DELL POWER SOLUTIONS

TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

NOVEMBER 2006 • $4.95

Securing the Enterprise

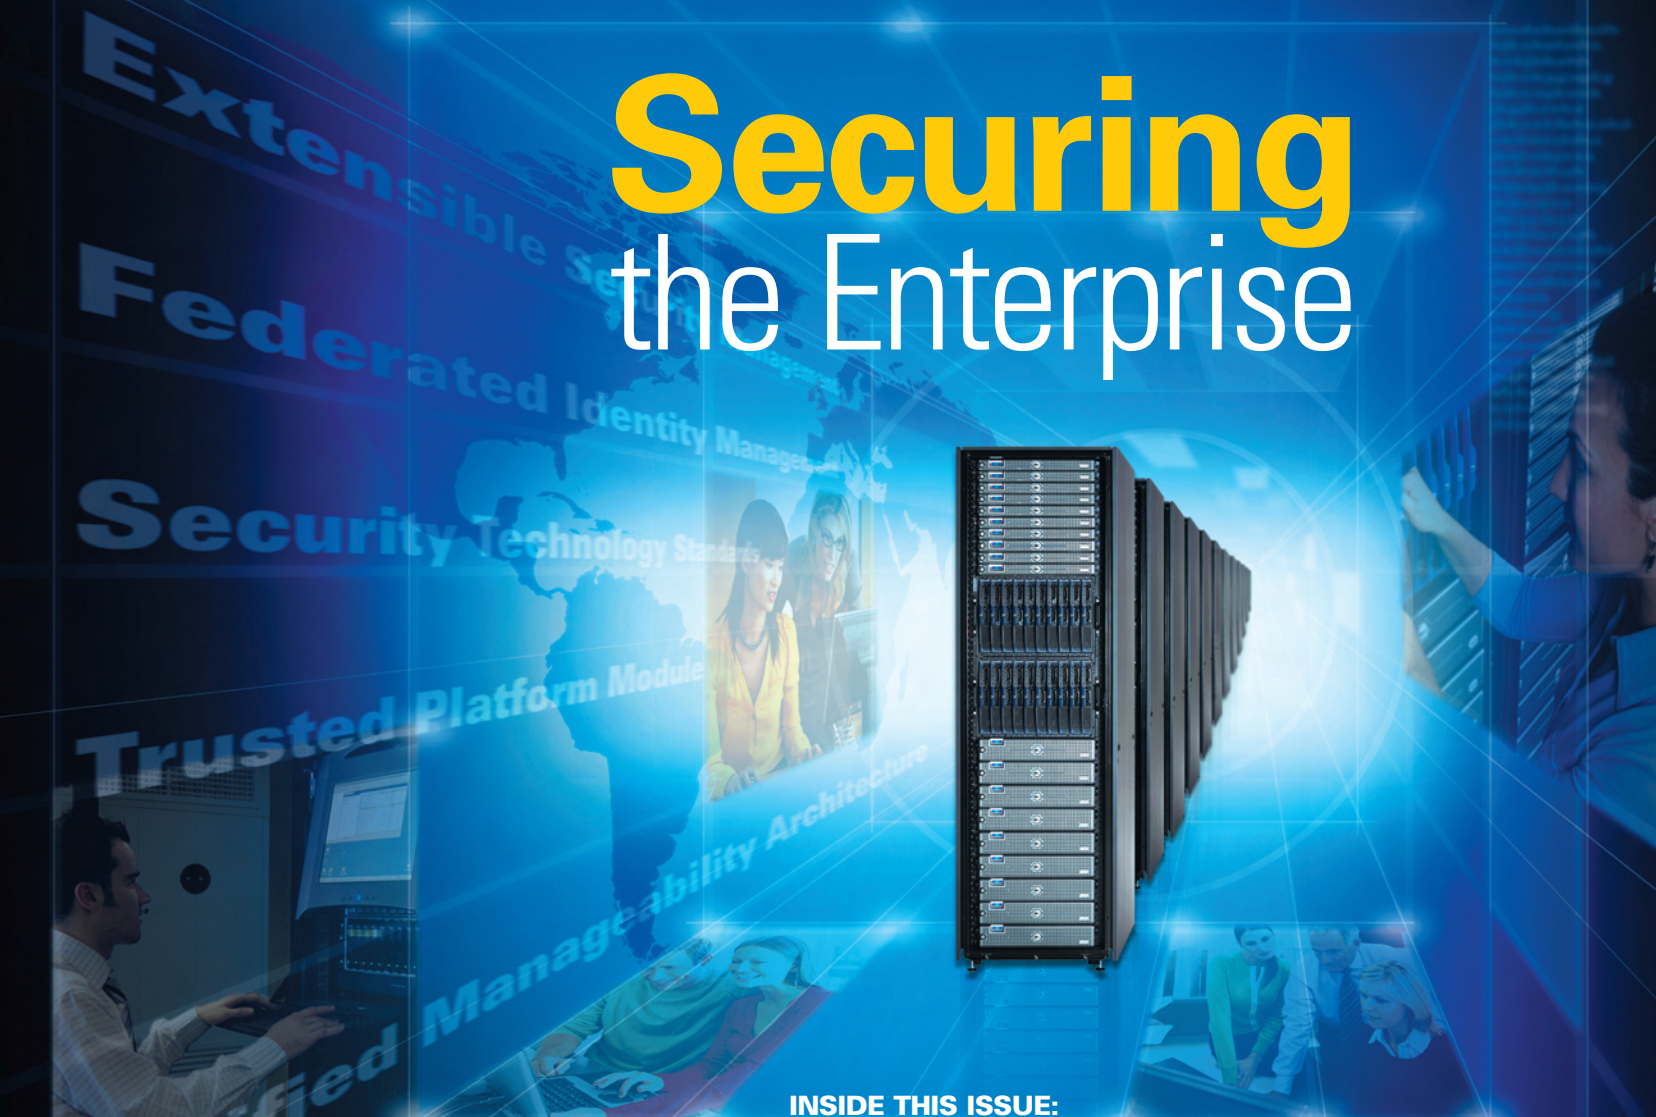

INSIDE THIS ISSUE:

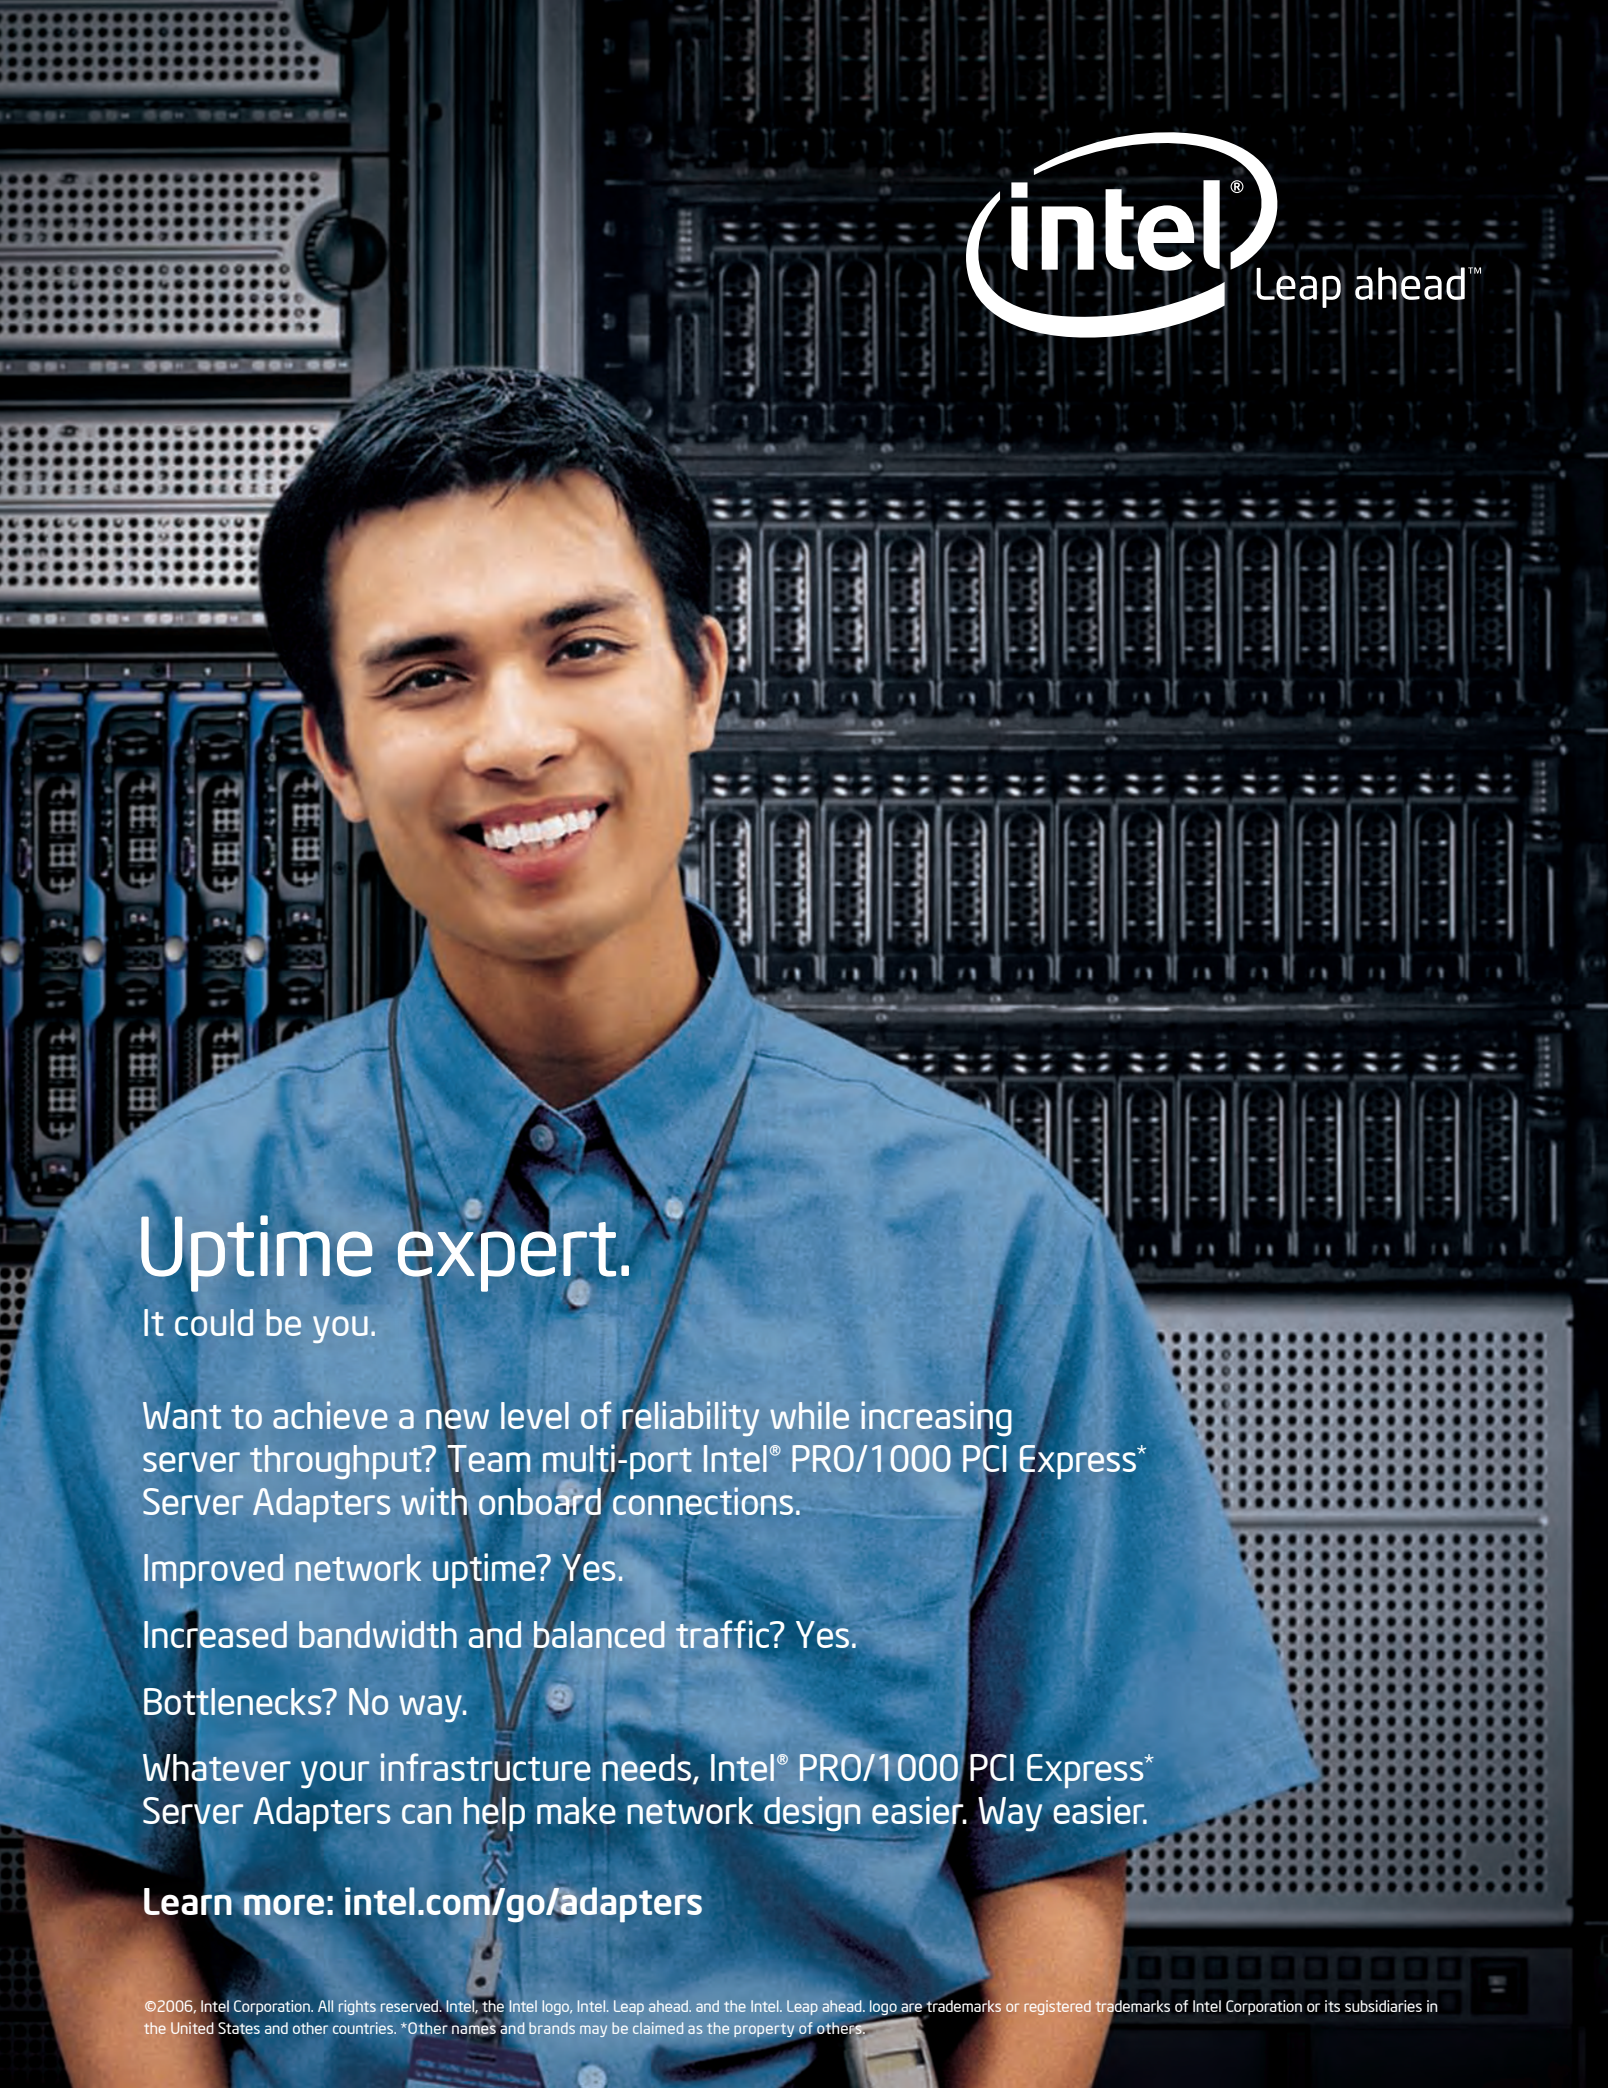
intel
Leap ahead™
Uptime expert.
It could be you.
Want to achieve a new level of reliability while increasing server throughput? Team multi-port Intel® PRO/1000 PCI Express* Server Adapters with onboard connections.
Improved network uptime? Yes.
Increased bandwidth and balanced traffic? Yes.
Bottlenecks? No way.
Whatever your infrastructure needs, Intel® PRO/1000 PCI Express* Server Adapters can help make network design easier. Way easier.
Learn more: intel.com/go/adapters
©2006, 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.

POWER SOLUTIONS

TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

November 2006

COVER STORY | PAGE 8

Securing the Enterprise

By Frank Molsberry

The Dell scalable enterprise strategy enables integrated, end-to-end protection using industry-standard data center components and a unified management framework that is designed to be inherently secure from the ground up.

NOVEMBER ISSUE | PAGE 89

Dell OpenManage Newsletter

Dell OpenManage IT Assistant 8.0 and the Dell Unified Manageability Architecture extend standards-based management for heterogeneous hardware and device instrumentation to help increase choice, reduce complexity, and enhance interoperability.

DELL POWER SOLUTIONS TRANSFORMING YOUR IT FRAMEWORK INTO A SCALABLE ENTERPRISE

WWW.DELL.COM/POWERSOLUTIONS

EDITORIAL
EDITOR-IN-CHIEF | Tom Kolnowski
MANAGING EDITOR | Debra McDonald
FEATURES EDITOR | Kathryn White
ASSOCIATE MANAGING EDITOR | Liza Graffeo
SENIOR EDITOR | Jim Duncan
EDITORIAL ASSISTANT | Amy Hargraves
CONTRIBUTING AUTHORS | Tim Abels; Barry L. Ader; Rizwan Ali; Andrew Bachler; Brian Berger; Ken Bignell; Sumner Blount; Sriranjan Bose; Bill Bryce; Winston Bumpus; Sajal Dam; Donnie Davis; Lauren Duda; Yung-Chin Fang; David S. Frankel; Subbu Ganesan; Richard K. Golasky; Aziz Gulbeden; Baris Guler; Rinku Gupta; Kelly Harriman-Polanski; Munira Hussain; Saeed Iqbal, Ph.D.; Dave Jaffe, Ph.D.; Suresh John; John L. Jones; Kevin Kline; Garima Kochhar; Jacob Liberman; Tong Liu; Victor Mashayekhi, Ph.D.; Jon Mercado; Rick Merino; Frank Molsberry; Todd Muirhead; Mark Norwalk; Patrick Ortiz; Ron Pepper; Craig Phelps; Ramesh Radhakrishnan, Ph.D.; Arun Rajan; Edward Reynolds; Vishvesh Sahasrabudhe; Amina Saify; Chad Sakac; Suman Kumar Singh; Brad Steckline; Prathap Thathireddy; David Ulbrich; Bharath Vasudevan; David Weber; and Lee Weiner

ART
ART DIRECTOR | Iva Frank
DESIGNER AND ILLUSTRATOR | Cynthia Webb
COVER DESIGN | David Chan
CONTRIBUTING ARTIST | Dennis Chatham

ONLINE
WEB DESIGN | Natanya Anderson
WEB PRODUCTION | Brad Klenzendorf

MARKETING
SPECIAL INSERTS MANAGERS | Ruby Halipoto and Erin Stolle

SUBSCRIPTIONS
Subscriptions are complimentary to qualified readers who complete the online subscription form. To sign up as a new subscriber, renew an existing subscription, change your address, or cancel your subscription, access the online Subscription Center forms at **www.dell.com/powersolutions**. For other subscription services, please e-mail **us_power_solutions@dell.com**.

ABOUT DELL
Dell Inc., headquartered in Round Rock, Texas, near Austin, is the world's leading direct computer systems company. Dell is one of the fastest growing among all major computer systems companies worldwide, with approximately 75,000 employees around the globe. Dell uses the direct business model to sell its high-performance computer systems, workstations, and storage products to all types of enterprises. For more information, please visit our Web site at **www.dell.com**.

Dell Power Solutions is published quarterly by Dell Inc., *Dell Power Solutions,* One Dell Way, Mail Stop RR5-03, Round Rock, TX 78682, U.S.A. This publication is also available online at **www.dell.com/powersolutions**, with syndicated RSS feeds available via **www.dell.com/rss**. No part of this publication may be reprinted or otherwise reproduced without permission from the Editor-in-Chief. Dell does not provide any warranty as to the accuracy of any information provided through *Dell Power Solutions.* Opinions expressed in this magazine may not be those of Dell. The information in this publication is subject to change without notice. Any reliance by the end user on the information contained herein is at the end user's risk. Dell will not be liable for information in any way, including but not limited to its accuracy or completeness. Dell does not accept responsibility for the advertising content of the magazine or for any claims, actions, or losses arising therefrom. Goods, services, and/or advertisements within this publication other than those of Dell are not endorsed by or in any way connected with Dell Inc.

Printed on recycled paper containing 10 percent post-consumer waste.

November 2006

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 having to work over the weekend meant going in to the office?

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 local coffee shop, your data center can always be within your reach.**

Digital Availability. Easy as DELL

www.dell.com

ADVERTISER INDEX

ONLINE EXTRAS

See It Here First!

Check the *Dell Power Solutions* Web site for our late-breaking exclusives, how-to's, case studies, and tips you won't find anywhere else. Plus: These *Dell Power Solutions* articles are available only online at **www.dell.com/powersolutions**.

Integrating EMC MirrorView with VMware ESX Server for Business Continuity and Disaster Recovery

By Jacob Liberman and David Ulbrich

Site-level disaster recovery can be achieved by integrating VMware ESX Server on Dell PowerEdge servers with EMC MirrorView in a Dell/EMC storage area network.

Multiple Ways to Efficiently Monitor and Manage Dell PowerEdge Servers

By Prathap Thathireddy and Sriranjan Bose

The Dell OpenManage suite, Dell Remote Access Controllers, and baseboard management controllers can facilitate local and remote IT management.

Remote Performance Monitoring Using Dell OpenManage IT Assistant 8.0

By Prathap Thathireddy and Sriranjan Bose

The Dell OpenManage IT Assistant 8.0 performance monitoring feature is designed to increase the efficiency and flexibility of enterprise IT management.

Using the Virtual Router Redundancy Protocol for Gateway Redundancy

By John L. Jones

Dell PowerConnect 6024 Gigabit Ethernet routing switches support the Virtual Router Redundancy Protocol, which enables gateway redundancy on LAN segments.

Introducing the McDATA 4416 Fibre Channel Switch Module for Dell PowerEdge 1855 and PowerEdge 1955 Servers

By Richard K. Golasky

The McDATA 4416 switch module can support a storage area network for Dell PowerEdge blade servers, enhancing storage management.

Best Practices for Microsoft Windows Installation on Dell PowerEdge Servers with Broadcom NetXtreme Devices

By Ken Bignell, Subbu Ganesan, and Donnie Davis

This article discusses best practices for manual and automated Microsoft Windows OS installations involving Broadcom device drivers on Dell PowerEdge servers.

Subscribe to RSS feeds: www.dell.com/rss

Online Related Categories index: www.dell.com/powersolutions

Catch the wave in affordable, easy-to-use SANs.

SANbox® Express
1400 series

SANbox® 5600 FC
Stackable switches

QLE2460/62 HBAS

QLA®4050 TOE

QLOGIC®

QLogic has everything you need— switches, HBAs, routers, Storage Services Platform and software — so that everyone in your organization can ride a radical wave of productivity with Fibre Channel and iSCSI storage area networks. And the same reasons you choose Dell — great prices, blazing performance, easy installation, and the confidence that comes from going with the recognized leader — are also the reasons to choose QLogic. Catch the wave at **www.qlogic.com.**

Got CFD? Get Job, Get Published

What does the science of computational fluid dynamics (CFD) have to do with 500 engineering jobs in central Texas? That question was answered recently in a big way when Dell unveiled a 14-by-48-foot billboard along a busy highway in Austin adorned with a CFD flow diagram and the words "UNDERSTAND THIS? Dell is Hiring 500 Experienced Engineers in Texas." Already Dell's largest product development center, the Austin Design Center is not only recruiting experienced engineers that understand the complexity of CFD, but also seeking credentialed engineers across a wide variety of disciplines—including diverse roles in electrical, mechanical, quality, and software engineering.

Are you an expert in your engineering field? If so, visit www.dell.com/careers and let Dell know about it. As a member of the Austin Design Center team, you would join forces with engineers in the company's Bangalore, China, Taiwan, and Singapore Design Centers as Dell continues to broaden its product and services portfolio with the new Dell 2.0 initiatives. You may also have an opportunity to get published: much of the leading-edge technical content in *Dell Power Solutions* is authored by Dell engineering experts in the worldwide design centers.

Astute *Dell Power Solutions* readers may have a case of déjà vu upon viewing the billboard photo: a virtually identical CFD diagram was published as part of an article on power and cooling in our August 2005 issue.[1] And as you might expect, the article was authored by a Dell engineer from the Austin Design Center.

For the November 2006 issue of *Dell Power Solutions*, we collaborated with the Austin Design Center's office of the Dell CTO to develop our featured section on enterprise security. In our cover story on page 8, "Securing the Enterprise," the lead security technologist at Dell describes how a scalable, standards-based architecture can enable a strong security strategy for today and tomorrow. We also worked with the Trusted Computing Group to understand how it is helping drive global security standards (page 13).

And there is much more content for technical browsing in this issue: check out our special section on high-performance computing (beginning on page 38), the second installment of the Dell OpenManage Newsletter (page 89), and the inside scoop on how the Dell IT team uses virtualized server farms (page 65). We hope this issue helps keep you informed on ways to advance your IT department.

Tom Kolnowski
Editor-in-Chief
tom_kolnowski@dell.com
www.dell.com/powersolutions

[1] For more information, see "Guidelines for Assessing Power and Cooling Requirements in the Data Center," by David Moss, *Dell Power Solutions,* August 2005, www.dell.com/downloads/global/power/ps3q05-20050115-Moss.pdf.

It's like a big "undo" button for your entire business.

A new virus. A careless user. A software installation gone wrong. You never know what's going to cause a major system problem. But with Symantec Backup Exec™ System Recovery (formerly LiveState Recovery), you'll know exactly how to undo it. Its disk-to-disk recovery technology lets you quickly restore Windows® systems anytime, from anywhere to virtually any device. You can bring your system back to its pre-problem state in minutes instead of hours, even to dissimilar hardware or to virtual environments. This means significantly less downtime for your company. To learn how a better system recovery solution can keep your business moving in the right direction, visit www.backupexec.com or contact your Symantec Certified Partner.

BE FEARLESS.

symantec™

Securing the

Reactionary security add-ons that have been implemented piecemeal are no match for increasingly sophisticated cyberthreats, let alone today's complex regulatory requirements. The Dell scalable enterprise strategy enables integrated, end-to-end protection using industry-standard data center components and a unified management framework that is designed to be inherently secure from the ground up.

Related Categories:

Data security

Enterprise security

Identity management

Regulatory compliance

Scalable enterprise

Security

Security software

Threat management

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

BY FRANK MOLSBERRY

In today's globally internetworked enterprise, just one nasty e-mail virus or one database breach has the power to launch a maelstrom of cleanup, recovery, and damage control—not to mention the furor that erupts over lost business, damage to reputation, and possible litigation. Afraid to risk negative publicity that could damage business even further than the intrusion itself, many enterprises are reluctant to report security breaches to law enforcement agencies. It is small wonder that high-stakes cybercrime and the mushrooming regulatory environment are driving companies to reach deep into the corporate coffers to protect vital business interests.

For example, a joint study from the Computer Security Institute (CSI) and the Federal Bureau of Investigation (FBI) reported that in 2006, firms with annual sales under US$10 million spent an average of US$1,349 per employee on computer security—a 210 percent increase over the average per-employee expenditure in 2005.[1] While large companies are able to achieve economies of scale, the same CSI/FBI study reported that organizations with annual sales between US$10 million and US$99 million

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

increased per-employee computer security expenditures by 327 percent compared to the previous year—to an average of US$461 per employee in 2006.

Unfortunately, the plethora of security technologies and options available today has created an extremely confusing situation for IT organizations charged with securing the enterprise. But this much is clear: reactionary security add-ons that have been implemented piecemeal are no match for increasingly sophisticated cyberthreats, let alone today's complex regulatory environment. Dell is driving an open, standards-based foundation for end-to-end protection that is highly scalable. In this approach, the IT infrastructure is designed to be inherently secure as it is deployed and scaled to meet evolving business requirements.

Aligning security with business needs

The sheer number of current security offerings is staggering. In addition to physical security measures such as asset tags, security cameras, and fingerprint access mechanisms, methods to protect data and networks range from passwords and smart cards to firewall and antivirus software, and from antispam software to anomaly detection systems.

The problem is that these are typically ad hoc building blocks that have been tacked onto systems as stopgap measures. Every day, organizations are being bombarded with information about how to respond to the latest threat, and they rarely have the opportunity to step back and see the big picture—to develop sensible policies for protecting business assets and to find cost-effective solutions for meeting specific business needs.

The fragmented marketplace is especially challenging for small businesses. Many simply do not have the expertise, resources, or time to manage a patchwork of temporary add-ons that routinely require integration, updating, and support. Yet these organizations are still affected by regulatory requirements and the threat du jour, and still must contend with competitive factors that affect how they run their businesses.

Meanwhile, corporate scandals and widely publicized breaches in data privacy have shaken consumer confidence on many fronts,

Checklist: Proactive security measures

- ✔ **Keep up with emerging government regulations.** Even if they do not apply to a business today, they set the tone for what may be coming.
- ✔ **Evaluate standards compliance of security solutions.** Consider emerging security and management standards from organizations like the Distributed Management Task Force, OASIS (Organization for the Advancement of Structured Information Standards), and the Trusted Computing Group.
- ✔ **Make security part of the design and development process.** As OS and application providers evolve their development and update processes to enhance security solutions, the weakest link may be internally developed applications. Every internally written and deployed application should be scrutinized for potential vulnerabilities.
- ✔ **Remember the human element.** People are usually the weakest link in the chain. Security solutions should be simple enough that people can easily use them. And although technology can help set and manage corporate policies, effective education and strong enforcement are also key elements in a security program.

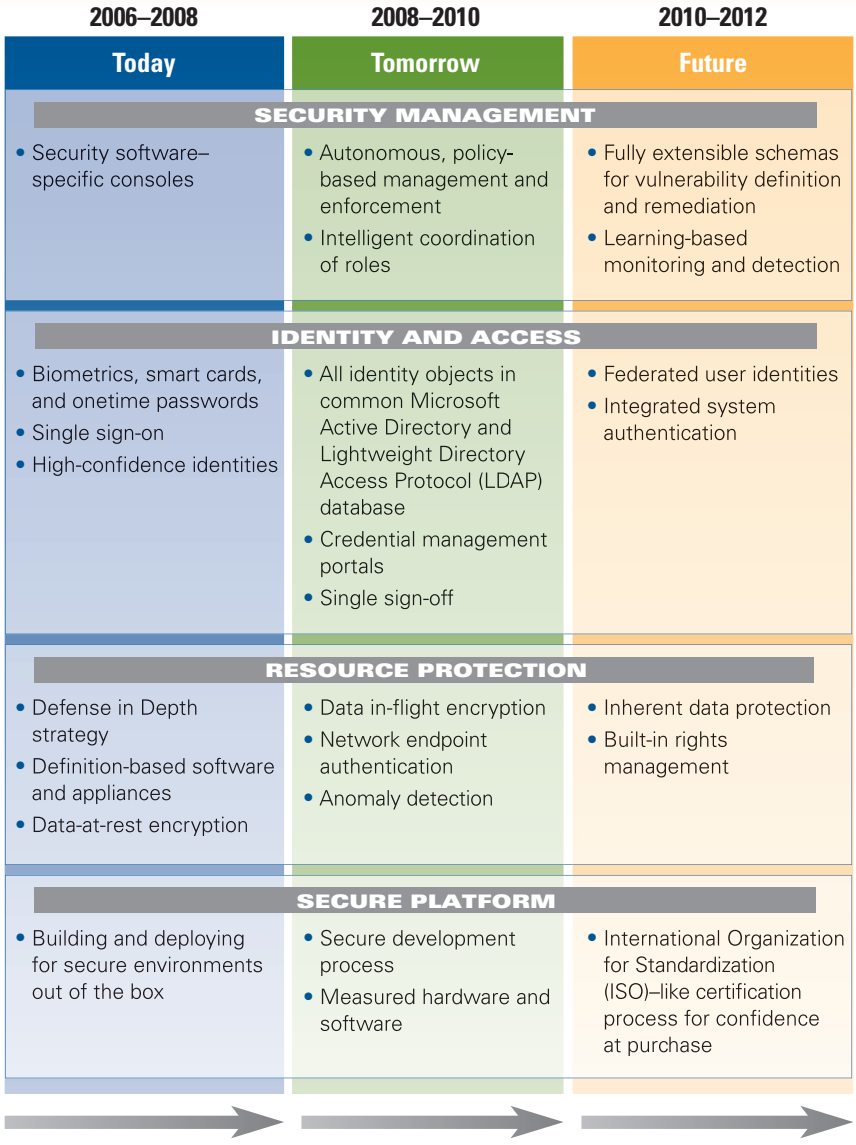

Figure 1. Planning the progression toward integrated, end-to-end security

especially in regard to online transactions and electronic data interchange. The question is, How can enterprises regain consumer confidence in e-commerce and secure electronic transactions with business partners and vendors while fortifying the all-important IT foundation on which the enterprise is built?

Building a high-level security framework

Dell believes the answers lie in the cooperation of the computer industry as a whole to categorically define and promote end-to-end security that addresses every aspect of the IT environment. Building a high-level framework around security helps enterprises not only to protect resources and identities, but also to provision and manage hardware and software with built-in security components simply and cost-effectively.

Moving forward, security must be inherent in every aspect of the IT development process, not an afterthought. Security components should be built into each phase of the life cycle—through design, manufacturing, deployment, provisioning, and retirement. Then, security becomes an integral part of the preferred enterprise management framework.

Furthermore, whether users are attached through a home network, a corporate network, or the Internet, the network and all endpoints—including clients, servers, and storage—must be able to self-monitor for security exposures. In the near future, these networks will likely be designed to avoid reactive user events by updating, isolating, or correcting themselves in an automated, policy-based fashion. Once environments can effectively police themselves, policies can be established to validate users, systems, and data as well as to define and enforce actions and events.

That is the vision IT should be painting today, but the path is not always clear. Dell believes a phased approach, based on existing and emerging standards, can provide the best possible route to the secure destination. Figure 1 shows Dell's pragmatic, phased approach to securing the scalable enterprise.

Protecting mobile endpoints

As more users access the Internet and workforces become increasingly mobile, devices such as notebooks and smartphones must be able to move on and off the network easily while remaining secure at all times. When a system reaches the end of its life, a user changes roles, or information becomes outdated, assets must be removed without risking the release of any personal information or confidential data.

Dell is calling on the industry to design all hardware components and software with a comprehensive, integrated level of protection—throughout the hardware stack, the middleware, the OS, and the applications. Today these elements exist in silos with multiple security components tied to each, including network appliances, antivirus, antispam, and identity management software as well as access controls, auditing forensics for regulatory compliance, and offline storage with additional encryption and life cycle management.

To achieve fully integrated protection, the industry must move toward a standards-based management framework that facilitates all security elements, including information life cycle management; global policy setting; integration with directory services such as Microsoft® Active Directory® and Novell® eDirectory™ directory services; credential management; and federated identities, which allow trusted credentials to be shared between organizations, companies, or Web sites. And this framework must provide automated ways of updating information, generating alerts, and establishing autonomous control.

Facilitating the flow of information

A major step toward the goal of fully integrated protection is to put all of the object information, the database information, and the information flow from solution silos into a common format in a federated repository, enabling communication among them. Then, each solution silo can plug into an overall enterprise management framework where administrators can apply business policy rules and service-level agreements to define actions and security policies.

Because business transactions today rely so heavily on third-party solutions and information, it is also important that OS middleware and application software be regularly updated. When a higher-level management framework is in overall control, these updates can be coordinated and scheduled to minimize user disruption. No one can tolerate gaping holes in the enterprise, which is why Dell continues to support low-level building-block architectures and standards for security elements such as biometrics and smart cards as well as standards for high-level management frameworks.

Dell is committed to developing a highly scalable, standards-based architecture and solutions for end-to-end security. As a result of Dell's work with middleware, software, and virtualization providers and strong alliances with partners who offer state-of-the-industry best practices, Dell™ hardware already integrates a variety of security features. For example, a Trusted Platform Module (TPM) is included in many Dell client systems for organizations that require security solutions with multifactor authentication or hardware-backed secure storage of digital keys, certificates, and passwords.

Laying the foundation with industry standards

Dell has assumed a leadership role in vendor-neutral standards organizations such as the Trusted Computing Group (TCG). TCG is developing specifications for trusted computing and security technologies that are designed to make security inherent in every aspect of the IT infrastructure as well as the overall management framework—including hardware components and software interface specifications across a range of platforms and operating environments.[2] For example, TCG specifications define, among other things, the standards for creating TPMs, which are microcontrollers incorporated into computing devices to provide hardware protection for security tasks and authentication information.

Dell is also working with industry standards bodies such as the Distributed Management Task Force to define the various characteristics of security objects so these objects can fit into large management frameworks. This effort includes creating common information flow formats for issues such as identity management and vulnerability definitions.

By defining and creating common information schemas, security objects and information can be shared among security solutions—allowing organizations the flexibility to choose the most appropriate application environment and management framework for their particular business needs. In addition, Dell is creating a federated repository model for flowing information outside the organization. This model, which also integrates into the management framework, can help ensure that electronic data interchange is protected.[3]

"We can't solve problems by using the same kind of thinking we used when we created them."

—Albert Einstein

The recently announced Dell Unified Manageability Architecture (UMA) is yet another example of Dell's commitment to help reduce security management complexity while helping ensure the best possible protection. Figure 2 depicts UMA as part of a broader security architecture that also includes other schemas and scalable enterprise elements. UMA is designed to help enable well-defined, widely accepted systems management standards that promote interoperability and flexibility in enterprise computing environments. The layered design of UMA has built-in security management so that security objects can be fully integrated into systems management tools.[4]

[2] For more information about TCG initiatives, see "Enhancing IT Security with Trusted Computing Group Standards," by Frank Molsberry and Brian Berger, in *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20070160-TCG.pdf.

[3] For more information, see "Dell Scalable Enterprise Architecture," by Jimmy Pike and Tim Abels, Dell Inc., August 2005, www.dell.com/downloads/global/vectors/2005_scalable_enterprise.pdf.

[4] For more information, see "Dell Unified Manageability Architecture: Blueprint for an Open Management Framework," by Winston Bumpus, in the Dell OpenManage Newsletter, *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20070141-OpenManageNews.pdf. To learn about the Dell UMA specification and implementation examples, visit www.dell.com/standards.

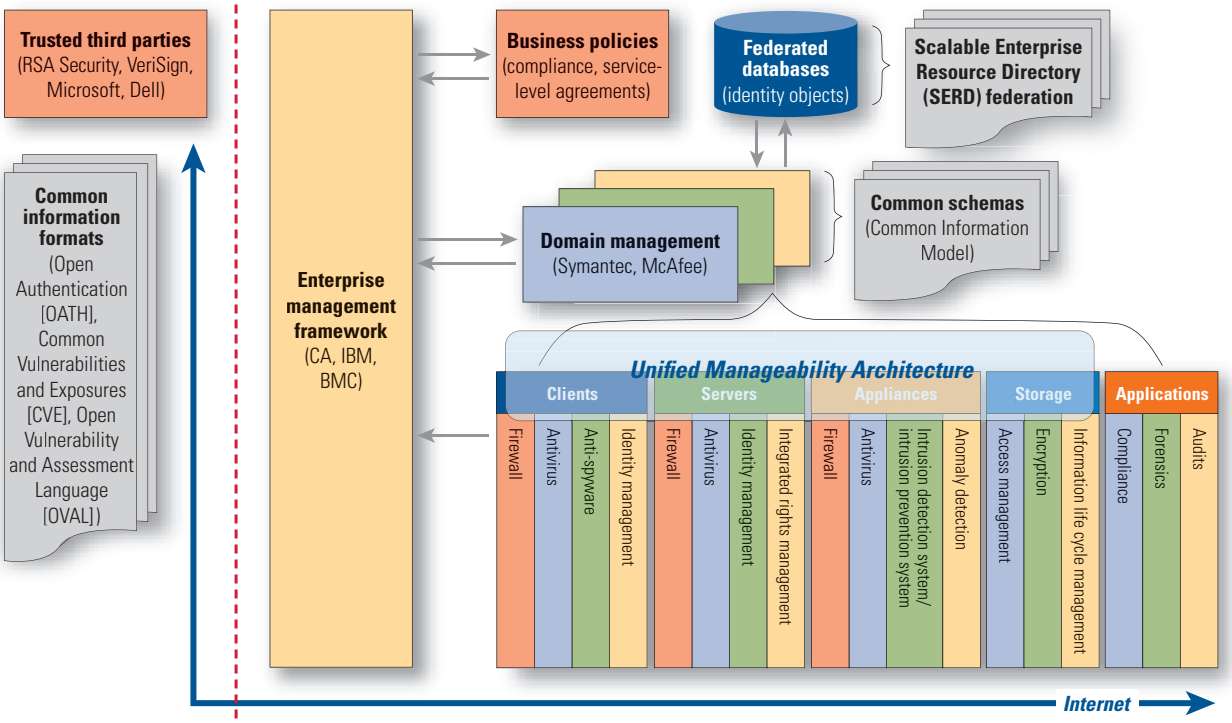

Figure 2. Creating a standards-based model enabling inherent security within a unified management framework

An additional example of how Dell is advancing security through standards is the Dell Secure Exchange Reference Architecture.[5] This architecture comprises industry-standard components that can help provide data protection and security for Microsoft Exchange messaging environments. Dell Secure Exchange solutions incorporate Dell hardware and Symantec security software to help support and protect the Exchange application platform.

Advancing toward self-policing, plug-and-play security

Stand-alone security functions are being phased out by a network-oriented approach in which security objects become part of overall management policies governing the IT infrastructure. To that end, Dell is playing an industry-leading role by actively participating in the development of security technology standards, building security-enabled platforms, and forging strategic partnerships to advance integrated, end-to-end security solutions.

At the end of the day, however, it is still important for administrators to remember that as the computer industry strives toward inherently secure solutions, other weak links in the chain cannot be overlooked. The human aspect—disgruntled employees, lack of security education, lack of security policies and enforcement, or poorly tested internal code and software—presents significant risks and is an equally important part of the security equation.

Most security experts agree that cybercrime is here to stay and the threats will only become more sophisticated and potentially more detrimental. However, Dell's initiative to push forward with standards-based architectures for managed security is designed to reduce the complexity and confusion of protecting against these threats.

Frank Molsberry is the lead security technologist in the office of the chief technology officer at Dell, and serves as the Dell representative to TCG. He has more than 20 years of experience in advanced systems software development and PC system architectures. Frank is a member of the Computer Security Institute and has a B.A. in Computer Science from the University of Texas at Austin.

FOR MORE INFORMATION

Dell security solutions:
www.dell.com/security

[5] For more information, see "Implementing the Dell Secure Exchange Reference Architecture," by Suman Kumar Singh and Bharath Vasudevan, *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20060452-Singh.pdf.

Enhancing IT Security

with Trusted Computing Group Standards

An increasingly interconnected global computing environment brings with it myriad threats for enterprises to guard against, including software attacks and theft of both data and physical devices. This article discusses Trusted Computing Group™ security standards that can help enable IT organizations to effectively respond to these challenges.

BY FRANK MOLSBERRY AND BRIAN BERGER

Related Categories:

Enterprise security

Security

Standards

Trusted Platform Module (TPM)

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

Concerns about the security of communications, transactions, and wireless networks—including problems such as data exposure, software attacks, identity theft, and even physical theft of mobile devices—can make it difficult to realize the potential benefits associated with pervasive connectivity and e-commerce. Standards-based security measures can address these issues by helping reduce security risks while also providing interoperability and protecting privacy.

The Trusted Computing Group (TCG) was formed in 2003 to respond to this challenge. TCG is a not-for-profit corporation with international membership and broad industry participation, including more than 135 members. The purpose of TCG is to develop, define, and promote open, vendor-neutral industry specifications for trusted computing and security technologies, including hardware building blocks and software interface specifications across multiple platforms and operating environments. Implementation of these specifications can help enterprises protect their information assets (data, passwords, certificates, keys, digital identities, credit card information, and so on) from software attacks and physical theft; provide mechanisms for proactively establishing trusted relationships for remote access through secure user authentication and computer authentication and attestation; and enable secure computing environments while avoiding compromises in functional integrity, privacy, and individual rights.

Protecting critical information

Products developed based on TCG specifications can help meet the challenges associated with software attacks from sophisticated and automated attack tools, increasing numbers of vulnerabilities, and widespread

user mobility. These problems can contribute to the risks of electronic theft of valuable personal or enterprise data—including identity or authentication information that can give hackers access to multiple systems and accounts, thereby compounding the potential damage from the attacks—and to the risks of physical theft of mobile user systems such as notebooks, which can provide another route to sensitive data.

Software-only security mechanisms may not be sufficient to protect information assets. Even firewalls protecting intranet environments can prove inadequate, especially when software attacks bypass the firewall (for example, through e-mail attachments) or originate from internal users. Hardware-based embedded security solutions are therefore an increasingly important element of secure environments. The goal of TCG is to make these protections available across a broad range of computing devices with common software interfaces to facilitate application development and interoperability.

Figure 1. Trusted Computing Group standards as part of a comprehensive security infrastructure

Defining TCG specifications

TCG provides hardware and software interface specifications along with white papers, marketing programs, and other materials to promote awareness, understanding, and adoption of these specifications. Key TCG policies related to specification development include the following:

- **Open platform development model:** TCG is committed to preserving the open development model that enables any party to develop hardware, software, or system platforms based on TCG specifications, and to preserving consumer freedom of choice.
- **Platform owner and user control:** TCG is committed to enabling owners and users of computing platforms to remain in control of their platform, and to requiring platform owners to opt in to enable TCG features.
- **Privacy capabilities:** TCG is committed to including capabilities for securing personally identifiable data in its specifications.

The primary TCG specifications rely on the Trusted Platform Module (TPM) hardware component, which is in widespread deployment, and the TCG Software Stack (TSS), which developers can use as a foundation for various applications. TCG has also released the Trusted Network Connect specifications for network security implementations, and has created additional work groups to develop security standards for storage, mobile devices, servers, infrastructures, and peripherals. Figure 1 illustrates the comprehensive nature of TCG standards as part of a security infrastructure.

Trusted Platform Module

The TPM specification defines the architecture and other standards for TPMs—microcontrollers designed to securely store digital keys, certificates, and passwords to help maintain data confidentiality. TPMs are typically affixed to PC motherboards, but can potentially be used in any computing device that requires these functions. They are designed to protect key operations and other security tasks that might otherwise be performed on unprotected interfaces in unprotected communications, and to protect platform and user authentication information and un-encrypted keys from software-based attacks. TPMs from various semiconductor vendors are included on enterprise desktop and notebook systems from Dell and other vendors.

TCG Software Stack

The TSS specifies a standard software interface to TPM functions that facilitates application development and interoperability across platform types. The TSS includes functions that developers can use to create interfaces for existing cryptography application programming interfaces (APIs), such as Microsoft® CryptoAPI, the Intel® Common Data Security Architecture, and the RSA Security Public-Key Cryptography Standard #11. In this way, the TSS helps enable TPM support for applications using these APIs. Application developers

can use the TSS to create interoperable client applications designed to improve tamper-resistant computing by taking advantage of TPM capabilities such as key backup, key migration, platform authentication, and attestation.

Trusted Network Connect

The TNC specifications define an open solution architecture designed to help network administrators protect networks by allowing them to audit endpoint configurations and impose enterprise security policies before establishing network connectivity. The TNC architecture builds on existing industry standards and defines new standards as necessary, with the objective of enabling nonproprietary, interoperable solutions within multivendor environments.

TNC provides a method of measuring and attesting to the characteristics of endpoint devices as they attempt to connect to a network. This method involves collecting endpoint configuration data and user authentication information for comparison with predefined organization access criteria—thereby helping create a security, or *safe computing,* profile for a system—and providing an appropriate level of network access based on the detected level of policy compliance, including full access, partial or directed access, or no access.

TCG work groups

To extend its specifications beyond PCs, TCG has created work groups to define implementation architectures for storage, mobile devices, servers, infrastructures, and peripherals. The Storage Work Group, for example, plans to build on existing TCG technologies and address standards for security services on dedicated storage systems, such as disk drives, removable media drives, flash storage, and multiple storage device systems. One objective is to develop standards and practices for defining the same security services across dedicated storage controller interfaces, including ATA, Serial ATA, SCSI, Internet SCSI (iSCSI), Fibre Channel, USB storage, IEEE 1394, and TCP/IP network attached storage. The Storage Work Group also acts as the TCG liaison to other industry groups that have jurisdiction over these storage interface standards to promote the adoption of TCG technology.

Implementing TCG specifications

When implemented in motherboards, desktop and notebook PCs, servers, and other computing systems, TCG specifications can help provide several important elements of a secure, integrated environment, including the following:

- Secure storage of files, personally identifiable information, and digital secrets, helping protect both data and identities from external software attacks or physical theft
- Strong multifactor user authentication through components such as security tokens, smart cards, passphrases, fingerprint readers, proximity badges, Subscriber Identity Modules, and so on
- Network access control in which an IT organization can control user access based on the organization's policies and security procedures, helping ensure that only secure client systems can access the network
- Exploitation of the latest OS features, such as the BitLocker feature of the Microsoft Windows Vista™ OS, which uses TPMs to measure boot process attributes and store keys for full-volume data encryption

Dell includes TPMs and Wave Systems EMBASSY Trust Suite software on many Dell™ Latitude™ notebooks, Dell OptiPlex™ desktops, and Dell Precision™ workstations. Dell also anticipates eventually incorporating TPM architectures on its servers and storage.

