DELL™
AUGUST 2006 · $12.95

POWER SOLUTIONS

THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS

TEAR-OUT POSTER BONUS! STORAGE SOLUTIONS

Scaling Business Applications with New Servers and Storage

Dell PowerEdge Servers:
The Ninth Generation

Dell/EMC Storage:
The CX3 UltraScale Series Arrays

Virtualization:
Inside VMware ESX Server 3

Premier Issue:
Dell OpenManage Newsletter


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.

DELL™
POWER SOLUTIONS
THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS
AUGUST 2006

COVER STORY | PAGE 10

Scaling Business Applications with New Servers and Storage

By Mark Nickerson, Joe Pollock, Stori Waugh, and Stacy Hower

Ninth-generation Dell PowerEdge servers and Dell/EMC CX3 UltraScale series storage arrays offer a world of choices through open management, interoperability, and virtualization. By embracing standards to simplify operations, increase resource utilization, and scale out cost-effectively, these high-performance platforms help grow scalable enterprise applications responsively and manage change flexibly.

FEATURE SECTION: NEW-GENERATION SERVERS AND STORAGE

PURE DELL

Premier Issue Inside!
Dell OpenManage Newsletter | Flexible Management for the Scalable Enterprise

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

DELL ENTERPRISE COMMAND CENTER | PAGE 60

Nothing Basic About Next-Generation Enterprise Support

At the Dell Enterprise Command Center, specialists rigorously monitor Dell support operations and provide communication logistics anywhere in the world at a moment's notice—no matter what. Now, the new Platinum Plus Enterprise Support service provides secure access to a Web-based tracking window that lets customers monitor logistics and the status of Dell service dispatches in real time using Google Earth Pro.

Dell/EMC CX3-40 storage array

DELL POWER SOLUTIONS | THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS

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; Ahmad Ali; Rizwan Ali; Brad Anderson; Paul Barcoe-Walsh; Janet Berlind; Charles Butler; Balasubramanian Chandrasekaran; RadhaKrishna Dasari; Alan Daughetee; Drew Engstrom; Joshua Giles; Richard Goodwin; Kevin Guinn; Baris Guler; Pankaj Gupta; Ian Hameroff; Robert Hays; Stacy Hower; John Hull; Harish Jayakumar; Rezwanul Kabir; Sangram Kadam; Zain Kazim; Tim Lambert; Allen Light; Cynthia Lovin; Shabana M.; Linda Martinez; Jon McGary; Gajanan Mudaliar; Todd Muirhead; Tesfamariam Michael; Mark Nickerson; Laurie Nylund; Dean Oliver; Seema Padghan; Eric Pan; Weimin Pan; Jimmy D. Pike; Joe Pollock; Ramesh Radhakrishnan, Ph.D.; Edward Reynolds; Vishvesh Sahasrabudhe; Ananda Sankaran; David Schmidt; Sonya R. Sexton; Rohit Sharma; Pramada Singireddy; Scott Stanford; Cindy Stap; Stori Waugh; David Weber; Tony Yaptangco; Jianwen Yin, Ph.D.; and Haihong Zhuo

ART

ART DIRECTOR | Iva Frank
DESIGNER AND ILLUSTRATOR | Cynthia Webb
CONTRIBUTING ARTIST | Mark Mastroianni
COVER DESIGN | Iva Frank

MARKETING

SPECIAL INSERTS MANAGER | Erin Stolle

ONLINE

WEB DESIGN | Natanya Anderson
WEB PRODUCTION | Brad Klenzendorf

Subscriptions are free 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 47,800 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 cannot be responsible for errors in typography or photography. Dell, the Dell logo, Dell OpenManage, Dell Precision, ImageWatch, PowerConnect, PowerEdge, PowerVault, and ServiceSystem are trademarks of Dell Inc. Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

Dell Power Solutions is published quarterly by the Dell Product Group, Dell Inc. *Dell Power Solutions,* Mail Stop RR5-03, Dell Inc., One Dell Way, Round Rock, TX 78682, U.S.A. This publication is also available online at **www.dell.com/powersolutions**.

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.

Copyright © 2006 Dell Inc. All rights reserved. Printed in the U.S.A.

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

August 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**.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

PURE INSIGHT TO SYSTEM STATUS

NEW DELL POWEREDGE SERVERS WITH BUILT-IN LCD DISPLAYS.

Introducing the new PowerEdge family. It was born to reduce complexity. System status is available at a glance, allowing you to diagnose, maintain and conquer datacenter chaos like never before. For a 360° online view of this pure leap forward, visit www.dell.com/poweredge. Streamline without sacrifice. That's the direct path to success. That's pure Dell.

PURE DELL

www.dell.com/poweredge | 1.866.217.3056

Dell cannot be responsible for errors in typography or photography. Dell, the Dell logo and PowerEdge are trademarks of Dell Inc. © 2006 Dell Inc. All rights reserved.

Dell OpenManage Newsletter

Flexible Management for the Scalable Enterprise

Faced with legions of management tools that don't talk to each other, few IT administrators can see the big picture. This launch issue explores Dell's strategy for monitoring industry-standard platforms with best-of-breed management tools from a central console, Microsoft's initiatives for moving closer to the dynamic data center—and how these efforts can help systems management tools work better for you.


Dell Precision workstations

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

McDATA® | Mobilizing the World's Data™

McDATA's Full 4 Gb/s Complement to Dell's new Dell/EMC CX3 Series Arrays

The perfect SAN combination: Dell's new 3rd generation Dell/EMC CX storage arrays, combined with McDATA's 4 Gb/s SAN switches form the core of today's scalable enterprise infrastructure that's easy to deploy, manage, and upgrade.

4416 Embedded FC Switch for the Dell PowerEdge 1855 Blade Server

Whether you're establishing your first SAN, meeting exponential storage growth, or implementing a backup and archive strategy, McDATA FC switches optimize the SAN by helping you contain and control the "Top 3" data center costs:

- **Real Estate:** The McDATA 4400 offers the smallest 4 Gb/s switch footprint in the industry at 1U and half-rack width
- **Utilities:** McDATA switches have the lowest power consumption and cooling requirements in the industry
- **Maintenance:** The McDATA 4400 is the industry's only entry-level fabric switch with a redundant external power supply to limit the cost of downtime in the event of a power supply failure – the most common failure in any IT hardware product.

16 to 32 "FlexPort" 4700

8 to 16 "FlexPort" 4400 (½ rack!)

Download McDATA's latest white paper on moving from Direct Attached Storage (DAS) to a Storage Area Network (SAN) at *http://www.mcdata.com/info/das2san/*

©2006 McDATA Corporation. All rights reserved. OEM and Reseller logos are the property of such parties and are reprinted with limited use permission. All other trademarks are the property of their respective companies. All specifications subject to change.

Dell PowerEdge 1955 blade server

ADVERTISER INDEX

ONLINE XTRA

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.

Your Information, Your Way

Can't wait? For real-time access to content as soon as it's posted, subscribe your favorite RSS (really simple syndication) reader to the new feed for *Dell Power Solutions* online at **www.dell.com/powersolutions** or **www.dell.com/rss**. Searching for specific article content? Visit our Related Categories index online at **www.dell.com/powersolutions**.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Servers

Storage

Systems Management

Services

Software

ORACLE®

Don't miss **Michael Dell's** keynote at **Oracle OpenWorld 2006** in San Francisco on October 23.

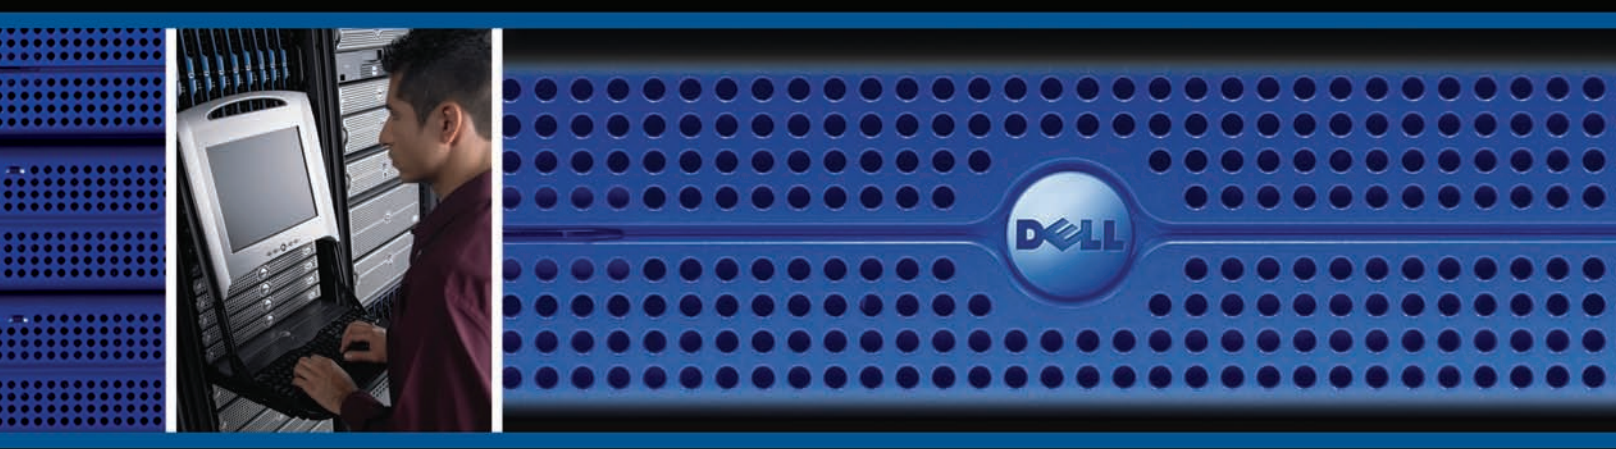

Dell Tests and Validates Solutions for Oracle Database 10*g*

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's Scalable Enterprise vision perfectly complements Oracle's Grid strategy, both focused on bringing you a flexible, scalable, cost-effective database solution.

DELL

Visit www.dell.com/oraclemag to see Dell's Tested & Validated Solutions for Oracle Database 10*g*.

Dell cannot be responsible for errors in typography or photography. Dell and the Dell logo are trademarks of Dell Inc. Windows is a registered trademark of Microsoft Corporation. Oracle is a registered trademark of Oracle Corporation. Linux is a registered trademark of Linus Torvalds. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others. © 2006 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the written permission of Dell is strictly forbidden.

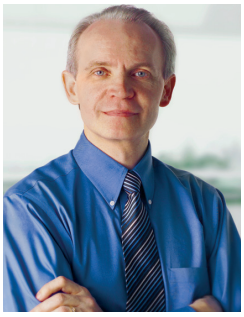

Through the Eyes of the 50

The *Dell Power Solutions* editorial team is, by design, quite small. So, in order to cover the full scope of enterprise IT issues in depth, we depend on an extensive web of subject-matter experts from diverse disciplines that live, breathe, and exercise the latest technology. For a typical issue of the magazine, we tap into the collective intellect and unique perspectives of over 50 such contributing authors (at last query, our database of *Dell Power Solutions* contributing authors had topped 750).

In planning the August 2006 issue, we connected with a particularly strong knowledge base to cover the very latest developments in server, storage, and workstation technology available from Dell. In our eight-page cover story, "Scaling Business Applications with New Servers and Storage," a quartet of product managers helps articulate the business value of the recently announced ninth-generation Dell™ PowerEdge™ servers, Dell/EMC CX3 UltraScale™ series storage arrays, and Dell Precision™ workstations in the context of scaling enterprise applications. Plus, you can read about how search-engine provider FAST uses the Dell PowerEdge 2950 server to strategic advantage.

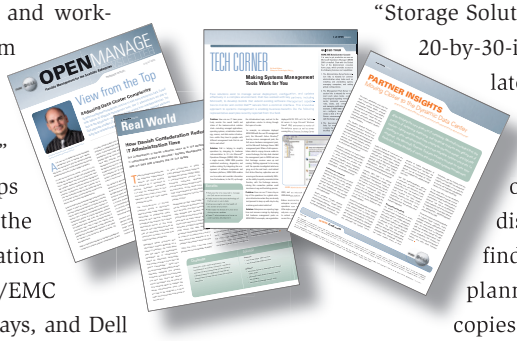

Staying true to our roots as a technical journal, we also engaged with many of the chief technologists, engineers, and planners responsible for designing and delivering the new Dell PowerEdge servers and Dell/EMC storage, asking them to take an editorial respite from their normal environs of CAD consoles and white lab coats. Their efforts resulted in over 24 pages of fresh drill-down content in the feature section, including need-to-know coverage on ninth-generation Dell PowerEdge server performance tuning, TCP/IP Offload Engine technology, management via Intelligent Platform Management Interface, the Dell Remote Access Controller 5, continuous data protection on Fibre Channel storage, and much more.

In recognition of the growing concern around enterprise systems management, we are very pleased to launch the *Dell OpenManage Newsletter*—a unique one-on-one dialogue with individuals involved in architecting Dell's flexible, open approach to managing a scalable enterprise, as well as insights from Dell customers, partners, and technical support teams. Starting on page 81, the premier issue focuses on Dell's strategy for reducing data center management complexity and integration with Microsoft technologies.

Facing page 48, you will find a special tear-out poster: "Storage Solutions for the Scalable Enterprise." This 20-by-30-inch, four-color schematic shows the latest Dell/EMC storage, Dell PowerVault™ storage, and Dell PowerEdge servers configured in small office/branch office, data center, and regional office/disaster recovery scenarios. We hope you find this helpful in your enterprise storage planning efforts. To download additional copies of the poster in PDF format, visit *Dell Power Solutions* online at www.dell.com/powersolutions.

Hats off to the 58 expert advisers and contributing authors featured in this issue. We depend on your peer-to-peer participation as we work together to advance best practices for enterprise IT.

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

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

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 and the Symantec logo are U.S. registered trademarks. Backup Exec and LiveState are trademarks of Symantec Corporation. Copyright ©2006 Symantec Corporation. All rights reserved.

symantec™

Scaling Business Applications

with New Servers and Storage

Today's IT organizations face unprecedented challenges brought on by globalization, rapid expansion, limited resources, and increasing prominence of core enterprise applications. Ninth-generation Dell™ PowerEdge™ servers and Dell/EMC CX3 UltraScale™ series storage arrays offer a world of choices through open management, interoperability, and virtualization. By embracing standards to help simplify operations, increase resource utilization, and scale out cost-effectively, these high-performance platforms grow scalable enterprise applications responsively and manage change flexibly.

BY MARK NICKERSON, JOE POLLOCK, STORI WAUGH, AND STACY HOWER

Related Categories:

Case study

Dell ninth-generation servers

Dell PowerEdge blade servers

Dell PowerEdge servers

Dell Precision workstations

Dell/EMC storage

Scalable enterprise

Storage management

Systems management

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

In a world of rapid growth and dizzying change, core IT solutions such as Microsoft® Exchange, Microsoft SQL Server™ 2005, mySAP™ Business Suite, and Oracle® Database 10g—together with the advent of virtualization—are forcing a paradigm shift in today's business environment. The proliferation of servers and storage systems deployed to support explosive data growth while meeting stringent regulatory requirements has led to mounting complexity. At the same

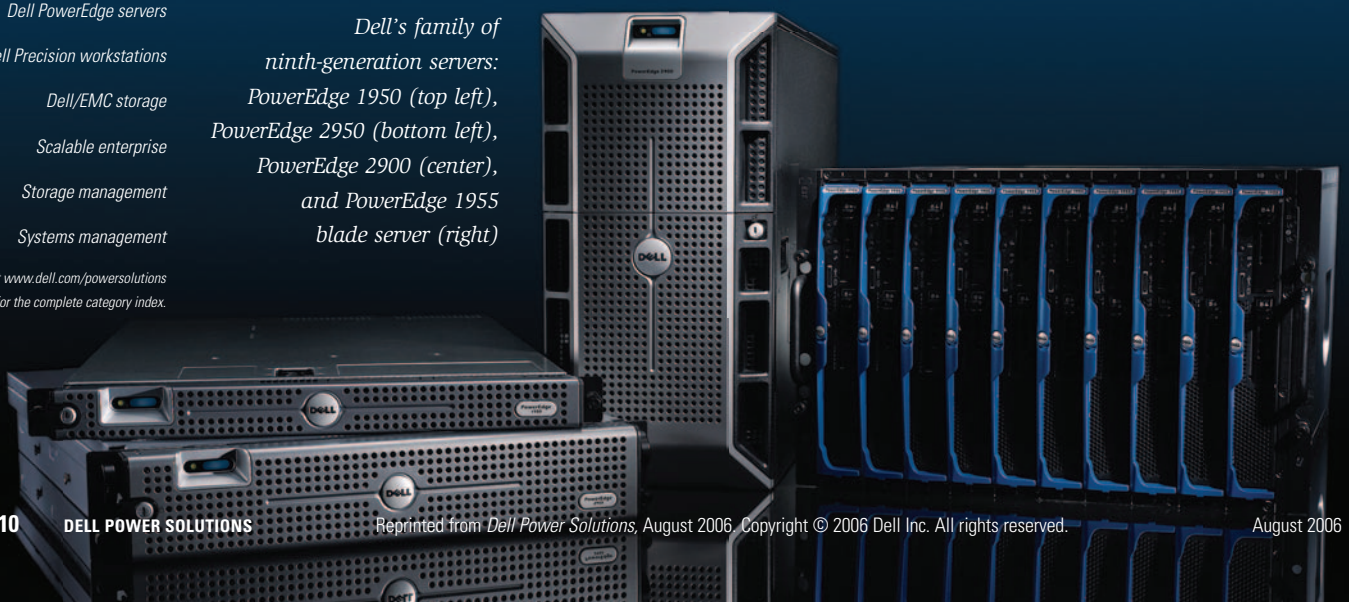

Dell's family of ninth-generation servers: PowerEdge 1950 (top left), PowerEdge 2950 (bottom left), PowerEdge 2900 (center), and PowerEdge 1955 blade server (right)

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

time, many administrators must still manually manage change, one server or storage array at a time, using a number of different tools and interfaces to oversee diverse application environments.

The proliferation of systems management tools often results from a one-application-per-server mindset that leads to wasted resources. All too often, when spikes in system demand require maximum compute power and availability, high-performance servers are not architected in a way that enables fast, flexible workload balancing to make the most of available data center resources.

Meanwhile, enterprises continue to demand more from their IT organizations, either holding IT staff levels constant (do more with the same resources) or decreasing headcount (do more with fewer resources). Such constraints present a major challenge for IT as pressure mounts to increase productivity and scale out business applications quickly, flexibly, and cost-effectively.

To address these critical business concerns, Dell has developed a new generation of server, storage, and workstation products that advance the scalable enterprise strategy—continuing the drive to standardize core infrastructure elements and reduce the complexity of the computing infrastructure so that enterprises can simplify operations, improve resource utilization, and scale out in pragmatic, cost-effective increments. Because Dell develops its products and solutions around open standards, this approach is designed to result in interoperability that enables businesses to create high-power, end-to-end solutions from standard hardware and software building blocks. Through a high level of interoperability, as well as innovative hardware and software design, Dell's new server, storage, and workstation products help enterprises meet business requirements today with platforms that are designed from the ground up for future growth. And, on the IT front, this highly standardized approach can also help scale human resources more effectively.

Ninth-generation servers: Reducing complexity by design

To efficiently support demanding business applications, ninth-generation Dell PowerEdge servers are designed with the latest performance technologies, including Intel® Xeon® processors, fully buffered dual in-line memory modules (DIMMs), Serial Attached SCSI (SAS) hard drives, and TCP/IP Offload Engines. For detailed comparisons of the new PowerEdge servers, see the sidebar "Innovative New Platforms Built to Mix Well with Others."

Hardware and software commonality

Dell has made significant hardware changes to increase commonality across platforms, which helps to reduce the complexity of server management. The resulting simplification in server management helps boost productivity, enabling IT staff to spend less time on routine server management tasks and more time on value-added projects that contribute to the bottom line.

Dell/EMC CX3-80 storage array

To help ensure consistency across the new models as well as future generations, Dell has developed a behavioral specification that defines interaction and design guidelines for the servers. Handles, latches, and even the placement of components are common across multiple platforms. For example, once an administrator knows how to service a PowerEdge 1950 server, the learning curve is sharply reduced for servicing the PowerEdge 2950 server, PowerEdge 2900 server, and PowerEdge 1955 blade server. Clear, consistent labeling, including all hot-pluggable hard drives, helps simplify the servicing and upgrading of these systems. In addition, an always-visible LCD status panel is now standard across ninth-generation rack and tower servers, featuring administrator-level programmability for enhanced staff productivity.

Following the path set in earlier generations, Dell also continues to reduce the complexity of software updates through its industry-leading software commonality. For example, administrators typically spend a lot of time managing a server's master system image, which comprises the BIOS, system drivers, OS, and applications. Administrators update and then copy this image to all production servers when software components change. In the ninth generation, Dell has reduced complexity so that Dell PowerEdge 1950, PowerEdge 2950, and PowerEdge 2900 servers can be maintained with one system image—compared to other server offerings that require a unique image for each server platform because of an unshared BIOS. By reducing the number of system images to maintain, organizations can reduce server administration time. In addition, the Dell ImageWatch™ program notifies administrators of upcoming changes, helping to proactively manage system images and reduce the number of changes. And through its block release process, Dell groups software changes to help minimize the total number of image changes that administrators need to manage.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Search and Deploy

FAST quickly turns to Dell PowerEdge 2950 servers as a way to meet the computationally-intense needs of its enterprise search infrastructure.

When asked what Norway is known for, many people might say Edvard Munch and his famous painting, *The Scream*. But if the engineers at FAST get their way, Norway will become more famous for screaming-fast enterprise search. That is because the Oslo-based company hosts many of the enterprise search applications it delivers for businesses and organizations worldwide. FAST's operations span six continents and serve more than 3,500 organizations, including many of the world's largest companies.

Searching for a way to meet enormous processing requirements

FAST offers a wide range of search-driven solutions, from applications solving specific vertical business challenges, such as online merchandizing, to solutions that apply to most industries, such as competitor monitoring. Some companies have even built their businesses on FAST, such as online directory providers. The common theme is the need to search and index billions of documents or mixed media, regardless of location, data type, language, or formatting—and to deliver answers or services based on the results of these searches.

Customers rely on the FAST data center to collectively support millions of queries per second to keep their mission-critical applications delivering to expectation. In turn, FAST relies on Dell PowerEdge servers and storage devices to deliver the performance, scalability, and reliability that FAST's customers demand. Dell products form the foundation of the FAST IT infrastructure. FAST continues to depend on Dell because the value and performance delivered by Dell products has helped the company scale to become a runaway success in the enterprise search marketplace.

Promoting server consolidation

As part of the Dell evaluation program, FAST was provided with PowerEdge 2950 servers before the general release in order to take them for a test ride. "We compared the PowerEdge 2950 against our current servers in a race to index one million documents,"

Enhanced systems management tools

Continuing to drive industry-standard management, Dell offers enhanced systems management tools that are designed to enable enterprises to manage Dell platforms with great flexibility, security, and control. Two new hardware features help increase the manageability of ninth-generation Dell PowerEdge servers:

- **Intelligent Platform Management Interface (IPMI) 2.0 support:** Embedded server management now adheres to the IPMI 2.0 specification, helping improve administrator efficiency by providing industry-standard tools for managing hardware, monitoring systems, and communicating faults.[1]
- **Dell Remote Access Controller 5 (DRAC 5):** The new DRAC 5 delivers enhanced features and functionality for administrators who require a remote console connection regardless of the status of the server. DRAC 5 builds on DRAC 4 features, such as continuous video, virtual media, and integration with the Microsoft Active Directory® directory service, and introduces important features such as USB virtual media support, Secure Sockets Layer (SSL)–encrypted virtual media, and industry-standard Systems Management Architecture for Server Hardware Command-Line Protocol (SMASH CLP).[2]

The Dell OpenManage™ 5 suite of comprehensive deployment, monitoring, and change-management tools provides excellent flexibility and security for managing Dell PowerEdge servers. Dell OpenManage 5 software supports industry-leading management standards and frameworks and can now manage printers, Dell/EMC storage, and network-attached tape automation systems.

Dell has also integrated its systems management tools with leading enterprise management applications and infrastructures—including Microsoft Operations Manager (MOM) 2005, Microsoft Systems Management Server (SMS) 2003, and Altiris® Server Management Suite™—to further extend the versatility and scope of management capability for Dell servers.

On the Linux® OS front, Dell and Novell have developed Novell® ZENworks® Linux Management–Dell Edition, the first fully integrated Linux deployment and change-management solution for Dell servers. Using ZENworks Linux Management–Dell Edition, administrators can update and patch the entire Dell-Linux infrastructure from a single console.

[1] For more information, see "Improved Management Using the LAN and Serial Interfaces in Ninth-Generation Dell PowerEdge Servers" by Tim Lambert, Rohit Sharma, and Harish Jayakumar in *Dell Power Solutions,* August 2006, www.dell.com/downloads/global/power/ps3q06-20050237-Lambert.pdf.

[2] For more information, see "Exploring the DRAC 5: The Next-Generation Dell Remote Access Controller" by Jon McGary and Weimin Pan in *Dell Power Solutions,* August 2006, www.dell.com/downloads/global/power/ps3q06-20060118-McGary.pdf.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

explains Demos Skipitaris, senior operations architect at FAST. "The PowerEdge 2950 consistently outperformed our current servers. As our older PowerEdge 2650 servers reach the end of their life cycle, we think we can replace 10 of them with around eight of the new PowerEdge 2950 servers based on this performance boost."

Enabling storage consolidation

In addition to promoting server consolidation, the presence of expanded disk capacity in the new PowerEdge 2950 servers helps promote storage consolidation by reducing the dependence on direct attach storage devices. "Although we currently use and like Dell PowerVault enclosures for added storage capacity, the new PowerEdge 2950 servers—with eight internal drives—will quickly become very valuable to us," explains Skipitaris. "When the 146 GB, 2.5-inch SAS drives become available, a single PowerEdge server will have the capability to provide an application with more than 1 TB of internal data storage. As processor speeds increase, the added disk space will make it possible to store more documents on a server, so the cost per searchable document will go down."

And that is not all the new PowerEdge 2950 servers have to offer. Dell makes migrating from earlier platforms as simple as possible by maintaining hardware consistency across generations. Through incremental but significant improvements to the design such as the programmable LCD screen that is visible with the bezel on, easy-open draw-latch, and unique color-coding that indicates which items are hot-pluggable, Dell continues to deliver on the promise to help reduce complexity and ease manageability. Skipitaris says, "I noticed many enhancements. One that I really like is that the PowerEdge 2950 servers have aligned the ports on the left and the power supplies on the right, making our cabling neat and clean. Plus, there is a cable management clip to hold all the cables together. This may seem like a small thing, but it makes a big difference to us."

Helping to reduce data center costs

By offering the potential for server and storage consolidation, PowerEdge 2950 servers help enterprises keep procurement costs and real estate–related expenses low. Plus, the new Intel Xeon processor was designed to reduce power requirements and cooling costs. Ultimately, the new PowerEdge 2950 servers have been designed to help enterprises get more performance while lowering data center costs. As Skipitaris concludes, "Dell is already the basis of our infrastructure, and with the release of the PowerEdge 2950 server, we have even more reason to stick with Dell."

New Dell/EMC storage: Easing SAN deployment

Dell lowered the barrier to entry for growing organizations requiring enterprise-caliber networked storage by pioneering the easy-to-deploy Dell/EMC AX150 and AX150i—storage area network (SAN) arrays starting at under $10,000 that can be configured with only a few clicks of the mouse. Now Dell extends the ease of deployment to its new generation of Dell/EMC CX3 UltraScale series SAN arrays.

The Dell/EMC CX3 series provides an end-to-end 4 Gbps architecture designed to lead the midrange storage array market in performance and scalability. Third-generation CX3 series arrays offer breakthrough architecture with an expansive range of scalability options that provide a platform for cost-effective growth throughout the SAN life cycle. For specifications of the CX3 series arrays, see the sidebar "Innovative New Platforms Built to Mix Well with Others."

For the introduction of the CX3 series, Dell worked to streamline the deployment and management of SANs through a series of integration efforts. As a result, Dell PowerEdge servers and Dell/EMC storage arrays can be even more closely coupled than before. The first step led to servers that are customized with SAN components—such as Fibre Channel host bus adapters (HBAs) and storage management software—preinstalled during manufacturing through Dell's build-to-order process. The HBAs enable direct connectivity to a storage array or SAN fabric through a Fibre Channel switch. The storage management software includes the EMC® Navisphere® Server Utility and the EMC PowerPath® path management tool. Preinstallation of these SAN components helps reduce SAN deployment time and the risk of installation errors—simplifying the initial SAN deployment and minimizing the time to deploy additional servers to an existing SAN.

The next step in Dell's effort to reduce complexity led to the integration of Dell and EMC management software. Version 8.0 of Dell OpenManage IT Assistant, which provides an integrated view of Dell's comprehensive suite of server monitoring and reporting tools, is now integrated with EMC Navisphere software—enabling complete management of a Dell/EMC or EMC CLARiiON® storage environment. Together, the two software products enable top-level administration of key SAN elements through a single management console, helping to simplify overall SAN management.

Dell PowerVault solutions: Extending the storage domain

For applications configured with direct attach storage, the recently enhanced Dell PowerVault™ line includes the Dell PowerVault

(Continued on page 16)

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Innovative New Platforms Built to Mix Well with Others

Dell PowerEdge 1950 server (top left), PowerEdge 2950 server (bottom left), and PowerEdge 2900 server (right)

High-performance, industry-standard design and software commonality promote interoperability and scalability of Dell's new servers, storage systems, and workstations. Plus, flexible management software and options for configuring IPMI settings together with DRAC 5 functionality enhance the efficiency of server administration—helping businesses do more, more easily.

Ninth-generation Dell PowerEdge servers

Figure A shows application environments suitable for ninth-generation PowerEdge server models, which are designed not only to help increase flat-out performance but also to help improve price/performance and performance per watt:

- **Dual-core Intel Xeon processors:** The new PowerEdge servers use dual-core Intel Xeon processors, increasing performance up to 69 percent with the 5000 series[a] compared to eighth-generation Dell servers; with the 5100 series, performance increases up to 152 percent[b] while power consumption decreases by as much as 25 percent[c] compared to eighth-generation Dell servers. The resulting decrease in heat helps businesses scale their data centers without increasing power input or thermal output.
- **Fully buffered DIMMs:** Fully buffered DIMMs provide a point-to-point serial memory interface for fast, reliable data transmission—up to three times greater throughput and up to three times greater capacity compared to previous Intel Xeon–based servers[d]—plus excellent signal integrity and error detection.
- **SAS hard drives:** SAS drives optimize data transmission through serialization while enhancing reliability through the use of small hard-drive form factors and cables—both of which are designed to consume less power and reduce thermal output compared to traditional SCSI hard drives.
- **TCP/IP Offload Engine:** A dedicated processor on the servers' embedded Broadcom 5708 Gigabit Ethernet[e] network interface cards (NICs) offloads TCP/IP traffic from the host processor—a boost for database and backup server applications as well as other applications with large I/O packets.
- **PCI Express:** PCI Express provides a high-performance Ethernet, RAID, InfiniBand, and Fibre Channel interconnect.

Dell PowerEdge storage servers

Dell PowerEdge 2900 and PowerEdge 2950 servers featuring Microsoft® Windows® Storage Server 2003 R2 are scalable network-attached

Model	Features	Suitability to task
PowerEdge 1950	Dual-socket, rack-dense 1U server	Virtualization, Web applications, SAN compute node, network infrastructure, database front end, e-mail and messaging
PowerEdge 2950	Dual-socket, rack-mount 2U server	Virtualization, network infrastructure, database, e-mail and messaging, Internet applications
PowerEdge 1955	Dual-socket, ultra-rack-dense blade server	Distributed Web environments, network infrastructure, terminal services, virtualization
PowerEdge 2900	Dual-socket 5U tower	E-mail and messaging, application server, database server, Web server

Figure A. Ninth-generation Dell PowerEdge servers

[a] Based on the SPECjbb2005 benchmark test performed by Dell Labs in December 2005 and April–May 2006 on a PowerEdge 2950 server with two dual-core Intel Xeon 5080 processors at 3.73 GHz (Dempsey); 8 GB, 533 MHz fully buffered DIMM (FBD) memory; one Serial ATA (SATA) 80 GB, 7,200 rpm hard disk drive; and Windows Server 2003 Enterprise x64 Edition OS, as compared to a PowerEdge 2850 server with two dual-core Intel Xeon processors at 2.8 GHz (Paxville); 8 GB, 400 MHz double data rate 2 (DDR2) memory; one SCSI 36 GB, 15,000 rpm hard disk drive; and Windows Server 2003 Standard Edition OS with Service Pack 1 (SP1).

