resources more efficiently than previous-generation systems—helping minimize system overhead, accelerate TCP/IP processing, and efficiently manage data copies within system memory.

Intel I/OAT enables a significant increase in server performance, as shown in recent system tests conducted by Intel Labs on both Red Hat® Enterprise Linux® 4 and Microsoft® Windows Server® 2003 with Service Pack 1 (SP1) and the Microsoft Scalable Networking Pack (SNP). In tests running on these operating systems, Intel Xeon processor–based servers with Intel I/OAT demonstrated up to 50 percent less processor load and up to 2.5 times higher throughput compared with a previous-generation Intel Xeon processor–based server.[2]

Enhancing scalability and reliability

Intel I/OAT helps improve network I/O performance by scaling seamlessly across up to six Gigabit Ethernet ports. In this way, Intel I/OAT—together with Intel PRO multi-port Gigabit Ethernet adapters, which are designed to leverage the full performance benefits of the PCI Express I/O standard and to provide high port density for slot-constrained servers—helps reduce total cost of ownership by enabling server consolidation.

Although the primary performance benefit of Intel I/OAT is its ability to move network data at optimal efficiency through Quad-Core and Dual-Core Intel Xeon processor–based servers, maintaining the integrity of the OS network stack is critical. Native support for popular server operating systems such as Red Hat Enterprise Linux 4 and Windows Server 2003 with the SNP enables Intel I/OAT to preserve critical network configurations such as teaming and failover. Intel I/OAT is designed to accomplish this by maintaining control of the network stack execution within the processor. Because it allows IT departments to rely on OS updates instead of third-party software to provide this functionality, Intel I/OAT helps reduce support risks—another factor that contributes to low total cost of ownership.

more Online
www.dell.com/**powersolutions**

QUICK LINK

Intel I/O Acceleration Technology:
www.intel.com/go/ioat

[1] This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

[2] Performance results for both processor load and throughput are based on an Ixia IxChariot 6.0 benchmark test performed by Intel Labs in January 2007 on a Dell PowerEdge 2900 server featuring Intel I/OAT and configured with two Dual-Core Intel Xeon processors (Woodcrest) at 3.0 GHz, 4 GB of 667 MHz RAM, three Intel PRO/1000 PT Dual Port Server Adapters, either Red Hat Enterprise Linux 4 (kernel 2.6.18) or Windows Server 2003 with SP1 and the SNP, and six Dell PowerEdge 750 servers per port under test as clients, each configured with an Intel Pentium® 4 processor at 3.4 GHz, Windows Server 2003 with SP1, and an Intel PRO/1000 Gigabit Ethernet server adapter. The performance baseline was established in February 2006 using a previous-generation Dell PowerEdge 2800 server configured with two Intel Xeon processors at 3.6 GHz, 4 GB of RAM, three Intel PRO/1000 PT Dual Port Server Adapters, and either Red Hat Enterprise Linux 4 Update 3 or Windows Server 2003 with SP1. Actual performance will vary based on configuration, usage, and manufacturing variability.

Proactive Maintenance and Power Management with Dell OpenManage and VMware Virtualization

BY BALASUBRAMANIAN CHANDRASEKARAN

PUNEET DHAWAN

Administrators can combine the Dell OpenManage™ systems management suite with the cluster resources of VMware® Infrastructure 3 to achieve proactive maintenance that enhances service continuity along with adaptive power utilization that helps reduce data center power and cooling costs.

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Power and cooling

Virtualization

VMware

Visit www.dell.com/powersolutions for the complete category index.

Cluster virtualization enables enterprises to consolidate existing workloads, reduce data center power and cooling costs, and respond to changing business needs. Administrators can take advantage of the Dell OpenManage suite and VMware Infrastructure 3 to manage virtualized cluster environments on Dell™ PowerEdge™ servers. Combining the tools available with Dell OpenManage and VMware Infrastructure 3 can provide proactive maintenance that enhances service continuity along with adaptive power utilization that helps reduce data center power and cooling costs.

Dell OpenManage and VMware Infrastructure 3 tools

Dell OpenManage and VMware Infrastructure 3 software are designed to simplify the management of virtualized environments running on Dell PowerEdge servers. Key tools for managing virtualized clusters include Dell OpenManage Server Administrator (OMSA), Dell OpenManage IT Assistant, VMware Distributed Resource Scheduler (DRS), and the VMware Infrastructure Software Development Kit (SDK).

Dell OpenManage

The Dell OpenManage suite consists of systems management applications for managing Dell PowerEdge servers, offering a comprehensive set of standards-based and interoperable tools for server deployment, monitoring, and change management. The two Dell OpenManage applications that are most relevant for the virtualization management described in this article are OMSA and IT Assistant.

Dell OpenManage Server Administrator. OMSA enables administrators to easily manage individual servers and internal storage arrays by performing tasks such as reviewing server status and inventory, configuring BIOS and RAID, setting actions based on events, and powering servers up and down. OMSA is fully qualified to run within a VMware Infrastructure 3 environment.

Dell OpenManage IT Assistant. IT Assistant provides a comprehensive, standards-based, one-view console for managing Dell servers, storage, tape libraries, network switches, printers, and clients systems. Among other features, IT Assistant enables administrators to capture events and alerts generated by Dell servers running OMSA, configure actions based on these events and alerts, and monitor server performance statistics.

VMware Infrastructure 3

The VMware Infrastructure 3 suite includes enterprise-class virtualization software to enable consolidation, management, resource optimization, and high availability for IT data centers. The software supports Internet SCSI (iSCSI), network attached storage, 64-bit virtual machines (VMs), four-way symmetric multiprocessing for VMs, and up to 16 GB of VM memory, along with features like VMware High Availability (VMware HA) and DRS. VMware Infrastructure 3 is fully qualified to run on Dell PowerEdge servers, including those that offer enhanced processor technologies such as quad-core processors, Intel® Virtualization Technology, and AMD™ Virtualization.

"Administrators can integrate Dell OpenManage systems monitoring with VMware VirtualCenter to help improve cluster fault tolerance and enable proactive maintenance."

Two key components of VMware Infrastructure 3 relevant to enabling proactive maintenance and adaptive power management are DRS and the VMware Infrastructure SDK. DRS provides dynamic resource scheduling and physical resource optimization; the SDK enables third-party applications to manage and control ESX Server hosts and VMs.

VMware Distributed Resource Scheduler. VMware Infrastructure 3 introduced the concept of an ESX Server cluster, a group of loosely tied ESX Server hosts that can be managed as a single entity. As the name suggests, DRS groups distributed computing resources on physical servers into a single pool, and schedules VMs on servers that can best serve the resource requirements of these VMs. It is built on VMware VMotion™ technology. DRS provides the following major features:

- Automatic initial placement of VMs on a "best-fit" cluster host
- Automatic resource optimization and relocation based on changes in a cluster's computing resources, such as the addition or removal of a host
- Automatic relocation of VMs based on resource requirements

VMware Infrastructure 3 also introduced the maintenance-mode host status, which migrates all VMs from a particular ESX Server host to other hosts. DRS automatically checks that the VMs are relocated among the remaining hosts in a way that helps maximize resource optimization.

VMware Infrastructure Software Development Kit. The VMware Infrastructure SDK allows developers to build custom Simple Object Access Protocol (SOAP)–based applications to manage ESX Server hosts and VMs. It also allows administrators to integrate existing management applications with VMware Infrastructure 3 and automate cloning and configuration of VMs, performance reporting, and other tasks.

Dell OpenManage and VMware Infrastructure 3 integration

Administrators can integrate different components of the Dell OpenManage suite and VMware Infrastructure 3 to enable comprehensive physical and virtual infrastructure management and task automation. Two examples can illustrate this process: creating a proactive response to hardware faults and configuring adaptive power management.

The scripts and program files to implement these examples are available at www.dell.com/downloads/global/solutions/prctv.zip.[1] These scripts and programs provide a framework to integrate systems management with VMware VirtualCenter by using the VMware Infrastructure SDK. Administrators can modify the code to fit their environment.

Creating a proactive response to hardware faults

Administrators can integrate Dell OpenManage systems monitoring with VMware VirtualCenter to help improve cluster fault tolerance and enable proactive maintenance. During an event such as a server hardware fault, the VMs from the faulty server can be proactively migrated to other healthy servers in the DRS cluster. This response helps avoid VM downtime from any additional hardware faults.

The steps described in this section are based on a Dell white paper[2] about using Dell OpenManage with ESX Server 2 and VirtualCenter 1; that paper discusses algorithms to choose target servers for migration based on the processor load on each of the candidate servers in the migration pool. However, such an approach is complex, requiring developers to build an optimal algorithm to choose target servers and make decisions for each migrating VM.

This article extends the same concepts to DRS clusters and takes advantage of the fully automated VMotion capabilities for load balancing and managing VM resource requirements. An action at the IT Assistant server layer based on a hardware fault, such as a loss of power redundancy reported by an OMSA agent, can automatically put ESX Server hosts into maintenance mode and migrate VMs to another healthy host.

"The VMware Infrastructure 3 suite includes enterprise-class virtualization software to enable consolidation, management, resource optimization, and high availability for IT data centers."

[1] These scripts and programs are provided as is, without any implied support or warranty.

[2] "Implementing Fault-Tolerance Through Dell OpenManage and the VMware Software Development Kit," by Dave Jaffe and Todd Muirhead, Dell Enterprise Product Group, September 2005, www.dell.com/downloads/global/solutions/OM-VMware-Integration.pdf.

Figure 1 illustrates the following sequence of actions that take place when a hardware fault occurs:

1. OMSA sends a Simple Network Management Protocol (SNMP) trap to the IT Assistant server that contains information about the server and the alert.
2. The IT Assistant server filters the traps as configured by administrators, and invokes a Java program by passing the server name and severity as arguments.
3. For any alerts selected by administrators, the Java program connects to the VirtualCenter server using the VMware Infrastructure SDK and issues a command to put the faulty ESX Server host in maintenance mode.
4. All VMs from the faulty host are migrated to other hosts in the cluster, with placement decided by the DRS algorithm.

The administrators can now look at the faulty server and perform maintenance actions, and the running VMs have avoided downtime. Once server health is restored, the following sequence of actions occurs:

1. OMSA sends an SNMP normal alert (an SNMP trap with the severity level set to Normal) to the IT Assistant server.
2. The IT Assistant server filters the traps as configured by administrators, and invokes the Java program by passing the server name and severity as arguments.
3. For normal alerts, the Java program sends an SNMP query to the server for global health information, to help ensure that other server subsystems are also healthy. If the global status of the server is healthy, the Java program connects to the VirtualCenter server using the VMware Infrastructure SDK and issues a command to remove the server from maintenance mode.
4. The DRS service discovers the addition of the new server into the cluster and redistributes the VMs to balance the cluster load.

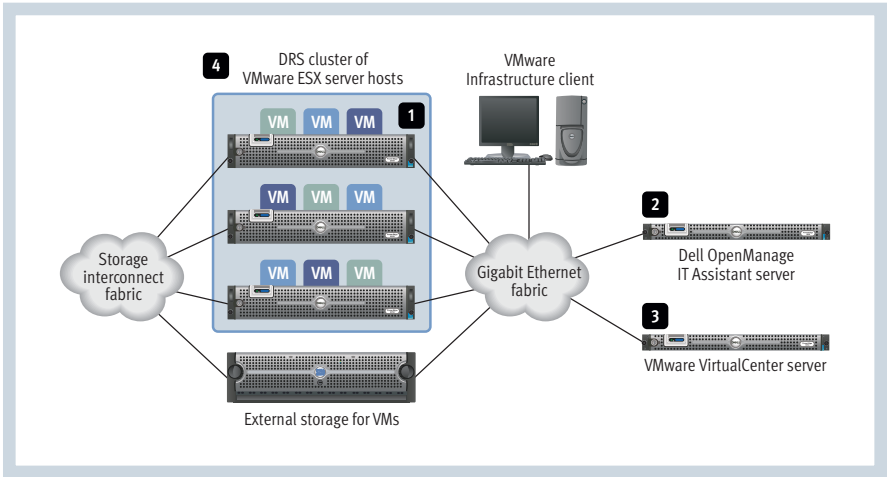

Figure 1. *Steps for a proactive response to hardware faults*

> "Using power efficiently and containing infrastructure costs are important elements of effective data center management."

Configuring adaptive power management

Using power efficiently and containing infrastructure costs are important elements of effective data center management. According to a Gartner report, most large enterprise IT organizations spend approximately 5 percent of their total IT budgets on energy, and this could double or triple within five years. The report also estimates that most enterprise data centers waste more than 60 percent of the energy used to cool their equipment.[3] These expenditures put large data centers under constant pressure to tackle increasing power and cooling demands. Server virtualization enables enterprises to consolidate a large number of underutilized servers, helping mitigate these costs.