Evolving to meet ongoing security challenges

The open hardware building blocks and software interface specifications developed and promoted by TCG are designed to increase security and trust in computing platforms through hardware-based cryptographic functions, protected storage of user data, mechanisms for secure storage and platform integrity reporting, and platform authentication with multiple attestation identities. Organizations can prepare for a trusted enterprise model by deploying technologies that support TCG standards as they are developed. As threats from software attacks, theft, and other sources increase, TCG anticipates that trusted computing and security technologies will evolve to meet these threats, and plans to work with the IT industry to continue enhancing computing security.

Frank Molsberry is the lead security technologist in the office of the chief technology officer at Dell, and serves as the Dell representative to TCG. He has more than 20 years of experience in advanced systems software development and PC system architectures. Frank is a member of the Computer Security Institute and has a B.A. in Computer Science from the University of Texas at Austin.

Brian Berger is an executive vice president of marketing and sales at Wave Systems as well as a TCG director and chairman of marketing. Brian has a B.A. from California State University, Northridge, and attended the Harvard Business School Executive Education program.

FOR MORE INFORMATION

Trusted Computing Group:
www.trustedcomputinggroup.org

Protecting Enterprise Assets

with Identity Management Solutions from Dell and IdentiPHI

The strength of enterprise security measures can depend not only on how these measures are implemented, but also on how users interact with them. Dell and IdentiPHI have created a comprehensive identity management solution conforming to the guidelines and implementation standards used by federal government agencies. IdentiPHI™ software—including the Enterprise Security Suite, Advanced Authentication, and Single Sign-On—is designed to help organizations plan, deploy, and manage their identity management programs and protect critical information and other assets.

BY MARK NORWALK AND CRAIG PHELPS

Related Categories:

Enterprise security

IdentiPHI

Identity management

Regulatory compliance

Security

Security software

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

Steady rises in cyber-crime, widespread worms and viruses, the threat of cyber-terrorism, and regulatory requirements continue to increase enterprises' need for powerful security technologies such as antivirus and anti-spyware software, virtual private networks (VPNs), strong authentication, storage encryption, and secure e-mail. But system vulnerabilities do not necessarily arise from technology or a lack thereof; they can also be created by end-user behavior.

These types of vulnerabilities often fall under the categories of *identity management* and *user authentication*. In simplified terms, identity management is the ability to verify, within an acceptable level of risk, the identity of individuals accessing organizational resources, and the ability to appropriately manage necessary changes within the identity management infrastructure. Because identity management is a key component of several protection strategies, it has become an important element of enterprise security initiatives. Dell and IdentiPHI have been working together to create solutions designed to help enterprises implement comprehensive, secure identity management for protection from both internal and external vulnerabilities.

Assessing enterprise security goals

Typical goals of security-conscious enterprises include protecting assets, information, and employees; complying with regulatory and organizational requirements; and

establishing a common security strategy across the organizational structure. An enterprise might determine, for example, that a single user credential such as a smart card could allow building access; hold certificates for access to networks, remote VPNs, or encrypted hard drive information; digitally sign e-mails; and allow access to other information.

But even strong measures may provide little or no security if not implemented properly. Enterprises must combine security technologies—such as encryption algorithms, antivirus and anti-spyware software, firewalls, and user authentication—with solid organizational procedures and process baselines to help maximize the potential of these technologies. An enterprise with separate management and technology silos may have different groups responsible for access to physical sites, VPNs, networks, databases, and so on, all of which have different requirements, budgets, and goals. Consistent guidance from management is a necessary element of creating a secure, integrated environment.

The first step in aligning enterprise objectives with secure identity controls begins with an internal operational assessment. Operational assessments take into account business objectives, current security practices and products, and both short- and long-term goals to identify specific security needs and offer an implementation strategy to meet those needs over a defined time period. When both business and technology groups participate, this assessment process can help provide cohesion between management and IT groups along with enhanced security, user convenience, and system cost-effectiveness. Third-party assessments can also help eliminate past biases and provide objective results. IdentiPHI performs such assessments for both private- and public-sector organizations.

IdentiPHI assessments and the IdentiPHI CompliSoft tool can help enterprises conduct effective identity management assessments. CompliSoft is an automated methodology toolkit populated with best practices and regulatory compliance frameworks such as those of the Health Insurance Portability and Accountability Act (HIPAA), Federal Information Security Management Act (FISMA), Gramm-Leach-Bliley Act (GLBA), and joint International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) 17799 specification. CompliSoft can enhance and streamline operational assessments, provide recommendations for specific technology deployments, and track implementation progress.

Implementing identity management: The HSPD-12 mandate

The White House issued Homeland Security Presidential Directive 12 (HSPD-12) on August 27, 2004, with the goal of improving the U.S. federal government's identity management programs, including changing the inconsistent and potentially insecure identification methods traditionally used to access federal information systems and physical facilities.[1] A set of published guidelines and technical requirements were subsequently released to define strategies for addressing the challenges inherent in institutional identity management. The same objectives, guidelines, and strategies adopted by federal agencies in response to HSPD-12 can also help commercial enterprises improve their security systems.

One of the primary objectives of HSPD-12 is to deploy common security strategies across different federal agencies. These measures are intended to help increase overall security, mitigate identity fraud, and improve the efficiency with which the government provides access to resources across agencies and departments. To successfully combine technological and process requirements into a solution with a user interface that is both intuitive and cryptographically secure demands either a large integration effort or a tool specifically designed for the task.

Following the release of HSPD-12, the National Institute of Standards and Technology published Federal Information Processing Standard (FIPS) 201, "Personal Identity Verification (PIV) of Federal Employees and Contractors," a standard that outlines specific means of meeting HSPD-12 requirements.[2] Several government and private-sector organizations contributed to this document, which specifies interoperable smart cards as the primary mechanism for secure government authentication and access control, with biometrics and passwords as complementary mechanisms. FIPS 201 also addresses two principal aspects of deployment: the technical specifications of the cards and their content, and the processes necessary to help ensure a consistent level of trust between authorities issuing and relying on the cards.

The FIPS 201 approach can help provide clear security and interoperability benefits for both government agencies and commercial enterprises. The next critical step is to find technologies that enable organizations to meet government mandates while keeping budgets under control.

Cost-effective solutions for meeting HSPD-12 requirements

Without cost-effective technologies to help secure and automate security processes, the management costs of meeting HSPD-12 and FIPS 201 requirements could become prohibitive—for both government agencies and commercial enterprises. Dell and IdentiPHI are working together and with additional partners, including ActivIdentity, to develop solutions that can help simplify the deployment and management of technology components by integrating user enrollment, document capture, biometric capture, card production, credentialing, issuance, printing, acquisition, and life-cycle management within

[1] To read the full text of HSPD-12, visit www.whitehouse.gov/news/releases/2004/08/20040827-8.html.

[2] FIPS 201 is available at csrc.nist.gov/publications/fips/fips201-1/FIPS-201-1-chng1.pdf.

the IdentiPHI suite of applications. An experienced team, which includes Dell, IdentiPHI, and other leading identity management system integrators and vendors, can help enterprises obtain products and services designed to provide efficient and effective identity management deployments.

Various technologies and vendors are required to provide a comprehensive identify management solution. Figure 1 illustrates the typical events and technology components of an HSPD-12-compliant personal identity verification (PIV) workflow, including the following operational roles defined by FIPS 201:

- **Applicant:** Requests PIV card and credentials
- **Sponsor:** Proposes applicants
- **Registrar:** Completes applications, enrolls applicants, and requests cards
- **Digital signatory:** Approves card requests
- **Issuer:** Issues cards
- **Authentication certification authority:** Signs and issues PIV authentication certificates

The Dell and IdentiPHI solution addresses the FIPS 201 requirements from initial user provisioning requests through the issuance and life cycle of user credentials. The solution has the flexibility to work with existing identity proofing methods, or a new vetting process can be determined during FIPS 201 solution development using the tools in the suite. Following this process, the card production request is passed along to a card management system along with the PIV data required for issuance. Several types of credentials are placed on the card, the sources of which may include certificate authorities, hardware security modules, biometric capture devices, databases, and Lightweight Directory Access Protocol (LDAP) directories. Card production can be either distributed with a desktop card production facility, or centralized using a hosted card management system model or shared service provider.

Essential to smart card flexibility and security are the applets loaded onto the card during manufacture or issuance. The ActivIdentity applets used by the Dell and IdentiPHI solution have been validated on smart cards from leading vendors such as Oberthur, and similar applets are widely used in smart card implementations, including the Department of Defense Common Access Card. In addition, ActivIdentity ActivClient software can natively enable smart card–based secure remote access, local and network login, physical access, and e-mail and other applications. After issuance, PIV cards verify the cardholder's identity and provide credentials for both physical (offices and other facilities) and logical (computers and networks) access.

Deploying IdentiPHI enterprise software

IdentiPHI enterprise software is an integral part of the comprehensive Dell and IdentiPHI identity management solution, and is designed to allow organizations to use smart cards, biometrics, and tokens to help achieve the following goals:

- **Enhance data security:** Multifactor authentication using smart cards, biometrics, and tokens can help substantiate user identity while strengthening network security.
- **Improve user convenience and reduce support calls:** Freeing users from using and remembering network access passwords can help reduce support calls related to password resets and consequently support center costs.
- **Comply with government standards:** Meeting the requirements of directives such as HSPD-12 and FIPS 201 can help protect data from unauthorized access.
- **Enhance administration:** The IdentiPHI directory-based policy inheritance model for user access can help provide authentication management aligned with enterprise infrastructures. IdentiPHI software interfaces are highly configurable, allowing authentication management at the domain, organization unit, server, desktop, and user levels.

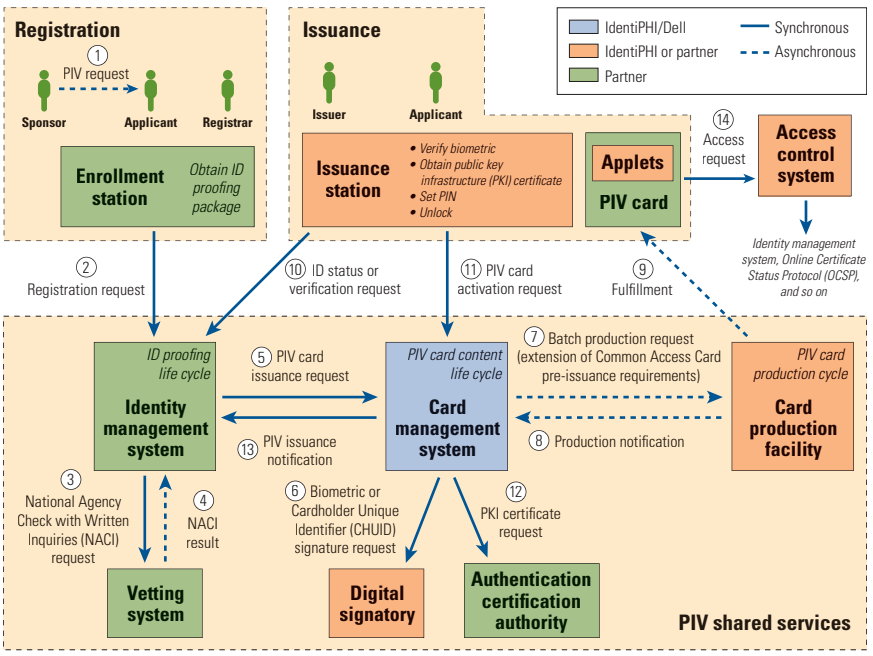

Figure 1. HSPD-12-compliant personal identity verification workflow

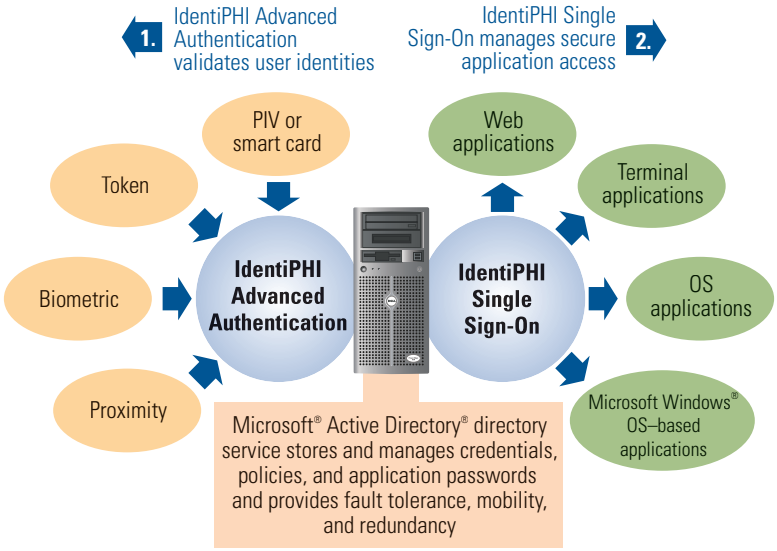

Figure 2. Identity management with IdentiPHI Advanced Authentication and Single Sign-On

- **Perform extensive auditing:** IdentiPHI configurable Simple Network Management Protocol (SNMP) and Simple Mail Transfer Protocol (SMTP) statistics can generate comprehensive real-time reports to help enable effective security management.
- **Reduce deployment overhead:** Rapid, efficient, and automated device self-enrollment can help significantly reduce deployment costs.

IdentiPHI enterprise software can help provide the performance, security, scalability, and flexibility required for successful identity management programs. Based on open standards to enable interoperability, this software can scale to high-volume environments and provide comprehensive security and auditing features; customizable workflows, user interfaces, and card profiles; and post-issuance management capabilities. A modular, vendor-agnostic approach enables this software to work with multiple vendors' products, so enterprises that have already partially deployed other solutions and want to avoid being locked into proprietary hardware or software can still benefit from it.

The IdentiPHI Enterprise Security Suite is designed as a modular security solution encompassing both IdentiPHI Advanced Authentication and IdentiPHI Single Sign-On, and includes a server component and management console. It enables enterprises to easily enforce and manage network security and password management policies across the organization, and can take advantage of existing directory services infrastructures, thereby helping reduce the need for a costly investment in additional hardware.

Enterprises can deploy Advanced Authentication and Single Sign-On independently or in combination to enable a comprehensive, integrated security solution. Advanced Authentication can help significantly increase network security by requiring users to prove their identity at network login according to enterprise security and authentication policies; providing a single point of authentication can also help simplify network and application access management. Single Sign-On can help manage the numerous application login IDs, passwords, and credentials that users must remember, and can help reduce repetitive logins by providing secure, simple application access.

Figure 2 illustrates how Advanced Authentication and Single Sign-On can work together to provide comprehensive identity management.

Protecting enterprise assets with identity management

As security threats continue to grow, so do enterprises' need to combat those threats, and creating a comprehensive identity management program—particularly one based on HSPD-12 and FIPS 201 requirements—can be a critical step toward doing so. IdentiPHI CompliSoft can help enterprises evaluate their needs by providing operational assessments and deployment recommendations, and the IdentiPHI Enterprise Security Suite, including IdentiPHI Advanced Authentication and IdentiPHI Single Sign-On software, can provide important elements of a secure environment, including such security services as file encryption; multifactor authentication using smart cards, biometrics, and tokens; and card management systems. When implemented as part of the Dell and IdentiPHI comprehensive identity management solution, these components can help enterprises create a secure environment to protect their critical information and other assets.

Mark Norwalk is the chief technology officer at IdentiPHI and oversees IdentiPHI security services and products. He has more than a decade of experience in management and a variety of technical leadership positions.

Craig Phelps is a security strategist in the Dell Enterprise Product Group, where he works across enterprise product teams to enable Dell core product and solution sets to meet the security needs of Dell enterprise customers. He is a Certified Information Systems Security Professional (CISSP) and received his B.A., B.S., and M.B.A. from Brigham Young University.

FOR MORE INFORMATION

IdentiPHI:
www.identiphi.net

Dell security solutions:
www.dell.com/security

Protecting Mobile and Remote Microsoft Windows–based Computers from Crimeware with

Symantec Client Security

The same mobile and remote computers that can be key to business success, such as notebooks, can, if not properly protected, open the door for hackers collecting corporate passwords, credit card information, or any other data from which they can profit. Symantec, a world leader in providing information security, availability, and integrity in Microsoft® Windows® OS environments, offers Symantec® Client Security software to help provide integrated, comprehensive, and proactive protection for Windows-based client systems against the aggressive threat of crimeware.

BY LAUREN DUDA

Related Categories:

Desktops

Enterprise security

Intrusion prevention

Notebooks

Security

Symantec

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

Today's IT threat landscape includes not only individual glory-seeking hackers, but also criminal organizations making concerted and financially motivated efforts to steal confidential information from specific organizations. These hackers can create back-door entrances to networks by infecting or otherwise compromising notebook computers, smartphones, remote connections to virtual private networks (VPNs), or instant messaging sessions. Although such attacks can start small, such as a Trojan horse that gathers user keystrokes and personal passwords or tracks Web site visits, the ultimate target can be important confidential data—credit card and bank account information, business plans, employee records, or corporate financial data prior to an IPO or an announcement that could affect a company's stock price. An attack could also have the sole purpose of tarnishing a company's image or weakening its brand.

Successful infections can affect the productivity of businesses and other enterprises. When spyware or adware infects network endpoints, system speed and user productivity can plummet, and help desks can be inundated with support calls from unhappy users unable to access information or run critical applications. IT administrators, meanwhile, may not have enough time or staff to continually track down, quarantine, and repair the infected endpoints.

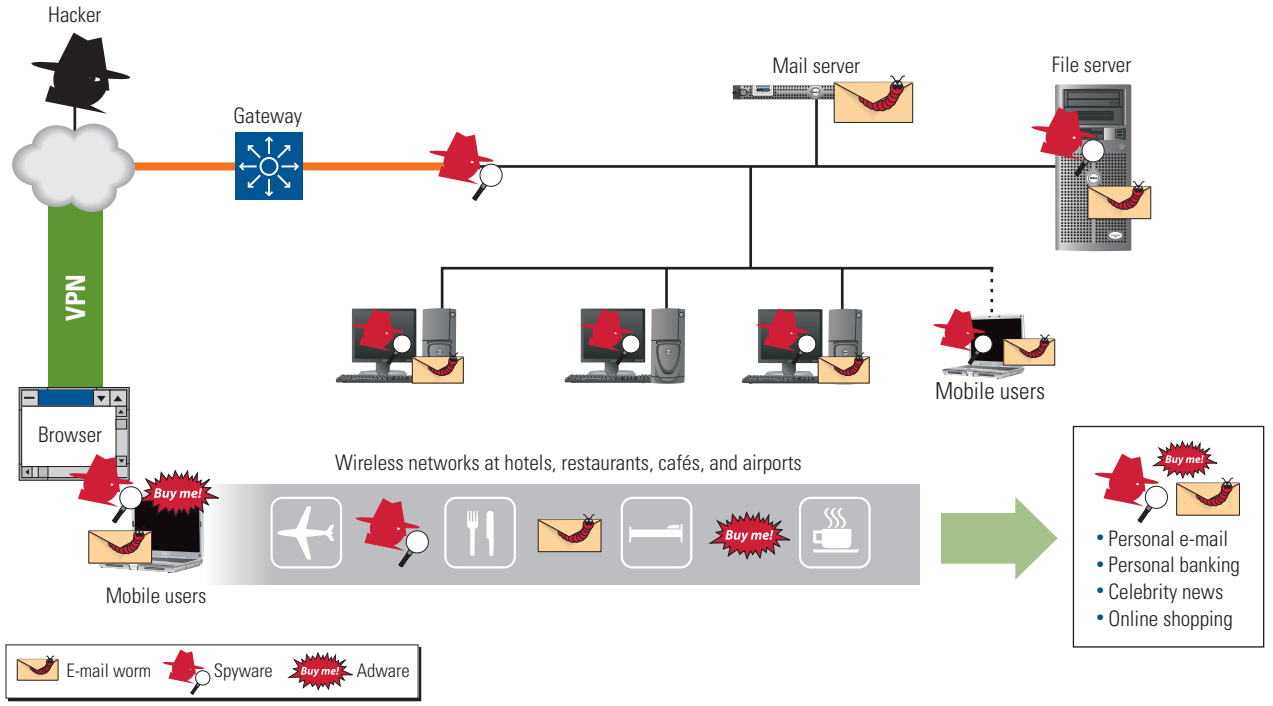

Figure 1. Potential security vulnerabilities created by mobile and remote users

Although antivirus and anti-spyware technology, such as that incorporated into Symantec AntiVirus™ software, can play an important role in defending standard desktops, this technology must be joined by a coordinated, multilayered defense that includes proactive vulnerability- and signature-based intrusion prevention and firewall control of both inbound and outbound traffic. Symantec Client Security software can help provide such integrated, comprehensive security for Microsoft Windows–based systems.

Mobile workers: Opening the back doors

Notebook computers can be critical for salespeople, executives, and other employees that are constantly on the move—without notebooks, their productivity could come to a standstill. But these notebooks can also become serious vulnerabilities without proper protection.

When mobile workers take their notebooks on the road and connect to networks in hotels, airports, or Internet cafés, those notebooks are no longer protected by the organization's office network perimeter defenses. Although a few mobile workers might use their notebooks strictly for business, many will likely also use them for personal activities, such as banking, shopping, checking stocks, reading celebrity news, downloading ring tones, accessing personal e-mail, playing online games, downloading MP3s, file sharing, and visiting chat rooms.

Trojan horses, bots, spyware, and worms can infect notebooks undetected through these sites, downloads, and e-mails; when the notebooks then return to the office and reconnect to the network, they can open a back door into the network. Figure 1 shows some of the ways mobile and remote users can create security vulnerabilities that malicious programs can exploit to infiltrate the network and spread to other computers. A February 2005 study conducted by the Enterprise Strategy Group found that 43 percent of worm attacks originated from notebooks carried into the network confines by employees, and 34 percent came from notebooks brought in by nonemployees. The study also found that the fourth most common source for worm attacks, at 27 percent, was home systems connecting to the network through the organization's VPN.[1]

Symantec: Closing the back doors

Employing multilayered security on all clients—including notebooks, desktops, and other remote computers—can enable organizations to minimize vulnerabilities by closing back-door breaches and helping protect against crimeware threats. The primary elements of such a client security system must be integrated with one another and include virus protection and remediation, spyware protection and remediation, vulnerability- and signature-based intrusion prevention, and firewall control of both inbound and outbound traffic.

[1] "Network Security and Intrusion Prevention," by Jon Oltsik, Enterprise Strategy Group, February 2005.

Features		Benefits
Client firewall	Inbound traffic control	• *Mobile notebooks:* Helps protect notebooks from peer-to-peer attacks from other wireless network users, potential network infections, and Internet attacks • *Network desktops:* Helps protect desktops from peer-to-peer attacks from other network users and potential network infections
	Outbound traffic control	• Helps contain infected computers and prevent them from spreading malware (initially coming from a USB drive, CD, personal e-mail, or similar source) across the network • Helps keep confidential information (passwords, bank account numbers, and so on) from leaving computers through malicious code
Intrusion prevention	Vulnerability-based protection against unknown threats (generic exploit blocking)	• Helps proactively protect against exploits and variants trying to take advantage of a vulnerability • Quickly distributes vulnerability-based signatures after a vulnerability announcement (typically within 24 hours)
	Signature-based protection against known threats	• Helps detect known threats on a computer, a network, and the Internet • Works in collaboration with the antivirus, anti-spyware, and firewall functions to help neutralize malicious activity

Figure 2. Symantec Client Security features

Integration is a critical aspect of client security. The virus protection, spyware protection, intrusion prevention, and firewall traffic control elements must be able to communicate with one another and work together to help protect the client system. Nonintegrated systems can require frequent manual intervention, which can weaken the system's ability to combat threats. In addition to providing a coordinated defense, integrated client security can be more easily managed than a collection of individual products that do not work together—for example, by allowing centralized management from a single console rather than multiple consoles.

Using individual antivirus, anti-spyware, intrusion prevention, and firewall products from four different vendors can also result in four different sets of licensing terms, four different service contracts, four different support centers, and four different sets of update subscription terms. An integrated client security system therefore not only helps simplify IT security management, but can also help significantly reduce licensing costs. Organizations that own a current Symantec antivirus product, such as Symantec AntiVirus, may find that an integrated client security tool such as Symantec Client Security can help improve security in a cost-effective way.

Employing multilayered security on all clients can enable organizations to minimize vulnerabilities by closing back-door breaches and helping protect against crimeware threats.

Symantec Client Security can help keep enterprise client systems safe by providing comprehensive and proactive protection against crimeware. This software can automatically detect and repair the effects of viruses, spyware, adware, and other malicious intrusions in real time. Its vulnerability-based detection works with antivirus and traffic control tools to help detect and block known, unknown, and emerging vulnerabilities to help keep systems safe and protect valuable and confidential information. Figure 2 lists some of the ways that Symantec Client Security can help protect systems and information.

Symantec Client Security: Providing comprehensive protection for Windows-based systems

Neither network perimeter defenses nor basic antivirus technology can typically provide sufficient protection against an organized, targeted crimeware attack. To help provide such protection, organizations need multilayered client security designed for this threat landscape, integrating virus protection, spyware protection, vulnerability- and signature-based intrusion prevention, and firewall control of both inbound and outbound traffic. Symantec Client Security integrates these functions into a comprehensive Windows-based client security solution.

Lauren Duda is a product marketing manager for the Endpoint Security team at Symantec. Lauren currently has global product marketing responsibility for Symantec AntiVirus and Symantec Client Security. She has a B.A. from the University of California, Los Angeles, and an M.B.A. from California State University, Long Beach.

FOR MORE INFORMATION

Symantec:
www.symantec.com

Symantec Client Security:
www.symantec.com/Products/enterprise?c=prodinfo&refId=839

Symantec Corporation. "Protecting Client Systems from the Crimeware Invasion." August 2006. eval.veritas.com/mktginfo/enterprise/white_papers/ent-protecting_client_systems_wp_08_2006.en-us.pdf

Securing and Archiving Instant Messages:

A Critical Step for Securing Microsoft Messaging Environments

As part of a secure and productive messaging environment where users can take advantage of the latest communication tools, Symantec® IM Manager can help organizations control instant messaging (IM) while complying with legal regulations and corporate policies. IM Manager supports both public and enterprise IM networks and helps manage, secure, log, and archive IM traffic.

BY LEE WEINER AND CRAIG PHELPS

Related Categories:

Enterprise security

Instant messaging

Security

Symantec

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

Instant messaging (IM) has become a key tool for enterprise communication, enabling employees to share information and collaborate in real time with colleagues, partners, and customers around the world. As a result, enterprises must find a way to safely enable IM while simultaneously satisfying the management policies, security needs, and regulatory compliance requirements associated with its use.

Dell and Symantec recognize the need for secure, highly available messaging environments supporting both e-mail and IM. For example, Figure 1 illustrates how important enterprises consider such an environment for e-mail. The first step in developing such an environment is to secure, protect, and archive Microsoft® Exchange Server deployments. This can be accomplished by implementing the Dell™ Secure Exchange Reference Architecture, which uses validated industry-standard components to simplify the deployment and scalability of secure enterprise messaging environments.[1]

Question: How important is it to implement a solution that integrates e-mail security, availability, backup, and archiving all together?

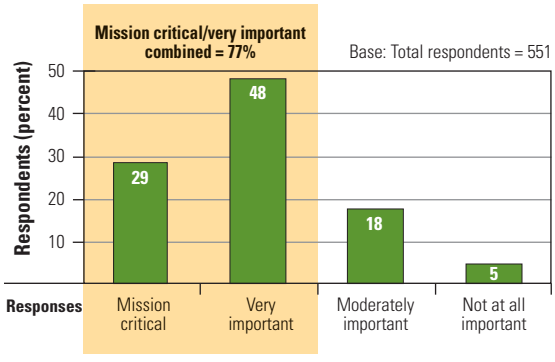

Source: Survey of technology decision makers conducted by Ziff Davis Media on behalf of Symantec, March 2006. For more information, see "A Single Solution for Messaging Management and Security," by Ziff Davis Media, 2006, www.interop.com/newyork/pdfs/symantec-white-paper.pdf.

Figure 1. Survey responses on the importance of integrated e-mail solutions

[1] For more information about this architecture, see "Implementing the Dell Secure Exchange Reference Architecture," by Suman Kumar Singh and Bharath Vasudevan, *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20060452-Singh.pdf.

The second critical step is to do the same for other messaging platforms, including IM, by deploying Symantec layered messaging security along with Symantec IM Manager software.

Understanding IM security and regulations

Because enterprise IM deployments are growing rapidly and can often be unmanaged and unmonitored, IM use can expose organizations to numerous security risks, including the following:

- Blended threats that use IM to bypass traditional security software
- Identity theft, spoofing, and phishing over IM
- Advanced spyware and spam over IM
- Proprietary information security leaks over IM
- Targeted IM attacks on enterprise domains

Widespread enterprise IM use can also mean that, in some cases, archiving requirements for IM are the same as those for e-mail and other enterprise messaging systems. Important regulatory requirements relevant to IM include the following:

- **Securities and Exchange Commission (SEC) rules 17a-3 and 17a-4:** Require firms to retain all Internet communications pertaining to their business, which includes IM
- **NASD rules 3010 and 3110:** Require firms to supervise, review, and demonstrate compliance procedures for electronic correspondence, which includes IM
- **New York Stock Exchange (NYSE) rules 342 and 440:** Explicitly include IM in NYSE information memo 03-7 as a type of communication that must be archived under SEC regulations
- **Department of Defense directive 5015.2:** Sets standards for records retention, which includes IM
- **Sarbanes-Oxley Act section 404:** Includes extensive requirements for monitoring and reporting financial communication and documentation
- **Health Insurance Portability and Accountability Act:** Requires medical and pharmaceutical companies to retain patient records during clinical trials and provide for the records' privacy, which includes information shared over IM
- **Federal Energy Regulatory Commission (FERC) regulations:** Require logging and auditing transaction-related information, which includes IM

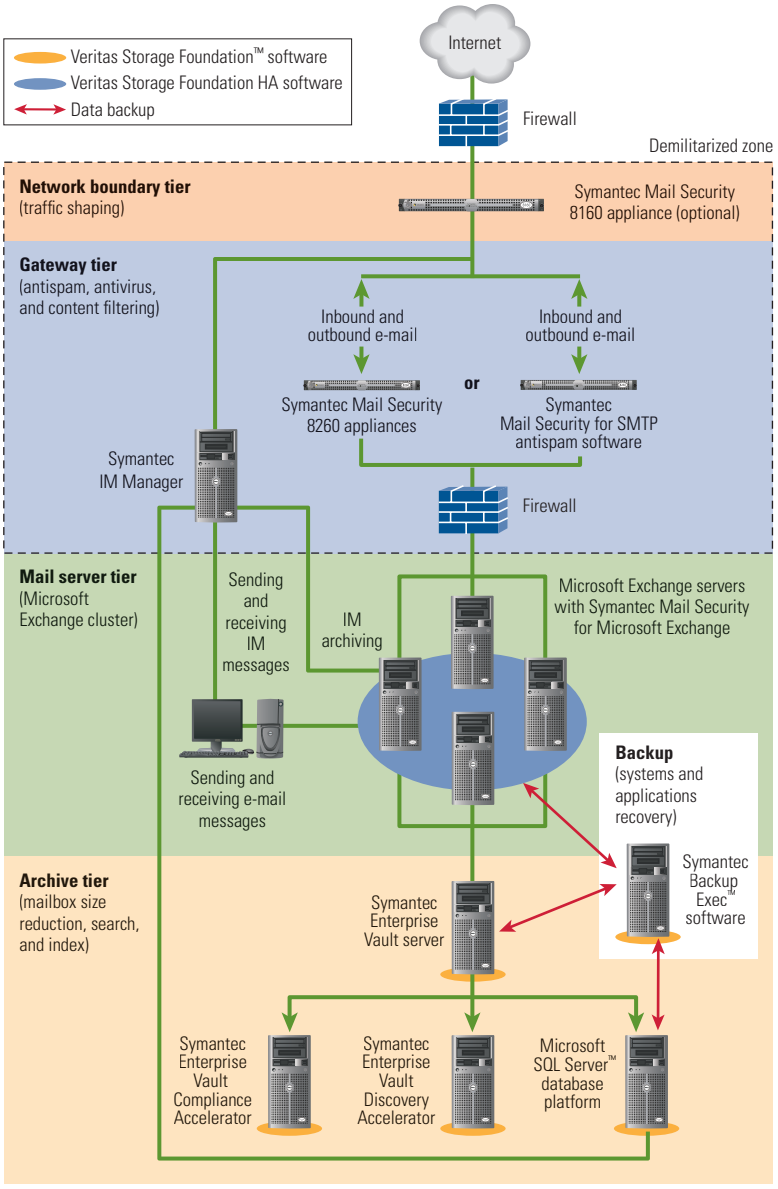

Figure 2. Symantec layered messaging security architecture

- **Federal Communications Commission (FCC) regulations:** Require extensive record keeping and storage, including supervising and indexing books and records
- **Corporate IM communication policies:** May require monitoring and controlling IM as part of general employee communications policies

Implementing Symantec layered messaging security and Symantec IM Manager can help protect organizations against IM-based threats such as viruses, worms, and malware and help enable compliance with legal and corporate requirements.

Implementing Symantec layered messaging security

Symantec layered messaging security deploys different types of protection at defined tiers inside the messaging architecture (see Figure 2):

- **Network boundary tier:** Reduces spam volume outside the network
- **Gateway tier:** Filters e-mail and IM messages outside the network, at the messaging environment perimeter
- **Mail server tier:** Filters e-mail messages inside the network

The separate but interdependent aspects of the messaging infrastructure enable layered functions to provide mutually reinforcing protections. Symantec recommends removing unwanted content from the messaging system at the earliest possible point; the critical interception points are entry and departure points for external e-mail and IM messages (in the gateway tier) and distribution points for internal e-mail and IM messages (in the mail server tier).

Securing IM with Symantec IM Manager

Symantec IM Manager is designed to help enterprises manage, secure, log, and archive IM traffic. It can help deliver real-time threat protection; rapid deployment; enterprise-class scalability, reliability, and management; and regulatory compliance for enterprise IM use. IM Manager offers comprehensive support for public and enterprise IM networks—including certified integrations with the IM software of industry leaders such as Microsoft, IBM, AOL, ICQ, Reuters, Yahoo!, and Jabber—and includes granular policy controls for text messaging, file transfers, audio, video, voice over IP (VoIP), application sharing, and other real-time communication capabilities associated with IM. Figure 3 provides an overview of key Symantec IM Manager features.

IM Manager can also help provide preemptive, automatic threat identification and protection against IM viruses, worms, and malware through the patent-pending Symantec Real-Time Threat Protection System (RTTPS). RTTPS IM threat protection goes beyond traditional reactive security systems and safeguards. Instead, it monitors enterprise IM traffic and searches for network anomalies and potential malicious behavior. Once a potential threat is recognized, the RTTPS predictive protection filter can identify the new threat signature and stop the potential outbreak by blocking it at the point of propagation.

Delivering comprehensive enterprise messaging management

Deploying Symantec layered security in conjunction with the Dell Secure Exchange Reference Architecture and Symantec IM Manager can help provide a comprehensive, secure, and highly available messaging environment that incorporates antivirus, antispam, archiving, backup, and recovery capabilities. Symantec IM Manager can help provide these capabilities for IM deployments by enabling organizations to manage IM traffic; improve security by scanning IM messages for viruses, worms, malware, and other threats; and comply with regulatory requirements for IM tracking and archiving. Dell and Symantec plan to continually enhance these tools and the Dell Secure Exchange program as enterprise messaging requirements evolve.

Function	Symantec IM Manager features
Managing IM traffic	• *User management and access control:* Controls IM user, group, and domain access to disparate IM systems, including integration with enterprise directory structures • *Priority-based policy enforcement:* Establishes consistent IM usage policy enforcement, including real-time content filtering, granular file transfer, and advanced client feature controls • *Real-time analytics and reporting:* Tracks and analyzes IM usage and growth patterns with real-time alerting and notifications, trend reporting, and custom monitoring
Controlling IM security and usage	• *Zero-day protection:* Helps detect and protect against zero-day attacks with patent-pending technology • *Automatic threat updates:* Automatically updates virus and spam signatures from the Symantec Security Response Team • *Virus scanning and file transfer control:* Scans file transfers and uses the Symantec AntiVirus™ Scan Engine to help prevent infected or confidential files from traversing networks
Complying with legal and corporate IM requirements	• *Rich message archive:* Selectively captures and retains IM conversations with direct links to employee data from the corporate directory for enhanced retention and discovery • *Integration with Symantec Enterprise Vault software:* Integrates with Enterprise Vault for enterprise retention and discovery and comprehensive messaging management • *Real-time content filtering:* Blocks messages and notifies administrators when messages containing restricted phrases or inappropriate content are sent

Figure 3. Key Symantec IM Manager features

Lee Weiner is a senior product manager in the Enterprise Messaging Management Group at Symantec.

Craig Phelps is a security strategist in the Dell Enterprise Product Group. He is a Certified Information Systems Security Professional (CISSP) and received his B.A., B.S., and M.B.A. from Brigham Young University.

FOR MORE INFORMATION

Dell and Symantec:
www.dell.com/symantec

Dell Secure Exchange:
www.dell.com/secure_exchange

Zurcher, Werner, and Garrett P. Jones. "Providing Multi-Tiered Security for Microsoft Exchange Environments." *Dell Power Solutions,* May 2006. www.dell.com/downloads/global/power/ps2q06-20060298-Symantec.pdf

Managing IT Security Costs

with Identity and Access Management

Identity and access management (IAM) enables organizations to align their security management strategies with their business goals by centralizing and automating the management of user identities and access to protected resources and services, enforcing privacy and security policies, and monitoring and auditing the security environment for compliance with legal and corporate mandates. CA® software can help administrators implement IAM in their own organizations.

BY SUMNER BLOUNT

Related Categories:

Access management

CA

Identity management

Security

Security software

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

Enterprise IT managers today face a wide array of pressures. Not only must they provide a secure environment to protect enterprise assets and industry reputation, but they also are often asked to do so at a lower cost than in the past. This pressure to increase IT efficiency and manage costs exists alongside demands to enhance and expand applications and services. These demands spring from several emerging trends, including the following:

- **Increased need for regulatory compliance:** The burdens of creating effective internal security controls to comply with regulations such as the Sarbanes-Oxley Act and the Health Insurance Portability and Accountability Act can fall heavily on IT security groups.
- **Increased merger and acquisition activity:** As organizations grow through acquisitions or mergers and incorporate new user populations, applications, and heterogeneous legacy systems into an existing infrastructure, the complexity of IT security can increase.
- **Increased user populations:** Extending applications to partners and online customers can expand IT demands significantly.

These factors, among others, can cause enterprise IT groups to search for ways to streamline infrastructure management, but the competing requirements to reduce costs and increase services can present huge challenges. Implementing integrated identity and access management (IAM) can help streamline IT security

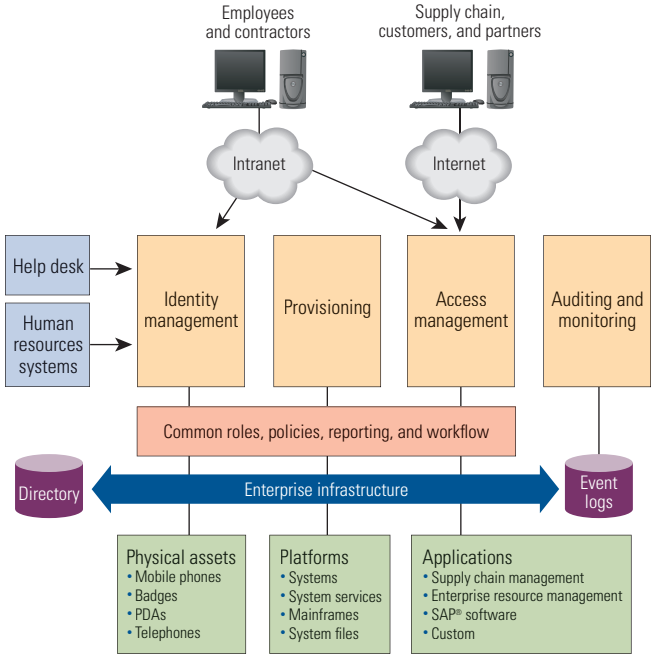

Figure 1. Integrated identity and access management platform

management, improve operational efficiency, and increase productivity, thereby helping reduce overall IT costs. CA offers several IAM software solutions to help administrators realize these benefits (see the "Implementing identity and access management" sidebar in this article).

Understanding identity and access management

User identities and access privileges are a core element of e-business. Successful IAM—which encompasses processes and systems that determine who has access to which applications, databases, and platforms, and the conditions under which that access should be granted—can help IT administrators effectively align security with business goals, protect vital assets, streamline operations, and achieve regulatory compliance by providing the following key capabilities:

- **Identity administration:** Enables the creation and administration of user identities and profile information
- **User provisioning:** Allocates the appropriate accounts and resource access rights to each user, and de-provisions them at the appropriate time (for example, when a user leaves the organization)
- **Access management:** Helps ensure that the organization maintains information and application integrity by preventing unauthorized access to critical resources, including Web and enterprise applications, systems and critical services, databases, and repositories
- **Monitoring and auditing:** Provides aggregation, filtering, analysis, and correlation of security events across all components within the environment

Sound IAM practices and systems can provide a foundation for security management by addressing identity-related system exposures, enforcing a consistent security policy across the organization, and delegating administrative access power. It can also help reduce administrative costs through integrated auditing and automated management. Figure 1 illustrates the key elements of an integrated IAM platform and the range of resources that must be protected.

Improving operational efficiency

The challenge to "do more with less" can be met by improving operational efficiency. To "do more" means increasing the productivity of each employee. Inefficient internal processes, excessive manual procedures, and dealing with issues unrelated to job function can reduce productivity. To achieve increased productivity "with less" may require a reduction in overall IT security management costs, which can be done by eliminating needless processes, making users self-sufficient, and automating IT administrative tasks that require excessive amounts of time to perform manually.

Sound IAM practices and systems can provide a foundation for security management by addressing identity-related system exposures, enforcing a consistent security policy across the organization, and delegating administrative access power.

Security administration is one of the most important areas in which administrators can do more with less. Administrators may be performing tasks manually or multiple times across different systems and applications—or both. Such tasks can include creating and managing user identities and access rights, provisioning users, managing security events, and administering applications. An integrated IAM implementation can help increase administrative efficiency and reduce administrative costs for these tasks.

Creating and managing user identities and access rights

Managing multiple identities for each user, especially when these identities are scattered in different places and managed individually for each user, can be time-consuming, inefficient, and expensive. Enforcing access rights within each application can lead to wasted administrative expense and inconsistent access rights

across applications. And each separately stored identity—the management of which involves not only creating the identity but also updating it as the user profile changes—entails additional expense. An IAM implementation allows administrators to create and manage user identities and access rights centrally, and can even allow users to manage some of their own profile attributes, both of which can help improve efficiency and reduce costs.

Substantial administrative savings can result from reducing 3 million events requiring analysis to 8 events requiring action, without even considering the security advantages that such an implementation can provide.

Organizations can calculate savings from centralized user identity creation and management using the metrics of average time to create an identity, expected number of identities created (or updated) per unit of time, and average number of places that an identity must be stored. Deploying a centralized IAM implementation can help reduce the average time to create or update a user profile and to approve an access request, the total time spent correcting access rights discrepancies across systems and applications, and the proportion of access requests that deviate from established access request processes or established user-role definitions.

Provisioning users

Provisioning users—the process of granting and managing access to systems or applications—can be one of the most time-consuming administrative tasks, especially when providing new users with all their required accounts and applications across multiple systems. An automated provisioning service allows administrators to perform this process centrally and in some cases without direct administrator intervention, which can help reduce IT costs. For example, accounts on target systems and access rights to protected applications can be set up automatically based on the user's role or organization group membership. This service can help reduce costs based on the number and arrival rate of new users, the number of accounts and applications that typically require access provisioning, and the time required to grant access to each of these accounts or applications (which depends heavily on the type of account and the system where the account resides).

Managing security events

One of the biggest problems IT administrators face is security information overload, which occurs when each component in a large IT environment produces large audit logs of events within that subsystem. Such components include Microsoft® Windows® OS–based systems, UNIX® OS–based systems, intrusion detection systems, firewalls, antivirus software, and many other components that generate log information.

Administrators must have a way to make sense of all this data. Not only can this task consume large amounts of resources, but manual analysis can also lead to security holes because it is difficult to correlate and draw conclusions from seemingly unrelated events to uncover security problems. Comprehensive security event management can enable administrators to perform the following tasks:

- Centralize collection and aggregation of security event information across environment components
- Normalize and filter security log entries
- Correlate apparently independent events to identify relationships that might indicate breach attempts
- Perform visual analysis of the status of target system security attributes
- Produce customized reports to provide each potential reader with specific information

For example, a large IT environment might generate 3 million log entries per day; any reasonable manual analysis of this massive amount of information could take a similarly massive amount of time. Figure 2 illustrates the potential benefits of filtering and correlating these security events automatically: substantial administrative savings can result from reducing 3 million events requiring analysis to 8 events requiring action, without even considering the security advantages that such an implementation can provide.

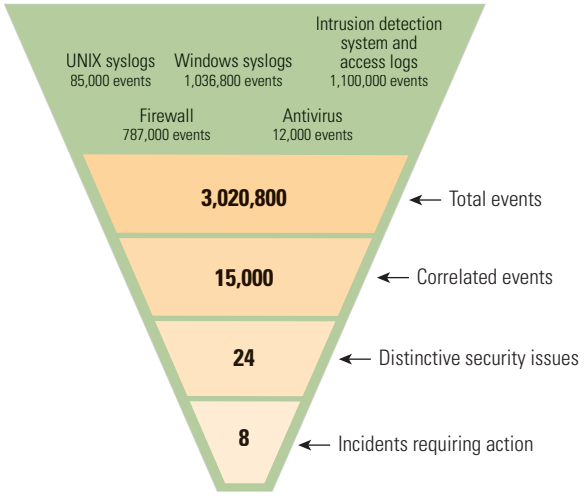

Figure 2. Example results of automatically filtering and correlating daily security log traffic

Administering applications

Enforcing access rights security within each application can incur significant application development and maintenance costs, sometimes simply to implement similar or identical security modules across multiple applications. By separating security enforcement from applications in a separate access management service, an IAM implementation helps administrators to simplify applications and avoid rigorously testing large amounts of security code, which can help reduce or eliminate the development and maintenance costs associated with these components.

The potential savings offered by this approach are specific to each environment and depend on the amount of access enforcement code within each application, the number of applications, the security testing overhead for the application security modules, and ongoing module maintenance.

Increasing productivity

In addition to increasing operational efficiency, an IAM implementation can help increase productivity and reduce or eliminate some of the hidden costs in IT environments. Two problems in particular can result from not automating user identity management: delays in providing new users with system and application access, and excessive management overhead for access request approvals.

IMPLEMENTING IDENTITY AND ACCESS MANAGEMENT

CA is the worldwide market revenue leader in IAM software for 2005* and offers a range of comprehensive and integrated IAM solutions. The table below summarizes some of the areas where integrated IAM can help streamline IT security management along with the CA software that can provide these capabilities.

IAM component	CA solution
Identity administration	CA Identity Manager
Automated user provisioning	CA Identity Manager
Automated security event management	*e*Trust® Security Command Center
Access management and single sign-on for Web applications	*e*Trust SiteMinder®
Access management for systems, files, and databases	*e*Trust Access Control
Automated account removal for mainframes	*e*Trust Cleanup

*"Worldwide Identity and Access Management 2006–2010 Forecast," by IDC, Doc #202728, August 2006.

Expediting system and application access for new users

New users can wait days or even weeks for full access to the system accounts, applications, physical resources, and information necessary for their job duties, and the potential impact of this delay can be substantial. For example, for a company with a 40-hour delay for account allocation and resource access, hiring an employee with a $65,000 salary can have a resulting initial productivity loss of up to $1,250 for that employee—a week's salary paid even if the employee cannot perform his or her duties. A comprehensive user provisioning implementation can help solve this problem by automating the process and thereby helping substantially improve provisioning efficiency.

An effective IAM implementation can help simplify and increase the security of the entire process of managing users and their access to protected resources.

Streamlining access request approvals

Paper-based access request approval can have a significant impact on management productivity. Automating this process can free administrators to focus on more important tasks and provide an audit trail of the approval process for later analysis. User provisioning with full workflow capabilities can also enable administrators to define complex approval dependencies, so that the complete approval process can be replicated within the provisioning implementation.

Creating effective enterprise identity and access management

Implementing identity and access management can help improve operational efficiency and reduce the costs associated with helpdesk support, system administration, and application development and maintenance. It can also help increase user productivity by making resources available quickly and streamlining long approval processes. An effective IAM implementation can help simplify and increase the security of the entire process of managing users and their access to protected resources.

Sumner Blount is the director of security solutions at CA.

FOR MORE INFORMATION

CA identity and access management software:
www.ca.com/iam

Implementing the Dell Secure Exchange Reference Architecture

The industry standards–based Dell Secure Exchange Reference Architecture is a comprehensive architecture for Microsoft® Exchange Server designed to help improve security, availability, and scalability in messaging environments. This article discusses the components and best practices of this reference architecture.

BY SUMAN KUMAR SINGH AND BHARATH VASUDEVAN

Related Categories:

Dell PowerEdge servers

Microsoft Exchange

Security

Symantec

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

Whether in the office or on the road, communicating internally or closing business with customers, enterprise IT users depend on e-mail to get their jobs done—and e-mail traffic is increasing rapidly. Managing this growing traffic while minimizing virus and spam threats can be a daunting task. A base messaging infrastructure that is secure, flexible, and scalable can help enterprises meet this challenge.

Dell Secure Exchange Reference Architecture

The Dell Secure Exchange Reference Architecture is based on industry-standard components that help provide e-mail data protection and scalability to support large Microsoft Exchange deployments. This approach allows enterprise IT organizations to choose the appropriate tools for different data center functions and enables features to evolve in a way that benefits the entire IT industry. The architecture components have been validated for interoperability, and partnering with industry leaders Microsoft and Symantec has allowed Dell to provide a comprehensive architecture for messaging environments, incorporating security and archiving in addition to basic e-mail functions.

Figure 1 illustrates the components of the Dell Secure Exchange Reference Architecture. This article discusses the perimeter network, front-end Exchange servers, back-end Exchange servers, and storage and tape backup architecture components.