[b] Based on the SPECjbb2005 benchmark test performed by Dell Labs in December 2005 and April–May 2006 on a PowerEdge 1950 server with two dual-core Intel Xeon 5160 processors at 3.0 GHz (Woodcrest); 8 GB, 667 MHz FBD memory; one SATA 80 GB, 7,200 rpm hard disk drive; and Windows Server 2003 Enterprise x64 Edition OS, as compared to a PowerEdge 2850 server with two dual-core Intel Xeon processors at 2.8 GHz (Paxville); 8 GB, 400 MHz DDR2 memory; one SCSI 36 GB, 15,000 rpm hard disk drive; and Windows Server 2003 Standard Edition OS with SP1.

[c] Based on testing performed by Dell Labs in May 2006 using the SPECjbb2005 benchmark on a PowerEdge 2950 server with two dual-core Intel Xeon 5160 processors at 3.0 GHz (Woodcrest) and then with two dual-core Intel Xeon 5080 processors at 3.73 GHz (Dempsey); 4 GB, 667 MHz and 533 MHz FBD memory; two SAS 73 GB, 15,000 rpm hard disk drives; and Windows Server 2003 Enterprise x64 Edition OS, as compared to a PowerEdge 2850 server with two dual-core Intel Xeon processors at 2.8 GHz (Paxville); 4 GB, 400 MHz DDR2 memory; two SCSI 36 GB, 15,000 rpm hard disk drives; and Windows Server 2003 Enterprise x64 Edition OS.

[d] Source: "Fully-Buffered DIMM Technology Moves Enterprise Platforms to the Next Level" by Jon Haas and Pete Vogt, *Technology@Intel Magazine,* March 2005, www.intel.com/technology/magazine/computing/fully-buffered-dimm-0305.pdf. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests.

[e] 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.

[f] Based on testing performed by Dell Labs in January 2006 using the SPEC CPU2000 (integer-rate base) benchmark on a Dell Precision 690 workstation configured with two dual-core Intel Xeon 5150 processors at 2.66 GHz; 4 MB level 2 (L2) cache; 1,333 MHz FSB (no Intel Hyper-Threading Technology available); 8 GB, 533 MHz DDR2 error-correcting code (ECC) FBD memory; SATA 160 GB, 7,200 rpm hard disk drive; and Microsoft Windows XP Pro SP2 as compared to a Dell Precision 670 workstation with two single-core Intel Xeon processors at 3.8 GHz; 2×2 MB L2 cache; 800 MHz FSB (Intel Hyper-Threading Technology off); 2 GB DDR2 ECC memory; SATA 160 GB, 7,200 rpm hard disk drive; and Microsoft Windows XP Pro SP2. **NOTE:** For performance improvement claims cited in footnotes a, b, c, and f, actual performance will vary based on configuration, usage, and manufacturing variability.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Model	Features	Suitability to task
CX3-20	• Provides four front-end Fibre Channel optical ports and/or eight Internet SCSI (iSCSI) ports that are capable of connecting up to 12 direct attach or 128 redundantly SAN-connected servers • Supports up to eight expansion enclosures, for a maximum of 120 back-end drives (up to 59 TB raw capacity) • Simultaneous connection of Fibre Channel and iSCSI servers	E-mail and messaging, database, disk-based backup and restore, audio-video streaming
CX3-40	• Provides four front-end Fibre Channel optical ports and/or eight iSCSI ports that are capable of connecting up to 12 direct attach or 128 redundantly SAN-connected servers • Supports up to 16 expansion enclosures, for a maximum of 240 back-end drives (up to 119 TB raw capacity) • Simultaneous connection of Fibre Channel and iSCSI servers	Large databases, e-mail and messaging, remote replication, online transaction processing
CX3-80	• Provides eight front-end Fibre Channel optical ports for direct connection of up to eight hosts • Supports up to 256 dual-connected hosts in a SAN • Supports up to 32 expansion enclosures, for a maximum of 480 drives (up to 239 TB raw capacity)	Data warehousing, remote replication, very large-scale backup to disk, seismic data processing, computational chemistry
DAE4P	• Provides a 3U enclosure • Supports up to fifteen 4 Gbps or 2 Gbps drives in the same array • Supports both standard Fibre Channel and low-cost Fibre Channel (LCFC) drives for maximum choice in performance and capacity	Additional expansion for CX3 arrays

Figure B. Third-generation Dell/EMC CX3 series storage arrays and enclosures

storage servers, specially designed to help midsize and enterprise environments deploy file sharing and other essential services quickly:

- Optimized for file sharing; preloaded with advanced capabilities for efficient storage management
- Ready to go right out of the box, in a matter of minutes
- Deployable anywhere on the Ethernet network
- Manageable remotely from any Microsoft Windows–based desktop through a convenient administrative console

The new storage servers provide the usability of a storage appliance with the enhanced price/performance and flexibility of the ninth-generation PowerEdge 2900 and PowerEdge 2950 servers—backed by Dell PowerVault and Dell/EMC storage systems.

Dell/EMC CX3-40 storage array

Model	Features	Suitability to task
Precision 490	• Up to 32 GB FBD ECC memory • Up to three hard disk drives, for a potential capacity of 1.5 TB • Small footprint • Support for a wide range of high-performance graphics cards, from 2-D to OpenGL 3-D, with optional dual-monitor capability	Compute-intensive applications in crowded, power-sensitive environments, such as financial trading floors, performance clusters, and render farms
Precision 690	• Up to 64 GB FBD ECC memory • Up to five hard disk drives, for a potential capacity of 2 TB • Innovative, highly expandable chassis • Dual-graphics platform that can support quad-monitor configurations as well as deliver ultra-high 3-D graphics using NVIDIA Scalable Link Interface (SLI) technology	Compute-intensive, graphically demanding applications such as high-end computer-aided design, life sciences imaging, and real-time processing of huge data sets for the oil and gas industry

Figure C. New Dell Precision workstations

Dell/EMC CX3 UltraScale series storage arrays

The Dell/EMC CX3 UltraScale series of networked storage arrays (see Figure B) form the core of a scalable infrastructure providing an end-to-end 4 Gbps architecture enabling the following benefits:

- Dramatically improved performance compared to previous-generation CX series arrays, enabling faster data access and shorter response times
- Up to double the capacity of previous-generation CX series arrays
- State-of-the-art storage processors, 4 Gbps host interconnect, 4 Gbps back-end topology, and 4 Gbps drives
- Easy integration with Gigabit Ethernet, 2 Gbps, and 4 Gbps SAN components
- Easy mix and match of storage tiers within the array to optimize performance or capacity
- Enhanced data integrity and availability

Dell Precision workstations

The new Dell Precision™ 490 and Precision 690 workstations (see Figure C) offer high performance, excellent graphics, and very large memory for memory-intensive applications, multi-threaded applications, multi-tasking environments, or demanding single-threaded applications:

Dell Precision 690

- **Dual-core Intel Xeon processors:** When equipped with two dual-core Intel Xeon processors, a Dell Precision 690 workstation can deliver performance increases of up to 155 percent in multi-threaded applications, compared to the previous generation.[f]
- **Dual independent frontside buses (FSBs):** Dual independent FSBs, each capable of up to 1,333 MHz, offer extremely high-speed data transfers.
- **Fully buffered DIMMs:** Four fully buffered DIMM channels provide enormous bandwidth for data movement between memory and the system.
- **Innovative chassis:** New enclosures enable excellent flexibility and scalability.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

(Continued from page 13)

MD1000 Disk Expansion Enclosure, a 15-drive system that can be outfitted with the latest SAS or SATA II disk drives. Designed and tested to work seamlessly with ninth-generation PowerEdge servers, the MD1000 enclosure can store up to 7.5 TB when fully equipped with fifteen 500 GB, 7,200 rpm SATA II drives.

Backup, recovery, and archiving are also chief concerns for IT organizations, and the latest tape automation systems from Dell are designed to scale in step with growing business requirements. For example, the modular Dell PowerVault ML6000 Tape Library series offers versatile, enterprise-level SAN backup with scalability up to 87 TB (native) and future scalability up to 161 TB (native). Also new on the scene is the Dell Tape Automation Sizing Tool (www.dell.com/tapesizing), which can help administrators plan effectively for tape automation needs.

New workstations: Boosting performance cost-effectively

The Dell Precision 490 and Precision 690 workstations are based on industry standards, providing the latest performance-enhancing technologies—such as multi-core processing, 64-bit computing, and OpenGL graphics—at an affordable price. Dell designed its new workstations to deliver excellent performance in highly threaded applications and in complex multi-tasking environments such as financial trading, digital content creation, oil and gas, medical imaging, and high-end computer-aided design. The sidebar "Innovative New Platforms Built to Mix Well with Others" describes the new workstations in more detail.

Scalable enterprise: Creating a stable foundation

Designed with software commonality and industry-leading interoperability in mind, the new Dell server, storage, and workstation platforms can help businesses achieve the benefits of the scalable enterprise strategy: simplified operations, improved resource utilization, and cost-effective growth in pragmatic, planned phases.

Simplifying operations in the IT infrastructure

Through innovative design and engineering, Dell has improved the ease of operation and standardization of its ninth-generation platforms compared to earlier Dell generations. Software commonality and integration help make maintenance tasks predictable and efficient. This leads to a reduction in complexity that helps maximize productivity—allowing IT staff to spend more time on contributions that enhance the value of business. For example:

- **SAN scale-out:** Integration between servers and storage systems streamlines SAN deployment and maintenance.
- **System administration:** Integration of EMC Navisphere with Dell OpenManage IT Assistant; tightly coupled plug-ins for MOM 2005, SMS 2003, and Altiris Server Management Suite; and Novell ZENworks Linux Management–Dell Edition help reduce the number of consoles required to keep systems running smoothly—minimizing the complexity of ongoing operations.
- **Server maintenance and deployment:** Reduction in the number of system images to manage on Dell servers allows administrators to spend less time updating system software and tracking BIOS or driver updates.

Improving utilization of system resources

By implementing server and storage consolidation, businesses can enhance management control and utilization of their IT infrastructures. Although consolidation is a way to collect disparate computing resources, consolidating onto complex platforms typically moves the pressure point from one spot to another. The benefits of consolidation clearly emerge when businesses consolidate on industry-standard platforms such as the new Dell servers and storage, which are designed to reduce complexity and promote integration.

The adoption of technologies such as virtualization also helps enterprises improve resource utilization. For more information

STORAGE SOLUTIONS FOR THE SCALABLE ENTERPRISE

Dell's highly scalable server, storage, and management platforms are based on industry standards to enable optimal performance, availability, and interoperability in any size enterprise. For schematics of highly scalable storage solutions in small office/branch office, data center, and regional office/disaster recovery scenarios, see the tear-out poster facing page 48. To download additional copies of the poster, visit www.dell.com/powersolutions.

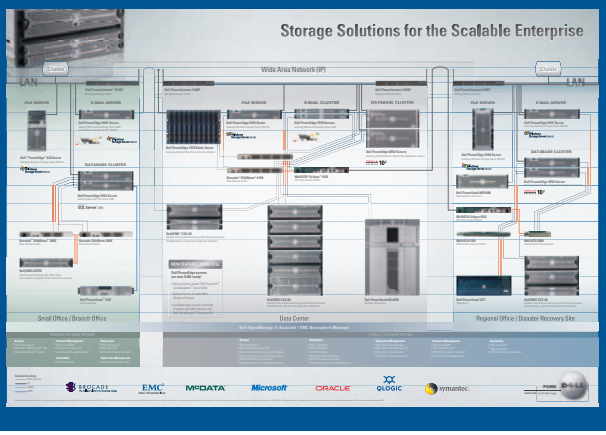

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

about the attributes and benefits that server virtualization may bring to the scalable enterprise, see "Server Virtualization in the Scalable Enterprise" by Jimmy D. Pike and Drew Engstrom in *Dell Power Solutions,* August 2006, www.dell.com/downloads/global/power/ps3q06-20060381-Pike.pdf.

Scaling out to meet evolving needs

Simplifying operations and improving resource utilization leads to the third opportunity enabled by the scalable enterprise: cost-effective scale-out. Using industry-standard building blocks and integrated systems management tools, enterprises can grow the IT infrastructure in practical, affordable increments—allocating computer and human resources as they are needed.

The modularity of Dell servers and storage systems facilitates the management of unexpected growth. Spikes in business can quickly be handled by adding incremental capacity, such as additional servers and storage, or by scaling to more advanced technology, such as the forthcoming quad-core Intel Xeon processors or 4 Gbps Fibre Channel infrastructures. Industry-standard data center components enable modular expansion in practical, planned phases—helping organizations to eliminate overbuying patterns for up-front investments based on projected needs and instead to maximize their budget by growing the IT infrastructure in pace with actual business requirements.

Through the flexibility of the new ninth-generation server architecture, features such as SAS storage, PCI Express, and fully buffered DIMMs provide the capacity to meet future business needs within existing systems. From a storage perspective, the ability to deploy early components for a 4 Gbps infrastructure today allows businesses to continue using existing SAN architectures while simultaneously preparing for the future.

Planning for growth and stability

Because one size does not fit all, Dell has designed its Server Advisor tool to help enterprises zero in on the most suitable server platform for their specific application environment. This online tool asks a brief series of questions about the number of users; performance and availability needs; network infrastructure; print-and-file service; e-mail, messaging, and collaboration systems; business applications and Web services to be hosted; and firewall, load balancing, and high-performance clustering requirements. For more information about the Dell Server Advisor, visit www.dell.com/serveradvisor.

In addition, Dell offers a broad range of expert IT infrastructure services that specialize in the planning, implementation, and maintenance of standards-based systems and environments. For more information about the Dell IT services suite, visit www.dell.com/services.

Complementing the ninth-generation server and storage launch is Dell's Platinum Plus Enterprise Support service, which helps ensure cost-effective uptime for business-critical server and storage infrastructures. The new service proactively monitors operations, provides communication logistics, and coordinates end-to-end support the world over. For more information, see "Nothing Basic About Next-Generation Enterprise Support," *Dell Power Solutions,* August 2006, www.dell.com/downloads/global/power/ps3q06-50060413-DellECC.pdf.

Business agility: Responding with flexibility and choice

Dell has evolved its server, storage, workstation, and management platforms to address the need for simplification, increased resource utilization, and cost-effective scale-out. Comprehensive enhancements in system architecture, processor and memory technology, and physical design characterize this new generation of products. Particularly in combination with advances in integrated systems management solutions and virtualization technology, Dell's new servers, storage arrays, and workstations enable organizations to better align the IT infrastructure with vital business processes, meet demanding enterprise requirements quickly and flexibly, and take control of the change-management process.

Mark Nickerson is the Dell/EMC business director for Asia Pacific and Japan. He has been with Dell for eight years as a technical consultant and in enterprise marketing. He attended The Ohio State University.

Joe Pollock is a storage marketing manager in the Dell Enterprise Product Group. He has a B.S. in Electrical Engineering from the University of Florida.

Stori Waugh is a server marketing manager in the Dell Enterprise Product Group. She has more than 10 years of marketing experience in the IT business. Stori has a B.A. in Communications from the University of Houston.

Stacy Hower is a Dell Precision workstations marketing manager. She joined Dell in 2005 with over 12 years of marketing experience. She has a B.A. in Journalism with an emphasis in public relations and advertising from the University of Houston.

FOR MORE INFORMATION

Dell scalable enterprise:
www.dell.com/enterprise

Dell PowerEdge servers:
www.dell.com/poweredge

Dell/EMC CX3 and CX series SAN arrays and Dell PowerVault storage:
www.dell.com/storage

Dell Precision workstations:
www.dell.com/precision

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Boosting Data Transfer with TCP Offload Engine Technology

on Ninth-Generation Dell PowerEdge Servers

TCP/IP Offload Engine (TOE) technology makes its debut in the ninth generation of Dell™ PowerEdge™ servers, and is expected to boost data transfer performance for many key enterprise applications. This article explores the anticipated benefits of TOE technology being developed by Broadcom, Dell, and Microsoft.

BY PANKAJ GUPTA, ALLEN LIGHT, AND IAN HAMEROFF

Related Categories:

Broadcom

Data networking

Dell ninth-generation servers

Internet Protocol (IP)

Microsoft Windows

Network fabric

Network interface card (NIC)

Scalable Networking Pack (SNP)

TCP/IP Offload Engine (TOE)

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

The online economy—particularly for e-business, entertainment, and collaboration—is dramatically increasing the amount of Internet traffic to and from enterprise servers. Most of this data is going through the TCP/IP stack and Ethernet controllers. Additionally, many IT organizations are moving storage, high-performance computing (HPC), and database cluster networks to Ethernet. As a result, Ethernet controllers are experiencing heavy network traffic, which requires increased system resources to process network packets. Because TCP/IP consumes a significant amount of host CPU processing cycles, a heavy TCP/IP load may leave few system resources available for other applications. This situation can create a bottleneck for applications involving significant network traffic or limit the amount of CPU cycles available for actual application processing. Applications that exhibit this behavior include file-oriented storage, block-oriented storage, backups, database transactions, and tightly coupled distributed applications such as HPC.

The TCP/IP Offload Engine (TOE) model is designed to improve data transfer performance over IP networks by relieving much of the overhead when processing TCP/IP from the host CPU. TOE allows the OS to offload all TCP/IP traffic to specialized hardware on the network adapter while leaving TCP/IP control decisions to the host server. By relieving the host processor bottleneck, TOE can help deliver the performance benefits administrators expect from applications running across high-speed network links. TOE is also cost-effective because it processes the TCP/IP stack on a high-speed network device that requires less processing power than a general-purpose high-performance CPU.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

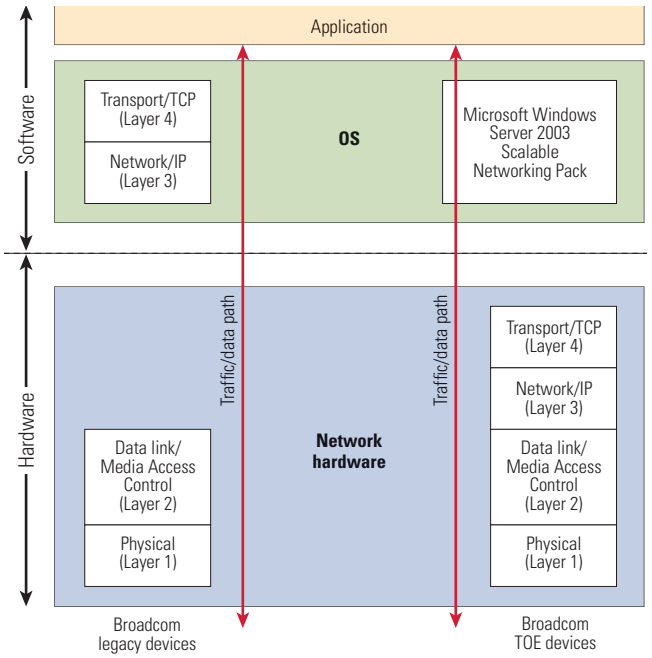

Figure 1. Networking layers implemented in a TOE NIC and a standard Gigabit Ethernet NIC

TOE technology from Dell, Broadcom, and Microsoft

Dell, Broadcom, and Microsoft have developed a TOE implementation that will be available on the Dell PowerEdge server product line. This implementation is based on Broadcom® NetXtreme® II network devices—such as the BCM5708C and BCM5708S 10/100/1,000 Mbps Ethernet controllers, which are integrated on the system board in dual-socket PowerEdge servers and available as optional add-in PCI-e adapters on other platforms—and the Microsoft® Windows Server® 2003 Scalable Networking Pack (SNP). Both include TOE and stateless offload functionality. The Broadcom devices implement Open System Interconnection (OSI) Layers 1, 2, 3, and 4 functionality and in the future may include some Layer 5 functionality. NetXtreme II devices handle the data processing of TCP/IP connections offloaded to them by the OS. The OS handles control and management of TCP/IP processing such as connection setup, connection termination, system resource allocation, prevention of denial-of-service attacks, and error and exception handling. Figure 1 provides a comparison of the various OSI layers implemented in a TOE network interface card (NIC) and a standard Gigabit Ethernet NIC.

To take advantage of the TOE feature in Dell PowerEdge servers equipped with Broadcom NetXtreme II devices, administrators must install the Microsoft Windows Server 2003 SNP. The SNP supports networking hardware enhancements designed to promote scalability and performance improvements on servers running Windows Server 2003. The architectural innovations provided in the SNP include TCP Chimney Offload, which enhances the existing software implementation of TCP/IP version 4 currently in Windows Server 2003.

TCP Chimney Offload creates a direct connection between the top of the protocol stack and the software drivers to enable partial offload of the protocol stack. The term *chimney* was used because the data is transferred to the top of the protocol stack without moving through the intermediate protocol layers, like smoke rising through a chimney to the top of a building without having to travel through the intermediate floors. Figure 2 shows the TCP Chimney Offload architecture.

The SNP also supports Receive Side Scaling (RSS), a network technology that enables the efficient distribution of network receive processing across multiple CPUs in multiprocessor systems.

The NetXtreme II devices are designed to support 1,024 simultaneous offloaded connections. The decision to offload the connection or pass it through the software protocol stack is made by the OS and is seamlessly handled by a software switch, as shown in Figure 2. A connection can be uploaded—that is, moved back from the hardware to the OS—for a variety of reasons. For example, the system automatically switches to the software protocol stack after 1,024 connections are offloaded per Ethernet port. The host or offload target—in this case, the NetXtreme II device—can also request uploading or termination of an offloaded connection. In addition, the Network Driver Interface Specification (NDIS) 5.2 intermediate driver can upload a connection indirectly by requesting the host TCP/IP stack to upload the connection. If the link fails, the event is reported to the host stack through an NDIS Plug and Play (PnP) event and the host determines the most appropriate course of action.

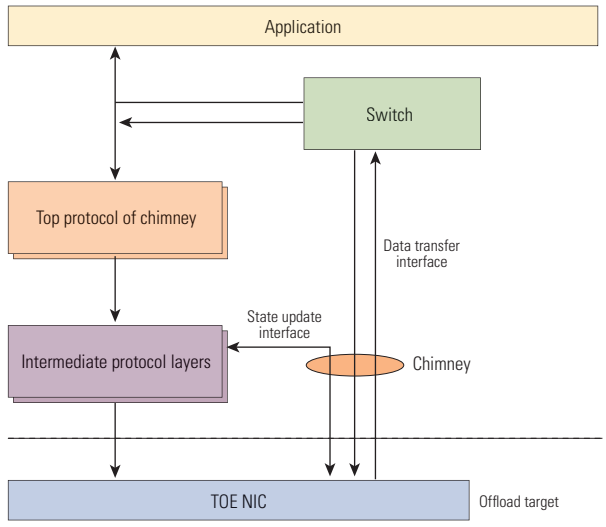

Figure 2. TCP Chimney Offload architecture

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

TOE components and dependencies

The minimum system requirements for TOE functionality are as follows:

- Microsoft Windows Server 2003 with Service Pack 1 (or later)
- Microsoft Windows Server 2003 SNP (KB91222)—factory installed on supported Dell PowerEdge servers for a TOE LAN on Motherboard (LOM) or installed as part of a TOE add-in NIC customer kit
- Broadcom BCM5708C or BCM5708S add-in NIC or integrated LOM
- TOE license key installed prior to boot on LOMs (the TOE license is built into add-in NICs and provided by default on a system with Broadcom BCM5708C or BCM5708S LOMs)
- Broadcom NetXtreme II drivers (NDIS 5.2 drivers)

The NIC operates in non-TOE mode if any of the preceding components are missing or if an OS other than Microsoft Windows Server 2003 with Service Pack 1 (or later) is running.

Dell platforms such as the PowerEdge 1900, PowerEdge 1950, PowerEdge 1955, PowerEdge 2900, and PowerEdge 2950 servers have integrated TOE capability with the LOM ports. TOE add-in NICs are also available for these and other Dell PowerEdge systems.[1] TOE licensing for Dell PowerEdge servers with Broadcom BCM5708 LOMs is automatically enabled via a hardware key plug-in module. The TOE enablement key is factory installed in these PowerEdge servers and add-in NICs. No additional work is required to enable the hardware feature. TOE software components—the Microsoft Windows Server 2003 SNP and Broadcom NetXtreme II drivers—are preinstalled with factory-installed Microsoft Windows Server 2003; otherwise, the software is provided for installation. Administrators should refer to the user's guide included with the network driver package for information about installing the necessary TOE software components.

TOE performance expectations

Network performance improvements gained from TOE technology can be determined by measuring either the increase in absolute network throughput or the reduction in system resources such as CPU utilization. TOE performance benefits vary with the type of applications being run. Applications with a small network packet size may experience gains in network throughput, while applications with a large network packet size may not show significant network throughput improvements with TOE but may experience a significant reduction in CPU utilization—thereby helping to keep CPU processing cycles available for other business-critical applications such as database, backup storage, media streaming, and file server applications. Applications that require extensive network utilization—such as network backups, network attached storage, file servers, and media streaming—typically benefit the most from TOE technology.

Applications with a small network packet size may experience gains in network throughput, while applications with a large network packet size may not show significant network throughput improvements with TOE but may experience a significant reduction in CPU utilization—thereby helping to keep CPU processing cycles available for other business-critical applications.

Configuring the test environment

A team of engineers from Dell, Broadcom, and Microsoft tested the TOE implementation during the product development process from May 2005 to May 2006. The test configuration included a Dell PowerEdge 2900 server equipped with 1 GB of double data rate 2 (DDR2) memory and a dual-core Intel® Xeon® processor at 3.2 GHz with a 1,066 MHz frontside bus. Intel Hyper-Threading Technology was enabled on the processor. Tests were performed with on-board Broadcom BCM5708C and BCM5708C-based add-in NICs.

The test results in this article highlight the advantages of TOE; however, real-world performance will vary based on the specific configuration running in a given enterprise environment. To demonstrate the potential performance benefits of TOE, Dell and its partners ran tests with TOE enabled using the Broadcom TOE NIC and without TOE enabled using an otherwise comparable industry-standard Layer 2 Gigabit Ethernet[2] adapter. The same configuration was used for both the TOE and non-TOE test environments.

Testing the impact of TOE on application performance

Network throughput and CPU utilization are common ways to measure a network adapter's performance. Network throughput is the rate in megabits per second that packets are sent and received, while CPU utilization is the percentage of CPU capacity required to process the packets. The Dell, Microsoft, and Broadcom

[1] The Dell PowerEdge 830, PowerEdge 850, PowerEdge 1800, PowerEdge 1850, PowerEdge 2800, PowerEdge 2850, PowerEdge 6800, and PowerEdge 6850 servers also support TOE add-in NICs.

[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.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

team used benchmarking tools such as IxChariot 6.20 from Ixia[3] and WebBench 5.0[4] to measure network throughput and CPU utilization in the test environment.

The ratio of network throughput to CPU utilization can be expressed as the performance efficiency (PE) index, which was originally developed by *PC Week* (now *eWEEK*) in 1995. The PE index is a commonly used performance metric for evaluating network adapters. For example, high PE indexes indicate high throughput with low CPU utilization, suggesting favorable overall system performance. As Figure 3 indicates, the PE index of the TOE NIC in the test environment described in this article exceeded the PE index for the legacy Layer 2 Gigabit Ethernet NIC for every network I/O size from 1 KB to 64 KB.

TOE-enabled network adapters are designed to achieve high throughput at small network packet sizes and low CPU utilization at large network packet sizes. Figure 4 shows results from the Ixia IxChariot 6.20 performance benchmarking tool using the IxChariot's high-performance script on the test configurations described in this article.

As the number of TOE-enabled ports in a system increases, network throughput can be expected to scale more effectively using TOE-enabled network adapters than standard Layer 2 network adapters. At the same time, the CPU utilization can be expected to remain low enough to allow other processes or applications to run. As shown in Figure 5, when using only four Gigabit Ethernet network controllers that are not TOE-enabled, nearly 60 percent of the PowerEdge 2900 server's CPU was utilized just processing network traffic. But with TOE enabled on the four Gigabit Ethernet network controllers, the CPU utilization was less than 20 percent in the Figure 5 test scenario.

The network scaling effect demonstrated in Figure 5 indicates that TOE-enabled Gigabit Ethernet controllers can help servers handle larger workloads than they can handle using a standard Layer 2 Gigabit Ethernet controller. For example, backup applications can benefit from significant performance improvements with TOE enabled, as shown in Figure 6. In this test scenario, a TOE-enabled network adapter allowed the PowerEdge 2900 server to back up the same amount of data (80 GB per stream) requiring approximately half the CPU utilization than when the PowerEdge 2900 was equipped with a standard Layer 2 network adapter, which means a backup server with a TOE-enabled NIC can support additional backup streams. This is an example of how TOE can enable Dell PowerEdge servers to deliver enhanced application performance.

RSS allows applications such as Web servers to process more requests with a large number of connections. For example, Figure 7 shows test results using the WebBench 5.0 benchmark tool. These

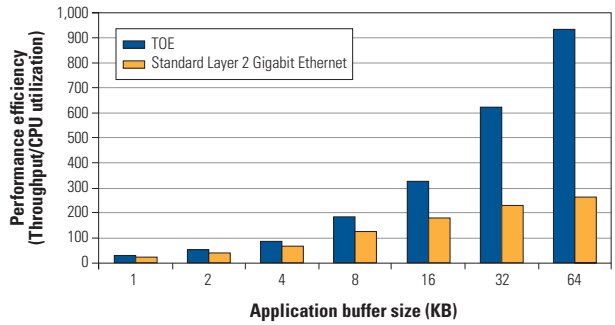

Figure 3. Performance efficiency: TOE-enabled Gigabit Ethernet NIC versus standard Layer 2 Gigabit Ethernet NIC

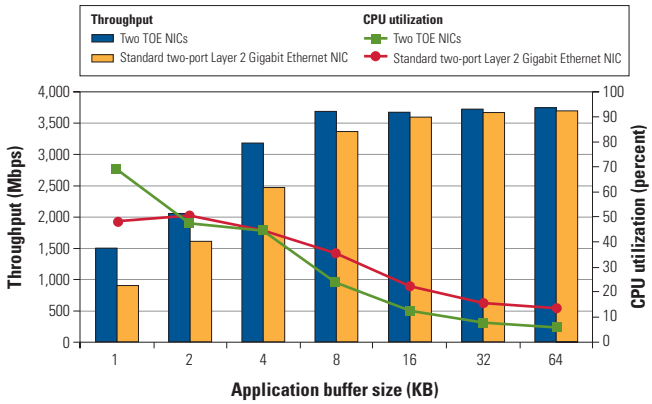

Figure 4. Two-port performance measured by the IxChariot 6.20 benchmark: TOE-enabled Gigabit Ethernet NIC versus standard Layer 2 Gigabit Ethernet NIC

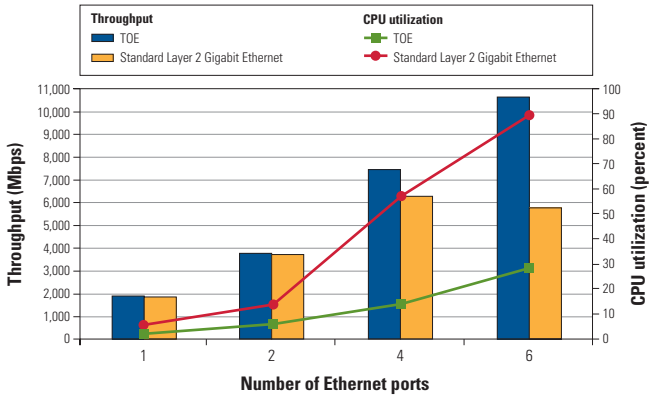

Figure 5. Throughput scaling and the effect on CPU utilization: TOE-enabled Gigabit Ethernet NIC versus standard Layer 2 Gigabit Ethernet NIC (without TOE enabled)

[3] For more information about Ixia IxChariot 6.20, visit www.ixiacom.com/ixchariot.