Because typical data center workloads have characteristic patterns of utilization peaks and troughs, keeping all cluster servers powered on all the time is rarely optimal; automatically consolidating workloads on fewer servers during off-peak hours helps avoid running servers unnecessarily and reduce associated costs. For example, resource utilization for workloads like internal data shares and e-mail, printing, and Web servers may peak during office hours but decrease during nights and weekends.

Administrators can combine the Dell OpenManage systems management suite with Virtual Infrastructure 3 DRS clusters to achieve adaptive power management. The resource utilization of the DRS cluster is constantly monitored for resource utilization using the VMware Infrastructure SDK; when the average cluster utilization falls below a set threshold, VMs are automatically consolidated to fewer servers, and the unneeded servers are automatically powered down using Dell Remote Access Controllers (DRACs) to save power. When the resource utilization of the VMs increases, servers are automatically powered on to meet the increase in demand (see Figure 2).

[3] "Why 'Going Green' Will Become Essential for Data Centers," by Rakesh Kumar, Gartner, Inc., October 10, 2006.

Adaptive power management algorithm. Adaptive power management aims to minimize the number of powered-up physical servers supporting the virtual infrastructure workload so that average cluster utilization remains within specified limits. The steps carried out by the adaptive power management algorithm are as follows:

1. The Java program polls the VirtualCenter server and measures cluster-level processor and memory utilization.
2. If the average processor or memory utilization falls below an administrator-configured minimum threshold, one server at a time is put into maintenance mode, and VMs are automatically migrated to other servers by the DRS service. Once the migration is complete, the server is powered down. This process is repeated until the cluster utilization rises above the minimum threshold.
3. If the average processor or memory utilization increases above an administrator-configured maximum threshold, one server at a time is powered on. The DRS service discovers the addition of the new server into the cluster and redistributes the VMs to balance the cluster load. This process is repeated until the cluster utilization falls below the maximum threshold.

Configurable parameters. Administrators can define a set of configurable parameters in the PowerSave.xml configuration file available at www.dell.com/downloads/global/solutions/prctv.zip:

- **Minimum utilization:** The cpuMin and memMin parameters define minimum limits on cluster processor and memory utilization as a percentage of total available resources. Servers can be automatically powered down when average cluster utilization falls below both of these values.
- **Maximum utilization:** The cpuMax and memMax parameters define maximum limits on cluster processor and memory utilization as a percentage of total available resources. Servers can be automatically powered up when average cluster utilization rises above either of these values.
- **Number of active VMs per server:** The VMMin and VMMax parameters define the limits on the number of VMs that can run on each physical server at one time. Administrators can use these parameters, for example, to prevent additional servers from being powered up when the average number of active VMs per server is below the minimum threshold, or to prevent servers from being powered down when the average number of VMs per server is above the maximum threshold. They can also avoid these limits by setting VMMin to 1 and VMMax to an extremely high value.
- **Timing:** The opTimeout parameter defines how often to poll for utilization information; the steadyState parameter defines the time between when a parameter threshold is passed and when the corresponding action begins, which helps ensure that random workload spikes and troughs do not cause servers to be powered up or down unnecessarily.

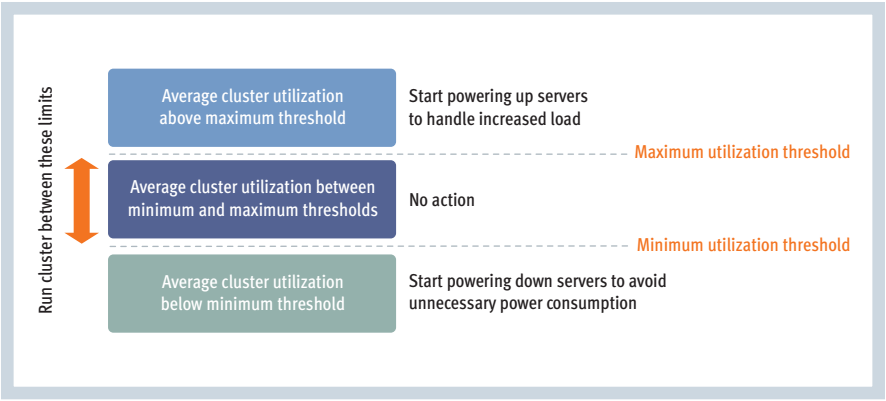

Figure 2. *Adaptive power management model*

Integrated management of virtualized clusters

Integrating the powerful systems management capabilities of the Dell OpenManage and VMware Infrastructure 3 suites can provide enhanced control over important cluster functionality such as proactive responses to hardware faults and adaptive power management. The processes described in this article are intended to help simplify management and optimize physical resource utilization, which can help reduce total data center costs.

Balasubramanian Chandrasekaran *is a systems engineer in the Dell Virtualization Solutions Engineering group. His research interests include data center virtualization, high-speed interconnects, and high-performance computing. Balasubramanian has an M.S. in Computer Science from the Ohio State University.*

Puneet Dhawan *is a systems engineer in the Dell Virtualization Solutions Engineering group. Puneet has a bachelor's degree in Electrical Engineering from Punjab Engineering College (PEC) and a master's degree in Computer Engineering from Texas A&M University.*

more Online
www.dell.com/**powersolutions**

QUICK LINKS

Dell and VMware:
www.dell.com/vmware

Supporting scripts and programs:
www.dell.com/downloads/global/solutions/prctv.zip

How to Jump-Start the Migration to Microsoft Windows Vista

BY CHRIS HARGET

IT groups can accelerate migration to the Microsoft® Windows Vista™ OS and 2007 Microsoft Office by using Citrix Presentation Server™ 4.5 software. Virtualizing applications allows delivery of legacy and new applications to existing clients and those running Windows Vista, while streaming applications enables local and offline operation to help eliminate compatibility issues.

Related Categories:

Citrix Systems

Microsoft Windows Vista

Virtualization

Visit www.dell.com/powersolutions for the complete category index.

For many organizations, the question is not *if* they will migrate to Microsoft Windows Vista and the 2007 Microsoft Office release, but *when*. When will all their mission-critical applications support Windows Vista? When will their clients be upgraded to meet system requirements?

What if the same applications in use today could be delivered to Windows Vista users—without requiring native Windows Vista support? What if clients that did not meet the system requirements for new applications such as 2007 Microsoft Office could still access those applications?

Application virtualization and application streaming, used together, can remove many obstacles to early adoption of, and migration to, Windows Vista and 2007 Microsoft Office. Individually, each technique has strengths and limitations. Combined, they provide well-balanced options to accelerate migration and help lower total cost of ownership.

Windows Vista includes significant changes that can affect application compatibility. This 64-bit OS does not support 16- and 32-bit drivers. New system application programming interfaces (APIs), User Account Control, Windows Resource Protection, and reduced rights for the Microsoft Internet Explorer® Web browser can cause compatibility issues. The 2007 Microsoft Office release, meanwhile, introduces new file formats, making it highly desirable for all users to run the same Microsoft Office version so that they can collaborate and share information effectively.

Separating applications from operating systems

In virtualized environments, desktop applications run on a server, and the user interface is virtually delivered to the end user over the network. These applications are not dependent on the local OS or device drivers, have minimal hardware requirements, and help reduce maintenance.

Although solutions such as Citrix Presentation Server 4.5 software offer advantages in control, security, and cost-effectiveness (see the sidebar on the next page), some trade-offs exist with application virtualization. Clients must be logged on to the network, limiting notebook users. Very long-distance connections may exhibit lag time. In the past, graphics-intensive 3D applications were challenging to deliver over networks.

Application streaming can deliver applications to clients for use anytime, even when they are not connected to a network. Applications are cached locally in an *isolation environment* rather than installed on the device. This approach helps eliminate application conflicts and reduce the need for extensive regression testing to certify applications. As a result, applications can be offered as an on-demand service, updated quickly, or de-provisioned just as quickly.

Application streaming helps make users independent with many of the manageability and cost-efficiency benefits of centralized applications, but there are some trade-offs. Potentially heavy bandwidth usage can occur when large applications are first streamed. Applications that require a real-time connection to a back-end database or are tightly coupled to the OS are not

HOW CITRIX PRESENTATION SERVER ACCELERATES THE MIGRATION PROCESS

IT planners exploring how to migrate to new applications such as 2007 Microsoft Office and operating systems such as Microsoft Windows Vista are looking at application virtualization and application streaming as strategies to help simplify and accelerate this migration, reduce administrative burden, and lower costs. Citrix Presentation Server 4.5 combines application virtualization and application streaming into a single, end-to-end application delivery solution. It provides flexible policies designed to dynamically choose the best delivery method based on the user, device, application, and network.

Virtualized applications run in the data center, helping reduce deployment time and costs. This approach allows legacy clients to access the latest applications, and Windows Vista clients to access legacy applications. When users are disconnected from the network and need a local version of an application, streaming allows applications to execute locally in a desktop isolation environment without installation. This approach helps eliminate compatibility issues, helps reduce regression testing time and cost, and allows centrally managed applications to run offline. Citrix Password Manager™ software provides enterprise single sign-on access to help streamline secure migration to new application versions.

Figure A shows how Presentation Server 4.5 helps ensure business continuity during Windows Vista migration by virtualizing and streaming Windows Vista–incompatible applications to Windows Vista clients, and streaming Windows Vista–compatible applications such as 2007 Microsoft Office to clients running the Microsoft Windows® XP or Windows 2000 operating systems. For more information, visit www.citrix.com/presentationserver.

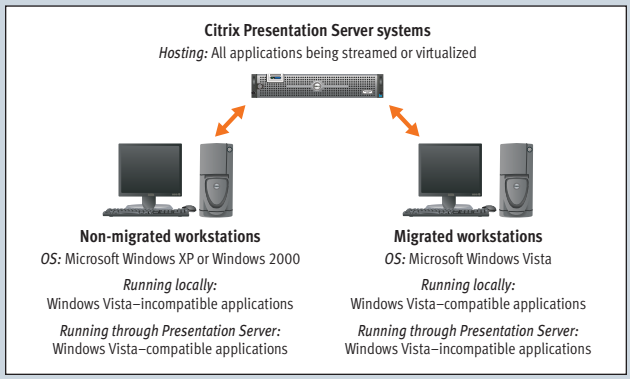

Figure A. *Microsoft Windows Vista migration using Citrix Presentation Server 4.5*

good candidates for streaming. Combining application virtualization and application streaming can give IT groups the flexibility to dynamically define how to deliver applications under various scenarios.

Planning strategies for fast rollout

Organizations typically migrate to a new OS either all at once or in stages coinciding with hardware refresh cycles. Migrating all at once can require a substantial up-front investment of personnel and budget and can increase strain on IT staff as users abruptly change to new applications, but the overall transition is shorter than a gradual migration.

For a simultaneous migration, virtualization can serve applications that do not yet support Windows Vista, allowing an early migration date. Old and new applications can be offered concurrently to help smooth user transition. For applications that should run locally, application streaming can facilitate reliable automatic updates for patches.

Organizations that migrate in planned phases face different opportunities and challenges. For example, gradual migration can maximize the value of older client systems before retirement, and can help reduce budget spikes. However, it can also create delays when a large fraction of users have migrated to Windows Vista. A common challenge is supporting multiple versions of the same applications across multiple operating systems.

For gradual migrations, application virtualization can serve new applications to old clients and old applications to new clients. This flexibility allows end users to migrate applications at their own pace, independently of the OS on the client. Virtualization helps maximize business continuity. Streaming helps simplify deployment and updating of locally run applications, as well as helping prevent application conflicts on old and new operating systems.

Many organizations are expected to transition users to the 2007 Microsoft Office release all at once because of the new file formats. Virtualization and streaming allow client systems that do not meet the minimum system requirements or that run a previous-generation OS to use 2007 Microsoft Office, helping organizations migrate sooner and at a lower cost than they could without virtualization and streaming.

Smoothing the way to Windows Vista

IT organizations should develop a comprehensive strategy for migrating to Microsoft Windows Vista and the 2007 Microsoft Office release that takes into account application compatibility, budgets, and other resources available for their particular environments. Regardless of the specific strategy, by separating applications from operating systems, application virtualization and application streaming can help remove many obstacles to early adoption of Windows Vista and 2007 Microsoft Office.

Chris Harget *is a product marketing manager at Citrix.*

Presenting Dell Digital Cinema

Coming Soon to a Theater Near You

BY FRANKLIN FLINT

Movie theaters worldwide are rapidly switching from film-based movie projection to digital playback technology. To help theaters manage this transition, the Dell OEM Industry Solutions Group has developed a cost-effective, high-performance digital cinema solution that is designed to be reliable and easy to support.

Dell's digital cinema solution offers significant cost advantages, with hassle-free support for theater owners

For over a century, the movie industry has used film as the medium on which to distribute and play back movies. Film is easy to use, and movie theaters have experience in owning and operating film projection systems. But film is also expensive. A single film print, for example, can cost as much as US$2,000 to create and weigh over 50 pounds. A movie in wide release is generally shown on more than 4,000 screens, which can mean a cost of US$8 million in prints plus the cost of shipping heavy film canisters to movie theaters around the globe. Furthermore, a film print lasts only for roughly 30–40 showings, so shipping and print costs for successful movies can easily double or triple as film prints wear out and need to be replaced.