Perimeter network

A perimeter network is typically the network segment closest to the Internet gateway, and is therefore the first network encountered by incoming traffic. If the network design includes a firewall, the perimeter network is typically part of that firewall. In the absence of a perimeter network, the front-end Exchange servers handle the

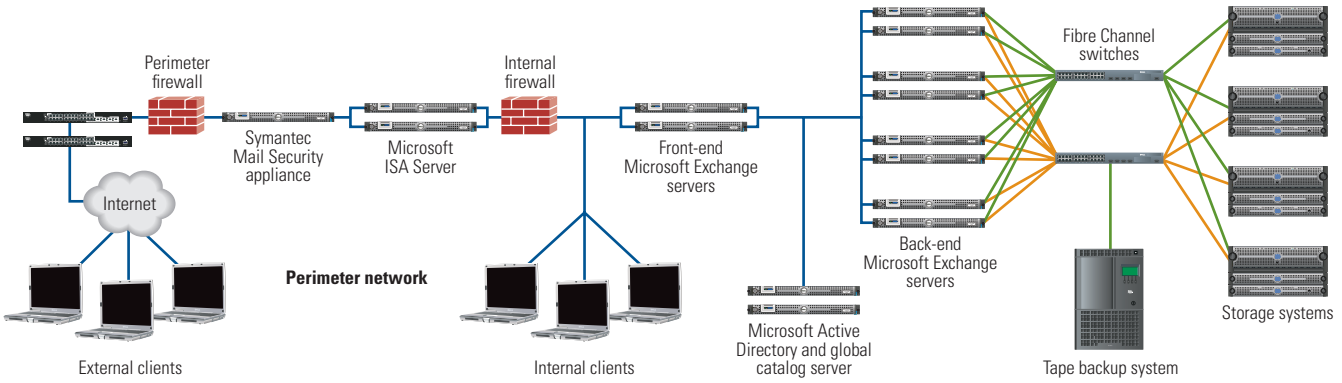

Figure 1. Dell Secure Exchange Reference Architecture

majority of incoming e-mail traffic, which can expose these servers to e-mail threats and viruses. Implementing a perimeter network can help enterprises control security by controlling external access to internal resources.

The amount of spam and the threat of viruses is ever-growing, and enterprises must prevent this spam from entering their internal networks. For deployments with more than 2,000 users, the amount of e-mail traffic can warrant a two-tiered approach to reducing unwanted e-mail traffic: traffic shaping and content filtering. For smaller deployments, both can still be implemented, but content filtering alone may be sufficient to reach performance goals.

Equally important to secure perimeter network deployments is restricting access to e-mail servers by unauthorized external clients or intruders. Microsoft Internet Security and Acceleration (ISA) Server can help provide this functionality by controlling traffic entering internal networks and traffic leaving messaging environments.

Traffic shaping involves intercepting and scanning incoming e-mail to determine its authenticity and relevance. This approach can help reduce unnecessary traffic before it reaches the internal network.

Traffic shaping

Traffic shaping involves intercepting and scanning incoming e-mail to determine its authenticity and relevance. This approach can help reduce unnecessary traffic before it reaches the internal network and can thereby help improve internal network performance and security. Organizations can implement traffic shaping by deploying antispam devices in the perimeter network.

The Symantec Mail Security 8160 appliance is a traffic shaper that integrates Symantec software with a Dell™ PowerEdge™ 1850 server. This appliance acts as a router, inspecting incoming Simple Mail Transfer Protocol (SMTP) traffic in real time and providing a mechanism to help reduce the amount of spam entering the internal network. It is designed to shape e-mail traffic at the TCP level and prevent spammers from forcing unwanted e-mail into the network. The appliance determines a sender's reputation based on a cumulative history and can automatically shape traffic based on that reputation.

Content filtering

In addition to traffic shaping, enterprises may want to control outbound content to help ensure that the organization is not perceived as a source of inappropriate or malicious content. Symantec Mail Security 8200 series appliances and Symantec Brightmail AntiSpam software can help address these challenges. Both products include an integrated virus and spam signatures update mechanism, which is frequently and automatically updated to provide the latest antivirus policies and rules, similar to antivirus clients running on a desktop computer.

The Symantec Mail Security 8260 appliance—which, like the Mail Security 8160, integrates Symantec software on a Dell PowerEdge 1850 server—is designed for environments with more than 1,000 users. It delivers antispam, antivirus, content filtering, e-mail firewall, and quarantine capabilities to help prevent unwanted traffic from entering internal networks.

Microsoft ISA Server

Enterprises can use Microsoft ISA Server to configure rules to securely publish internal mail services to external users by allowing access using specified protocols such as Messaging Application Programming Interface (MAPI), Post Office Protocol 3 (POP3), and Internet Messaging Access Protocol 4 (IMAP4). ISA Server can also handle inbound requests from client applications such as Microsoft Outlook, Outlook Web Access, and other POP3 e-mail clients and route them to the appropriate Exchange server on the internal network, thereby helping protect the internal mail servers from direct communication with external clients.

ISA Server is designed to protect the Exchange server by acting as a proxy to receive all requests for the Exchange server. Using ISA Server to handle inbound requests from client applications means that front-end Exchange servers can be moved from the perimeter network to the internal network, helping provide an additional layer of security for these servers. Because of these security benefits, best practices recommend deploying ISA Server in the perimeter network even if traffic shaping and content filtering appliances are already in place.

Front-end Microsoft Exchange servers

Microsoft Exchange Server 2003 supports a two-tiered architecture consisting of front-end and back-end servers. Front-end servers accept requests from clients and send them to the appropriate back-end server for processing. This architecture is recommended if the clients use multiple access protocols to access the Exchange server or if the Exchange environment includes multiple back-end servers; however, it can also be used in an environment with a single back-end server. Front-end servers should generally be deployed behind the internal firewall.

Front-end servers can provide benefits such as the following:

- **Single namespace:** Users can access their mailboxes using a single name even if the mailbox moves from one server to another or new servers are added to the back-end infrastructure. Outlook Web Access, POP3, and IMAP4 clients can also access the mailbox using the same URL.
- **Off-loaded processing:** Front-end servers can help improve messaging system performance by off-loading processing tasks typically performed by the back-end server in the absence of a front-end server, such as managing encryption and decryption processing of incoming and outgoing e-mail traffic.
- **Strengthened security:** The front-end server provides a single point of access for all incoming requests and traffic. Because it does not store user information, it helps provide an additional layer of security for mailboxes. It can also authenticate requests before sending them to the back-end server, thereby helping protect against security breaches, and eliminates the need to open Remote Procedure Call (RPC) ports from the perimeter network into the internal network.
- **Simplified scalability:** Because front-end servers provide a single namespace for all users, they allow enterprises to increase or decrease the number of front-end or back-end servers without disrupting users, helping simplify scalability.

Selecting the appropriate front-end server is critical to messaging system performance. This selection depends on multiple factors, including the number of users, number of back-end servers, protocols used, and functions performed by the front-end server. Because mailboxes do not reside on front-end servers, these servers typically do not require large or fast disk storage. However, they do typically require more processing power than back-end servers.

A typical front-end server would support 4 GB of memory and include scalable processor options while providing the network bandwidth required to host Exchange. The Dell PowerEdge 2950 server, when configured with up to two dual-core Intel® Xeon® processors and up to 32 GB of fully buffered memory, can provide the necessary processing power for front-end server functionality. Depending on the number and types of users, large Exchange deployments may require multiple front-end servers.

Back-end Microsoft Exchange servers

Back-end Microsoft Exchange servers host mailboxes and public folders. End-user performance and messaging infrastructure availability depend heavily on the selection and design of the back-end Exchange infrastructure, and particularly on choosing the appropriate hardware and mailbox design. Providing sufficient processor, disk, memory, and network resources can help prevent the back-end infrastructure from becoming a bottleneck for the entire messaging system.

Some general guidelines for back-end servers are as follows:

- **Processors:** Back-end server tasks are not typically processor intensive. However, if a server is running other applications, such as antivirus or antispam software, it should have sufficient processing power to support those applications. If the back-end server processor utilization is consistently greater than 60 to 70 percent, organizations should closely monitor the server's performance, because it may require additional processing power.

If a server is running other applications, such as antivirus or antispam software, it should have sufficient processing power to support those applications.

- **Disks:** Exchange Server 2003 is an I/O-intensive application: All client activity causes updates to the Exchange database, which produces I/O operations to disk. The disk subsystem should be able to meet the required I/O performance, not just mailbox capacity requirements. External storage options such as direct attach storage or storage area networks (SANs) can help provide a scalable disk I/O subsystem.
- **Memory:** Exchange Server 2003 is a 32-bit application, meaning that the maximum amount of memory it can use efficiently is limited to 4 GB. Deploying back-end servers with 4 GB of physical RAM can help take advantage of Exchange

capabilities; deploying servers with more than 4 GB of RAM can have a negative performance impact.

- **Network:** Exchange servers and messaging clients access the Microsoft Active Directory® directory service when logging on to a network, connecting to a mailbox, or accessing server-based address lists. Because these activities can generate a large amount of network traffic between servers, organizations should provide sufficient network bandwidth between servers and client computers.

For a medium-size environment of about 2,000 users, the Dell PowerEdge 2950 server can help meet these requirements, providing sufficient processing power to host 2,000 mailboxes while also providing expandable network and disk options. Organizations with larger environments can add additional mailbox servers to the back-end Exchange infrastructure. For more specific information, visit www.dell.com/exchange and use the Dell Exchange Advisor Tool. This automated tool asks simple questions and converts the answers into Dell-specific server, storage, and software recommendations.

Mailbox server security

Back-end Exchange server security is an important aspect of Exchange infrastructures, because these servers host critical data such as user mailboxes and public folders. Even with strong perimeter network security in place, the mailbox servers require protection against spam and viruses, which can enter the servers from such sources as Web mail, USB drives, and other removable storage media.

Symantec Mail Security for Microsoft Exchange software is designed to provide integrated mail protection against virus threats, spam, and other unwanted content. It enables administrators to inspect content in real time as e-mail is being committed to and accessed from the Exchange database. Attachment and subject-line blocking capabilities provide responses to known threats, and support for hourly definition updates enables organizations to respond quickly to emerging threats. Administrators can also conduct scheduled or on-demand scans to identify, detect, and quarantine inappropriate content or potential viruses.

High-availability clustering

Because the back-end Exchange servers host critical data, organizations should cluster these servers to help provide high availability. High-availability clustering enables the Exchange application to restart on another designated server in the cluster following a server failure, which can help ensure that the physical server hosting the Exchange mailboxes is not a single point of failure.

Dell high-availability clusters built using Microsoft Cluster Server software are part of the Dell Secure Exchange Reference Architecture and are designed to avoid a single point of failure.[1] High-availability clustering using Microsoft Cluster Server requires shared storage, because every node in the cluster needs access to the Exchange data.

Storage systems for Microsoft Exchange environments

The Dell Secure Exchange Reference Architecture allows storage to be shared over Fibre Channel–based SANs, which can help provide high bandwidth and low latency. When deployed with the emerging 4 Gbps Fibre Channel architecture, these SANs can provide sufficient bandwidth to support I/O-intensive applications. In addition, using redundant fabrics can provide multiple paths to the data and thereby help improve availability, and using logical units can help abstract the storage systems efficiently and eliminate physical dependencies between applications and data. Eliminating these dependences can also help improve storage system scalability, because storage can be added or removed without disrupting applications.

The Dell Secure Exchange Reference Architecture allows storage to be shared over Fibre Channel–based SANs, which can help provide high bandwidth and low latency.

Because Exchange is an I/O-intensive application, the disk system can potentially become a bottleneck. Using a large number of small-capacity drives instead of a small number of large-capacity drives can help improve disk subsystem performance. Organizations with environments of approximately 2,000 users typically should use external storage, such as direct attach storage or SANs. Dell/EMC Fibre Channel storage arrays offer highly scalable networked storage and advanced management, and can be well suited for Exchange mailbox stores. The Dell Exchange Advisor Tool can provide specific deployment recommendations.

Tape backup methods

The first line of defense in protecting e-mail data is usually to back up critical information to tape or disk. The Dell Secure Exchange Reference Architecture recommends using a SAN-based backup model; however, other methods can be used depending on the particular environment and requirements. As shown in Figure 2, the SAN-based backup model interconnects all messaging subsystem components—Exchange servers, master backup server, storage systems, and tape library—on the Fibre Channel SAN fabric. Data traffic can be routed from the Exchange servers through a high-speed Fibre Channel switch and written directly to the tape library.

[1] For more information about Dell high-availability clustering, visit www.dell.com/ha.

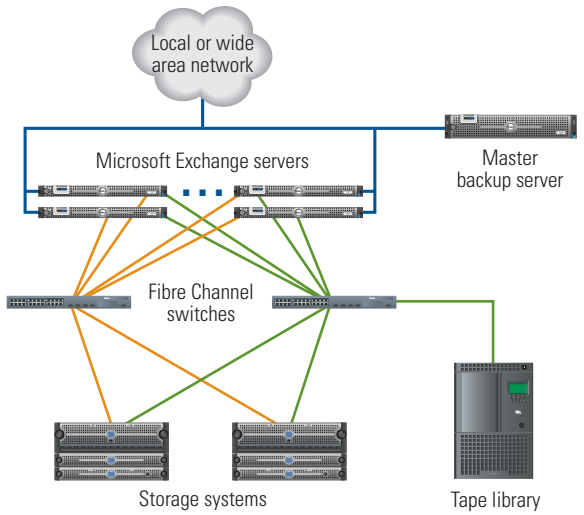

Figure 2. SAN-based tape backup model

After administrators have configured hardware and software components in the backup infrastructure and identified critical backup data, they must implement a backup strategy. Exchange works with any of the following backup methods or a combination of them: full backup, differential backup, incremental backup, and mirror backup.

Full backup. A full backup is designed to store all data, including Exchange database files and transaction logs. This approach helps simplify the recovery process because it saves all data files and transaction log files in a single backup session. However, a full backup consumes the most bandwidth and requires the most storage space of the four backup methods. Best practices therefore recommend organizations perform a full backup at regular intervals, but in conjunction with another backup method.

Differential backup. A differential backup contains only the Exchange transaction log files that have changed since the last full backup; the database files are not copied. Because all the transaction logs since the last full backup are necessary for a restore operation, circular logging cannot be enabled during a differential backup. Recovery requires both the last full backup and the last differential backup. If this method is used, best practices recommend that organizations perform a full backup at regular intervals and supplement it with daily differential backups.

Incremental backup. An incremental backup contains the Exchange transaction log files that have changed since the last full, differential, or incremental backup. Of these three methods, an incremental backup is the fastest, and thus it can be suitable for large Exchange databases with a high volume of daily activity. The drawback to the incremental approach is that recovery requires the last full backup and all subsequent incremental backups. If this method is used, best practices recommend that organizations perform a full backup at regular intervals and supplement it with daily incremental backups.

Mirror backup. A mirror backup is similar to a full backup, except that no file marking is performed. Mirror backups typically are not used for recovery. Organizations can use this method to make a full copy of the Exchange database without disrupting incremental or differential backup procedures.

Symantec Backup Exec

Dell PowerVault™ tape libraries and Symantec Backup Exec software together can help provide a reliable hardware and software platform to protect critical e-mail data against application- or hardware-based failures. The PowerVault tape libraries are designed to be scalable to help meet both current and future storage capacity requirements.

The Symantec Backup Exec Agent for Microsoft Exchange Server provides administrators with the tools to help enable fast, flexible, comprehensive protection of Exchange servers. Its customizable options allow users to perform individual mailbox or even individual message backup or restore, and the support for SAN-based backups can also help improve backup and recovery processes.

Archiving

The ever-growing volume of enterprise e-mail means that archiving should also be part of a comprehensive Exchange architecture, to help meet enterprise needs, technical requirements, or both. Archiving can also help enable simplified mailbox management, .pst migration, and rapid retrieval.

Symantec Enterprise Vault software provides a flexible archiving framework to enable the discovery of content in e-mail, file system, and collaborative environments while helping reduce storage costs and simplify management. Enterprise Vault manages content through policy-controlled archiving to online stores for active retention and seamless information retrieval. It provides powerful search and discovery capabilities to enable end users to access e-mail content. Dell and Symantec have worked together to validate and test Enterprise Vault on Dell servers and storage to help ensure interoperability and performance and provide an end-to-end Exchange solution.[2]

Standards-based architecture for secure, high-availability messaging

As Microsoft Exchange and messaging become pivotal enterprise applications, enterprise IT organizations must assess and implement an integrated strategy for mail security, data protection, and archiving. To safeguard business-critical workloads, industry leaders

[2] To learn more about Symantec Enterprise Vault, see the Yellow Book "Symantec Enterprise Messaging Management for Microsoft Exchange," by Symantec Corporation, ses.symantec.com/YB_EMMMSE.

should collaborate on providing secure Exchange implementations for any size organization—including small and medium businesses, government entities, and commercial enterprises. To that end, Symantec, Microsoft, and Dell have partnered to deliver the Dell Secure Exchange Reference Architecture.

Dell believes that standardization is the key to continued data center evolution. All the components of the Dell Secure Exchange Reference Architecture are based on industry standards and are designed for scalability, to accommodate organizational changes. Enterprise IT organizations can also add or remove capacity without having to change the base messaging system architecture. Using this reference architecture can help IT organizations design a highly available Exchange infrastructure incorporating both security and archiving within a comprehensive messaging environment.

Suman Kumar Singh is a systems engineer in the High-Availability Systems Group at Dell. He specializes in messaging systems architecture and sizing. His other interests include storage area networks, virtualization, and security. He has published and presented several papers at industry conferences.

Bharath Vasudevan currently manages the High-Availability Cluster Group at Dell. He has previously designed server hardware and served as a lead engineer for multiple cluster releases. His current interests include application performance characterization and storage technologies. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

FOR MORE INFORMATION

Dell Secure Exchange:
www.dell.com/secure_exchange

Dell and Microsoft Exchange:
www.dell.com/exchange

Dell and Symantec:
www.dell.com/symantec

Microsoft Exchange:
www.microsoft.com/exchange

Symantec enterprise products:
www.symantec.com/enterprise

Symantec Yellow Books:
www.symantec.com/enterprise/yellowbooks/index.jsp

Strengthening Communications with Dell Secure Exchange

As the scope of business communication expands, Dell greets the challenges of a changing messaging infrastructure with comprehensive solutions featuring best-of-breed hardware, software, and services.

No doubt about it—ubiquitous business messaging has become imperative for successful company performance, reputation, and profitability. And many businesses are now demanding rock-solid messaging infrastructures with advanced capabilities to help keep information flowing securely, efficiently, and cost-effectively.

Dell can help enterprises meet this challenge by assisting in the deployment of or migration to Microsoft® Exchange—recognized as an industry-leading platform for e-mail messaging and collaboration. Dell has amassed vast experience in this arena, completing over 1,000 Exchange migration projects and managing more than 500 major enterprise migrations and consolidations per year. Thus far, Dell has migrated over 5 million Exchange user mailboxes—including 450,000 users for one of the world's five largest companies.

Dell has expanded its offerings to comprehensively address today's dynamic communications needs. Based on high-powered Dell™ hardware and the Microsoft Exchange platform, Dell Secure Exchange solutions leverage best-of-breed Symantec software and the unique Dell Secure Exchange Reference Architecture to deliver cutting-edge environments for security, data archiving, business continuity, and remote assistance.[1] The result is a robust, end-to-end solution that extends the functionality of the Exchange infrastructure with enhanced interoperability, performance, and availability. To assess readiness for these Dell Secure Exchange solutions and then design, implement, and support them, Dell provides various types of Exchange-specific services. Following are highlights of the services offered with Dell Secure Exchange solutions.

Exchange Security

Without stringent messaging security in place, enterprises face serious risks when employees communicate across public and private networks—including a rampant influx of viruses and spam that can lead to data loss, downtime, clogged networks, fraud, and theft. The Exchange Security service involves developing a custom design for the Dell Secure Exchange solution and then installing, configuring, and tuning the solution. The Dell Services team also provides knowledge transfer and documentation.

The Exchange Security service leverages Symantec Mail Security software or Symantec Mail Security 8200 series appliances to mitigate risks and maintain the integrity of vital business communications. In conjunction with Dell server hardware and the Exchange platform, these Symantec products apply over 20 industry-leading antispam technologies that can help prevent costly downtime from viruses or overloaded networks and mailboxes. And Symantec Mail Security 8200 series appliances are designed to offer near-pinpoint accuracy, effectively detecting 95 percent of spam messages and avoiding false positives approximately 99 percent of the time.[2]

Content-compliance features, which monitor inbound and outbound e-mail content according to customer-defined policies, can help further reduce virus threats, while innovative e-mail firewall technologies restrict unwanted connections from spam-sending servers. And automatic spam filter and virus definition updates help simplify management for ongoing security to help minimize overall infrastructure costs.

Exchange Archiving and Mailbox Management

In the post-Enron world, many businesses now are challenged by complex data archiving requirements to comply with stringent government and industry regulations. These regulations act as watchdogs over records management and storage processes—and they can be very useful as the volume of data

[1] For more information, see "Implementing the Dell Secure Exchange Reference Architecture," by Suman Kumar Singh and Bharath Vasudevan, *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20060452-Singh.pdf.

[2] "Symantec Mail Security 8200 Series: Feature Summary," by Symantec Corporation, eval.veritas.com/mktginfo/enterprise/other_resources/ent-other_resources_mailsecurity8200_092005.en-us.pdf.

employees produce and transfer over networks continues to grow. The reasoning is simple: without archive-friendly infrastructures built to support information throughout its life cycle, businesses run the risk of fraud accusations, liability lawsuits, and government fines—and executives face the possibility of imprisonment.

Under the Exchange Archiving and Mailbox Management service, Dell professionals walk customers step-by-step through the design and configuration of an archiving solution featuring Symantec Enterprise Vault software. The Dell Services team analyzes customer archiving requirements; maps the requirements to Enterprise Vault policy and configuration settings; installs, configures, and tunes the Dell Secure Exchange solution; and provides knowledge transfer and documentation.

With tailored policies for managing and maintaining content held within e-mail, file system, and collaborative environments—along with specialized client applications for corporate governance, risk management, and legal protection—the Exchange Archiving and Mailbox Management service effectively promotes regulatory compliance. Easy and rapid search and retrieval of content lets users tap into organizational knowledge. And the storage optimization provided by Enterprise Vault software is designed to reduce Exchange message stores by 50 percent or more compared to implementations not using Enterprise Vault software.[3]

Furthermore, enterprises can take this service beyond compliance to achieve core business benefits as well. For example, simplified processes for archiving and accessing critical information can help decrease costs associated with both storage and management. Moreover, the stable environment achieved by combining Dell hardware, the Exchange platform, and Symantec software enables companies to turn the data archive into a functional "information warehouse" that can be mined as a knowledge resource using built-in index and search technologies.

Exchange Backup and Recovery

Often equally as important to compliance as data archiving, effective backup and recovery procedures provide a means of protecting information throughout its life cycle—from cradle to grave. When companies lack the infrastructure to support reliable backup and recovery procedures, they risk unplanned downtime, service interruption, breached content integrity, or even permanent data loss.

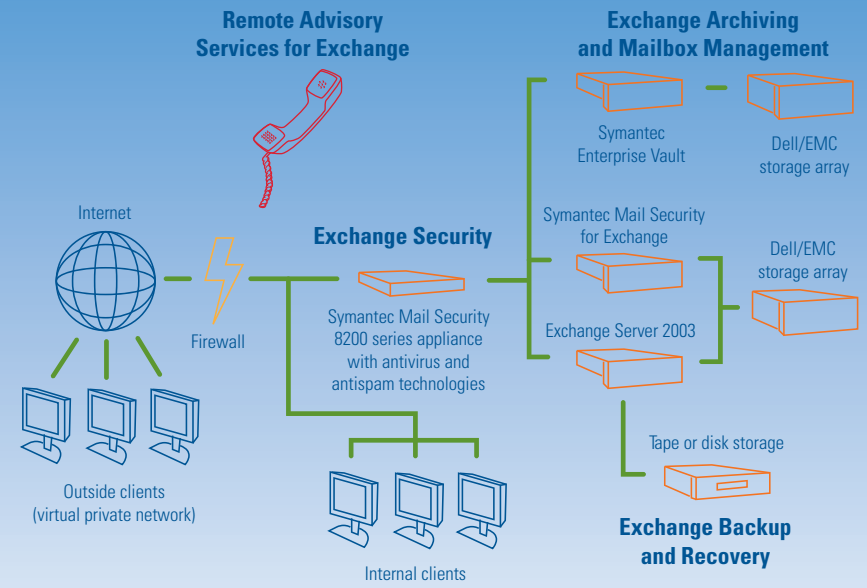

Deploying Dell Secure Exchange across the enterprise

The Exchange Backup and Recovery service provides a comprehensive solution design, assessment of any interactions with server applications, a definition of tape usage and strategy, and validation of the storage area network (SAN). This service also includes assessments of design and site readiness as well as solution implementation and testing—including configuration of the tape library and library resources and installation of applicable software. The Dell Services team then conducts a product orientation session and provides the associated documentation.

With the Exchange Backup and Recovery service, Dell experts deploy Symantec Backup Exec 10d for Windows Servers to provide comprehensive backup and recovery for servers, desktops, and laptops running Microsoft Windows® operating systems. Complete disk-to-disk-to-tape protection, along with centralized administration and SAN support, promote scalable management of distributed backup and remote servers. Additionally, high-performance agents and options provide fast, flexible, and granular protection of servers.

Because Dell Secure Exchange solutions are designed to simplify data management, eliminate backup windows, and decrease the risk of downtime, they can help enterprises rein in costs and focus administrative resources on business-critical tasks.

Remote Advisory Services for Exchange

Rounding out Dell Secure Exchange solutions are the Remote Advisory Services for Exchange, which provide expert how-to support for many common tasks—such as adding user mailboxes, changing security settings, or increasing mailbox size. Developers simply call in with questions and receive personal, professional Dell expertise.

Remote Advisory Services are available during normal business hours (which vary by customer and time zone). Alternatively, enterprises can select a "scheduled service" option, where Dell professionals assign and prioritize issues on a case-by-case basis. By facilitating remote assistance, Dell helps streamline management of the messaging infrastructure to help enterprises further cut administrative costs and reduce time spent on troubleshooting or support.

Comprehensive solution for complete communications

Bolstered by comprehensive services, Dell Secure Exchange solutions can help enterprises migrate to the appropriate Exchange platform, implement customized end-to-end messaging services, and enjoy the advantages of high-performance Dell hardware.

To learn more about Dell Secure Exchange, visit www.dell.com/secure_exchange.

[3] "Veritas Enterprise Vault for Microsoft Exchange," by Symantec Corporation, eval.veritas.com/mktginfo/products/Datasheets/Data_Protection/ent-datasheet_ent_vault_microsoft_exchange_v6_02-2006.en-us.pdf.

Deploying Microsoft Windows Compute Cluster Server 2003

on Dell PowerEdge Servers

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

BY RON PEPPER AND VICTOR MASHAYEKHI, PH.D.

Related Categories:

Dell PowerEdge servers

High-performance computing (HPC)

Microsoft Windows Compute Cluster Server 2003

Microsoft Windows Server 2003

System deployment

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

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

CCS includes two components: the Windows Server 2003 Compute Cluster Edition OS and the Compute Cluster Pack (CCP). The Compute Cluster Edition is a limited version of Windows Server 2003 that does not allow some server services to function.[1] The CCP contains the necessary components to create server clusters, along with a job scheduler, MPICH MPI library, and Microsoft Windows Remote Installation Services (RIS) extensions.

Configuring cluster hardware

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

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

[1] For more information about these limitations, see the Windows Server 2003 Compute Cluster Edition end-user license agreement.

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

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

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

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

Preparing the head node for Compute Cluster Pack installation

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

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

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

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

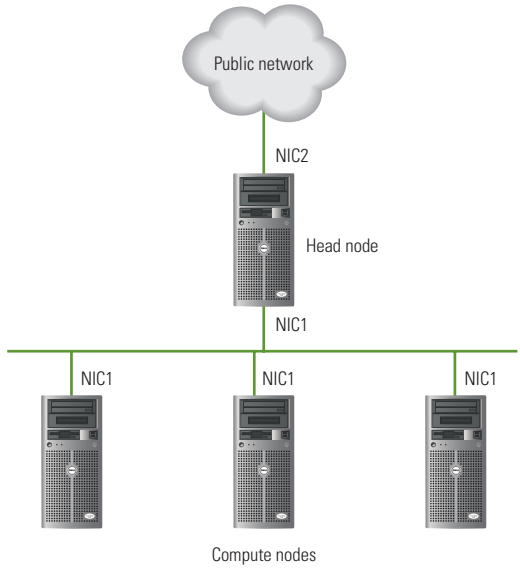

Figure 1. Example Microsoft Windows Compute Cluster Server 2003–based cluster configuration

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

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

Installing the Compute Cluster Pack

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

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

Configuring the Compute Cluster Pack

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

[2] For more information about installing a domain controller, visit www.microsoft.com/technet/prodtechnol/windowsserver2003/technologies/directory/activedirectory/stepbystep/domcntrl.mspx.

Networking task section

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

The Manage Firewall Settings wizard enables or disables public network firewall settings, which should typically be enabled. Administrators can later provide firewall access to individual services as needed.

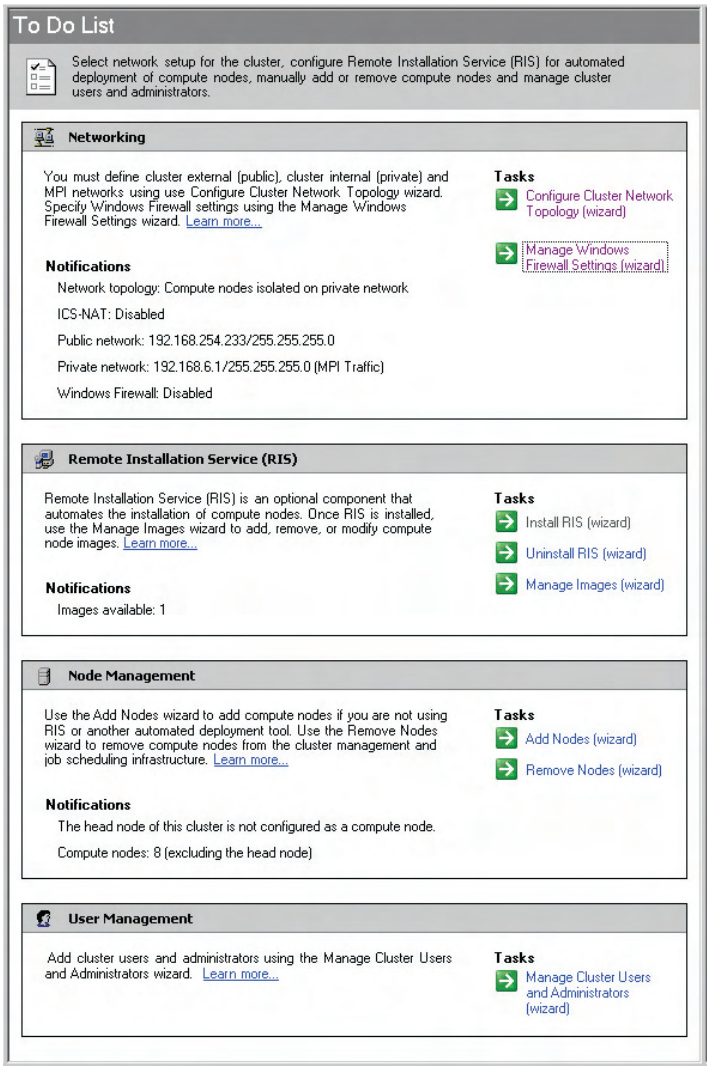

Figure 2. To Do List screen following Compute Cluster Pack installation

RIS task section

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

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

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

Node Management task section

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

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

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

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

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

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

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

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

    Add the following line to the `Unattended` section:

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

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

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

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

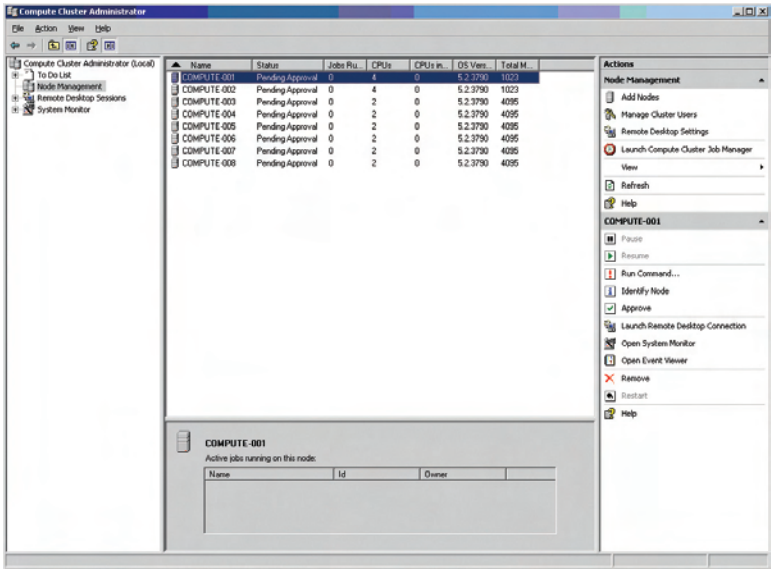

Figure 4. Compute Cluster Administrator Node Management screen showing newly installed compute nodes pending approval

User Management task section

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

Approving installed compute nodes

As a final step before the cluster can run jobs, administrators must launch Compute Cluster Administrator and select "Node Management." They can then approve and un-pause the newly installed cluster compute nodes by selecting the compute nodes from the list and clicking "Approve" and "Resume" in the Actions pane (see Figure 4).

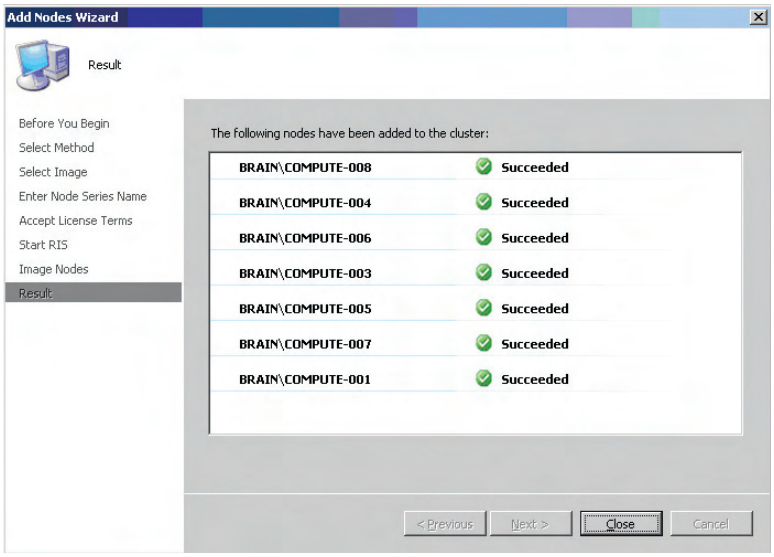

Figure 3. Result screen following completion of the Compute Cluster Pack Add Nodes wizard

Enabling simplified cluster installation and management

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

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

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

Servers

Storage

Systems Management

Services

Software

ORACLE

How Dell Does IT

Dell™ understands that mission critical databases require powerful performance, highly reliable and available hardware platforms, and scalability that allows you to grow with your database needs. Dell understands this so well we run our own supply chain management system, which supports hundreds of suppliers and manages nearly a billion parts per year, on powerful and scalable Dell PowerEdge™ servers running Oracle® Database 10*g*. Dell offers pre-engineered, tested, and validated solutions for Oracle Database 10*g* on Linux® and Windows®. The entire solution stack is tested and supported by Dell – the servers, storage, switches, and software, including the Oracle Database software and the operating system – and Dell offers services to help you accelerate deployment.

DELL

Go to dell.com/oraclemag to get the full story on How Dell IT uses Oracle Database 10*g* in our supply chain management system.

Evaluating Scalability and Power Benefits

of Ninth-Generation Dell PowerEdge Servers in an HPC Environment

Energy efficiency and scalability have become increasingly important to many enterprises. This article discusses the benefits of Intel® Xeon® 51*xx* processors for high-performance computing cluster environments by comparing performance/watt and cluster scalability results using compute-intensive applications.

BY RIZWAN ALI; BARIS GULER; RAMESH RADHAKRISHNAN, PH.D.; AND VISHVESH SAHASRABUDHE

Related Categories:

Characterization

Dell ninth-generation servers

Dell PowerEdge servers

High-performance computing (HPC)

HVAC

Performance

Power consumption

System deployment

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

Ninth-generation Dell™ PowerEdge™ 1950 servers use Intel Xeon 50*xx* processors based on the Intel NetBurst® microarchitecture and Intel Xeon 51*xx* processors based on the Intel Core™ microarchitecture. The PowerEdge 1950 uses the Intel 5000X chipset with dual frontside bus (FSB), quad-channel fully buffered dual in-line memory modules (DIMMs), and PCI Express I/O architecture. The Intel Core microarchitecture combines the energy-efficient philosophy of the Intel mobile microarchitecture with the existing Intel NetBurst microarchitecture, and adds a number of significant performance innovation designs. This microarchitecture is designed to optimize the performance, energy efficiency, and scalability of multi-core processors.[1]

The high-performance computing (HPC) community is becoming increasingly aware of the impact that power consumption has on the total cost of an HPC cluster. Operational costs incurred in running the cluster and cooling it in a large room, or maintaining specialized buildings to prevent system failures, increase as server power consumption levels rise.

The Intel Core microarchitecture provides an interesting design choice for HPC applications. A team of Dell engineers in July 2006 used HPC benchmarks—Linpack; NAS (NASA Advanced Supercomputing) Parallel Benchmarks (NPB); MIMD (multiple instruction, multiple data) Lattice Computation (MILC); FLUENT; and OOCORE—to measure performance/watt

[1] For more information about the Intel Core microarchitecture, visit www.intel.com/technology/architecture/coremicro/index.htm.

Servers	One Dell PowerEdge 1850 server	One Dell PowerEdge 1950 server
Processors	• Intel Xeon processor at 3.6 GHz with 1 MB cache and 800 MHz FSB (Nocona) • Intel Xeon processor at 3.6 GHz with 2 MB cache and 800 MHz FSB (Irwindale) • Dual-core Intel Xeon 7030 processor at 2.8 GHz with two 2 MB caches and 800 MHz FSB (Paxville)	• Dual-core Intel Xeon 5080 processor at 3.73 GHz with two 2 MB L2 caches and 1,333 MHz FSB (Dempsey) • Dual-core Intel Xeon 5150 processor at 2.67 GHz with 4 MB L2 cache and 1,333 MHz FSB (Woodcrest)
Memory	Four 1 GB PC2-3200 double data rate 2 (DDR2)–400 DIMMs	Four 1 GB PC2-5300 fully buffered DIMMs
Disk	SCSI	Serial Attached SCSI (SAS)
OS	Red Hat® Enterprise Linux® 4 Update 2 OS with support for Intel Extended Memory 64 Technology (EM64T)	
Compilers	Intel C, C++, and Fortran Compilers version 9.1, build 20060323 (used where source code available)	

Figure 1. Performance/watt test configuration

Servers	Four Dell PowerEdge 1950 servers (dual-socket, dual-core)
Processor	Intel Xeon 5150 processor at 2.67 GHz with 4 MB L2 cache and 1,333 MHz FSB
Interconnects	24-port Dell PowerConnect™ 5324 Gigabit Ethernet switch; Cisco SFS 7000 InfiniBand switch with PCI Express–based MemFree InfiniBand host channel adapters
Message Passing Interface (MPI)	MVAPICH (Cisco driver version 3.2.0-67); MPICH version 1.2.6
Memory	PC2-5300 fully buffered DIMMs from various vendors (each server had 4 GB of RAM)
Disk	SAS
OS	Red Hat Enterprise Linux 4 Update 2 OS with support for Intel Extended Memory 64 Technology (EM64T)
Compilers	Intel C, C++, and Fortran Compilers version 9.1, build 20060323 (used where source code available)

Figure 2. Scalability test cluster configuration

and cluster scalability of Dell PowerEdge 1950 servers using Intel Xeon 51*xx* processors.

Benchmarks for evaluating cluster performance

The following synthetic and application benchmarks were used in the Dell tests. These benchmarks and applications represent a broad spectrum of HPC workloads.

Linpack. This is a popular benchmark for HPC environments. The High-Performance Linpack (HPL)[2] implementation is commonly used to rank supercomputers on the TOP500 Supercomputer Sites list.

NAS Parallel Benchmarks. NPB is an application-centric suite of benchmarks that has been widely used to measure and compare the performance of parallel-processing computers.[3] The Dell team used the IS (Integer Sort) and LU (Lower-Upper Diagonal) B Class programs, and then calculated the sum of the IS and LU results to evaluate performance/watt.

MILC. The code developed by the MILC Collaboration is used in high-energy physics for simulations of 4-D special unitary (SU) lattice gauge theory on MIMD parallel-processing systems.[4]

FLUENT. A popular computational fluid dynamics (CFD) application suite, FLUENT is commonly used in HPC environments. The FLUENT applications allow users to perform CFD analysis around their particular models.[5] Several benchmark data sets (workloads) available from Fluent Inc. were used in the Dell tests.

OOCORE. An out-of-core matrix solver, OOCORE handles matrix equations that are too large for the cache. This benchmark writes large amounts of data to the disk and thus also tests the disk I/O performance of the server.[6]

Test environment for Dell HPC cluster

The performance/watt test environment was based on a Dell PowerEdge 1850 server and PowerEdge 1950 server using different Intel Xeon processors. The scalability test environment was based on a cluster comprising four Dell PowerEdge 1950 servers using non-blocking Gigabit Ethernet and InfiniBand interconnects. The cluster was installed using Platform Open Cluster Stack (OCS).[7] Figures 1 and 2 describe the hardware and software the Dell team used for testing.

Results and analysis of performance/watt and cluster interconnect tests

The study focused on two sets of tests: one comparing performance/watt for Intel Xeon 51*xx* processors against previous-generation Intel

[2] For more information about HPL, visit www.netlib.org/benchmark/hpl.

[3] For more information about NAS Parallel Benchmarks, visit www.nas.nasa.gov/Software/NPB.

[4] For more information about the MILC code, visit www.physics.utah.edu/~detar/milc/milcv6.html.

[5] For more information about FLUENT, visit www.fluent.com/software/fluent/index.htm.

[6] For more information about OOCORE, visit www.nsf.gov/pubs/2006/nsf0605/nsf0605.jsp.

[7] For more information, see "Platform Rocks: A Cluster Software Package for Dell HPC Platforms," by Rizwan Ali, Rinku Gupta, Garima Kochhar, and Bill Bryce, *Dell Power Solutions,* November 2005, www.dell.com/downloads/global/power/ps4q05-20050227-Ali.pdf.

Xeon processors, and one comparing cluster performance using Woodcrest processors with Gigabit Ethernet interconnects against InfiniBand interconnects.

Performance/watt tests

The tests compared performance/watt for eighth- and ninth-generation Dell PowerEdge servers using different types of Intel Xeon processors. The processor cores in eighth-generation servers are based on the Intel NetBurst microarchitecture. The Intel Xeon 50*xx* (Dempsey) processors used in ninth-generation servers are also based on the Intel NetBurst microarchitecture, which allows for higher processor frequencies compared to previous architectures while incurring a high penalty on power consumption. The Intel Xeon 51*xx* (Woodcrest) processors, which are also used in ninth-generation servers, are based on the Intel Core microarchitecture, which is designed to provide higher performance and lower power consumption compared to the Intel NetBurst microarchitecture.

Figure 3 compares the measured performance/watt on eighth- and ninth-generation Dell servers using these different processor architectures. A higher performance/watt metric indicates that the system is providing higher performance while consuming less power. The baseline is an eighth-generation server (PowerEdge 1850) using the Intel Xeon Nocona processor at 3.6 GHz. The next two bars show performance/watt for an eighth-generation server using the Intel Xeon Irwindale processor and the dual-core Intel Xeon 7030 Paxville processor, respectively; the remaining two bars show performance/watt for a ninth-generation server (PowerEdge 1950) using the dual-core Dempsey and Woodcrest processors, respectively.

The Paxville processor exhibited significant improvement over the Nocona and Irwindale processors because of the additional processing cores designed to enhance performance. The Dempsey processor, which is based on the same microarchitecture as the Nocona, Irwindale, and Paxville processors, did not provide any performance gains compared to the Paxville processor except with the NPB benchmark. However, performance/watt for the Woodcrest processor was significantly higher compared with all other processors—the increase ranged from 57 percent (FLUENT) to 138 percent (OOCORE) compared with the Dempsey processor—making it the most energy-efficient choice for HPC cluster environments among the tested processors.

Scalability tests

The first set of tests focused on the performance and energy efficiency derived from the Intel Core microarchitecture; the second set focused on the scalability of the microarchitecture when the same HPC benchmarks were run in a cluster using Gigabit Ethernet and InfiniBand interconnects. The testing environment is described in Figure 2.

To obtain optimal performance from each cluster node, all the physical processors in each server were utilized when running the benchmarks. Each benchmark was run using 4, 8, and 16 processors in one, two, and four servers, respectively. Figure 4 shows the results when running the benchmark with these configurations using either Gigabit Ethernet or InfiniBand interconnects. The 1 × 4 (one node with four processors) configuration results serve as the baseline, and the results for the 2 × 4 and 4 × 4 configurations are shown as relative speedups.

In Figure 4, the set of lines on the left shows the scalability of the benchmarks when Gigabit Ethernet was used as the internode communication fabric. As can be seen from these scalability lines, benchmarks bound by computation and local I/O but less dependent on internode communication scaled better than communication-intensive benchmarks. For example, the HPL and FLUENT (large) benchmarks, which are highly computation intensive and have less internode communication, scaled well. On the other hand, the NPB-IS benchmark involves intensive all-to-all communications with a mixture of small and very large messages. NPB-IS is sensitive to communication latency and bandwidth;

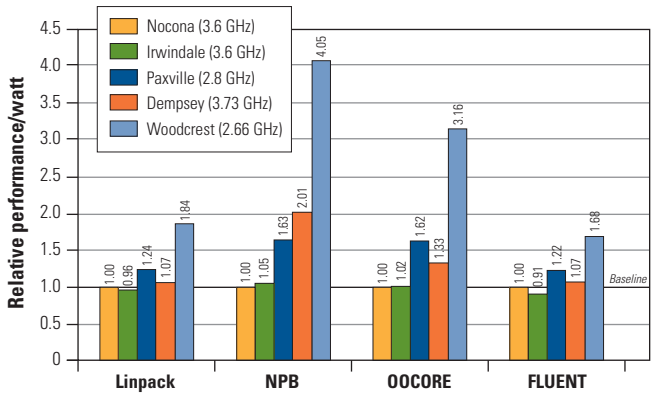

Figure 3. Relative performance/watt of Intel Xeon processors on eighth- and ninth-generation Dell PowerEdge servers

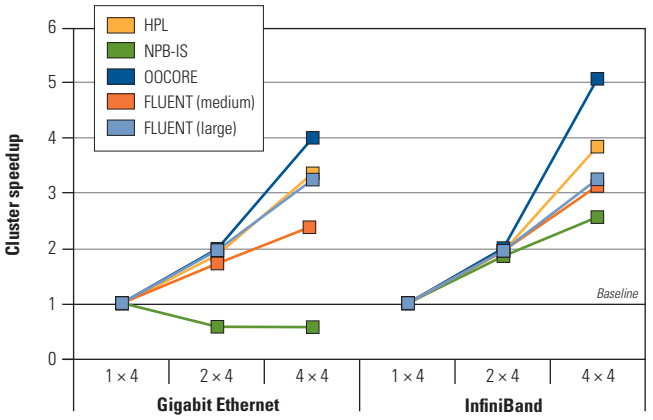

Figure 4. Cluster speedup using Gigabit Ethernet and InfiniBand interconnects

therefore, it experienced a performance degradation when run in a Gigabit Ethernet cluster.

The set of lines on the right in Figure 4 shows the scalability of the same set of benchmarks when InfiniBand was used as the internode communication fabric. NPB-IS, which experienced a degradation in performance when using Gigabit Ethernet, scaled up to approximately 2.5 times the baseline performance when using InfiniBand in a 4 × 4 configuration. In addition, scalability for every other benchmark was better when using InfiniBand as the interconnect as compared to Gigabit Ethernet. An interesting observation is that OOCORE exhibited superscalar performance on the 4 × 4 configuration. The problem size for all node configurations was the same, and hence the amount of data written to disk decreased significantly when run on four nodes. Thus, OOCORE benefited not only from improved interconnect performance but also from the decreased local I/O traffic on each node, leading to superscalar performance improvement.

As Figure 4 shows, the benchmarks scaled better on a cluster with InfiniBand fabric than on one using Gigabit Ethernet. Moreover, as Figure 5 shows, the actual performance of those benchmarks for the same number of nodes was also better when run using InfiniBand. This figure compares the performance of the FLUENT (large and medium), HPL, OOCORE, and MILC benchmarks for the 4 × 4 configuration running over InfiniBand relative to the results for the same problem size running over Gigabit Ethernet. The results for the FLUENT (large) benchmark showed a slight degradation. As stated earlier in this section, this is a compute-intensive benchmark, and hence the performance benefit from using a high-bandwidth, low-latency interconnect is not readily apparent. All the other benchmarks experienced a performance improvement with InfiniBand as the cluster fabric, with OOCORE and MILC showing the largest improvement. When communication-sensitive benchmarks run in a cluster larger than this 4 × 4 configuration, the benefits of InfiniBand over Gigabit Ethernet should be more apparent.

The Intel Xeon 51*xx* processors employing the Intel Core microarchitecture are designed to enable significant power savings while providing a higher level of performance compared to previous-generation processors.

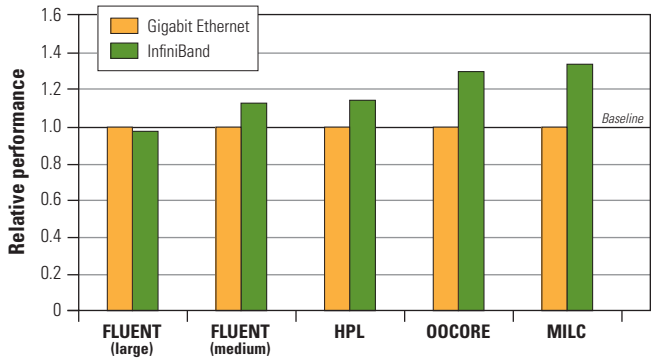

Figure 5. Relative performance of Gigabit Ethernet and InfiniBand interconnects for a four-node cluster

Changing landscape for HPC

For decades, the focus when choosing an HPC cluster has been on performance and occasionally price/performance. However, the primary focus for enterprise IT organizations has become a substantial reduction in HPC system power consumption to provide the best performance/watt possible. As is evident from the tests described in this article, the Intel Xeon 51*xx* processors employing the Intel Core microarchitecture are designed to enable significant power savings while providing a higher level of performance compared to previous-generation processors. The enhanced scalability afforded by the Intel Core microarchitecture and the increase in the number of available processors per server necessitate using high-bandwidth, low-latency interconnects such as InfiniBand for communication-sensitive applications.