[4] For more information about WebBench 5.0, visit www.veritest.com/lionbridge.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

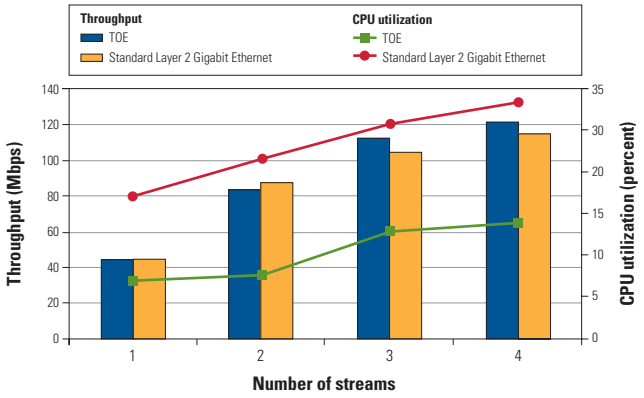

Figure 6. Backup performance: TOE-enabled Gigabit Ethernet NIC versus standard Layer 2 Gigabit Ethernet NIC

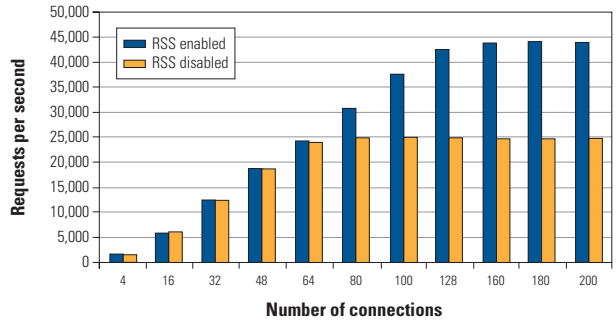

Figure 7. WebBench 5.0 benchmark performance: TOE-disabled Gigabit Ethernet NIC with RSS enabled and standard Layer 2 Gigabit Ethernet NIC with no RSS support

results highlight the performance advantage of enabling RSS on a TOE-disabled adapter compared to a standard Layer 2 Gigabit Ethernet adapter.

TOE can significantly benefit typical Web server environments in which connections are long lived, such as configurations for load-balancing and online shopping. Applications such as ad servers may not show the same benefits of using TOE-enabled NICs because of the overhead for uploading and offloading many short-lived connections. However, the majority of Web applications maintain connections over longer periods of time than the benchmarks used to measure performance in the example scenarios tested in this article.

Increased throughput, reduced CPU utilization for enhanced application performance

Dell PowerEdge servers such as the PowerEdge 1900, PowerEdge 1950, PowerEdge 1955, PowerEdge 2900, and PowerEdge 2950 have integrated TOE capability with the LOM ports. TOE-enabled PCI-e add-in adapters are also available for these and other Dell PowerEdge servers. To enable TOE functionality, servers must run Microsoft Windows Server 2003 with Service Pack 1 (or later) using the Microsoft Windows Server 2003 Scalable Networking Pack and Broadcom NetXtreme II drivers (NDIS 5.2). TOE technology is designed to enhance application performance by increasing network throughput and reducing CPU utilization for applications requiring extensive network activity. Applications such as backup, storage, file servers, and media streaming are likely to benefit most from TOE-enabled systems.

Pankaj Gupta works in the Dell Product Group and led the TOE project. His interests include communication, networking, operating systems, and the development of system software. He has a master's degree in Computer Science from Polytechnic University of New York.

Allen Light is the product line manager for the Broadcom NetXtreme II C-NIC Ethernet controllers at Broadcom Corporation, an Irvine, California–based provider of integrated circuits for wired and wireless communications. He has worked in marketing and application engineering at Intel Corporation and Philips Semiconductors. Allen has a B.S.E.E. from the University of California at Davis, and an M.B.A. from the University of New Mexico.

Ian Hameroff is the senior product manager in the Security and Access Product Marketing Group at Microsoft Corporation, where he is responsible for the product management and marketing of the Microsoft Windows Server platform networking technologies. In this role, Ian drives the go-to-market strategy for Windows Server networking and solutions focused on key security and networking initiatives, such as server and domain isolation, scalable networking, and IPv6 adoption. He has a bachelor's degree in Computer Science with a minor in Mathematics from Plattsburgh State University in New York and is a Certified Information Systems Security Professional (CISSP).

FOR MORE INFORMATION

Senapathi, Sandhya, and Rich Hernandez. "Introduction to TCP Offload Engines." *Dell Power Solutions,* March 2004. www.dell.com/downloads/global/power/1q04-her.pdf

Microsoft Windows Server 2003 Scalable Networking Pack: www.microsoft.com/snp

Broadcom NetXtreme II network devices:
www.broadcom.com/products/Enterprise-Small-Office/Gigabit-Ethernet-Controllers/BCM5708C
www.broadcom.com/products/Enterprise-Small-Office/Gigabit-Ethernet-Controllers/BCM5708S

Network teaming technical brief:
support.dell.com/support/edocs/network/582dx/teaming.htm

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

IPMI Configuration

on Ninth-Generation Dell PowerEdge Servers

Administrators can choose from a number of software options when configuring Intelligent Platform Management Interface settings on ninth-generation Dell™ PowerEdge™ servers. This article discusses these programs as well as the environments and locations in which they can be used.

BY WEIMIN PAN AND HAIHONG ZHUO

Related Categories:

Dell ninth-generation servers

Dell PowerEdge servers

Intelligent Platform Management Interface (IPMI)

Remote management

Systems management

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

The Intelligent Platform Management Interface (IPMI) is an industry standard that defines interfaces to the platform management subsystem. All ninth-generation Dell PowerEdge servers incorporate, in their base configuration, support for systems management capabilities defined by IPMI 2.0. Such functionalities include remotely accessing the platform through standard messaging channels such as IPMI Over Serial, LAN, or Serial Over LAN (SOL) connections, as well as generating and sending a Platform Event Trap (PET) when a monitored event condition occurs.

IPMI configuration on ninth-generation PowerEdge servers can be carried out by different software programs and utilities within the Dell OpenManage™ systems management software suite, with both graphical user interface (GUI) and command-line interface (CLI) options available. Configuration can be divided into the following categories:

- **Configuration of a LAN connection for IPMI messaging and alerting:** Includes enabling IPMI communication over a LAN connection, selecting the source for the remote-access IP address (statically assigned or obtained using the Dynamic Host Configuration Protocol), assigning a static IP address, choosing the highest privilege level allowed on the connection, and configuring virtual LAN (VLAN) settings and the baseboard management controller encryption key
- **Configuration of a serial connection for IPMI messaging and alerting:** Includes setting the connection mode, baud rate, flow control, terminal mode, and highest privilege level allowed on the connection
- **Configuration of a SOL connection:** Includes enabling SOL and configuring the baud rate and settings for how characters should be accumulated and sent
- **Configuration of IPMI remote-access users:** Includes enabling users; assigning usernames, passwords, and privileges on each connection channel (LAN and serial); and enabling SOL payloads on the LAN connection
- **Configuration of Platform Event Filters (PEFs):** Includes enabling PEF alerting on the LAN connection, selecting actions to take, and configuring alert destinations
- **Restoration of IPMI configuration parameters:** Returns these parameters to the preconfigured default values

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

The configuration programs run within either pre-OS or OS-present environments and on either the managed platform or a remote management console. Figure 1 lists these programs and shows the different locations and environments in which they can be used.

IPMI configuration in a pre-OS environment

Administrators can use the Dell Remote Access Configuration Utility or Dell OpenManage Deployment Toolkit to configure IPMI settings in a pre-OS environment on ninth-generation PowerEdge servers. The Dell Remote Access Controller (DRAC) GUI and CLI can also be used in both pre-OS and OS-present environments; these interfaces are discussed in the "IPMI configuration with the Dell Remote Access Controller 5" section in this article.

Dell Remote Access Configuration Utility

The Dell Remote Access Configuration Utility (see Figure 2) is a BIOS setup tool that can be accessed by pressing Ctrl + E when prompted during the system power-on self-test (POST). This utility enables administrators to manage a subset of the IPMI configuration—such as basic IPMI Over LAN parameters, administrator user settings, and the first destination of PETs—assisting administrators in quickly configuring the most commonly used IPMI settings during system startup in a pre-OS environment.

Dell OpenManage Deployment Toolkit

The Dell OpenManage Deployment Toolkit provides a CLI for configuring IPMI settings in a pre-OS environment during system deployment. Administrators can use this CLI to configure nearly all IPMI remote-access parameters supported on ninth-generation PowerEdge servers. For example, to assign a static IP address to the IPMI Over LAN connection, administrators can issue the command `bmccfg lancfgparameters -ipaddrsrc=`*`ipaddress`*.

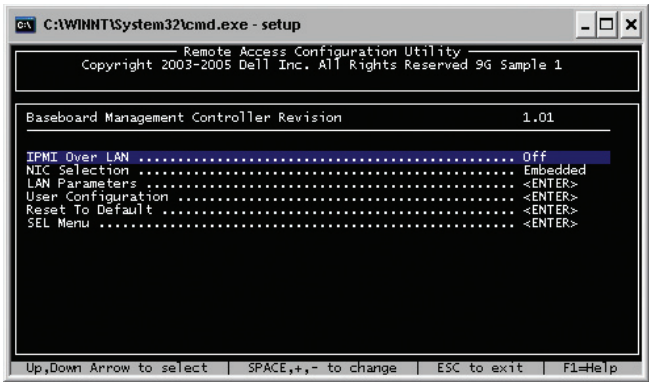

Figure 2. Dell Remote Access Configuration Utility

IPMI configuration in an OS-present environment

Once the system boots to an OS environment, administrators can use either the CLI or the Web-based GUI offered by Dell OpenManage Server Administrator (OMSA) to configure IPMI settings on the system. The GUI and CLI support nearly the same set of operations, including configuration of all IPMI remote-access parameters supported on ninth-generation PowerEdge servers. Although the OMSA CLI program must run locally on the managed server, the OMSA GUI program can run on either the managed server or a remote management console—on the latter, the GUI accesses the managed server through Ethernet if the managed server is booted to an OS environment.

Dell OpenManage Server Administrator GUI

In the OMSA GUI, most IPMI remote-access configuration settings for the managed server can be accessed by selecting System > Main System Chassis > Remote Access on the left control panel, as shown in Figure 3. The Configuration tab displays information and configuration settings for the IPMI Over LAN, IPMI Over Serial, and SOL connections for IPMI communication; the Users tab displays information and configuration settings for IPMI remote-access users.

To access information and configuration settings for PEFs, administrators can click "System" in the left control panel of the OMSA GUI, and then select the Alert Management tab followed by the Platform Events tab.

Dell OpenManage Server Administrator CLI

The OMSA CLI provides various IPMI configuration commands. For example, to assign a static IP address

Software	Program location		Environment on managed server	
	Managed server	Remote management console	Pre-OS	OS present
Dell Remote Access Configuration Utility	✔		✔	
Dell OpenManage Deployment Toolkit	✔		✔	
Dell OpenManage Server Administrator CLI	✔			✔
Dell OpenManage Server Administrator GUI	✔	✔		✔
DRAC GUI	✔	✔	✔	✔
DRAC CLI	✔	✔	✔	✔

Figure 1. Comparison of Dell software that can be used for IPMI configuration

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

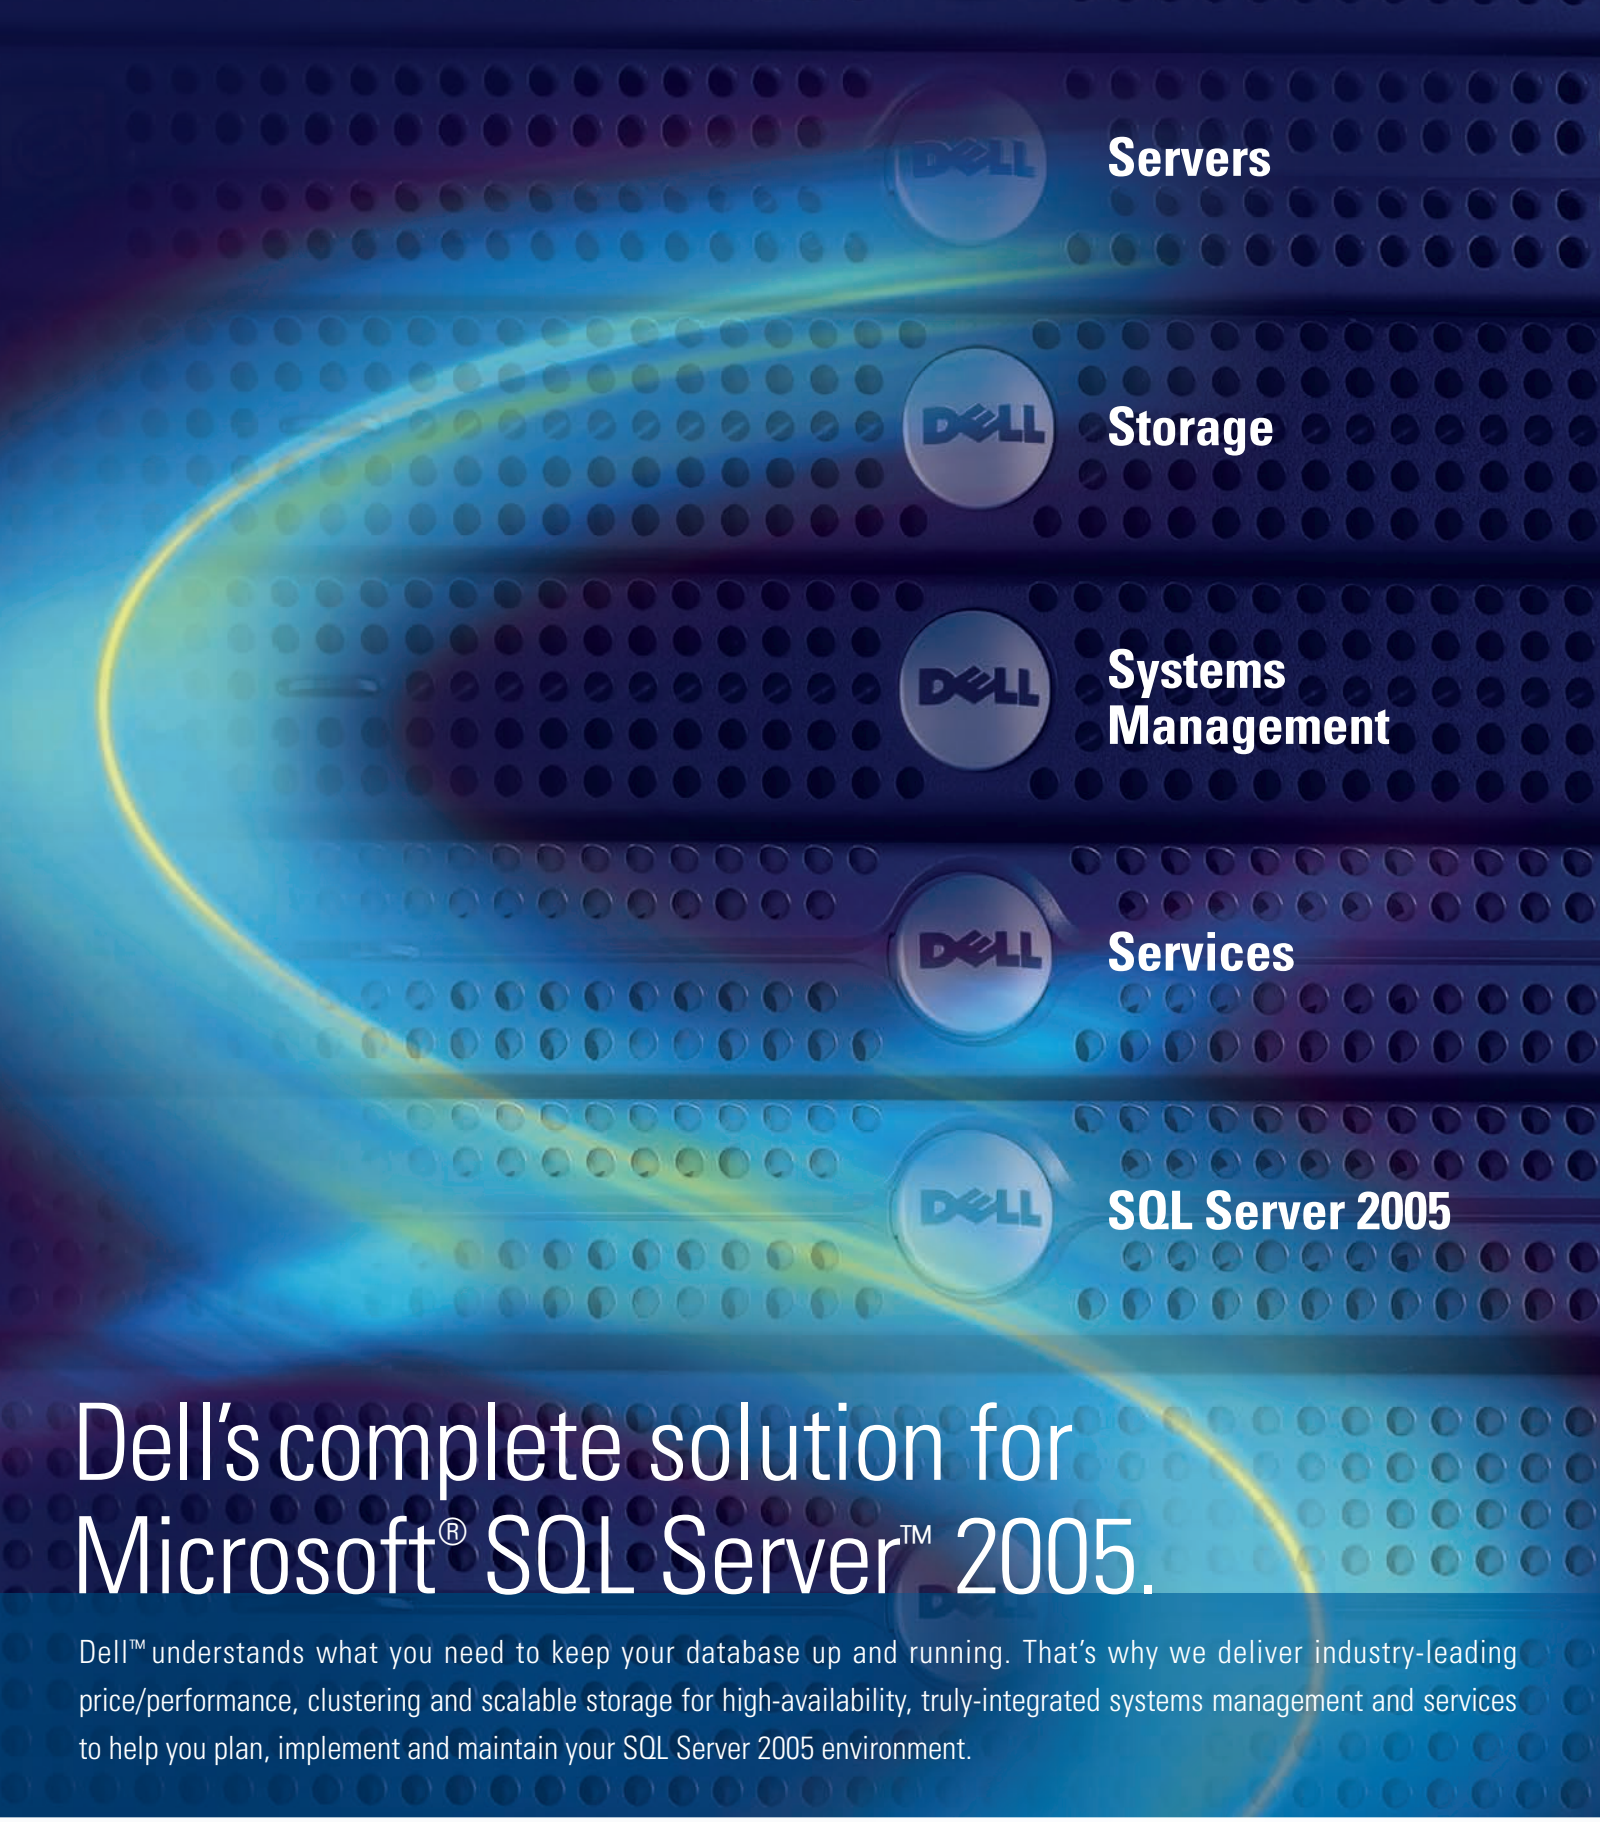

Dell's complete solution for Microsoft® SQL Server™ 2005.

Dell™ understands what you need to keep your database up and running. That's why we deliver industry-leading price/performance, clustering and scalable storage for high-availability, truly-integrated systems management and services to help you plan, implement and maintain your SQL Server 2005 environment.

Visit www.dell.com/sqlmag for more information, whitepapers and "The Definitive Guide to Scaling Out with SQL Server 2005" e-book.

Dell cannot be responsible for errors in typography or photography. Dell and the Dell logo are trademarks of Dell Inc. Microsoft and SQL Server are trademarks or registered trademarks of Microsoft Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others. © 2005 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without the written permission of Dell is strictly forbidden. December 2005.

to the IPMI Over LAN connection, administrators can issue the command `omconfig chassis remoteaccess config=nic ipaddress=`*`ipaddress`*.

IPMI configuration with the Dell Remote Access Controller 5

The DRAC 5 is the fifth generation of Dell Remote Access Controllers and can be installed on ninth-generation Dell PowerEdge servers. It offers enhanced performance and advanced remote-management capabilities, including full IPMI 2.0 remote-access functionality such as IPMI Over LAN, IPMI Over Serial, SOL, and PEFs and PETs.

If a DRAC 5 is installed, administrators can configure the IPMI settings using the Web-based GUI, Telnet or Secure Shell console, or CLI. Both the DRAC GUI and CLI can run on the managed server within an OS; they can also run from a remote management console by accessing the managed server through Ethernet, even when the managed server is not running an OS.

DRAC GUI

In the DRAC GUI, most IPMI remote-access settings for the managed server can be found by selecting System > Remote Access on the left control panel, as shown in Figure 4. The Configuration tab displays information and configuration settings for the IPMI Over LAN, IPMI Over Serial, and SOL connections for IPMI communication, as well as information and configuration settings for IPMI remote-access administrators.

To access information and configuration settings for PEFs, administrators can click "System" in the left control panel and then select the Alert Management tab.

DRAC CLI

The DRAC CLI racadm utility can run locally on the DRAC 5 or remotely from a management console. The racadm command syntax to configure the IPMI settings on the system is `racadm config -g` *`group`* `-o` *`object`*. For example, to enable an IPMI Over LAN connection, administrators should enter `racadm config -g cfgIpmiLan -o cfgIpmiLanEnable 1`. They can type `racadm help` or `racadm help` *`subcommand`* to obtain a list of available subcommands or detailed information about a specific subcommand, respectively.

Flexible options for IPMI configuration

Several server management software options are available for IPMI configuration of Dell PowerEdge servers in both pre-OS and OS-present environments, including GUIs and CLIs that run on either the managed server or a remote management console. Using these interfaces can help provide administrators with flexibility and efficiency in managing ninth-generation Dell PowerEdge servers.

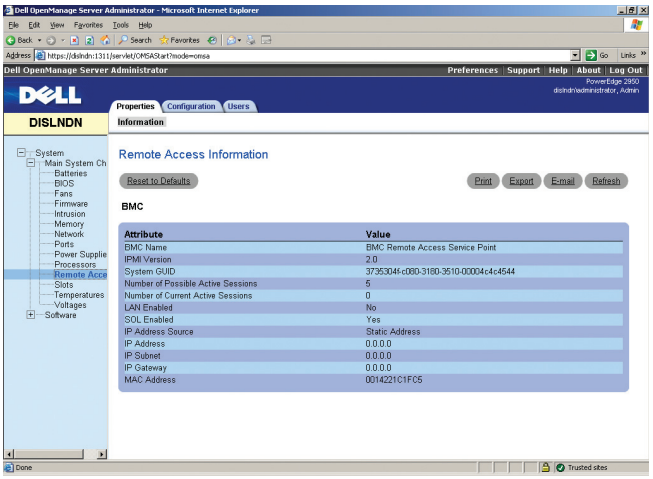

Figure 3. Remote Access Information screen in the Dell OpenManage Server Administrator GUI

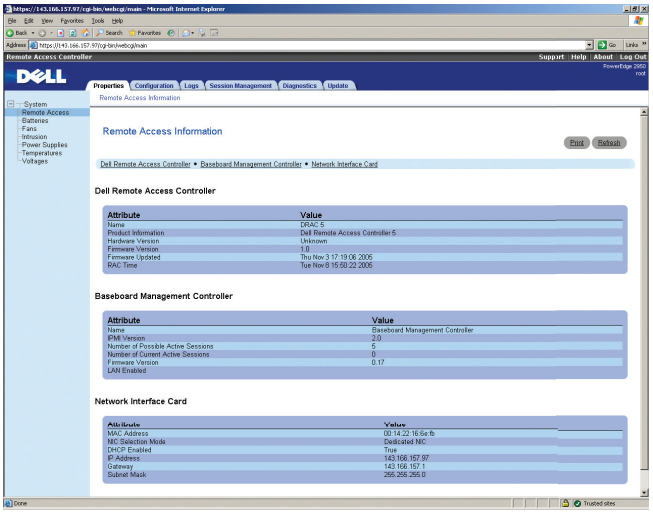

Figure 4. Remote Access Information screen in the DRAC GUI

Weimin Pan is a software engineer strategist in the Dell Remote Management Group. He has worked as a senior systems engineer in the Dell Storage Enclosure Subsystem Group. Weimin has an M.S. in Electrical Engineering from the University of Utah and an M.S. in Computer Engineering from Shanghai Jiao Tong University in China.

Haihong Zhuo is a software engineer senior consultant in the Dell Enterprise Software Development Group. She has worked on systems management console solutions and is currently with the Systems Management Instrumentation team. Haihong has a B.S. in Electrical Engineering from Tsinghua University in China and an M.S. in Computer Engineering from The University of Texas at Austin.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Exploring the DRAC 5:

The Next-Generation Dell Remote Access Controller

Dell™ remote access controllers provide administrators with the necessary tools and functionality to monitor, troubleshoot, and repair servers whether they are around the corner or around the world. This article discusses the features and functionality of the Dell Remote Access Controller 5—available in ninth-generation Dell PowerEdge™ servers—and explores how administrators can help reduce the time required to manage servers, enable fast recovery of remote servers, and lower total cost of ownership.

BY JON McGARY AND WEIMIN PAN

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Dell Remote Access Controller (DRAC)

Remote access controllers (RACs)

Remote management

Systems management

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

Chief among concerns for many IT professionals is managing change—particularly when large numbers of remote servers are involved. Dell remote access controllers (RACs) can help administrators cope efficiently with change, and the RAC available in ninth-generation Dell PowerEdge servers, the Dell Remote Access Controller 5 (DRAC 5), provides enhancements over previous-generation RACs.

The baseline Dell remote access architecture in the Dell OpenManage™ suite consists of hardware and software components that enable administrators to do the following:

- Access a server after a server failure, power outage, or loss of network connection using a network interface card (NIC)
- Remotely view server internal event logs for diagnostic purposes
- Manage servers at multiple locations from a central location
- Manage servers by redirecting the console output (graphics and text) to a remote console
- Perform an orderly shutdown of a server for maintenance tasks
- Diagnose a server failure and restart the server
- Receive alerts through e-mail or Simple Network Management Protocol (SNMP) traps when the server detects an error

Hardware for remote access

The DRAC 5 is an optional hardware controller powered by an AMD Alchemy Au1550 333 MHz processor. It can be integrated into a Dell PowerEdge server as a daughter-card that connects to the system motherboard, and it is compatible with all Dell PowerEdge ninth-generation servers (*x*9*x*0 model numbers).

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

The DRAC 5 interfaces with the baseboard management controller (BMC) chip on the server motherboard. The BMC is based on the Intelligent Platform Management Interface (IPMI) 2.0 standard, which helps Dell provide cost-effective remote management capabilities.

The DRAC 5 uses a flash file system that allows various alert configurations and up to 16 defined local administrators. As an alternative, an administrator may choose to implement the Microsoft® Active Directory® directory service to manage security. The DRAC 5 supports an integrated Web server that allows up to four DRAC administrators to be connected at the same time using a supported Web browser; at any given time, two administrators with redirection privileges may use the console redirection feature and one administrator with virtual media privileges may use the virtual media feature.

Accessing the DRAC 5

Multiple interfaces can be used to access the DRAC 5. These interfaces include the following:

- **Dell OpenManage Server Administrator:** Installed on the managed server, the Server Administrator application provides a comprehensive Web-based graphical user interface (GUI) to configure and launch the DRAC 5 GUI.
- **DRAC 5 GUI:** The DRAC 5 provides a dedicated Web-based GUI to configure the RAC and monitor the server through the DRAC 5 network adapter. Although access to this user interface is provided in Dell OpenManage Server Administrator, it does not need to be installed on the managed server or the management station server. This interface allows stand-alone operation using any supported Web browser.
- **Racadm command-line interface (CLI):** The racadm interface provides a scriptable CLI that enables an administrator to configure the RAC locally or remotely through the DRAC 5 network adapter. This interface requires installation of a small client-side executable, which then interacts directly with the DRAC 5 across the network. No other software is required on the managed server.
- **Dell OpenManage IT Assistant:** IT Assistant can configure and launch the DRAC 5 GUI and operates much like Dell OpenManage Server Administrator, except that it is installed on the management station and displays groups of servers installed on the network.
- **Serial and Telnet console:** Serial and Telnet console commands allow administrators to power up, power down, power cycle, and reset servers; view logs and sensor status; and configure the DRAC 5. The serial and Telnet console supports the `racadm` command, which is useful for scripting one serial client connection and up to four Telnet client connections at one time.
- **Secure Shell (SSH):** SSH provides a fully encrypted username and password LAN interface to the DRAC 5. This access method allows administrators to connect to the DRAC 5 and perform actions on the server, and is similar to the serial and Telnet interfaces. Up to four SSH connections are supported at one time.
- **IPMI:** IPMI out-of-band interfaces, including IPMI Over LAN, IPMI Over Serial, and IPMI Serial Over LAN (SOL), are supported on the DRAC 5. The ipmish BMC management utility and IPMItool open source utility may also be used to manage a system.
- **Systems Management Architecture for Server Hardware Command-Line Protocol (SMASH CLP):** The industry-standard Server Management Working Group SMASH CLP interface defines a CLI to remotely monitor and manage hardware resources.

Returning features from the DRAC 4

Features from the DRAC 4 that remain in the DRAC 5 include the following:

- Ability to configure the RAC network, alerts, administrators, and security settings from the RAC GUI
- Ability to assign RAC administrator groups and permissions
- OS-independent console, keyboard, and mouse redirection functionality
- Virtual floppy and virtual CD media functionality
- Support for Active Directory authentication
- Secure connections using Secure Socket Layer (SSL) and SSH interfaces
- Automatic registration of DRAC name to Domain Name System (DNS)

Features and enhancements introduced in the DRAC 5

The DRAC 5 introduces the following features and enhancements:

- Enhanced GUI
- IPMI out-of-band interfaces
- Enhanced virtual media functionality
- Enhanced virtual KVM (keyboard, video, mouse) functionality
- Enhanced network connectivity
- Enhanced connectivity using the SMASH CLI
- Enhanced security features

Enhanced GUI

The DRAC 5 GUI includes additional usability features compared to previous-generation RACs. The interface offers a combination of tabs and links to logically divide common functionalities between system-level activities and remote management activities. Each

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

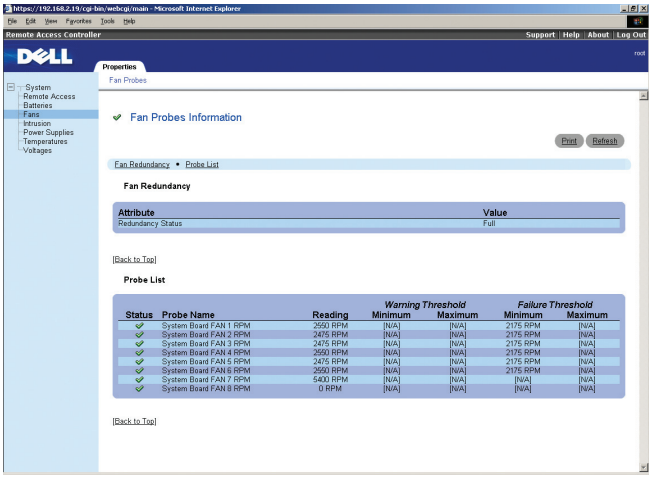

Figure 1. Viewing the status of server hardware components with the DRAC 5 GUI

major tab allows administrators to select various submenus that display status information or an assortment of available actions. One of the many enhancements is the ability to view the status of the server hardware components such as batteries, fans, or power supplies with a single click from the GUI main page (see Figure 1). Another enhancement enables administrators to view the username, connection method, and IP address of each administrator logged in to the DRAC 5—and with the appropriate permission level, they can even disconnect the remote session.