Now, the high cost of movie distribution is about to change as the movie industry transitions to the digital distribution and playback of movies. With digital cinema, movies are distributed via hard drive or satellite, avoiding the need to create and ship expensive, heavy film prints. Once at the theater, movies are stored digitally and played back using digital projectors. Movies stored on digital media do not wear out as film prints do, so they do not need to be replaced during long runs. Overall, the movie industry anticipates such tremendous savings from the switch to digital media distribution that it is expected to convert entirely from film to digital media in the next few years.

Digital cinema offers significant benefits to theater owners as well. Unlike film prints, digital cinema does not limit the number of screens showing a given movie, so each theater complex needs only one copy of a particular movie. In addition, theater owners can use digital projection equipment to show non-movie entertainment such as live music and sports.

However, movie theaters must undergo a technology makeover for digital cinema to work—changing from tried-and-true film-based projection systems to digital storage and projection. Many theater owners are reluctant to switch to a digital infrastructure because they are concerned about the cost and manageability of such an environment. To help movie theaters make the transition to digital technology, the Dell OEM Industry Solutions Group is working with digital cinema vendors to develop a robust, high-performance solution that is also cost-effective and easy to support.

Meeting strict industry standards for digital cinema quality and reliability

Although digital cinema has the potential to deliver tremendous cost benefits compared with film, the movie industry requires any digital cinema playback solution to deliver the same level of quality and reliability as film. To address this concern, six major motion picture studios—Walt Disney, Fox, Paramount, Sony Pictures Entertainment, Universal, and Warner Bros.—formed Digital Cinema Initiatives (DCI). This joint venture was created to establish voluntary specifications for an open architecture that is designed to ensure a uniform, high level of technical performance, reliability, and quality control for digital cinema. To receive movies digitally, theaters must implement a digital cinema storage and playback solution that meets stringent DCI standards for picture and sound quality, storage capacity, network performance, and reliability.

For example, theater-quality digital movies require a significant amount of storage—potentially as much as 500 GB for a full-length feature. As a result, DCI standards specify that any digital cinema solution make available a minimum of 1 TB of storage per screen. Also, movies require high-performance network bandwidth, so the DCI standards specify a minimum network throughput between central storage and each individual screen of approximately 400 MB/sec. For actual projection of movies, the DCI specifies that the image and sound quality of a digital projector be comparable to that of 35 mm film.

Perhaps the most important quality of any digital cinema playback solution is reliability. It is essential to the moviegoing experience that a movie not be interrupted, so DCI standards regarding continuity and reliability are especially stringent. According to the DCI standards, a movie in progress should not be interrupted during presentation for any reason except in the case of a catastrophe, such as a natural disaster. Also, DCI standards stipulate that any failure within the system be repaired within two hours, and that the mean time between failures should be no less than 10,000 hours. A digital cinema solution must have the redundancy and reliability to meet these strict standards.

Movie theater operators are generally not used to working with digital equipment, not to mention meeting strict DCI standards. A digital cinema solution needs to be easy to install, use, manage, support, and upgrade. Also, because theater owners are used to purchasing relatively inexpensive film projection technology, a digital cinema solution must be cost-effective.

Smoothing the transition to digital cinema

For many years, the Dell OEM Industry Solutions Group has worked with companies and vendors to design, deploy, and support vertical original equipment manufacturer (OEM) solutions in a wide range of industries. (For more information, see the "Ticket to Growth" sidebar in this article.) To help movie theaters make the transition to digital cinema, the Dell OEM Industry Solutions Group is working with several digital cinema vendors to design and develop solutions that meet the stringent DCI standards for quality, performance, and reliability in a cost-effective, manageable way.

The Dell OEM Industry Solutions Group offers a cost-effective digital cinema infrastructure based on Dell servers, storage, and networking products that is designed to be robust, reliable, and easy to manage. Because movie theater complexes vary in size and requirements, Dell does not offer a single, specific digital cinema implementation but rather a general design that can be scaled up or down to meet the individual requirements of a particular theater complex.

Deploying a scalable digital cinema solution: How it works

Storage is at the heart of any digital cinema solution. To meet the massive storage requirements of digital cinema, the Dell OEM Industry Solutions Group recommends a Dell™ PowerVault™ MD3000 modular disk storage array supported by one or two Dell PowerVault NX1950 networked storage systems running the Microsoft® Windows® Unified Data Storage Server 2003 OS (see Figure 1). The PowerVault MD3000 is a Serial Attached SCSI (SAS) array designed to support up to 15 SAS or Serial ATA (SATA) II hard drives, for a maximum capacity of 7.5 TB of storage—enough to support up to seven screens. For theater complexes that require additional storage, up

Ticket to Growth

The Dell OEM Industry Solutions Group offers digital cinema vendors and their customers extensive services—including configuration, deployment, support, and upgrade assistance—that make it easy for a theater owner to operate a digital cinema solution. All services, such as two-hour on-site repair service, are available through a single point of contact so customers do not have to worry about whom to call for what issue.

For many years, the Dell OEM Industry Solutions Group has worked with companies and vendors to design, deploy, and support vertical OEM solutions in industries ranging from aerospace to telecommunications. By leveraging Dell's trusted technology and flexible, cost-efficient manufacturing processes, the Dell OEM Industry Solutions Group enables customers to focus on their competitive strengths, decrease time to market, and broaden product offerings without adding complexity or sacrificing quality.

For more information about the Dell OEM Industry Solutions Group, visit www.dell.com/OEM.

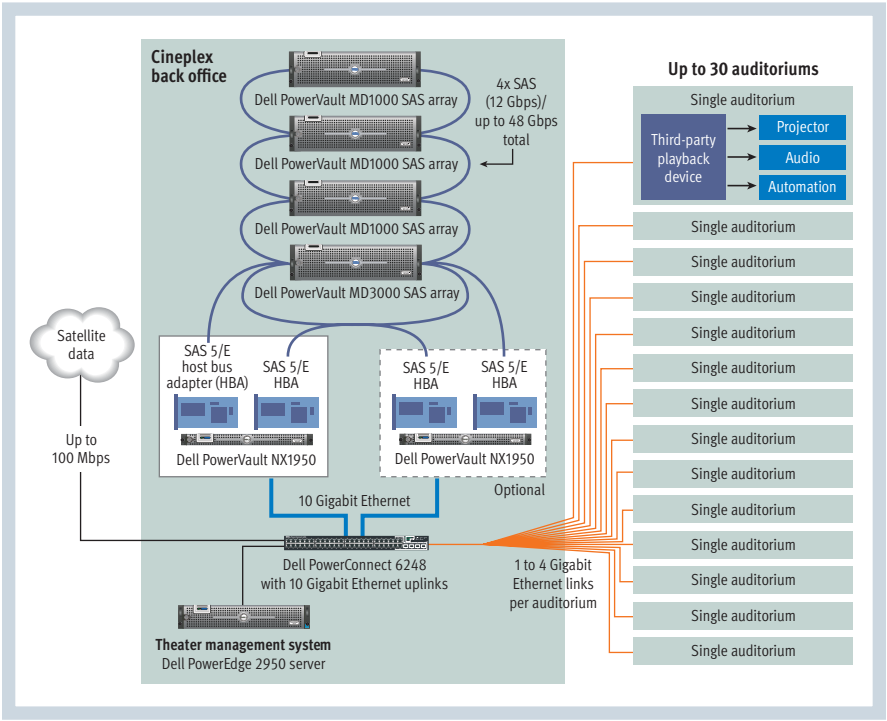

Figure 1. *Best-practices digital cinema architecture is based on highly scalable Dell servers, storage, and networking components to meet the individual requirements of each theater complex*

to three Dell PowerVault MD1000 disk expansion enclosures can be added for a maximum storage capacity of 60 TB—enough to support up to 60 screens.

Besides providing massive storage capacity, a configuration based on the PowerVault MD3000 modular disk storage array and PowerVault NX1950 networked storage server is designed for reliability and ease of use. For example, the PowerVault MD3000 SAS array features multiple redundancies throughout the system to help ensure continuity even if an individual component fails. Additionally, the PowerVault NX1950 storage server is designed to work with the PowerVault MD3000 SAS array to facilitate storage deployment and help simplify storage management. Together, the PowerVault MD3000 SAS arrays and PowerVault NX1950 networked storage servers offer outstanding storage capacity and performance.

To provide for the network bandwidth required by digital cinema, Dell recommends a 10 Gigabit Ethernet network connected to one or more Dell PowerConnect™ 6248 switches, which provide optional 10 Gigabit Ethernet uplinks and support advanced Layer 3 routing and multicast protocols to help reduce congestion and manage network traffic. Configuring one to four Gigabit Ethernet links to each auditorium exceeds the per-screen bandwidth requirements specified by the DCI and also allows for simultaneous screening of multiple movies from the central storage array.

In addition to the storage and network infrastructure, digital cinema vendors must also provide a theater management system that allows theater personnel to perform basic functions such as routing movies to individual screens and beginning playback of a movie. To support each vendor's theater management system, Dell recommends a Dell PowerEdge™ 2950 server. (Figure 1 does not show digital projection systems, which also are expected to be provided by digital cinema vendors.)

Designing for scalability and ease of management

In addition to meeting the performance and reliability standards of the DCI, a best-practices digital cinema solution from the Dell OEM Industry Solutions Group is designed to be flexible and easy to manage. For example, all Dell digital cinema solutions are built using scalable, standards-based Dell components that can be custom configured to suit the unique needs of individual theater owners. Additionally, Dell components can be easily integrated with other systems and easily expanded to enhance performance and functionality, minimizing the cost of expensive integrations and upgrades.

Digital cinema solutions have the potential to revolutionize the way the movie industry shoots, distributes, and plays back movies. The Dell OEM Industry Solutions Group is committed to helping theater owners join this digital revolution by offering cost-effective, high-performance digital cinema solutions that are designed to be reliable and easy to support. Let the show go on!

Franklin Flint *is a systems consultant in the Advanced Systems Group supporting the Dell OEM Industry Solutions Group. He has been at Dell for 12 years, and currently focuses on audio/video technology in emerging business verticals such as IP television, voice over IP, digital cinema, video surveillance, and audio recording.*

more Online
www.dell.com/**powersolutions**

QUICK LINKS

Dell Business Solutions Center:
www.dell.com/solutions

Dell PowerVault storage:
www.dell.com/powervault

Dell PowerConnect 6248 switch:
www.dell.com/content/products/productdetails.aspx/pwcnt_6248?c=us&cs=555&l=en&s=biz

Digital Cinema Initiatives:
www.dcimovies.com

Benchmarking SAP Customer Relationship Management Software on Dell PowerEdge Servers

BY MORTEN LODERUP

Implementing appropriate hardware and software configurations is key to optimizing performance in complex data centers. This article describes results from the two-tier SAP® Interaction Center Benchmark for the mySAP™ Customer Relationship Management 2005 application on Dell™ PowerEdge™ servers.

Related Categories:

Application servers

Benchmarks

Dell PowerEdge servers

Dell/EMC storage

Dual-core technology

SAP

Visit www.dell.com/powersolutions for the complete category index.

SAP customer relationship management software enables enterprises to help increase efficiency by providing rapid access to critical real-time information and tight integration with the components of a customer relationship cycle, including telephone, e-mail, and Internet systems; mobile clients and handhelds; and existing SAP software. It is built on the SAP NetWeaver® platform and the SAP enterprise service-oriented architecture (SOA), a business process–driven approach designed to enhance competitive advantage, dynamically adjust to change, and help ensure consistent productivity. Specific SAP customer relationship management software is available for marketing, sales, service, e-commerce, interaction center operations and management, and channel management.

To help demonstrate the hardware requirements and performance of SAP customer relationship management software, in September 2006 the Dell SAP Competence Center performed tests using the two-tier SAP Interaction Center Benchmark for the mySAP Customer Relationship Management (mySAP CRM) 2005 application. The tests showed that three variables can have a significant effect on mySAP CRM performance in a Dell hardware environment: the number of instances, the number of dialog work processes, and the sequential memory setting.

SAP Interaction Center Benchmark

To help enterprises understand the performance of a particular SAP application in representative hardware environments, SAP provides benchmarking and certification programs. After SAP receives a benchmark and certification from a hardware partner and the results are approved, SAP makes the results available on its Web site at www.sap.com/benchmark. Enterprises can use this data to help create an optimal environment in which a large number of users can access the hardware with low response times.

The SAP Interaction Center Benchmark for the mySAP CRM 2005 application (hereafter referred to as the IC Benchmark) focuses on typical interaction center activities based on SAP Interaction Center WebClient functions. It includes the following scenarios:

- Incoming call with display of contact history and creation of interaction record
- Contact with follow-up activity
- Search of contact history and display of historical record
- Inbound sale with creation of sales order
- Service call with creation of service ticket

Each user carries out one scenario during each test, with five user types running five different scenarios in parallel. The IC Benchmark was certified by SAP on October 4, 2006 (certification number 2006078). For additional details on the steps carried out in each scenario, visit www.sap.com/solutions/benchmark/ic.epx.