Rizwan Ali is a member of the Scalable Systems Group at Dell. His current research interests include performance benchmarking, cluster architecture, parallel applications, and high-speed interconnects. He has a B.S. in Electrical Engineering from the University of Minnesota.

Baris Guler is an HPC application specialist in the Scalable Systems Group at Dell. His current research interests are parallel processing, diskless HPC clusters, performance benchmarking, reservoir engineering and simulation, and numerical methods. Baris has a B.S. in Petroleum and Natural Gas Engineering (PNGE) from the Middle East Technical University in Turkey and an M.S. in PNGE from Pennsylvania State University.

Ramesh Radhakrishnan, Ph.D., is a member of the Scalable Systems Group at Dell. His interests include performance analysis and characterization of enterprise-level benchmarks. Ramesh has a Ph.D. in Computer Engineering from the University of Texas at Austin.

Vishvesh Sahasrabudhe is a member of the Scalable Systems Group at Dell. His current research interests include high-speed interconnects and performance benchmarks. He has a B.Tech. in Electrical Engineering from the Indian Institute of Technology in Bombay and an M.S. in Computer Science and Engineering from the Ohio State University.

Secure HPC Cluster Management

with Ninth-Generation Dell PowerEdge Servers

The increasing use of high-performance computing (HPC) clusters for critical or sensitive high-performance computations has created a need for secure node management and monitoring capabilities. This article discusses the available tools for creating a secure HPC environment with the Dell OpenManage™ suite and ninth-generation Dell™ PowerEdge™ servers.

BY ARUN RAJAN, TONG LIU, YUNG-CHIN FANG, AND SAEED IQBAL, PH.D.

Related Categories:

Cluster management

Clustering

Dell OpenManage

Dell ninth-generation servers

Dell PowerEdge servers

High-performance computing (HPC)

Intelligent Platform Management Interface (IPMI)

Security

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

High-performance computing (HPC) clusters require secure management to protect against internal and external misuse or abuse of compute resources. In the past, HPC cluster security typically was not a significant problem because clusters were built for private use by a handful of dedicated users. However, many enterprises are now using large-scale clusters that are often shared across departments with easy public access—and increased vulnerability has become an issue. Vulnerable features of HPC clusters include high-bandwidth connections (which can be used to launch denial-of-service attacks on other sites, for example), massive computational power (which can be used to execute parallelized password-cracking tools), and extensive storage (which can be hacked and used to save copyrighted or illegal information).[1]

Because of this increased vulnerability, substantial efforts have been made to develop HPC management and monitoring tools. One way to provide security in an HPC cluster environment is to enforce secure communication—that is, only nodes and users that are meant to access certain resources should be allowed access. Secure communication can be provided by adhering to secure protocols, providing firewall-like features such as port blocking, allowing role-based authentication for individual users, and providing credible login authentication procedures and sufficient data encryption. Dell HPC cluster environments offer

[1] For more information about cluster security, see "Cluster Security as a Unique Problem with Emergent Properties: Issues and Techniques," by William Yurcik, Gregory A. Koenig, Xin Meng, and Joseph Greenseid, 5th Linux Clusters Institute Conference, May 17–20, 2004, www.linuxclustersinstitute.org/Linux-HPC-Revolution/Archive/PDF04/07-Yurcik_W.pdf.

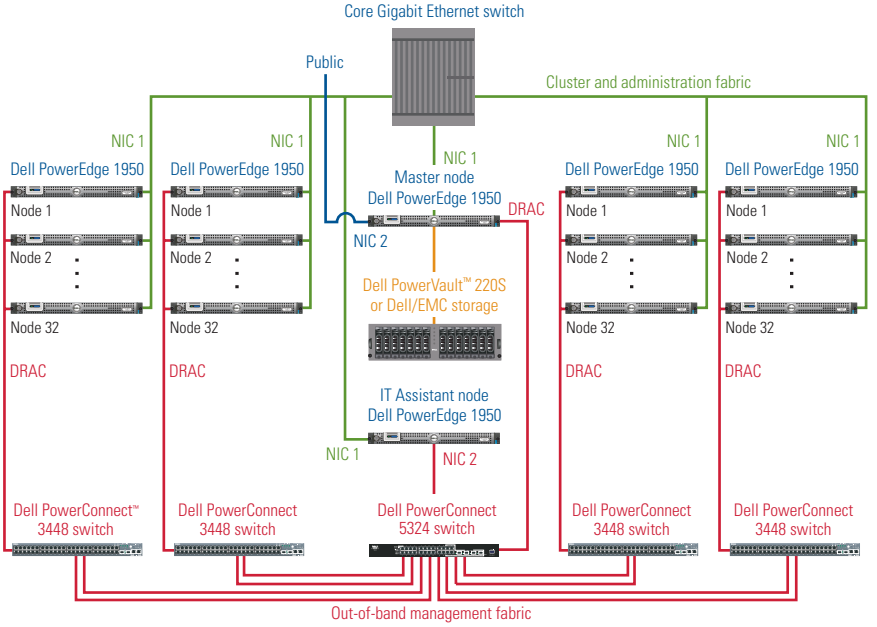

Figure 1. Dell HPC cluster architecture

such measures to help provide high levels of security while utilizing a wide range of management components.

Using ninth-generation Dell PowerEdge servers in HPC clusters

Dell HPC clusters can comprise both eighth- and ninth-generation Dell PowerEdge servers. Ninth-generation PowerEdge servers introduce the Intel® 5000X and 5000P chipsets with 1,066 MHz and 1,333 MHz dual frontside buses, enhanced processor-memory bandwidth, quad-channel double data rate 2 (DDR2) fully buffered dual in-line memory modules (DIMMs), and PCI Express I/O architecture. These servers use the dual-core Intel Xeon® 5*xxx* processors, which are based on the Intel Core™ microarchitecture.

The major hardware management components in ninth-generation Dell PowerEdge servers are the Dell Remote Access Controller 5 (DRAC 5) and the baseboard management controller (BMC). The DRAC 5 is used for out-of-band management and provides security features for its command-line interface (CLI) and Web browser–based graphical user interface (GUI) by using Secure Sockets Layer (SSL). The BMC and the associated BMC Management Utility are compliant with the Intelligent Platform Management Interface (IPMI) 2.0 specification with enhanced Remote Management Control Protocol + (RMCP +) authentication.

Figure 1 shows a typical Dell HPC cluster architecture with public (external) access through network interface card 2 (NIC2), the out-of-band management channel through the DRAC 5 fabric, and the in-band management channel through the cluster and administration fabric over NIC1. NIC1 also supports out-of-band management through the BMC. Figure 2 shows an overview of the security protocols used in the Dell HPC management infrastructure.

Tools for managing Dell HPC cluster nodes

The Dell OpenManage suite enables administrators to monitor and manage Dell PowerEdge servers remotely, streamlining node management in HPC clusters. This suite includes several components: Dell OpenManage IT Assistant, Dell OpenManage Server Administrator, Dell OpenManage Storage Services, the BMC Management Utility, the Dell Update Package, the Software Update Utility, the Dell OpenManage Deployment Toolkit, Dell OpenManage Server Assistant, and Dell OpenManage Online Diagnostics. Open source products such as OpenIPMI, IPMItool, Ganglia, and Cluster Monitoring (Clumon) are also available to manage Dell HPC clusters.

The Platform Open Cluster Stack (OCS) software stack—based on the San Diego Supercomputer Center (SDSC) National Partnership for Advanced Computational Infrastructure (NPACI) Rocks—is a comprehensive HPC cluster toolkit, designed to simplify the deployment and management of large-scale Linux® OS–based clusters. Developed by Platform Computing, Platform OCS can be used along with Dell OpenManage to provide comprehensive node management capabilities—for example, to configure cluster nodes for remote management through IPMI and console redirection.

Implementing secure change management with SNMP

Simple Network Management Protocol (SNMP) is a systems management standard originally designed for network management. It

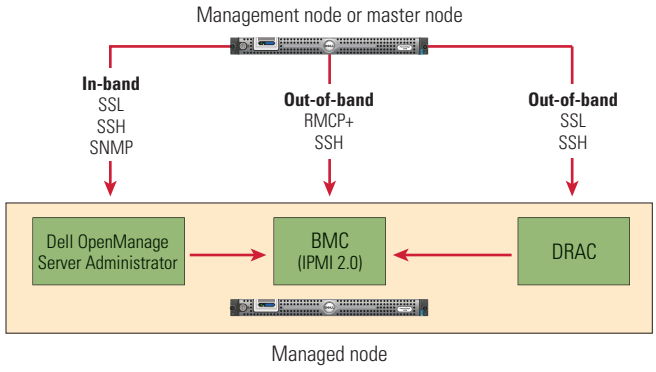

Figure 2. Security protocols for Dell HPC management

is an application-layer protocol that is part of the TCP/IP protocol suite and facilitates the exchange of management information between network devices. SNMP standards are defined by the Internet Engineering Task Force.

SNMP-managed systems provide data to a management system through SNMP agents. At each SNMP agent, a community string is configured that is transmitted as part of all SNMP request messages. Using these community strings means that the requester cannot be verified as part of the specified community, resulting in limited security. Because Get requests are noninvasive, limited security levels can be permitted. However, this limited security is insufficient for Set operations, which can result in intrusive actions such as RAID reconfiguration, system power down, and so on; for this reason, certain vendors do not support SNMP Set operations at all. Dell OpenManage Server Administrator supports SNMP Set operations in a secure manner by forcing its SNMP agents to implement a hash/digest mechanism to help prevent unauthorized operations; however, only Server Administrator SNMP management applications support this mechanism. In an HPC cluster, Server Administrator is usually used to monitor and diagnose remote node health and support OS-level services such as storage management.

Secure communication can be provided by adhering to secure protocols, providing firewall-like features such as port blocking, allowing role-based authentication for individual users, and providing credible login authentication procedures and sufficient data encryption.

The management information base (MIB) is a type of database used to store management data, and is part of a standard that describes the managed objects contained in the MIB. For example, the MIB provides management data, and the Cluster Group defines attributes such as the number of systems in the cluster and the cluster's capabilities, type, and name. Server Administrator provides various MIBs for different purposes, such as the Instrumentation MIB for instrumentation data, the Remote Access MIB for in-band information on remote hardware, and the Change Management MIB for monitoring the inventory of devices and applications with SNMP management applications. The SNMP master transmits and receives SNMP requests, and the SNMP extension agent registers MIB objects. SNMP traps (asynchronous events generated when some significant action has taken place) are used to preserve network bandwidth, by preventing the need for frequent polling over critical parameters. In an HPC cluster, MIB implementations provide interoperable manageability among various management frameworks.

Transitioning to IPMI 2.0

IPMI, a joint effort promoted by Dell, Intel, Hewlett-Packard, and NEC, is a hardware-level interface specification that defines a common, abstracted, message-based interface for platform monitoring and control functions. IPMI defines common interfaces to intelligent hardware used to monitor a server's physical health characteristics, such as temperature, voltage, fans, power supplies, and chassis. Since IPMI 1.0 was introduced in 1998, more than 170 companies across the industry have adopted IPMI, including system and motherboard original equipment manufacturers (OEMs), silicon vendors, and embedded computer manufacturers. IPMI 2.0 contains significant enhancements over IPMI 1.5; it is backward-compatible with previous versions and introduces important features and security enhancements that can be applied to help reduce security concerns for HPC management, including the following:

- **Enhanced authentication:** Extensions to the protocols for IPMI Over IP, collectively referred to as RMCP+, support algorithms that provide robust key-exchange processes for establishing sessions and authenticating users. RMCP+ incorporates authentication based on Secure Hash Algorithm 1 (SHA-1) and supports the Advanced Encryption Standard (AES).
- **Virtual LAN (VLAN):** VLANs work with VLAN-aware routers and switches to allow a physical network to be partitioned into virtual networks in which a group of devices on different physical LAN segments can communicate with one another as if they were all on the same physical LAN segment. Administrators can use this technology to set up a management VLAN in which only devices that are members of that VLAN receive packets related to management, and conversely these devices are isolated from network traffic for other VLANs.
- **Serial Over LAN (SOL):** SOL provides a mechanism that enables the serial controller of a managed system to be redirected over an IPMI Over IP session. SOL is implemented as a payload type under the payload capability in RMCP+.
- **Payloads:** RMCP+ enables IPMI Over IP sessions for other types of traffic in addition to IPMI messages, which include both standard payload types defined in the IPMI specification (such as SOL) and OEM value-added payload types.
- **Encryption:** IPMI messages and other payloads carried over RMCP+ can be encrypted, enabling confidential remote configuration of parameters such as user passwords and transfer of sensitive payload data over SOL.

- **Extended user login options:** Extended options support role-only logins for simple environments in which administrators can enable logins according to a given privilege level, without needing to assign or configure usernames. Support for two-key logins, which require both a user-specific and BMC-specific key to connect to a given BMC, enables administrators to configure BMCs for a robust environment.

Performing secure remote management with RMCP

LAN interface specifications define how IPMI messages can be sent to and from the BMC encapsulated in RMCP packets as datagrams, a capability also referred to as IPMI Over LAN. IPMI 2.0 defines an extended packet format and capabilities that are collectively referred to as RMCP+. RMCP+ uses the RMCP packet format, but defines extensions to the fields defined under the IPMI message class data carried within the RMCP packet. These extensions support enhanced authentication, encryption, discovery, and the ability to carry additional types of traffic (payloads) in addition to IPMI messages over an IPMI session, whereas IPMI 1.5 supported carrying only IPMI messages.

Enhancements with RMCP+

RMCP+ provides support for multiple payload types over an IPMI session, including both standard payloads (such as the payload for the SOL capability defined in this specification) and OEM payloads. It incorporates enhanced user authentication algorithms, including session-setup and key-handling algorithms that are more robust than those for IPMI Over LAN in IPMI 1.5. IPMI messages and other payloads can be encrypted under an IPMI session, enabling confidentiality for remote operations such as setting user passwords and for SOL.

RMCP+ follows many of the packet format and authentication elements defined for RMCP as specified in the Distributed Management Task Force (DMTF) Alert Standard Format (ASF) 2.0 specification. RMCP+ supports encrypted/unencrypted and authenticated/unauthenticated traffic on a single connection; encryption and authentication are handled at the IPMI message class level, which means that encrypted and authenticated sessions can be established on any User Datagram Protocol (UDP) port, including port 26Fh. IPMI allows a BMC to be configured so that authentication and encryption are utilized only when the payload or privilege level of operation requires it, eliminating the need to have all traffic authenticated or encrypted on a connection when only a small portion of the traffic may require that level of security. This configuration can provide a significant performance benefit when using inexpensive microcontrollers for BMCs.

Setting the RMCP+ encryption key	`omconfig chassis remoteaccess config=`*`nic`* `encryptkey=`*`text`* `confirmencryptkey=`*`text`*
Setting the remote access password	`omconfig chassis remoteaccess config=`*`user`* `id=`*`number`* `newpw=`*`text`* `confirmnewpw=`*`text`*
Setting the remote access user role for serial access	`omconfig chassis remoteaccess config=`*`user`* `id=`*`number`* `serialaccesslevel=administrator | operator | user | none`
Setting the remote access user role for LAN access	`omconfig chassis remoteaccess config=`*`user`* `id=`*`number`* `lanaccesslevel=administrator | operator | user| none`
Setting the remote access user role for DRAC access	`omconfig chassis remoteaccess config=`*`user`* `id=`*`number`* `dracusergroup=admin | poweruser | guest | testalert | custom | none`

Figure 3. Dell OpenManage Server Administrator commands for configuring remote access settings

Secure management using RMCP+

To utilize the enhanced security of RMCP+, administrators set the encryption key. One way to set the RMCP+ key is to invoke the BIOS utility menu for remote access by pressing Ctrl+E during system startup and entering the key manually in the BMC configuration options. Another way to set the key and assign user privileges is on the managed node through Dell OpenManage Server Administrator using omconfig utility commands (see Figure 3).

Once the RMCP+ key is set along with the user ID and password, accessing the BMC requires users to specify all three parameters, including the hexadecimal RMCP+ encryption key. For example, using the ipmish CLI utility (available with the BMC Management Utility suite) in a Dell HPC cluster with a one-to-one IPMI session to a specific remote server to obtain the system status would require the following command:

```
ipmish -ip IP address -u username -p password
   -k hexadecimal encryption key sysinfo
```

These commands may be scripted (including the RMCP+ key) for easy access to the server BMCs in the HPC cluster.

Securing out-of-band management with the DRAC 5

The DRAC 5 provides out-of-band remote management capabilities, crashed system recovery, and power control functions for Dell PowerEdge servers. The DRAC 5 communicates directly with the BMC and can be used to configure e-mail alerts and log event data. The DRAC 5, powered by the system in which it is installed and utilizing its own processor and memory subsystems, can be connected with a LAN cable through its RJ-45 connector to provide an out-of-band fabric for external communication in an HPC environment. It also provides several useful features such as virtual media, console redirection, and a comprehensive GUI to add alerts and modify configuration parameters.

To prevent unauthorized access to remote systems, the DRAC 5 provides IP address filtering, which defines a specific range of IP addresses that can access the DRAC 5, as well as IP address blocking, which limits the number of failed login attempts from a specific IP address. The racadm CLI enables administrators to locally or remotely configure and manage the DRAC 5. This CLI runs on the management station and the managed system and is included on the Dell Systems Management Consoles CD; it is also available through the serial, Telnet, and Secure Shell (SSH) console to the DRAC subsystem using the DRAC IP address. To restrict the login to a single IP addresses (for example, 192.168.0.57), administrators should use the full mask, as shown in the following example commands:

```
racadm config -g cfgRacTuning
   -o cfgRacTuneIpRangeEnable 1
racadm config -g cfgRacTuning
   -o cfgRacTuneIpRangeAddr 192.168.0.57
racadm config -g cfgRacTuning
   -o cfgRacTuneIpRangeMask 255.255.255.255
```

The following example commands prevent a client IP address from establishing a session for five minutes if that client has failed its five login attempts in a one-minute period of time:

```
racadm config -g cfgRacTuning
   -o cfgRacTuneIpRangeEnable 1
racadm config -g cfgRacTuning
   -o cfgRacTuneIpBlkFailCount 5
racadm config -g cfgRacTuning
   -o cfgRacTuneIpBlkFailWindows 60
racadm config -g cfgRacTuning
   -o cfgRacTuneIpBlkPenaltyTime 300
```

The Server Management Working Group's Server Management Command-Line Protocol (SM CLP) specification is a subcomponent of the overall DMTF Systems Management Architecture for Server Hardware (SMASH) standard. The DRAC 5 provides support for SM CLP, which is hosted from the DRAC 5 controller firmware. SM CLP supports Telnet-, SSH-, and serial-based interfaces. The DRAC 5 SM CLP interface is based on SM CLP Version 1.0, which enables users to perform server power management (powering up, powering down, or rebooting the system) as well as system event log management (displaying or clearing system event log records). The SM CLP interface with the DRAC 5 makes the DRAC hardware independent and provides for secure interoperability. The DRAC 5 also provides the following security features:

- Racadm CLI and GUI operation, which supports 128-bit and 40-bit SSL encryption (for countries where 128-bit is not acceptable), except when using Telnet
- User ID and password configuration through the secure GUI or racadm CLI
- Session time-out configuration (in seconds) through the GUI or racadm CLI
- IP port configuration (where applicable)
- SSH, which uses an encrypted transport layer for enhanced security

The DRAC 5 helps provide a reliable one-to-one out-of-band management fabric for Dell HPC clusters, and the security enhancements can help improve cluster security for critical management tasks under an architecture built on SMASH and SM CLP standards.

Using Dell OpenManage security features

The Dell OpenManage suite provides the following security features:

- Role-based authority that allows administrators to configure specific privileges for each user
- User ID and password configuration through the GUI or CLI in most cases
- SSL 128-bit and 40-bit encryption (for countries where 128-bit is not acceptable), except when using Telnet
- Session time-out configuration (in minutes) through the GUI or CLI
- Configuration of many commonly known ports

Features such as role-based access control, authentication, and encryption in both GUIs and CLIs help ensure security through access administration. Security is enforced by restricting the operations that can be executed by users in specific roles, with some users assigned multiple roles; this feature can help enable security administration to resemble any organization's structure. Group privileges decide access rights for each Dell OpenManage Administrator user. The three levels are User (can view information), Power User (can set warning threshold values, run diagnostic tests, and configure alert actions), and Administrator (can perform power-down actions, configure auto-recovery, clear logs, and send e-mail).

Dell OpenManage software also helps provide security through authentication and encryption.

Authentication. The Dell OpenManage Server Administrator authentication scheme validates the context of the current process under which the CLI or GUI is executed, to help ensure that Server Administrator functions are properly authenticated. For authentication mechanisms, Server Administrator uses Integrated Windows Authentication and the Pluggable Authentication Modules (PAM) library on the Microsoft® Windows® and Red Hat® Enterprise Linux operating systems, respectively. PAM is a documented library of functions that allows administrators to determine how individual applications authenticate users.

Encryption. Dell OpenManage Server Administrators uses an HTTP Over SSL (HTTPS) connection to help protect the managed system. The Windows, Red Hat Enterprise Linux, and Novell® SUSE® Linux Enterprise Server operating systems use Java Secure Socket Extension (JSSE) to help protect user credentials and other sensitive data transmitted when a user accesses the Server Administrator GUI over the socket connection.

In addition to security protocols, the Dell OpenManage suite provides OS-level agents for monitoring software on compute nodes, which use encryption algorithms and authentication mechanisms to enhance cluster security against external attacks.

Features such as authentication, encryption, role-based authority, and port configuration have been incorporated into the Dell OpenManage suite. These features are critical to helping provide a secure operating environment for managing Dell HPC clusters. Administrators can incorporate a hierarchical role distribution (and apply various other security settings) to help prevent unauthorized users from intentionally or unintentionally misusing cluster components and capabilities.

Managing Dell HPC clusters in a secure environment

Dell offers a comprehensive management package for ninth-generation Dell PowerEdge servers through its Dell OpenManage suite, which can be utilized for HPC clusters. In addition to security protocols, the Dell OpenManage suite provides OS-level agents for monitoring software on compute nodes, which use encryption algorithms and authentication mechanisms to enhance cluster security against external attacks. The DRAC 5 subsystem, which provides a hardware- and software-based fabric for out-of-band management, also includes security features such as SSL for the GUI, SSH access to the racadm CLI, configurable RMCP+ encryption keys, and role-based authentication—which enables an easily configurable, hierarchical access system. Used together, these components can help enhance security of Dell HPC clusters.

Arun Rajan is a systems engineer in the Scalable Systems Group at Dell. His current interests and responsibilities include HPC cluster management, cluster computing packages, performance benchmarking, and product development. He has a B.E. in Electronics and Communications Engineering from the National Institute of Technology, Tiruchirappalli, in India, and an M.S. in Computer and Information Science from the Ohio State University.

Tong Liu is a systems engineer in the Scalable Systems Group at Dell. His current research interests are HPC cluster management, high-availability HPC clusters, and parallel file systems. Tong serves as a program committee member of several conferences and working groups on cluster computing. Before joining Dell, he was an architect and lead developer of High Availability Open Source Cluster Application Resources (HA-OSCAR). Tong has an M.S. in Computer Science from Louisiana Tech University.

Yung-Chin Fang is a senior consultant in the Scalable Systems Group at Dell. He specializes in HPC systems, advanced HPC architecture, and cyberinfrastructure management. Yung-Chin has published dozens of conference papers and articles on these topics. He also participates in industry-standard organizations, academic conferences, and HPC cluster–related open source communities as a Dell representative.

Saeed Iqbal, Ph.D., is a systems engineer and consultant in the Scalable Systems Group at Dell. He is one of the lead engineers for the Beowulf HPC cluster program based on ninth-generation Dell servers. His current work involves evaluation of resource managers and job schedulers used in HPC clusters in industry and academia. He is also involved in performance analysis of HPC clusters. Saeed has a B.S. in Electrical Engineering and an M.S. in Computer Engineering from the University of Engineering and Technology in Lahore, Pakistan, and a Ph.D. in Computer Engineering from the University of Texas at Austin.

FOR MORE INFORMATION

Dell systems management documentation:
support.dell.com/support/systemsinfo/documentation.aspx?~cat=6&~subcat=111

IPMI 2.0 specification:
www.intel.com/design/servers/ipmi

Ali, Rizwan, Rinku Gupta, Garima Kochhar, and Bill Bryce. "Platform Rocks: A Cluster Software Package for Dell HPC Platforms." *Dell Power Solutions,* November 2005. www.dell.com/downloads/global/power/ps4q05-20050227-Ali.pdf

Fang, Yung-Chin; Arun Rajan; Monica Kashyap; Saeed Iqbal, Ph.D.; and Tong Liu. "Dell OpenManage Tools for High-Performance Computing Cluster Management." *Dell Power Solutions,* November 2005. www.dell.com/downloads/global/power/ps4q05-20050239-Fang.pdf

Kochhar, Garima, Rizwan Ali, and Arun Rajan. "Configuring the BMC and BIOS on Dell Platforms in HPC Cluster Environments." *Dell Power Solutions,* November 2005. www.dell.com/downloads/global/power/ps4q05-20050222-Kochhar.pdf

Rajan, Arun, Tong Liu, Yung-Chin Fang, Garima Kochhar, and Ron Pepper. "Customizing Management of HPC Clusters." *Dell Power Solutions,* November 2006. www.dell.com/downloads/global/power/ps4q06-20060427-Fang.pdf

Platform Open Cluster Stack:

An Enhanced Cluster Software Package for Dell HPC Platforms

Platform™ Open Cluster Stack (OCS) is an open source, standards-based cluster software stack designed to help administrators efficiently deploy and manage high-performance computing clusters. This article discusses some of the features and enhancements introduced in the latest Platform OCS version.

BY BILL BRYCE, GARIMA KOCHHAR, RINKU GUPTA, AND RIZWAN ALI

Related Categories:

High-performance computing (HPC)

Platform Computing

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

Linux® OS–based high-performance computing (HPC) clusters on industry-standard servers have emerged as a viable alternative to supercomputers and other architectures running proprietary operating systems. The appropriate combination of hardware, interconnects, storage, and software can enable organizations to quickly build an HPC cluster. However, choosing, building, and testing these components can be complex and time-consuming.

Platform Computing and the Dell HPC Cluster team created Platform Open Cluster Stack (OCS) based on the San Diego Supercomputer Center (SDSC) Cluster Toolkit for building stable, manageable, and scalable clusters.[1] Platform OCS is designed to provide a comprehensive application development and deployment stack, improving on previous cluster management toolkits by providing additional functionality and enhanced support to help organizations deploy, manage, and maintain Linux-based HPC clusters.

Platform OCS architecture

In the past, deploying a UNIX® OS–based HPC system typically involved purchasing both the hardware and software from a single vendor. Organizations typically chose their hardware based on recommendations from an application vendor that ported and qualified its applications onto proprietary UNIX systems, or, alternatively, they obtained the necessary software to implement an HPC cluster from hardware vendors or integrators.

With the emergence of the open source Linux OS, the software stack is typically no longer completely delivered by the hardware vendor. Organizations now have multiple choices, both open source and commercial, for each software stack component—for cluster management, workload management, file systems, and other software. Although this change has provided increased flexibility, organizations must prudently select among the available software components and configure them correctly to work together to obtain optimal usage and performance

[1] This product contains software developed by the Rocks Cluster Group at the San Diego Supercomputer Center at the University of California, San Diego, and its contributors.

Cluster component	Platform OCS component
Dell servers	• *Front-end nodes:* Support for Dell PowerEdge 1850, PowerEdge 1950, PowerEdge 2850, and PowerEdge 2950 servers • *Compute nodes:* Support for Dell PowerEdge SC1425, PowerEdge 1850, PowerEdge 1855, PowerEdge 1950, PowerEdge 1955, and PowerEdge 2950 servers
Management hardware	Support for Dell Remote Access Controllers and baseboard management controllers
OS	Red Hat Enterprise Linux 4
Cluster installation	Red Hat Anaconda
Node and cluster management	Dell OpenManage, Ganglia, National Center for Supercomputing Applications (NCSA) Cluster Monitoring (Clumon)
Network and node file systems	IBRIX Fusion, Network File System, Parallel Virtual File System, standard Linux file systems
Resource management	Platform Enterprise Grid Orchestrator® grid platform
Application workload management	Platform Lava, Platform LSF HPC
Development tools	Intel compilers and tools, GNU compilers and tools
Middleware application programming interfaces	Parallel Virtual Machine (PVM), several MPI implementations (including MPICH, Intel MPI, LAM/MPI, and MVAPICH)
Benchmark tests	IOzone, Bonnie, High-Performance Linpack
High-speed interconnects	Support for Myricom Myrinet and Cisco Topspin InfiniBand

Figure 1. Platform OCS hardware and software components

from their cluster. Application vendors must also select the appropriate hardware, OS, and software environment on which to test and qualify their software.

Platform OCS is designed to help simplify these choices. Organizations can use the default open source components shipped with Platform OCS or replace some with commercial tools. Platform OCS is designed to streamline cluster building and management by simplifying the process of adding or removing these components and automatically configuring them when installed, which can help reduce IT costs by limiting custom development. For application vendors, Platform OCS also provides a consistent software stack from the OS up to the application level, helping provide a stable testing environment.

Platform Computing and Dell test and qualify hardware and software for HPC deployments. Figure 1 shows the primary Platform OCS components, including a list of currently supported Dell™ PowerEdge™ servers. Figure 2 illustrates the Platform OCS software stack.

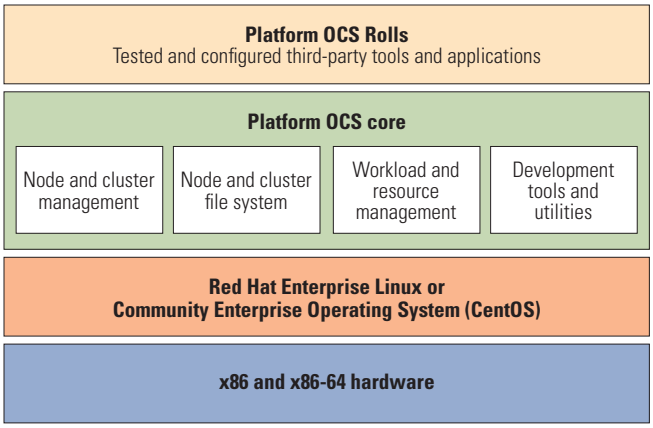

Figure 2. Platform OCS software stack

Using Platform OCS with Dell servers and storage

The latest release of Platform OCS (4.1.1-1.0) supports ninth-generation Dell PowerEdge servers with dual-core Intel® Xeon® 5000 or 5100 series processors and eighth-generation PowerEdge servers with previous-generation Intel Xeon processors, as shown in Figure 1. It provides the Red Hat® Enterprise Linux 4 Update 3 OS with support for Intel Extended Memory 64 Technology (EM64T)—including Red Hat Enterprise Linux AS for front-end nodes and Red Hat Enterprise Linux WS for compute nodes—and includes qualified Dell drivers for supported Dell servers and storage, such as the Dell PowerVault™ MD1000 Serial Attached SCSI (SAS) disk expansion enclosure and the associated Dell PowerEdge Expandable RAID Controller 5. The built-in support and hardware drivers can enable organizations to deploy a cluster quickly and easily.

Platform OCS 4.1.1-1.0 also includes the following features:

- **Single installation DVD:** A single DVD helps simplify and automate installation, eliminating the need to use multiple CDs.
- **Fixes for Red Hat Enterprise Linux and SDSC Cluster Toolkit:** The Platform OCS package integrates fixes for known issues with Red Hat Enterprise Linux 4 Update 3 and SDSC Cluster Toolkit.
- **Dell scripts:** These scripts enable server BIOS and baseboard management controller (BMC) configuration for Intelligent Platform Management Interface (IPMI) traffic, console redirection, and remote management, as well as configuration of features such as Intel Hyper-Threading Technology and Demand-Based Switching. These scripts use omconfig, a Linux-based utility provided by the Dell OpenManage™ suite. The scripts also allow PowerEdge SC1425 BMC configuration using the Linux IPMItool utility.[2]

[2] For more information about these scripts, see "Configuring the BMC and BIOS on Dell Platforms in HPC Cluster Environments," by Garima Kochhar, Rizwan Ali, and Arun Rajan, *Dell Power Solutions,* November 2005, www.dell.com/downloads/global/power/ps4q05-20050222-Kochhar.pdf.

EXAMPLE PLATFORM OCS COMMANDS

The table below lists some commands administrators can use with the Platform OCS rocks-update, custom-partition, rocks-compute, and rollops tools.

`rocks-update -d emacs`	Downloads the emacs package to the front-end node, updates the front-end repository, and increments the repository version number
`rocks-update -f emacs`	Installs the new emacs version of the package on the front-end node
`rocks-update -c emacs`	Updates compute nodes with the new emacs package
`custom-partition -r 12000 -b`	Sets the root partition size for compute nodes to 12,000 MB and rebuilds the distribution so that subsequent compute node installations also use this partitioning scheme
`rocks-compute -a -p /tmp/downloads/myrpm-1.1.x86_64.rpm -b`	Adds the myrpm package to the extend-compute.xml file with the appropriate syntax and rebuilds the distribution so that subsequent compute node installations also include this package
`rollops -l`	Lists the Rolls currently installed on the system, including their version number and architecture type

- **IBRIX file system compatibility:** Platform OCS is compatible with the IBRIX Roll for the IBRIX Fusion file system.
- **Cisco and Myricom high-speed interconnect support:** The Cisco Topspin InfiniBand Roll contains drivers for supported InfiniBand configurations, including PCI Express–based single-port memory-free InfiniBand host channel adapters. The Myricom Myrinet Roll includes drivers for Myrinet D cards.
- **SDSC Rolls compatibility:** Rolls developed for versions 4.0 and 4.1 of SDSC Cluster Toolkit typically should be compatible with Platform OCS as well.

Platform OCS features and enhancements

Platform OCS 4.1.1-1.0 introduces or enhances several features and tools, including the rocks-update, custom-partition, rocks-compute, and rollops tools; front-end partitioning options; the Environment Modules utility; benchmarking applications; Intel software tools; and cluster installation features. For more information, including usage details, refer to the man pages on the front-end node.

Rocks-update, custom-partition, rocks-compute, and rollops tools

The rocks-update, custom-partition, rocks-compute, and rollops tools provided with Platform OCS can help administrators perform important cluster management tasks. The "Example Platform OCS commands" sidebar in this article lists some commands for these tools.

Rocks-update. Rocks-update enables administrators to automatically download Red Hat Package Manager (RPM™) cluster update packages and their dependencies from Red Hat Network (RHN) using the up2date tool, provided they have an RHN subscription. They can then update the front-end and compute nodes without reinstalling their OS and software, and can also update the front-end RPM distribution so that future compute node installations are up-to-date. Version control is included to help distinguish updated and non-updated nodes. The front-end node also maintains the download and update history.

Rocks-update does not allow administrators to update the running kernel or any kernel-dependent RPM packages, because doing so on a production cluster could result in missing or improperly tested kernel drivers, which can cause the hardware to stop functioning properly. Administrators can use rocks-update to install new packages (those that are not already installed on the cluster); however, the RPM package name must be known, because rocks-update will not apply all available RHN updates without explicit instruction. This features gives administrators control over software installed on the cluster and can help protect the cluster from unwanted automatic updates. Administrators should first verify all patches on a separate node before updating the entire cluster.

Custom-partition. Administrators can use the custom-partition command-line interface to change the compute node root and swap partition sizes. Without this utility, administrators must modify the extend-auto-partition.xml file manually, check for syntax errors, and rebuild the distribution. Custom-partition handles these tasks and helps simplify the process of changing partition sizes. Platform OCS maintains a log of executed custom-partition commands in the directory where the command was run.

Rocks-compute. Administrators can use rocks-compute to extend the definition of a compute node by adding RPM packages and post-installation scripts to the extend-compute.xml file. Without this utility, administrators must modify this file manually, check for syntax errors, and rebuild the distribution. Rocks-compute handles these tasks and helps simplify the process of customizing compute nodes for specific applications and tasks. Platform OCS maintains a log of executed rocks-compute commands in the directory where the command was run.

Rollops. Rollops enables administrators to add, remove, and upgrade Rolls on the front-end node without reinstalling its OS and software. Rollops applies the Roll on the front-end node and updates the distribution. Rolls can also be enabled or disabled using the access keyword. If a Roll is disabled, it will not be applied to compute nodes even if the Roll is installed on the front-end node. Once the distribution is updated, administrators must reinstall Platform OCS on the compute nodes for the Roll changes to take effect.

The rollops upgrade feature is available for the Dell Roll—administrators can upgrade this Roll to a new version in a single operation. This feature enables administrators to easily add support for new hardware and incorporate Dell driver updates into a cluster; for Rolls without this feature, administrators must first remove the existing Roll and then add the new version.

Front-end partitioning options

Unlike previous Platform cluster toolkit releases, which erased the front-end node's hard drive each time administrators reinstalled the toolkit, Platform OCS 4.1.1-1.0 allows administrators to preserve existing partitions and install Platform OCS on existing free hard drive space. If Platform OCS detects a Microsoft® Windows® OS partition during front-end installation, it provides the option of preserving this partition and installing Platform OCS on the rest of the disk. If it detects a Linux partition, it also provides the option of preserving this partition; however, after installation, administrators must manually remount the partition on the front-end node.

Environment Modules utility

The Environment Modules utility, an open source tool that enables administrators to dynamically load different environments in their shells, is now available as a Platform OCS Roll on the base DVD.[3] The Roll provides a set of default *modulefiles* that contain information to configure different Message Passing Interface (MPI) environments. Once administrators set up the modules package, they can use these files to load and unload environments for different applications, which can be especially useful for administrators using multiple MPI libraries. Many administrators can share configuration files in a cluster, which can help provide a consistent working environment. Administrators can check which modules are available with the `module avail` command, and load a particular module with `module add` *`module name`*.

Benchmarking applications

Platform OCS provides several benchmarking tools to help determine cluster reliability and overall performance, including IOzone, Bonnie, and High-Performance Linpack (HPL).[4] HPL, for example, determines raw cluster performance and is commonly used to rank supercomputers on the TOP500 Supercomputer Sites list;[5] once administrators have generated reasonable HPL results for a cluster, it is typically ready for real applications.

Intel software tools

Platform OCS supports the Intel Software Tools Roll, available from the Platform Open Source Grid Development Centre in Singapore.[6] This Roll contains license-based tools such as the Intel C, C ++, and Fortran Compilers; Intel MPI; Intel Math Kernel Library (MKL) Cluster Edition; Intel Integrated Performance Primitives; Intel Trace Analyzer/Collector; Intel Application Debugger; and MPICH and LAM (Local Area Multicomputer)/MPI compiled with Intel Compilers. For dual-core Intel Xeon 5100 series processor–based clusters, administrators should use Intel Compilers version 9.1 and MKL version 8.1. The Intel MPI Library Runtime Environment Kit, which contains the software necessary to run applications developed using the Intel MPI Library, is also packaged as a Roll and available on the Platform OCS base DVD.

Cluster installation features

Platform OCS 4.1.1-1.0 contains several cluster installation features, including an enhanced installer and kickstart file caching.

Enhanced installer. SDSC Cluster Toolkit includes an updated installer, one feature of which is to split the kickstart generation process between the front-end and compute nodes. When a compute node requests a kickstart file, the front-end node traverses the kickstart graph and creates the XML kickstart file, as in previous SDSC Cluster Toolkit versions. But now, this XML file is handed over to the compute node, which completes the process and generates the finished kickstart file. This change helps reduce the kickstart generation load on the front-end node, allowing it to serve additional kickstarting clients at the same time and helping improve installer scalability.

Administrators can also now use BitTorrent to distribute packages to installing nodes instead of HTTP or HTTP over Secure Sockets Layer (HTTPS). With BitTorrent enabled, the front-end node defines installing peers, and nodes installing at the same time can use any of the peers to pull RPM packages instead of using the front-end node. This process can help reduce the front-end bottleneck and distribute the load over the peers. Because this feature is not typically beneficial for clusters with fewer than 64 nodes, by default Platform OCS uses HTTPS on the front-end node to distribute packages. Platform OCS provides a command-line utility to enable or disable BitTorrent: to set it up and prepare the

[3] For more information about Environment Modules, visit modules.sourceforge.net.

[4] For more information about IOzone, Bonnie, and HPL, visit www.iozone.org, www.textuality.com/bonnie, and www.netlib.org/benchmark/hpl.

[5] For more information about TOP500, visit www.top500.org.

[6] For more information about the Intel Software Tools Roll, visit www.platform.com/osgdc.

front-end distribution for BitTorrent transfers, administrators can run the command `/opt/rocks/sbin/rocks-bittorrent`.

Kickstart file caching. When administrators install multiple nodes of the same type, the kickstart file differs slightly from node to node—in the node name, IP address, and other unique configuration information—but much of the file is the same for each node. Platform OCS can utilize these similarities by caching the kickstart file for a given node type, then using this cached file when the front-end node receives its next kickstart file request for this node type. Unique configuration settings are generated for each kickstart request, which can help ensure that nodes are installed correctly. Because the front-end node no longer needs to carry out all generation steps for every kickstart file request, the cached file can help reduce front-end load and improve installer scalability. This feature is enabled and disabled depending on whether the /home/install/sbin/cache/disable-cache file is present on the front-end node.

Optional Platform OCS features

Optional Platform OCS features include the Platform Load Sharing Facility (LSF®) HPC system, Optional Rolls, and support subscriptions.

Platform LSF HPC

Platform LSF HPC is a commercial workload management system that organizations can purchase as part of Platform OCS. LSF HPC 6.2 is designed to improve LSF HPC installation and maintenance; for example, previous LSF HPC Rolls required manual post-installation configuration, but Platform OCS 4.1.1-1.0 can automatically add or remove compute nodes from the appropriate LSF HPC configuration files. Once a compute node installation is complete, the node can join the LSF HPC cluster.

Platform LSF HPC can also automatically fail over to any designated master candidate host in the system. Platform OCS supports LSF HPC failover to a master candidate, and the shared transaction log file can be stored on a shared file system, helping ensure the system remains operational even during a front-end node failure.

Optional Rolls

The Platform OCS base DVD contains two categories of Rolls: Required Rolls and Optional Rolls. Required Rolls are essential for cluster functioning, while Optional Rolls allow administrators to customize a cluster to their specific requirements. Optional Rolls include the Ntop, Clumon, Ganglia, Extra Tools, Platform Lava, Myricom Myrinet, and Cisco Topspin InfiniBand Rolls.

Support subscriptions

Platform OCS includes an optional support subscription for commercial customers that includes 24/7 e-mail support; business-day phone support; emergency pager support; and a Web site allowing customers to log and view support tickets, download software and patches, e-mail developers, and obtain updated documentation. Supported customers can also register for the Platform OCS e-mail list to receive notifications of critical updates or security patches.

Efficient HPC cluster management

Platform OCS can help provide application developers with a stable test and deployment platform and enable administrators to efficiently deploy and manage Linux-based HPC clusters. Platform Computing and Dell have collaborated to provide a comprehensive, open source, standards-based cluster software stack that they plan to continue improving in the future to benefit both application developers and IT administrators.

Bill Bryce is the senior product manager for Platform OCS at Platform Computing, where he has worked for the past 10 years. His interests include distributed computing, parallel programming, communications protocols, and operating systems. Bill has a B.Math (Computer Science) from the University of Waterloo in Canada.

Garima Kochhar is a systems engineer in the Scalable Systems Group at Dell. She has a B.S. in Computer Science and Physics from Birla Institute of Technology and Science in Pilani, India, and an M.S. in Computer Science from the Ohio State University, where she worked on job scheduling.

Rinku Gupta is a systems engineer in the Scalable Systems Group at Dell. Her current research interests are middleware libraries, parallel processing, performance, and interconnect benchmarking. Rinku has a B.E. in Computer Engineering from the University of Mumbai in India and an M.S. in Computer Information Science from the Ohio State University.

Rizwan Ali is a systems engineer in the Scalable Systems Group at Dell. His current research interests include performance benchmarking, cluster architecture, parallel applications, and high-speed interconnects. Rizwan has a B.S. in Electrical Engineering from the University of Minnesota.

FOR MORE INFORMATION

Platform Open Cluster Stack:
www.platform.com/openclusterstack

Ali, Rizwan, Rinku Gupta, Garima Kochhar, and Bill Bryce. "Platform Rocks: A Cluster Software Package for Dell HPC Platforms." *Dell Power Solutions,* November 2005. www.dell.com/downloads/global/power/ps4q05-20050227-Ali.pdf

Using OpenFabrics InfiniBand for HPC Clusters

The OpenFabrics Alliance was formed to resolve issues with hardware and software interoperability and to deliver open source software for Remote Direct Memory Access fabric technologies like the InfiniBand architecture. Major initiatives by this alliance have resulted in the creation of an open source Linux® OS–based software stack for InfiniBand, which is rapidly gaining acceptance in the high-performance computing field.

BY MUNIRA HUSSAIN, RINKU GUPTA, AND TONG LIU

Related Categories:

Dell PowerEdge servers

High-performance computing (HPC)

InfiniBand

Interconnects

OpenFabrics

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

Effective server resource utilization requires a balance among CPU speed, memory bandwidth, and I/O performance. However, traditional bus-based architecture is limited in its ability to meet fast-growing CPU performance. To eliminate the I/O bus architecture bottleneck, the InfiniBand Trade Association (IBTA), consisting of leading computer technology companies like Dell, Intel, Sun, and AMD, was formed to develop specifications for the InfiniBand architecture. By the end of 2000, the IBTA had defined InfiniBand technology, which is a switched fabric I/O architecture that replaces the traditional server PCI bus with a high-speed, low-latency serial I/O interconnect to deliver a point-to-point, scalable interconnect infrastructure. In contrast to proprietary interconnects, InfiniBand is based on industry standards, enabling it to evolve and scale.

InfiniBand has shown itself to be a compelling high-performance interconnect technology for performance-focused cluster environments. Because of its performance and scalability,[1] InfiniBand has been adopted in high-performance computing (HPC) cluster environments for communication-intensive applications, including those in the fields of aerospace and defense, oil and gas exploration, fluid dynamics, weather modeling, and biology and life sciences. According to the June 2006 TOP500 Supercomputer Sites list of the world's most powerful computers, InfiniBand is rapidly gaining acceptance as a high-performing cluster interconnect, and the number of InfiniBand-based supercomputers has grown 33 percent since November 2005.[2]

In addition to its benefits for Interprocess Communication (IPC), InfiniBand has been successfully recognized and utilized in the storage systems field. By delivering a converged communication and storage interconnect, InfiniBand can help simplify data center management and help reduce total cost of ownership. The adoption of

[1] For more information, see "Exploring InfiniBand as an HPC Cluster Interconnect," by Onur Celebioglu, Ramesh Rajagopalan, and Rizwan Ali, *Dell Power Solutions,* October 2004, www.dell.com/downloads/global/power/ps4q04-20040140-Celebioglu.pdf; and "Using PCI Express Technology in High-Performance Computing Clusters," by Rinku Gupta; Saeed Iqbal, Ph.D.; and Andrew Bachler, *Dell Power Solutions,* November 2005, www.dell.com/downloads/global/power/ps4q05-20050226-Gupta.pdf.

[2] November 2005 and June 2006 "Interconnect" lists, TOP500.org, www.top500.org/stats/26/conn and www.top500.org/stats/27/conn.

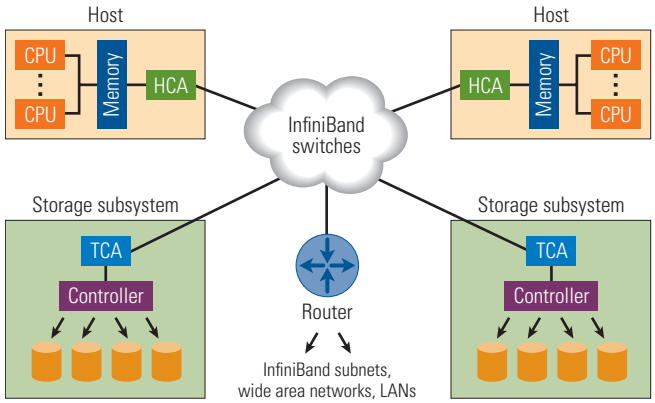

Figure 1. InfiniBand hardware architecture

InfiniBand in the server and storage market, in turn, is driving the demand for development of a complete software stack to support InfiniBand products from different vendors. To fulfill this demand, the OpenFabrics Alliance (formerly the OpenIB Alliance) was formed in June 2004 to create a standard open source InfiniBand software stack. The availability of a standard open source software stack can help promote wide adoption and interoperability across heterogeneous computing architectures.

InfiniBand hardware architecture

Established as an industry standard, the InfiniBand architecture develops a different approach toward I/O efficiency by replacing the traditional shared bus architecture with a switched fabric.[3] As Figure 1 shows, the unified InfiniBand I/O fabric includes three basic components: host channel adapter (HCA), target channel adapter (TCA), and InfiniBand switches. HCAs are connection interfaces to CPUs; TCAs are links to storage, Fibre Channel networks, and other I/O nodes. InfiniBand switches provide the switched connections between HCAs and TCAs. InfiniBand routers can be used to connect multiple switched subnets.

The InfiniBand protocol is an OS bypass protocol; it provides direct access to the InfiniBand HCA and can reduce the number of user-kernel context switches and memory copies. It off-loads data movement from the server CPUs to the InfiniBand HCA, enabling communication between devices and hosts without the traditional system resource overhead associated with network protocols. With Remote Direct Memory Access (RDMA) enabled, InfiniBand adopts a zero-copy approach to transferring data between the sender's memory and the receiver's memory without involving the host CPUs. RDMA has the potential to be useful in communication-intensive applications.

InfiniBand offers traffic management through virtual lanes (VLs), creating multiple virtual links within a single physical link. This mechanism allows a pair of linked devices to isolate communication interference from other connected devices. Each InfiniBand link can accommodate 2 to 16 VLs, which includes 1 VL for traffic management and others for packet transmission.

InfiniBand delivers high bandwidth by utilizing full duplex bidirectional links between devices. It is designed to support three data transfer rates: 1X, 4X, and 12X, with a base data rate (1X) of 2.5 Gbps, thus theoretically providing full duplex single data rate (SDR) bandwidths of 5 Gbps, 20 Gbps, and 60 Gbps. Ongoing development of InfiniBand technology with double date rate (DDR) and quad data rate (QDR) operations may help enable significantly increased throughput.