IPMI out-of-band interfaces

In addition to the interfaces used in previous-generation RACs, the DRAC 5 can use the following interfaces to connect to the server:

- **IPMI Over LAN with Remote Management Control Protocol+:** Allows an administrator to remotely manage a system by sending industry-standard IPMI commands to the DRAC 5 over the LAN
- **IPMI Over Serial:** Supports both basic and terminal IPMI modes over a serial connection
- **IPMI SOL:** Provides a mechanism to redirect a text serial console over a LAN session

Enhanced virtual media functionality

The DRAC 5 virtual media feature has been redesigned from an IDE interface to a USB 2.0 interface (see Figure 2). The USB interface supports the functionality of a virtual floppy and virtual CD, and is designed to provide the following benefits:

- Virtual media performance with up to 1.5 MB/sec transfer speeds
- Data storage through remote and local 16 MB USB keys
- Support for ISO images
- Support for persistence boot images
- SSL encryption of data transferred over the LAN, for enhanced security

Enhanced virtual KVM functionality

The DRAC 5 virtual KVM functionality has been enhanced with the following improvements:

- Improved maximum supported screen resolution from 1,024 × 768 to 1,280 × 1,024
- Console redirection that supports a localized keyboard (including English, French, German, and Spanish) and programmable function keys
- Console redirection viewer that supports a native plug-in for both the Linux® and Microsoft Windows® operating systems, removing any requirement for Java to be installed on the client system
- SSL encryption of KVM data transferred over the LAN, for enhanced security

Enhanced network connectivity

The DRAC 5 has the following network enhancements (see Figure 3):

- **Support for virtual LANs (VLANs):** This feature allows a single physical network cable to be partitioned into multiple logical networks, helping administrators enhance network performance by limiting the number of broadcasts to a small group of devices. VLANs also help improve network security by limiting network access to subscribers defined in the VLAN group, even if the devices are on the same physical LAN.
- **Shared NIC:** The server and the DRAC 5 share the same physical network cable, which allows the administrator to

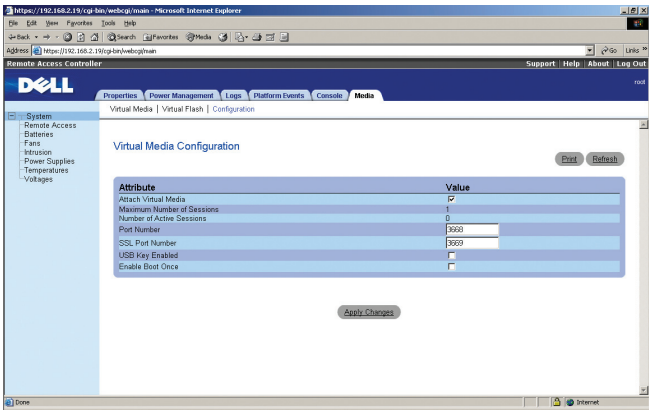

Figure 2. Configuring virtual media with the DRAC 5 GUI

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

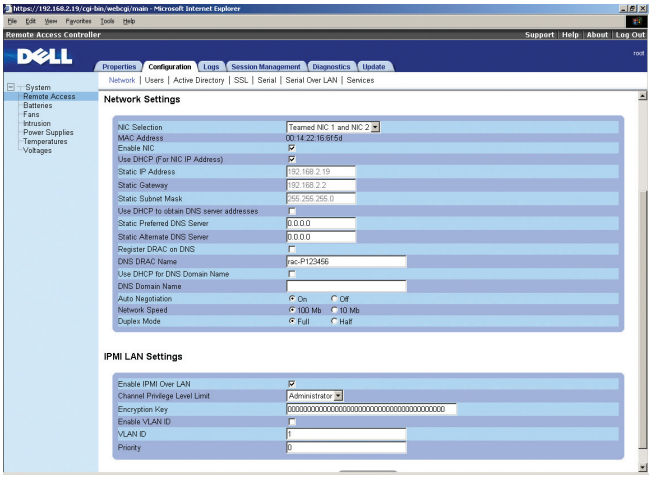

Figure 3. Configuring network settings with the DRAC 5 GUI

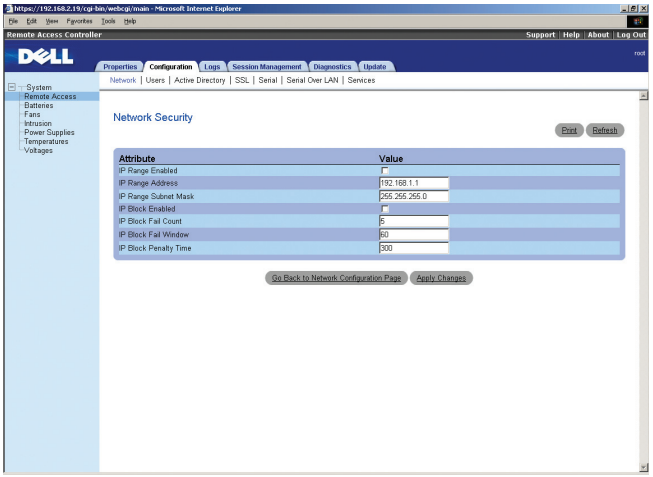

Figure 4. Configuring network security with the DRAC 5 GUI

connect to the DRAC 5 through the system NIC instead of the dedicated NIC built into the DRAC 5. This feature is advantageous because it allows administrators to install only a single network cable connected to each server instead of requiring two cables, one to the server and one to the DRAC 5.

- **Shared NIC with failover:** NIC failover is provided by the DRAC 5 as a way of providing fault tolerance in servers. Fault tolerance is provided by allowing a failed or inactive network link to fail over to the second active link.

Enhanced connectivity using the SMASH CLI

The DRAC 5 supports a Distributed Management Task Force (DMTF) SMASH-compliant CLI, accessible with the Telnet, SSH, and serial interfaces. The DRAC 5 SMASH CLI is designed to help reduce total cost of ownership by providing an industry-standard interface that enables interoperability over a large heterogeneous hardware environment. The DRAC 5 SMASH CLI supports industry-standard commands that allow administrators to do the following:

- View the system event log
- View the server power status
- Power up a server
- Power down a server
- Reset a server

Enhanced security features

Security is an important aspect in any remote management application that is transmitting potentially sensitive data over the Internet. The DRAC 5 enhances security through the following features (see Figure 4):

- SSL encryption of all virtual media data
- SSL encryption of all KVM console redirection data
- 128-bit encryption of all data, which can be turned on or off dynamically by the administrator to help improve performance
- Use of VLAN groups to limit network access to devices subscribed to the VLAN group
- Ability to configure all network port numbers
- Ability to limit remote access control by specifying a limited range of valid IP addresses
- Ability to limit login retries to prevent retry attacks

Centralized remote access for managing global networks

In distributed computing environments, the ability to remotely manage a large number of servers is not simply a helpful option—it is a fundamental requirement. With administrators managing large networks of servers globally through centralized support centers, the need for improved remote management capabilities plays a critical role in the success of an IT organization. The functionality provided by the DRAC 5 to remotely manage, upgrade, troubleshoot, and repair systems is increasingly important to enterprises whose goals include reducing time spent on maintenance activities.

Jon McGary is a senior software developer in the Dell Remote Management Group. Prior to joining Dell, Jon was employed by Tandem Computers and specialized in remote management of fault-tolerant computers. He has a B.S. from Texas A&M University.

Weimin Pan is a senior development engineer in the Dell Remote Management Group. He has worked as a senior systems engineer in the Dell Storage Enclosure Subsystem Group. Weimin has an M.S. in Electrical Engineering from the University of Utah and an M.S. in Computer Engineering from Shanghai Jiao Tong University in China.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Improved Management

Using the LAN and Serial Interfaces in Ninth-Generation Dell PowerEdge Servers

This article discusses the enhanced server management functionalities of ninth-generation Dell™ PowerEdge™ servers, which can be accessed through LAN and serial interfaces to enable flexibility and uniformity in server management.

BY TIM LAMBERT, ROHIT SHARMA, AND HARISH JAYAKUMAR

Related Categories:

Dell ninth-generation servers

Dell PowerEdge servers

Servers

Systems management

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

Effective server management is key to maintaining high availability and reliability in enterprise IT environments. The industry standard for server management is the Intelligent Platform Management Interface (IPMI). The term *intelligent platform management* refers to autonomous monitoring and recovery features implemented directly in platform management hardware and firmware. IPMI functionalities are a key component in providing enterprise-class management for high-availability systems.

This article discusses the two interfaces available in ninth-generation Dell PowerEdge servers for out-of-band (OOB) communication: LAN and serial. It also explores the advantages of each interface as well as scenarios in which one interface is preferable to the other.

LAN interface

Ninth-generation Dell PowerEdge servers offer more server management connectivity mechanisms than previous-generation Dell servers. When using the LAN interface, administrators have three options for connecting to and managing servers: dedicated NIC mode, shared NIC mode, and failover mode.

Dedicated NIC mode

A dedicated network interface is provided via the Dell Remote Access Controller 5 (DRAC 5). This interface allows access using Telnet, a Web-based graphical user interface (GUI), and Secure Shell. Administrators who wish to isolate management traffic from in-band traffic should choose this connectivity path.

Administrators can select the DRAC dedicated network interface card (NIC) option in the baseboard management controller (BMC) BIOS binary option ROM by pressing Ctrl + E during the power-on self-test (POST) and then selecting "Dedicated" in the NIC Selection field (see Figure 1). *Note:* This option is available only if a DRAC 5 is present.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

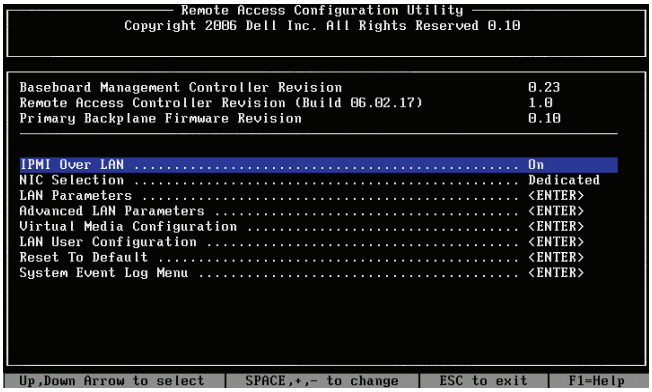

Figure 1. NIC Selection field in the BMC BIOS binary option ROM

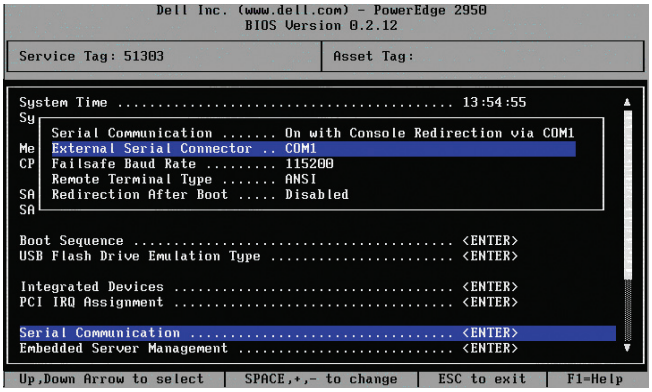

Figure 2. Serial Communication menu options in the BIOS

Shared NIC mode

For base system management, administrators can connect to the BMC through the LAN on Motherboards (LOMs) in shared NIC mode. As an enhancement over previous generations of Dell PowerEdge servers, shared NIC mode enables connection to the BMC through either LOM, not just one.

Shared NIC mode is useful for administrators seeking fewer network router ports and less cabling between the server and network infrastructure compared to dedicated mode. In shared NIC mode, server management traffic is on the same network as in-band traffic to the host OS. The LOMs isolate management traffic by filtering out IPMI packets to the side-band interface for communication with the BMC. Administrators can select the shared NIC mode in the BMC BIOS binary option ROM by selecting "Shared" in the NIC Selection field (see Figure 1).

Shared NIC mode is useful for administrators seeking fewer network router ports and less cabling between the server and network infrastructure compared to dedicated mode.

While in shared NIC mode with a DRAC 5 installed, network connectivity to the DRAC 5 is available through the LOMs. All functions that are available through the DRAC 5 dedicated NIC are also available in shared NIC mode, including Telnet, the Web GUI, remote console redirection, and virtual media.

Failover mode

Dual-LOM systems support failover for in-band traffic in case a problem arises with one of the LOMs or its network connectivity. In failover mode, the same is true for management traffic that goes through the LOMs to the BMC or DRAC 5.

In this mode, both LOMs are teamed using IEEE 802.3ad teaming and are fault tolerant. Administrators seeking high network uptime and throughput should use this option, because it is designed to increase network uptime by protecting the network in the event of a failure or lost connectivity. Administrators can select the failover mode in the BMC BIOS binary option ROM by selecting "Failover" in the NIC Selection field (see Figure 1).

The network connectivity options for managing ninth-generation Dell PowerEdge servers are designed to provide flexibility for administrators. Administrators can isolate management traffic to a separate network using the dedicated NIC interface or save an extra port by merging management traffic with in-band LOM traffic through the shared NIC and failover modes.

Serial interface

Using the serial interface, administrators can communicate with the BMC or DRAC 5 through a direct serial or external modem connection. Ninth-generation Dell PowerEdge servers have one external serial DB-9 connector on the chassis. This connector is shared between the two built-in host COM ports (COM 1 and COM 2) and the BMC and DRAC 5 serial management ports. This sharing is achieved by means of a mechanism called serial port sharing, which uses a multiplexer to switch between different serial ports.

The serial interface uses two modes of communication: IPMI terminal mode and IPMI basic mode. Terminal mode is intended primarily for direct serial connection operations. This mode is designed so that a simple terminal or terminal emulator can be used to generate requests and receive responses from the BMC. It also supports ASCII text commands, although these are limited to a few commands (such as power control) and system reset options. For programmatic access, basic mode supports a binary interface used by utilities such as ipmish (available on the Dell OpenManage™ systems management software CDs shipped with the Dell PowerEdge server).

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

NOT THAT YOU WOULD...

... but you could.

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 just enjoying low tide, your data center can always be within your reach.**

Digital Availability. Easy as DELL™

Dell and the Dell logo are trademarks of Dell Inc. © 2006 Dell, Inc. All rights reserved.

www.dell.com

Administrators can use the system BIOS settings to configure serial port sharing, which controls how serial ports use the external serial connector. To change these settings, administrators can select the Serial Communication menu (see Figure 2) from the BIOS setup screen (accessed by pressing F2 during the POST). In this menu, the Serial Communication and External Serial Connector settings can be used to achieve flexibility and manageability in serial port usage. The following sections examine example scenarios for using these settings (see Figure 3).

Case 1

In this scenario, the external serial port is connected to the system, and the administrator receives the system console output on this port. The administrator can connect a serial cable and use clients such as HyperTerminal and minicom to manage the system through the serial port.

Case 2

This scenario is similar to case 1: The administrator receives the system console output on this serial port, and can connect a serial cable and use clients such as HyperTerminal and minicom to manage the system through this port. The added advantage is that the administrator can dynamically switch and communicate with the system BMC, view the system health status, and so on. Because the BMC is snooping all the traffic, when the BMC is in terminal mode, the administrator can switch to communicate with the BMC using the escape sequence Esc + Shift + 9. This mode gives administrators the advantage of managing both the host system and the server BMC, because they can now dynamically switch between the two.

Case 3

This scenario is nearly identical to case 2, except that the administrator is initially connected to the BMC. To switch to the host console, the administrator must enter the escape sequence Esc + Q.

The behavior in cases 2 and 3 are slightly different if a DRAC 5 is present on the system. When the system is connected to the BMC and DRAC 5, depending on the DRAC 5 serial console attribute setting, the administrator has the option of using the BMC console

	External Serial Connector setting	Serial Communication setting
Case 1	COM1	On with Console Redirection via COM1
Case 2	COM2	On with Console Redirection via COM2
Case 3	Remote Access	On with Console Redirection via COM2
Case 4	Remote Access	N/A

Figure 3. Example settings for the Serial Communication menu in the BIOS

DRAC 5 serial console	User interface presented
Disabled	IPMI terminal or basic mode
Enabled	DRAC CLI

Figure 4. User interface options when a DRAC 5 is present

(in IPMI terminal or basic mode) or the DRAC 5 serial console. The user interface options are listed in Figure 4.

When using the case 1 or 2 settings, the administrator can also gain access to the system console through a Serial Over LAN (SOL) session. The SOL session preempts any existing serial session (that is, control passes to the SOL session). When the SOL session ends, control returns to the serial session.

Case 4

In this scenario, administrators can manage the system using a command-line interface (CLI) utility called ipmish. This utility is a CLI console application and has no GUI; its commands and options are specified using command-line arguments. It supports OOB access (over a LAN or through the serial port) to a single system, and allows administrators to perform simple server management tasks such as viewing the current power status, viewing the system event log, and performing power control.

For ipmish to work, the BMC must be in basic mode, which enables the utility to programmatically access the BMC. If a DRAC 5 is present in the system, the DRAC 5 serial console option should be disabled using the OOB GUI or the racadm utility.

Flexible server management

Ninth-generation Dell PowerEdge servers offer administrators various interfaces and modes, each with its own advantages. Using these tools along with the Serial Communication settings can help provide flexibility and uniformity in server management.

Tim Lambert is a development engineer consultant in the Dell Server Development Group concentrating on server designs and systems management. He has a B.S. in Electrical Engineering from The University of Texas at Austin and an M.B.A. from Texas State University.

Rohit Sharma is a test engineer in the Dell OpenManage Software Development Group. He has a B.S. in Computer Engineering from the University of Mumbai in India and an M.S. in Computer Science from North Carolina State University.

Harish Jayakumar is a development engineer in the Platform Embedded Software team. He has a B.S. in Computer Science from the University of Madras in India and an M.S. in Computer Science from Arizona State University.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Configuring Ninth-Generation Dell PowerEdge Servers

for High-Performance Computing Environments

Dell engineers tested performance settings for ninth-generation Dell™ PowerEdge™ servers using representative computation- and memory-intensive applications. Benchmark performance under various BIOS settings and memory configurations was also measured.

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

Related Categories:

Benchmarks

Dell ninth-generation servers

Dell PowerEdge servers

High-performance computing (HPC)

Parallel systems

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

Ninth-generation Dell PowerEdge servers feature the Intel® 5000X or 5000P chipsets, each of which has a fast dual frontside bus (FSB) to enable increased CPU and memory bandwidth, quad-channel double data rate 2 (DDR2) memory in fully buffered dual in-line memory modules (DIMMs), and PCI Express I/O architecture. A team of Dell engineers in June 2006 used high-performance computing (HPC) benchmarks—Linpack; LS-DYNA; STREAM; NAS (NASA Advanced Supercomputing) Parallel Benchmarks (NPB); MIMD (multiple instruction, multiple data) Lattice Computation (MILC); FLUENT; and OOCORE—comparing various settings to help determine optimal configurations for HPC environments.

System performance is sensitive to several factors that depend on an application's behavior and how it interacts with the underlying hardware. Other than compiler and software optimizations, the hardware configuration (including memory, CPU speed, and disk speed) and BIOS-modifiable performance flags have an impact on application performance. This article focuses on two of these aspects: BIOS settings and memory configuration. All the results described in this article were achieved on a single server; the effect of these settings may vary in an HPC cluster. Furthermore, the results may also vary based on application data sets and parameters. Figure 1 describes the hardware and software environment that the Dell team used for testing performed on the Intel 5000X–based ninth-generation Dell PowerEdge 1950 server.

Benchmarks for evaluating HPC performance

The following synthetic and application-centric 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)[1]

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

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Chipset	Intel 5000X
CPU	*BIOS tests:* Intel Xeon 5060 processor at 3.2 GHz with 2×2 MB L2 cache and 1,066 MHz FSB *Memory tests:* Intel Xeon 5080 processor at 3.73 GHz with 2×2 MB L2 cache and 1,066 MHz FSB
Memory	PC2-5300 fully buffered DIMMs from various vendors
Disk	Serial Attached SCSI (SAS)
OS	Red Hat® Enterprise Linux 4 Update 2 with support for Intel Extended Memory 64 Technology (EM64T)
Compilers	Intel C, C++, and Fortran Compilers version 9.0, build 20050914 (used where source code available)
Message Passing Interface (MPI)	MPICH version 1.2.6
FLUENT	Version 6.2.16
LS-DYNA	Version mpp970s and revision 6763.205

Figure 1. System configuration for BIOS and memory tests on a ninth-generation Dell PowerEdge 1950 server

implementation is commonly used to rank supercomputers on the TOP500 Supercomputer Sites list.

LS-DYNA. This general-purpose, transient dynamic finite-element program can simulate complex real-world problems.[2] LS-DYNA is optimized for shared and distributed memory on UNIX®, Linux®, and Microsoft® Windows® platforms. The specific LS-DYNA workload used in the Dell study was neon_refined.[3]

STREAM. The STREAM benchmark is a synthetic benchmark that measures sustainable memory bandwidth.[4] The Dell test team used the average of the results from the COPY, SCALE, ADD, and TRIAD kernels.

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.[5] The NPB suite consists of a set of eight programs, which are derived from computational fluid dynamics (CFD) code. 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.

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. The workload used for this application was partially based on the MILC Collaboration's public lattice gauge theory code.[6]

FLUENT. A popular CFD application suite, FLUENT is used commonly in HPC environments. The FLUENT applications allow users to perform CFD analysis around their particular models. Several benchmark data sets (workloads) available from Fluent Inc. were used in the Dell tests.[7] To represent FLUENT benchmark performance, the Dell team took the geometric mean of the results.

OOCORE. An out-of-core matrix solver, OOCORE handles matrix equations that are too large for the cache.[8] This benchmark writes large amounts of data to the disk and thus also tests the disk I/O performance of the server.

Performance tests using various BIOS settings

The first part of the Dell study examined the effects that various BIOS settings have on HPC application performance. As part of this study, the test team ran benchmark tests on a Dell PowerEdge 1950 server equipped with the Intel 5000X chipset, two dual-core Intel Xeon® 5060 processors at 3.2 GHz, and eight 1 GB PC2-5300 fully buffered DIMMs. The BIOS settings in this study controlled Intel Hyper-Threading Technology (HT Technology), snoop filter, and processor prefetch capabilities.

Intel Hyper-Threading Technology. Available in Intel CPUs, HT Technology provides thread-level parallelism on each processor to enable efficient use of processor resources, high processing throughput, and enhanced performance for multi-threaded software.

Snoop filter. The Intel 5000X chipset uses a snoop filter to reduce the number of snoop cycles (which are generated to maintain cache coherency) on the processor FSB. Also known as a cache coherency filter, the snoop filter is a cache structure that exists on the chipset to help reduce unnecessary snoop traffic.

Prefetching. Fetching instructions or data from the memory well before the processor needs it is defined as prefetch. The prefetched instructions or data may simply be the next items in the program fetched while the current instruction is being executed. Prefetching may also be part of a complex speculative prediction algorithm, in which the processor tries to anticipate and fetch the appropriate instructions or data in advance.

In this study, the Dell team used four permutations of these three BIOS settings. Configuration A represents the baseline setting; the other three configurations are derived by starting from configuration A

[2] For more information about LS-DYNA, visit www.lstc.com.

[3] For more information about the neon_refined workload, visit www.topcrunch.org/benchmark_details.sfe?query=3&id=60.

[4] For more information about STREAM, visit www.cs.virginia.edu/stream/ref.html.

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

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

[7] For more information about these benchmark data sets, visit www.fluent.com/software/fluent/fl5bench/flbench_6.2.

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

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

and then changing one setting at a time. The configurations used in the study were as follows:

- **Configuration A:** Baseline setting—HT Technology off, prefetch on, and snoop filter off
- **Configuration B:** HT Technology enabled (prefetch on and snoop filter off)
- **Configuration C:** Snoop filter enabled (HT Technology off and prefetch on)
- **Configuration D:** Prefetch disabled (HT Technology off and snoop filter off)

The results of comparing configurations B through D against the baseline setting (configuration A) are shown in Figure 2. For example, the first set of bar charts shows the relative performance of each of the seven benchmarks when run under configuration B and compared against configuration A.

Evaluating the impact of Intel Hyper-Threading Technology

The first set of bars in Figure 2 show the effect of enabling HT Technology. NPB and OOCORE experienced a slight degradation in performance when HT Technology was enabled, whereas the other benchmarks showed little or no difference. Disabling HT Technology is recommended unless a significant performance benefit can be obtained from it.

Determining the effect of the snoop filter

The second set of bars in Figure 2 show the effect of enabling the snoop filter. The STREAM benchmark experienced more than 30 percent improvement with the snoop filter enabled. Similarly, the performance of NPB and MILC also improved. However, Linpack, LS-DYNA, and OOCORE experienced a slight degradation in performance. The snoop filter can cause some degradation in performance depending on the workload, so testing an application using both snoop filter settings is recommended to determine the optimal setting.

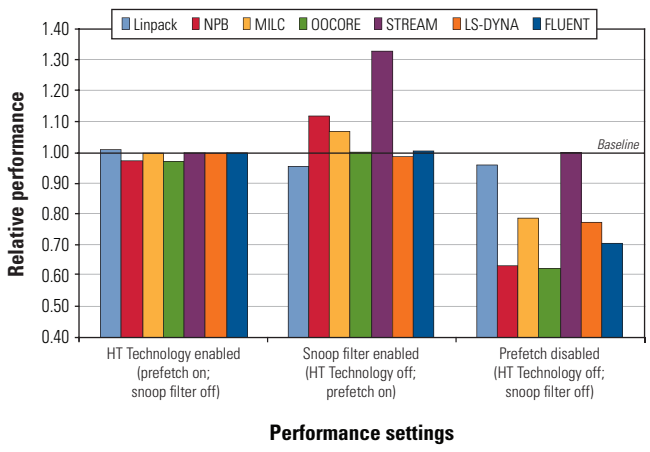

Figure 2. Relative performance of HPC applications for various BIOS configurations

Examining the benefits of prefetching

The last set of bars in Figure 2 show the effect of disabling prefetch settings in the BIOS. Except for STREAM, which was not affected, all the benchmarks showed significant degradation in performance when prefetch was disabled. Thus, enabling the processor prefetch settings is recommended for optimal performance of compute-intensive applications.

Performance tests using different memory configurations

The second part of the Dell study focused on comparing various memory configurations. Just as in the BIOS tests, the Dell test team used the seven benchmarks—Linpack, NPB, MILC, OOCORE, STREAM, LS-DYNA, and FLUENT. As part of this study, the Dell team ran benchmark tests on a PowerEdge 1950 server equipped with the Intel 5000X chipset and two dual-core Intel Xeon 5080 processors at 3.73 GHz. The team ran the tests on three memory sizes, each of which were configured with varying types and numbers of DIMMs:

- **2 GB configuration:** Two 1 GB DIMMs and four 512 MB DIMMs
- **4 GB configuration:** Two 2 GB DIMMs, four 1 GB DIMMs, and eight 512 MB DIMMs
- **8 GB configuration:** Four 2 GB DIMMs and eight 1 GB DIMMs

In addition to running each benchmark on the different memory configurations, the Dell team measured the average performance for all benchmarks on each configuration, as shown in Figures 3 and 4. The results represented in Figures 3 and 4 were obtained with the snoop filter enabled and disabled, respectively. For both figures, the baseline for the results is the system performance when using dual 1 GB DIMMs.

For the 2 GB configuration, performance improved on average by 30 percent when using four 512 MB DIMMs compared to two 1 GB DIMMs. Similarly for the 4 GB configuration, the four 1 GB DIMMs outperformed the two 2 GB DIMMs. This demonstrates the performance advantage of populating at least four DIMMs in the system so that each memory channel has at least one memory DIMM.

The results for four 1 GB DIMMs, eight 512 MB DIMMs, four 2 GB DIMMs, and eight 1 GB DIMMs show that configurations with one DIMM per channel performed similarly to configurations with two DIMMs per channel, especially when the snoop filter was disabled. Another interesting observation is that, with the snoop filter enabled, a higher delta in performance occurred between the baseline configuration—two 1 GB DIMMs—and the

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

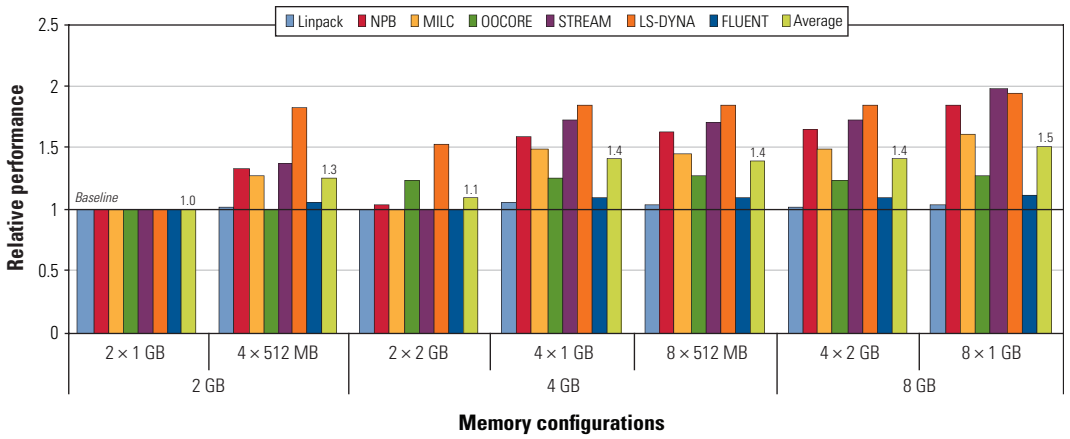

Figure 3. Memory configuration results with snoop filter enabled

rest of the configurations. With the snoop filter disabled, the average performance delta was around 30 percent; however, with the snoop filter enabled, the delta in performance was 40 percent. Furthermore, with the snoop filter enabled, there was approximately a 10 percent difference in performance between the two 8 GB configurations: the eight 1 GB DIMMs outperformed the four 2 GB DIMMs. Therefore, if an application benefits from using the snoop filter, it may show further improvements with fully populated DIMMs.

Optimal settings for ninth-generation Dell servers in HPC environments

The Dell study evaluated the performance effects of multiple BIOS settings and memory configurations for HPC benchmark applications running on a ninth-generation Dell PowerEdge 1950 server. As demonstrated in this study, using the appropriate BIOS settings and memory configurations can help achieve optimal server performance in HPC environments. Similarly, at least four DIMM slots should be populated on ninth-generation Dell servers to take advantage of the four memory channels offered on the Intel 5000X and 5000P chipsets. Heeding these guidelines can enable enterprise IT organizations to take advantage of the enhanced performance capabilities of the new-generation Dell PowerEdge servers.