Test environment

The hardware used for the IC Benchmark included a Dell PowerEdge 6850 server with dual-core Intel® Xeon® 7040

Server	**Model**	Dell PowerEdge 6850
	Processors	Four dual-core Intel Xeon 7040 processors at 3.0 GHz, each with one 16 KB L1 cache and one 2 MB L2 cache per core
	Memory	32 GB
	OS	Microsoft Windows Server 2003 Enterprise x64 Edition
Client	**Model**	Dell PowerEdge 2500
	Processor	Intel Pentium III processor at 993 MHz
	Memory	4 GB
	OS	Microsoft Windows Server 2003 Enterprise Edition
Storage array	Dell/EMC CX600	
HBAs	Two QLogic QLE2460 HBAs	
Switches	Two Brocade SilkWorm 3800 Fibre Channel switches	
LUN RAID level	RAID-10	
Software	• mySAP CRM 2005 • Microsoft SQL Server 2005 • EMC PowerPath	

Figure 1. *Hardware and software used in the test environment*

processors and a Dell/EMC storage area network (SAN). The PowerEdge 6850 server ran the Microsoft® Windows Server® 2003 Enterprise x64 Edition OS along with mySAP CRM 2005 and Microsoft SQL Server™ 2005 software. A Dell PowerEdge 2500 server was used as the client system. Figure 1 summarizes the hardware and software used in the test environment.

The hardware test architecture (see Figure 2) used a two-tier configuration in which the database and applications ran on a central application server (the PowerEdge 6850) and the mySAP CRM front end ran on a client (the PowerEdge 2500). The client connected to the server through a Gigabit[1] Ethernet connection. The application server sent database requests to the SQL Server 2005 database residing on a RAID-10 logical unit (LUN) on the Dell/EMC CX600 storage array through redundant QLogic QLE2460 host bus adapters (HBAs) and two Brocade SilkWorm 3800 Fibre Channel switches. The switches were zoned so that only one initiator (HBA) and multiple targets (LUNs) existed in each zone. The storage array also connected to the same physical switch and the same logical zone. EMC® PowerPath® software on the application server provided optimized HBA load balancing and failover. The multiple paths between the storage array and each configured LUN helped ensure that the requesting application still had access to its data even following a component failure.

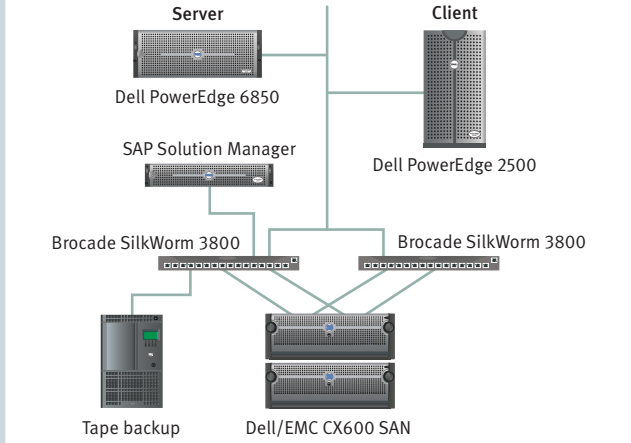

Figure 2. *Hardware architecture used in the test environment*

Test results

The IC Benchmark tests revealed several factors that can affect the response time and processor utilization when using the mySAP CRM application, most significantly the number of instances, the number of dialog work processes, and the sequential memory setting.

Number of instances

An instance consists of a runtime environment to process user requests. It could be equivalent to an application server (the physical host), but an application server could also have several instances running simultaneously. The runtime environment includes dialog work processes and a dispatcher, and may also have update, spool, and background work processes. Because each of these processes consumes server resources, setting an appropriate number of instances helps optimize both the server hardware resources and throughput.

To determine the optimal number of instances in the test environment, the test team ran the IC Benchmark with a user load of 300, 13 dialog work processes, sequential memory access disabled, and from one to four instances, and tracked response time and processor utilization for each configuration. As Figure 3 shows, with one instance, the response time was slow relative to tests with more instances, and the processor utilization was low. The response time decreased for both two and three instances, and then increased with the addition of a fourth instance. The processor utilization increased with each additional instance. The optimal number of instances in the test environment was three, which provided the fastest response time and good processor utilization.

Number of dialog work processes

Dialog work processes handle screen changes and interaction with the mySAP CRM graphical user interface (through the dispatcher). By default, two dialog work processes are assigned to each instance; the maximum number of dialog work processes is 99 per instance. Because running too

[1]This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

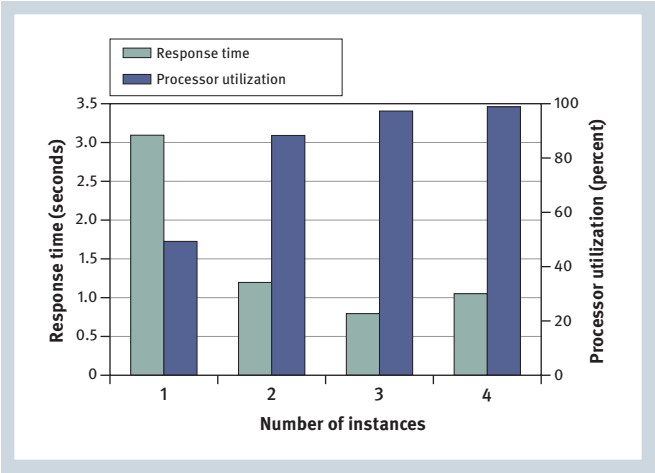

Figure 3. *Response time and processor utilization for different numbers of instances*

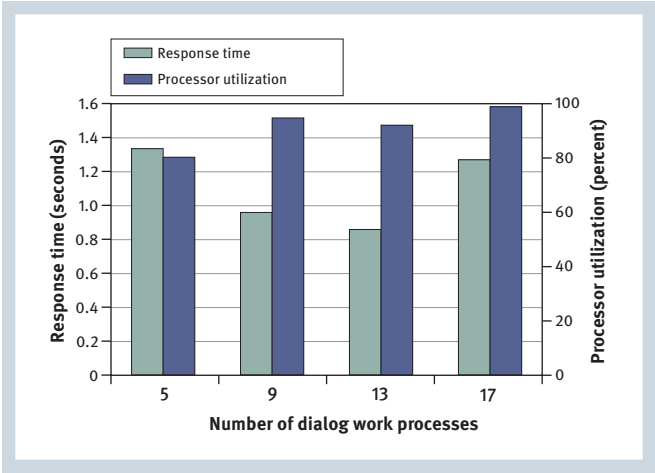

Figure 4. *Response time and processor utilization for different numbers of dialog work processes*

few processes does not make optimal use of available hardware, and running too many can interfere with the performance of critical applications, configuring the appropriate number is key to optimizing the environment.

To determine the optimal number of instances in the test environment, the test team ran the IC Benchmark with a user load of 300, three instances, sequential memory access disabled, and from 5 to 17 processes, and tracked response time and processor utilization for each configuration. As Figure 4 shows, the optimum number of processes for the test environment was 13, which provided the fastest response time and good processor utilization.

Sequential memory setting

Sequential memory access is set in the BIOS processor settings, and allows prefetching to bring data into memory or the cache before it is processed. It is typically enabled by default. When enabled, it optimizes the system for applications that require sequential memory access. If data is usually accessed in the same order, enabling it may be the best choice; however, if data is accessed randomly (for example, in an SAP database), disabling it may increase performance.

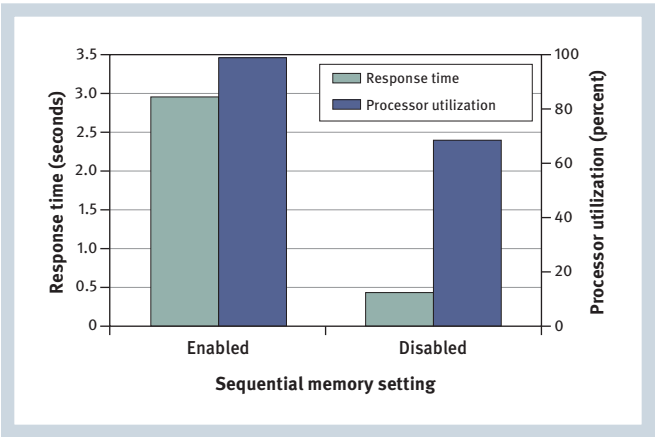

Figure 5. *Response time and processor utilization for enabled and disabled sequential memory access*

The test team ran the IC Benchmark with a user load of 250, three instances, 13 dialog work processes, and sequential memory access both enabled and disabled to determine which setting provided the best results for mySAP CRM 2005 in the test environment. As Figure 5 shows, disabling sequential memory access did increase performance by decreasing both response time and processor utilization.

Optimized customer relationship management

When implementing SAP software, enterprises should consider how both hardware and software configuration affect performance. The SAP Interaction Center Benchmark tests performed by the Dell SAP Competence Center showed that tuning certain parameters in mySAP CRM 2005–based systems—particularly the number of instances, the number of dialog work processes, and the sequential memory setting—can deliver significant performance increases, helping improve response times while maintaining an efficient level of processor utilization. Enterprises can take this data into account when deploying SAP customer relationship management software to help optimize performance and thereby increase their return on investment.

Morten Loderup *is a solution architect in the Dell SAP Competence Center. His interests include high-availability clustering, SANs, and SAP solutions.*

more online
www.dell.com/**powersolutions**

QUICK LINK

Dell and SAP:
www.dell.com/sap

How Dell Streamlined Authentication and Identity Management Using Quest's Vintela Authentication Services

BY DAVE WILSON AND ROMMEL MERCADO

IMPLEMENTATION STUDY

CHALLENGE

Integrate authentication and identity management for Microsoft Windows, UNIX, and Linux platforms into Microsoft Active Directory across the entire Dell infrastructure

SOLUTION

Vintela Authentication Services from Quest Software deployed with the help of Quest Professional Services

BENEFITS

- Streamlined Active Directory–based system helps simplify authentication and identity management and frees IT staff to focus on other projects, helping increase operational efficiency
- Consistent cross-platform approach helps eliminate common compliance problems and reduce audit costs without requiring additional infrastructure
- Advanced features enable directory consolidation and identity migration at the pace—and according to the needs—of each organization's unique circumstances
- Automated provisioning and de-provisioning can help increase security in the future
- Implementing the solution in-house allows Dell to confidently recommend the solution to its own customers facing the same challenges

Related Categories:

Authentication, Linux, Microsoft Active Directory, Quest Software, Security

Visit www.dell.com/powersolutions for the complete category index.

When security, management, and compliance demands required Dell authentication and identity management systems to be consolidated into a single common directory, the Dell IT group turned to Quest Software's Vintela® Authentication Services, which enabled the integration of Microsoft® Windows®, UNIX®, and Linux® platforms with the Microsoft Active Directory® directory service.

Compliance concerns can drive enterprises to reexamine the way they handle identity and authentication. At the core of many regulations—including the Sarbanes-Oxley Act (SOX), the Health Insurance Portability and Accountability Act (HIPAA), and the Graham-Leach-Bliley Act—is the mandate to securely authenticate users, grant appropriate access to data and resources, and track the activities surrounding that authentication and access. As enterprises become increasingly complex and diverse, the challenges presented by authentication and access are mounting. Compliance requires effective solutions to the challenges of identity across multiple, disparate operating systems. Dell is no exception: the Dell IT group faced such challenges when consolidating its heterogeneous authentication and identity management systems.

Managing a diverse environment

"Along with Microsoft Windows as our core platform, we also have a wide mix of systems running UNIX and Linux," says David Taylor, principal Linux engineer at Dell. "We have a significant installed base of servers running three different versions of Linux from two different vendors, plus a very small subset of systems using different versions of AIX and Solaris. All in all, we have about 2,200 combined UNIX- and Linux-based systems. Primarily, these platforms act as database servers for systems like our core Oracle business systems, but they do fill other functions as well."

But why not just base everything on a single platform—like Windows—and eliminate the problems heterogeneous systems present? For many organizations, the answers lie in the financial advantages of a multi-vendor approach and the fact that many mission-critical applications require specific platforms. In addition, organic growth can often introduce diverse systems. Although those answers are valid for Dell as well, the company has another compelling reason.

"We need to stand on our corporate principles," says Taylor. "And that means proving that Dell works on Dell. In accordance with this strategy, we have been migrating our UNIX-based systems to Linux for a number of years. Today, our standard Oracle database server runs Linux, but we still have a handful of legacy AIX- and Solaris-based systems running as well."