OpenFabrics software stack

The InfiniBand architecture is a complex and hierarchical structure comprising many components, libraries, programming interfaces, and modules along with multiple protocols that can be used based on application requirements. The InfiniBand interconnect gained popularity in the open source community as research and scientific computing labs adopted it. As a result, multiple versions of the software stack became available. To standardize the InfiniBand software stack, the OpenFabrics Alliance (then named the OpenIB Alliance) was formed to create a standardized Linux OS–based InfiniBand software stack that would be hardware and vendor independent. Key members of the alliance include Dell, Intel, Cisco Systems, and AMD; InfiniBand-focused companies like Mellanox Technologies, SilverStorm Technologies, and Voltaire; and research and scientific computing labs such as Lawrence Livermore National Laboratory, Los Alamos National Laboratory, and Sandia National Laboratory.

The maturity of the InfiniBand stack and the efforts of this alliance have resulted in the adoption of InfiniBand into the upstream Linux kernel. As a result, core components of InfiniBand have been accepted in the Linux 2.6 kernels (2.6.11 and later). Since its inception, the OpenFabrics Alliance has broadened its charter to include other RDMA-capable data center fabrics. The OpenFabrics Alliance now supports iWARP (RDMA on Ethernet) along with InfiniBand.

InfiniBand software architecture

The InfiniBand software architecture (see Figure 2) consists of both kernel-level and user-level components, which complement each other to help provide an end-to-end solution. In the software stack, the low-level InfiniBand kernel driver module is hardware specific and ported on top of the associated hardware. The rest of the architecture is hardware agnostic and is divided into kernel-space software and user-space software.

[3] Additional details about the InfiniBand architecture can be found in the InfiniBand specification at www.infinibandta.org/specs.

Kernel-space software

The kernel-space software consists of the core InfiniBand modules and upper-layer protocols.

Core InfiniBand modules. The core InfiniBand module layer consists of the following:

- **InfiniBand verbs:** These are defined semantically in the InfiniBand specification and describe the action or function to take place. The actual translation of the semantics of a verb to the form of an application programming interface (API) is carried out on top of the hardware driver in the kernel-level software.
- **Management Datagram (MAD) services and agents:** These define an entry point for the upper layers to interact with the HCA driver and hardware.

The InfiniBand specification defines the Subnet Manager, which is responsible for topology discovery and management. It also defines a group of managers called General Services Managers, which are responsible for management and operation of the InfiniBand hardware; examples of such managers include the Performance Manager, which is responsible for retrieving performance statistics from devices, and the Communication Manager, which is responsible for managing connections between devices. Each General Services Manager typically has a corresponding agent on the InfiniBand device (for example, the Performance Manager Agent or Device Manager Agent) and uses MAD packets to request information or operations.

The MAD services layer in the kernel includes components like the Subnet Management Interface (SMI) and General Services Interface (GSI). The SMI is a unique queue pair (QP, the InfiniBand mechanism to send and receive data), called QP0, associated with InfiniBand ports to send and receive information and operation request-response packets between the Subnet Manager and its agents. Similarly, the GSI is a unique QP1 associated with InfiniBand ports to communicate with the various General Services Managers and their agents.

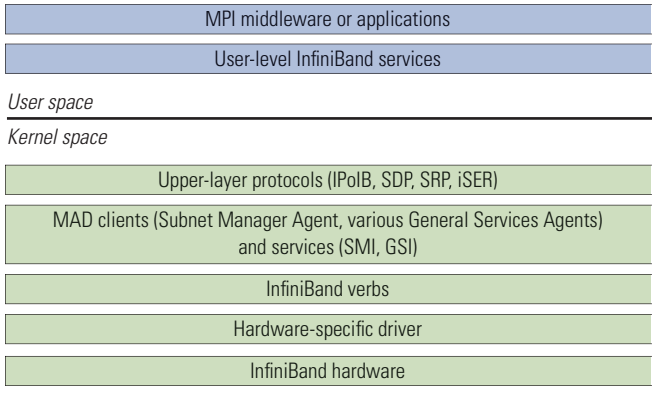

Figure 2. InfiniBand software architecture and component stack

Upper-layer protocols. The rest of the kernel-space software consists of upper-layer protocols like IP Over InfiniBand (IPoIB), Socket Direct Protocol (SDP), SCSI RDMA Protocol (SRP), and Internet SCSI Extensions for RDMA (iSER) in the kernel space.

User-space software

The user-space software consists of access to certain user-level InfiniBand services. InfiniBand allows user-level access to the InfiniBand hardware by providing user-space APIs. The user-level verbs API (semantically similar to the kernel-space verbs) allows upper-level middleware like Message Passing Interface (MPI) software and other applications to bypass the kernel and OS to obtain direct access to the InfiniBand hardware.

OpenFabrics Enterprise Distribution

The OpenFabrics Enterprise Distribution (OFED) is an effort by the OpenFabrics Enterprise working group—a group composed of select hardware vendors providing products based on the OpenFabrics stack—to create a commercial-quality, solution-level OpenFabrics software distribution for end users. It enables end users to deploy InfiniBand devices from various hardware vendors by using a generic software stack. The OFED distribution is a superset of the OpenFabrics software, and contains a snapshot of the OpenFabrics software and components such as MPI software (which falls outside the scope of OpenFabrics) along with the latest kernel modules taken from the OpenFabrics repository.

As of June 2006, the first release of the OFED distribution includes a comprehensive list of kernel modules and software packages such as IPoIB, SDP, SRP, User Direct Access Programming Library (uDAPL), MPI implementations like Open MPI and MVAPICH, subnet managers like OpenSM, and various performance and diagnostic tools.

InfiniBand interoperability and manageability

The OpenFabrics software stack enables hardware and software interoperability and manageability between heterogeneous InfiniBand clusters. Such capabilities have helped this software stack to contribute immensely toward the adoption of InfiniBand. For more information, visit www.openfabrics.com.

Munira Hussain is a systems engineer adviser in the High-Performance Computing Group at Dell.

Rinku Gupta is a systems engineer and consultant in the High-Performance Computing Group at Dell.

Tong Liu is a systems engineer in the High-Performance Computing Group at Dell.

Serial Attached SCSI Storage

for High-Performance Computing

Serial Attached SCSI (SAS) is the successor of SCSI technology, and is designed to meet the performance requirements of new-generation servers by taking advantage of serial connection to help improve performance and scaling. This article discusses deployment of SAS storage in high-performance computing clusters.

BY AZIZ GULBEDEN, AMINA SAIFY, ANDREW BACHLER, AND RAMESH RADHAKRISHNAN, PH.D.

Related Categories:

Characterization

Cluster management

Dell PowerEdge RAID Controller (PERC)

Dell PowerEdge servers

Dell PowerVault storage

High-performance computing (HPC)

Performance

RAID controllers

Serial Attached SCSI (SAS)

Storage

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

Serial protocols are becoming widespread as performance requirements advance beyond what traditional bus-based systems can provide. Traditionally, the SCSI disks are connected to a shared parallel bus. However, the accuracy of parallel connections decreases at high speeds, and performance requirements limit the bus length.

The Serial Attached SCSI (SAS) protocol is designed to help avoid these shortcomings. Serial connections between devices can provide much higher throughput than parallel bus-based connection schemes. The current SAS technology uses point-to-point connections with a maximum 300 MB/sec capacity per disk—in contrast, SCSI disks share a bus that has a maximum 320 MB/sec capacity. The SCSI Trade Association expects the performance of SAS links to double with SAS 600 (expected to be available in 2007) and double again with SAS 1200 (expected to be available in 2010).[1] Additionally, SAS overcomes the SCSI limit of 16 devices per channel. As a result of these improvements, even though SAS uses the same set of SCSI commands, it is not backward compatible with SCSI.

Storage needs in high-performance computing clusters

In a high-performance computing (HPC) environment, storage is a key component. HPC applications require high performance and highly available access to large storage. For example, a fluid dynamics simulation starts parallel processes on several compute nodes that produce a large amount of data. The data is shared among different processes through a distributed file system on the shared storage, and the stored data is reused and updated as the simulation proceeds.

[1] "Serial Attached SCSI – Roadmap," by the SCSI Trade Association, January 2004, www.scsita.org/aboutscsi/sas/SAS_roadmap2004.html.

In addition to high performance and high availability, the storage systems must also be manageable and expandable. Manageability allows setting the appropriate configuration for the storage and enables hardware-related problems to be fixed easily. Furthermore, the storage can be tuned for various applications so that they benefit from custom cache- or block-size settings. Extensibility is necessary to allow growth as requirements change. These factors make using external disk storage attractive for clusters because the external disk arrays are optimized to address these challenges.

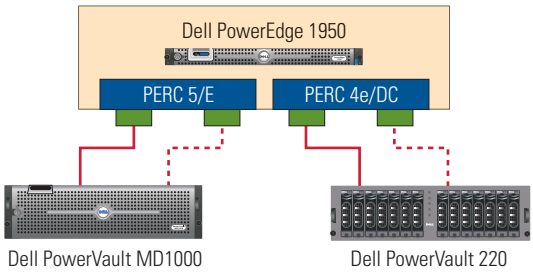

Figure 1. Test environment for comparing SAS and SCSI storage

Test environment for comparing SAS and SCSI

To compare the performance of SAS storage with SCSI storage, a team of Dell engineers ran benchmarks on two Dell™ PowerVault™ storage enclosures in April 2006. The test environment used one Dell PowerVault MD1000 SAS storage enclosure and one Dell PowerVault 220 SCSI enclosure; both were connected to a Dell PowerEdge™ 1950 server with dual 3.2 GHz processors (see Figure 1).

PowerVault MD1000 configuration: External SAS storage

The PowerVault MD1000 is an external disk storage enclosure utilizing SAS disks; it is classified as a JBOD (Just a Bunch of Disks). RAID configuration on the storage is managed by a Dell PowerEdge Expandable RAID Controller (PERC) 5/E adapter that resides inside a PowerEdge server. The PERC 5/E is required on the server to communicate with the PowerVault MD1000. The PERC 5/E is a PCI Express card designed to deliver four times the bandwidth of PCI Extended (PCI-X); it supports RAID-0, RAID-1, RAID-5, RAID-10, and RAID-50. One x4 SAS cable, which can provide bandwidth of up to 1.2 GB/sec, connects the PERC 5/E to the PowerVault MD1000. The PERC 5/E settings can be changed with the Dell OpenManage™ Server Administrator application or from its BIOS, which is accessed by pressing Ctrl + R during system startup.

In a high-performance computing (HPC) environment, storage is a key component. HPC applications require high performance and highly available access to large storage.

The setup used one enclosure with 14 disks; however, for achieving higher performance than what was demonstrated in this environment, up to three enclosures can be daisy-chained, which allows a server to connect to a maximum of 45 disks through one port when all enclosures are fully populated. For maximum storage capacity, both of the ports on the PERC 5/E can be used to connect the server to two sets of daisy-chained PowerVault MD1000 enclosures containing 45 disks each, which allows a maximum of 90 SAS disks connected to one server.

PowerVault 220 configuration: External SCSI storage

Similar to the PowerVault MD1000, the PowerVault 220 is an external disk storage enclosure that uses SCSI disks; it is managed by a PERC 4e/DC adapter that supports the same RAID levels as the PERC 5/E. The PowerVault 220 can host a maximum of 14 disks.

In the setup, both storage enclosures contained fourteen 73 GB, 15,000 rpm hard drives. The PowerVault MD1000 was deployed with fourteen 73 GB Fujitsu SAS disks, and the PowerVault 220 with fourteen 73 GB Seagate Cheetah disks.

Both the PERC 4e/DC and PERC 5/E are dual-port controllers. The Dell HPC team ran I/O benchmarks on two configurations: single channel (one port on the controller connected to the storage) and dual channel (both ports on the controller connected to the storage). For the PERC 5/E, the results include only a single channel because no significant performance difference was observed between single-channel and dual-channel configurations for the PERC 5/E using either the IOzone benchmark or OOCORE application. Other applications with different data-access patterns may benefit from the two-channel configuration with the PowerVault MD1000.

Performance tuning and benchmark study

The Dell HPC team used the IOzone benchmark and OOCORE application to compare SAS and SCSI performance. Before doing so, however, the team conducted initial performance tuning to determine the appropriate settings for the test environment.

Performance study with the Linux® OS read-ahead cache showed that this setting affects the disk storage performance significantly, especially for streaming reads on fast storage devices such as SCSI, SAS, and external disk storage. When a block is accessed, the read-ahead algorithm prefetches and caches the subsequent blocks for enhanced performance. The algorithm turns itself off when it finds that the data is being accessed randomly.

The highest performance was achieved with a 4 MB read-ahead cache; therefore, for all the benchmark tests, the cache was set to 4 MB (8,192 blocks of 512-byte sectors) using the command

`blockdev -setra 8192 /dev/device` or, alternatively, `hdparm -a 8192 /dev/device`. In the Linux 2.6 kernel, this value is stored under /sys/block/*device*/queue/read_ahead_kb in kilobytes.

IOzone benchmark study

The IOzone benchmark[2] was used to measure the performance of both storage enclosures for sequential write and sequential read operations. The test file size was set to 12 GB to eliminate the cache effect. Figure 2 shows the performance of different configurations relative to the PERC 4e/DC single-channel RAID-0 write performance.

For sequential write and read operations, SAS storage provided much higher performance than SCSI storage with the same number of similar hard drives. Read operations in particular performed two to three times faster on SAS storage than on SCSI storage.

OOCORE application study

The OOCORE application,[3] an out-of-core matrix solver, was used for application study. *Out-of-core matrix solving* refers to solving the matrices that do not fit into the CPU memory. The OOCORE application uses the disk to store the matrix and temporary files generated during the process of solving the matrix, which makes the benchmark I/O-intensive. The configuration parameters were adjusted to increase the number of disk I/O operations, and the amount of RAM was reduced to help ensure that the benchmark performance reflected the storage performance.

Figure 3 shows the performance results relative to the PERC 4e/DC single-channel RAID-0 configuration performance. The results show that SAS storage outperformed SCSI storage in all test scenarios. For example, in a RAID-5 configuration, the SAS storage performed 47 percent better than the one-channel PERC 4e/DC SCSI storage and 14 percent better than the two-channel PERC 4e/DC SCSI storage.

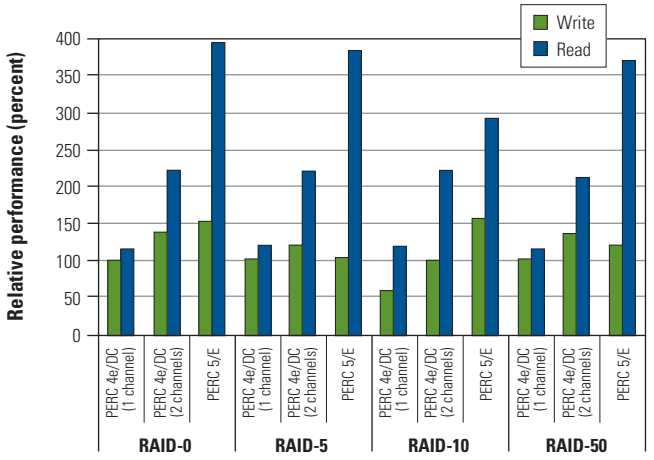

Figure 2. IOzone sequential-access performance: SAS versus SCSI storage

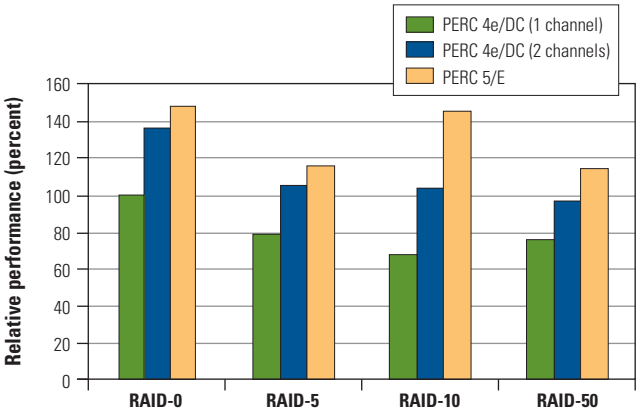

Figure 3. OOCORE performance: SAS versus SCSI storage

Enhanced storage performance

SAS technology with serial architecture can help provide high performance for HPC applications that use storage extensively. The IOzone benchmark and OOCORE application performance results described in this article show that SAS storage consistently provided better throughput than SCSI storage in similar configurations. Performance can also be enhanced by taking advantage of the large number of disks supported with SAS by connecting additional enclosures.

Aziz Gulbeden is a systems engineer in the Scalable Systems Group at Dell. His current areas of focus include scalable high-performance file and storage systems. He has a B.S. in Computer Engineering from Bilkent University in Turkey and an M.S. in Computer Science from the University of California at Santa Barbara.

Amina Saify is a member of the Scalable Systems Group at Dell. Amina has a bachelor's degree in Computer Science from Devi Ahilya University in India and a master's degree in Computer and Information Science from the Ohio State University.

Andrew Bachler is a systems engineer in the Scalable Systems Group at Dell. He has an associate's degree in Electronic Engineering and 12 years of UNIX® and Linux OS experience.

Ramesh Radhakrishnan, Ph.D., is a member of the Scalable Systems Group at Dell. His interests include performance analysis and characterization of enterprise-level benchmarks. Ramesh has a Ph.D. in Computer Engineering from the University of Texas at Austin.

[2] The IOzone benchmark can be downloaded from www.iozone.org.

[3] The OOCORE application can be downloaded from www.nsf.gov/publications/pub_summ.jsp?ods_key=nsf0605.

Scalable Enterprise Implementation Study:

How Dell IT Uses Virtualization to Enable Test and Development

The Dell IT group uses a server farm running virtualization software to provide more than 1,000 test and development environments on fewer than 100 physical servers. This farm enables Dell IT to manage the test and development environments with a small team of administrators, allowing engineers and developers to focus on a wide range of internal projects and minimize time spent setting up test environments.

BY TODD MUIRHEAD; RICK MERINO; DAVE JAFFE, PH.D.; AND JON MERCADO

Related Categories:

Case study

Dell Scalable Enterprise Technology Center

Virtualization

VMware

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

As Dell continues to grow globally, requirements for new applications and updates to existing applications place increasing demands on the Dell IT group infrastructure. Dell IT must be able to scale its computing resources while containing the capital costs and staffing resources needed to efficiently manage these systems. Using a virtual infrastructure for test and development helps Dell to maximize its computing resources with improved server utilization, operational cost-effectiveness, and productivity through rapid provisioning.

This article discusses the virtualization best practices Dell IT has developed as well as the architecture of its large server farm—which is an example of how Dell uses a scalable enterprise architecture based on industry-standard hardware and software to achieve tangible benefits today. It also discusses future plans and expected directions for virtualization technology used by Dell.

Virtualization and the scalable enterprise

Server virtualization has been possible for decades on large, proprietary mainframe systems, enabling these expensive systems to be partitioned and used for multiple purposes at the same time. In recent years, the same type of technology has become available for Intel® processor–based servers, with VMware® ESX Server software leading the way. Using virtualization, a single Dell™ PowerEdge™ server can host multiple virtual machines (VMs) concurrently, with each VM potentially running different applications. The virtualization

technology that was once available only on large mainframes can now be used by enterprises of any size.

Dell IT has implemented virtualization by using ESX Server to support more than 1,000 test and development environments on fewer than 100 physical servers. The environments' design and management is an implementation of the Dell scalable enterprise strategy—that is, using an architecture based on industry standards and helping enable simplified management, improved utilization, and cost-effective scaling.

Ultimately, the scalable enterprise vision leads to an automated data center based entirely on industry standards.[1] Today, enterprises can achieve limited automation in data centers, but as more standards evolve, Dell anticipates that increasing levels of automation will be possible. Dell IT has implemented the first stages of an automated virtualization infrastructure and plans to continue building an automated data center based on industry standards.

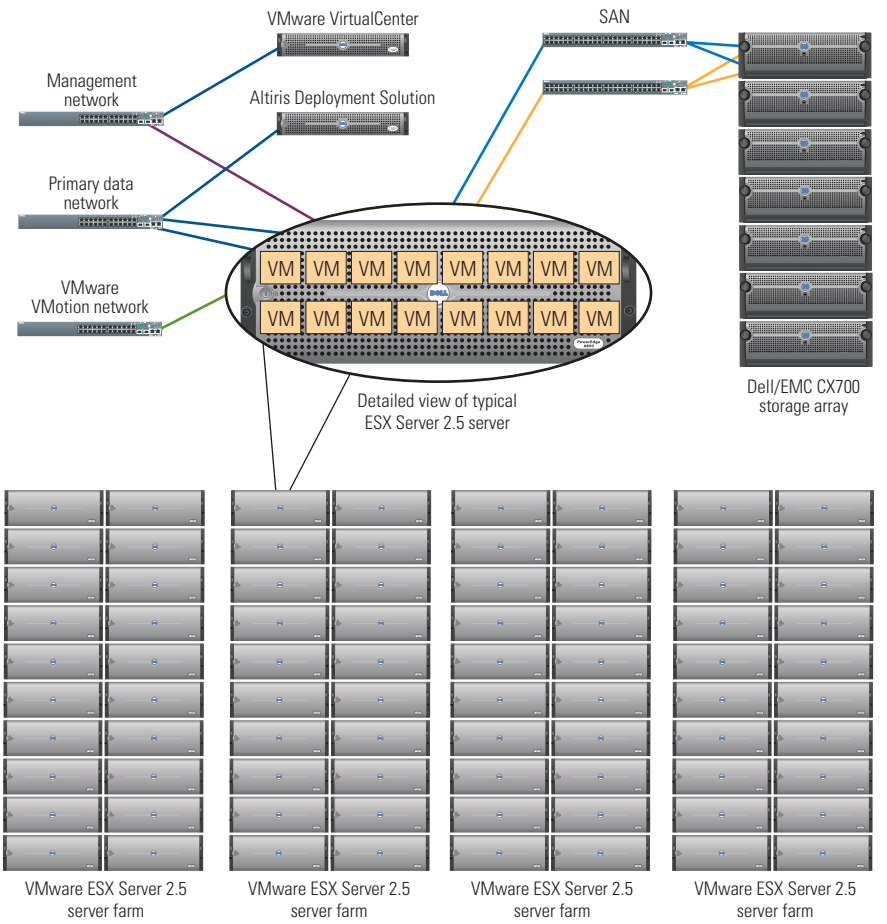

Figure 1. Dell IT virtualization server farm used for test and development

Dell IT virtualization server farm

The Dell IT virtualization server farm was standardized on Dell PowerEdge 6650 servers. These servers each have four Intel Xeon® processors, 16 GB of RAM, two Emulex host bus adapters for connection to the back-end storage area network (SAN), and four Intel Gigabit Ethernet[2] network interface cards (NICs) for network connections. One NIC is dedicated for access to the ESX Server service console, one is dedicated for the VMware VMotion™ feature, and the remaining two are teamed and dedicated for use by the VMs. The SAN is a Dell/EMC CX700 storage array, with most of the VM disk files residing on RAID-5 logical units. Figure 1 shows the Dell IT virtualization server farm.

All servers in the farm run VMware ESX Server 2.5. The servers are divided into groups of 20 for manageability, but they share the same SAN and are managed by a single PowerEdge 2650 server running VMware VirtualCenter 1.2. The VirtualCenter database resides on a clustered instance of the Microsoft® SQL Server™ platform. If a failure occurs on the VirtualCenter server, a VM is installed with VirtualCenter and can be attached to the clustered database instance. Because the VirtualCenter application cannot be clustered, this process helps mitigate the amount of possible downtime caused by server-level failure. The redundancy at the database level is much higher than at the server level, because the database runs on a clustered SQL Server instance.

Virtualization best practices

As virtualization technology for industry-standard servers began to mature, Dell IT evaluated it for use in various roles. It became clear early on that test and development environments were well suited for virtualization. In fact, some groups within Dell IT had already started projects to use VMware products for that very purpose.

To gain control of this technology and use it efficiently, Dell IT developed two virtualization policies. The first policy was to deploy

[1] For more information about the Dell scalable enterprise vision, see "Dell Scalable Enterprise Architecture," by Jimmy D. Pike and Tim Abels, www.dell.com/downloads/global/vectors/2005_scalable_enterprise.pdf; and "Server Virtualization in the Scalable Enterprise," by Jimmy D. Pike and Drew Engstrom, www.dell.com/downloads/global/solutions/server_virtualization.pdf.

[2] 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.

virtualization in a controlled, managed server farm. The second policy was to use existing processes and systems whenever possible, and to only create new processes that could take advantage of virtualization capabilities or features. Dell IT used these two policies to implement test and development environments in its server farm.

No backups for test and development virtual machines

In the Dell IT server farm, virtualization is not implemented as a backup or disaster recovery solution; instead, it is used to efficiently support and enable the test and development process. Existing systems are in place for backup and code management. To help simplify the management of the server farm, none of the VMs have a backup. Engineers and developers using the VMs are informed that these systems are not file shares or code repositories; the existing file shares and code repositories are accessible from the VMs over the network and should be used. Both the file servers and code repository servers have backups performed on a regular basis.

Even though the VMs are not backed up, the VMs themselves reside on the SAN, which is highly available. In the few instances when a server failure has occurred, all affected VMs were recovered onto other servers in the farm within a few hours.

Server provisioning with Altiris Deployment Solution

For all Microsoft Windows® OS–based server installations, Dell IT uses a scripted installation process that includes the most recent security patches, antivirus software, systems management agents, and other corporate standards for servers. To help ensure that all VMs have the same level of security patches and adhere to all other corporate software standards, Dell IT uses the Altiris® Deployment Solution™ scripted installation to install the master VMs.

Clone copies are then made from these master VMs using either the VirtualCenter wizard for single copies or the vmclone script (available at www.dell.com/downloads/global/solutions/vmclone.zip) to automatically create large numbers of copies. To help ensure that the security level of the master VMs remains up-to-date, they are left powered on when not being cloned. This procedure allows the standard Microsoft Systems Management Server (SMS) updates to be applied as they are pushed out to the master VMs, so that the next time a clone is made it includes the latest updates. Additionally, to help ensure that the master VMs have the latest builds, they are rebuilt with the Altiris scripted installation every quarter.

Dell IT has implemented virtualization by using ESX Server to support more than 1,000 test and development environments on fewer than 100 physical servers.

In addition to provisioning the VMs, Altiris Deployment Solution is used to deploy the physical servers running ESX Server. The Altiris job installs ESX Server using a scripted installation and includes the Dell OpenManage™ Server Administrator application for hardware management and monitoring. This process helps make deployment of new servers in the server farm quick and easy. Figure 2 illustrates the master VM update process.

Review committee

For several years, a review committee within Dell IT has been responsible for approving requests for new hardware. This committee now includes VMs as part of its review and approval process. The default implementation for any test and development request is a VM, but exceptions are made when a case can be made for acquiring new hardware.

Not all requests for VMs are approved, however, because a virtualized test and development environment still requires server and storage resources. All groups that benefit from the virtualization farm also contribute to the farm's budget. These groups have

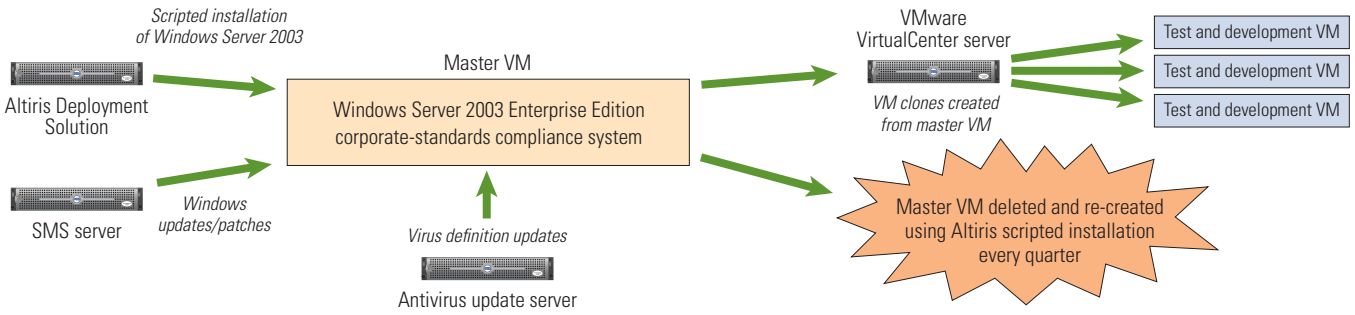

Figure 2. Master VM update process

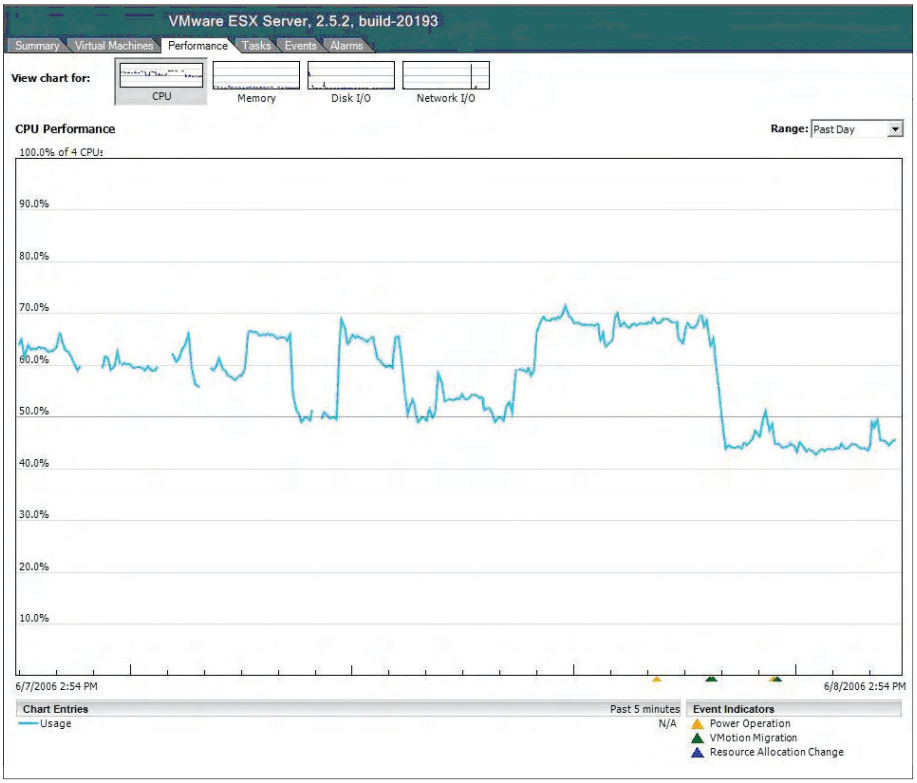

Figure 3. VMware VirtualCenter interface showing CPU performance for a Dell IT server running ESX Server 2.5

representation on the review committee, which allows them to control how their money is allocated.

Standardized virtual machine configuration

To help simplify operations, Dell IT has defined a standard VM configuration to be used unless the review committee grants an exception. The standard VM configuration incorporates dual virtual processors, 512 MB of RAM, a 16 GB virtual disk that is split evenly onto the C:\ and D:\ drives, and the Microsoft Windows Server® 2003 Enterprise Edition OS with the latest security patches and service packs.

Using a standard VM configuration helps simplify cloning and recovery. The vmclone script can quickly create a large number of the same type of VMs, which can help dramatically reduce the administrative time required to fulfill requests for new VMs. If a server failure occurs, VM recovery is also simplified when they are all the same configuration, because administrators do not need to track which VMs were assigned different amounts of RAM or different hard disk sizes.

Workload monitoring

One of the key advantages of using ESX Server is the ability to use the VMotion live-migration feature to move an active VM from one physical server running ESX Server to another. This feature allows the load to be redistributed across the farm by moving VMs from heavily loaded servers to less-loaded servers. Dell IT uses VirtualCenter to identify the heavily loaded servers and to initiate VMotion as needed to move VMs around the server farm. Figure 3 shows the effects on CPU performance when VMotion is used to reduce the load on one of the Dell IT servers running ESX Server 2.5. As VMotion is initiated (denoted by the green triangles along the bottom of the graph), the load on this system is reduced because some of the busy VMs are moved to less-loaded servers.

Dell IT's standard sizing for the farm's four-processor servers is 4 VMs per processor or 16 VMs per server. The number of VMs on individual servers usually changes as the servers are monitored and VMs are redistributed based on load, but in general the number is kept as close to 16 per server as possible.

Plans for emerging virtualization technologies

Dell IT plans to take advantage of emerging virtualization technologies as part of an initiative to expand the usage of virtualization within Dell IT. These emerging technologies include automated physical-to-virtual server conversion, automatic load balancing, and disaster recovery.

Automated physical-to-virtual server conversion

Many older servers in Dell data centers around the world are well suited for virtualization, but legacy applications are difficult to move directly onto new hardware for a variety of reasons. Physical-to-virtual server conversion tools from VMware and third parties such as PlateSpin can help accelerate the server consolidation effort underway in Dell data centers by allowing these systems to be quickly converted into VMs. These conversion programs use many of the tools that standard imaging programs use to capture the physical server image. Typically, the physical server is Preboot Execution Environment (PXE) booted to an imaging server, which then downloads and boots a Windows Preinstallation Environment (WinPE) or Linux® OS kernel on the server to be converted. Next, the disk image is captured, and any drivers needed to run under ESX Server are added to the image. The final image is copied to the data store managed by ESX Server, which starts up the image

as a VM. The entire process may be scripted and run on multiple servers at once. Other variants of this technology can be used for virtual-to-physical conversion, which assists in creating a physical production server from a virtual development server, and virtual-to-virtual conversion, which facilitates converting VMs from one virtualization server to another (for example, from a server running Microsoft Virtual Server to one running ESX Server).

Automatic load balancing

To help decrease the amount of manual intervention required to load balance the test and development server farm, Dell IT has begun testing with the next version of ESX Server, ESX Server 3. This version introduces the Distributed Resource Scheduler (DRS) feature, which can be used to manage virtual server loads across the many physical servers in a server farm. DRS is a policy-based tool that uses VMotion to move VMs to underutilized servers. Administrators can set the degree of automation desired, from totally automated (with DRS moving the VMs without any administrator intervention) to totally manual (with DRS notifying the administrator when an ESX Server host is overloaded and recommending another physical server for the high-utilization VMs).

Disaster recovery

To help improve VM availability on the server farm, Dell IT is evaluating another feature introduced in ESX Server 3, VMware High Availability (VMware HA). This feature detects physical server errors and warnings and moves VMs from a failing ESX Server host to other hosts in the farm. As with DRS, administrators can set the degree of automation. Dell IT anticipates that this feature could be useful in the test and development environments currently targeted by Dell IT for consolidation.

Scalable virtualization for the future

Dell IT has implemented a scalable solution for a large test and development environment using virtualization. This implementation has helped greatly reduce the number of physical servers required while also increasing the speed at which new test and development projects can be started compared with previous IT environments. As additional virtualization capabilities are developed, Dell IT intends to look for areas beyond test and development where virtualization can further increase cost-effectiveness and operational benefits.

TALK BACK

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

Todd Muirhead is a senior engineering consultant on the Dell Scalable 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.

Rick Merino is a senior systems engineer on the Dell IT Global Technology Engineering Services Core Engineering team.

Dave Jaffe, Ph.D., is a senior consultant on the Dell Scalable 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 at San Diego.

Jon Mercado is a senior systems engineer in the Global Platforms Engineering group at Dell. He specializes in server consolidation through virtualization and is the primary engineer for Dell IT virtualization infrastructure.

FOR MORE INFORMATION

Dell Scalable Enterprise Technology Center:
www.dell.com/setc

Dell virtualization solutions:
www.dell.com/virtualization

VMware:
www.vmware.com

Microsoft SQL Server 2005 Virtualization in the Dell Scalable Enterprise

The Dell scalable enterprise architecture uses industry-standard components to help simplify operations, improve resource utilization, and scale cost-effectively. Virtualization technology is advancing IT infrastructure standards by abstracting software from hardware, which can facilitate deployment and management of software in the scalable enterprise. To demonstrate this, the Dell Scalable Enterprise Technology Center deployed Microsoft® SQL Server™ 2005 software in a virtualized environment.

BY TODD MUIRHEAD

Related Categories:

Dell Scalable Enterprise Technology Center

Microsoft SQL Server 2005

Scalable enterprise

Virtualization

VMware

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

Virtualization can enable enterprise IT organizations to easily create new virtual systems in minutes, but such systems should still be deployed in a controlled and manageable manner, similar to how physical systems are deployed. The Dell Scalable Enterprise Technology Center used a farm of servers running VMware® ESX Server software to manage a deployment of Microsoft SQL Server 2005 virtual machines (VMs). This article provides an introduction to current virtualization products and SQL Server 2005, and examines techniques for managing SQL Server 2005 in a virtualized environment. It also includes best practices for managing a virtualized server farm, example usage scenarios for virtualization and SQL Server 2005, and basic performance results from testing SQL Server 2005 in a virtualized environment.

Defining the Dell scalable enterprise

The Dell scalable enterprise strategy focuses on standardizing core elements of IT infrastructures to help simplify operations, improve utilization, and enable cost-effective scaling. Using virtualization technology to implement Microsoft SQL Server 2005 can help deliver all three of these benefits:

- **Simplified operations:** Managing VMs can be, in some ways, much easier than managing physical

systems. Adding a virtualized layer provides a key advantage—portability. A VM can be moved from one server to another without any changes in configuration because the virtualization layer remains the same. Additionally, deploying VMs is a software-only task once the physical hosts running the virtualization layer are set up. Administrators can also quickly create a clone of an existing VM to take advantage of an existing application setup without having to perform a second installation and configuration on another system.

- **Improved utilization:** One common use of virtualization is server consolidation, which is intended to improve utilization of existing server resources. The ability to run many VMs on a single physical server can help greatly reduce the number of physical servers while improving utilization.
- **Cost-effective scaling:** Virtualization can help make adding capacity easy and cost-effective. The high level of utilization of existing resources and the ease of management allow each new server added to the virtualization server farm to be fully utilized quickly—which, because of the portability advantages, enables IT functions using virtualization to grow quickly and take advantage of the latest technology.

Understanding virtualization technology

Different forms of server virtualization have been available on expensive proprietary systems for decades. In the last five years, server virtualization technology for industry-standard Intel® processor–based servers has matured and is now widely used. The availability of virtualization technology on industry-standard Intel Xeon® processor–based servers has made this cost-effective technology accessible to enterprises of all sizes.

Microsoft and VMware are leading vendors of virtualization products today. Both have products that are installed as an application on top of an existing OS that then allows the creation of VMs on top of that OS. The Microsoft product, Microsoft Virtual Server, and the comparable VMware product, VMware Server, are both available as free downloads.

VMware also has a virtualization product called ESX Server, which loads directly onto the hardware instead of being hosted on an OS. Administrators can use ESX Server to create VMs that host guest operating systems without the overhead of a host OS. ESX Server provides features such as wizard-based cloning, centralized management of VMs, and live migration of VMs from one physical server to another.

ESX Server 3 is sold as part of a bundle called Virtual Infrastructure 3 Enterprise Edition, which includes ESX Server 3, VirtualCenter 2, the VMotion™ feature, the Virtual SMP™ feature, resource load balancing, consolidated backup, and high availability. This bundle of products is designed to allow flexible use of virtualization and provides many features needed to support production-level applications.

TALK BACK

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

Using Microsoft SQL Server 2005 in a virtualized environment

Microsoft SQL Server 2005 introduces many features and tools that administrators may find useful. As with any migration of critical data, however, administrators should thoroughly test these features before upgrading existing SQL Server 2000 databases or deploying new databases. Key features included with SQL Server 2005 are enhanced support for business intelligence, tight integration with Microsoft Visual Studio® 2005 software, and database mirroring. Virtualization can be used to create a test and development environment for SQL Server 2005 where these features can be tested.

To demonstrate the advantages of managing virtualized SQL Server 2005 on Dell™ PowerEdge™ servers, the Scalable Enterprise Technology Center set up VMs (see Figure 1) to test scenarios such as creating test and development environments, implementing disaster recovery, and consolidating servers. This section discusses how such environments can be monitored; the "Studying example virtualization usage cases" section describes each of these scenarios.

Management tools for virtual machines

Microsoft and VMware both provide tools to manage the VMs created using their respective products. The Microsoft Virtual Server administrative Web console provides tools for creating and modifying

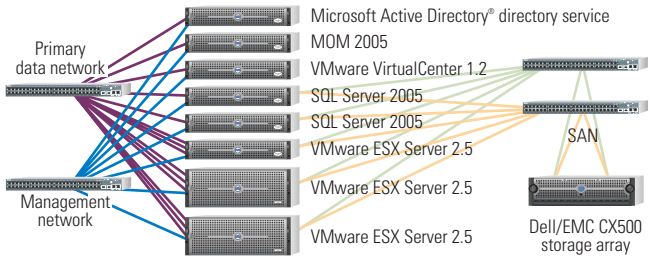

Figure 1. Dell Scalable Enterprise Technology Center virtualization architecture

the VMs on an individual server. VMware Server comes with a similar administrative console to manage VMs on an individual server; ESX Server also has a Web interface for this purpose. Additionally, VMware offers VirtualCenter, which can be used to manage and monitor VMware-based VMs across an enterprise. VirtualCenter offers more advanced management capabilities than the previously mentioned management tools, such as the abilities to clone VMs and to move VMs from one physical server to another.

One of the primary functions of MOM is to monitor the virtualized operating systems and applications running on the VMs. It can also monitor the virtualization layer itself with the addition of Virtual Server or VMware management packs.

The Dell server management tool, Dell OpenManage™ Server Administrator (OMSA), can be loaded on servers using Microsoft Virtual Server, VMware Server, or VMware ESX Server. With the first two products, OMSA is installed on the host OS and interfaces with the Dell server hardware as usual. In the case of ESX Server, the Linux® OS version of OMSA is used, but a few changes have been made to account for the differences between Linux and ESX Server.[1]

Microsoft Operations Manager (MOM) also plays a significant role in managing these virtual environments. One of the primary functions of MOM is to monitor the virtualized operating systems and applications running on the VMs. It can also monitor the virtualization layer itself with the addition of Virtual Server or VMware management packs. Administrators should handle many of the detailed virtualization operations using the tools included with the individual products, but should use MOM to integrate the virtualization products into the same management process as other systems. Doing so can provide administrators with a comprehensive view of the entire IT infrastructure or data center.

Microsoft Operations Manager and associated management packs

The management tools that are included with ESX Server and Virtual Server are strictly for managing the virtual infrastructure used to run the VMs. To manage the operating systems and applications running on these VMs, administrators should apply the same methodology used in the rest of the scalable enterprise architecture: a single enterprise management console, standard deployment practices, and structured monitoring and alerting.

MOM provides a platform for administrators to manage not only the virtualization infrastructure (the hardware and software supporting virtualization), but also the applications running within the virtualized environment. The Scalable Enterprise Technology Center configured MOM with the Virtual Server management pack and the eXc Software Virtual Agent with VMware management pack to monitor its environment. The SQL Server management pack was also loaded to support monitoring of the SQL Server 2005 databases loaded onto the VMs.

In addition to these management packs, the Dell OpenManage management pack allows MOM to receive alerts based on server hardware health. This functionality enables MOM to be aware of hardware-level issues in addition to OS- and application-level information, enabling administrators to obtain a comprehensive view of their environment. Figure 2 shows the management packs used in the Scalable Enterprise Technology Center for this project.

Planning and policies for virtual machines

Creating and deploying a new VM is relatively easy, and can cost much less to deploy than a physical server. These advantages do not remove the need for management, nor eliminate the cost of the physical server that is hosting the VM. IT organizations should deploy virtualization in a controlled, managed manner and treat VMs just like physical servers except when there is a significant reason to take advantage of a virtualization feature. The following policies can be established for a virtualization server farm:

- Justification for new VMs should be approved using the same process as for new physical servers.
- MOM should be used to monitor ongoing operations on every VM.

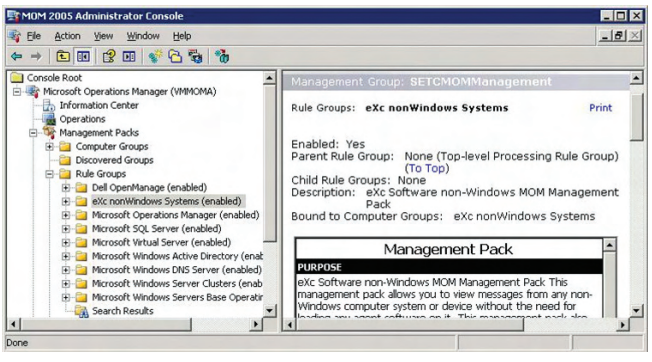

Figure 2. MOM 2005 Administrator console showing management packs for virtualization, SQL Server, and more

[1] For more information, see the *VMware ESX Server 2.5.2 Software for Dell PowerEdge Servers Deployment Guide,* by Dell Inc., www.dell.com/downloads/global/solutions/vmware_252_deployment_guide.pdf.

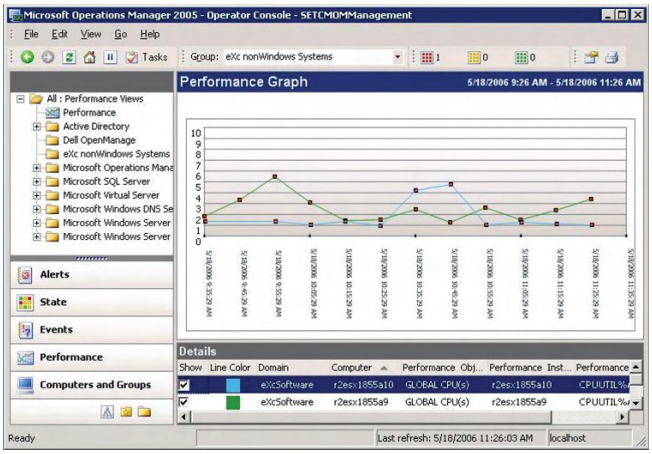

Figure 3. MOM 2005 Operator console showing CPU utilization of two servers running ESX Server

- VM cloning should be used whenever possible to help reduce the time required to configure a system.
- VMs should be the default option for a new server deployment, with exceptions made only when sufficient justification for a new physical server is provided.

Performance monitoring

Administrators can use several tools to monitor the performance of a virtualized environment. MOM provides a tool for overall performance monitoring, and can show performance for servers running ESX Server in the context of the entire environment (see Figure 3). VirtualCenter provides more detailed performance information than MOM, and should be used to obtain information about specific VMs and host servers. Additionally, administrators can use Microsoft Windows® Performance Monitor, the Linux version of the top tool, or other traditional OS-level tools within a VM to monitor application performance.

Studying example virtualization usage cases

Many usage models take advantage of the strengths of virtualization, including test and development, disaster recovery, and server consolidation. This section discusses these usage cases with SQL Server 2005 as the example application.

Test and development

SQL Server 2005 introduces several features, including integration with Visual Studio 2005 and support for the .NET common language runtime (CLR) within the database. This support allows database developers to use Transact-SQL, as with previous versions of SQL Server, as well as CLR languages such as the Microsoft Visual Basic® .NET or Visual C#® .NET languages to create their database applications.

To quickly create a large number of SQL Server 2005 test and development environments, the Scalable Enterprise Technology Center used a farm of servers running ESX Server to host a virtualized environment. A new clone was created in about 15 minutes using the existing Microsoft Windows Server® 2003 R2 OS master VM. SQL Server 2005 was then installed on the new VM, creating a SQL Server 2005 master VM to be used for creating all future SQL Server 2005 VMs. Installing SQL Server 2005 and creating the test database took approximately two hours. A script was then used to clone the SQL Server 2005 VM 16 times. Each clone creation took about 15 minutes to complete, with all 16 finishing in about four hours. To create and replicate a test and development environment for SQL Server 2005, all steps were completed in less than eight hours, with about half the time being handled by the vmclone script (available at www.dell.com/downloads/global/solutions/vmclone.zip).

A major feature introduced in SQL Server 2005, available with Service Pack 1 (SP1), is database mirroring. The Scalable Enterprise Technology Center set up and tested database mirroring with three SQL Server 2005 VMs running on the ESX Server farm. Database mirroring can be configured in several ways; in this example, a fully synchronous mirror with a witness was configured, which allows a SQL Server 2005 database to have changes written to both the primary and secondary systems at the same time and a third witness SQL Server 2005 instance to watch for failures and initiate automatic failover if necessary. If SQL Server is moved into production on a physical server, it may be beneficial in some environments to leave the secondary and witness databases on VMs. Doing so can help maintain high availability while using less hardware than if all three systems were moved to physical servers.

Creating and deploying a new VM is relatively easy, and can cost much less to deploy than a physical server. These advantages do not remove the need for management, nor eliminate the cost of the physical server that is hosting the VM.

Disaster recovery

Each VM is stored on the server as a set of files: a configuration file, a disk file, and potentially a redo file. The configuration file contains information such as the amount of RAM, number of network interface cards (NICs), and disk file location. The disk file is basically the hard disk of the VM. In the case of servers running ESX Server, disk files are typically given a .vmdk or .dsk

file extension. If the ESX Server incremental changes feature is used, a redo file stores all the changes made to the VM disk since the last time a snapshot was made. Because a complete VM is represented by these two or three files, it can easily be backed up.

But even more important for a disaster recovery scenario, the virtualization layer provides enhanced flexibility for the restore. Any server that runs ESX Server can be used as the server to which VMs are restored. Additionally, a server running ESX Server can be set up and waiting as a recovery server that can support systems that fail, which can help improve recovery time. Even if ESX Server must be installed first, the installation generally adds approximately 30 minutes to the recovery time.

SQL Server 2005 VMs are no different than other VMs in the case of disaster recovery. VMs can also use traditional backup agents and perform backups just like physical servers. These backups are typically performed for production databases to conform to existing policies and procedures already established in a data center for database backups. Additional LAN and storage area network (SAN) backup options are available for ESX Server, which can off-load server processing and help reduce the complexity that can be caused by backup agents installed on all VMs.[2]

Any server that runs ESX Server can be used as the server to which VMs are restored. Additionally, a server running ESX Server can be set up and waiting as a recovery server that can support systems that fail, which can help improve recovery time.

The features and capabilities of VMware High Availability (VMware HA), part of Virtual Infrastructure 3, allows for automatic restarting of VMs that are affected by the failure of an ESX Server system. Affected VMs can be automatically restarted on a remaining ESX Server system in the server farm with spare capacity, which can help improve availability.

Server consolidation

Server consolidation is perhaps the most obvious benefit of running multiple VMs on a single physical server. Likely candidates for this type of consolidation are existing low-utilization servers and applications that have low to moderate performance requirements. Many virtualization efforts start with these systems, and the reduction in total hardware requirements can help improve data center cost-effectiveness.