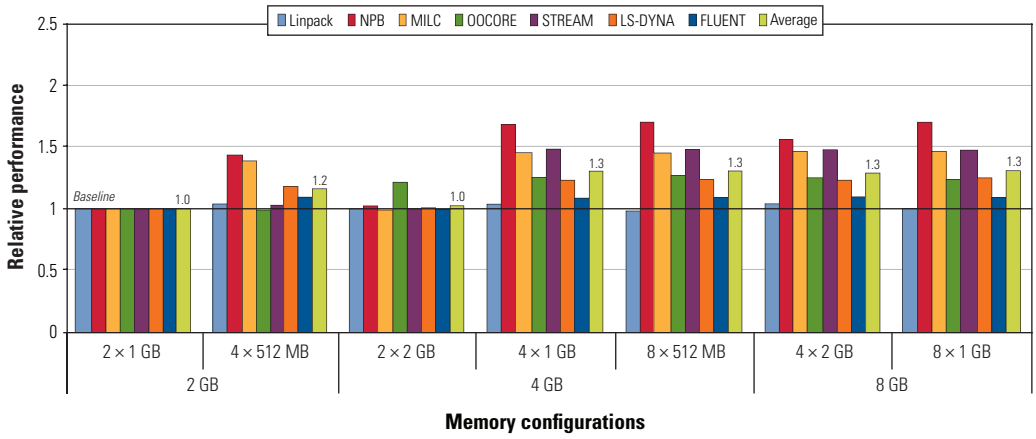

Figure 4. Memory configuration results with snoop filter disabled

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. 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 systems engineer and adviser in the Scalable Systems Group at Dell. He has a bachelor of technology 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.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Array-based Data Protection:

Making Data Recovery Fast, Easy, and Cost-Effective

Disk-based data protection is increasingly essential as a supplement to tape-based protection to meet rapid recovery requirements for business-critical enterprise applications. Among the various disk-based solutions, array-based data protection offers distinct benefits that take advantage of the latest technologies to provide fast, affordable recovery for any size organization.

BY SONYA R. SEXTON

Related Categories:

Business continuity

Dell/EMC storage

Disaster recovery

Replication

Storage

Storage software

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

Across enterprises of all types and sizes, not only are online operations generating more data, but more of that data is linked to vital applications that cannot fail without incurring serious consequences, including loss of revenue, decreased productivity, a damaged reputation, and even legal penalties. For a growing number of applications, the service-level requirements for data recovery time are now measured in minutes or hours, not days. Because traditional tape-based recovery can be time-consuming, enterprises must deploy disk-based protection technologies extensively to meet increasing requirements for rapid recovery.

Heightened demand for rapid recovery is driving innovation in disk-based data protection technology. The evolution of external disk storage devices to support low-cost ATA and Serial ATA (SATA) media is a key development that has helped reduce the cost of implementing disk-based solutions. Additional developments in disk-based data protection include enhancements to user environments, resulting in friendlier, more efficient interfaces, and application programming interface (API)–level integration with complementary programs to help increase operational effectiveness.

Types of data protection

Data protection functionality consists of two broad types: remote replication to a second array, whether in a second data center at the same location or at a distant site, and point-in-time copying within a single system. These functions can be used individually or in combination. Remote replication, often called remote mirroring, replicates file or block changes to application data to one or more remote nodes in real time. Point-in-time copying (see Figure 1) creates copies of volumes at a particular point in time; these copies are often referred to as snapshots or clones. A snapshot is a virtual copy comprising pointers to the base volume and an index of blocks that have changed in the original volume since the snapshot creation; as new data is written to the source volume, old data is copied to the snapshot index. A clone is a complete block-for-block copy of a volume at a particular point in time.

The range of data protection technologies is best surveyed based on where the software services run in the I/O path. Most implementations run on either host platforms or the controllers of external storage arrays. A third implementation is also emerging—one that runs in the network, either on a dedicated appliance or in the network switch.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

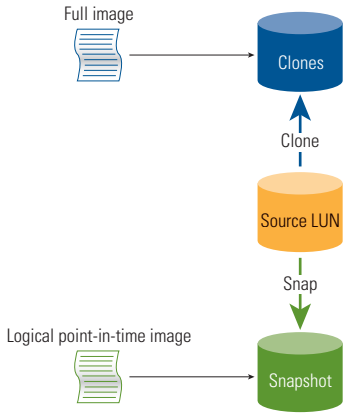

Figure 1. Point-in-time copying using clones and snapshots

Host-based technologies tend to be the most cost-effective of available data protection technologies, but they are point solutions, providing one type of data protection service (either remote mirroring or local point-in-time copying) for a specific application type. Also, because they must be deployed for individual servers and are often modified to follow OS or application revisions, management overhead can be high in environments with multiple servers. These solutions compete with production applications for processing resources on both the primary and target nodes.

Network-based technologies have had limited deployments. Available as either an appliance or a blade option with storage area network (SAN) switches, they are deployed directly in the I/O path between the hosts and the external storage devices or as a node on the network. They are intended to operate between disparate disk systems. Note that supported systems each require discrete interoperability engineering.

Array-based technologies are gaining wider acceptance, particularly with organizations following consolidation strategies to effectively manage growth. These data protection implementations allow organizations to use the same solution for their entire environment and provide a wide range of mature rapid recovery mechanisms, including remote replication, local point-in-time copying, and combinations of the two. Cost-effectiveness can be enhanced through one-time, device-based licensing (as opposed to capacity-based licensing) combined with tiering support, both within arrays (by means of drive intermix) and across arrays (by means of interoperability between the full range of storage platforms, from entry-level to enterprise). Most leading storage arrays provide a basic set of data protection options; a smaller subset provide premium options. Given the increasing importance of rapid recovery, premium options can be a key factor when evaluating storage consolidation platforms. For more information about Dell/EMC arrays and data protection software, see the "Dell/EMC CX3 series and array-based replication" sidebar in this article.

Beyond software functionality, hardware differences also have a direct impact on the effectiveness of a given solution. Among the most important distinctions is the ability of the hardware architecture to sustain demanding performance levels for production applications in tandem with extensive deployment of data protection services. Choice products offer powerful performance architectures designed for consolidation with challenging requirements for rapid recovery. These products generally support more active copy and replication sessions than less-advanced platforms.

Array-based data protection: Exploring the range of functionality

Array-based data protection provides the most extensive functionality of any of the available data protection technologies. This expanded functionality allows administrators to tailor recovery service levels for diverse applications. Choice products complement extensive functionality with elegant tools for administering the replication services to enable both efficiency and operational excellence. Choice products also provide flexibility in configuring disk resources for use as replication repositories—for example, the ability to configure global repositories using low-cost disks within the array, which can centrally support all data protection operations with efficient use of disk resources.

Remote mirroring

Replication between a source array and one or more remote arrays is generally supported in two operational modes: synchronous and asynchronous. In synchronous mode, both the source and target arrays must acknowledge a write operation before the next write can occur. This mode provides recovery at the exact point of failure, with near-zero data loss. In asynchronous mode, writes can continue on the source volume while the target volumes are updated in the background.

The interconnect channels between the source and target arrays are a key part of any remote mirroring implementation. Standard Fibre Channel links are practical for distances of up to 10 km (about 6 miles) for most applications; beyond that range, channel latency and interference can become problems. Use of optical line drivers and repeaters or dense wavelength division multiplexing (DWDM) can extend Fibre Channel range up to 100 km (about 60 miles). For longer distances, remote communication services must be provided through an internal IP option in the array or an external protocol conversion of Fibre Channel to IP. These conversion products are readily available from most switch providers.

In synchronous mode, the application is dependent on the bandwidth of the link connecting the source and target. The connection bandwidth should be planned to exceed the maximum peak of the application I/O load. Because high bandwidth over long distances can be expensive, synchronous mode is often used with targets at campus- or municipal-area ranges for quick recovery from room, building, or site disasters. Many enterprises

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

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

Dell is a trademark of Dell Inc.

EMC²: EMC, Navisphere, DiskXtender, EmailXtender and where information lives are registered trademarks of EMC Corporation. Centera, SAN Copy and SnapView are trademarks of EMC Corporation. All other trademarks used herein are the property of their respective owners. © 2006 EMC Corporation.

© 2006 Dell Inc. All rights reserved.

consider synchronous mirroring mandatory for critical applications that have low tolerance for downtime and near-zero tolerance for data loss.

In asynchronous mode, applications can be less dependent on link bandwidth than in synchronous mode, so the connection bandwidth can be planned to accommodate the average application I/O load, as opposed to the peak load. As a result, asynchronous mode is more cost-effective than synchronous mode for providing rapid recovery from regional or geographic disasters, particularly if the application has some tolerance for data loss.

Both synchronous and asynchronous mirroring modes are supported on most arrays; beyond operational modes, however, array-based products become markedly different. These differences include the range of mirroring functions supported as well as the types of administrative aids that facilitate operational effectiveness and efficiency.

Premium remote mirroring features

Choice arrays provide advanced levels of functionality and manageability that enable high levels of disaster tolerance, cost-effective mechanisms for rapid recovery, and unique applications that enhance the business value of data and improve storage administration. The following premium mirroring features can add significant value to array platforms.

Integration with point-in-time copying. Best practices for rapid recovery call for at least two versions of production data at a recovery site: the active mirror and a full-volume point-in-time copy (see Figure 2). The ability to combine mirroring and point-in-time copying helps provide data recoverability for the broadest range of contingencies. This feature also provides application benefits by enabling a secondary server to mount a fully consistent readable and writable copy of the mirrored data and then perform operations such as decision analysis, application testing, or backup.

Suspend and resume capabilities with fast and reverse resynchronization. Giving administrators the ability to control mirroring operations and quickly establish a desired state of consistency at both the source and target brings flexibility that benefits both

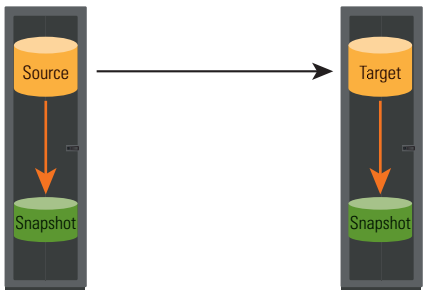

Figure 2. Integrating remote mirroring and local point-in-time copying

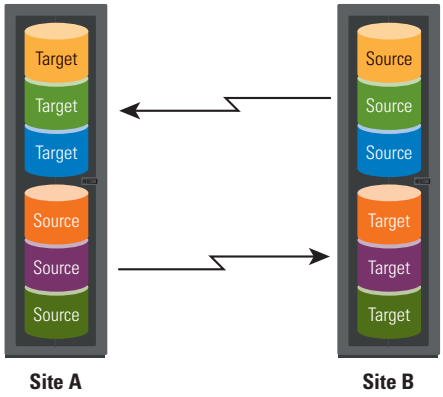

Figure 3. Bidirectional mirroring

recovery and operational effectiveness. Advanced products also provide a throttling mechanism that allows administrators to control the performance impact of resynchronization.

Session persistence. This feature provides rapid recovery of the replication sessions themselves. It uses a write intent log on the source array to facilitate fast resynchronization of the target in the event of a temporary session disruption. In the event of a line disruption, less-advanced products require administrators to reestablish replication sessions from the beginning.

Fan-in and fan-out mirroring. Fan-in mirroring (sometimes called many-to-one mirroring) replicates data from multiple sources to a single target, enhancing the economics of disaster recovery by allowing multiple sites to use a single array—potentially a low-cost platform—as a copy repository. When used in combination with suspend and resume functions, fan-in mirroring helps enable consolidated backups, simplified failover, and consolidated remote processing activities, among other applications. In contrast, fan-out mirroring (sometimes called one-to-many or *n*-way mirroring) replicates data to multiple targets, enhancing disaster tolerance by allowing multiple recovery copies and facilitating nondisruptive data migration.

Bidirectional mirroring. Bidirectional mirroring (see Figure 3) allows arrays to function as both sources and targets, helping reduce the cost of disaster recovery in distributed environments and providing increased flexibility in enabling rapid recovery.

Discrete drive types. This feature gives administrators flexibility when optimizing disk resources at the target (for example, mirroring from RAID-1 Fibre Channel drives to RAID-5 SATA drives). Additionally, it allows administrators to respond nimbly to changing service-level objectives with data migration projects (for example, reallocating an application to a secure site and from RAID-5 SATA drives to RAID-1 Fibre Channel drives).

Consistency groups. During a rolling disaster, data can become inconsistent at the target. Applications such as databases that require a certain sequence of write operations to multiple volumes, or *dependent writes,* are particularly vulnerable. If consistency is not

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

DELL/EMC CX3 SERIES AND ARRAY-BASED REPLICATION

The Dell/EMC CX3 UltraScale™ series of storage systems is designed to support demanding production I/O workloads in combination with the full breadth of array-based replication services. This series features architectural enhancements that are designed to significantly enhance replication performance—including up to two and a half times more processing power plus significant bandwidth increases on the internal buses and I/O interfaces when compared with previous-generation Dell/EMC CX series arrays.[1] Equipped for powerful processing and high throughput, each array model in the CX3 series supports both remote mirroring and point-in-time copy operations. Hundreds of mirroring and point-in-time copy sessions can be performed simultaneously by a single array.

Optional replication programs that run on the CX3 series arrays include the EMC® MirrorView™ application for remote replication, the EMC SnapView™ application for local point-in-time copying, and the EMC SAN Copy™ application for cross-array data migration. Reflecting the rich heritage of EMC as a pioneer of array-based replication technologies, all three programs combine advanced functionality with a remarkably accessible user environment and full API-level integration across all services. In addition, the programs work with any supported data types and OS platforms and are production-proven, with thousands of deployments operational across all major industries. These software distinctions, in combination with the uniquely powerful processing architecture of the CX3 series, deliver exceptional capabilities for array-based replication.

Remote replication: MirrorView

MirrorView performs both synchronous and asynchronous remote mirroring over Fibre Channel links, including Fibre Channel tunneling over IP links for regional and geographic disaster recovery implementations. The software delivers a full set of premium replication and operational management features in an easy-to-use package that is both flexible and affordable. It works on all CX3 series platforms and facilitates mirroring between different array models and generations. With storage tiering supported across all CX products, IT organizations can take advantage of cost-effective disk media to bring disaster recovery within budget for organizations and applications within the enterprise.

Local point-in-time copying: SnapView

SnapView combines the ability to create local snapshots and clones in a single cost-effective package. In addition to copy and restore functions, SnapView includes advanced tools for automating copy operations and helping provide data consistency across versions. Tightly integrated with MirrorView, SnapView can be used to help optimize rapid recovery at disaster recovery sites, to support secondary processing operations such as backup and archiving, and to perform testing and mining. SnapView functionality is enhanced by the Replication Manager application for Microsoft® Exchange software, which provides key integration to automate disk-based backup and restore operations for Exchange stores. SnapView has unique flexibility to copy data to different RAID formats and disk media, enabling storage administrators to cost-effectively implement disk-to-disk backup and restore.

Cross-array data migration: SAN Copy

SAN Copy allows full-volume copy operations between Dell/EMC arrays, including both CX and AX platforms. It enables administrators to perform data migrations nondisruptively over any distance using Fibre Channel and Fibre Channel–to–IP tunneling technologies. The software comes with a full set of advanced features that enhance operational efficiency and give administrators flexibility to copy, relocate, distribute, and back up data from one storage system to another as requirements change. The software leverages the built-in storage-tiering capabilities of Dell/EMC array platforms to help contain storage costs.

[1] These performance improvements are for a Dell/EMC CX3 array with an Intel® Pentium® 4 Xeon® processor at 2.88 GHz when compared with a Dell/EMC CX300 array with an Intel Pentium III processor at 800 MHz. Actual performance will vary based on configuration, usage, and manufacturing variability.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

maintained at the target, rapid recovery can be compromised. In most cases, a full recovery from backup is required. By establishing consistency groups, administrators can provide write order integrity at the target, because this feature requires writes to all volumes in the group to complete in the proper order before an acknowledgment is sent to the application.

Bandwidth optimization for asynchronous mode. Techniques for grouping write changes for a definable period of time (referred to as delta sets) and intelligently filtering repeated changes so that only the most current are transmitted to the target can greatly help to reduce bandwidth requirements, enabling administrators to configure the system for the average write workload rather than the peak workload—which can provide significant cost savings in wide area network (WAN) replication.

Point-in-time copying

The ability to create either virtual copies (snapshots) or complete block-for-block copies (clones) of volumes is crucial to rapid recovery from logical or user errors and some forms of data corruption, such as viruses. Point-in-time copies can also be used for offline processing operations, such as backups, software testing, and data mining. These services typically generate copies locally, within the same array as the source volume. Because snapshots do not generate a complete one-for-one volume copy, they save space and can be created nearly instantly. However, generating a snapshot increases write overhead because every new write to the volume is accompanied by a copy of the old data to the snapshot index. A read access to the snapshot can also impose a performance penalty, particularly if the read requires access to the old data still residing in the source. By contrast, full-volume copies require a repository that is equal to or greater in size than the source volume, and they take more time to generate than snapshots do; however, they can be produced as a background operation and therefore have less impact on production operations than snapshots do. Also, access to volume copies does not impose performance penalties on operations to the source volume.

Premium point-in-time copying features

Choice arrays provide advanced tools for both creating copies and managing copy operations effectively. The following premium point-in-time copying features can add significant value to array platforms.

Copy repository right-sizing. Setting up snapshot repositories to effectively accommodate copy-on-write operations while making optimum use of valuable disk resources can be particularly challenging because write activity varies for individual applications. The ability to monitor write overhead on a particular logical unit (LUN) and then intelligently right-size the snapshot repository is a tremendous boon.

Remote copying. Although less-advanced array platforms permit only local copies, the ability of more-advanced platforms to create point-in-time copies on a remote array gives administrators powerful options for nondisruptive data migration and offline processing by a secondary server, which can use remote resources effectively to mitigate the performance overhead of offline operations from production nodes.

Versioning with scripting. This feature allows administrators to automatically generate multiple copies of individual LUNs in succession, providing a series of time-stamped images to use for rapid recovery. This is particularly crucial with logical corruption and viruses, which may not be detected immediately.

Application integration. The ability to integrate point-in-time copying operations with application hot-backup routines and then orchestrate the copy process through the application can address crucial consistency issues with backing up database applications and helps automate a traditionally labor-intensive procedure.

Log-based resynchronization. Administrators can quickly update full-volume point-in-time images by applying changes logged since the last image was created, allowing fast, low-overhead synchronization of full-volume images.

Reverse resynchronization. The ability to establish consistency on a production LUN using a simple resynchronization process not only helps speed up restores, but also helps reduce the potential for administrative error. This feature can also complement offline processing when a reprocessed LUN represents a new desired state of consistency for a source LUN.

Consistent split and restore. This feature helps provide consistent resynchronization for applications such as databases that interact with multiple LUNs.

Flexible and efficient data protection

As online operations continue to proliferate and dictate new requirements for rapid recovery across a growing base of applications, evaluating the benefits of array-based data protection becomes increasingly important. Array-based data protection extends the benefits of centralized operations to all levels of business, which can help improve efficiencies in both management and cost of ownership and provide the flexibility to tailor recovery scenarios for a complete range of disaster situations. In addition, array-based data protection can provide valuable management aids for administrators. Although basic features are commonly available among popular array platforms, choice array platforms deliver premium features that enhance rapid recovery capabilities and enable administrators to meet high levels of operational excellence.

Sonya R. Sexton is a competitive intelligence analyst, focused on storage systems, for the Enterprise Marketing Operations Group at Dell.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

VIRTUALIZATION GETS REAL

A virtualized IT infrastructure can deliver significant operational and economic benefits—boosting availability while slashing cost of ownership. Dell, Intel, VMware, Altiris and EMC provide industry-leading virtualization solutions that companies use today to gain a competitive edge.

When it comes to quantifiable business benefits, there's nothing "virtual" about a virtualized IT infrastructure. The payback is real—just ask Bob Neuberger, the Server, Storage and Database Manager for National Semiconductor. Santa Clara, Calif.-based National is the industry's leading manufacturer of high-performance analog devices and subsystems, reporting sales of $1.91 billion in its most recent fiscal year.

Like many firms, National experienced years of growth in the demand for Windows-based business applications. To provide the isolation necessary to run these programs reliably, National followed the common practice of deploying each on a dedicated server. Eventually the number of Windows-based servers and the cost of maintaining them became overwhelming—and most were very significantly underutilized.

"We had this beautiful farm of Dell PowerEdge servers," recalls Neuberger, "but many were only running at 5 percent of capacity."

Neuberger and his team chose an optimization strategy—based on Dell servers and VMware virtual infrastructure—that allowed them to consolidate up to eight physical servers as virtual machines running on a single Dell server.

"Our goal was to implement VMware software without purchasing any additional servers," Neuberger says. "We've easily accomplished that, despite continually implementing new Windows applications to support our internal customers."

www.dell.com/virtualizationtoday

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

A HIGHER STANDARD OF INTEGRATION

The close integration that distinguishes server virtualization solutions from Dell and competing offerings is nowhere better illustrated than in the partnership and joint development initiative between VMware and Altiris. The two companies have joined forces to address the challenges of concurrently managing both a host server infrastructure and a fleet of VMware virtual machines running on it.

Together, the two companies have created a hybrid administrative discipline called "Managed Virtualization." This combines Altiris' simplified approach to IT life-cycle management with the power of VMware virtualization. This innovative approach fosters efficient ways to manage not only physical servers, but also VMware virtual machines running as virtualized servers.

Managed Virtualization combines VMware technology with Altiris Server Deployment Suite and Altiris Server Management Suite. This combination provides IT administrators with enhanced capabilities to provision, deploy, manage and update Dell servers based on 64-bit Intel® Xeon® processors and VMware ESX Server running on that hardware.

The benefits that can result from this integration are clear and compelling: faster deployment of new services, streamlined management, reduced administration and improved responsiveness.

The payback? An expected 33 percent reduction in costs.

BELIEVE THE HYPE

National's experience is being duplicated in data centers around the globe. Analysts at Framingham, Mass.-based IDC believe, based on a 2005 study of server virtualization, that companies are rapidly adopting virtualized servers by partitioning x86 systems to contain costs and handle growing workloads. This study further found that survey respondents currently using server virtualization technologies expect 45 percent of new servers purchased this year to be virtualized.

IDC expects spending on virtualization to approach $15 billion worldwide by 2009. Analysts at Stamford, Conn.-based Gartner, Inc., concur. At Gartner Symposium/ITxpo 2005, the firm classified virtualization as a "megatrend," saying that it's the best tool that companies have right now to increase efficiency and drive up server utilization.

By decoupling workloads from physical server and storage resources through virtualization, enterprises can also:

- Enable dynamic capacity allocation
- Increase availability
- Simplify business continuity planning
- Reduce the cost of service delivery
- Lower total cost of ownership

The bottom line is that virtualization enables a more flexible and agile infrastructure—one that can respond readily to business change.

Regardless of the benefits, IT professionals may initially be wary of trading familiar challenges for new ones. Indeed, managing the transition from a physical data center environment to a virtual one can be a seemingly complex undertaking, preventing many organizations from realizing the financial and operational benefits. Consider the steps involved in a typical virtual infrastructure rollout:

- Determine hardware requirements and select a virtual infrastructure partner to provide planning and assessment services
- Implement the appropriate supporting storage environment
- Select and deploy virtualization software, as well as management tools for both physical and virtual resources
- Manage the cutover with minimal downtime

There are also cultural issues to consider. For instance, some customers and managers are accustomed to "owning" a physical server. In some cases, they may need to be convinced that a virtual environment can deliver the same levels of security, performance, reliability and scalability as a physical one.

Objections aside, however, the simple truth is that the operational and economic benefits of a virtualized IT infrastructure are far too significant to ignore. Companies hoping to leverage the advantages of virtualization must find a viable and economical way to address the issues.

MAKING VIRTUALIZATION MANAGEABLE

With these challenges in mind, powerful industry leaders—comprised of Dell, Intel, VMware, Altiris and EMC—are joining forces to bring standards-based innovation, volume production, and direct distribution to a virtualized IT infrastructure space.

Dell and its partners have worked

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

together to assemble and integrate a solution stack that makes the transition from a physical to virtual infrastructure swift and practical. The result is an easy-to-implement solution that packages leading hardware and software technologies, delivers them directly to the end user, and wraps them in the security of single-source support from the industry's most efficient provider of high-performance IT solutions.

Virtualization is an integral component of Dell's Scalable Enterprise Architecture, a framework for a highly efficient data center infrastructure. This framework combines standardized hardware and software components and a services organization capable of delivering automated, policy-based management in the future. It will enable dynamic real-time resource allocation in response to changing business conditions.

Dell brings together best-of-breed solutions, including server systems, SANs, services, virtual infrastructure software, virtual environment management software, and systems deployment and management software. Several key features set these solutions apart from the competition, including:

- A single, global point of contact for every aspect of assessment, design, delivery, deployment and support.
- Integration between management tools for both the physical and virtual environments.
- A full range of customizable virtualization solutions for small, medium and large companies.

Organizations worldwide are deploying these standards-based platforms today due to their unique combination of performance, tight integration, ease of deployment and unmatched value. Dell virtualization solutions address issues of low utilization and high management costs of existing environments, while preparing the IT infrastructure to scale and adapt more efficiently in the future.

DELIVERING PERFORMANCE AND VALUE

The core components of virtualization solutions available through Dell are standards-based hardware, software and services that deliver optimal performance and integration without locking the user into proprietary technologies that inflate costs or limit flexibility. These components include:

- **Dell PowerEdge Servers** offer a complete portfolio of blade and rack systems, in 2-way and 4-way configurations powered by Dual-Core Intel® Xeon® processors. Equipped with advanced reliability features for business-critical applications, Dell servers are designed to reduce the cost and complexity of enterprise environments by providing a standardized, high-performance platform for flexible, highly available and easily manageable scale-out architectures.

- **Intel Virtualization Technology** provides silicon-level software support specifically architected for virtual machine monitors that improves the dependability and interoperability of Dell PowerEdge servers with leading virtual infrastructure solutions such as VMware ESX and VMware Virtual Center. Intel also provides core server platform technologies such as Dual-Core Intel® Xeon® processors that enhance server performance and conserve power, and provide the reliability and scalability to help grow your business, reduce costs and mitigate risk.

- **VMware Virtual Infrastructure Technology** empowers IT organizations to match IT infrastructure—servers, storage and networks—to the information access and computing needs of their business. In a VMware virtual infrastructure the servers, storage and networking infrastructure can be treated as a common resource pool of flexible capacity. This is highly beneficial as a virtual infrastructure provides capabilities to rapidly deliver applications

THE TRUE VALUE OF A PACKAGED SOLUTION

To fully appreciate the value of Dell's comprehensive and proven virtualization solutions, consider the technical and organizational complexity of designing and implementing a solution. Tasks include:

- Confront and overcome the cultural objections to a virtual infrastructure solution. Don't be surprised to find business owners deeply attached to the "one application, one server" paradigm.
- Assess the current level of resource utilization and the available underutilized capacity.
- Assess the physical requirements of the intended virtual infrastructure.
- Select and deploy a virtualization software solution that is compatible with the physical hardware and software components.
- Select and deploy appropriate software tools to monitor and manage the physical resource pool.
- Select and deploy appropriate software tools to monitor and manage the virtual environment.
- Plan and manage the cutover from physical to virtual resources without interrupting service delivery or negatively impacting normal business processes.

When viewed in this light—particularly with respect to the internal resources required to complete the tasks on this list—the net worth of a fully integrated, tested and proven solution rises significantly.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

and services on the fly to users across the enterprise without the common effort and time required to procure and configure new IT resources.

VMware has a full range of software technologies that include a platform for migration of live virtual machine workloads across x86 host servers in a virtual infrastructure; systems management of virtual infrastructures; utilities that enable physical-to-virtual migration and simple server partitioning to manage multiple OS/application workloads on a single x86 host server. VMware virtualization technology provides many benefits to customers and lowers risk in the IT environment.

■ **Altiris Enterprise Systems Deployment and Management Tools** give IT administrators a comprehensive solution for monitoring and managing both virtual services and the physical servers that support them. Designed to simplify the entire server life cycle, the Altiris Server Management Suite provides deployment, management and monitoring functions from a centralized console to automate operations, improve system availability, and reduce overall infrastructure costs. It significantly reduces management complexity with comprehensive tools that streamline and automate provisioning, deployment, patch management, desired-state management and recovery services.

■ **Dell/EMC Storage Systems.** To eliminate virtual server dependencies on specific physical storage resources, consolidated storage is often necessary. Dell/EMC storage solutions combine EMC's data management technology leadership with Dell's unsurpassed capabilities in system design, manufacturing and distribution, providing a high-performance solution to support any size virtual environment.

■ **Dell Services** help companies realize the benefits of virtualization with solutions that deliver flexibility, ease of integration and deployment options that match business and budget requirements. Offerings are grouped into three phases: assessment, design, and implementation. During assessment, Dell service professionals determine current and future system requirements and provide ROI and TCO analyses. Next, a scalable solution is crafted with minimal end-user disruption in mind. For the implementation phase, designs are validated, a pilot consolidation is conducted, and data is migrated to the new system. Finally, Dell Enterprise Support Services provide proactive support with 24/7 rapid response and resolution. By delivering complete, end-to-end solutions, Dell Services ensure smooth transitions and speed time-to-benefit.

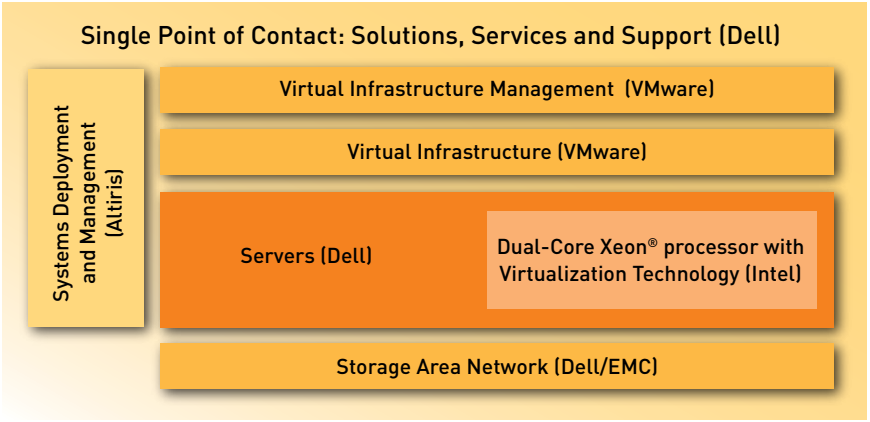

DELL VIRTUALIZED I.T. INFRASTRUCTURE SOLUTIONS INCLUDE:

- Dell PowerEdge servers powered by 64-bit, dual core Intel® Xeon® processors
- VMware virtual infrastructure and management software
- Dell/EMC storage area network solutions
- Altiris systems deployment and management tools
- Dell Services for planning, implementing and maintaining virtual physical environments

VIRTUALIZATION: A PRACTICAL BUSINESS STRATEGY

Dell brings together the industry's foremost standards-based server, silicon, storage, virtualization and systems management technologies in a tightly integrated package that is business proven, deployment-ready and that offers immediate benefits and positive ROI. These virtualization solutions are helping organizations worldwide unlock the power, potential and capacity of existing IT infrastructure, reduce costs, and respond more rapidly and profitably to new opportunities.

Just ask National's Neuberger. "When there's a hot new initiative, we can provision new servers in minutes," he explains, "and that helps reduce product development cycles. One of the main functions of IT at National is to enable the business to produce faster and more efficiently. The Dell-VMware combination is helping advance that goal by giving us tremendous on-demand capacity to meet the needs of our business."

By providing the business customer with a single source of sales and service responsibility, Dell is helping to bring the financial and operational advantages of a virtualized IT infrastructure to a wider range of potential beneficiaries. Whether your goal is consolidation, cost reduction, improved utilization, business continuity or increased agility, Dell server virtualization solutions hold the key to success.

www.dell.com/virtualizationtoday

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Server Virtualization in the Scalable Enterprise

Server virtualization can help bring measurable benefits to enterprises, including simplified operations, improved resource utilization, and cost-effective scalability to meet business demands. As this important technology evolves, Dell intends to continue driving standards development and delivering standards-based products that offer choice and control for enterprise IT organizations.