Providing centralized authentication and granting secure access in this diverse mix of operating systems would be a challenge for any organization. Dell's traditional approach to

> "In fairly short order, we have eliminated our security, compliance, and management concerns with cross-platform identities. And I think we are not even using Vintela Authentication Services to its fullest yet."
>
> —David Taylor
> Principal Linux engineer at Dell
> February 2007

authentication across heterogeneous systems was to build a single domain and authenticate through the Server Message Block (SMB) protocol. Unfortunately, this approach requires distributed accounts across all systems. "Passwords were held in the domain, but for that technology to work, you still must have a local account on every computer," says Taylor.

This approach was adequate for the short term because Dell had relatively low personnel turnover, but it still presented management and security concerns. "With our old approach, any kind of turnover meant you had to touch all of the systems to modify the user lists," says Taylor. "No one is going to be perfect in this area—or at least we were not perfect about it—so we would end up leaving user accounts out there for people that have moved on to other jobs when we should have been terminating their access immediately. We were creating a security problem as well as a management problem. We built some scripts to de-provision from a centralized location, but it still is not a very effective way to manage user accounts."

Consolidating authentication and identity management systems

Dell realized that security, management, and compliance demands would require centralizing authentication and bringing all user accounts into a single common directory for the entire enterprise. So the company launched a project called Multi-Platform Management Integration (MPMI). Its goal was to make the Microsoft Active Directory directory service the authoritative authentication system and master source for all user accounts across all systems within Dell—those running Microsoft Windows, IBM® AIX, Sun Solaris, and the various Linux operating systems.

"We could have created another directory and either passed authentication through it or run some sort of synchronization," says Taylor. "But that would not be as simple and elegant as going directly into Active Directory."

Adds Tony The, the MPMI project manager, "Active Directory is already the company standard directory, and every user already has an Active Directory account. Active Directory supports our current needs, and from a management perspective, our people already know and understand it."

With the decision made to expand the influence of Active Directory to include UNIX- and Linux-based systems, the next challenge Taylor and The faced was how to execute the project. Active Directory authentication is based on the Kerberos encryption standard and the Lightweight Directory Access Protocol (LDAP). Because Active Directory is based on industry standards, several open source technologies are available for integrating UNIX and Linux with Active Directory.

"A couple of years ago I looked at integrating our UNIX- and Linux-based systems with Active Directory using open source technologies based on Kerberos and LDAP," says Taylor. "But with the size of our Active Directory structure, it really was not a workable solution. It would have required LDAP to constantly query the directory for group membership, creating a significant burden on the network. Without a local caching mechanism, open source solutions really are not a viable option for an organization of our size. We also considered using Winbind, but that would not work either, because it requires a mapping database on every server."

After ruling out open source alternatives, Taylor began a search for a third-party solution that provided the functionality the MPMI project demanded as well as the support and stability of a commercial product. One of the project architects began doing Internet research looking for vendors that offered the type of solution Dell required. One solution quickly rose to the top.

Finding the solution: Vintela Authentication Services from Quest

"Our research led us to Quest Software and a product called Vintela Authentication Services," says Taylor. "It offered the features we needed and allowed our UNIX- and Linux-based systems to join Active Directory. With a very tight time frame—we needed to implement a solution as quickly as possible—we gave Vintela Authentication Services a two-day proof-of-concept test in our lab, purchased the solution, and began a rollout plan."

The Vintela Authentication Services software is installed on UNIX- and Linux-based systems and integrates the native identity and authentication mechanisms of each OS with the Kerberos and LDAP components of Active Directory. Fundamentally, Vintela Authentication Services allows the AIX, Solaris, and various Linux platforms to act as full citizens in Active Directory. It helps eliminate the need for local accounts on each non-Windows system, leverages the secure authentication already present in Active Directory for UNIX and Linux, and allows other advanced Active Directory functionality, such as Group Policy, password policies, Windows security policies, and single sign-on. Vintela Authentication Services extends to many popular and widely deployed UNIX and Linux platforms.

"Quest Professional Services guided us through the design of the solution," says Taylor. "With so many differences across platforms and all that we wanted to accomplish, it was a more involved process than it originally appeared. Quest helped us design the Group Policy Objects we used in our rollout. The deployment began in July 2006. Today, the Vintela Authentication Services client is installed on all of our UNIX- and Linux-based servers, and they are authenticating

> "Quest Software delivered a comprehensive authentication solution that did not require any additional infrastructure. It mirrors Windows authentication as closely as we could have hoped."
>
> —Tony The
> Multi-Platform Management Integration project manager at Dell
> February 2007

against Active Directory. We also are now preinstalling the client on every new Linux-based system provisioned within Dell."

Planning for the future

With the UNIX and Linux platforms joined to the Active Directory domain, the next step in the Dell MPMI project is to migrate the local user accounts from each non-Windows system into Active Directory. Vintela Authentication Services offers several advanced features to enable directory consolidation and identity migration at the pace—and according to the needs—of each organization's unique circumstances. These options range from simply leveraging Active Directory for passwords and authentication while maintaining existing UNIX and Linux structures, to moving existing UNIX and Linux structures into Active Directory as a subset of the Active Directory user account, to fully migrating from multiple, disparate identities to a single Active Directory–based identity for all systems. "In the future, we expect to migrate all of the accounts into Active Directory," says The.

Adds Taylor, "We plan to set up our own import file. Our accounts are pretty straightforward; they are either local or already use Active Directory. Next, we plan to use the Vintela Authentication Services Ownership Alignment Tool to resolve conflicting file ownerships as we move from multiple UNIX and Linux accounts to a single Active Directory account.

"When MPMI is complete, we should no longer need local user accounts on our UNIX- or Linux-based systems—just required system accounts such as root and bin," says Taylor. "All the systems that are provisioned these days have no local user accounts on them, and that is the way they stay. Everything going forward is based on Active Directory, so systems start out clean. In fairly short order, we have eliminated our security, compliance, and management concerns with cross-platform identities. And I think we are not even using Vintela Authentication Services to its fullest yet. I would estimate that just streamlining operations through directory consolidation and centralizing authentication is freeing up one or two people each year. We can put those people on more important and interesting projects than managing authentication and user accounts."

"Beyond the obvious operational expense savings this solution has brought us, an even more important benefit is cost avoidance," says The. "We are going to be growing our Linux environment quite a bit going forward, and Vintela Authentication Services can help us avoid a lot of cost simply by eliminating one of the areas that may cause problems in our compliance audits. The benefit to Dell is a combination of cost avoidance and operational efficiency."

In the future, Dell plans to gain additional benefits from its integrated identity environment by leveraging Active Directory Group Policy to control security-related parameters in the UNIX and Linux environments. In addition, future projects may involve automating provisioning and de-provisioning for all systems based on Active Directory and extending the benefits of Vintela Authentication Services to additional UNIX applications.

"My focus is security, and from that perspective, Vintela Authentication Services did exactly what we needed it to do," says The. "Quest Software delivered a comprehensive authentication solution that did not require any additional infrastructure. It mirrors Windows authentication as closely as we could have hoped."

Building a relationship

According to Mark Witucki, the account manager at Quest Software who managed the relationship with Dell, "The fact that this project was extremely successful has really catapulted the relationship to the next level. We have felt how important this was to the Dell field."

"Quest and Dell have a relationship that extends well beyond the simple vendor-customer relationship that our MPMI project introduced," says Taylor. "Many Dell customers are running into the same challenges we encountered—compliance, security, and managing identity in a heterogeneous environment. People are always asking how we manage our own systems. It is nice to be able to say that we have implemented a centralized solution for the same problem these Dell customers face. The fact that it works in-house at Dell, and works so well, allows our sales force to confidently offer a similar solution to Dell customers.

"Vintela Authentication Services is the best product we have found on the market," concludes Taylor. "It satisfies our needs and can help us expand where we need to in the future."

Dave Wilson *is the vice president of identity management and interoperability at Quest Software.*

Rommel Mercado *is the senior manager of the IT Core Platform Engineering team at Dell.*

more Online
www.dell.com/**powersolutions**

QUICK LINK

Quest identity management solutions:
www.quest.com/IdM_Dell

DELL OPENMANAGE NEWSLETTER
Flexible Management for the Scalable Enterprise
May 2007

View from the Top

Simplifying IT Management: Two Choices, No Compromise

Dell offers two ways to help simplify IT management without compromising functionality: Dell™ systems management tools built on the Altiris® infrastructure or the systems management tool of your choice from a Dell OpenManage™ Certified Partner.

By Edward Reynolds, Senior Manager, Systems Management Product Marketing, Dell Inc.

Systems management tools are supposed to simplify IT management. But today's complex IT environments require so many tools—tools for managing clients, tools for managing servers, tools for managing storage, and so on. Sometimes it seems like management tools are adding complexity instead of taking it away. At Dell, we're building two different ways to reduce the number of tools required to manage your IT environment without sacrificing management functionality.

Dell OpenManage Certified Partner tools

If you would like the flexibility to use your own systems management vendor, the Dell OpenManage Certified Partner Program certifies systems management products that integrate Dell hardware management—allowing you to seamlessly manage Dell hardware from the Dell-certified management tool of your choice.

In particular, we provide Dell hardware management developer toolkits that enable participating vendors to add functionality such as deployment and change management for Dell servers to their applications. When the integration is complete, Dell engineers test the applications to certify that they deliver the required functionality. Vendors whose applications pass are designated as Dell OpenManage Certified Partners and granted a Dell OpenManage Certified Partner logo to affix to certified products.

Already, a wide range of systems management vendors such as Microsoft, Quest Software, Novell, BMC Software, and CA are in discussions with Dell about participating in the program. Altiris and LANDesk have signed and are scheduling their certification tests. The Dell OpenManage Certified Partner Program gives our customers not only the capability of managing their Dell hardware without having to deploy a separate tool, but also the flexibility to choose the systems management console that fits their environment the best.

Dell integrated systems management suite

If you prefer a systems management solution from Dell, we offer the comprehensive Dell OpenManage portfolio of deployment, monitoring, and update tools for Dell hardware. And in November 2006, we announced our intent to build our next-generation console in partnership with Altiris. For basic administration, this suite will remotely discover, inventory, monitor, configure, and patch Dell hardware assets including clients, servers, storage, printers, and more at no additional charge. For advanced functionality such as deployment and imaging, software delivery, and backup and recovery management, you will be able to purchase plug-in modules that seamlessly integrate into the Dell console.

Furthermore, because this suite will leverage the Altiris infrastructure, advanced feature modules will be readily available from Altiris such as asset management, OS/application patching, virtualization management, and help-desk capability.

The choice is yours

Dell is committed to helping you simplify the task of IT management. Choose a Dell Certified Partner and manage your Dell hardware assets seamlessly without having to use a separate tool. Or choose Dell systems management tools and manage your IT environment in a simple, scalable, pay-as-you-grow way—all from a single console. Either way, the choice is yours. ■

TECH CORNER

By David Weber, Enterprise Technologist, Dell Inc.

Unified Manageability Architecture: A Blueprint for Flexible Systems Management

Currently, most systems management architectures are proprietary and monolithic—and increasingly unsuited to the challenges posed by today's complex, heterogeneous IT environments. Unified Manageability Architecture is designed to facilitate systems management by providing a layered, standards-based blueprint for the development of flexible, scalable, and open systems management infrastructures.

Today's systems management marketplace is filled with proprietary, best-of-breed tools that are good at what they do, but don't easily integrate with one another. Part of the reason for this is that there has never been a universally accepted systems management model that standardizes systems management terminology, functionality, and interfaces in a single, unified architecture. The result is that, in addition to ending up with a plethora of individual tools that cannot integrate with one another, IT organizations have no standard way of modeling the systems management requirements of their IT infrastructures as a whole, and no way of assessing whether the tools in place deliver on those requirements.

Until now.

Dell has developed Unified Manageability Architecture (UMA) based on standards and technologies from the Distributed Management Task Force (DMTF) and other standards bodies. UMA is a layered, standards-based systems management architecture that codifies systems management terminology, functionality, and interfaces. UMA is designed to simplify systems management by providing a blueprint that allows IT organizations to comprehensively model their systems management requirements and enables vendors to develop products with open, standards-based systems management instrumentation that seamlessly interfaces with standards-based management tools or consoles.

Systems management layers

UMA partitions systems management infrastructures into six layers, as shown in Figure 1. The first four layers—platform, logical mapping, aggregation, and access—address aspects of the managed device, while the remaining two—resource management and orchestration—pertain to the overall management system.

At the managed device level, the *platform layer* describes the actual physical elements being managed, such as processors, disk drives, and memory, and includes environmental information such as temperature, enclosures, and location. The *logical mapping layer* consists of the software drivers and providers used to map logical representations to the physical elements within the platform layer. The *aggregation layer* is responsible for aggregating logical representations to various data models used to correlate and describe the managed elements. Finally, the *access layer* is the communication interface between the managed device and the management system.

At the management system level, the *resource management layer* encompasses most of the functionality currently found in management applications, such as servicing user and management application requests, requesting and parsing responses for resource management functions, and creating and terminating sessions with managed nodes. The *orchestration layer* is an optional layer that provides a level of capability beyond that of traditional management systems. This layer enables automated, policy-based management through a policy engine with defined service levels.