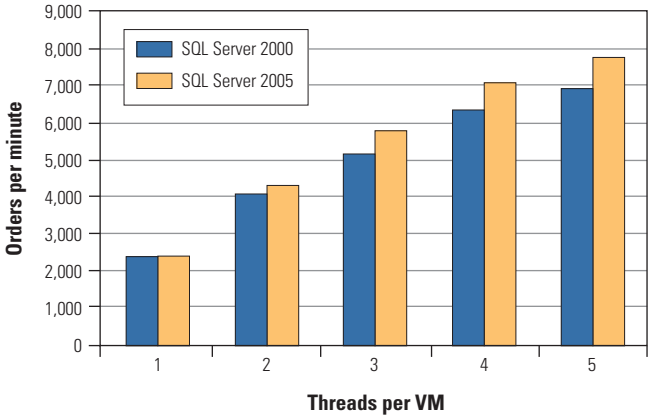

Figure 4. Orders per minute for SQL Server 2000 and SQL Server 2005

ESX Server provides a virtualized environment that can be used to run VMs at a reasonable level of performance. Dell has published several performance-based studies using ESX Server and a range of different types of VMs running a variety of applications that show high performance in a virtual environment.[3]

To investigate the performance of SQL Server 2005 in a server consolidation environment, in May 2006 the Scalable Enterprise Technology Center compared the performance of 15 VMs running SQL Server 2000 with 15 VMs running SQL Server 2005 on a Dell PowerEdge 2850 server with two dual-core Intel Xeon processors at 2.8 GHz and 8 GB of RAM. Each VM had the same configuration except for the version of SQL Server, with 512 MB of RAM, a 10 GB hard disk, and a vmxnet virtual Gigabit Ethernet[4] NIC. ESX Server 2.5.2 was installed on the PowerEdge 2850 server, and Windows Server 2003 Enterprise Edition with SP1 was the guest OS on all the VMs. The number of VMs running on the server was set at 15, which was based on findings in the Dell white paper "VMware ESX Server Performance Gains on Dell PowerEdge 2850 Dual Core Servers."[5]

The application used for testing was Dell DVD Store, which is available at linux.dell.com/dvdstore and has been released under the GNU General Public License. The DVD Store application includes build scripts and driver programs for testing on several databases, including SQL Server.

[2] For more information about these backup options, see "Using VMware ESX Server System and VMware Virtual Infrastructure for Backup, Restoration, and Disaster Recovery," by VMware, Inc., www.vmware.com/pdf/esx_backup_wp.pdf.

[3] For more information, visit www.dell.com/vmware and select "White Papers."

[4] 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.

[5] "VMware ESX Server Performance Gains on Dell PowerEdge 2850 Dual Core Servers," by Todd Muirhead and Dave Jaffe, Ph.D., www.dell.com/downloads/global/solutions/esx_2850_dualcore.pdf.

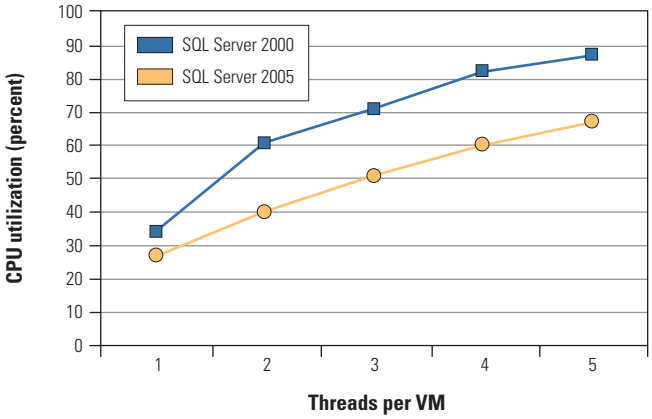

Figure 5. CPU utilization for SQL Server 2000 and SQL Server 2005

To simulate multiple small databases, the medium-size (approximately 1 GB) version of the DVD Store database was created. A separate PowerEdge 2650 server was used as the load driver system. A separate driver program thread was started for each VM simultaneously, and the results from all drivers were captured and totaled to obtain the orders per minute for the server. The ESX Server CPU utilization was monitored with the esxtop utility.

The results of this test show that both SQL Server 2000 and SQL Server 2005 achieved a high rate of orders (or transactions) per minute, but SQL Server 2005 did so more efficiently than SQL Server 2000. Under the same load, the SQL Server 2005 VMs processed more orders per minute at a lower CPU utilization than the SQL Server 2000 VMs. Figures 4 and 5 show orders per minute and CPU utilization for the 15 VMs running simultaneously on a single PowerEdge 2850 server.

Implementing virtualization in scalable enterprise architectures

Virtualization can provide a key infrastructure layer that enables scalable enterprise architectures, and the ability to easily provision and copy environments is well suited for test and development endeavors. Running multiple VMs on the same server can help improve efficiency of resource use, which makes server consolidation a key motivation for implementing virtualization. Crucial to server virtualization is the ability to use management tools such as Microsoft Operations Manager and VMware VirtualCenter to centralize and consolidate views of the environment, which can help improve management and overall cost-effectiveness.

Todd Muirhead is a senior engineering consultant on the Dell Scalable 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.

Reprinted from *Dell Power Solutions,* November 2006.

Scalable Enterprise Implementation Study:

How Dell IT Implements Microsoft SQL Server 2005 with Database Mirroring

The Microsoft® SQL Server™ database platform is an important component of the scalable Dell infrastructure and data centers. This implementation study discusses the manageability, availability, and performance advantages the Dell IT group gained by moving from SQL Server 2000 to SQL Server 2005.

BY TODD MUIRHEAD, SAJAL DAM, AND PATRICK ORTIZ

Related Categories:

Database

Microsoft SQL Server 2005

Scalable enterprise

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

Dell uses Microsoft SQL Server to support many applications throughout the company, including those that are integral to the Dell online presence and are used heavily by customers through the Internet. Because of its direct-customer model and large number of customers, Dell requires large, well-implemented IT infrastructures and data centers. Dell IT meets these requirements with careful planning and implementation of scalable architectures. Following the release of SQL Server 2005, the Dell IT group determined that the manageability, availability, and performance of its SQL Server–based applications could be greatly improved by moving from SQL Server 2000 to SQL Server 2005.

Microsoft SQL Server 2005 and the Dell scalable enterprise

Microsoft SQL Server 2005 introduces features and performance enhancements over the previous version of this software, SQL Server 2000, which enables it to better support scalable enterprise architectures. Running SQL Server 2005 on Dell™ PowerEdge™ servers is an example of using industry-standard components to build an effectively managed and highly available IT environment. Using such a scalable enterprise architecture can help provide simplified management, improved utilization, and cost-effective scalability.

Dell IT achieved gains in all three of these areas by implementing SQL Server 2005. The Microsoft Operations Manager (MOM) management pack for SQL Server 2005 provided Dell IT with detailed database monitoring capabilities, which help simplify management. SQL Server 2005 Service Pack 1 (SP1) introduces a database mirroring feature that helps improve availability without complex setup or difficult management. In addition, SQL Server 2005 has been streamlined, providing more efficient execution and increased throughput compared with SQL Server 2000—which can help improve resource utilization and in turn enable organizations to scale cost-effectively. The following sections focus on each of these three areas: simplified management with MOM management packs, improved availability with database mirroring, and enhanced performance over SQL Server 2000.

Simplified management with MOM management packs

Dell IT had managed previous implementations of SQL Server 2000 with various SQL Server–specific management tools that provided a range of information and varying levels of detail. Using these tools meant that even minor enhancements—for example, monitoring a database in a new way or using a new type of alert—required Dell IT to submit a request to the software vendor and then wait for the vendor to add the enhancement. This process was frustrating and resulted in slow responses to database problems, ultimately leading to increased support costs for the SQL Server databases.

Database mirroring allows synchronous or asynchronous mirroring of a SQL Server database across a standard TCP/IP connection without requiring shared storage between the two sides of the mirror.

As part of the SQL Server 2005 implementation, Dell IT decided to use MOM 2005 with the SQL Server management pack as its database monitoring tool. MOM 2005 provides a monitoring and management architecture implemented in and executed at a detailed level by individual management packs. The SQL Server management pack includes by default some of the simple database health monitoring capabilities that Dell IT wanted. But the key to MOM 2005 is the way that management packs can be easily changed and enhanced. The packs are based on open and accessible MOM scripts, allowing organizations to modify them to provide capabilities beyond those included by default. Dell IT wanted to monitor the expiration of SQL logins, additional details about database mirroring failures, and the size of the entire database rather than just individual files—so the group enhanced the SQL Server management pack by writing scripts to provide these capabilities.

TALK BACK

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

Improved availability with database mirroring

Database mirroring, which was introduced with SQL Server 2005 SP1, is the feature that prompted Dell IT to quickly implement the software. This feature allows synchronous or asynchronous mirroring of a SQL Server database across a standard TCP/IP connection without requiring shared storage between the two sides of the mirror. The most compelling detail of this feature for Dell IT was not its fast failover time, but its equally fast and seamless failback time.

SQL Server 2000 allows organizations to use log shipping or replication with a remote disaster recovery site. The Dell IT SQL Server 2000 implementation worked well, but had two major drawbacks. The first was the possibility for data loss between the failure at the primary site and the last log shipped to the disaster recovery site—so the potential for data loss was set by the interval between log updates or the rate of replicating queries. The second problem was the extremely long failback time. To fail back, Dell IT had to recover the primary site and then apply all changes since the failover. This process took many hours to complete and had to be repeated to bring the disaster recovery site back online and ready for future failovers. In addition, the setup and configuration processes were very complex.

By using SQL Server 2005 and database mirroring, Dell IT reduced the window for potential data loss to less than one second and shortened the failback process to less than one hour. The reason failback is so much faster than it was with SQL Server 2000 is that the process of synchronizing the two sides of the mirror is now managed by SQL Server 2005, meaning that the failback time is limited by the amount of data added between the time of failure and the time that the primary site is brought back online. Once the primary site is reconnected, the mirror is reestablished and the changes are synchronized by SQL Server 2005. After this step, a manual failover can be initiated, which takes only a few seconds to complete.

The applications using these databases automatically connect to the active side of the mirror based on the SQL Server connection string. Applications must be made aware of the mirror, which can

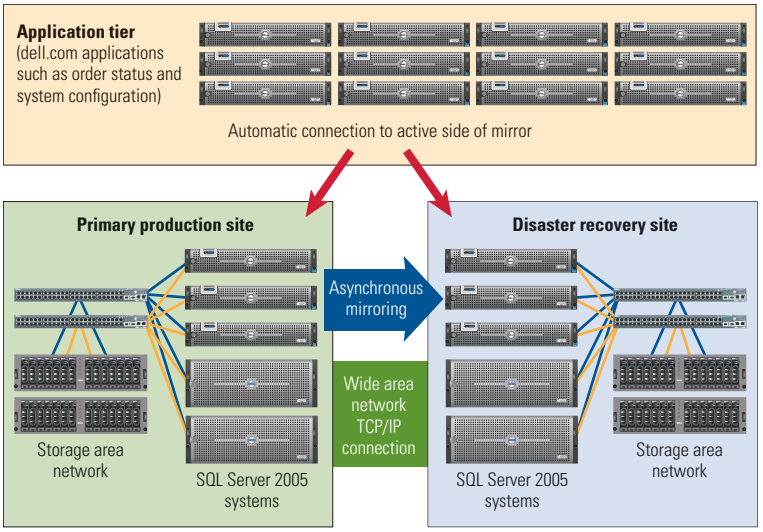

Figure 1. High-level architecture of the Dell IT SQL Server 2005 database mirroring deployment

be done simply by adding a failover partner parameter that specifies the other side of the database mirror in the connection string. Figure 1 shows the high-level architecture of the Dell IT SQL Server 2005 database mirroring deployment.

Enhanced performance over SQL Server 2000

An additional concern for Dell IT in moving to SQL Server 2005 was performance. To determine the relative performance of SQL Server 2000 and SQL Server 2005, in April 2006 Dell IT carried out performance tests using the Dell DVD Store open source test application and database (available at linux.dell.com/dvdstore) on Dell PowerEdge 2850 and PowerEdge 6850 servers. The DVD Store application and database represent a typical online store where users are logging in, searching for products, filling a shopping cart, and purchasing products. The DVD Store download includes a database, Web application, and load driver programs. In these tests, the Web application tier was not set up, but was simulated using one of the included load driver programs that run directly against the database.

For comparative purposes, Dell IT loaded the DVD Store database on PowerEdge 2850 and PowerEdge 6850 servers on both SQL Server 2000 and SQL Server 2005. Because the objective of the test was to measure performance increases when using SQL Server 2005 as compared with SQL Server 2000, the same configuration of PowerEdge 2850 and PowerEdge 6850 servers with external storage was used for both versions of the database software. Figure 2 provides detailed system configuration information.

These systems were put under load using the driver program, and both transactions per second (TPS) and response time were measured. The results show that SQL Server 2005 achieved between 30 percent and 67 percent higher performance than SQL Server 2000 on the same hardware.

The improvement range depends on the performance metric. Figure 3 shows the TPS results by percentage of CPU utilization for the DVD Store user database. At 40 percent CPU utilization on a PowerEdge 6850 server, SQL Server 2000 processed 58 TPS, while SQL Server 2005 processed 76 TPS—approximately a 30 percent increase in performance. A similar comparison could not be made with the data collected for the PowerEdge 2850 because the small overlap in common CPU utilization occurred after SQL Server 2000 was generating a very high response time.

Figure 3 also clearly shows that SQL Server 2005 can achieve higher levels of CPU utilization and more TPS than SQL Server 2000 on the same hardware—a result that is particularly pronounced on the PowerEdge 6850, where the utilization when running SQL Server 2000 only reached slightly more than 50 percent. One reason for these results is that SQL Server 2005 is more efficient at managing I/O and can avoid the I/O bottleneck that SQL Server 2000 typically encounters, and this increased efficiency results in higher TPS at the same CPU utilization as well as higher maximum TPS than SQL Server 2000.

Dell IT also tested TPS based on a given response time; the results are shown in Figure 4. At half a second for response

	Dell PowerEdge 2850	Dell PowerEdge 6850
Processor	Two Intel® Xeon® processors at 3.2 GHz with 1 MB L2 cache	Four Intel Xeon processors at 3.0 GHz with 8 MB L3 cache
Memory	8 GB	32 GB
Internal disks	Two 73 GB, 10,000 rpm Ultra320 SCSI disks	Two 73 GB, 10,000 rpm Ultra320 SCSI disks
Network interface cards (NICs)	Two 10/100/1,000 Mbps NICs (internal)	Two 10/100/1,000 Mbps NICs (internal)
Disk controller	PowerEdge RAID Controller (PERC) 4/ei	PERC 4e/Di
Fibre Channel host bus adapters (HBAs)	Two QLogic QLA2340 HBAs	Two QLogic QLA2340 HBAs
Storage area network logical units (all RAID-10)	• *Data:* Four arrays of 16 disks • *Temp:* One array of 16 disks • *Log:* One array of 6 disks	• *Data:* Four arrays of 16 disks • *Temp:* One array of 16 disks • *Log:* One array of 6 disks

Figure 2. Server and storage hardware configuration used in Dell IT performance tests with SQL Server 2000 and SQL Server 2005

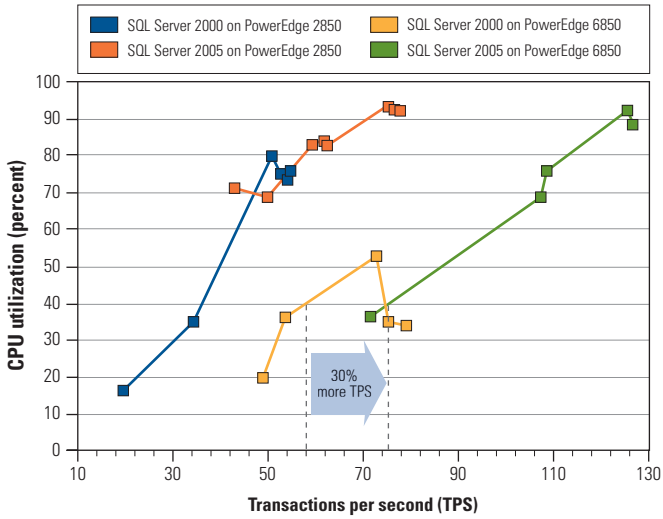

Figure 3. Transactions per second by percentage of CPU utilization for SQL Server 2000 and SQL Server 2005 on Dell PowerEdge 2850 and PowerEdge 6850 servers

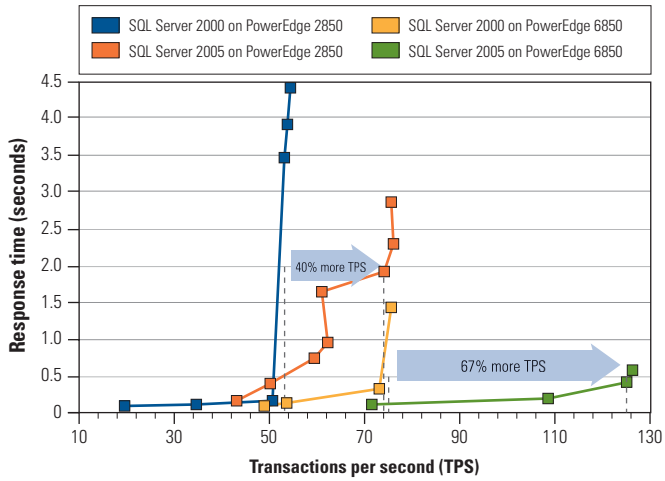

Figure 4. Transactions per second by response time for SQL Server 2000 and SQL Server 2005 on Dell PowerEdge 2850 and PowerEdge 6850 servers

time on the PowerEdge 6850, SQL Server 2000 processed 75 TPS, while SQL Server 2005 processed 125 TPS—a 67 percent increase in performance. Similarly, at two seconds for response time on the PowerEdge 2850, SQL Server 2000 processed 53 TPS, while SQL Server 2005 processed 74 TPS, for a 40 percent increase in performance.

In Figure 4, the sharp turn upward in each line clearly shows when the maximum TPS was reached for each of the four test cases. As additional load was added past this point, only response time increased, leaving TPS at approximately the same level. The figure also shows that, on both the PowerEdge 2850 and PowerEdge 6850 servers, SQL Server 2005 processed more TPS than SQL Server 2000.

Enhanced manageability, availability, and performance with SQL Server 2005

The move from SQL Server 2000 to SQL Server 2005 has greatly benefited Dell IT. The enhanced ability to monitor and manage SQL Server 2005 with MOM 2005 has provided database and systems administrators with tools that can help them increase database uptime. SQL Server 2005 database mirroring helps improve database availability and reduce support costs compared with traditional log shipping and replication implementations, because of streamlined failover and failback abilities. In addition, the Dell IT tests demonstrated that SQL Server 2005 can offer clear performance advantages over SQL Server 2000 when running on the same server and storage hardware. These benefits exemplify the advantages that can be gained when using industry-standard hardware and software to run an IT infrastructure.

Todd Muirhead is a senior engineering consultant on the Dell Scalable 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.

Sajal Dam is a senior manager in Dell IT. Sajal has an M.B.A. from Duke University and an M.Tech. in Computer Science from the Indian Institute of Science in Bangalore. He is the author of the book *SQL Server Query Performance Tuning Distilled.*

Patrick Ortiz is a senior database administrator in Dell IT. Patrick has a B.A. in Mathematics from the University of Texas at Austin. He leads the team of database administrators that manages the production databases.

FOR MORE INFORMATION

Dell Scalable Enterprise Technology Center:
www.dell.com/setc

Microsoft SQL Server 2005:
www.microsoft.com/sql

MOM 2005 management packs:
www.microsoft.com/technet/prodtechnol/mom/mom2005/catalog.aspx

MOM 2005 scripts:
www.microsoft.com/technet/scriptcenter/hubs/mom.mspx

Implementing Database Mirroring with Microsoft SQL Server 2005

The database mirroring feature of the Microsoft® SQL Server™ platform can help enterprises of all sizes make their databases highly available. This article discusses the architecture and management of SQL Server 2005 and the relative performance of different types of database mirroring.

BY TODD MUIRHEAD

Related Categories:

Database mirroring

Microsoft SQL Server 2005

Scalable enterprise

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

Microsoft SQL Server 2005 Service Pack 1 (SP1) introduces database mirroring, a software-based replication feature that can help provide flexible high-availability databases to enterprises of all sizes. Deploying this feature on Dell™ PowerEdge™ servers can help IT organizations improve database availability without requiring specialized, dedicated storage equipment.

Using SQL Server 2005 database mirroring in the scalable enterprise

The Dell Scalable Enterprise Technology Center implements the Dell scalable enterprise architecture strategy to demonstrate how enterprises can use this strategy to help achieve tangible benefits today. The core of the strategy focuses on using standardized hardware and software to help provide simplified operations, improved utilization, and cost-effective scaling. Deploying SQL Server 2005 database mirroring can help provide all three of these benefits when implemented and managed using a scalable enterprise approach:

- **Simplified operations:** Database mirroring configuration and management are straightforward. In addition, database mirroring is a much more streamlined and integrated part of SQL Server 2005 as compared with the log shipping and replication implementations used with previous versions of SQL Server.
- **Improved utilization:** The simple nature of the database mirroring feature allows failover and failback to occur quickly and easily, which can help improve database availability and allow

enterprises to allocate more server resources for database processing and less for recovery and failover.

- **Cost-effective scaling:** Database mirroring has a low performance overhead and does not require additional specialized storage hardware, which can allow even small enterprises to take advantage of this feature using internal or heterogeneous storage systems.

TALK BACK

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

Database mirroring architectures: Synchronous and asynchronous

Before implementing SQL Server 2005 database mirroring on Dell PowerEdge servers, enterprises should understand the different setup and configuration options. SQL Server 2005 database mirroring supports both synchronous and asynchronous mirroring. Both configurations use a primary server where all clients access the database, and a mirror server that maintains a secondary copy of the database (see Figure 1). All communication and mirroring activity between the servers occurs over a standard TCP/IP connection. The mirror can function regardless of the type of storage used with either the primary or mirror server, but performance can be affected by both the type of storage and the number of disks used on both sides of the mirror.

With synchronous mirroring, the primary server does not commit an operation to the database until it has also been written to the mirror server, meaning no data should be lost if the primary server fails. The mirror then returns a write confirmation back to the primary server. The disadvantage of synchronous mirroring is that the primary server's performance can be limited by the time it takes for a database operation to be written to the mirror server. When the latency between the two servers is low, however, the performance impact is usually also low.

Synchronous mirroring also allows using a third SQL Server 2005 system as a witness to provide automatic failover, an option not available with asynchronous mirroring. This witness server monitors the primary server and initiates a failover if the witness and the mirror server lose their connection to the primary server but the witness remains connected to the mirror server. When database clients reconnect, they are automatically rerouted to the mirror server, which has become the new primary server.

With asynchronous mirroring, the primary server does not wait for the mirror server to commit operations to disk. The disadvantage of asynchronous mirroring is that data can be lost because of the delay between the primary server writing data to disk and the mirror server writing that data to disk. However, this delay is usually small, so the amount of data potentially lost in such a scenario is typically minimal.

Example scalable database mirroring implementation

To demonstrate how to create a scalable enterprise architecture, the Dell Scalable Enterprise Technology Center configured SQL Server 2005 database mirroring on a group of Dell PowerEdge servers. This setup used three Dell/EMC storage arrays to demonstrate how previous-generation storage can be used for the mirror and more powerful new-generation arrays can be used for the primary data store. Figure 2 shows the setup, including fundamental infrastructure systems that are part of the Dell scalable enterprise architecture and support the SQL Server 2005 implementation. For simplicity, this figure does not show redundancy at the level of the Microsoft Active Directory® directory service or Microsoft Operations Manager (MOM) 2005, although these levels of redundancy were implemented.

Storage for the database on the Dell/EMC CX3-80 and CX700 arrays used a RAID-10 configuration to help improve performance. The Dell/EMC CX500 array used a RAID-5 configuration to provide a

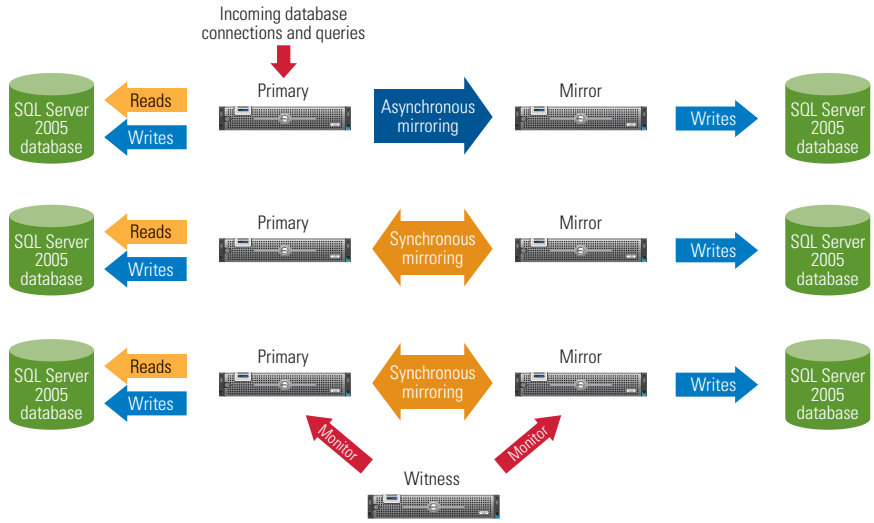

Figure 1. SQL Server 2005 synchronous and asynchronous database mirroring architecture

lower-performing but higher-capacity configuration compared with the RAID-10 setup. Although all the storage resided in the same storage area network (SAN), all mirroring was implemented at the host level through SQL Server 2005 database mirroring; alternatively, the storage could have been locally attached or located on three different SANs.

Other options for creating high-availability databases

Besides database mirroring, enterprises have several other options to enable high availability with SQL Server 2005, including log shipping and external storage array software such as the EMC® MirrorView™ application.

Log shipping. Log shipping is a more complex process than database mirroring, involving three steps to back up, copy, and restore the database logs from the primary server to the secondary server. Log shipping offers two advantages over database mirroring: it allows for multiple secondary servers, and any secondary server can be used for reporting. Because of these advantages, some enterprises may want to implement log shipping in combination with database mirroring. The disadvantage of log shipping is that the amount of data that can be lost is dependent on the interval between logs shipped to the secondary server, which is typically a much larger amount than asynchronous mirroring would lose.

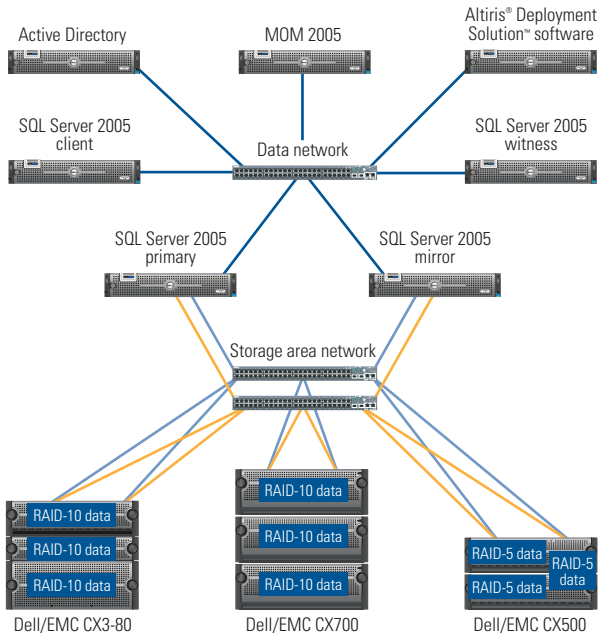

Figure 2. Dell Scalable Enterprise Technology Center setup for SQL Server 2005 database mirroring

External storage array software. Some external storage arrays, such as the Dell/EMC CX series of Fibre Channel storage, offer high-end mirroring software for SQL Server databases that are managed at the storage array level. Because the host is not directly involved in mirroring, these implementations can free up server resources for other tasks. EMC MirrorView, available on the Dell/EMC CX series, allows both synchronous and asynchronous mirroring with SQL Server 2005.

Managing SQL Server 2005 database mirroring

Management is a key component of a scalable enterprise architecture. MOM 2005 enables administrators to consolidate hardware, OS, and application management into a single console. MOM relies on management packs installed on the MOM server that contain rules and product knowledge for monitoring, managing, and resolving product-specific issues. MOM then applies these management packs to the appropriate systems based on discovery rules. For example, the Dell OpenManage™ suite management pack for Dell servers and the SQL Server management pack are automatically distributed to the MOM agent running on the servers in the mirror.[1]

Using event rules for database mirroring in the SQL Server 2005 MOM management pack

Shortly after SQL Server 2005 SP1 was released, an update to the MOM management pack for SQL Server added 18 default event rules specific to database mirroring. These rules, listed in Figure 3, provide monitoring for failover and synchronization states. The rules allow enterprises to monitor the standard types of events that occur with database mirroring and allow MOM to have a centralized log of all events related to database mirroring. Although these default rules are a good starting point, they do not include any predefined response action or a severity level higher than informational (except "Database mirroring has been terminated," which is predefined as a critical error). The MOM administrator determines which events warrant a response or a higher severity level.

Customizing MOM for database mirroring

Administrators can extend the basic MOM capabilities in a number of ways to help manage database mirroring, including adding tasks to MOM and customizing existing MOM event rules.

Adding a task to MOM. MOM tasks define a command, script, or program to be executed on the MOM management console or MOM-managed system. These operations can be started in context based on the system being managed from the MOM Operator console—for example, one task in the SQL Server 2005 MOM management pack allows administrators to open the SQL Server

[1] For more information about using MOM to manage SQL Server, see "Managing Microsoft SQL Server 2005 with Microsoft Operations Manager 2005 in a Dell Scalable Enterprise Architecture," by Todd Muirhead, *Dell Power Solutions*, August 2006, www.dell.com/downloads/global/power/ps3q06-20060360-Muirhead.pdf.

- Automatic failover
- Connection with mirror lost
- Connection with principal lost
- Database mirroring has been terminated
- Database mirroring is active
- Database mirroring is inactive
- Manual failover
- Mirroring suspended
- Monitor default instance
- No quorum
- Null notification
- Principal running exposed
- Synchronized mirror with witness
- Synchronized mirror without witness
- Synchronized principal with witness
- Synchronized principal without witness
- Synchronizing mirror
- Synchronizing principal

Figure 3. Event rules for database mirroring in the SQL Server 2005 MOM management pack

Management Studio on a managed system. Parameters can also be passed when starting a task based on the system currently being managed on the MOM Operator console. SQL Server 2005 database mirroring allows administrators to cause a manual failover by issuing a single SQL command:

```
ALTER DATABASE database name SET PARTNER FAILOVER
```

Administrators can create a task using the MOM Administrator console; Figure 4 shows the properties of an example task added in this console. This figure shows a command-line task created in the SQL Server 2005 group that specifies sqlcmd.exe as the executable and provides the necessary SQL statement to cause the mirroring failover. This custom task can then be executed from the MOM Operator console against any MOM agent–managed system to initiate database mirroring failover, which can be useful for an administrator preparing a system for upgrading or rebooting.

Customizing an existing MOM event rule. Administrators can also customize existing MOM event rules simply by modifying the rule properties. Administrators should make a copy of the original rule before editing, and disable the original version to help prevent duplicate alert events from appearing. Figure 5 shows an example

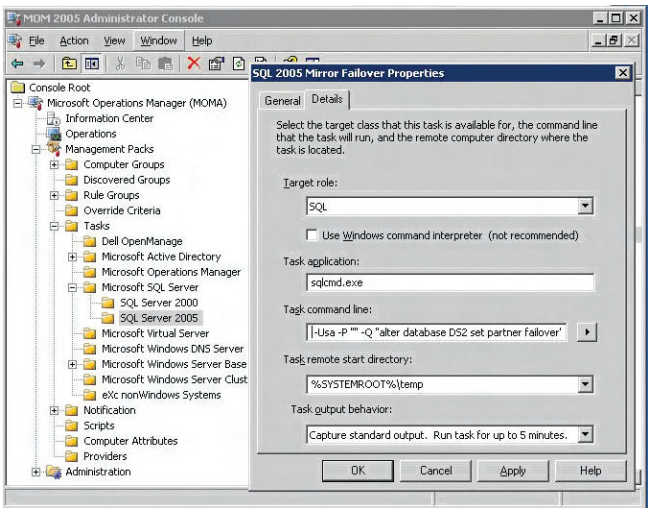

Figure 4. MOM Administrator console showing a new task that initiates database mirroring failover

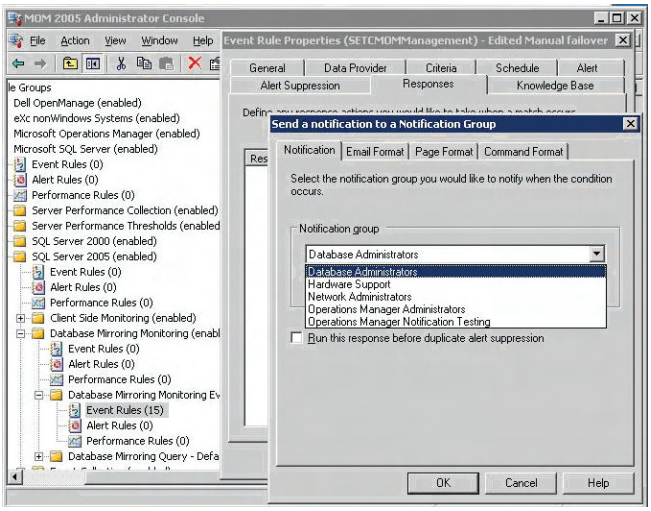

Figure 5. MOM Administrator console showing a notification event added to an existing event rule

rule modified in the MOM Administrator console to notify database administrators of a manual failover.

Administrators can easily raise the alert severity level for this event from informational to warning or critical on the Alert tab of the event rule. Increasing this level causes the entry in the MOM Operator console to appear with a yellow warning or red critical icon, which helps important events be displayed prominently.

Assessing SQL Server 2005 database mirroring performance

Whether to use either synchronous or asynchronous database mirroring is an important decision. For enterprises that cannot

tolerate any data loss, synchronous mirroring is preferable. For enterprises that can tolerate a small amount of data loss in exchange for the additional availability that database mirroring can provide, asynchronous mirroring can be a viable option.

While planning a database mirroring implementation, enterprises should evaluate the potential performance impact of choosing synchronous or asynchronous database mirroring. In July 2006, the Dell Scalable Enterprise Technology Center ran tests to provide sample performance results for synchronous and asynchronous database mirroring with a specific application, a specific set of servers, and three types of storage subsystem configuration.[2]

Test configuration

The test configuration used two Dell PowerEdge 2850 servers, each with two Intel® Xeon® processors at 3.6 GHz, 8 GB of RAM, dual QLogic QLA2340 Fibre Channel host bus adapters, and the Microsoft Windows Server® 2003 Release 2 (R2) Enterprise x64 Edition OS. A PowerEdge 6650 server was used as a driver system to stress the database during testing.

Before deploying database mirroring, enterprises should understand the differences between synchronous and asynchronous mirroring, and balance their performance requirements and sensitivity to data loss against potential improvements in availability.

The tests also used three Dell/EMC storage arrays to simulate differences that may exist between primary and disaster recovery data centers. The primary site used a Dell/EMC CX3-80 array with 15,000 rpm disks configured into three RAID-10 logical units (LUNs) with 10 disks each for data, and two RAID-1 LUNs with 2 disks each for logs. The first of the two mirror server arrays used a Dell/EMC CX700 array with the same LUN configuration as the CX3-80 but with 10,000 rpm disks. The second mirror server array used a Dell/EMC CX500 array with 10,000 rpm disks configured into three RAID-5 LUNs with five disks each and two RAID-1 LUNs with two disks each. The primary site had the highest-performing storage of the three arrays, the CX3-80; the two mirror site configurations represented two lower-performing configurations.

The test application used for performance evaluation was the Dell DVD Store (available at linux.dell.com/dvdstore), an open source test database and application that simulates an online DVD store. The DVD Store package includes database build scripts, database load scripts to populate the database, programs to create the data to be loaded, and a set of programs that simulate customers searching for and purchasing items.

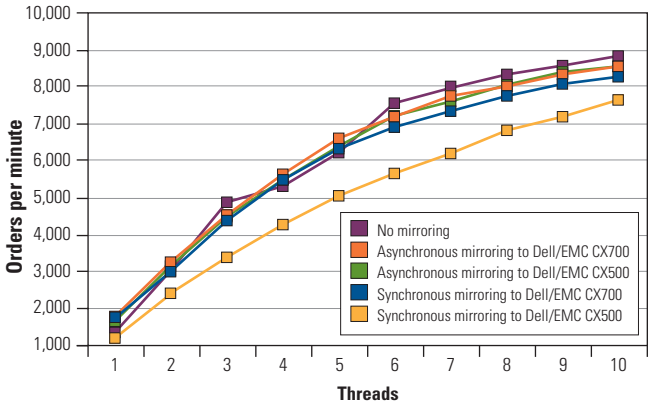

Figure 6. Orders per minute by number of threads with SQL Server 2005 database mirroring

These tests used the large version of the DVD Store database, which is approximately 100 GB in size and comprises 1 million products, 200 million customers, and one year of order history data containing about 600 million rows. The database was built on the primary server using the DVD Store build scripts. Before creating a mirror with the other server, a baseline test was run to provide a point of comparison for the tests with mirroring enabled. A simple script was created that executed the DVD Store driver program using from 1 to 10 threads. Each test with a given number of threads was repeated three times, and then a database cleanup was run. The three results were then averaged to calculate the orders per minute with a given number of threads.

To create the mirror, the test team backed up the database and restored it with the No Recovery option to the mirror server. Both synchronous and asynchronous mirror relationships were created and tested using the same script as the baseline test, resulting in orders per minute for 1 to 10 threads.

Test results

Figure 6 shows the results for all five test scenarios. Although the differences between the test configurations were relatively small, the test results demonstrate a clear and predictable performance hierarchy depending on the LUN and mirroring types. With six or more threads, the baseline test with no mirroring maintained the highest performance. The two asynchronous mirroring configurations provided the next-highest performance, with the CX700 mirror outperforming the CX500. Finally, the two synchronous

[2] Although these results may be similar to what others would find with the same type of tests, changing any of the individual elements could change the results. These tests serve only as an example and may not reflect actual performance in different environments.

configurations had the lowest performance, again with the CX700 outperforming the CX500.

The interesting part of these test results is how close in performance the five configurations were—even the synchronous CX500 RAID-5 mirror configuration is close, which is surprising given that it has approximately half the disks of the higher-performing CX700 RAID-10 mirror configuration. This similarity in performance can be ascribed to two factors that helped improve mirroring performance. The first is that both servers in the mirror were connected to the same switch residing in the same room, which led to very low latency and allowed the synchronous mirroring to avoid what could have been a large latency penalty if the primary and mirror servers were located hundreds or thousands of miles apart. The second factor is the percentage of writes in the workload. The DVD Store application simulates customers searching for titles and looks up customers' previous purchases, both of which are read-only queries that consume a significant amount of resources on the primary server. However, these read-only queries are not run on the mirror; only the writes are replicated. Enterprises using databases with a low write percentage may be able to support the mirror server with fewer or lower-performing disks than they use with the primary server—although doing so could cause problems following a failover in which the mirror server becomes the primary server.

Creating high-availability databases with SQL Server 2005 database mirroring

The Microsoft SQL Server 2005 SP1 database mirroring feature can help enterprises of all sizes make their databases highly available, and MOM 2005 with the SQL Server and Dell OpenManage management packs can help them create a highly manageable solution with integrated tasks and notification capabilities. Before deploying database mirroring, however, enterprises should understand the differences between synchronous and asynchronous mirroring, and balance their performance requirements and sensitivity to data loss against potential improvements in availability.

Todd Muirhead is a senior engineering consultant on the Dell Scalable 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.

Online Book Excerpt:

Managing Windows and Virtualization with MOM 2005

Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise documents best practices, operational techniques, and practical advice for using MOM 2005 to manage Microsoft® Windows® OS–based solutions on industry-standard platforms such as Dell™ PowerEdge™ servers and Dell/EMC storage. This Chapter 2 excerpt features a guided tour that shows how MOM resolves a Dell OpenManage™ event.

BY TIM ABELS

Related Categories:

Dell PowerEdge servers

Dell Scalable Enterprise Technology Center

Enterprise management

Microsoft Operations Manager (MOM)

Microsoft Windows

Systems management

Virtualization

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

By optimizing management systems and implementing best practices, administrators can manage large-scale data centers using only a fraction of the total number of nodes for resource management and orchestration. For example, a data center encompassing tens of thousands of managed nodes and storage can be controlled using only a few hundred servers. Following the same best practices can lead to similar advantages for any size organization, avoiding rework and inefficiency while enhancing scalability.

Chapter 2 of *Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise*, "MOM 2005 Overview: Managing Windows and Virtualization," describes best practices for achieving these benefits by using MOM to manage Microsoft Windows–based systems and virtualization environments enabled by Microsoft Virtual Server and VMware® ESX Server software.

Best practices for MOM 2005

The Dell Scalable Enterprise Technology Center Series reveals best practices for the integration of complete solution stacks—including managed application, OS, virtualization, server, and data center infrastructure—to help simplify operations, improve resource utilization, and scale out the IT infrastructure quickly, flexibly, and cost-effectively.

To learn more about the Dell Scalable Enterprise Technology Center and download the full text of Chapters 1 and 2 from *Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise,* visit www.dell.com/setc.

Implementing the Dell Unified Manageability Architecture

The Dell Unified Manageability Architecture (UMA) models systems management of agent-managed platforms including servers, storage, networks, clients, printers, and their subcomponents within an embedded platform. Six UMA layers define a comprehensive taxonomy of MOM management, emphasizing both the integration points between MOM and other systems managers and between MOM and the managed resources. Two UMA layers, resource management and optional orchestration, address

DELL AND COX ARE KEEPING

6.6 MILLION CUSTOMERS TUNED IN, IN TOUCH AND CONNECTED.

COX COMMUNICATIONS' DATACENTER OF THE FUTURE.

When Cox, an industry-leading broadband communications company, needed a new IT infrastructure to handle their mission-critical operations, they partnered with Dell. The Dell solution handles the most sensitive data Cox uses to run its business from core financial records, to supply chain management, to compliance and more. With Dell, Cox Communications is getting the technology and services they need to connect with their customers. Isn't it time we did the same for you?

Click www.dell.com/coxcommunications

Call (toll free) 1.866.212.9335

Guided tour: MOM resolution of a Dell OpenManage event

Figure 1 shows the MOM management flow following a loss of redundant power, such as an accidental unplugging. MOM helps automate and accelerate the detection, notification, diagnostic, containment, repair, and reporting processes.

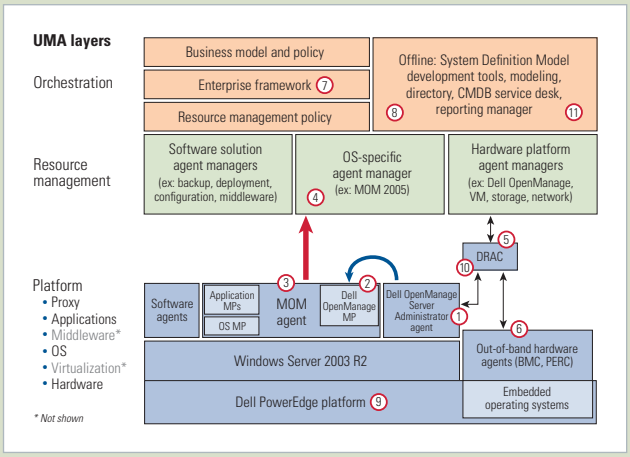

Figure 1. MOM management flow in the Dell Unified Manageability Architecture orchestration, resource management, and platform layers

1. The **Dell OpenManage Server Administrator agent** detects error events—for example, the loss of redundant power.
2. The **Dell OpenManage MP** translates error events into MOM alerts and provides useful scripts.
3. The **MOM agent** notifies the OS-specific agent manager, in this case the MOM management group server.
4. The **OS-specific agent manager (MOM 2005)** sends a notification to the operator.
5. The operator uses the MOM Operator console to launch the **Dell Remote Access Controller (DRAC)** console and Dell OpenManage Server Administrator for remote diagnostic confirmation. The operator also sets the alert status to New in the Operator console.
6. If the remaining power was lost, the DRAC performs its diagnosis using the **out-of-band hardware agents**.
7. An **enterprise framework** connect using the MOM Connector Framework (MCF) may correlate this MOM alert with other resource manager alerts and determine that the event is isolated.
8. The operator launches a third-party **configuration management database (CMDB) service desk** trouble ticket, using the MCF or Skywire Software iWave Integrator, to dispatch a local administrator to reconnect the power, and updates the alert status to Acknowledged.
9. The local administrator physically reconnects the redundant power on the **Dell PowerEdge platform**.
10. The operator repeats step 5 to verify the resolution with the **DRAC** and Dell OpenManage Server Administrator, closes the alert by setting its status to Resolved, and adds any useful process insights to the Operator console under the Company Knowledge tab.
11. A **reporting manager**, such as Microsoft System Center Reporting Manager, may publish a report on most of the preceding steps, which is then available to subscribers.

the management system aspect; four UMA layers address managed-resource aspects, including device-level instrumentation: platform, logical mapping, aggregation, and access. MOM management groups are part of the resource management layer, while MOM agents are hosted in the platform layer and use management packs (MPs) for gathering alerts from platform sublayers. The "Guided tour" sidebar in this article illustrates how these elements function when using MOM to resolve a Dell OpenManage event.

Managing virtualization with MOM

Virtualization is becoming an essential technology for optimizing resource utilization and data center management. Adding hardware virtualization as a platform layer enables organizations to benefit from several management scenarios:

- Consolidating servers by running multiple operating systems on the same production system to help improve resource utilization and contain future resource requirements
- Optimizing development and testing with multiple snapshots, simulating complex multi-tier applications, and performing copy-exact operations throughout the production life cycle
- Simplifying deployment and management with a single virtual machine (VM) image file that can run across a broad range of server configurations

When implementing the preceding scenarios, enterprises should follow best practices for managing both physical and virtual resources and minimize differences between the two.

Planning for the future

Microsoft has announced plans for future versions of MOM and its strategy for the System Center family, which includes many new products and upgrades. For more information about how the Microsoft System Center management architecture relates to the Dell UMA model, download the full text of Chapter 2 at www.dell.com/downloads/global/power/mom_2005_dellse_abels_ch2.pdf.

Tim Abels is a senior software architect on the Dell Scalable Enterprise Technology Center team and Dell 2005 Co-Inventor of the Year. Tim has an M.S. in Computer Science from Purdue University.

PURE DELL OPENMANAGE NEWSLETTER

Flexible Management for the Scalable Enterprise

November 2006

By Edward Reynolds, Senior Manager, Systems Management Product Marketing, Dell Inc.

View from the Top

Extending the Scope of Standards-based Management

Dell OpenManage™ IT Assistant 8.0 combines standards-based performance monitoring and device support to give you visibility into the health and status of heterogeneous servers, Dell/EMC networked storage devices, clients, and printers—making it a capable management application for small to medium businesses as well as departments within large enterprises. And the price is right: No additional charge.

Server management is all about stability, availability, and performance. You should know the status of every subsystem and schedule software updates to minimize downtime. Sending alerts about potential and actual failures, as well as when thresholds are exceeded, is critical to effective operations management. On the client side, managing software to protect against viruses and system failures is of paramount importance. In any environment, the ability to discover, inventory, and monitor devices from a central console—and to drill down for in-depth information about and control over problem conditions—can be crucial for meeting service-level agreements and managing change.

Part of the Dell OpenManage suite, IT Assistant 8.0 marks a significant advance. Using the Intelligent Platform Management Interface (IPMI) to discover systems, and Windows Management Instrumentation and Secure Shell to monitor performance, enables IT Assistant to support heterogeneous server environments. Additional support for Dell/EMC networked storage and printers extends the coverage of IT Assistant to most major components of the IT infrastructure.

Bird's-eye view

IT Assistant 8.0 recognizes servers that integrate a baseboard management controller (BMC) or service processor supporting IPMI 1.5 or later, even if they are not Dell™ servers and even if the server agent has not been installed. Using IPMI, you can perform many operations and status checking tasks regardless of whether the server is powered up or the OS is running. In addition, IT Assistant 8.0 can deploy the Dell server agent—Dell OpenManage Server Administrator. IT Assistant can discover a Dell server using IPMI; configure the BIOS to enable remote, unattended deployment of the OS and applications; and then deploy the server agent.*

Enhanced device support in IT Assistant 8.0 lets you discover and monitor systems through an intuitive interface that displays system health and status at a glance. If more information is required, such as a configuration change, IT Assistant offers the appropriate options in a convenient right-click menu.

Performance monitoring and enhanced systems control

Perhaps the most significant IT Assistant 8.0 enhancement is standards-based performance monitoring. IT Assistant can help you monitor server performance for a period of time so you can create a baseline and set thresholds. Alerts can then be sent if a system exceeds these thresholds. Performance monitoring is supported for all major server subsystems, including CPU, memory, internal hard drives, and network interfaces. This type of granularity enhances systems control and helps identify where you have surplus and where you need to add capacity. And because IT Assistant 8.0 monitors performance using standard protocols, it can monitor performance of non-Dell equipment too.

IT Assistant 8.0 is part of the Dell unified systems management strategy, which is designed to help simplify operations, increase resource utilization, and scale out cost-effectively to meet evolving business needs. So if you are one of the many administrators charged with accomplishing more with fewer resources, you can download IT Assistant 8.0 and begin using it straightaway. And unlike many other tools, there is no additional charge.

* For more information, see "Monitoring and Managing Agentless Servers Using Dell OpenManage IT Assistant 8.0 with IPMI," by Suresh John, *Dell Power Solutions,* November 2006, www.dell.com/downloads/global/power/ps4q06-20070158-John.pdf.

INSIDE THIS ISSUE:

Bossier Parish Schools Get the Message

Dell™ PowerEdge™ 2650 servers, a Dell/EMC CX300 storage area network, and the Dell OpenManage™ suite form the foundation of a unified messaging system at Bossier Parish Schools

In addition to challenges in the classroom, teachers face a barrage of communications every day—memos from the school district, e-mails from other staff and administrators, calls from parents, and input from the community. Bossier Parish Schools, a district that serves approximately 19,000 students in Bossier Parish, Louisiana, has developed an extensive messaging system designed to help teachers handle communications efficiently. "Teachers receive hundreds of communications each day," explains Bill Allred, director of technology for Bossier Parish Schools. "We use messaging technologies to help organize those communications, which cuts down on classroom disruptions and helps teachers maximize the value of the time they spend with students."