BY JIMMY D. PIKE AND DREW ENGSTROM

Related Categories:

Scalable enterprise

Virtualization

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

Virtualization is among the most significant topics in computing today. It is also a technology in the midst of a multiyear evolution from its origins as a discrete layer of middleware to part of the base infrastructure of industry-standard servers and operating systems. This article explains the role of server virtualization within the broad context of enterprise computing and explores the important attributes server virtualization brings to the Dell scalable enterprise strategy.

Dell scalable enterprise strategy

The Dell scalable enterprise strategy is a standards-based, multipurpose architecture that enables optimized use of resources, anticipates incremental growth, and supports extensive expansion. It focuses on standardizing core infrastructure elements, which can help reduce complexity and increase value. Systemic standardization, and the resulting interoperability, enables scaled-out solutions to be built from standard hardware building blocks and best-of-breed software components.

Why is standardization important? Mainstream, standardized technologies—whose development has been fueled by market volume and competitive pressures—can help deliver optimized, sustainable technology innovation with increased value. A simple look at where the investments and innovations are taking place in the IT industry reinforces the power of standards, interoperability, and integration.

Dell is unlike many of the world's top IT providers in that its scalable enterprise strategy is built upon, and demands, industry standards. The strategy charts a phased, practical road map for building the IT infrastructure of the future based on standardized technologies. Thus, organizations can avoid closed architectures that rely on proprietary shortcuts, which can increase costs while limiting flexibility and control. Dell, in contrast, works with industry-leading partners to develop an approach in which enterprises do not sacrifice openness to gain the value-added benefits of an advanced architecture.

The net result of the Dell scalable enterprise strategy is that enterprises can gain organizational benefits through simplified operations, improved utilization, and the ability to cost-effectively scale their IT infrastructure to meet changing business demands.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Basics of virtualization

Virtualization is a core enabling technology of the scalable enterprise strategy. Simply put, virtualization decouples software from hardware and presents a logical view of physical hardware to software, allowing a single server to act as multiple, independent servers.

Virtualization is a basic tenet of computing. It is used in various forms to abstract the interface of a resource from its physical implementation. Examples include storage (RAID or logical volumes), networking (virtual LAN), and Web-service application programming interfaces. Server virtualization applies the same concept to industry-standard servers.

Server virtualization is enabled primarily by software that sits between the server hardware and the OS. Often called the hypervisor, this software provides a container that presents a logical hardware interface to an OS. This logical interface enables hardware resources to be partitioned and shared among multiple guest operating systems, which can help improve server utilization. Additionally, each guest OS and its associated applications and data—known collectively as a virtual machine (VM)—are stored as a large file on the physical system. Administrators can easily transfer a VM from one physical system to another or save it on shared storage so that multiple systems can access it as needed.

Virtualization's role in IT infrastructure

Server virtualization can help simplify operations, improve utilization and uptime, and enable cost-effective scaling.

Simplifying operations

Server virtualization effectively hides hardware details from software, allowing the hardware to be truly interchangeable without affecting the software. Virtualization can help simplify operations by doing the following:

- Allowing workload portability across multiple servers, which includes the ability to "rehost" software (including legacy operating systems that are no longer supported)
- Streamlining application development and platform certification by certifying on a common virtual interface rather than multiple implementations of physical hardware
- Encapsulating complex configurations into a file that is easily replicated and provisioned

Improving utilization and uptime

Server virtualization can help organizations make optimal use of existing resources. A single physical server with multiple VMs can host several workloads that, without virtualization, would require multiple physical servers. Additionally, because workloads can be relocated or replicated easily in a virtualized environment, administrators can move them when performing maintenance without affecting service levels and can quickly provision new servers as needed.

Virtualization can help improve utilization and uptime by doing the following:

- Enabling safe resource sharing on industry-standard servers, so that if one VM fails, the other VMs on the physical server are not affected
- Providing the ability to migrate workloads dynamically from one physical server to another, so that workload service-level agreements can automatically match demand with capacity and system maintenance can be performed without disrupting enterprise services
- Empowering disaster recovery operations by restoring lost services regardless of the target physical platforms providing the services

Enabling cost-effective scaling

The management and utilization benefits provided by server virtualization facilitate cost-effective, "pay-as-you-grow" scalability across a common, standards-based infrastructure. Virtualization enables cost-effective scaling by doing the following:

- Separating the OS and application workload from the physical server and enabling dynamic scaling through resource sharing and rapid provisioning
- Leveraging and optimizing industry-standard hardware, which can help deliver sustained, incremental performance improvements at consistent, competitive prices

Efforts to standardize virtualization technology

Server virtualization is still in the early stages of a multiyear evolution that Dell believes will likely culminate with the technology being integrated as a standard component of the base IT infrastructure. However, virtualization technologies can be deployed effectively now to help deliver near-term benefits (see the "Usage scenarios" sidebar in this article for example cases).

Influence for today

Dell is currently engaged in a number of activities to help provide effective deployment of virtualization technologies today:

- **Dell labs:** Several Dell labs are working on practical implementation details related to virtualization, helping enterprises adopt virtualization quickly and confidently. The Solutions Engineering lab works closely with Dell software partners to certify interoperability between Dell hardware and required drivers and virtualization software. The Scalable Enterprise Technology Center creates scalable enterprise reference

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

architectures, including virtualized environments, to show administrators how to build and maintain these implementations. And the Technology Showcase gives in-depth technical briefings to organizations that want to learn more about virtualization and how it should be implemented in their environment.

- **Practical publications:** Dell delivers a set of simple, pragmatic publications to help enterprises understand a variety of subjects associated with virtualization (see the supplemental online section of this article at www.dell.com/powersolutions).
- **Platform management standards:** One of the key objectives influencing virtualization is management of the underlying platform. Dell supports platform management as defined by the Distributed Management Task Force (DMTF) and its Server Management Work Group. Specifically, Dell has been a key contributor to the development of Systems Management Architecture for Server Hardware (SMASH)—which defines architectural semantics, industry-standard protocols, and profiles to unify data center management—and the recently ratified SMASH Command-Line Protocol (SMASH CLP) specification.[1] SMASH CLP enables simple and intuitive management of heterogeneous servers independent of machine state, OS state, server system topology, or access method, thereby facilitating local and remote management of server hardware in out-of-service and out-of-band environments. Extensions supporting the management of VMs and their connection to physical resources are under development. Further work on VM and platform management and standardization is being conducted by the DMTF System Virtualization, Partitioning, and Clustering Work Group, which is charged with defining schema for managing VM resources and their dependencies.[2]

This focus on the management of virtualization, and integration into general systems management tools, is important to consider. Virtualization introduces capabilities requiring new management interfaces, but the use of and integration into standard tools and existing processes are key. For example, if an administrator uses virtualization to enable a single server to accomplish what 10 servers used to do, that administrator has 9 fewer physical servers to manage—but still has 10 operating systems, 10 applications, and the virtualization layer to manage. Integrated, standardized systems management options can help address these issues.

Impact for the future

Dell intends to focus on several key areas it deems essential to streamlining the evolution of virtualization:

- **Standard interfaces between the VM resource manager and VM operational interfaces:** Every managed element within a virtualized environment is associated with a resource manager. Typical functions consist of local operations within a VM or the management of multiple VMs. As such, the interface must be standardized and appropriate for multiple types of VMs.
- **VM monitor interoperability layer:** The VM monitor is the interface that links each VM to the underlying hardware resource. To date, successful virtualization schemes have either replaced or emulated instructions to enable virtualization. Some CPU and OS suppliers have begun adding virtualization support designed to eliminate this need. Although a single common hypervisor implementation is unlikely, development of a common interoperability mechanism that minimizes differences is important.
- **Virtual hard disk interoperability:** Today, there are at least two competing standards for the virtual hard disk format. One common standard must emerge for all virtualization schemes to support heterogeneous guest OS environments.

As these standards evolve, server virtualization as a technology may also affect the following:

- **Software distribution:** By creating a file-based distribution of server configurations, virtualization could dramatically simplify how servers are configured and provisioned.

[1] For more information about SMASH, visit www.dmtf.org/standards/smash.

[2] For more information about the DMTF System Virtualization, Partitioning, and Clustering Work Group, visit www.dmtf.org/about/committees/SVPC_Charter.pdf.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Intel® PRO Network Connections

Slot constrained?

Increase throughput by teaming embedded NICs with Intel® multi-port server adapters

intel Leap ahead™

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. Copyright© 2006, Intel Corporation. All rights reserved.

USAGE SCENARIOS

Although server virtualization is still an evolving technology, it can and should be incorporated into today's IT infrastructure to help increase effectiveness. The following usage cases describe how virtualization can be deployed effectively today.

Legacy OS support. Virtualization allows for continued support of guest operating systems and applications and can help provide long-term stability and support in a static environment. By decoupling hardware from the host OS, virtualization allows each to evolve without disrupting the other.

Software development and testing. Virtualization enables enterprises to create an environment for software development and testing without adding physical servers. It provides a private environment in which to develop software and a pristine environment in which to deploy it for testing. Multiple VMs can be used to create a multi-server environment for distributed testing on a single physical server. Additionally, IT departments can test patches and upgrades to exact copies of their environment without fear of corruption or end-user downtime. VMs, because they are stored in a flat file, can be set up in minutes and used multiple times. Certain software products also have an undo option to roll back changes to a VM.

Workload consolidation. Virtualization can accommodate many applications that do not coexist well in a conventional environment on a single platform, because the isolation between multiple guest operating systems provides distinct execution environments that do not overlap. The resulting improved utilization can help increase return on assets.

Software distribution. Virtualization can provide a reliable and portable distribution mechanism for software that is supported on multiple hardware platforms. Virtualization software combines the image of an OS and application and creates a portable "wrapper"—allowing for application deployment on any server supporting the guest OS. The host and guest support helps to eliminate recertification of applications on multiple platforms, provided the guest OS is supported by the host.

Load balancing. Virtualization can help create a highly reliable mechanism for load balancing across a pool of cost-effective resources. Load balancing is achieved by allowing guests to be transported among multiple physical servers, helping provide a robust and cost-effective environment and helping minimize exposure to hardware failure.

Disaster recovery. Virtualization can help create a highly reliable mechanism for reprovisioning and restoring execution environments if disaster occurs. The encapsulation of VMs provides distinct execution environments that can be renewed as needed. Several products are available that support this cost-effective environment for disaster recovery.

- **Dynamic, policy-driven resource allocation:** Virtualization enables workload portability between servers. Because VMs can be migrated from one physical server to another without interrupting operations, resources can be rebalanced dynamically to meet computing needs.

Virtualization standards can help enterprises avoid vendor lock-in, retain choice and control over their IT environment, improve access to resources to deploy and maintain a virtualized environment, and ultimately increase value as market forces compete for customers' virtualization dollars.

Virtualization: An essential data center technology

Server virtualization is a powerful technology that can significantly improve manageability and resource utilization. Although the technology is still maturing, current technologies can help deliver tangible benefits, particularly in supporting server consolidation, streamlining development and test environments, improving business continuity, and increasing IT responsiveness.

Dell believes that virtualization can become a standard feature of every industry-standard server system deployed, and that it is one of the essential technologies enabling the scalable enterprise strategy. Organizations should begin actively evaluating and implementing current technologies to realize these benefits now and to lay the foundation for broad-scale implementation in the future.

Dell is taking an active role in bringing virtualization to the mainstream by working closely with industry leaders to enable current technologies and help develop interface and management standards. Dell delivers tested virtualization infrastructure solutions along with services designed to streamline implementation and support. Dell also provides pragmatic guidance and best practices to help enterprises with deployment. By focusing on standards and avoiding proprietary shortcuts, Dell enables organizations to achieve a dynamic, scalable IT infrastructure while retaining choice and control.

Jimmy D. Pike is a director and distinguished engineer in the Dell Server Architecture and Technology Group, where he is responsible for the strategic system architecture of the enterprise product line.

Drew Engstrom manages the Scalable Enterprise Solutions team in the Dell Product Group, where he focuses on server virtualization and advanced data center architecture.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Extending the Virtualization Infrastructure

on Dell PowerEdge Servers Using VMware ESX Server 3 Advanced Features

Virtualization, which allows several operating systems to run simultaneously on a single physical server, is a key technology for achieving a scalable enterprise. Ninth-generation Dell™ PowerEdge™ servers also can help improve performance and scalability in enterprise data centers, especially when combined with VMware® ESX Server™ 3 virtualization software.

BY DAVID SCHMIDT, SCOTT STANFORD, AND BALASUBRAMANIAN CHANDRASEKARAN

Related Categories:

Dell ninth-generation servers

Dell PowerEdge servers

Virtual data center

Virtualization

VMware

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

VMware ESX Server 3 introduces features and capabilities not available in previous versions of this virtualization platform, including expanded support for storage technologies, advanced guest OS features, and other performance enhancements. When combined with ninth-generation Dell PowerEdge servers—which offer enhanced processor, memory, and connectivity options—these virtualization capabilities can be further extended and can help build a highly scalable data center.

Features introduced in VMware ESX Server 3

The features introduced in VMware ESX Server 3—support for Internet SCSI (iSCSI) and network attached storage (NAS), VMware ESX Server File System version 3 (VMFS-3), the four-way VMware Virtual SMP™ feature, 16 GB memory for virtual machines (VMs), and experimental support for 64-bit guest operating systems—make a compelling case for implementing virtualization in data centers, branch offices, and remote field locations. These features allow large, resource-intensive applications to be hosted on virtual platforms, particularly applications that require Physical Address Extension[1] or hugemem[2] support for optimal performance. When combined and appropriately leveraged, these features can enable a robust, high-performance, and scalable virtualization framework at the individual ESX Server level.

Internet SCSI and network attached storage support

Past versions of ESX Server supported locally attached storage and Fibre Channel storage area networks (SANs)

[1] For more information about Physical Address Extension, see "Planning Considerations for Intel Extended Memory 64 Technology on Servers and Workstations" by John Coombs and John Fruehe in *Dell Power Solutions,* June 2004, www.dell.com/downloads/global/power/ps2q04-018.pdf.

[2] For more information about hugemem, see "RHEL4 x86 kernel with more than 4 GB of RAM" by Matt Domsch, e-mail to *Linux-PowerEdge* mailing list, August 20, 2005, lists.us.dell.com/pipermail/linux-poweredge/2005-August/022327.html.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

for hosting and storing VM configuration files, guest OS virtual disk files, and application-level data within the guest OS flat file. To take advantage of the advanced VMware VMotion™ feature, previous-generation ESX Server hosts required shared storage supported by SANs.

In addition to supporting direct attach storage and Fibre Channel SANs, ESX Server 3 introduces support for iSCSI SANs and NAS devices. Both of these technologies provide storage across traditional Ethernet networks, which helps simplify deployment and management of storage fabrics and provides system administrators and storage architects with flexibility when designing, deploying, and upgrading ESX Server storage infrastructures. Although the initial list of supported NAS and iSCSI devices is limited, these technologies have the potential to deliver a new paradigm of enterprise storage.

Although still considered a new technology, iSCSI shows great promise as an enterprise-class storage fabric because of its cost-effectiveness and flexibility. Advanced Ethernet technologies and future iSCSI storage products will likely enhance the performance of iSCSI networks and lead to widespread adoption of this technology in the data center, where it can seamlessly scale and support the dynamics introduced by virtualization.

In addition to supporting direct attach storage and Fibre Channel SANs, ESX Server 3 introduces support for iSCSI SANs and NAS devices.

Like iSCSI, NAS provides flexible, cost-effective, high-volume storage. One NAS feature not commonly found in traditional monolithic file systems or storage devices is multi-protocol and multi–file system support. NAS devices provide simultaneous support for industry-standard file systems such as Network File System and Common Internet File System, making these multifunctional storage devices well suited for ESX Server, Linux®, UNIX®, or Microsoft® Windows® environments. In addition, many NAS systems can be seamlessly integrated into Fibre Channel SANs, serving as front-end gateways for ESX Server hosts that need to connect to back-end SAN disk resources over an IP-based transport layer.

VMware ESX Server File System version 3

VMFS-3 allows ESX Server 3 hosts to manage the entire VM on a single file system. In previous ESX Server versions, the VM swap file and configuration file were stored in a separate file system; VMFS-3 supports a directory structure in the file system, and all configuration and swap files are stored in the data store. This structure allows every component to be located on a single enterprise data store, which enables VMware VirtualCenter features such as Distributed Resource Scheduler and VMware High Availability to operate at the file system level.

VMware Virtual SMP

Four-way Virtual SMP is another key feature introduced with ESX Server 3. This feature provides support for four-processor symmetric multiprocessing (SMP) VMs, allowing guest operating systems to utilize up to four virtual processors. As in previous versions of ESX Server, Virtual SMP is enabled and scheduled at the VMkernel level for guest operating systems supporting applications that can benefit from or require SMP capabilities. Software that can benefit from four-way Virtual SMP include database, messaging, Web, streaming media, and other multi-threaded applications whose workloads are either processor-centric or most efficient when running in an SMP environment.

16 GB memory for virtual machines

In ESX Server 3, up to 16 GB of memory can be allocated to a VM, increasing the amount of memory that can be allocated to an individual VM by more than 75 percent compared with ESX Server 2.*x*. For applications such as databases that are memory sensitive, this increase can help yield substantial performance improvements, because much of the application-level workload can be contained within a VM's system memory. Because frequently requested application data can be held in a VM's memory, fewer calls are necessary to the VM virtual disk subsystem—meaning that in a VMware ESX Server 3 environment, the ESX Server VMkernel intercepts, translates, and manages fewer calls to the underlying physical disk controller and attached disk devices compared to an environment with less VM memory. For memory-sensitive applications, the end result is an efficient use of physical host-server system memory, reduced guest OS requests for data from the virtualized disk subsystem, and, most important, fast response times for database queries from end users.

Support for 64-bit guest operating systems

In the physical server environment, 64-bit processors and memory subsystems allow operating systems and applications to take advantage of wider and more efficient paths to system memory and processors compared to 32-bit hardware. As a result, more work can be done in the system memory than is possible on a 32-bit platform, resulting in fewer disk accesses and larger chunks of data processed per processor cycle.[3]

[3] For more information about the benefits of 64-bit processors and efficiencies gained by keeping work in system memory, see "Evaluating the Performance of 64-bit Oracle9*i* Database Release 2 on the Intel EM64T-based Platform" by Ramesh Radhakrishnan, Ph.D., in *Dell Power Solutions,* October 2004, www.dell.com/downloads/global/power/ps4q04-20040167-Radhakrishnan.pdf; "The Effect of L3 Cache Size on MMB2 Workloads" by Scott Stanford in *Dell Power Solutions,* February 2003, www1.us.dell.com/content/topics/global.aspx/power/en/ps1q03_stanford?c=us&cs=555&l=en&s=biz; and "VMware ESX Server Performance on Dell PowerEdge 2850 and PowerEdge 6850 Servers" by Todd Muirhead; Dave Jaffe, Ph.D.; and Scott Stanford in *Dell Power Solutions,* February 2006, www.dell.com/downloads/global/power/ps1q06-20050312-Muirhead.pdf.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Introducing VMware® Infrastructure 3 | The Next Generation in Virtualization

VMware delivers enterprise-class virtualization with VMware Infrastructure 3 Data Center Management and Optimization Suite

VMware Infrastructure 3 software suite features the award-winning, VMware ESX Server 3 with enhanced functionality for resource-intensive enterprise workloads and increased hardware interoperability. Revolutionary new distributed services features for increased business continuity include: VMware High Availability, VMware Distributed Resource Scheduler and VMware Consolidated Backup.

VMware Infrastructure 3 transforms computing, storage and networking hardware into a resource pool of IT power. This resource pool can be dynamically provisioned, migrated and managed to directly respond to the needs of organizations and users within enterprises, institutions and agencies.

VMware Infrastructure 3 combined with Dell PowerEdge servers delivers the Power of Virtualization. Dell customers, of all sizes—large, medium and small—realize significant benefits from the Power of Virtualization:

- Speed through accelerated responsiveness and flexibility to organizational needs
- Economic Savings from reduced operating and administrative costs
- Minimized Risks via increased application availability and rapid system recovery processes

Get The Power of Virtualization Today

Upgrade to VMware Infrastructure 3—contact your Dell sales team to inquire about eligibility for a limited-time special upgrade offer.

Evaluate VMware software: visit www.vmware.com/POV Click on Request VMware Evaluation Software

The Power of Virtualization

Dell and VMware jointly test and certify VMware Infrastructure 3 software on Dell PowerEdge Servers featuring Intel dual-core Intel Xeon Processors

vmware®

© 2006 VMware, Inc. All rights reserved. Protected by one or more of U.S. Patent Nos. 6,397,242, 6,496,847, 6,704,925, 6,711,672, 6,725,289, 6,735,601, 6,785,886, 6,789,156 and 6,795,966; patents pending. VMware, the VMware "boxes" logo and design, Virtual SMP and VMotion are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions. Dell, the Dell logo and PowerEdge are trademarks of Dell Inc. All other marks and names mentioned herein may be trademarks of their respective companies.

The concept of wider data paths and more work per processor cycle carries over into the virtualized infrastructure with ESX Server 3. Guest operating systems and applications that are 64-bit aware can take advantage of the wider data paths on 64-bit platforms. For ESX Server 3, 64-bit guest OS support is experimental. Application, OS, and database development, testing, and quality assurance activities are no longer limited to the 32-bit OS or application realm.

A 32-bit OS employs methods that enable 64-bit-like access for some applications, but do so with a potential performance penalty given the overhead required to remap and present system memory beyond the 4 GB memory level to applications as if it were part of the normal 32-bit memory address space. When formal, nonexperimental support for 64-bit guest operating systems is offered in future ESX Server versions, production-level workloads should be able to run in a complete 64-bit stack.

Virtual data centers with VirtualCenter 2.0

Data centers are often viewed as collections of individual or isolated processing elements in which software services or tasks are tied to specific hardware to achieve guaranteed service levels. Virtualization breaks this tight bond between software and hardware by mobilizing the software components. This mobilization is accomplished primarily through the migration of VMs within the data center. Administrators can use ESX Server 3 features to build virtual data centers on pools of versatile, scalable ESX Server hosts.

Figure 1. VMware Distributed Resource Scheduler and VMotion migrating virtual machines to a new location to achieve load balancing

In previous versions of VirtualCenter and ESX Server, VMotion events were primarily performed manually. To help improve on some of the manual processes that inherently limit dynamic systems reallocation in data centers, VirtualCenter 2.0 introduces a feature called Distributed Resource Scheduler (DRS). DRS provides a mechanism by which VMs can automatically migrate to a new location to achieve load balancing or avoid resource contention, as shown in Figure 1. DRS enables system administrators and architects to design application and workload containers on resource pools composed of VMs.

Because VMs are not tied to a particular server, DRS-based policies and rules enable VMs to autonomously and automatically migrate as data-center resource demands ebb and flow throughout the day or work cycle. Some key benefits provided by this autonomous and automatic migration feature are as follows:

- Physical data center resources can be used efficiently.
- Granular service-level agreement (SLA) guarantees can be achieved by configuring the VM resource pools for permanent or temporal assignment to specific hardware resource pools using an affinitization model or policy-based migration scheduling.
- Administrators can view the data center as a pool of resources instead of a collection of discrete servers.
- Where the software service is running becomes less relevant than it was without DRS.
- How to design and properly configure VM and hardware resource pools becomes more relevant than it was without DRS.
- Chief information officers have a robust and agile tool set with which to meet ever-changing business demands or strategic high-level IT initiatives.

Another key feature in VirtualCenter 2.0 is VMware High Availability (HA). VMware HA detects server failures and restarts the VMs on other servers, as shown in Figure 2. This method is similar to the uptime levels provided by Microsoft Cluster Server (MSCS) software, except that DRS can restart the VM on any other available ESX Server host in the appropriate hardware resource pool or VMware HA cluster. MSCS can be used for mission-critical applications, whereas DRS can be used for applications that can withstand minor downtime.

DRS in some configurations may be less complicated to set up than MSCS. Note that even though DRS is a VirtualCenter feature, once set up it can work independently of VirtualCenter. The ability to function independently is similar to an MSCS deployment model known as a majority node set server cluster.[4] For VMware HA, this

[4] For more information about majority node set server clusters, visit technet2.microsoft.com/WindowsServer/en/Library/c20dd042-5d52-49b2-889f-f163e0e112751033.mspx.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

ability means that any failure in the VirtualCenter infrastructure does not affect the ability of servers to work together to achieve high availability.[5] Both DRS and VMware HA are available as add-on services to VirtualCenter 2.0, and when used effectively, can help IT organizations achieve high levels of efficiency and flexibility in data centers.

Features introduced in ninth-generation Dell PowerEdge servers

Supporting the resource pools, dynamic provisioning, and flexibility offered by VirtualCenter 2.0 DRS and VMware HA, and the robust ESX Server 3 features at the individual compute-node hardware layer, are ninth-generation Dell PowerEdge servers.

Dell PowerEdge 1950, PowerEdge 2950, and PowerEdge 6850 servers include such advanced features as virtualization-specific processor technology, large memory support, high-bandwidth peripheral interconnect buses, and high-density local storage. These advanced features, although appropriate for any supported OS, are particularly well suited to enabling stable, dense, high-performance ESX Server 3 hosts that can easily become part of large resource pools in the VirtualCenter 2.0 virtual infrastructure.

Processors with Intel Virtualization Technology

Ninth-generation PowerEdge servers offer both the Intel® Xeon® 5000 series and Intel Xeon 5100 series dual-core processors and associated chipsets. Intel Xeon processors in Dell ninth-generation servers support dual-core processing—multiple physical processing units per processor die—and enable the Intel Virtualization Technology (VT) hardware feature.

VT provides virtualization-specific instructions at the processor level that can be used by the VM monitor for context-switching operations between VMs. This avoids unnecessary binary translation of the privileged instructions executed by the VMs. VT is required for 64-bit guest operating systems.

Large memory support

Ninth-generation PowerEdge servers support larger memory capacity than eighth-generation PowerEdge systems. The PowerEdge 2950, for example, supports up to 32 GB of memory in an 8 × 4 GB fully buffered dual in-line memory module (DIMM)[6] configuration. The benefits of large system memory amounts in a 2U system such as the PowerEdge 2950 for ESX Server 3 deployments include the following:

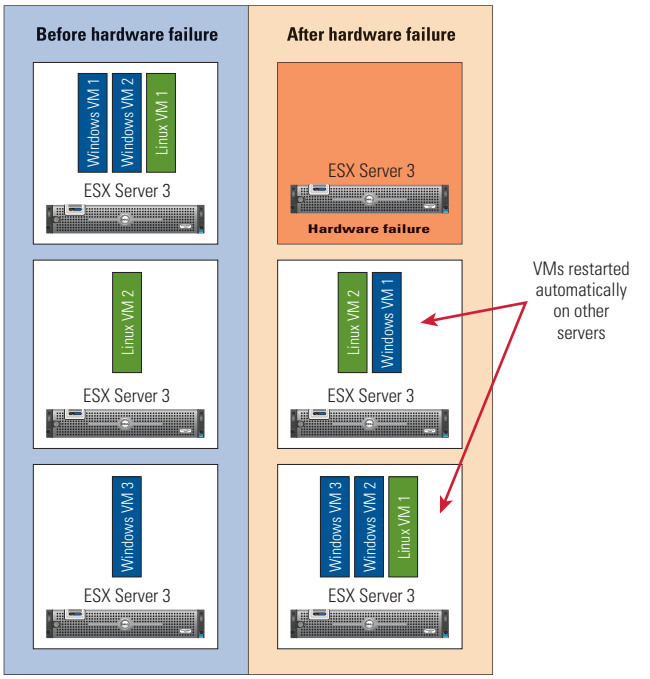

Figure 2. VMware High Availability Services restarting virtual machines on other servers following a hardware failure

- Support for more than 30 VMs, assuming a 1 GB active memory footprint per VM
- ESX Server host memory-system oversubscription models that take advantage of memory ballooning features provided by ESX Server and VMtools drivers

Ninth-generation PowerEdge systems with ESX Server 3 hosts also offer error-correction code (ECC) and single-device data-correction (SDDC)[7] features. ECC and SDDC implement single- and multi-bit memory-level error checking and correction functionality. These features are critical components designed to help VMs run in a reliable and stable memory subsystem.

PCI Express

Ninth-generation PowerEdge servers support port-dense, high-throughput, and low-latency PCI Express–based network interface card and host bus adapter configurations. High-performing PCI Express devices—with greater port density and higher capacity and throughput capabilities compared with previous-generation

[5] For more information about robust VirtualCenter infrastructure configurations, see "Architectural Considerations for Creating High-Availability VMware VirtualCenter Infrastructures" by Scott Stanford, Simone Shumate, and Balasubramanian Chandrasekaran in *Dell Power Solutions,* November 2005, www.dell.com/downloads/global/power/ps4q05-20050285-Stanford.pdf; and the VMware white paper "VMware VirtualCenter Technical Best Practices" at www.vmware.com/pdf/vc_technical_best.pdf.

[6] For more information about fully buffered DIMMs, see "Fully-Buffered DIMM Technology Moves Enterprise Platforms to the Next Level" by Jon Haas and Pete Vogt in *Technology@Intel Magazine,* March 2005, www.intel.com/technology/magazine/computing/Fully-buffered-DIMM-0305.htm.

[7] For more information about SDDC, visit www.intel.com/design/chipsets/applnots/29227401.pdf.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

PCI technologies—allow ESX Server 3 to support larger numbers of virtual network switches and virtual SCSI devices and ultimately more VM-based I/O traffic than ESX Server hosts configured with previous-generation PCI technologies. Because more VMs can access the PCI Express–based devices through the virtualized hardware presented by ESX Server without incurring I/O bandwidth limitations found in legacy PCI-enabled buses, administrators can theoretically deploy more VMs per ESX Server host, assuming that application and SLA metrics can be met while hosting a larger number of VMs than was possible with eighth-generation servers.

Serial Attached SCSI

The Serial Attached SCSI (SAS) protocol is designed to support full-duplex transfer rates of up to 3 Gbps. SAS controller small-form factors are well-suited for use as on-board server boot controllers or as add-in PCI Express–based cards to support external storage devices. In addition, SAS controllers support the use of smaller disk-drive form factors than traditional parallel SCSI devices. Ninth-generation PowerEdge servers leverage these features and offer a variety of high-density disk drives that can be configured in fault-tolerant yet high-performance configurations. SAS storage configurations are particularly well suited for deploying ESX Server configurations that require high capacity or performing locally attached VMFS-3 or ESX Server boot partitions.

By combining ESX Server 3 with ninth-generation PowerEdge servers, administrators can use virtualization to enhance data center flexibility, reliability, and scalability.

Dell OpenManage 5

ESX Server 3 supports Dell OpenManage™ 5 systems management software. Dell OpenManage is a suite of systems management applications for managing Dell servers. Dell OpenManage 5 supports ninth-generation PowerEdge servers, with management support for SAS storage controllers and the next generation of the Dell Remote Access Controller.

As in previous versions of ESX Server, the ESX Server 3 service console supports Red Hat™ Package Manager (RPM™)–based installation of Dell OpenManage 5. After installing the software, administrators can access systems management information through the Dell OpenManage Server Administrator interface. Administrators can also use the Intelligent Platform Management Interface to control PowerEdge servers.

Dell certified configurations

Dell designs, tests, validates, and certifies Dell PowerEdge, Dell PowerVault™, and Dell/EMC storage configurations to maintain precise alignment with Dell's fundamental principles of customer focus, supportability, simplicity, and reliability. ESX Server software and VirtualCenter can be factory-installed on Dell systems. In addition, installation guides, systems management papers, white papers, and best practices guides for Dell-VMware configurations can be found at www.dell.com/vmware. The resources on this Web site are updated based on customer feedback and as new options become available for supported configurations.

Flexible, scalable virtual data centers

VMware ESX Server 3 offers advanced features designed to enable a dynamic and scalable enterprise data center, and the advanced industry-standard features of ninth-generation Dell PowerEdge servers can help improve data center performance and reliability. In addition, ninth-generation PowerEdge server hardware with Intel Virtualization Technology can align with ESX Server 3 features such as 16 GB memory support and experimental 64-bit guest OS support. By combining ESX Server 3 with ninth-generation PowerEdge servers, administrators can use virtualization to enhance data center flexibility, reliability, and scalability.