At all levels of the architecture, UMA takes advantage of current and emerging industry standards whenever possible. Standards such as the Common Information Model (CIM) for data, the Systems Management Architecture for Server Hardware (SMASH) Server Management Command-Line Protocol (SM CLP), the Desktop and Mobile Architecture for System Hardware (DASH), and the WS-Management Web services specification provide IT administrators with a frame of reference. In this way, organizations can map their environments and provide vendors with open standards for developing flexible, interoperable management interfaces.

Standards-based product instrumentation

At Dell, we're committed to simplifying systems management. In fact, our engineers and partners are applying the principles of UMA right now to help standardize the management instrumentation on Dell products such as servers, storage, switches, client systems, and printers. These efforts allow Dell managed devices to work seamlessly with third-party, standards-based management tools such as Microsoft System Center. We know IT infrastructures are increasingly complex, but we don't think they should be complicated to manage. ■

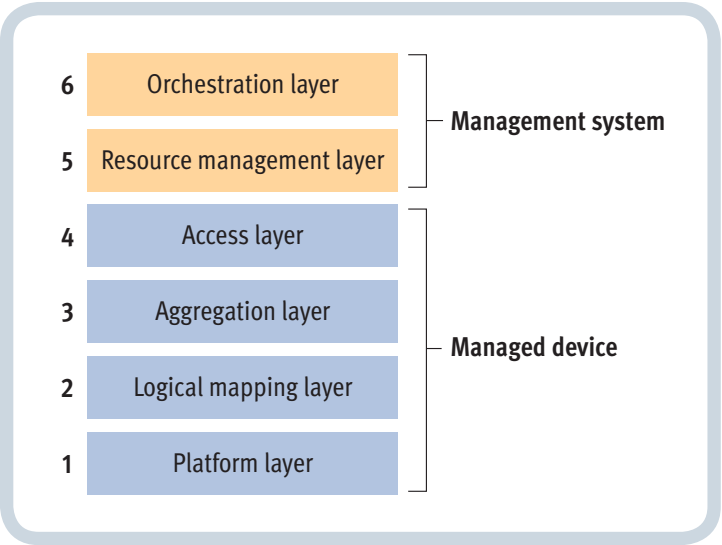

Figure 1. Unified Manageability Architecture model

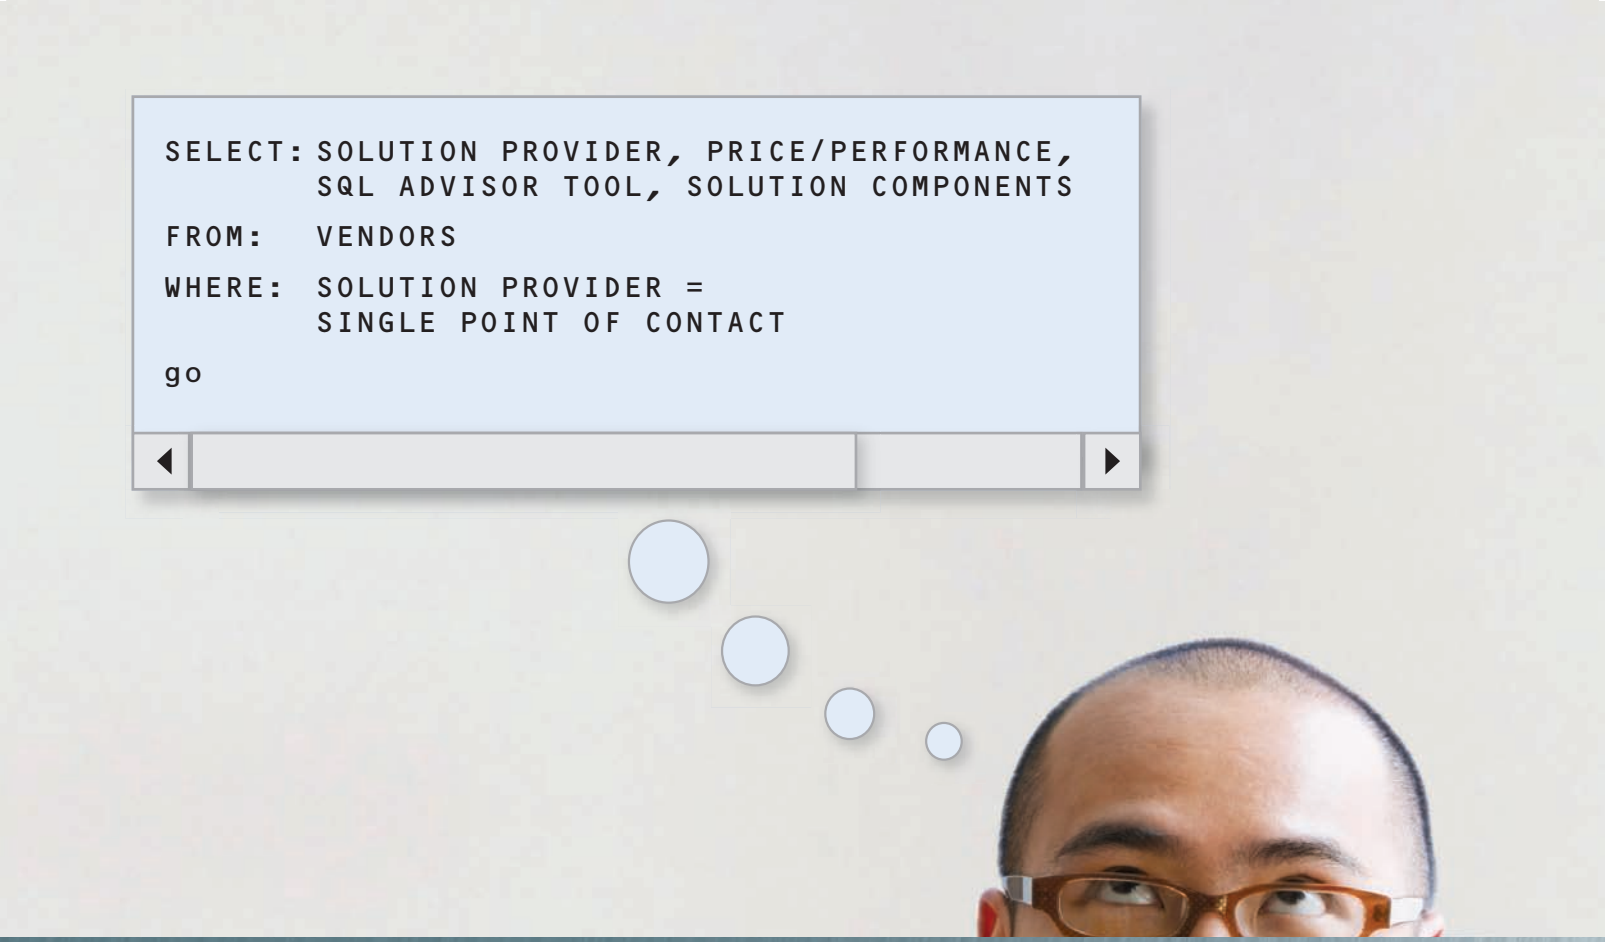

RESULT: DELL DATABASE SOLUTIONS FOR SQL SERVER 2005

Think Smart. Choose Dell.

You require a database with robust performance and high availability. But sometimes choosing the platform to run your Microsoft® SQL Server™ 2005 database can be like a complex query. Dell provides a simple answer: **Dell Database Solutions for SQL Server 2005**. Dell holds the **top five TPC-C price/performance leadership positions** for SQL Server. We offer an **easy-to-use SQL Advisor Tool** to help create an architecture for your SQL Server 2005 deployment. Plus, we provide a **single point of contact** for the entire solution — the servers, storage, systems management, services and SQL Server 2005 software.

Visit www.dell.com/sql to access the Dell SQL Advisor Tool, input your specifications for analysis and receive a SQL Server configuration based on performance sizing and Dell recommended best practices.

Microsoft SQL Server 2005

DELL

Disabling Local Video Consoles in Ninth-Generation Dell PowerEdge Servers

BY ABHAY SALUNKE
KALYANI KHOBRAGADE
THIRUMALAA SRINIVAS

The Dell™ Remote Access Controller 5 (DRAC 5) enables administrators to remotely turn off the local video consoles of ninth-generation Dell PowerEdge™ servers to help increase server management security and flexibility.

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Dell Remote Access Controller (DRAC)

Remote access controllers (RACs)

Remote management

Visit www.dell.com/powersolutions for the complete category index.

When an administrator starts a remote virtual KVM (keyboard, video, mouse) console redirection session through a Dell Remote Access Controller (DRAC), the local video console displays the remote actions being performed. Anyone who can see the screen can therefore observe all of these actions.

The Dell Remote Access Controller 5 (DRAC 5) now allows administrators to remotely disable local video consoles on ninth-generation Dell PowerEdge servers, helping them establish private remote virtual KVM console redirection sessions and easily configure servers for headless operation. This feature operates independently of remote virtual KVM sessions and the server OS, and administrators can still use other DRAC 5 features normally.

Understanding system requirements

Controlling the local video console remotely requires the following:

- The managed server must be a ninth-generation Dell PowerEdge rack or tower server with a DRAC 5 (Dell PowerEdge 1955 blade servers do not support the local video feature).
- The DRAC 5 firmware must be version 1.20 or later and support the local video feature, and the baseboard management controller (BMC) firmware must be version 1.29 or later and support this feature.
- The server must have an LCD screen to display notifications about the local video.
- The administrator must have virtual KVM access privileges.

Enabling and disabling local video consoles

Administrators can enable and disable local video consoles from both the DRAC 5 graphical user interface (GUI) and DRAC 5 command-line interface (CLI). In the GUI, they can use the Local Server Video check box in the Console Redirect Configuration screen, as shown in Figure A in the supplemental online section of this article, available at www.dell.com/powersolutions. After clearing this box and clicking "Apply Changes" to disable the console, a notification box appears on both the management station console and local video console indicating that local video will be turned off in 15 seconds; this notification provides anyone using the local video console time to take any necessary action before the console is disabled. Conversely, enabling the local video console through the GUI takes effect immediately. Figure 1 shows the DRAC 5 Console Redirect Configuration settings with local video enabled.

Administrators can also use the DRAC 5 CLI to disable a local video console by issuing the following command:

```
racadm config -g cfgRacTuning -o
   cfgRacTuneLocalServerVideo 0
```

Like the DRAC 5 GUI, after this command is entered, the CLI displays a message on both the management

station console and local video console indicating that local video will be turned off in 15 seconds. Administrators can re-enable local video with the following command:

```
racadm config -g cfgRacTuning -o
   cfgRacTuneLocalServerVideo 1
```

Also like the DRAC 5 GUI, the local video console is enabled immediately, with no notification message.

Administrators can view the current local video status by issuing the following command, which displays all objects in the cfgRacTuning group:

```
racadm getconfig -g cfgRacTuning
```

The hardware configuration enabling these interactions between a remote user, DRAC 5, and local user is shown in Figure B in the supplemental online section of this article, available at www.dell.com/powersolutions. Because the BMC controls the console video switch on the motherboard, these settings persist across reboot and power cycles.

Viewing server LCDs

The server LCD is a key component that reflects the different local video console states. When an administrator remotely disables local video, the LCD displays the message "Video off in 15" until the 15-second transition time has elapsed; once the video has been disabled, it displays the "Video off" message. If administrators discover that the local video is off on a particular server, they should first check the server LCD for this message. They should keep in mind that video messages can be masked by medium- and high-priority messages.

Enhancing remote security and flexibility

The ability to disable local video consoles with the DRAC 5 can enhance remote security and flexibility. Taking advantage of this feature can provide multiple benefits for data center administrators, including helping simplify headless server configuration, helping reduce unnecessary monitor power consumption, and allowing private remote virtual KVM console redirection sessions to prevent unwanted local monitoring.