In 1997, Bossier Parish Schools began using a messaging system to deliver e-mails and voice mail messages to users' e-mail in-boxes. However, as the volume of communications expanded, the district began experiencing slowdowns in message delivery—the system had reached its capacity. Allred's group decided that it was time for strategic changes to the district's messaging infrastructure.

The district migrated from a proprietary e-mail application to Microsoft® Outlook® software using a centralized repository on a Microsoft Exchange server. To provide redundant failover support for the Exchange database, Dell PowerEdge 2650 servers with dual Intel® Xeon® processors and Microsoft Windows Server® 2003 were deployed in a clustered configuration. In addition to the server infrastructure, a storage system was deployed using a Dell/EMC CX300 Fibre Channel storage area network (SAN) and a Dell PowerVault™ 775N network attached storage (NAS) server. This storage system helps provide fast, reliable delivery of messaging data, and uses Veritas software from Symantec to back up application and user data. The Dell infrastructure gives the district the ability to scale its Exchange database in step with the number of users on the system.

In addition to the PowerEdge 2650 servers that host the Exchange environment, Dell PowerEdge 2500, PowerEdge 2600, and PowerEdge 2800 servers support other applications—including reading enhancement, library inventory, and automation software. These application servers are also used as file and print servers to provide printing, copying, and scanning services. In addition, Allred is planning to install VMware® virtualization software on a high-end PowerEdge 6850 server to consolidate and centralize several applications on a single physical system, including the accelerated reading program and video streaming of educational content. To manage the district's server environment, Allred relies on Dell OpenManage software. "Dell OpenManage gives us up-to-the-minute information about our servers and the ability to troubleshoot remotely—meaning we can often fix problems before users even realize an issue exists."

Bossier Parish Schools relied on Dell Deployment and Exchange Migration Services to install the SAN and the Exchange server configuration. The district also receives Dell Platinum Enterprise Support for its SAN and PowerEdge systems running the Exchange server, which helps protect continuity with active monitoring.

> *"Now that we have adequate processing power and storage capacity to handle our messaging loads, users can get their messages and move on to more important things—like teaching."*
>
> *—Bill Allred*
> *Director of Technology*
> *Bossier Parish Schools*

As a result of the Exchange migration, the district no longer experiences the messaging slowdowns that previously plagued its system. Thanks to the performance and reliability of the unified messaging system built on Dell servers and storage and maintained by Dell Platinum Enterprise Support, Bossier Parish Schools teachers and staff can communicate more effectively—and are therefore better able to do their primary jobs. ■

Challenge

Slow, unreliable messaging solution inhibits teacher productivity, disrupts time spent with students, and forces IT personnel to waste time creating temporary fixes

Solution

Hardware:
- Dell PowerEdge 2650 servers
- Dell/EMC CX300 Fibre Channel SAN
- Dell PowerVault 775N NAS server

Software:
- Microsoft Windows Server 2003
- Microsoft Exchange
- Dell OpenManage
- VMware ESX Server

Services:
- Dell SAN Installation Services
- Dell Exchange Migration Services
- Dell Platinum Enterprise Support

TECH CORNER

By David Weber, Enterprise Technologist, Dell Inc.

Taking Systems Management to the Next Level

Maintaining an IT environment is a little like maintaining a car—you have multiple systems to monitor and fix, and it often takes a long-overdue tune-up before you realize that mileage is poor and not all cylinders are firing properly. Here's how you can take a proactive approach to bring systems management to the next level.

For many IT departments, optimizing performance is simply a question of getting ahead of all the specific tasks and responsibilities involved in daily systems management—including notifications, inventory, asset management, hardware and software configurations, and server updates. There are essentially four stages to implementing a proactive monitoring system that effectively handles the volume of incoming information and manages alerts. At each of these stages, Dell provides the tools that enable you to move into a proactive mode, get ahead of the curve, and reapply resources to more strategic tasks.

Stage 1: User-driven monitoring

In this common scenario, the IT department is not running an enterprise monitoring solution and instead relies on end users to call in whenever they discover a problem. This usually results in numerous calls but very little information.

How Dell moves you forward: The first step in moving away from stage 1 is to instrument your systems and develop an understanding of the underlying factors that affect uptime. While Dell designs for reliability, we also understand that, by their very nature, moving parts sometimes fail. For that reason, Dell provides you with the tools to stay ahead. Dell OpenManage™ Server Administrator captures changing conditions and alerts and sends the information to your console of choice. In addition, Dell OpenManage IT Assistant takes information from the hardware in the environment—servers, desktops, notebooks, switches, printers, and networked storage—and consolidates it into a single, integrated health and status view.

Stage 2: Service-level monitoring

Once you deploy instrumentation to monitor service levels, you can receive proactive failure alerts that tie into specific hardware such as a disk or network controller. In addition, you gain visibility into network and CPU utilization, which enables you to identify potential causes for failures, such as high utilization or a rise in temperature.

How Dell moves you forward: As you collect all this information with Dell OpenManage Server Administrator, you can determine the performance characteristics of your environment and develop a picture of the life cycle of each system. At this stage, you can secure your instrumentation and integrate it into your infrastructure. Dell OpenManage Server Administrator and Dell™ Remote Access Controllers can fully integrate into your authentication process so you no longer have to manage lots of different account connections, and only users who are already trusted and authenticated in your environment can gain access. Furthermore, IT Assistant enables you to gather performance information from your Microsoft® Windows® OS–based and Linux® OS–based servers and use that to generate a picture of what the OS is doing on those servers so you can begin to identify potential issues.

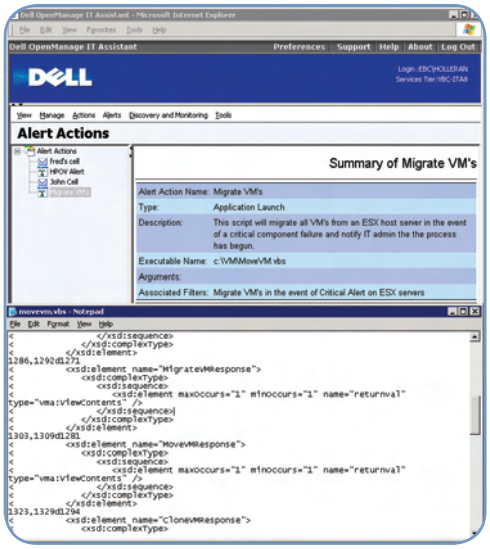

Policy-based management using Dell OpenManage IT Administrator

Stage 3: Proactive systems management

At this stage, you can track system performance and take proactive steps, such as upgrades and new server provisioning, to avoid system failures and saturation points. This puts you slightly ahead of the curve by enabling you to use the information you have collected to conduct predictive analysis of failure, overutilization, and other conditions or situations that have historically kept you in a reactive mode.

How Dell moves you forward: Dell systems management tools provide you with information about your IT environment that helps determine whether any of your systems need servicing or whether new equipment should be provisioned before a problem occurs and end users start calling.

Stage 4: Policy-based management

Moving even further ahead of the curve, now you can put policies in place that automatically take predetermined steps to respond to certain events and conditions. For instance, if a disk fails on a SAN disk array, a notification is sent to Dell and a hard drive is automatically ordered for replacement. In virtualized environments, a typical example of policy-based management might be that, when a usage threshold in an application server farm is exceeded, policies are in place to automatically provision a new virtual machine and redistribute traffic among virtual machines, transparently handling the load.

How Dell moves you forward: The policy-based automation capabilities of Dell OpenManage Server Administrator enable you to use policies from any one of your choice of enterprise tools, including Microsoft Operations Manager, Altiris® Server Management Suite™, and Novell® ZENworks® software.

The bottom line is that no matter where you are on the curve, Dell can help you take a proactive approach to systems management. Our standards-based tools are designed to help you move quickly to the next level while leveraging existing investments. By implementing Dell OpenManage Server Administrator and IT Assistant, you can understand how well your cylinders are firing when your next tune-up is due.

Benefits

- Dell infrastructure supports high-performance, reliable access to messages through multiple media
- A highly scalable SAN promotes smooth performance growth as messaging volume increases
- Convenient access to both phone and e-mail communication helps teachers focus on students

MANAGEMENT INSIGHTS

By Winston Bumpus, Director of Systems Management Architecture, Dell Inc.

Dell Unified Manageability Architecture: Blueprint for an Open Management Framework

Few climbers have made it up the side of Mount Everest without a detailed plan of action and a reliable guide. Leading the way toward comprehensive, industry-standard systems management solutions, the Dell Unified Manageability Architecture offers both the plan and the guide for an open, integrated systems management infrastructure.

The Dell Unified Manageability Architecture is based on existing and emerging standards, established by global standards bodies such as the Distributed Management Task Force (DMTF) and the Storage Networking Industry Association (SNIA). Adhering to the guiding principles of scalability, flexibility, rapid integration, open interfaces, and support for a wide range of devices, this architecture is designed to simplify operations, improve resource utilization, and scale out to thousands of nodes cost-effectively—providing a blueprint for secure, flexible growth and change management.

The Dell unified systems management framework is organized into six layers. Four layers address aspects of the managed node and device-level instrumentation for manageability: platform, logical mapping, aggregation, and access. Two layers deal with aspects of the management system such as Dell OpenManage™, Microsoft® Systems Management Server, or Microsoft Operations Manager software: resource management and optional orchestration. This layered organization is designed to help the IT industry standardize terminology, isolate functionality and interfaces, and address the various aspects of systems management in a consistent manner.

Dell is leading the industry with a standards-based blueprint that allows partners to design device instrumentation for open manageability. In this way, the Dell™ Unified Manageability Architecture enables a standard, highly scalable management framework across Dell and partner products—helping to increase choice, reduce complexity, and enhance interoperability among PDAs, notebooks, printers, servers, and storage.

Object-oriented model

Equally important to the Dell Unified Manageability Architecture's layered organization is its adherence to the Common Information Model (CIM), an object-oriented approach for describing computer and networking environments. All managed elements are positioned within CIM to clarify semantics and streamline integration. Furthermore, CIM facilitates data reuse, delivering consistent information across all systems to help increase efficiency and productivity.

In today's distributed environment, Web services are instrumental in leveraging the technology of the Web and enabling advanced interoperability. Therefore, Dell has included a large number of Web-based protocols in the Unified Manageability Architecture, including DMTF Web Services for Management (WS-Management), CIM-XML, Systems Management Architecture for Server Hardware (SMASH), and many others.

These protocols enable platform-independent information access and exchange and direct device management. Dell anticipates that over time more devices—such as printers, displays, and projectors—will include embedded management instrumentation based on Web services standards so they also can be managed natively. For enhanced flexibility and functionality, the Dell Unified Manageability Architecture specification includes in-band, out-of-band, and proxy management configurations.

Platform-native management

Already Dell engineering teams and Dell partners are mapping technologies onto this blueprint to help simplify manageability. Their goal is to make standards-based management infrastructure native to the platform, rather than something that is grafted on after the fact. Besides moving the mountain toward open systems management, the Dell Unified Manageability Architecture can help make it a lot easier to climb. ■

For more information about the Dell Unified Manageability Architecture specification and implementation examples, visit www.dell.com/standards.

November 2006

Monitoring and Managing Agentless Servers

Using Dell OpenManage IT Assistant 8.0 with IPMI

Dell OpenManage™ IT Assistant 8.0 introduces several features, including support for system discovery using the Intelligent Platform Management Interface (IPMI) protocol. This article explores how IT Assistant IPMI capabilities can be used to remotely monitor and manage Dell™ and non-Dell agentless servers.

BY SURESH JOHN

Related Categories:

Benchmarks

Dell OpenManage

Dell PowerEdge servers

Intelligent Platform Management Interface (IPMI)

Remote management

Systems management

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

Dell OpenManage IT Assistant 8.0 enables administrators to manage Dell and non-Dell agentless servers that support Intelligent Platform Management Interface (IPMI) 1.5 and later. IPMI provides autonomous monitoring and recovery features implemented directly in platform management hardware and firmware, without requiring additional instrumentation agents. Its inventory, monitoring, logging, and recovery control functions are available independent of a system's main processors, BIOS, and OS, so platform management functions can be available when the system is powered down, and administrators can obtain platform status information and initiate recovery actions when systems management software and normal in-band management mechanisms are unavailable.

IT Assistant 8.0 uses IPMI Over LAN Remote Management and Control Protocol (RMCP) messaging in conjunction with the Simple Network Management Protocol (SNMP) and Common Information Model (CIM) to connect to baseboard management controllers (BMCs). The IPMI out-of-band channel can be available even when the system is powered down or has no OS installed.

Using IT Assistant, administrators can discover systems, monitor system health, view hardware logs, and perform power control tasks. The example usages of IT Assistant described in this article do not require third-party components (for example, Dell OpenManage Server Administrator) to be installed on the managed systems.

Discovering systems

Administrators can configure IT Assistant to obtain system inventory, status, and hardware log information through the IPMI protocol based on the system and BMC IP addresses. This information adds to that provided by other in-band agents, such as the Dell OpenManage Server Administrator agent, SNMP management information base 2 (MIB2), and Windows Management Instrumentation (WMI) provider.

After configuration, IT Assistant can use the credentials provided in the discovery range to retrieve system information from the BMC using IPMI Over LAN. This process does not require a Server Administrator agent on the server, but if one has been installed, IT Assistant can collect information from the BMC and agent asynchronously,

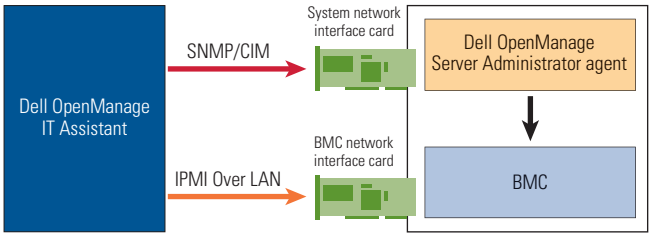

Note: The BMC network interface card could be a dedicated card or a LAN on Motherboard in shared mode.

Figure 1. Discovering a system with Dell OpenManage IT Assistant using IPMI

then use the host name or service tag to map the information from both sources to a single server. Figure 1 illustrates this discovery process. (Although Server Administrator agents provide more information than the BMC, that information is not available if the system is powered down.)

Administrators can configure IT Assistant to discover a system through IPMI using the New Discovery Wizard, which they can access by selecting Discovery and Monitoring > Ranges, right-clicking "Include Ranges" in the left pane, and selecting "New Include Range." They can then add the BMC IP address to the discovery ranges, enable IPMI discovery, and provide the BMC username and password (see Figure 2). For systems supporting IPMI 2.0 (such as ninth-generation Dell PowerEdge™ servers), they can also provide a key generator (KG) key—a security mechanism that enables encrypted traffic between the BMC and management station.

Note: IPMI Over LAN uses the UDP-based RMCP, which communicates over port 623. For firewall-protected systems, administrators should ensure that this port is available for IPMI Over LAN traffic.

Monitoring system health

In addition to basic inventory information such as service tag, host name, OS, and firmware version, IT Assistant can monitor system health over IPMI—for example, if the server temperature is above normal, the health status displays a warning. In addition to obtaining regular system health updates using scheduled status polls, IT Assistant can obtain health updates dynamically whenever the health status changes. To use this feature, administrators must set the BMC Platform Event Trap (PET) destination to the management station IP address. After receiving a PET, IT Assistant carries out a status polling of the system, connects to the BMC to retrieve the chassis status, and displays a warning with the device health status.

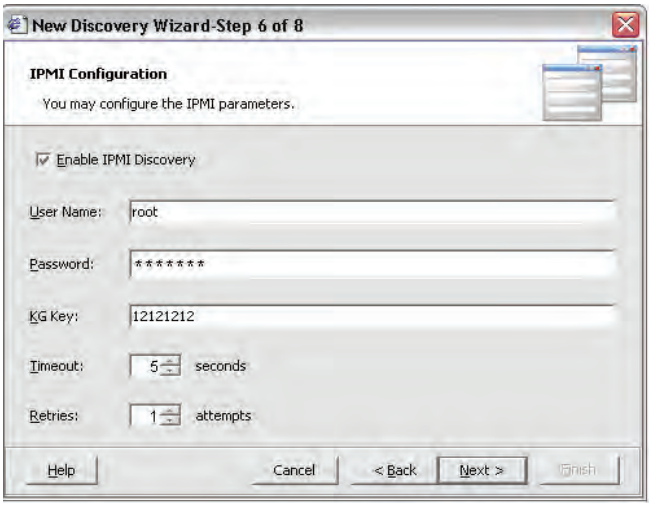

Figure 2. Configuring IPMI discovery with the Dell OpenManage IT Assistant New Discovery Wizard

Note: Systems using IPMI 1.5 have limited status capabilities, including no critical state. IPMI 1.5 also does not take redundancy into account, so it might show a system as healthy even when one of the redundant power supplies has failed. Systems using IPMI 2.0 can obtain comprehensive device status information.

Viewing hardware logs

IT Assistant allows administrators to use IPMI to view BMC system event logs (SELs) and sensor data records (SDRs), which contain information about the type and number of sensors in the platform, sensor threshold support, event-generation capabilities, and the types of readings the sensor provides. Every managed node event creates a SEL entry. Managing these logging functions through IPMI can help ensure that this information is available following a failure that disables the system processors.

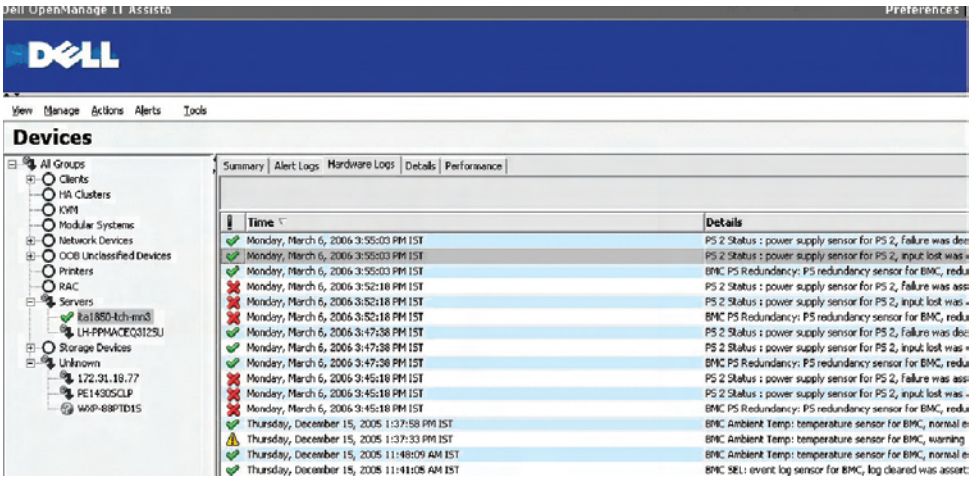

Figure 3. Viewing hardware logs in Dell OpenManage IT Assistant using IPMI

IT Assistant correlates SELs and SDRs in the hardware logs. Administrators can easily view an event by clicking the Hardware Logs tab for the device (see Figure 3).

Performing power control tasks

Administrators can use IT Assistant with IPMI to perform tasks such as power cycling, power resetting, and powering up or down. IT Assistant uses the ipmish utility to carry out scheduled power control tasks; for this utility to be available, the BMC management utility must be installed on the management station.

Administrators can create a power operations task using the New Task Wizard by selecting Manage > Tasks, right-clicking "Command Line Task" in the left pane, and selecting "New Task." They can then select "IPMI Command Line" from the Type menu (see Figure 4) and complete the wizard, choosing whether to schedule the task or run it immediately, and providing the BMC username and password along with the KG key for systems supporting IPMI 2.0.

Troubleshooting IPMI connectivity

Administrators may occasionally encounter IPMI connectivity problems such as the following, which can occur when IT Assistant cannot contact the device using RMCP or when it is using invalid credentials:

- IT Assistant discovers a device, but displays it with no category or inventory information except IP address and name.
- IT Assistant discovers a device and categorizes it as a server, but clicking the Hardware Logs tab for the device returns an error.
- IT Assistant discovers and categorizes a device, but displays its status as unknown.
- IT Assistant fails to discover a device.

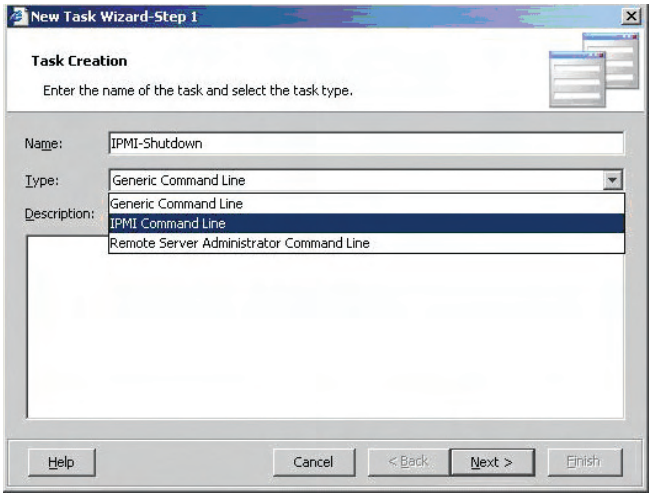

Figure 4. Creating an IPMI task with the Dell OpenManage IT Assistant New Task Wizard

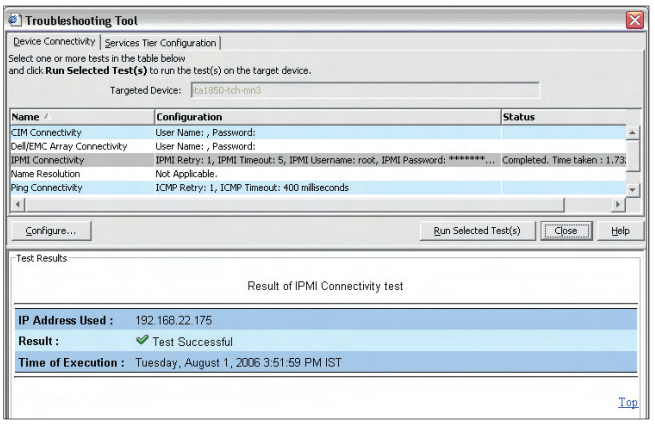

Figure 5. Using the IPMI Connectivity test in the Dell OpenManage IT Assistant Troubleshooting Tool

Administrators can use the IPMI Connectivity test, part of the IT Assistant Troubleshooting Tool, to help diagnose and resolve these problems (see Figure 5). This test sends an RMCP ping packet to the BMC to determine whether it can successfully reply. If it can, IT Assistant attempts to connect to the BMC using the credentials provided in the discovery range. If IT Assistant has discovered the device, administrators can launch the Troubleshooting Tool by right-clicking on the device and selecting "Troubleshooting Tool"; otherwise, they can do so by selecting Tools > Troubleshooting Tool.

Monitoring and managing servers flexibly

Dell OpenManage IT Assistant 8.0 enables administrators to carry out operations such as discovering systems, monitoring system health, viewing hardware logs, and performing power control tasks on both Dell and non-Dell agentless servers through IPMI. Using these features can provide administrators with additional flexibility in carrying out basic system monitoring and management tasks.

Suresh John is a senior engineering analyst with the Dell OpenManage Group at the Dell Bangalore Development Center, and has nine years of experience designing and developing systems management applications. Suresh has a bachelor's degree in Electronics and Communication from the University of Kerala in India.

FOR MORE INFORMATION

Dell OpenManage IT Assistant User's Guide:
support.dell.com/support/edocs/software/smitasst

IPMI specification:
www.intel.com/design/servers/ipmi

Optimizing Microsoft SQL Server 2005 Environments

with EMC Assessments and Quest Software

Migrating to or upgrading a Microsoft® SQL Server™ environment can be a challenging process. EMC provides assessments for SQL Server by using Spotlight® on SQL Server Enterprise from Quest Software to evaluate database performance. These assessments can help organizations understand their existing environment and plan their new deployment to meet their performance and availability goals.

BY CHAD SAKAC AND KEVIN KLINE

Related Categories:

EMC

Microsoft SQL Server 2005

Quest Software

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

Many organizations today struggle with complex administration, widely distributed data, system availability problems, and high overhead costs. The Microsoft SQL Server 2005 database platform can help mitigate or eliminate these problems, providing organizations of all sizes with an enterprise-grade database platform for critical online transaction processing and business applications that can help enhance availability, performance, scalability, and security.

Planning a migration or upgrade to SQL Server 2005, however, can be a challenging process. As organizations plan a SQL Server 2005 implementation, their first step should be to assess their current environment. EMC Corporation, a leader in information management and storage, offers assessments for SQL Server to help organizations understand their existing environment and plan for their migration or upgrade.

Carrying out EMC assessments for Microsoft SQL Server

An EMC assessment for Microsoft SQL Server evaluates an organization's existing SQL Server environment and provides recommendations for a migration or upgrade. EMC gathers information about the environment nonintrusively and then uses this information to provide guidance based on actual environmental metrics, including architectural insights and advice as well as evaluations of performance, capacity needs, forecasting, and design.

The first stage of an EMC assessment involves data collection and analysis. The process begins with the

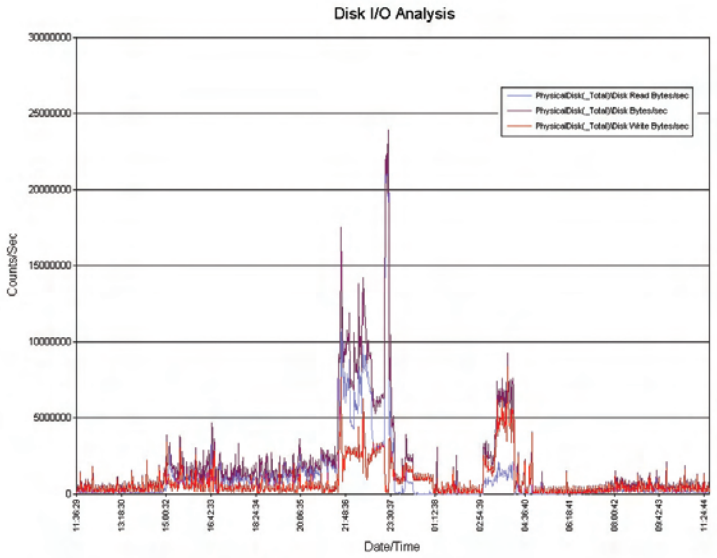

Figure 1. EMC analysis of data collected with Microsoft Windows Performance Monitor

organization completing an in-depth questionnaire about its SQL Server environment. EMC then uses a variety of data collection and analysis tools to provide comprehensive information about the environment. Data is collected for up to 72 hours, during periods of both high and low activity, to help provide an accurate analysis.

Within days following an assessment, the organization receives a final report that provides detailed information about the existing SQL Server environment, recommendations, and an action list and implementation plan for executing the recommendations. The recommendations cover system and storage infrastructure, including particular areas where migrations to SQL Server 2005 and a Microsoft Windows Server® 2003 x64 OS as well as SQL Server optimization can provide benefits. The summary report also includes recommendations to help maximize local backup and recovery capabilities, create test and development replicas, and protect the SQL Server environment from a site disaster.

Enhancing the assessment process with data collection and analysis tools

EMC uses multiple tools to help provide accurate assessment information. The most basic of these tools is Microsoft Windows® Performance Monitor (perfmon). Perfmon collects and records performance counters from the Windows OS and SQL Server database, measuring the amount of activity within a discrete area of the OS or database. However, perfmon does not correlate these discrete areas of performance with each other, so after this tool has collected up to 72 hours of activity on the database server, an EMC expert analyzes this data using internally developed tools and processes. Figure 1 shows a screenshot of such an analysis.

The perfmon part of the analysis can provide excellent insight into elements such as utilization, storage I/O profiling, and SQL Server OS (SQLOS) memory management—all critical from a server, OS, and storage design standpoint. However, EMC cannot rely on perfmon alone to provide a thorough assessment, particularly in the areas of system impact correlation to specific transactions, SQL Server query analysis, and optimization. Fortunately, powerful tools from Quest Software have allowed EMC to enhance the speed, accuracy, and comprehensiveness of its assessment process.

Using Spotlight on SQL Server Enterprise from Quest Software

When searching for tools to enhance its assessment process, EMC wanted to provide its field experts with immediate and simple analysis of OS and database performance; real-time, granular, root-cause analysis; correlation of performance metrics with SQL Server user activity; and a detailed knowledge base for understanding and diagnosing performance bottlenecks.

Spotlight on SQL Server Enterprise from Quest Software met these requirements. By providing a visual representation of the OS and database components through an intuitive graphical user interface, Spotlight enables quick and easy identification of performance bottlenecks and problem areas. The reports generated by Spotlight further augment the ability of EMC field experts to carry out a comprehensive assessment.

Figure 2. Top-level dashboard window in Spotlight on SQL Server Enterprise from Quest Software

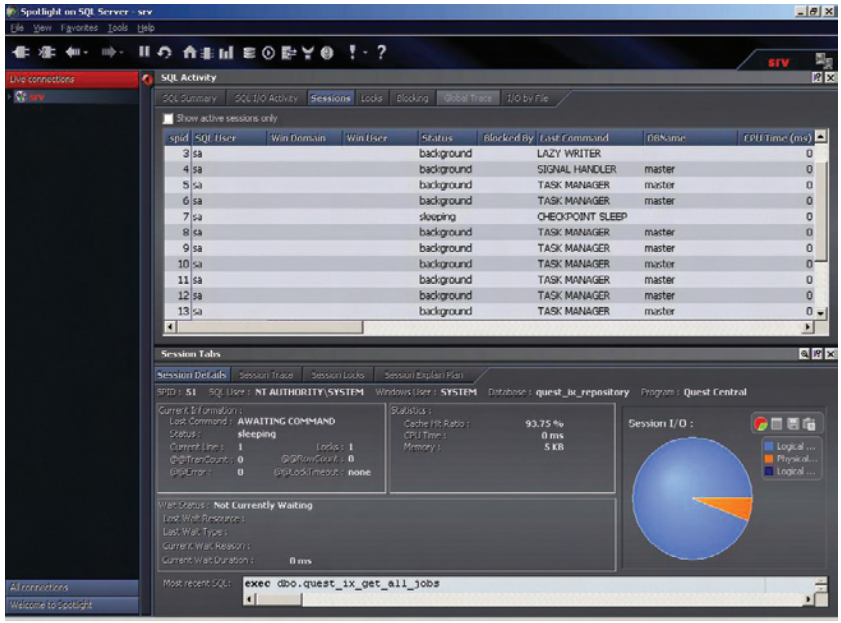

Figure 3. Details window in Spotlight on SQL Server Enterprise from Quest Software showing SQL Server sessions

Spotlight uses a dashboard approach to depict OS and database performance. Figure 2 shows the top-level window in Spotlight. A green gauge indicates that a given component of the database is functioning properly, a yellow gauge indicates a potential problem, and a red gauge indicates an outright problem, an error, or a serious bottleneck. When a dashboard gauge turns yellow or red, an EMC expert can immediately view the low-level details to perform a root-cause analysis.

The thresholds for each level (green, yellow, and red) are dynamically determined by Spotlight using an internal calibration process. Spotlight first assesses the normal activity for a particular server, then configures the thresholds that escalate a gauge from green to yellow and from yellow to red. (Users can also configure these settings manually.) For example, Figure 2 shows the CPU Usage gauge at 100 percent, so Spotlight has turned this gauge red and raised an alert. The EMC expert can then view the details of the CPU Usage gauge to see CPU activity, the processes using the CPUs, the processes waiting for the CPUs, and so forth.

Figure 3 shows an example of a typical Spotlight details window, in this case for SQL Server sessions. From the major panels on this screen, users can also view additional session details and other related information.

In addition to performing real-time monitoring and root-cause analysis, EMC experts can also research issues within the extensive Spotlight knowledge base. This knowledge base can be helpful when an expert does not have personal experience with a given problem.

Optimizing SQL Server with expert analysis and tools

Organizations can face many challenges when migrating to or upgrading Microsoft SQL Server environments. EMC, which has more than a decade of experience with SQL Server environments, offers organizations its expertise and knowledge of database and storage requirements through EMC assessments for Microsoft SQL Server. In carrying out these assessments, EMC employs industry standards, best practices, and useful tools such as Spotlight on SQL Server Enterprise from Quest Software, which provides real-time database performance information and helps EMC discover the root causes of potential and existing database problems and bottlenecks.

EMC assessments for Microsoft SQL Server can help organizations understand their existing environment before moving forward with a migration or upgrade, by providing real performance data, an analysis of the current environment, and recommendations for the future—with the ultimate goal of helping organizations enhance SQL Server performance and maintain their database investment.

Chad Sakac is a director for midsize enterprise solutions at EMC, where he is responsible for solution validation—creating, testing, benchmarking, and refining best practices for the primary applications deployed on EMC platforms. He also manages a global field team of application experts who work with customers to implement best practices for their critical application environments.

Kevin Kline is the director of technology for SQL Server at Quest Software, designing products for database administrators and developers. Kevin is the author or coauthor of four books, including *SQL in a Nutshell* and *Transact-SQL Programming,* and numerous magazine and online articles. Kevin also serves as president of the Professional Association for SQL Server and has been a Microsoft MVP for SQL Server since 2004.

FOR MORE INFORMATION

EMC services for midsize enterprises:
www.emc.com/stdforms/urlTrackServlet.jsp?evntCode=AMA00002357

Spotlight on SQL Server Enterprise from Quest Software:
www.quest.com/spotlight_on_sql_server_enterprise

Quest "Worth Upgrading to SQL Server 2005?" Webcast:
www.quest.com/upgrade

Scaling Business Process Platforms:

Identifying and Meeting the Challenges

Business process platforms can help simplify the process of developing new software by enabling the creation of services, processes, and applications from existing reusable components. This article discusses the key challenges of scaling these platforms and outlines how elements such as metadata and software product lines can help meet these challenges.

BY DAVID S. FRANKEL

Related Categories:

Application development

Business applications

Enterprise resource planning (ERP)

SAP

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

Business process platforms can help provide increased flexibility in business software systems by enabling developers to build new services, processes, and applications from existing reusable components. These platforms consist of two basic elements: a *technical software platform* and an *application platform*.

The technical software platform is a set of packaged technical capabilities designed to simplify the software development process; it includes database and network management systems as well as middleware that manages transactions, componentization, security, persistence, data transformation, Web services, and so on. The steady rise in the abstraction level of technical software platforms has enabled a similar rise in the abstraction level of programming languages and model-driven development tools.

The second part, the application platform, sits on top of the technical software platform and consists of frameworks of reusable, executable business-oriented services (such as a post-to-ledger service) and executable business processes composed of multiple services (such as an order-entry process). Application platforms also enable a rise in the abstraction level of model-driven tools that combine simple services into complex services, processes, and applications.

Identifying the principal challenges of scaling business process platforms

Several challenges can arise when scaling business process platforms, including three specific challenges related to the application platform: finding the appropriate components, building the appropriate components, and managing configuration.

Finding the appropriate components

Over time, business process platforms can offer increasingly large portfolios of reusable services and large numbers of processes composed from those services. Such an array of components can present several basic questions when building new services, processes, and applications from these components—for example, how can those who assemble solutions from components identify suitable components to reuse? And once they have identified

candidate components, how can they evaluate these candidates and determine whether they are compatible with one another?

Building the appropriate components

Building software from reusable components has often not worked as well in practice as envisioned. Anticipating the needs of those using these components can be difficult, particularly for large platforms intended to support multiple business needs, such as enterprise resource planning, customer relationship management, and supplier relationship management.

Managing configuration

Creating applications from reusable components does not solve all the traditional challenges of a software life cycle—in fact, some of these challenges can be exacerbated if not handled properly. Configuration management—including design-time configuration, deployment-time configuration, and version management—can be particularly complicated.

Design-time and deployment-time configuration. Each component may have a set of properties that the application modeler or developer sets during application design, such as a property of an accounts payable component that indicates whether to use cash or accrual accounting. A component might also have properties set during application deployment, such as a property of an accounts payable service that chooses a relational database source that the component uses for persistence, or a technically oriented property that determines the service communication protocol. Because each component can have its own set of configuration parameters, it can be daunting to debug errors that arise from subtle incompatibilities among multiple components' configuration properties.

Version management. Each application component has its own life cycle and version history, which can complicate software management. Installation procedures must deploy not only the correct application version at particular sites, but also the correct versions of all application components.

Creating a metadata-rich environment for business process platforms

To help meet the challenges of scaling business platforms and to simplify the process of finding and managing large numbers of components, business process platforms must provide a metadata-rich environment. Metadata provides machine-readable information about the platform components. This section outlines the foundational types of metadata required by business process platforms—semantically rich service contracts, quality-of-service constraints, configuration invariants, and version information—and briefly discusses the goal of integrated metadata management.

Semantically rich service contracts

A service contract consists of four fundamental elements:

- **Signature:** This defines the types of information that the service requires as input and produces as output.
- **Preconditions:** These are constraints that must be satisfied for the service to execute in a well-defined manner. A precondition of a money transfer service, for example, might be that the source and destination accounts must be owned by the same customer. Some money transfer services may not have this restriction, but for those that do, it is an important part of the service contract.
- **Postconditions:** These are constraints that must be satisfied when the service has completed execution; in essence, they describe the effects that result from that execution. A postcondition of a money transfer service is typically that the balance of the source account is decremented by the amount of the transfer, and the balance of the destination account is incremented by the same amount.
- **Information invariants:** These are constraints that apply to the information structures the service uses as its input and output parameters. An invariant for an ordinary checking account might require that the account balance never fall below zero (invalidating attempts to withdraw an amount greater than the balance), while an invariant for a checking account with overdraft protection might allow a balance of $1,000 below zero.

Signatures have typically been the only machine-readable part of service contracts. For example, the Web Services Description Language (WSDL) allows developers to define an operation's signature. Supplementing the contract with preconditions, postconditions, and information invariants is central to the Design by Contract approach, which uses such constraints to define a contract's semantics (meaning).[1] The semantics of a service contract can often be obscured in code, without even being captured informally in English or some other human language. Expressing these semantics using formal constraint languages makes them machine-readable. In a metadata-rich environment, a full machine-readable contract—including both the signature and the semantics—is an integral part of a component's manifest.

The Unified Modeling Language (UML) includes the Object Constraint Language (OCL), which supports writing machine-readable preconditions, postconditions, and information invariants. The United Nations Centre for Trade Facilitation and Electronic Business (UN/CEFACT) Modeling Methodology is an example of a methodology that uses UML in this way.[2]

[1] For more information on Design by Contract, see *Object-Oriented Software Construction,* by Bertrand Meyer, 2nd ed. (Prentice Hall, 1997).

[2] For more information about the UN/CEFACT Modeling Methodology, visit www.unece.org/cefact/umm/umm_index.htm.

There are other avenues besides UML, including various initiatives grouped under the name *Semantic Web Services,* which aim to use constraint languages in conjunction with WSDL and the Semantic Web languages RDF (Resource Description Framework) and OWL (Web Ontology Language). Emerging technologies can also represent machine-readable constraints in a human language, such as carefully structured English, so that nontechnical business analysts can read and write them.[3]

In addition, certain approaches to structuring business information can help reduce (but not eliminate) the need to express information invariants in constraint languages. The UN/CEFACT Core Components Technical Specification has an approach to organizing the definitions of *global data types* such that the very structure of a data element definition provides a map of the element's semantics.[4] This specification uses the joint International Organization for Standardization and International Electrotechnical Commission (ISO/IEC) 11179-5 approach to classifying data elements based on semantic structure—an approach that is also widely used in the Semantic Web community.[5] These semantic maps provide powerful metadata that can help enrich the service contract's semantic content when the data elements appear as service parameters. The industry is just beginning to learn how to exploit this type of metadata.

Including service contracts in the business process platform's metadata can provide several key benefits that can help address scalability problems:

- **Precise contract:** A precise contract specification with unambiguous, mathematically specified constraints can help simplify efforts to understand whether a component can satisfy specific requirements. Precision is particularly important when components must interact across organizational boundaries and human language barriers.
- **Constraint enforcement:** With some limitations, developers can use tools to enforce constraints. For example, they could generate code that checks whether a precondition is satisfied before invoking a service.
- **Collision detection:** Knowledge representation tools can detect some types of contract conflicts, such as mutually contradictory constraints—for example, a constraint on one component could require a customer's age to be 18 or under, while another requires it to be 21 or over. Such tools require that the languages used to define information structures be grounded in mathematical formalisms, as is the case with the Semantic Web's RDF and OWL languages. Although a knowledge representation tool cannot detect certain types of collisions with certainty,[6] in such cases it might be able to raise a warning flag, allow the human to decide how to proceed, and store that decision. The next time the same warning occurs, the tool can display previous decisions and eventually could suggest a default decision.
- **Semi-automated service composition:** Emerging tools can partially automate the identification of candidate services that may satisfy a requirement. The searcher expresses the requirement in a machine-readable fashion, and the tool scans the available services to match the requirement against the service contracts. These tools also require formally grounded languages.

Quality-of-service constraints

The service contracts discussed in the preceding section specify the functional behavior of a service; they do not specify nonfunctional, quality-of-service constraints. Yet quality-of-service constraints are important metadata too. A suitable service not only implements the necessary functional behavior, but also satisfies quality-of-service requirements.

Separating functional contract aspects from nonfunctional quality-of-service aspects can be useful. For example, standards bodies might want to codify the functional behavior contract for a common service, such as a post-to-ledger service, while leaving quality-of-service constraint specifications to negotiations among the service providers and clients. Implementers who provide the service can advertise the quality of service that the implementation offers.

Configuration invariants

Configuration invariants are a special type of information invariant that specify constraints on the values of a component's design-time or deployment-time configuration parameters. The value of one configuration parameter may constrain the values of others. Tools can enforce these kinds of constraints, with some limitations, and detect collisions resulting from specific component combinations; again, these collisions cannot always be detected with certainty.

Version information

A substantial amount of metadata can help developers effectively manage application and component versions, including helping anticipate the potential ramifications of a component change by tracking which applications use which components, down to the

[3] For example, see "Semantics of Business Vocabulary and Business Rules," by the Object Management Group, August 22, 2005, www.omg.org/docs/bei/05-08-01.pdf.

[4] For more information about this specification, see "Core Components Technical Specification – Part 8 of the ebXML Framework," by the United Nations Centre for Trade Facilitation and Electronic Business, November 15, 2003, www.unece.org/cefact/ebxml/CCTS_V2-01_Final.pdf.

[5] The ISO/IEC 11179-5 standard is available from the International Organization for Standardization Web site at www.iso.org/iso/en/CatalogueDetailPage.CatalogueDetail?CSNUMBER=35347.

[6] Although knowledge representation tools can typically detect with certainty collisions among the relatively simple constraints that make up description logics, certainty is generally unattainable for collisions among constraints using first-order logic.

version level. Such procedures can also help improve quality assurance processes.

Integrated metadata management

Using a wide variety of metadata can introduce its own challenges. Managing each type of metadata with different tools using different mechanisms can isolate metadata within separate silos. Integrated metadata management technologies such as the MetaObject Facility (MOF) specification, which is one of the core Model Driven Architecture standards, and the Eclipse Modeling Framework (EMF), essentially a type of MOF that underpins the Eclipse metadata management facilities, can help enterprises streamline metadata management.

Building components with software product lines

The Carnegie Mellon Software Engineering Institute (SEI), which produced the Capability Maturity Model (CMM) for Software, has defined an approach to organizing components for reuse called software product lines (SPLs). SPLs can help address one of the major problems with component-based development: scope. Constraining the scope to which a framework of components applies can help ease the task of creating truly reusable components.

The SEI defines an SPL as "a set of software-intensive systems that share a common, managed set of features satisfying the specific needs of a particular market segment or mission and that are developed from a common set of core assets in a prescribed way."[7] For example, an SPL could be a set of products that manage risk for portfolios of tradable financial derivatives or a set of products that provide role-based access security. A fairly narrow focus is typical of an SPL, because ensuring reusability is easier for a restricted, well-defined scope than for a broad, loosely defined scope.

The SPL approach divides software development into two distinct but related processes: core asset development and product development. Core asset development produces a framework of reusable assets for a product line and defines the framework architecture. Product development uses this framework to produce individual products. A production plan provides instructions for using the framework in accordance with the architecture to produce products. Figure 1 illustrates this approach, including how the two processes are divided by a *Sims water line*—an analogy developed by Oliver Sims to describe the separation, from the developer's viewpoint, between aspects of a system that appear "above the surface" and aspects that the infrastructure handles "below the surface."

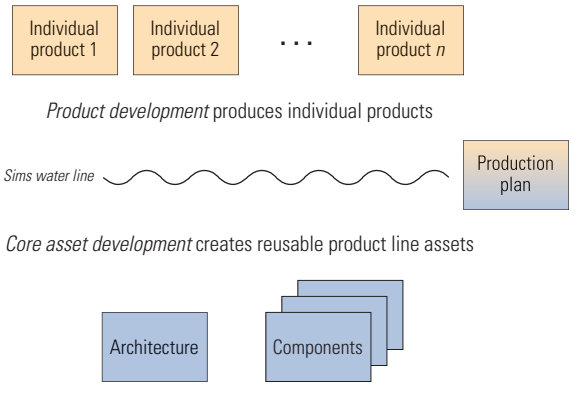

Figure 1. Software product line approach to development

Organizing a business process platform as a set of distinct but integrated product lines, each with a well-defined scope, can help simplify building reusable service and process components. Figure 2 illustrates such a platform.

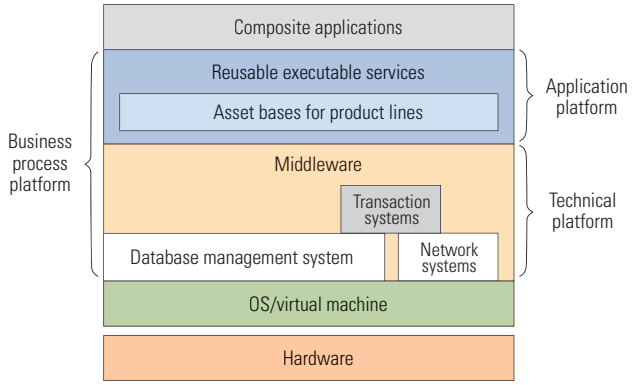

Figure 2. Business process platform with distinct but integrated product lines

Enhancing software development with scalable business process platforms

Business process platforms can help simplify software development by making it possible to build new services, processes, and applications from existing reusable components. Scaling such platforms presents numerous challenges. Addressing these challenges now can enable business process platforms to provide business value throughout their continuing evolution.

David S. Frankel has been in the software industry for more than 25 years as a software developer, architect, and technical strategist. He is the author of many published articles and the book *Model-Driven Architecture: Applying MDA to Enterprise Computing,* and was the lead editor of the book *The MDA Journal: Model Driven Architecture Straight from the Masters.* David is currently the lead standards architect for model-driven systems at SAP Labs.

[7] "Software Product Lines," by the Carnegie Mellon Software Engineering Institute, www.sei.cmu.edu/productlines/index.html.

EXPERT ADVICE. PEER SUPPORT. LAME JOKES.

How many experts does it take to solve a custom development problem? At sdn.sap.com, you'll find 500,000 developers, system managers and other insiders to help with your toughest applications challenges and coding snafus. Not to mention free sample downloads, advice from SAP staff and maybe even a few new punch lines.

// JOIN IN AT SDN.SAP.COM

THE BEST-RUN BUSINESSES RUN SAP™

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™ Expandable RAID Controllers (PERCs) and Dell PowerVault™ storage. To sign up for Dell Technical Updates, visit support.dell.com/support/notifications/technicalupdates.aspx.

Firmware updates for Dell PERCs and PowerVault direct attach storage

Product		PERC 3	PERC 4/DC, PERC 4/SC	PERC 4e/DC, PERC 4e/Di, PERC 4e/Si	PERC 4/Di	PowerVault 220S	
Firmware version		199A	352B	522A	422A, 252A*	E.18	E.19
Firmware release date		September 11, 2006	September 11, 2006	September 11, 2006	September 11, 2006	October 12, 2005	April 18, 2006
Criticality		Urgent	Urgent	Urgent	Urgent	Urgent	Recommended
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						■

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		1826, A12	53Y3, A04	5BG2, A01	2C00, A11	A16E, A09	5D5D, A14
Firmware release date		June 19, 2006	April 17, 2006	April 18, 2006	June 7, 2006	November 16, 2005	April 10, 2006
Criticality		Urgent	Recommended	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 picker issues			■			

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.

Reshaping Data Protection with

Recovery Management

Recovery management goes beyond the backup-and-restore paradigm to offer an efficient way to protect data and help ensure its continual availability. Using replication and snapshot technology to create a recovery tier within the storage environment, a recovery management implementation can provide enterprise IT organizations with uninterrupted access to data.

BY KELLY HARRIMAN-POLANSKI

Related Categories:

Business continuity

CommVault

Dell PowerVault storage

Dell/EMC storage

Disaster recovery

Storage

Storage software

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

When implementing data protection strategies, many IT organizations face a similar problem: too much data, too many applications, and not enough time to back everything up and restore it all. Recovery management offers a different approach to solving this problem—an approach that has the potential to protect any amount of data as often as necessary and to recover that data virtually instantaneously when needed. For e-mail applications and databases, such an approach could prevent data corruption and virus attacks from causing vital operations to go down for hours or even days. For a business, this approach could mean having continuous access to the data and information needed for analysis, decision making, and actions, leading to enhanced competitiveness.

An alternative to traditional backup and restore processes, recovery management involves creating and managing online replicas of production data. When a replica is online, it is immediately available and does not have to undergo a lengthy restore process before it can be used. This is a dramatic change from data backup—even backup-to-disk—in which the backup copy is not immediately usable but must first go through a restore process. Depending on how much data is involved, a typical restore process for a Microsoft® Exchange or Oracle® database can take several days—during which time the end users cannot use e-mail or process orders, grinding business operations to a halt.

Creating online replicas using snapshot and replication technologies has been possible for some time. But whereas in the past, these technologies were restricted to large, expensive storage devices, they are now becoming widely available. Dell currently offers entry-level and midtier storage devices that incorporate snapshot and replication technologies.