David Schmidt manages the Virtualization Solutions Engineering team within the Dell Global Solutions Engineering organization. Previously, David worked as a software developer in the systems management group, where he worked on the Dell OpenManage Deployment Toolkit. David has a B.S. in Computer Engineering from Texas A&M University.

Scott Stanford manages the Custom Solutions Engineering team within the Dell Global Solutions Engineering organization. Previously, he was a systems engineer in the Virtualization Solutions Engineering team. He has a B.S. from Texas A&M University and an M.S. in Community and Regional Planning from The University of Texas at Austin, and he is pursuing an M.S. in Computer Information Systems at St. Edward's University.

Balasubramanian Chandrasekaran is a systems engineer in the Virtualization Solutions Engineering team within the Dell Global Solutions Engineering organization. His research interests include virtualization of data centers, high-speed interconnects, and high-performance computing. Balasubramanian has an M.S. in Computer Science from The Ohio State University.

FOR MORE INFORMATION

Dell and VMware:
www.dell.com/vmware

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Intel PRO/1000 Multi-Port Server Adapters for PCI Express: Advancing Network Capacity and Performance

With up to four Gigabit Ethernet connections in a single adapter, Intel® PRO/1000 multi-port server adapters for PCI Express help increase network bandwidth for slot-constrained servers.

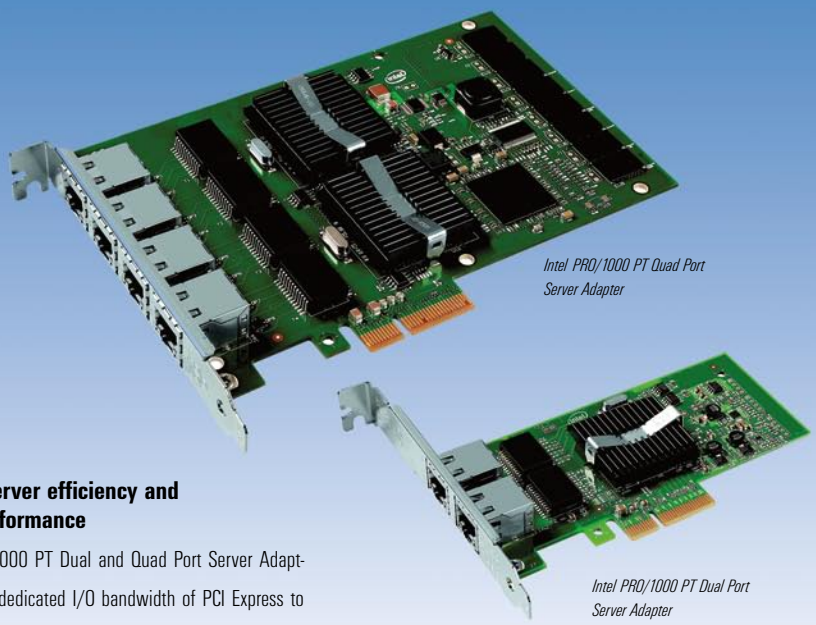

Intel PRO/1000 PT Quad Port Server Adapter

Intel PRO/1000 PT Dual Port Server Adapter

As business-critical applications push the envelope on performance, enterprises rely upon a high degree of network capacity, reliability, and availability. To meet these demands, network administrators leverage strategies such as server consolidation, virtualization, network segmentation, and multi-tier data center architecture. Enabling the IT infrastructure for these approaches is a key hardware element: the multi-port server adapter.

The Intel PRO/1000 PT Dual and Quad Port Server Adapters are part of the fifth generation of Intel network adapters for Gigabit Ethernet.[1] Besides conserving valuable server slots, Intel PRO/1000 multi-port server adapters are designed to leverage the full benefits of the PCI Express I/O standard, which enables packets to travel at full wire speed. As a result, Intel multi-port server adapters help minimize the time servers spend waiting for client responses.

For easy installation and maintenance, Intel PRO/1000 multi-port server adapters support Intel PRO Intelligent Install and the Intel PROSet Utility for Microsoft Device Manager. Intel PROSet helps simplify adapter implementation, including point-and-click configuration and management of Intel PRO network connections.

Enhanced server efficiency and network performance

The Intel PRO/1000 PT Dual and Quad Port Server Adapters support the dedicated I/O bandwidth of PCI Express to help ensure priority performance on each port without bus sharing. These Intel PRO/1000 adapters are designed to provide high performance in multiprocessor systems by efficiently balancing network loads across multiple CPUs when used with either Receive-Side Scaling on the Microsoft® Windows® OS or Scalable Network I/O on the Linux® OS.

Intel PRO/1000 PT Dual and Quad Port Server Adapters also support Intel I/O Acceleration Technology (Intel I/OAT), which moves network data efficiently through servers based on the new Dual-Core Intel® Xeon® processor to allow fast, scalable, and reliable networking. Intel I/OAT provides network acceleration that scales across multiple Ethernet ports and is a flexible choice because it is tightly integrated into popular operating systems.

Designed to be fully compliant with both Gigabit Ethernet and PCI Express standards, the Intel PRO/1000 PT Network Adapters are compatibility tested for interoperability with key network infrastructure components.

Dedicated ports for virtualization

For servers running VMware® server virtualization software, best practices recommend at least three ports to establish separate physical network connections for application production traffic, management, and virtual machine migration. Intel PRO/1000 PT Dual and Quad Port Server Adapters for PCI Express provide two or four Gigabit Ethernet connections in a single PCI Express card, increasing port density for slot-constrained servers. A single, integrated, dual-port Gigabit Ethernet controller chip enables high performance and reliability with low power consumption.

Minimized server downtime

Intel PRO/1000 multi-port server adapters for PCI Express include support for advanced server features, including adapter fault tolerance to provide redundant network links for server failover as well as adaptive load balancing and link aggregation for enhanced scalability and throughput. Intel Advanced Network Services (ANS), free software available with Intel PRO server adapters, helps administrators increase uptime using redundant, teamed connections.

For more information

Intel PRO/1000 multi-port server adapters:
www.intel.com/go/pcieadapters

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

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Nothing Basic About Next-Generation Enterprise Support

At the Dell Enterprise Command Center, specialists rigorously monitor Dell support operations and provide communication logistics anywhere in the world at a moment's notice—no matter what.

Dell Enterprise Command Center: Coordinating end-to-end support and crisis management

Dell Enterprise Support Services are designed to steer organizations out of harm's way before costly business disruptions occur. By executing proactive monitoring practices, repeatable problem-solving techniques, and field-tested operational procedures, Dell Enterprise Support Services specialists optimize the performance and availability of server and storage systems that support vital business functions.

When disasters occur, the Enterprise Support Services team instantly mobilizes the Dell global network of command, communications, control, and computer resources to expedite the recovery process and safeguard business continuity. Organizations faced with the dilemma of balancing uptime requirements with cost constraints may choose from a flexible and robust suite of Enterprise Support Services that can be tailored to a range of business requirements (see the "Enterprise Support Services" sidebar).

Dell Enterprise Command Center

The Gulf Coast's Hurricane Katrina, the London terrorist bombings, the Southeast Asia tsunamis, and the New York City transit strike–somber reminders of how natural and man-made circumstances can paralyze daily routines without warning. But in the midst of destruction and deadlock, the Dell Enterprise Command Centers (ECCs) were operating at full throttle.

While residents in the southern United States were fortifying homes and businesses in preparation for Hurricane Katrina, crisis management experts at the Dell ECC in Round Rock, Texas, were mobilizing rapid-response teams, moving resources to staging areas, and coordinating essential support activities. When the storm hit, the ECC was already working to restore critical systems for local U.S. Coast Guard stations and other emergency responders such as hospitals, police departments, fire departments, medical services, local government agencies, and power plants.

Dell created the first ECC in 2003 to meet growing business demands for proactive service response and continuous uptime. To choreograph its sophisticated logistics and tightly integrated network of resources, technology, and partners, Dell patterned the ECC after NASA Mission Control Centers, which coordinate every facet of manned rocket launches and flight control. Today, five ECCs are staffed around the clock, around the globe, for constant coordination of end-to-end support operations and critical incident management.

ECCs in the United States, Ireland, Japan, China, and Malaysia are at the ready to mobilize a network of 3,000 certified experts and 450 depot centers located the world over. Engineers and technicians monitor systems and activities 24 hours a day using state-of-the-art communication, computer, and data display equipment. ECC facilities display an intricate matrix of live data feeds, weather and news reports, maps, and real-time tracking panels.

How it works

The ECC, part of the Dell™ ServiceSystem™ delivery model (see sidebar), assimilates information from a wide network of sources and partners–weather services, traffic reports, news feeds, electronic communications, parts inventories–and provides a time window for arrival of on-site technicians, parts, and other dispatched resources.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Web-based ECC Real-Time Tracking Window

The ECC Real-Time Tracking Window integrates satellite mapping technology with Dell service monitoring to provide enterprises with direct access to inbound, in-process, and outbound logistics–integrating and updating the current status of each dispatch. Minute-by-minute movement details help organizations visually track progress, as if they were seated in an ECC.

The ECC Real-Time Tracking Window uses an HTTPS protocol and the Google Earth Pro application to securely display global information with point-and-click simplicity. To access the tool, enterprises must elect the Dell Platinum Plus Enterprise Support service level and purchase a license for Google Earth Pro (earth.google.com).

Operations Performance Benchmarking

Included in the new Dell Platinum Plus Enterprise Support service, the Operations Performance Benchmarking tool helps organizations continually assess how well their operations and server and storage systems are functioning–comparing critical IT performance metrics to industry benchmarks and historical results. In addition, this management tool illuminates internal performance comparisons by region. A customized dashboard gauges benchmark performance summaries of each location and analyzes time to resolution.

If a site is operating below par, a Dell Technical Account Manager identifies opportunities to improve performance, prescribes best practices, and recommends corrective actions to enhance operations and optimize uptime.

Logistical peace of mind

Cost-effective uptime need no longer present a pulse-pounding dilemma for administrators who must meet stringent performance and availability requirements for enterprise servers and storage. Dell Enterprise Support Services provide a flexible, comprehensive portfolio of plans, all rooted in a deep commitment to operational excellence, customer-led focus, and single-point accountability. This approach helps organizations protect their IT investments and find the calm before–and during–the storm.

Real-Time Tracking Window: Monitoring worldwide service dispatches

Enterprise Support Services

Silver, Gold, and Platinum Plus Enterprise Support offer increasingly higher levels of response and resolution, including the benefits of the Dell ServiceSystem delivery model. The new Platinum Plus Enterprise Support service is designed to be highly reliable and broadly available. In addition, it enables enterprises to access the Operations Performance Benchmarking tool and the Real-Time Tracking Window.

The Dell ServiceSystem is an integrated set of people, processes, tools, and infrastructure that forms the foundation for all Dell Enterprise Support Services:

- **Enterprise Command Centers:** A mission-critical networked infrastructure staffed by high-performance teams that continuously monitor and coordinate support operations, utilizing automated delivery tools and dispatching from 450 depot centers
- **Services Innovation Lab:** A combination of elite intellectual resources that drive meaningful innovations in process and tools designed to enable superior customer support
- **Services Specialist Network:** Highly coordinated and connected teams distributed around the world to help ensure responsiveness to a wide range of global operations
- **"Execute" Service Engagement Methodology:** Proven techniques and methods designed to drive the Dell operations teams to deliver services with reduced variance in execution

For more information

To learn about how the suite of Dell Enterprise Support Services can address specific business needs, visit **www.dell.com/services/enterprisesupport**.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Accelerating Dell PowerEdge Server Migration with

Symantec Backup Exec System Recovery

This article discusses some of the ways system administrators can use Symantec® Backup Exec™ System Recovery software for migration and deployment of new hardware into IT environments with previous-generation system architectures.

BY CHARLES BUTLER AND RICHARD GOODWIN

Related Categories:

Backup

Backup, recovery, and archiving (BURA)

Business continuity

Disaster recovery

Storage software

Symantec

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

The explosive growth of computing requirements across enterprises that use Microsoft® Windows® operating systems has increased hardware and processing demands on systems deployed to meet these requirements. Often the best way for administrators to cope with these demands is to have an ongoing process of migrating to and deploying the latest technology to relieve the strain placed on previous-generation system architectures. Although obtaining a new system that incorporates the most advanced hardware is a simple process, the actual deployment and integration with existing architecture pose a more complex problem: the difficult and time-consuming process of migrating the OS and software environment to the new hardware.

Symantec Backup Exec System Recovery Server Edition—with its Restore Anyware™ feature—allows administrators to quickly and easily create a complete system image and deploy it on a new system with dissimilar hardware.

Traditional server migration

To understand the power of Symantec Backup Exec System Recovery, administrators should first understand the tedious process of traditional server migration. This process requires backing up system data, configuring the new system hardware (including all new hardware drivers and peripheral device support), installing and updating the OS, installing and updating applications, and finally restoring the backup data from the old system onto the new system:

1. **Backing up system data.** Most environments include some method for data protection using either disk or tape as backup targets. Backups are run based on an optimal schedule to help protect system and application data. Administrators should close all applications running on the system and create a full backup (system image) of the system

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

to be replaced immediately before that system is taken offline to help ensure that the most current data on the system is migrated.

2. **Configuring the new system hardware.** This process typically involves customizing system BIOS settings and creating and initializing RAID configurations. In addition, the administrator must ensure that the appropriate low-level drivers, such as mass storage controllers, are available and working properly, and must be prepared with all the requisite high-level drivers to support the various hardware components needed by the OS.
3. **Installing and updating the OS.** This process is typically time-consuming, because administrators must locate and install the required OS updates and the latest versions of hardware drivers. Following a successful OS installation, administrators should configure the various system-specific configurations such as networking and security details.
4. **Installing and updating applications.** Following the OS installation, administrators must reload any system-specific software and install all applicable software updates. In addition, they must configure all of the custom application-specific settings, assuming the necessary records have been kept of all the configuration settings needed by each application on the previous system.
5. **Restoring the backup data.** The final step is restoring system and application data using traditional backup and recovery software. This process first involves cataloging the backup data sets contained on either disk or tape. Once cataloging is complete, administrators can restore the system data onto the new hardware. The new system is then ready for testing and subsequent deployment in production.

Figure 1 lists estimated protection and recovery times for a typical Windows-based server with a 25 GB system and application volume. As indicated by the times in the "Traditional backup and migration" column, this multistep process can easily consume more than six hours, even in a best-case scenario. Depending on the complexity of the system configuration and application roster, this time could be extended by hours or even days.

The second column lists estimated protection and recovery times using Symantec Backup Exec System Recovery instead of the traditional process. Performing system migration with Symantec Backup Exec System Recovery can save significant time, which can correlate to improved system uptime and business continuity.

Symantec Backup Exec System Recovery server migration

Symantec Backup Exec System Recovery combines the speed and reliability of disk-based, bare-metal Windows system recovery with innovative technologies for hardware-independent restoration. The result is the ability to perform seamless system migrations with a relatively simple process. The "Usage scenarios for Symantec Backup Exec System Recovery" sidebar in this article provides some examples of how administrators can use Symantec Backup Exec System Recovery.

The Symantec Backup Exec System Recovery Restore Anyware feature introduces an image-based system recovery capability for dissimilar hardware. The Restore Anyware feature helps make system migration to a completely different hardware platform simple and reliable. For example, administrators can restore an OS from a single-processor system onto a multiprocessor system and from SCSI storage onto Serial ATA or Serial Attached SCSI storage. In addition to these potential hardware differences, the Restore Anyware feature enables recovery to different hardware abstraction layers, chipsets, and kernel models.

Symantec Backup Exec System Recovery migration process

Before beginning the system migration process, administrators must gather any storage drivers required for the new system. The Symantec Recovery Disk contains a boot environment and the necessary drivers to allow administrators to perform system restores during the system migration process. The Symantec Recovery Disk also enables administrators to load drivers during the restore process if they are not already contained on this disk. Once the drivers are available, administrators can begin the system migration process.

Task	Traditional backup and migration (hours:minutes)	Symantec Backup Exec System Recovery backup and migration (hours:minutes)
Back up the system data with DAT 72 media	2:00	N/A
Create a recovery point	N/A	0:25
Configure the BIOS and RAID settings	0:10	0:10
Install the OS (from CD)	0:50	N/A
Update the OS	0:30	N/A
Install the applications (from CD)	0:15	N/A
Update the applications	0:15	N/A
Configure the system state and other unique settings or restore the recovery point	0:15	0:25
Restore the system data	2:00	N/A
Total	**6:15**	**1:00**

Figure 1. Estimated protection and recovery times for a typical Windows-based server with a 25 GB system and application volume

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

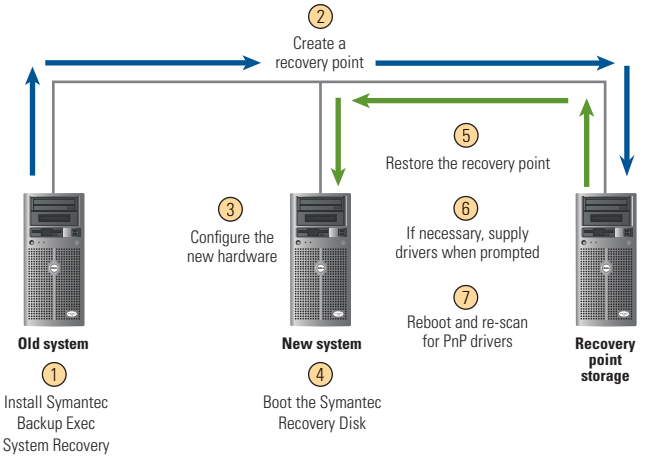

Figure 2. System migration process with Symantec Backup Exec System Recovery

Figure 2 illustrates the following steps for performing a system migration with Symantec Backup Exec System Recovery:

1. Install Symantec Backup Exec System Recovery on the old system from which the OS or data is being migrated.
2. Create a base recovery point of the old system and store it in a location that can be accessed by the new system, such as a network share.
3. Configure the new system hardware, such as the BIOS and RAID settings, if this has not already been completed.
4. Boot the new system into the Symantec Recovery Environment using the Symantec Recovery Disk.
5. Perform a recovery point restore using the recovery point created in step 2.
6. Supply storage drivers if prompted by the Restore Anyware retargeting procedure.
7. Upon reboot, allow Windows Plug and Play (PnP) to scan for additional devices and provide drivers if necessary.

Once these simple steps are completed, the new system is an exact replica of the previous system. This includes all OS components and application settings and the system state. By creating an exact replica, the amount of testing is minimized and the system can be rapidly deployed.

Symantec Backup Exec System Recovery helps simplify system migration by eliminating the software and OS reinstallation steps performed in the traditional system migration process. In addition, it is designed to preserve the unique system and software configuration settings of the environment being migrated. The disk-to-disk technology used in Symantec Backup Exec System Recovery also helps administrators to meet ambitious system migration time objectives.

USAGE SCENARIOS FOR SYMANTEC BACKUP EXEC SYSTEM RECOVERY

Symantec Backup Exec System Recovery is available in three versions: Server Edition, Windows Small Business Server Edition, and Desktop Edition. The powerful imaging technology utilized by Symantec Backup Exec System Recovery makes it a versatile tool that can be deployed in a variety of scenarios, including the following:

- **Dell™ PowerEdge™ server upgrades:** Administrators can perform migrations from previous-generation servers to the latest-generation technology.
- **Acquisition accounts:** Enterprises can easily migrate from Hewlett-Packard servers to Dell servers while minimizing the impact on resources.
- **Disaster recovery:** By protecting the OS and critical application and user data, Symantec Backup Exec System Recovery works with Symantec Backup Exec for Windows Servers to provide an end-to-end disaster recovery solution.
- **Windows upgrades:** Before applying OS upgrades, administrators can create a base recovery point that can be quickly restored if an OS environment fails after applying an update.

Advanced system protection and recovery

Symantec Backup Exec System Recovery is designed to offer advanced system protection by delivering a rapid, reliable approach to system recovery. It also provides a nonintrusive method for backing up mission-critical systems in organizations of any size.

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

Richard Goodwin is a technical product manager for the Backup Exec series of products at Symantec, and is responsible for incorporating customer and partner requirements into current and future releases. His experience in the storage industry spans almost a decade with Dell, VERITAS, and Symantec.

FOR MORE INFORMATION

Symantec Backup Exec:
www.backupexec.com/besr

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Managing Microsoft SQL Server 2005 with

Microsoft Operations Manager 2005

in a Dell Scalable Enterprise Architecture

The Microsoft® Operations Manager (MOM) 2005 tool enables enterprise IT organizations to consolidate server monitoring and management in one location. IT administrators can add hardware- and software-specific management packs to MOM 2005 to customize its capabilities. A team at the Dell Scalable Enterprise Technology Center implemented a representative scalable enterprise architecture—Dell™ PowerEdge™ servers running Microsoft SQL Server™ software in a cluster configuration—to demonstrate the capabilities of MOM 2005.

BY TODD MUIRHEAD

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Microsoft Operations Manager (MOM)

Microsoft SQL Server 2005

Microsoft Windows Server 2003

Scalable enterprise

Systems management

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

Built around industry standards, the Dell scalable enterprise architecture is designed to help organizations use standards-based products from Dell and Dell partners to build an architecture that can scale as requirements grow. As an example, the Dell Scalable Enterprise Technology Center built the Scalable Enterprise Reference Architecture model. This reference architecture uses Microsoft Operations Manager (MOM) 2005 integrated with Dell OpenManage™ tools to help simplify enterprise management by providing a unified management console.

To demonstrate the power of such an architecture, the Scalable Enterprise Technology Center team used MOM to manage a preemptive failover of a Microsoft SQL Server 2005 database in the event of a loss of hardware redundancy. MOM can respond to a hardware failure detected through Dell OpenManage by moving Microsoft Cluster Service (MSCS) from the server with a hardware failure to another node in the cluster. This article examines the deployment and configuration of the scalable enterprise architecture and the preemptive failover scenario.

Defining the scalable enterprise architecture

Dell has defined the scalable enterprise architecture as a method of using industry-standard components to create an infrastructure that is centrally managed and flexible enough to respond to changing requirements without requiring major changes in architecture. As elements of the data center standardize over time, centralized management and orchestration of resources becomes possible. The definition of the scalable enterprise includes multiple degrees of automation, culminating in comprehensive data center automation.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

With products that are available today, there are no clearly defined industry standards that can accomplish a truly centralized orchestration of all data center resources. However, a high degree of management centralization is achievable with some limited orchestration possible. The example of MOM automatically moving cluster resources in response to an event is a straightforward example, with broad application, of automation or orchestration.

The three principles of the scalable enterprise—simplified operations, improved resource utilization, and cost-effective scaling—guided the Scalable Enterprise Technology Center team's decisions for designing the example architecture described in this article. A single enterprise management console was combined with cost-effective, industry-standard hardware and software to create an environment that is highly manageable and highly scalable.

Understanding the MOM 2005 architecture

Microsoft Operations Manager 2005 is designed to manage and monitor an entire data center environment from a single console. MOM is scalable and can be deployed in highly available configurations that support the management and monitoring of tens of thousands of nodes.

MOM is designed around the unit of a management group, which consists of one MOM operations database, one or more management servers, and managed nodes. A MOM management group is controlled and configured through the MOM Administrator console, while day-to-day management and monitoring is performed through the MOM Operator console. MOM can be deployed in a highly available configuration by clustering the MOM database and using at least two MOM management servers. A highly available configuration of MOM 2005 was implemented in the Scalable Enterprise Technology Center, as shown in Figure 1.

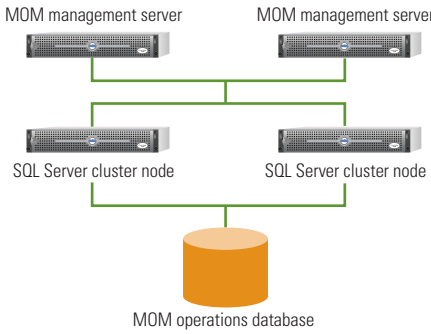

Figure 1. MOM high-availability configuration as set up in the Dell Scalable Enterprise Technology Center

The MOM operations database can be a clustered instance of Microsoft SQL Server 2000 which—when combined with multiple MOM management servers—can provide robust availability. All events and data that MOM gathers are stored in the operations database. A MOM management group can have more than two management servers, but to provide redundancy, two is the recommended minimum for any production deployment.

MOM management servers perform a discovery operation to find systems to manage. MOM can then manage the discovered systems by installing the MOM agent or by operating in an agentless mode. Installing the MOM agent provides a much higher level of functionality than the agentless mode; and best practices recommend using the agent, which runs as a Microsoft Windows® service, whenever possible. The MOM agent can be installed on the discovered systems via the MOM Administrator console.

Each system is assigned to all the appropriate rule groups that have been defined by the management packs installed in MOM. The MOM management server then pushes down the associated rules from each rule group to the agents as required by group membership of the servers in the MOM rule groups.

Management packs for MOM

MOM management packs are the method by which all monitoring and alerting is defined for MOM. A management pack provides predefined configuration checking, fault monitoring, performance alerts for a specific product or set of products, and product-specific expert knowledge. Microsoft makes a broad set of management packs available at no cost for almost all of Microsoft's enterprise products including Microsoft SQL Server™, Microsoft Exchange Server, Microsoft Active Directory® service, MSCS, and the Microsoft Windows server operating systems. Third parties also can develop and provide MOM management packs. Dell offers at no cost the Dell Management Pack for Microsoft Operations Manager, which links MOM and Dell OpenManage management tools.[1]

Installing the example architecture

The Scalable Enterprise Technology Center is dedicated to developing and testing example scalable enterprise architectures. Many infrastructure components are reused from project to project, including Active Directory, Domain Name System (DNS), Dynamic Host Configuration Protocol (DHCP), a Fibre Channel storage area network, and Altiris® software–based image deployment. This base platform was used as the starting point for the architecture described in this article.

For server deployment, the Scalable Enterprise Technology Center uses Altiris Deployment Solution™ software to perform

[1] For a comprehensive list of MOM management packs, visit www.microsoft.com/management/mma/catalog.aspx. The Dell Management Pack for MOM is available on this site as well as on the Dell support Web site (support.dell.com/support/downloads/download.aspx?fileid=127527).

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

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

SANbox® Express 1400 series

SANbox® 5600 FC Stackable switches

QLE2460/62 HBAS

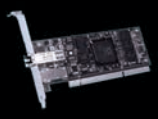

QLA®4050 HBAS

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.**

©2006 QLogic Corporation. Specifications are subject to change without notice. All rights reserved worldwide. QLogic and the QLogic logo are registered trademarks of QLogic Corporation. All other brands and product names are trademarks or registered trademarks of their respective owners.

image-based installations. In the example architecture, the Microsoft Windows Server® 2003 OS was deployed on industry-standard Intel® Xeon® processor–based Dell PowerEdge servers. One image of 32-bit Windows Server 2003 Release 2 (R2) and another image of 64-bit Windows Server 2003 R2—both of which could be deployed on Dell PowerEdge 1850, PowerEdge 1855, and PowerEdge 2850 servers—were used. Included in each image was Dell OpenManage Server Administrator (OMSA) as well as all of the preferred system setting configurations. These images enabled the Scalable Enterprise Technology Center team to quickly install the servers because deploying an image took only about 15 minutes. Figure 2 shows the components of the example architecture, including the MOM operations database, MOM management servers, and SQL Server 2005 cluster nodes.

MOM operations database

To create a highly available environment for the MOM operations database, the Scalable Enterprise Technology Center team clustered Microsoft SQL Server 2000 with Service Pack 4 on two Dell PowerEdge 2850 servers. Each PowerEdge 2850 server was deployed with the 32-bit image of Windows Server 2003 R2 and was connected to a Dell/EMC CX500 storage array via Emulex 9802 host bus adapter (HBA) PCI-X cards. The shared storage necessary for the SQL Server cluster was hosted on a five-disk (4 + 1) RAID-5 logical unit (LUN) on the CX500 array.

MOM management servers

MOM 2005 with Service Pack 1 was installed on a Dell PowerEdge 1850 server running the 32-bit image of Windows Server 2003 R2; this server acted as the primary MOM management server. The creation of a new management group was selected, which enabled a new MOM operations database to be created to support the new management group. The clustered SQL Server database was specified as the host for the MOM operations database during the installation. A new domain administration account, MOMAdmin, was created on the Windows domain and was specified as the MOM action account during the installation.

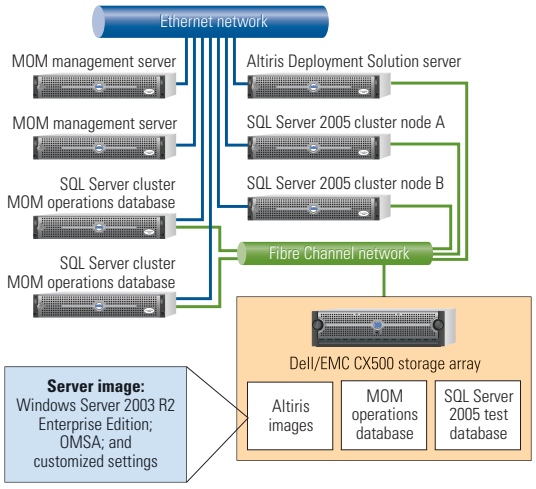

Figure 2. Components of the Dell Scalable Enterprise Technology Center example architecture

Following the installation of the initial MOM management server, a second MOM management server was installed on a virtual machine running on VMware® ESX Server™ virtualization software on a farm of two Dell PowerEdge 1855 servers. The option to create a new management server in an existing management group was selected and the same clustered SQL Server database used for the first MOM management server was specified as the location for the MOM operations database. The same MOM action account was also specified. Once the installation of the second MOM management server was complete, MOM automatically configured that server as the failover management server for all agents—meaning that, in the event that an agent is unable to reach the first management server, it fails over to the second management server.

MOM management packs were downloaded from the Microsoft MOM Web site (www.microsoft.com/mom) for Dell OpenManage, SQL Server, MSCS, DNS, Active Directory, and Windows server operating systems. Each management pack was installed on the local system and then imported from the MOM Administrator console.

To obtain additional useful tools for MOM, the Scalable Enterprise Technology Center team downloaded the MOM 2005 Resource Kit and the MOM 2005 Software Development Kit (SDK) from the MOM Web site and installed them on the first management server. The MOM script described in the "Using MOM to manage cluster failover" section in this article was developed with the help of tools from these kits.

SQL Server 2005 cluster nodes

For the managed nodes, the Scalable Enterprise Technology Center team installed the 64-bit image of Windows Server 2003 R2 Enterprise Edition on two Dell PowerEdge 2850 servers. SQL Server 2005 64-bit Enterprise Edition was then installed in a failover cluster on the two PowerEdge 2850 servers. Both of these servers were attached to the Dell/EMC CX500 storage array via Emulex 9802 HBAs.

Monitoring systems with MOM management packs

MOM management packs provide predefined rule definitions that MOM uses to monitor systems. Each management pack creates computer groups that have a formula for determining which systems are part of a certain computer group based on registry values. For example, MOM automatically places servers running SQL Server 2005 into the SQL Server 2005 computer group that was created by the SQL Server management pack. No manual setup is required

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

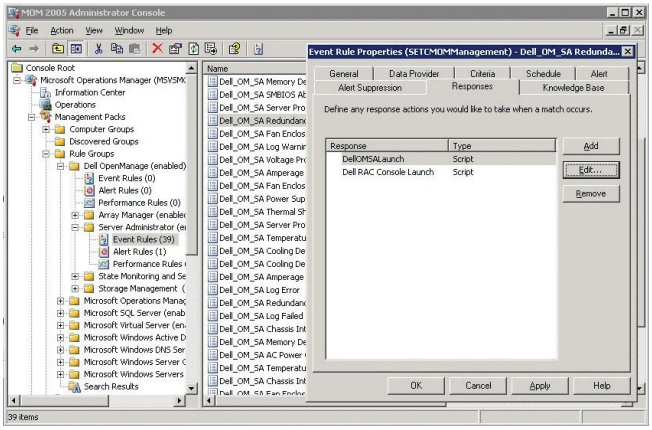

Figure 3. MOM Administrator console showing Dell Management Pack scripts to embed links for OMSA and the DRAC in a MOM event rule

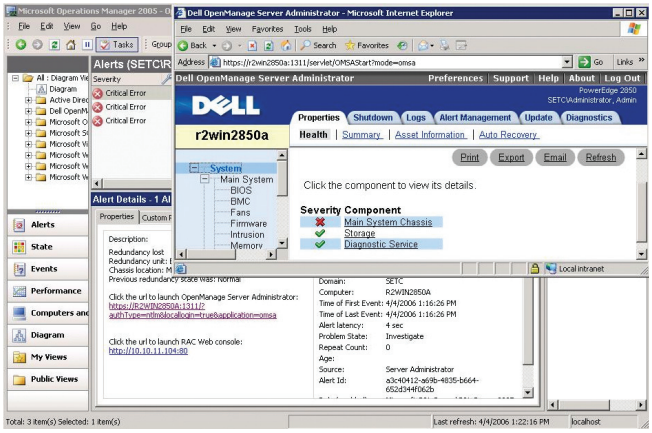

Figure 4. OMSA interface launched from an embedded link in an alert in the MOM Operator console

to add systems to the computer groups created by MOM management packs.