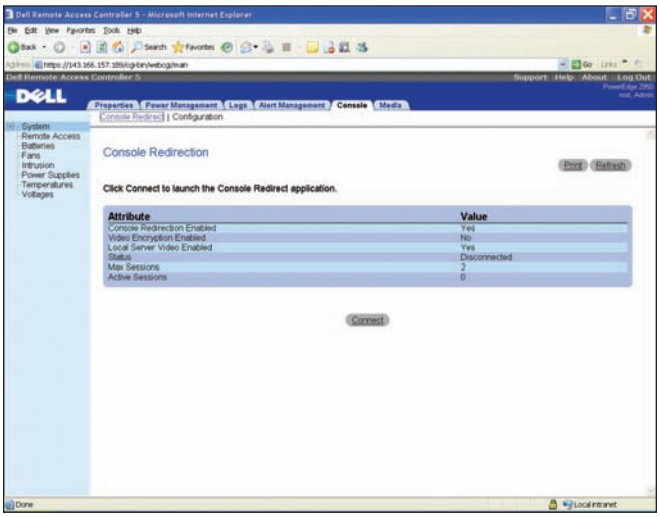

Figure 1. *Local video enabled in the DRAC 5 Console Redirect Configuration screen*

Abhay Salunke *is a senior software developer in the Remote Access Firmware Group at Dell. He is the technical lead for several DRAC projects and has contributed to various Dell OpenManage™ software development projects. Abhay has a bachelor's degree in Electronics Engineering from the University of Mumbai.*

Kalyani Khobragade *is a firmware engineer in the Remote Management Group at Dell. She has a B.S. and an M.S. in Electronics Engineering from Visvesvaraya Regional College of Engineering, and an M.S. in Electrical and Computer Engineering from Purdue University.*

Thirumalaa Srinivas *is a project engineer in the Enterprise System Test Group at Dell. His responsibilities include quality assurance for enterprise products, and his current interests include operating systems, virtualization, enterprise solutions, and product development. He has a B.E. in Computer Science from the University of Madras and an M.S. in Computer Science from the University of Texas at Dallas.*

ATI FireGL Accelerators from AMD Deliver the Graphics Power on Dell Precision Workstations

From 3D animation and industrial design to advanced medical diagnostics, ATI™ FireGL™ graphics accelerators from AMD are designed to handle the most complex 3D models, the largest data sets, and the highest-definition textures to meet the wide-ranging needs of today's graphics-intensive application environments.

Today, many production studios rely on workstation graphics to create 3D animation and cinematic special effects. The automotive industry turns to high-performance graphics for designing and rendering innovative vehicles and running simulations to help ensure they meet critical design requirements in performance, handling, aerodynamics, comfort, and safety. Architects depend on workstation graphics to create digitally modeled homes that enable prospective buyers to do a virtual walk-through before construction begins. Health care professionals demand flawless image fidelity in medical imaging applications in order to make critical diagnoses.

Lamborghini image designed and visualized using ICEM software. © 2007 ICEM. Reprinted with permission.

ATI FireGL graphics horsepower fuels innovative automotive designs

Architectural image designed and visualized using Bentley MicroStation. © 2007 Bentley Systems. Reprinted with permission.

ATI FireGL powers virtual architecture tours in real time

ATI FireGL workstation graphics accelerators from AMD are built on a scalable, ultra-threaded graphics architecture and innovative ring bus memory system designed to handle the most complex 3D models, the largest data sets, and the highest-definition textures used today. This efficient design approach to the family of FireGL products allows users to select an accelerator based on price and performance requirements without compromising capabilities or functionality. Featuring a 10-bit display pipeline and high dynamic range (HDR) 16-bit per RGB color component output, ATI FireGL accelerators are designed to produce over 1 billion colors for vibrant and accurate visual fidelity. This incredible image quality enables the precision required for specialized markets like visual simulation and medical imaging.

ATI FireGL V7200 workstation graphics accelerator from AMD

With world-class performance and image quality, it's no wonder designers, animators, and engineers rely on ATI FireGL graphics accelerators. IT managers and system administrators can also benefit substantially from the ATI FireGL accelerators' low total cost of ownership. The single, unified drivers that support the entire FireGL product family help reduce costs by simplifying system administration and maintenance for new software releases and updates. IT managers have the flexibility to select from a range of FireGL products—from entry-level to ultra-high-end—based on staff, application, and budget needs, knowing that one driver package supports the entire product line. Also, the unified driver architecture streamlines validation with professional applications across the entire FireGL product family. ATI FireGL accelerators are thoroughly tested and certified with major computer-aided design (CAD) and digital content creation (DCC) applications[1] to help ensure optimum performance and compatibility.

In addition, the FireGL drivers include an application detection feature that is designed to automatically recognize and load optimal profile settings for a specific application and to provide fully optimized driver performance when running multiple applications simultaneously. Plus, ATI offers direct customer access to a dedicated workstation technical support team should the need arise. These key support components provide peace of mind and help ensure that ATI FireGL accelerators are reliable and easy to maintain, enhancing productivity and minimizing the potential for downtime.

For outstanding high-end graphics performance, Dell Precision™ workstations are available with the ATI FireGL V7200 (256 MB) graphics accelerator, which features two Dual Link outputs that can drive two 30-inch Dell UltraSharp™ wide-screen monitors at up to 2560 × 1600 resolution each, producing a massive multi-monitor display for enhanced productivity. Providing exceptional value without compromising features or image quality, Dell Precision workstations can also be configured with the ATI FireGL V3400 (128 MB) graphics accelerator, which features two display outputs, including one Dual Link output to drive ultra-high resolutions on a single 30-inch Dell UltraSharp flat-panel monitor.

For more information on why ATI FireGL is a smart choice for your next Dell Precision workstation, visit **www.dell.com/precision**.

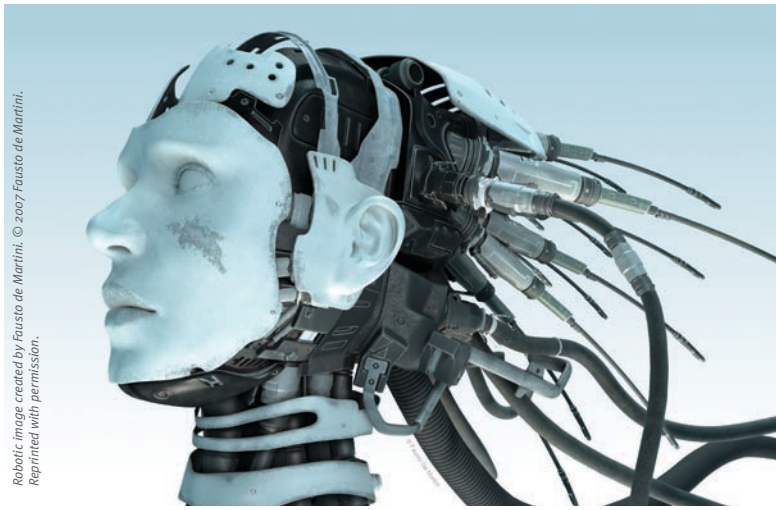

Robotic image created by Fausto de Martini. © 2007 Fausto de Martini. Reprinted with permission.

ATI FireGL delivers high-end rendering for 3D artists and animators

[1] For a list of ATI FireGL certifications, visit ati.amd.com/products/workstation/ISVCertsFireGL.pdf.

Related Categories:

Change management

Dell PowerEdge servers

Dell PowerVault storage

Firmware updates

Storage management

Visit www.dell.com/powersolutions for the complete category index.

MAINTAINING DELL PLATFORMS WITH DELL TECHNICAL UPDATES

Dell periodically releases updates for its servers, storage, and other enterprise platforms. These updates are classified as urgent, recommended, or optional depending on the enhancements they provide. Dell recommends customers keep their system firmware and software up-to-date to take advantage of these enhancements, which are designed to improve system functionality and minimize potential problems. The following tables summarize recent firmware updates for Dell™ PowerEdge™ components and Dell PowerVault™ storage. To sign up for Dell Technical Updates, visit support.dell.com/support/notifications/technicalupdates.aspx.

Firmware updates for Dell PowerEdge Expandable RAID Controllers (PERCs), PowerVault direct attach storage, and hard disk drives

Product		PERC 3	PERC 4/DC, PERC 4/SC	PERC 4e/DC, PERC 4e/Di, PERC 4e/Si	PERC 4/Di	PowerVault 220S		Maxtor Atlas 10KV (Genesis)
Firmware version		199A	352B	522A	422A, 252A*	E.18	E.19	BP005
Firmware release date		September 11, 2006	September 11, 2006	September 11, 2006	September 11, 2006	October 12, 2005	April 18, 2006	December 8, 2006
Criticality		Urgent	Urgent	Urgent	Urgent	Urgent	Recommended	Urgent
Enhancement	Fixes problems with EVPD inquiry commands	■	■	■	■			
	Fixes write cache policy reporting	■	■	■	■			
	Fixes problem with cluster mode setting for multiple PERC 4/DC cards on a PowerEdge 6650		■					
	Reduces loss of storage access or communications					■		
	Minimizes SCSI resets, SCSI Enclosure Services time-outs, and MegaRAID (MRAID) errors					■		
	Improves SCSI parity						■	
	Provides Self-Monitoring, Analysis, and Reporting Technology (SMART) and minimizes time-outs							■

Firmware updates for Dell PowerVault tape drives

Product		PowerVault 110T LTO-2-L Certance	PowerVault 110T LTO-2	PowerVault 110T LTO-3	PowerVault 110T DLT VS160	PowerVault 100T DAT72	PowerVault 110T SDLT 320
Firmware version		1914, A00	53Y3, A04	5BG2, A01	3200, A13	A16E, A09	5D5D, A15
Firmware release date		December 18, 2006	April 17, 2006	April 18, 2006	November 20, 2006	November 16, 2005	April 10, 2006
Criticality		Urgent	Recommended	Recommended	Urgent	Recommended	Recommended
Enhancement	Improves load/unload operation	■	■	■	■	■	■
	Improves head cleaning	■				■	
	Improves error recovery	■	■	■	■	■	
	Reduces read/write errors	■	■	■	■	■	
	Reduces media failures	■	■	■	■	■	■
	Reduces eject/insert errors	■	■	■	■	■	■

Firmware updates for Dell PowerVault tape autoloaders and libraries

Product		PowerVault 122T LTO-2	PowerVault 124T LTO-2-L, LTO-3, and DLT VS160	PowerVault 132T and PowerVault 136T LTO-2	PowerVault 132T and PowerVault 136T SDLT 320
Firmware version		53Y3, A03	Loader V31	53Y3, A04	5D5D
Firmware release date		October 7, 2005	June 19, 2006	June 22, 2005	April 10, 2006
Criticality		Recommended	Recommended	Recommended	Recommended
Enhancement	Improves load/unload operation	■	■	■	■
	Improves head cleaning		■		
	Improves error recovery	■	■	■	
	Reduces read/write errors	■	■	■	
	Reduces media failures	■	■	■	■
	Reduces eject/insert errors	■	■	■	■
	Fixes move issues		■	■	
	Fixes communication issues		■	■	
	Fixes picker issues		■	■	

*Firmware update 422A applies to PowerEdge 1750 servers, and 252A applies to PowerEdge 2600 servers.

By Challenging Traditional Thinking, Intel Accelerates I/O with New Platform Technology

By Indumathi Madhavan

Despite dramatic advances in network bandwidth and processor speeds, expected increases in server performance have not materialized. It turns out that the limiting factor is a set of network I/O bottlenecks. By addressing the primary causes of I/O bottlenecks, Intel® I/O Acceleration Technology streamlines the movement of network data, helping to accelerate network performance and improve scalability.

Network performance is a constant challenge for IT managers and administrators, given exponentially increasing amounts of data combined with growing user and application bandwidth requirements and the prevalence of processor- and I/O-intensive environments such as video and backup and restore operations. In response, the industry has moved from 10 Mbps Ethernet to 100 Mbps, then to Gigabit[1] Ethernet, and is transitioning to 10 Gigabit Ethernet. Correspondingly, processors have reached 3 GHz and increased performance.

Yet, with the speed of today's processors and the prevalence of Gigabit Ethernet, why are many server applications not benefiting from increased bandwidth? This is the question that drove research and development teams at Intel to develop Intel® I/O Acceleration Technology (Intel® I/OAT)—a comprehensive collection of platform-based technologies developed to move data efficiently across Quad-Core and Dual-Core Intel® Xeon® processor–based servers.

Eye-opening research leads to discovery

In approaching the I/O bottleneck problem, Intel first conducted exhaustive research and testing, starting with the prevalent idea that protocol processing is the root cause of I/O bottlenecks. The results confirmed what many had suspected: data movement is the real culprit—often causing the longest delays within a system.

[1] This term does not connote an actual operating speed of 1 Gbps. For high-speed transmission, connection to a Gigabit Ethernet server and network infrastructure is required.

Case Study
Dual-Core Intel® Xeon® Processor–based Servers and Intel® PRO/1000 Multi-Port Server Adapters

Intel® I/O Acceleration Technology Speeds Up Performance

Challenges

- Back up several petabytes of data weekly from Intel's large server and storage infrastructure
- Keep pace with growing business needs and database sizes by resolving performance bottlenecks in the enterprise backup and recovery system
- Scale network backup throughput and storage capacity without adding costly server hardware

Solutions

- Dual-Core Intel® Xeon® processor 5160–based servers
- Intel® multi-port network adapters
- Intel® I/O Acceleration Technology
- Veritas NetBackup* from Symantec

Benefits

- Enhanced the scalability of the backup and recovery system
- Reduced the number of servers required in the backup and recovery environment, thereby helping decrease total cost of ownership

Intel Global IT Group Solves Performance Challenges Using Intel® I/O Acceleration Technology

The Intel Global IT group designs, deploys, and supports the global compute infrastructure and applications needed to run Intel's core business. One of these crucial applications—enterprise backup and restore (EBaR)—backs up the company's critical data for rapid recovery in the event of system data loss or site-wide disaster. The EBaR system uses Veritas NetBackup* from Symantec and Gigabit Ethernet to pull data from the enterprise servers to Veritas NetBackup Media Server*, which then leverages Fibre Channel networks to write the data to tape.