Adding a recovery tier to the storage environment

Creation of online replicas is forming a *recovery tier* of near-line storage in data center environments. Whether this tier

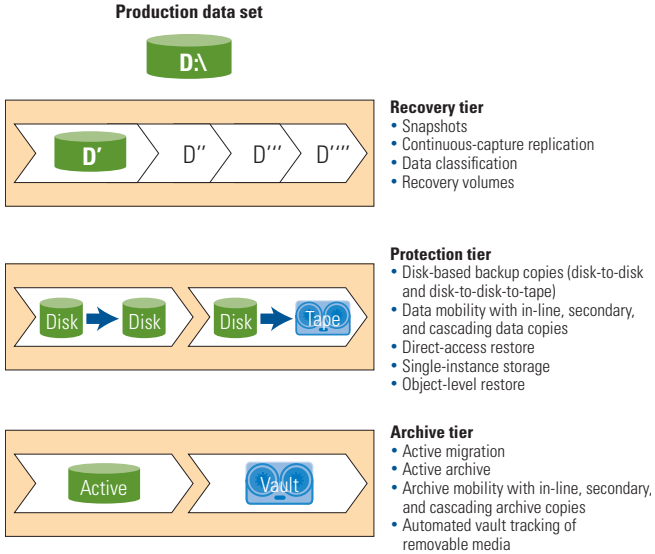

Figure 1. Storage tiers within the data center

is physical (residing on a separate storage device) or logical (residing in a device alongside production data, for example), it performs the same function: creating readily available online replicas of production data. These replicas can be created more rapidly than traditional backup copies because they maintain data in its native format, using snapshot, mirroring, or replication technologies. Because they are online and do not require a restore process, these replicas are more rapidly accessible than traditional backup copies.

In addition to providing online recovery, the recovery tier complements and enhances the *protection tier,* which has been present in most data centers in some form for a long while, and the *archive tier,* which has emerged more recently because of increased corporate regulations on retention and discovery of information. Each of these tiers offers different capabilities and service levels for data management, which in combination can provide a comprehensive management system (see Figure 1).

A recovery tier can help dramatically improve recovery point objective (RPO) and recovery time objective (RTO) service levels (see Figure 2). These improvements are necessary to enable recovery from virus attacks and other data-corrupting events that affect an entire Microsoft Exchange Information Store or other large database. What traditionally could take days to restore from backup copies can now take minutes to remount from an online mirror-image volume.

In addition to offering dramatic improvement in RPO and RTO service levels, data replicas can be used to create backup copies. Therefore, instead of having to work with production data directly—which can slow or interrupt application processing—data backups can be taken from the replica. With recovery management, the concept of trying to fit all data into a backup window is suddenly obsolete.

Implementing recovery management in the enterprise

CommVault and Dell products can help enterprises implement recovery management. Consider an example scenario in which a virus attack corrupts an enterprise's Microsoft Exchange Information Store. Traditionally, the enterprise would have no other option but to restore and rebuild the entire system. Depending on the size of the Information Store, this could take hours or even days to accomplish. And if multiple Exchange servers are involved, the recovery problem is compounded.

In such an environment, recovery management can offer quick recovery by enabling Exchange systems to be remounted to an online copy of the Information Store and then restarted immediately. Figure 3 shows how the components for recovery management can be deployed together in a storage area network (SAN) to provide comprehensive data management for Exchange and file server systems.

In Figure 3, CommVault Galaxy® backup software including the MediaAgent™ software runs on the ProxyHost, which also runs the Galaxy agents for the Exchange server. CommVault®

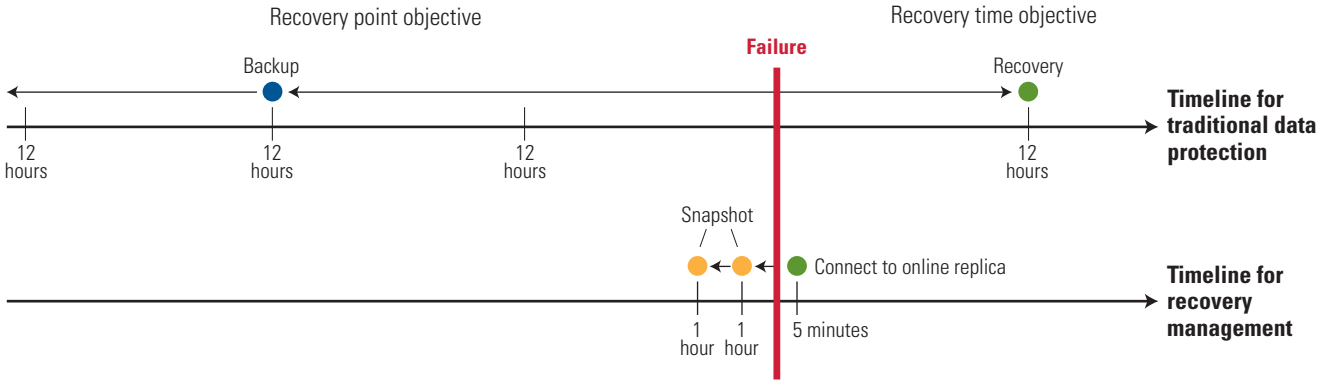

Figure 2. RPO and RTO service levels possible with recovery management as compared with traditional data protection

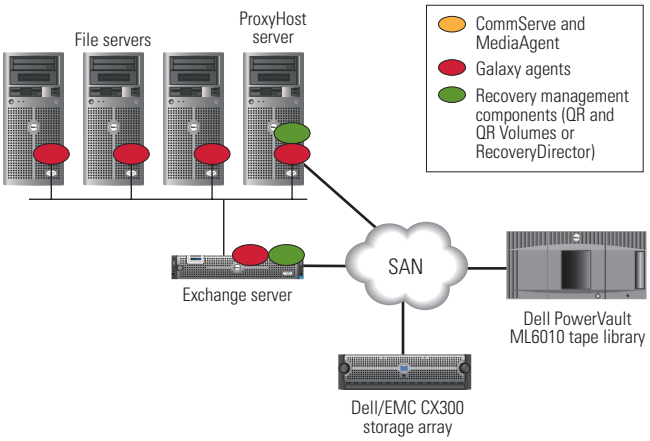

Figure 3. Example recovery management configuration for Microsoft Exchange with data protection and archive management capabilities

QuickRecovery™ (QR) software, part of the recovery management implementation, runs on the Exchange server to manage EMC® SnapView™ software on the Dell/EMC CX300 storage array. QR Volumes can easily be created for multiple point-in-time snapshots using SnapView software on the CX300 system. Once QR Volumes are created, they are available for remounting should a virus or other event cause the Exchange Information Store to become corrupted; remounting and continuing Exchange services is managed easily and quickly from the CommServe graphical user interface console. In this way, QR is designed to help dramatically improve RPO and RTO service levels.

To extend the benefits of creating QR Volumes from the recovery tier to backup, CommVault RecoveryDirector works with the ProxyHost, which makes creating Galaxy backup copies from the QR Volumes easy. This prevents Galaxy from having to back up the Exchange production data directly. Also, working from the ProxyHost avoids any impact on the Exchange server. In this way, the backup window is effectively eliminated; backup copies can be created at any time because there is no impact on the application server nor on the data.

Backup operations performed by the ProxyHost server use the SnapView images on the CX300 to create backup copies. The backup copies can be stored for the initial retention period on the CX300 and then moved to Dell™ PowerVault™ ML6010 tape devices.

Administrators in this example deployment scenario have several options for recovering the Exchange Information Store:

- Restart Exchange services with a QR Volume stored on the CX300; this method uses one of several point-in-time snapshots made every few hours
- Select the Exchange server and individual e-mail messages, e-mail boxes, or folders to restore from backup; Galaxy finds and restores the selected data in the production location on the CX300 or the recovery volume location, depending on which is in use
- Select and restore the entire Information Store from a backup copy

Although Figure 3 shows only a single Exchange server, the ProxyHost server and CX300 deployment scenario also can benefit environments with multiple Exchange servers. Each Exchange server would require QR and QR Volumes. The ProxyHost server would require only a single copy of RecoveryDirector, regardless of the number of Exchange servers in the deployment.

The recovery management capabilities described in this section for Exchange can also apply to Microsoft SQL Server™ and Oracle databases, and are particularly helpful for recovering very large databases when they become corrupted. Again, only a single ProxyHost server is required to provide recovery management benefits to any number of combined Exchange, SQL Server, and Oracle systems.

Providing a unified, efficient approach to data management

Managing recoveries in a unified approach is a highly efficient approach to data management. Because snapshot management is integrated with backups, the data is always coherent—and scheduling the backup and the snapshot can be accomplished in a single policy. The production data is affected only once—to produce the snapshot—adding to the efficiency of the process. Restoring is also efficient, because the CommVault ProxyHost makes the recovery process the same as if the backup were created from the production data; it involves the same steps of selecting the client, selecting the data to recover, and then executing the recovery process. Recovery management implemented with CommVault software and Dell hardware can provide an important asset for any data management deployment.

Kelly Harriman-Polanski is the director of product marketing for CommVault. Kelly recently joined CommVault and has worked in the storage software market for nearly 10 years, most recently at EMC/Legato. Her interests include data and information management; data archiving, retrieval, compliance, and classification; and integrated snapshot and replication management. Kelly graduated magna cum laude and Phi Beta Kappa from Augustana College in Illinois.

FOR MORE INFORMATION

CommVault recovery management:
www.commvault.com/recovery_management.asp

Implementing Cost-Effective Data Protection with Dell/EMC CX3 Series Storage

Implementing a business continuity solution to help protect and ensure the availability of critical data has become increasingly important for many enterprises. This article discusses how enterprises can evaluate their business continuity requirements and how Dell/EMC CX3 series storage arrays and EMC® software can help meet these requirements cost-effectively.

BY BRAD STECKLINE AND BARRY L. ADER

Related Categories:

Business continuity

Data security

EMC

Storage

Storage area network (SAN)

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

Although protecting critical IT systems and mitigating the risks associated with their disruption have become increasingly important for enterprises of all sizes, implementing a comprehensive business continuity solution throughout the organization has typically been financially impractical for all but the largest enterprises. Fortunately, as enterprises have begun to address these requirements, technology developments have made a growing number of cost-effective solutions available.

Evaluating business continuity requirements

Many factors can contribute to the importance of business continuity for a given enterprise: compressed business cycles and increased data volume brought about by online processes, demanding customer expectations for services to be available at all times, regulations such as the Sarbanes-Oxley Act that dictate data protection requirements, or a move to a centralized storage network with the goal of increasing IT efficiency. These factors can increase the chances of data loss and simultaneously reduce tolerance for it, causing enterprises to explore flexible business continuity options that can still meet constrained IT budgets.

Establishing business continuity requirements before making specific technology decisions can help enterprises implement successful configurations suited to their needs. Two questions regarding specific data and systems can help determine these requirements: First, how much data can the enterprise afford to lose if its applications fail? And second, what is the longest outage the enterprise can tolerate? The answers are, respectively, the enterprise's recovery point objective (RPO) and recovery time objective (RTO).

Enterprises should also carefully consider two factors that can help determine the criticality of specific data and applications: the monetary value of a specific application and its associated data, and the costs that could be incurred if this system becomes unavailable. Once this value is established, choosing the appropriate level of data protection can become an objective assessment of value versus costs and help enterprises meet their business continuity requirements cost-effectively.

Matching protection level to enterprise value

Once an enterprise understands its business continuity requirements, it can match the RPO and RTO to an

appropriate business continuity implementation by considering the costs for meeting each objective and making sure these costs are consistent with the value. If the IT organization has already established a catalog of data protection service levels, the enterprise can quickly determine whether it is willing to pay to meet the RPO and RTO.

Fundamentally, business continuity solutions protect production data by creating copies. How quickly a copy can be made dictates how often the copies can occur, which in turn dictates the maximum potential data loss; all transaction records that are captured after the last copy was made are potentially lost if a disruption occurs. The copy method determines the level at which the data is protected and the costs required to reach that level. Figure 1 illustrates the range of available data protection methods—tape backups, disk backups, snapshots, clones, asynchronous mirrors, and synchronous mirrors—and the hierarchy of their typical relative implementation costs.

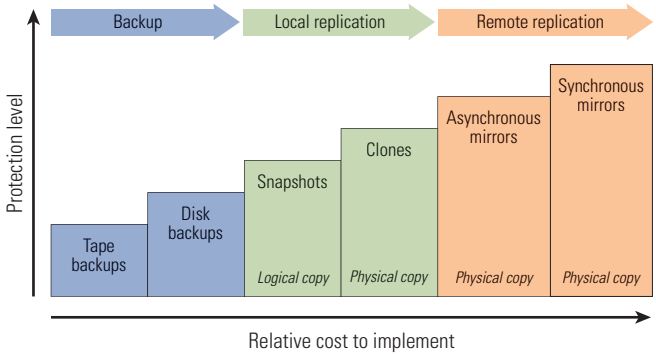

Figure 1. Data protection methods and the hierarchy of typical relative implementation costs

Tape backups. Of these data protection methods, tape backups provide the lowest level of protection and can result in the greatest potential data loss and longest time to restore following a disruption. A medium-size enterprise might use tape backups as its entire data protection strategy, while a large one might use them for some portion of its applications. Regardless of the storage environment size, enterprises can still face significant challenges in executing backup procedures. For example, as data volumes grow, these procedures can take increasing amounts of time to execute. Traditionally, performing backups meant taking production systems offline sometime in the middle of the night, but this procedure may not be acceptable for enterprises with online processes and 24/7/365 operations.

Disk backups. Disk backups, which can be implemented on cost-effective ATA-based storage arrays, can be a faster and more reliable data protection method than tape backups. They can therefore alleviate some of the challenges of tape backups while providing higher protection levels, although at a higher price.

Snapshots. Snapshots are not actual physical copies of data, but pointers to the original data as it was at the time of the snapshot. Snapshots allow logical copies to be made with much greater frequency than tape or disk backup methods allow, because snapshots can be instantly accessible; physical copies such as clones or mirrors, in contrast, must be fully synchronized before becoming accessible. Snapshots also consume less disk space than physical copies.

The disadvantage of this method is that, unlike full physical copies, disrupted data cannot be restored from a snapshot without significant data loss. Snapshots are therefore typically implemented in combination with full physical copies or replicas. An enterprise should consider the service level required for its environment when deciding whether or how to implement snapshot technology.

Clones. Clones, also called business continuance volumes, are separate physical copies of production data at a specific point in time, generated by one of two methods: creating an initial clone and then updating it by applying the ongoing, incremental changes to the production data, or "splitting off" or disconnecting one of the copies from the clone process.

Asynchronous mirrors. Mirrors are separate physical copies of production data that continuously track changes made to the data. They are typically created at a remote site—to decrease the chances that an event disrupts both the production and mirrored data—and can be created using either the application server or a storage array. Server-based replication can cost less than storage-based replication, but it can also reduce the number of server cycles dedicated to production workloads. Storage-based replication can help ensure the availability of critical production applications.

Mirrors can be either asynchronous or synchronous. Asynchronous mirrors accumulate changes to production data and then apply these changes at specified intervals. This method can help avoid the performance penalties of synchronous mirrors, but any transactions made between the copy intervals may be lost if the production data is disrupted.

Synchronous mirrors. Synchronous mirrors continuously apply the changes to production data to the remote copy of the data before the changes are committed to the production server. Because the copy is synchronized with the production data, no transactions should be lost following a disruption to production data; however, because the production server waits for the mirrored data set to change before committing a transaction, performance can decrease as the physical distance separating the production and mirrored data sets increases.

Extending protection levels for complex application environments

The growth of Web-based applications and services means that enterprise environments can involve multiple databases and interrelated applications, in which sales, manufacturing, e-commerce, and customer service records may all share common databases.

Before an application portfolio that uses multiple databases can be restarted following a disruption to a primary server or storage array, copies of each database must be in place that represent the same point in time at the transaction level; restarting with inconsistent copies can cause significant data integrity problems.

Advanced consistency software can help automate the restart process and help ensure the consistency of replicas made in these complex application environments. This software couples production data and its associated mirrors with the application servers to manage the restart process. When a primary site fails, it can seamlessly transfer control from the failed production storage to the mirrored storage and restart the application on the remote secondary server.

Implementing tiered data protection using Dell/EMC CX3 series storage

Enterprises typically cannot protect all their data assets with just one data protection method in a cost-effective way. Less-critical data may only require a simple tape or disk backup, while the most critical data may require a synchronous mirror coupled with automated application restart software. Using only one method to protect both types of data can cause either unnecessary risk of data loss when critical data is insufficiently protected, or unnecessary costs when less-critical data is overprotected. But combining these technologies into a tiered protection infrastructure can help provide appropriate data protection levels based on the data's value to the organization.

Tiered protection can also enable enterprises to not only protect their critical data assets, but also use the copies of production data to support parallel processing activities—such as backup, application development and testing, and data warehouse refreshes—while helping increase production system availability. This type of infrastructure can help enhance decision support by allowing the data warehouse to be refreshed frequently with current data without stopping user query activity, and can enable frequent backups to help reduce the potential for data loss. Enterprises can thoroughly test system upgrades, maintenance fixes, and configuration changes offline using up-to-date copies of production data, which can help reduce the chances of failure or data integrity problems when implementing the changes on production systems.

Mid-range storage platforms such as Dell/EMC CX3 series storage arrays can help large and medium-size enterprises implement cost-effective tiered data protection that only large organizations could previously afford. Enterprises can use high-performance disk drives to support critical production volumes while utilizing cost-effective high-capacity drives for less-critical environments (such as those for development, testing, or reporting) and tape-emulating disk libraries for backup and recovery operations. The CX3 series offers a range of end-to-end business continuity solutions that, when coupled with EMC Consulting Services, are designed to meet various business continuity needs.

In addition to working transparently with host-based replication software such as the EMC RepliStor® application, the CX3 series works with a range of array-based replication software

SCENARIO 1: SAFEGUARDING GOVERNMENT RECORDS USING EMC SNAPVIEW AND SAN COPY ACROSS A DELL/EMC SAN

An IT department for a county government serving several cities and hundreds of thousands of residents addressed its data protection needs by deploying EMC software and Dell/EMC CX series storage arrays.

Challenge: While serving residents' online information needs, the county's IT department experienced difficulty in implementing the appropriate level of data protection and in managing the growing pool of direct attach storage for the servers scattered across its facilities.

Solution: The IT department implemented a storage area network (SAN) based on two Dell/EMC CX series arrays to consolidate critical information and address storage provisioning challenges. EMC SnapView software is used to create a local clone of the Microsoft® SQL Server™ databases residing on a CX series array at one of the county's IT facilities. These databases hold registered deeds and geographic information system (GIS) records supporting mapping applications such as Enhanced 911, dispatch, appraisals, and elections. EMC SAN Copy software is then used to push copies of the clones across the SAN to another CX series array located a few miles away at the courthouse facility.

Benefit: The CX series arrays provide a manageable storage environment with advanced data protection and business continuity capabilities. The SAN enables the IT department to recover from server failures quickly—if a server fails, they can simply boot a second server from the SAN. The remote copies of data provided by the clones allow the IT department to test upgrades, fixes, and configuration changes without disrupting the production system, helping significantly improve the reliability and accuracy of these processes. And with critical data records expected to grow significantly and possible new mapping applications easily doubling that growth rate, these data protection and business continuity capabilities can help the storage environment scale to meet those needs.

When information comes together, business just keeps getting better.

DELL | EMC²

ALL THE RIGHT CONSOLIDATION, BACKUP AND ARCHIVE SOLUTIONS

Whether you need fast backup and complete protection or scalable and easy-to-manage storage consolidation for your enterprise, Dell/EMC brings you solutions that are high on results—and simple to use. That's because it's easier than ever to put premium software, robust storage, and world-class technical support to work solving your business's critical IT challenges.

Entry SAN Solution

- Dell/EMC AX150 Storage Platform
- iSCSI or Fibre Channel Connectivity
- EMC® Navisphere® SAN Management Software

SAN Windows Backup Solution

- Dell/EMC CX3-20 Storage Platform
- EMC® Navisphere® SAN Management Software
- EMC SnapView™ and EMC Replication Manager/SE Software
- EMC SAN Copy™ Software

Data Archiving Solution

- EMC Centera™ Storage Platform
- Windows File System Archive Edition with EMC DiskXtender® Software
- Governance Edition with EMC EmailXtender® and EMC DiskXtender® Software

BUSINESS SOLUTIONS FOR MIDSIZE ENTERPRISES

CALL 800.999.3355 **www.dell.com/emc**
toll free

that can deliver snapshots, clones, or mirrors without consuming valuable server cycles or LAN bandwidth, including EMC SnapView™, SAN Copy™, and MirrorView™ software:

- **SnapView:** This software runs within CX3 series arrays and can deliver either snapshots or full-volume clones that enterprises can use as the source data for fast, frequent, and nondisruptive backups or for development, testing, and data warehouse procedures.
- **SAN Copy:** This software also runs within CX3 series arrays, and is designed to move full or incremental copies between CX3 series arrays and other storage platforms, including EMC Symmetrix®, Hitachi, HP, IBM®, and Sun storage arrays. By integrating SAN Copy with SnapView, enterprises can use snapshots and clones as the source volume from which SAN Copy pulls data for enhanced remote recovery operations.
- **MirrorView:** This advanced, array-based replication software can provide synchronous or asynchronous mirrors between two or more CX3 series arrays. By integrating MirrorView with SnapView, enterprises can use copies of production data for parallel operational procedures from a secondary location. Enterprises can also use MirrorView replicas with advanced high-availability software to help provide near-instantaneous restart of failed servers at a remote location.

Both SnapView and MirrorView include application consistency technology to help ensure that copies made in environments with multiple databases and interrelated applications represent a consistent point in time, which can help provide data integrity for processes that use those copies. The two sidebars in this article describe how two enterprises—a county government IT department and a health care organization—implemented data protection solutions using Dell/EMC storage and EMC software.

SCENARIO 2: PROTECTING CRITICAL MEDICAL RECORDS USING EMC MIRRORVIEW/SYNCHRONOUS

A fast-growing health care organization used EMC software and Dell/EMC CX series storage arrays to help implement an aggressive plan to eliminate paperwork throughout its network of hospitals.

Challenge: As paper disappears, digital records grow and must be properly protected—not only to help ensure effective patient care in the event of data loss, but also to comply with strict federal regulations that govern information security. Included in these records are Picture Archiving and Communication System (PACS) images, each of which can be several gigabytes in size. Losing the PACS images could have a severe impact on patient care because the physicians reference them frequently in the course of treating patients.

Solution: The health care organization decided to store the PACS images on a high-performance Dell/EMC CX series Fibre Channel SAN in a primary data center. To further protect these vital medical records, it uses EMC MirrorView/Synchronous (MirrorView/S) software to continuously replicate the PACS images to a second CX series SAN at a disaster recovery site a few miles away.

Benefit: The CX series arrays and MirrorView/S software provide rapid access to actively used PACS images. Using MirrorView/S also enables the organization to verify that the data is intact when it reaches the disaster recovery site. In addition, the organization can recover these images quickly in the event of a disaster. When PACS images were stored on optical disk, restoring them took several hours—and possibly even days if mechanical problems occurred. Today, the organization can fail over to its disaster recovery site in a matter of minutes, and physicians can continue to access the PACS images with virtually no disruption to patient care.

Taking advantage of cost-effective data protection

Investment in business continuity is more than an insurance policy that sits idly until disaster strikes. Technology originally developed for disaster recovery has evolved so that it can also help increase the efficiency of operational procedures such as development, testing, and data warehouse refreshes. As a result, in addition to disaster recovery, these technologies can help increase the productivity of both IT staff and end users utilizing decision support systems—benefits that can help justify the required investments for medium-size enterprises.

Mid-range storage platforms such as the Dell/EMC CX3 series coupled with EMC software can help enterprises implement cost-effective tiered data protection that only large organizations could previously afford. After assessing both the risks and consequences of data loss, enterprises can take advantage of these technologies to develop appropriate implementation plans for protecting and maintaining the availability of critical production data and applications.

Brad Steckline is the senior manager of the EMC Backup, Recovery, and Archiving Initiative.

Barry L. Ader is the senior director of marketing for EMC CLARiiON® storage.

FOR MORE INFORMATION

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

EMC business continuity:
www.emc.com/solutions/continuity

THE GREAT VIRTUALIZATION MIGRATION

The benefits of infrastructure virtualization are many and proven. Increased operating efficiencies top the list.

New server deployment time—one hour. Server hardware cost savings—$331,000. Server process utilization rate—60 percent, up from 10 percent[1]. For Concord, Mass.-based Welch Foods, producer of Welch's juices, virtualization has paid enormous bottom-line dividends. It's also allowed the company to free up space in the data center by consolidating applications onto just 10 Dell PowerEdge servers running VMware ESX Server. The alternative would have been to purchase 100 standalone servers at a cost of $720,000.[2]

Today, with robust management software and native, on-chip support at the processor level, virtualization solutions are bringing operational efficiency to distributed computing environments like never before. As a result, IT executives are taking notice—and enterprise adoption is taking off. In fact, according to a May 2006 survey conducted by The Strategy Group for Ziff Davis Media (sponsored by Dell), over 80 percent of the 205 IT decision-makers polled have either implemented or are planning to implement virtualization in their environment.[3]

The survey further revealed the reasons why companies are embracing virtualization (see survey results graph, below). The top three—lower hardware costs, lower maintenance costs and higher utilization rates—were cited by a third or more of respondents. Clearly, IT executives are searching for greater operational efficiency. And virtualization delivers.

Virtualization is a key component of the Dell Scalable Enterprise strategy. Working with industry leaders Intel, VMware, Altiris and EMC, Dell offers mature, proven, best-of-breed solutions that can deliver dramatically higher levels of operational efficiency by helping to:

- Improve server utilization rates
- Streamline development and test environments
- Support legacy applications more effectively

Let's look at each of these cases more closely.

IMPROVING SERVER UTILIZATION

Imagine if factories operated their machinery at five or ten or even twenty percent of production capacity. They wouldn't stay in business long. And yet that has been the standard operating procedure for enterprise servers.

However, with server virtualization, efficiency is the new model. Diverse operating systems and applications that previously would have required many physical servers can now be encapsulated into virtual servers and consolidated quickly and easily on a single server. Moreover, virtualization lets administrators relocate or replicate workloads quickly and easily, so maintenance can be performed without affecting service levels or uptime. The effect is better utilization of resources—hardware, software and human—and greater operational efficiency across the board.

In addition, virtualization can provide the basis for outstanding scalability. By moving away from the one-application/one-server model, enterprises don't have to buy excess server capacity that locks them into older technology. By deploying applications on virtual machines, companies can make the most of the re-

TOP FACTORS DRIVING VIRTUALIZATION ADOPTION

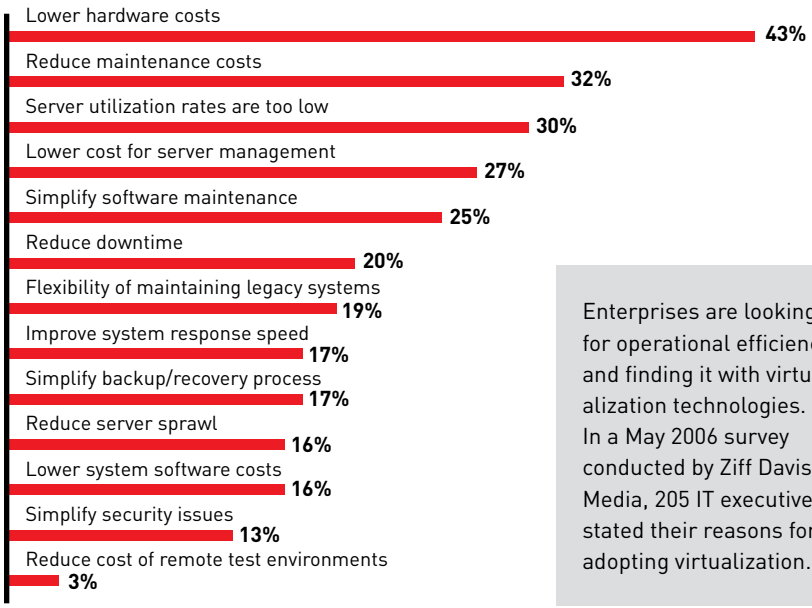

Enterprises are looking for operational efficiency and finding it with virtualization technologies. In a May 2006 survey conducted by Ziff Davis Media, 205 IT executives stated their reasons for adopting virtualization.

[1] Results not typical; [2] "Welch Foods: Out of a Jam with Virtualization," *Baseline,* April 2006; [3] Survey conducted in May 2006 by The Strategy Group for Ziff Davis Media on behalf of Dell

sources they have, then add more capacity as needs arise. Stability may also be enhanced because if one virtual machine fails while sharing physical resources, the other servers (virtual machines and physical hosts) aren't affected.

STREAMLINING SOFTWARE DEVELOPMENT AND TEST ENVIRONMENTS

Virtualization enables companies to test software in a controlled environment prior to deployment without disrupting production environments. Multiple virtual machines may be used to create a virtual multi-server environment for distributed testing on a single physical server.

This lets IT departments demonstrate new software as well as test patches and upgrades to copies of their exact environment without fear of corruption or end-user downtime. It also helps speed up platform certification because it allows for certifying to a common virtual interface rather than multiple implementations or generations of physical hardware.

Lastly, virtual machines can be set up in minutes and used multiple times. Applications and services are therefore far more likely to be developed on time and within budget—contributing to skillful, economical application lifecycle management.

SUPPORTING LEGACY APPLICATIONS

Virtualization can simplify operations by ensuring workload portability across multiple servers. This includes the ability to "re-host" software—even legacy operating systems that are no longer supported—on new servers. In essence, companies can migrate to new versions of Windows® and Linux®, while keeping old OSs and applications for as long as they serve a purpose. Platform-specific incompatibilities or "quirks" can be eliminated by encapsulation in virtual machines.

By supporting legacy OSs and applications, virtualization provides long-term stability and support in a static environment. By decoupling hardware from the host OS, each can evolve without disrupting the other environment. This type of deployment may be carried out today within environments where the guest OS is supported by the virtualization software, as is the case with the VMware virtual infrastructure environment.

CHOOSING THE PRACTICAL PATH

Every organization has its own reasons for migrating to virtualization, based on business needs, competitive challenges, existing infrastructure, and other factors. And as we've seen, virtual infrastructures that help companies run a leaner, more efficient data center can't be beaten.

Where do companies turn when seeking guidance on virtual infrastructure technologies and deployment solutions? The Strategy Group/Ziff Davis Media virtualization survey confirms that no clear leader has yet emerged. The top three server systems vendors were virtually tied for first place as hardware vendors likely to be considered for a virtualization project, and most preferred brands.

Dell and its partners—Intel, VMware, Altiris and EMC—aim to change the game with tightly integrated virtualization solutions that are proven and ready for deployment. By working together on the key components—servers, virtualization infrastructure software, management and storage—these partners deliver the best-of-breed answer to today's top IT challenges. And Dell Services for assessment, design and implementation help companies jump-start virtualization projects and keep them running successfully.

Forward-thinking companies around the globe are using these virtualization solutions to get the most out of what they already own, while positioning their infrastructure to respond to new business demands. For more information, visit www.virtualization.ziffdavis.com. ■

VIRTUALIZATION GIVES AD AGENCY COMPETITIVE EDGE

For GSD&M, a 750-person advertising agency in Austin, Texas, server virtualization was key to staying competitive. GSD&M boasts a demanding list of clients that includes Wal-Mart, Southwest Airlines, and the PGA Tour.

"We in the IT department see ourselves as enablers," explained GSD&M Chief Technology Officer Jerry Rios. "We allow our creative people to respond quicker and deliver work faster than our clients have been accustomed to in the past."

Besides enabling his creative people to work more efficiently, Rios and his staff are improving the reliability and ease of management of the IT environment through a solution that includes Dell PowerEdge servers, VMware virtual infrastructure software, Altiris management tools and a Dell/EMC storage area network. Most recently, GSD&M purchased a four-way Dell PowerEdge 6850 server—based on the 64-bit Intel® Xeon® processor—that features on-chip virtualization support at the processor level via Intel Virtualization Technology. Performance is outstanding even while handling many virtual machines containing multiple operating systems and applications.

"We're consolidating 18 virtual servers on one physical server," Rios said. "The PowerEdge 6850 is very powerful hardware that allowed us to clear out a huge area of our data center."

Advertising is a tough business: as agencies acquire new clients, they are expected to add staff virtually overnight to accommodate them. To help speed the deployment of data resources and improve the management of its computers, GSD&M uses Dell OpenManage™ along with the Altiris Management Suite. Virtualization enables GSD&M to encapsulate complex configurations onto virtual servers that can be easily replicated. And with Dell and Altiris software, new employees are provisioned automatically.

BETTER BUSINESS PROTECTION THROUGH VIRTUALIZATION

Virtualization can help make your business protection strategy comprehensive and affordable.

Downtime is not an inconvenience: It's a revenue killer.

A recent Infonetics Research study, *The Costs of Downtime: North American Medium Businesses 2006*, found that mid-market enterprises (101 to 1,000 employees) lose an average of 1 percent of their annual revenues to downtime. Even more startling, another study has found that up to 40-50 percent of businesses that have suffered major service interruptions never recover completely, and fail within two to five years.

Clearly, most enterprises can't survive without viable business continuity strategies, which must include means by which to provide both high system availability as well as disaster recovery mechanisms.

High availability keeps services running and continuously usable in the event of component or software failure. Disaster recovery helps ensure that all data and applications can be restored quickly in the event of a catastrophe—such as a virus, natural disaster, or user error—that destroys part or all of an organization's IT resources.

BUSINESS CONTINUITY CHALLENGES

Unfortunately, the cost of implementing such strategies via traditional means oftentimes was expensive and complex.

- **Expense.** Technologies such as replication and mirroring, which maintain and update copies of an organization's data and applications at a second remote site, are critical elements to comprehensive DR planning. These and other DR technologies require companies to make major investments in additional hardware and software.
- **Complexity.** These same solutions require new, specialized processes as well as considerable management and testing resources to ensure that they operate effectively. Many enterprises that implemented complex business continuity processes found that the latter ultimately were flawed, and failed when they were desperately needed.
- **Downtime.** Other organizations were simply forced to lower their expectations, accepting strategies with longer recovery time objectives even though the downtime they caused could cost the companies huge sums in lost productivity and revenue.

The good news is that, while business continuity strategies like replication, mirroring, and remote DC are still good ideas, complementary virtualization solutions can help address issues of expense, complexity, and downtime.

AFFORDABLE BUSINESS CONTINUITY

Virtualization is an enabling technology that not only may be used to help lower costs and simplify the IT infrastructure, but may also be used to extend business continuity planning to services that previously may not have justified the expense of a continuity plan.

In fact, according to a May 2006 Dell-sponsored survey conducted on behalf of Ziff Davis Media by The Strategy Group, more than 80 percent of 205 polled decision makers stated that they either have implemented or plan to implement virtualization in their IT environments. When asked to cite the top reasons for these decisions, 20 percent indicated downtime reduction, while 17 percent cited backup and recovery process simplification.

Virtualization solutions, such as those based on VMware Infrastructure 3 software, virtualize and aggregate industry-standard servers and their attached network and storage into unified resource pools. Complete environments including operating systems and applications are encapsulated in virtual machines that are independent from hardware. A set of virtualization-based distributed infrastructure services for virtual machines bring vastly improved levels of availability and capabilities for disaster recovery and system restarts. VMware solutions are fully compatible with and enhance traditional business continuity solutions like backup and server clustering, as well as redundant network and storage interface adapters.

LOWERED EXPENSES

Servers in today's enterprises are often underutilized, resulting in overinvestment in server hardware. By aggregating servers and attached network storage into unified resource pools, virtualization can help lower the cost of redundant equipment by hosting multiple back-up virtual machines on spare equipment, thereby providing business continuity for less-critical applications.

There also are other ways to improve business continuity via virtualization.

The first is to divide a few existing physical servers into multiple virtual machines. This strategy was employed by Houston-based financial services broker/dealer NEXT Financial Group, which lowered its hardware investment significantly by dividing just four physical servers at its remote site—a secondary disaster recovery center—into mul-

tiple virtual machines that mirrored the contents of all 22 Dell PowerEdge servers at its primary site. The remote site includes a Dell/EMC storage array (which mirrors the Dell/EMC storage area network at the primary data center to enable fast recovery in the event of failure) and Dell PowerEdge servers running VMware ESX Server virtualization software. This strategy saved thousands of dollars that NEXT otherwise would have had to spend on the purchase of 18 additional servers.

The second option is to operate the data center as a single pool of processing, storage, and networking power to be allocated and de-allocated on the fly to various software services. In a virtual infrastructure, users see resources as if they were dedicated to them, while the administrator manages and optimizes resources globally across the enterprise. Through a virtualization strategy, organizations have realized faster, more flexible, and more reliable disaster recovery, along with increased service levels, at lower costs, with little or no investment in additional servers.

The third option is to implement new technologies available in VMware Infrastructure 3 technology that helps ensure rapid, reliable failover, maximizes systems availability, and streamlines backup processes. Each of these components comprises a feature-rich suite that empowers data center administrators to realize a unique set of capabilities that enable a virtualized environment to be more dynamically responsive, highly available, and rapidly recoverable than traditional physical IT environments.

VIRTUALIZATION INFRASTRUCTURES AND SHARED STORAGE: PERFECT TOGETHER

Business continuity solutions rely on robust, fault-tolerant storage devices in addition to equally fault-tolerant virtualized servers. Most enterprises utilize shared storage in the form of a storage area network (SAN) or network attached storage (NAS).

Shared storage provides economies of scale for the virtual infrastructure by allowing scalable access to common storage arrays without constant hardware upgrades. Shared storage helps simplify backup preparations and expedite systems re-deployment times for disaster recovery.

Shared storage also helps make high availability functions like server clustering and workload balancing practical and less complex to implement. Virtual infrastructures and shared storage work together to deliver robust business continuity. Dell/EMC shared storage devices are tested and certified with VMware Infrastructure 3 software. Ease of creating, provisioning, and managing storage resources is a tremendous benefit to systems administrators. Additionally, shared storage can be matched to business needs via resource pools in a virtual infrastructure, thereby increasing IT resources to meet the needs of workloads for all parts of the organization.

Dell/EMC shared storage devices provide full support for features such as Distributed Resource Scheduler, High Availability, VMware VMFS, and VMware Consolidated Backup. Each of these features provides enhanced efficiencies in storage and data backup administration, and are building blocks for building the dynamic, automated, and self-optimizing data center.

DECREASED COMPLEXITY

Ensuring reliable recovery typically involves locating appropriate hardware, installing an operating system, installing backup agents, modifying system configurations, and starting a recovery process. Testing is complex and can take several hours to several days. Complexity and time can directly impact business and organizational productivity. A virtual infrastructure significantly reduces complexity, enables rapid recovery times, and facilitates high availability.

LESS DOWNTIME

Solutions such as VMware's cutting-edge VMware High Availability and VMware VMotion technology empower IT staff to perform maintenance functions without any service shutdowns. They also help reduce unplanned downtime by proactively moving running applications away from servers that generate alerts or cross certain defined management thresholds.

THE DELL SOLUTION

Thanks to its close relationships with Intel, VMware, Altiris, and EMC, Dell offers comprehensive virtualization solutions that can help corporations bolster business continuity efforts while lowering IT costs.

These solutions include powerful, scalable, and reliable Dell PowerEdge servers based on Dual-Core Intel® Xeon® processors featuring silicon-level support for virtualization through Intel® Virtualization Technology; VMware's market-leading virtual infrastructure software; Altiris' acclaimed enterprise systems deployment and management tools; and top-rated Dell/EMC storage systems. And global Dell Assessment, Design, and Implementation Services can help companies realize the full benefits of virtualization, and meet tomorrow's business challenges more effectively.

By offering a single source for proven, integrated, market-leading virtualization solutions, Dell helps make business protection feasible for virtually all organizations today. For more information, visit **www.virtualization.ziffdavis.com.** ■

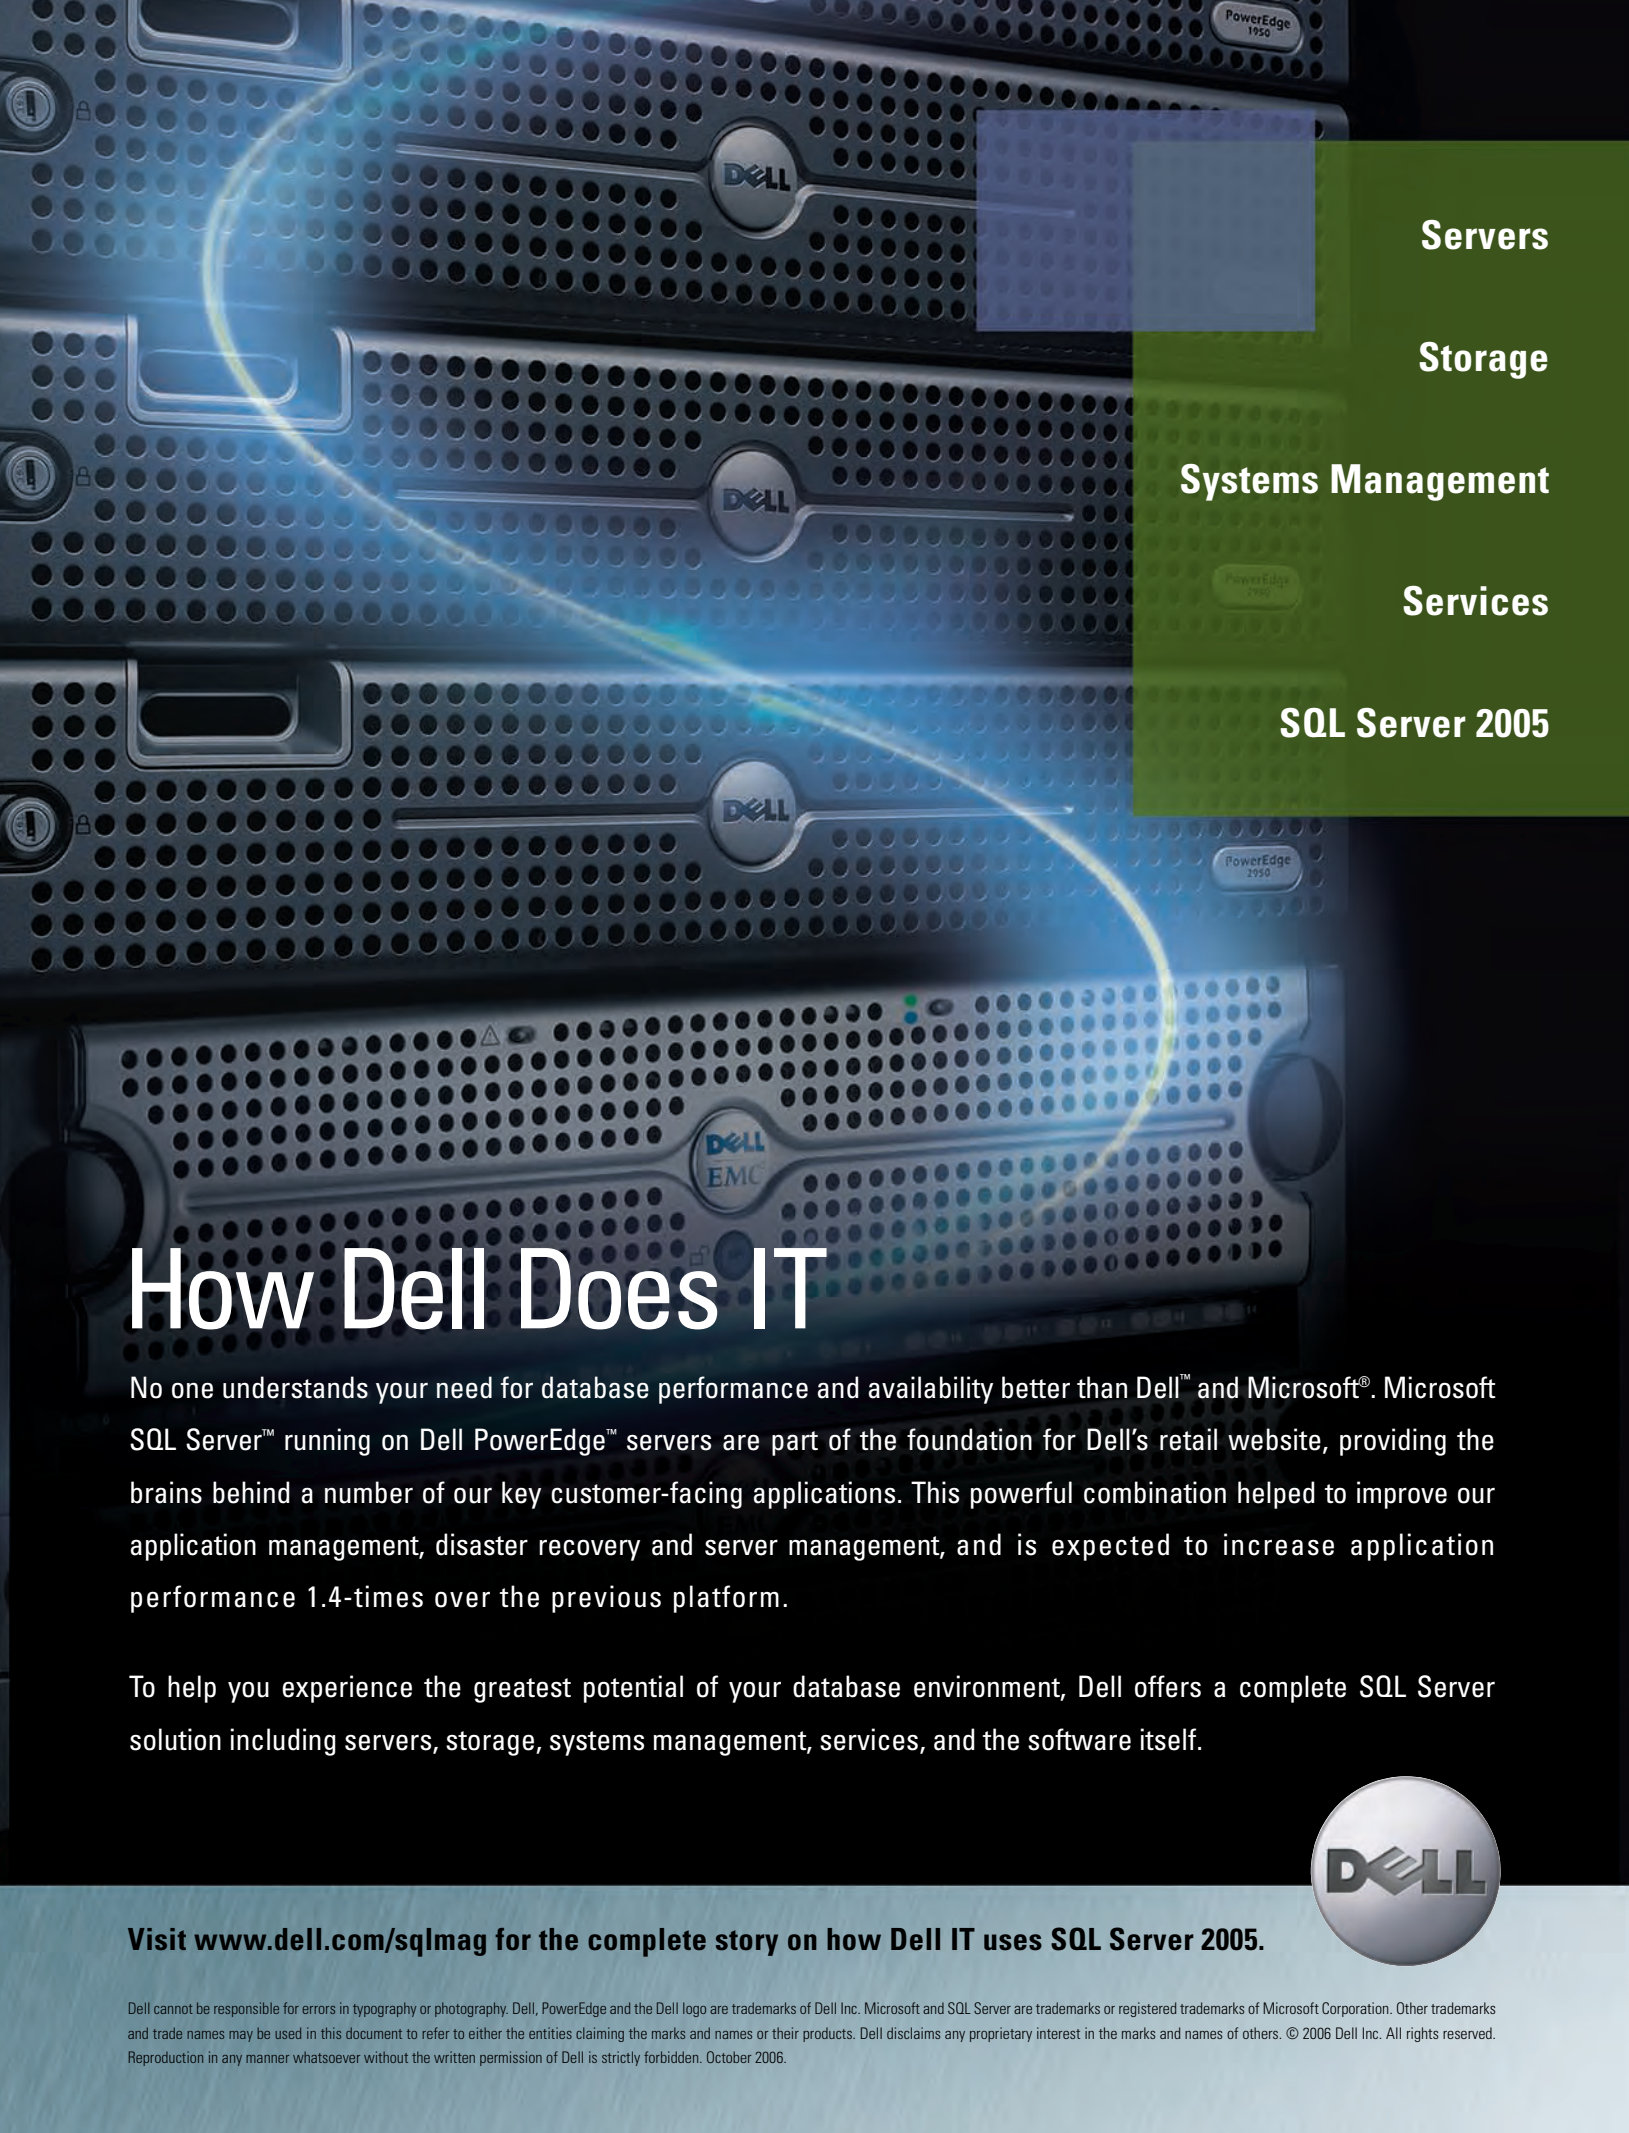

Servers

Storage

Systems Management

Services

SQL Server 2005

How Dell Does IT

No one understands your need for database performance and availability better than Dell™ and Microsoft®. Microsoft SQL Server™ running on Dell PowerEdge™ servers are part of the foundation for Dell's retail website, providing the brains behind a number of our key customer-facing applications. This powerful combination helped to improve our application management, disaster recovery and server management, and is expected to increase application performance 1.4-times over the previous platform.

To help you experience the greatest potential of your database environment, Dell offers a complete SQL Server solution including servers, storage, systems management, services, and the software itself.

Visit www.dell.com/sqlmag for the complete story on how Dell IT uses SQL Server 2005.