Over time, as the amount of data grew and the speed and performance of tape drives increased, the original EBaR system could not keep pace. The system became constrained by performance bottlenecks, limiting the maximum throughput per server to 125 MB/sec.

Previous attempts at solving performance challenge failed

In the original EBaR design, each NetBackup media server had a single Gigabit Ethernet port connected to a dedicated backup network switch. In optimal conditions, this switch could handle the throughput to four Quantum Super DLT* tape drives. When new higher-speed tape drives were introduced, the network performance bottleneck started negatively affecting performance. After determining that the bottleneck was not due to software limitations of NetBackup, Intel IT needed to speed up server I/O.

To keep up with demand, servers were added to scale the data backup system—a costly solution. A potential alternative that Intel IT explored was *teaming*, or aggregating multiple ports of Gigabit Ethernet network interface cards (NICs). However, teaming did not solve the bottleneck problem, because network traffic was serviced by only one processor rather than the load being shared by the other processors in the system. In fact, the high processor overhead degraded performance.

continued from previous page

Intel research teams found that each time the processor sends a read or write request to memory, it typically waits through hundreds of idle cycles for memory to finish the slower-speed memory access, or *fetch,* operation before the processor can continue. And because the data must move back and forth between system memory, processor cache, and network interface card buffers, the delay affects not just the processor, but the entire server platform.

Based on these findings, Intel researchers determined that a platform-wide, multi-component approach would be required to alleviate network I/O bottlenecks and help improve server performance. The result is Intel I/OAT, which takes advantage of innovative changes in the processor, chipset, motherboard, LAN silicon, and software that make up a server platform to help ensure that data moves to and from applications efficiently and reliably. Based on Intel® QuickData Technology, Intel I/OAT includes the following:

- An optimized protocol stack designed to reduce protocol processing cycles
- Header splitting to allow packet headers and payloads from multiple packets to be processed on parallel paths
- Interrupt moderation to help prevent excessive interrupts
- Asynchronous, low-cost data copy to help reduce the number of processor cycles consumed while waiting for a memory access to finish

Working with Intel's networking product experts, the Intel IT team learned the root causes behind the system's I/O performance bottlenecks: lack of scaling across processors, TCP/IP processing inefficiencies, and multiple memory copies. Dual-Core Intel® Xeon® processor 5100 series–based servers featuring Intel® I/O Acceleration Technology (Intel® I/OAT) could potentially solve the network performance barrier. These systems can balance the network load across multiple cores and processors, enabling NIC port teaming to provide much greater throughput than was previously possible. In addition, Intel I/OAT works in conjunction with improvements to the Microsoft Windows* and Linux* OS network stacks to process TCP/IP network traffic quickly and help reduce system overhead by moving data efficiently through the server platform.

The Intel IT team tested the new servers with the EBaR application, and the test was a resounding success. The new platform with Dual-Core Intel Xeon processor 5160–based servers broke through the performance barrier, achieving more than 400 MB/sec with NetBackup. Test results also demonstrated enhanced scalability and processor utilization thanks to network interface teaming. Combining the Dual-Core Intel Xeon processor 5160–based servers with Intel I/OAT, the Intel IT group can now team up to four Gigabit Ethernet network adapters and retain 400 MB/sec or greater backup throughput to a single machine.

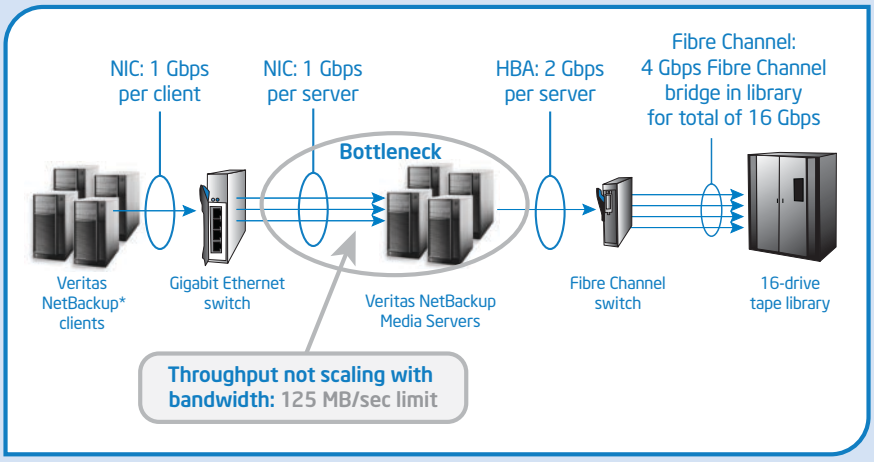

Throughput bottleneck in typical enterprise backup and restore configuration

Intel technology enables server consolidation and helps lower TCO

With the ability to improve performance and scalability using the Dual-Core Intel Xeon processor 5160–based server platform, Intel® multi-port network server adapters, Intel I/OAT, and the teaming features of the network adapters, Intel IT can significantly reduce the number of servers required for the EBaR system, which in turn helps reduce total cost of ownership for Intel data centers.

Aggregating several Gigabit Ethernet interfaces into a single virtual interface helps simplify the environment, and the performance optimizations in Intel I/OAT allow Intel IT to scale up the backup architecture without continually adding servers. With this performance breakthrough, backup windows are expected to shrink and fewer servers are necessary to support the data protection application. By leveraging the data movement efficiencies of these Intel technologies, the Intel Global IT group can cost-effectively scale the EBaR system as business needs and data sizes continue to expand.

Efficient data movement helps increase server performance and scalability

Ideal for I/O-intensive applications such as storage, data warehousing, data mining, and image access, Intel I/OAT is designed to improve network application responsiveness without requiring modification to applications.

In recent Intel Labs tests, Intel I/OAT enabled a significant increase in server performance on systems running the Red Hat Enterprise Linux 4* OS and the Microsoft Windows Server 2003* OS with Service Pack 1 (SP1) and the Microsoft Scalable Networking Pack (SNP). In tests running on these operating systems, Intel Xeon processor–based servers with Intel I/OAT demonstrated up to 50 percent less processor load and up to 2.5 times higher throughput compared with a previous-generation Intel Xeon processor–based server.[2]

Intel I/OAT enables fast, scalable network performance

Now a standard feature at no additional cost on server platforms based on Quad-Core and Dual-Core Intel Xeon processors that include Intel® LAN on Motherboard connections or Intel® PRO/1000 Server Adapters for PCI Express*, Intel I/OAT offers cost-effective scalability across up to six Gigabit Ethernet ports. At the same time, it exhibits power and thermal characteristics similar to those of a standard Gigabit Ethernet network adapter.

Quad-Core and Dual-Core Intel® Xeon® Processors Boost Server Performance

Ninth-generation Dell PowerEdge servers support Intel® I/OAT*

Building on a 20-year track record, Quad-Core and Dual-Core Intel® Xeon® processors help businesses stay competitive by providing enterprise-class performance, quality, and scalability. Ideal for compute-intense environments, 32-bit and 64-bit business-critical applications, and high-end workstations, the Dual-Core Intel Xeon processor 5100 series is designed to deliver up to three times the performance[3] and over three times the performance per watt[4] of previous-generation single-core Intel Xeon processors. The Quad-Core Intel Xeon processor 5300 series can provide up to another 50 percent increase in performance compared to the Dual-Core Intel Xeon processor 5100 series using roughly the same amount of power.[5]

Platforms based on the Quad-Core Intel Xeon processor 5300 series and Dual-Core Intel Xeon processor 5100 series support many Intel advanced server technologies designed to enhance operations, reduce costs, and improve business continuity. One of these technologies is Intel® I/O Acceleration Technology (Intel® I/OAT), which optimizes TCP/IP processing and helps reduce system overhead. Based on Intel® QuickData Technology, Intel I/OAT leverages architectural enhancements within the Intel Xeon processor, chipset, network controller, and firmware to minimize performance-limiting bottlenecks—helping to increase I/O throughput by freeing the host processor of network I/O–related load.

Intel I/OAT is now available as a standard feature on ninth-generation Dell PowerEdge* servers that include Quad-Core and Dual-Core Intel Xeon processors. To take advantage of the networking acceleration features of Intel I/OAT on these systems, an Intel® PRO/1000 Server Adapter for PCI Express* is required.

continued from previous page

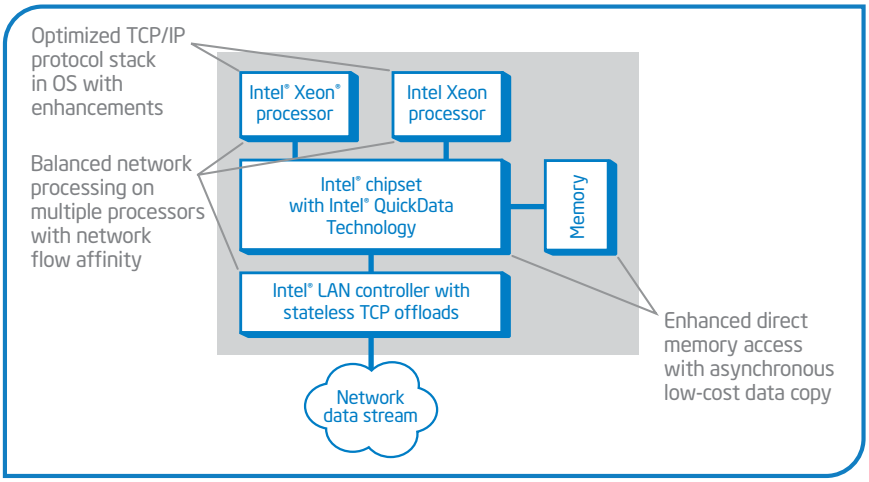

Server with Intel® I/O Acceleration Technology performance enhancements

Intel I/OAT is also tightly integrated with popular operating systems such as Microsoft Windows Server 2003 and Linux* and preserves critical network configurations by maintaining control of the network stack execution within the processor.

Intel I/OAT helps servers begin to take advantage of today's Gigabit Ethernet infrastructures and high-performance processors for fast, scalable networking. Intel I/OAT is available on ninth-generation Dell PowerEdge* servers starting in April 2007.

For more information

Intel I/O Acceleration Technology:
www.intel.com/go/ioat

Gigabit solutions from Intel and Dell:
www.intel.com/go/dellgig3

Indumathi Madhavan is a product marketing engineer in the Intel LAN Access Division, where she is responsible for Ethernet components and technology. Indu has more than 7 years of experience in the communications and networking industry, focusing mainly on Ethernet and associated technologies. She is currently supporting the development of Intel I/O Acceleration Technology and future Intel server platform technologies.

[2] Performance results for both processor load and throughput are based on an Ixia IxChariot 6.0 benchmark test performed by Intel Labs in January 2007 on a Dell PowerEdge 2900 server featuring Intel I/OAT and configured with two Dual-Core Intel Xeon processors (Woodcrest) at 3.0 GHz, 4 GB of 667 MHz RAM, three Intel® PRO/1000 PT Dual Port Server Adapters, either Red Hat Enterprise Linux 4 (kernel 2.6.18) or Windows Server 2003 with SP1 and the SNP, and six Dell PowerEdge 750 servers per port under test as clients, each configured with an Intel® Pentium® 4 processor at 3.4 GHz, Windows Server 2003 with SP1, and an Intel® PRO/1000 Gigabit Ethernet Server Adapter. The performance baseline was established in February 2006 using a previous-generation Dell PowerEdge 2800 server configured with two Intel Xeon processors at 3.6 GHz, 4 GB of RAM, three Intel PRO/1000 PT Dual Port Server Adapters, and either Red Hat Enterprise Linux 4 Update 3 or Windows Server 2003 with SP1. Actual performance will vary based on configuration, usage, and manufacturing variability.

[3] "Server Benchmarks: Integer Throughput," by Intel Corporation, *Intel Xeon Processor Server Performance*, www.intel.com/performance/server/xeon/intthru.htm.

[4,5] "Performance per Watt," by Intel Corporation, *Intel Xeon Processor Server Performance*, www.intel.com/performance/server/xeon/ppw.htm.

FROM DESKTOP TO DATACENTER.

FAST. SECURE. STREAMLINED.

DELL ANTICIPATED YOUR NEEDS.

Secure Communication Drives Business Success

Microsoft® Exchange™ Server 2007 running on Dell PowerEdge™ servers delivers a comprehensive, completely integrated end-to-end unified messaging solution – from the user to the datacenter.

Dell brings together aspects of a messaging environment – networking, servers, storage, services, tools and support – in one easy-to-use, cost-effective solution to reduce risk and increase performance and end-user productivity.

Rigorously tested and validated, it can reduce the costs and complexity of managing messaging environments. And since it is from Dell, you get scalable solutions built on standards-based products, and confidence in moving your business forward.

Microsoft® Exchange Server 2007

For more information, contact your sales representative or visit www.dell.com/exchange2007

DELL™