Rule groups created by management packs are predefined to be associated with the appropriate computer group. Each rule group can have event, performance, and alert rules defined. This section examines the Dell Management Pack for MOM and the Microsoft SQL Server management pack.

Dell Management Pack for MOM

The Dell Management Pack for MOM enables detailed, Dell-specific hardware information, including device failures and pre-failure alerts, to appear in MOM. Dell OMSA alerts and events are properly interpreted and the appropriate status changes are implemented for the affected server in the MOM Operator console. This means that Dell OpenManage critical and warning alerts change a Dell PowerEdge server's status in MOM to critical and warning, respectively. Dell OpenManage informational events are logged in the MOM operations database but do not change the server status.

The integration of OMSA and the Dell Remote Access Controller (DRAC) are also important components of the Dell Management Pack for MOM. In the event details section of a MOM alert, a link is included to launch OMSA or a DRAC session via a Web browser. Figures 3 and 4 show how these Dell tools integrate with the MOM Administrator console and MOM Operator console, respectively. This integration allows system administrators to quickly view the Dell server and investigate the problem in more detail if needed. MOM also can send e-mail notifications to designated groups of e-mail addresses. This allows MOM to be a central monitoring and alerting console for Dell servers when the Dell Management Pack is loaded.[2]

Microsoft SQL Server management pack

The Microsoft SQL Server management pack provides rules for detecting and sending alerts about critical events in SQL Server 2000 and SQL Server 2005. This management pack's goal is to help reduce or prevent SQL Server outages by proactively managing the configuration, performance, and security of SQL Server instances.

This management pack can alert administrators about many types of critical SQL Server conditions, including deadlock problems, blocking issues, and unavailability. Included in the alert is guidance for resolving the issue—configuration advice or problem resolution information based on the type of alert. Additionally, commonly monitored performance statistics for SQL Server with predefined thresholds are defined in the management pack as performance rules. These default thresholds can be modified, but they are a good starting point for monitoring the performance of a SQL Server instance. Administrators can add performance thresholds if there are additional metrics that need to be monitored closely.

Using notification groups, MOM sends out e-mail messages whenever any critical errors in SQL Server occur—similar to how the hardware errors detected by the Dell Management Pack can be used. In addition, notification e-mails can be sent from MOM if performance thresholds are exceeded. This means that MOM can log all events that occur on the SQL Server instances in the data center in a central console and then alert a defined list of database administrators when an alert is deemed critical enough.

Using MOM to manage cluster failover

Each MOM event can have a defined response, which can be as simple as sending a notification e-mail. MOM also allows for responses to be more active, such as a command-line execution

[2] For more information about the Dell Management Pack for MOM, see "Managing Dell PowerEdge Servers Using the Dell Management Pack for Microsoft Operations Manager" by Balasubramaniam J. in *Dell Power Solutions*, May 2006, www.dell.com/downloads/global/power/ps2q06-20060304-Microsoft.pdf.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

or a Windows script. Windows scripts use information from the event that triggered the response, enabling the script to programmatically take action based on the contents of an alert.

To demonstrate this powerful scripting capability, the Scalable Enterprise Technology Center team initiated a preemptive failover of MSCS. When a Dell PowerEdge server loses redundancy at the hardware level, such as after a power supply failure, a Dell OpenManage event is created. If the Dell Management Pack is loaded, MOM can detect this problem and change the status of the server to critical. A simple response is to send an e-mail notification to the server administrator. The team implemented a more advanced type of response: A Windows script that could determine whether the server was the active node for any MSCS-based cluster, and if so, could instruct those cluster services to be moved to another node not experiencing a hardware failure.

MOM configuration for demonstrating MSCS failover

For this example scenario, the configuration included MOM 2005 with SP1 and the Dell Management Pack for MOM; two Dell PowerEdge 2850 servers running Windows Server 2003 R2, MSCS, and SQL Server 2005 clustered; and a Windows script added as a MOM response to an alert event.

To help simplify the testing in MOM, a computer group called SQL 2005 Failover was created. This group contained only the two servers in the cluster—2850a and 2850b. This enabled the Scalable Enterprise Technology Center team to easily view these two systems in the MOM Operator console. A SQL 2005 Failover rule group was also created and applied to only the SQL 2005 Failover computer group so that, during testing and development of the script, only the two servers in the cluster would be involved.

The loss-of-redundancy rule from the Dell OpenManage rule group—Dell_OM_SA Redundancy Lost Warning—was copied to the SQL 2005 Failover rule group by copying and pasting in the MOM Administrator console. The name of the new rule was Cluster Proactive Failover. This rule monitored the event log of the server for an OMSA event with an ID of 1306. It also included two response scripts to insert links for OMSA and the DRAC into the MOM event generated when a loss in redundancy is detected (see Figure 5).

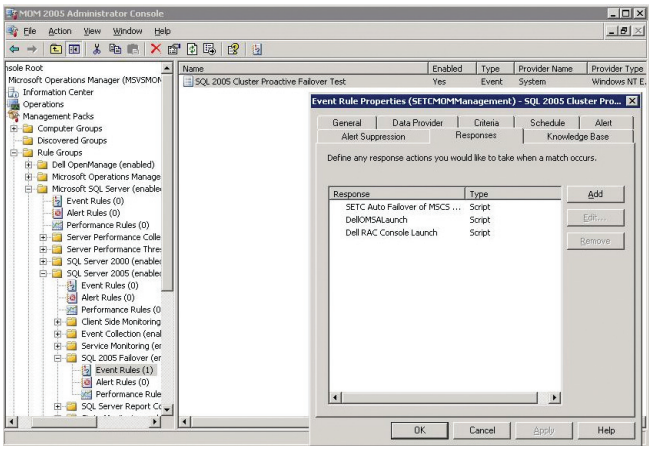

Figure 5. Responses for the SQL 2005 Cluster Proactive Failover event rule in the MOM Administrator console

MOM response script

To customize the failover rule, the Scalable Enterprise Technology Center team added a response that launches a Windows script. This script programmatically determines which cluster services are active on the server that has lost hardware redundancy and then moves those services to the other cluster node. The script was written with the Microsoft Visual Basic® development system and uses the Windows Management Interface (WMI), including special MOM classes, to accomplish the move.

> Microsoft Operations Manager 2005 and its management packs can create a powerful management tool for the enterprise data center.

Simple Visual Basic scripts that can be run with Windows Script Host (wscript.exe or cscript.exe) can be launched as a MOM response script. The major difference between Windows Script Host and MOM is that MOM scripts do not have access to the console because they are launched via the MOM process. For this reason, MOM provides a ScriptContext object, which provides the script with access to the MOM object that called the script and enables the creation of MOM events.

Microsoft provides several tools and documents to assist in the creation or development of Windows scripts for MOM. Some of these tools are included in the MOM 2005 SDK and MOM 2005 Resource Kit available on the MOM Web site. In addition, the Microsoft Scriptomatic tool and a tutorial on MOM scripting are available on the Microsoft scripting center Web site (www.microsoft.com/technet/scriptcenter).

Creating a non-MOM script. The first step in creating a MOM script should be to create a non-MOM version that can run from the command line with wscript.exe or cscript.exe. To create the correct WMI code to interrogate the desired information, Scriptomatic can automatically generate code that obtains the values for the necessary WMI objects. In the example scenario, Scriptomatic generated code that showed how to access the correct cluster information in WMI, specifically the cluster names and cluster group names that were associated with a given server. Using this code, the Scalable Enterprise Technology Center team wrote a command-line executable script—with a hard-coded host name—that could move active cluster groups to the other cluster node.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Converting to a MOM script. Using the MOM scripting documentation provided on the Microsoft Web site, the Scalable Enterprise Technology Center team was able to convert the script to a MOM script. This meant that all wscript.echo calls were removed and ScriptContext calls were inserted to obtain the name of the node that had experienced the problem and to create MOM events upon movement of the cluster services.

Testing the script with ResponseTest. Included in the MOM 2005 Resource Kit is a tool called ResponseTest, which allows for MOM scripts to be executed from the command line. To supply a MOM script with event data, ResponseTest accepts an XML file as an argument that defines all of the event data for the test. When ResponseTest is used to run the MOM script with the test event XML file, it simulates as closely as possible running a MOM script from the command line. This is extremely useful for debugging because tests can be rerun without having to actually trigger a MOM event.

Adding the script response in MOM. Once the script—when run with the ResponseTest tool—was able to successfully move the SQL Server 2005 cluster service, the script was copied into MOM and added as a response for the lost-redundancy alert event that was created and associated with the SQL 2005 Failover group. When this script successfully moves a cluster service, a MOM informational event is created so that all moves are logged and can be tracked. The script with comments can be found in the supplemental online section of this article at www.dell.com/powersolutions.

Failover in the example SQL Server 2005 cluster

With the MOM script added as an event response, the following steps occur automatically if a power supply fails on one of the Dell PowerEdge 2850 servers in the SQL Server 2005 cluster:

1. OMSA detects the failure and issues a failure event.
2. The MOM agent on the server detects that OMSA has issued a failure event and sends an alert to the MOM management server.
3. The SQL 2005 Failover group's predictive failover rule is triggered.
4. MOM creates an event that includes links to OMSA and the DRAC.
5. The predictive failover script is executed. It detects a SQL Server 2005 cluster instance running on the PowerEdge 2850 with the failed power supply and moves that instance to the other node in the cluster.

Figure 6 illustrates these steps. This is a simple example of the type of tasks that can be automated with MOM. Integrating the capabilities of MOM and Dell OpenManage in a centralized console can provide administrators with a breadth of information in one location.

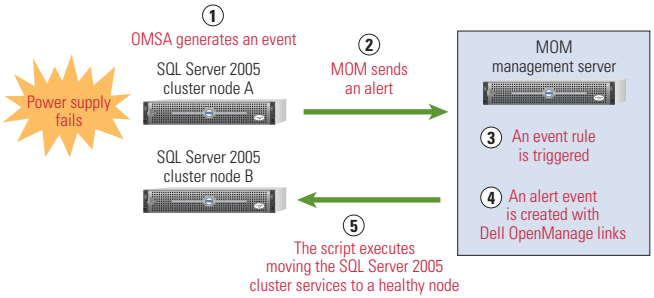

Figure 6. Failover in a MOM environment using a predictive failover script

Managing the Dell scalable enterprise through MOM

Microsoft Operations Manager 2005 and its management packs can create a powerful management tool for the enterprise data center. Specifically, the Dell Management Pack for Microsoft Operations Manager enables MOM to be aware of Dell-specific hardware events. Application-specific management packs, such as the Microsoft SQL Server management pack, also enable MOM to track and monitor application performance with appropriate performance counters. Combining all of this event data with the ability of MOM to respond with Windows scripts can enable functionality not possible with most stand-alone tools. The example scenario demonstrated by the Dell Scalable Enterprise Technology Center team—moving active cluster groups from a server that has detected a hardware problem—is possible only with a single management tool that aggregates information from many sources.

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.

FOR MORE INFORMATION

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

Dell scalable enterprise:
www.dell.com/enterprise

Dell OpenManage:
www.dell.com/openmanage

Dell Management Pack for Microsoft Operations Manager User's Guide:
support.dell.com/support/edocs/software/smdmpac/2.0/ug/ug.pdf

Microsoft Operations Manager:
www.microsoft.com/mom

Microsoft Scriptomatic:
www.microsoft.com/technet/scriptcenter/tools/scripto2.mspx

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Online Book Excerpt:

MOM 2005 Overview and Alert Flow

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. The MOM 2005 reference prescribes configurations tested in the Dell Scalable Enterprise Technology Center to integrate management systems and managed nodes for optimal performance and scalability. This Chapter 1 excerpt features a guided tour of the MOM alert flow.

BY TIM ABELS

Related Categories:

Dell PowerEdge servers

Dell Scalable Enterprise Technology Center

Enterprise management

Microsoft Operations Manager (MOM)

Microsoft Windows

Systems management

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

Outages of mission-critical systems can lead to lost revenue, disgruntled customers, and missed business opportunities. Outage frequency and recovery time are affected by manual, nonintegrated diagnostics and resolutions; unpredictable support costs; and staff turnover of undocumented knowledge and processes. To address these issues, the Dell Scalable Enterprise Technology Center team has written *Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise*. Chapter 1, "MOM 2005 Overview: The Scalable Foundation for Integrated

Best Practices for MOM 2005

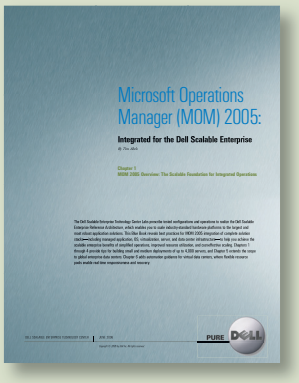

Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise marks the launch of the Dell Scalable Enterprise Technology Center reference series. This online book reveals best practices for MOM 2005 integration of complete solution stacks—including managed application, OS, virtualization, server, and data center infrastructure—to help simplify operations, improve resource utilization, and scale out the IT infrastructure quickly, flexibly, and cost-effectively.

For the full text of Chapter 1, "MOM 2005 Overview: The Scalable Foundation for Integrated Operations," go to www.dell.com/downloads/global/power/mom_2005_dellse_abels_ch1.pdf. For more information about the Dell Scalable Enterprise Technology Center, visit www.dell.com/setc.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Operations," explains the MOM process; management group components; versions, editions, and extensions; accounts and users; and console, management pack (MP), and agent components.

Organizing labs-tested advice

Microsoft Operations Manager (MOM) 2005: Integrated for the Dell Scalable Enterprise contains expert guidance from the Dell Scalable Enterprise Technology Center team, including *tips,* which share specific information to help organizations use MOM 2005 to their best advantage; *notes,* which point to important related information; and *cautions,* which point out potential insecurities or damage—and ways to avoid the problem. The following sample tips, notes, and cautions are excerpted from Chapter 1:

- **Tip:** To prevent management servers from discovering and managing each other, add an Exclude Managed Computer rule to the primary and then the secondary management servers.
- **Note:** If the Action account is disabled, deleted, or password expired, you need to reset the Action account; for more information, visit support.microsoft.com/default.aspx?scid=kb;en-us;883347. For example, for the Microsoft Active Directory® directory service MP, the Action account must be a domain account, which will probably require password expiration.
- **Caution:** To prevent MOM database overgrowth, set an alert rule to automatically notify and escalate on grooming job failures.

Keeping pace with data center evolution

Key to the evolution of the data center is standardization—of software, servers, storage, and switched fabrics. To help plan the future of advanced data centers, Dell has created the Scalable Enterprise Reference Architecture, a model that serves as a taxonomy for designing current and emerging data center environments. This reference architecture is built upon industry-standard applications and 64-bit Intel® Xeon® x86 multi-core processor–based servers, using virtualization to scale quickly and flexibly to meet unpredictable business needs and evolving business requirements.

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.

FOR MORE INFORMATION

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

Dell scalable enterprise:
www.dell.com/enterprise

Dell and Microsoft systems management:
www.dell.com/content/topics/global.aspx/alliances/en/microsoft_sms

Guided Tour: MOM alert flow

Figure 1 shows the flow of an alert through the MOM components with numbered tags 1 through 10 (items 1, 8, 9, and 10 handle alerts offline only, particularly item 1, the Administrator console).

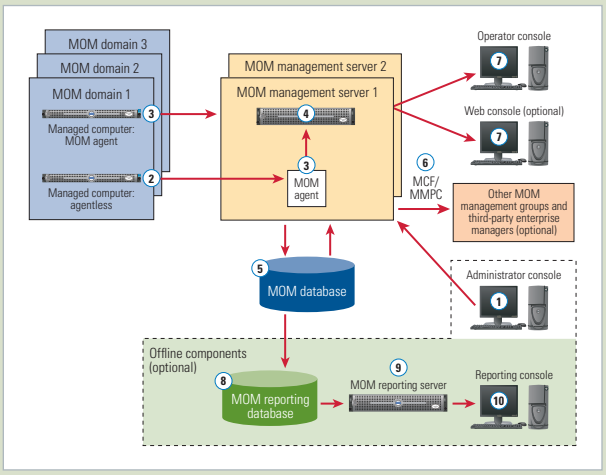

Figure 1. General processing flow of a MOM alert

1. Offline, administrators can define alerts at the **Administrator console**.
2. At runtime, **agentless managed computers**, which use an agent on the MOM management server, can generate alerts.
3. **Agent-managed computers**, which host **MOM agents**, can also generate alerts.
4. The managed computers send the alerts to the **MOM management server**, which passes the alert data to the MOM database.
5. The **MOM database** stores the alert data.
6. The **MOM Connector Framework (MCF)** or **MOM-to-MOM Product Connector (MMPC)** enables alerts to be shared in two-way communication with **other MOM management groups** or **third-party enterprise managers**.
7. Users can perform ongoing alert management from the **Operator console** or **Web console**.
8. For offline analysis and reporting on historical MOM alerts and other data, the **MOM reporting database** pulls data from the MOM database.
9. The **MOM reporting server** extracts and presents the alert data from the reporting database.
10. Users can manage the MOM reporting server through the **Reporting console**.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Applying Updates for Dell PowerEdge Servers Using

Microsoft Systems Management Server 2003 Part 1

Managing hardware updates has become a key feature of Microsoft® Systems Management Server (SMS) 2003 with the introduction of the SMS 2003 Inventory Tool for Dell Update. When designing and implementing a methodology to manage software updates, administrators can also plan for and select appropriate hardware updates, thereby consolidating both functions into a single process that can help keep environments stable and reliable.

Related Categories:

Change management

Dell OpenManage

Dell PowerEdge servers

Dell Update Packages

Microsoft Systems Management Server (SMS)

Storage management

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

The Microsoft Systems Management Server (SMS) 2003 Inventory Tool for Dell Update (ITDU) helps administrators update Dell™ PowerEdge™ servers with the latest system BIOS; system firmware, also known as Embedded Server Management (ESM) firmware; remote access controller firmware; RAID controller firmware; and device drivers. This tool utilizes the existing software update management feature of SMS 2003 Service Pack 1 (SP1) and later to help automate the update process for Dell systems, minimizing the resources required to transport Dell Update Packages across a network. Once the SMS 2003 ITDU software is installed, the SMS console interface seamlessly provides access to its set of Dell update capabilities (see Figure 1).

This article provides best-practices guidelines for those who manage, review, or approve hardware updates for Dell servers or oversee Dell server–based data centers and test environments. The update process described in this article is based on a high-level, four-stage patch management process model. The four phases are as follows:

- **Assess:** Administrators determine what systems are in the production environment, what security threats and vulnerabilities might affect those systems, and whether the organization is prepared to respond to updates.
- **Identify:** Administrators discover updates in a reliable way, determine whether the updates are relevant to the environment, and determine whether an update represents a standard response or an emergency change.
- **Evaluate and Plan:** Administrators decide whether to deploy the update, determine what is needed to deploy it, and test the update in a production-like environment to help confirm that it does not compromise critical systems and applications.
- **Deploy:** Administrators roll out the approved update to the production environment to meet the requirements of any deployment service-level agreements (SLAs).

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

This article discusses the Assess phase and the Identify phase; for the Evaluate and Plan phase and the Deploy phase, see "Applying Updates for Dell PowerEdge Servers Using Microsoft Systems Management Server 2003: Part 2" at www.dell.com/downloads/global/power/ps3q06-20060249-Microsoft-SOE.pdf.

Prerequisites for using the SMS 2003 Inventory Tool for Dell Update

Before using the SMS 2003 ITDU tools for patch management, administrators should have basic network and system administrative skills, including the ability to manage software and hardware update life cycles. In addition, they should be familiar with Microsoft SMS 2003 and SMS 2003 SP1; SMS 2003 software update scanning tools, consisting of the Security Update Inventory Tool and the Office Inventory Tool for Updates; the latest version of Microsoft Baseline Security Analyzer (MBSA); and Dell server hardware components and models.

The SMS 2003 ITDU is designed to scan and update supported servers[1] that meet all of the following criteria:

- Managed by SMS 2003 SP1 or later
- Running a supported OS: Microsoft Windows® 2000 Server SP3 or later; Microsoft Windows 2000 Advanced Server SP3 or later; or Microsoft Windows Server® 2003 Standard Edition or Enterprise Edition
- Using one of the supported OS languages within a supported SMS configuration

The SMS 2003 ITDU must be installed on an SMS 2003 SP1 site server, and it should be installed and configured from the most central SMS 2003 SP1 site in the hierarchy, allowing central administration and configuration of hardware updates.

The SMS 2003 ITDU uses components from a Dell Partner Development Kit (PDK), and support for these components does not follow the standard Microsoft enterprise product support life cycle. Microsoft intends to issue periodic updates to the SMS 2003 ITDU; only the most current version will be supported. Dell intends to issue periodic updates to the PDK, and the SMS 2003 ITDU includes a feature to automatically download compatible PDK versions; if this feature is not used, administrators must manually download the PDK at regular intervals. Only the current and two previous versions of the PDK support the SMS 2003 ITDU.

Key challenges for the hardware update process

Administrators should be aware of several key challenges in applying hardware updates to Dell server environments. For example, some Dell server models require multiple updates to be applied at the same time and in the correct sequence. Some servers may also not have a working SMS agent installed, meaning SMS may not be able to inventory or manage them. In addition, custom configurations may cause unexpected results during update deployment.

Assess phase

During the Assess phase, administrators establish baselines for the existing IT environment to prepare the environment for patching. This ongoing process helps determine which Dell server models and hardware components exist in the environment.

Activities carried out during the Assess phase include discovering and inventorying Dell servers, assessing Dell hardware configurations, assessing Dell hardware dependencies, and establishing a baseline for Dell server hardware.

Discovering and inventorying Dell servers

Administrators can use the SMS 2003 systems discovery mechanism to discover computers in the IT environment. They must then inventory the systems to identify Dell servers according to

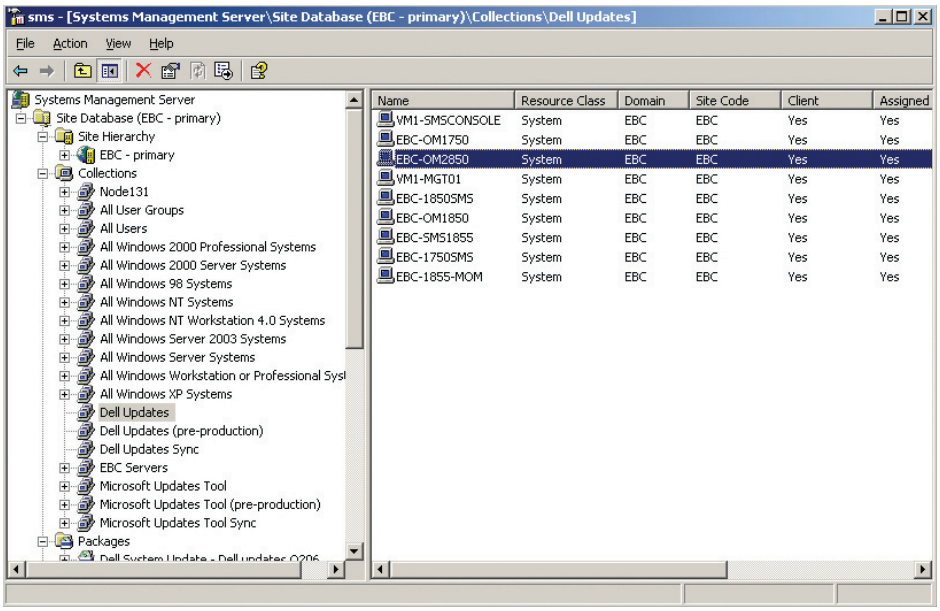

Figure 1. Access to Dell updates following installation of the SMS 2003 Inventory Tool for Dell Update within an SMS site

[1] For a list of Dell server models supported by the SMS 2003 ITDU, see the readme file at ftp.dell.com/cmsdk/PDK_readme.doc.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

the server model types and versions, OS types and versions, server roles, applications and versions, application dependencies, network connectivity, and applicable and installed hardware updates.[2]

The SMS 2003 ITDU assists with the Assess phase using the following:

- **Scan component:** Enables administrators to scan Dell servers for installed and missing updates, in the same way MBSA determines compliance for Microsoft security updates
- **Synchronization Host component:** Downloads the Dell Update Catalog from within SMS directly from the Dell FTP site on a recurring schedule
- **Dell-specific reports:** Show installed and missing Dell system and component updates

Assessing Dell hardware configurations

The next step in the Assess phase requires understanding how Dell servers have been configured, specifically whether a server is attached to an external storage device, such as a storage area network (SAN) or network attached storage (NAS) device. Depending on how an external device is configured with the server, a reboot may need to be suppressed until a local administrator is physically available on-site.

The SMS 2003 ITDU cannot determine the types of devices attached to a Dell server. Nevertheless, administrators should understand how external devices interact with the server and how the server manages and configures devices. They can use the SMS 2003 SP1 hardware inventory feature to inventory a server for applications or tools that are loaded when a disk controller or external device connectivity is set up.

Some items that should be considered when assessing Dell hardware configurations include the following:

- Attached devices and their configurations, including SAN, NAS, and SCSI
- Dell servers configured in an active/active or active/passive cluster
- Potential compatibility issues with Dell management tools such as the Dell OpenManage™ suite and Dell OpenManage Array Manager, which require specific hardware versions to be intact on Dell servers
- Customized BIOS settings, including BIOS passwords, boot sequence, and so on
- Multiple operating systems on a single Dell server
- Boot configuration order, including drives and directories
- SLAs associated with the specific server and the services it is providing
- Approved maintenance windows for applying updates and taking the server out of service

Although stand-alone servers may not need as much attention as servers with attached devices, administrators should use the SMS 2003 SP1 hardware and software inventory features to assess stand-alone configurations.

Assessing Dell hardware dependencies

Updating hardware components on Dell server models using the SMS 2003 ITDU requires an assessment of the following:

- OS version and service pack
- Server manufacturer and model
- Hardware component versions, including drivers for devices such as the Dell PowerEdge Expandable RAID Controller (PERC); firmware for devices such as the PERC, Cost Effective RAID Controller (CERC), Dell Remote Access Controller II (DRAC II), DRAC III, and so on; ESM firmware; and BIOS
- Dell OpenManage version, including Server Administrator and Array Manager
- SMS site maintenance tasks, particularly the Delete Aged Inventory History and Delete Inactive/Obsolete Client Discovery Data tasks
- Patch age date, located in *X*:\Program Files\Dell Update Inventory Tool\PkgSource\Scan.ini, where *X* represents the drive where the SMS 2003 ITDU was originally installed

Validity of assessment results. The SMS 2003 ITDU scans for applicable and installed hardware updates for supported Dell server models. Hardware component details reported to the SMS 2003 site database depend on the hardware component manufacturer's ability to publish this information to Windows Management Instrumentation (WMI) and on whether the vendor's hardware configuration software has been installed to pick up those additional hardware attributes in WMI. It is not always possible to capture all details, such as the asset tag, serial number (service tag), or system ID.

Note: Administrators should consider the time interval of Dell hardware updates, because hardware updates may not occur as frequently as OS and application updates. If hardware inventory does not run during a 20-day period, hardware updates are aged out of the WIN32_PatchStat_Extended WMI class. To help avoid this problem, administrators should modify the value of the PatchAge setting in the Scan.ini file to coincide with the hardware upgrade interval. If this setting is left at the default setting of 20, updates

[2] For more information about SMS 2003 systems discovery and inventory features, see the *Systems Management Server 2003 Operations Guide* at www.microsoft.com/technet/prodtechnol/sms/sms2003/opsguide/default.mspx and *Systems Management Server 2003 Concepts, Planning, and Deployment Guide* at www.microsoft.com/technet/prodtechnol/sms/sms2003/cpdg/default.mspx.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

may not report back properly within the SMS Administrator console, especially if the SMS site maintenance tasks are configured to occur frequently. After this file has been updated, administrators should refresh the distribution points to reflect the change and have the clients perform a scan with the updated file.

Dell OpenManage dependencies. The Dell OpenManage suite, including Server Administrator and Array Manager, can be dependent on a particular hardware component version. Administrators may need to upgrade the suite in conjunction with a hardware update—for example, BIOS, firmware, and driver updates are typically dependent on the latest Dell OpenManage version. If the updates are applied and do not correlate with the Dell OpenManage version, the suite may not operate correctly.

Once the SMS 2003 ITDU software has been installed in the site, the SMS console can indicate whether a server is running the required minimum version of Dell OpenManage (4.2 or later). Administrators also have several other ways to assess the Server Administrator and Array Manager components of Dell OpenManage 4.2 or later. First, they can create a query to collect the version information using the SMS hardware class Add/Remove Programs (assuming the application has published its information into Add/Remove Programs, which by default it does not do). Second, they can modify the Sms_def.mof file[3] and create a new hardware inventory class based on a registry provider for Server Administrator and Array Manager. Finally, they can use the SMS software inventory feature to capture the file versions.

Establishing a baseline for Dell server hardware

Establishing a baseline for each Dell server model is an important step prior to deploying a hardware update. With the help of the SMS 2003 ITDU, administrators can establish a baseline for Dell servers to a particular system update version, which helps simplify the update process by reducing what needs to be tracked and deployed. The specific update version depends on the environment's current needs, which can include the need to improve performance, increase server stability, or resolve known problems.

Prior to establishing a baseline, administrators should understand the current Dell support model for hardware update versions, which requires that an update be no more than six versions older than the most current update.

The steps to begin establishing a baseline are as follows:

1. Install the SMS 2003 ITDU software tools on the primary site.
2. Create an SMS 2003 collection organized by "All Dell Servers," then create subcollections by each model or hardware component. Optionally, administrators can also customize the newly created SMS collections.
3. Configure and deploy the SMS 2003 ITDU program and advertisement to occur on a schedule according to the organization's needs.
4. Verify the health of the Scan component package installation process using SMS reports and the SMS Advertisement Status Message Viewer. For more information about recommended reports for each patch management phase, see the second part of this article at www.dell.com/downloads/global/power/ps3q06-20060249-Microsoft-SOE.pdf.

To determine the existing servers in an environment, administrators can use SMS 2003 SP1 to run a discovery method, such as Active Directory System Discovery, and organize the results. After the hardware inventory runs, the results are propagated back to the SMS site database, and the Extended Software Updates hardware class is populated with a list of all Dell hardware updates that are both installed and applicable for each Dell server. The Scan Package Version hardware class indicates the name of the PDK scan package (Dell component update) and the last time the package was updated.

When establishing a baseline, administrators should also consider the following:

- **Monitoring failed deployments:** Administrators can use the SMS "Software update health" report to track failed SMS 2003 ITDU deployments across the update life cycle. The report provides a list of failures for catalog synchronization, scanning, and the patch installation agent.
- **Updating PDK files:** The PDK files used by the Scan component must always be the latest version because they are used to scan the server for hardware updates that appear in the Catalog.cab Dell Version Control Catalog file. If the latest PDK files are not available when the scan takes place, the server may not report applicable updates.

After establishing a baseline, administrators should determine which servers fall below the compliance level for particular server models. This level depends primarily on the conditions necessitating the hardware update. Administrators can use the powerful built-in SMS 2003 compliance reporting functions for Dell system updates to view which servers meet or fail the compliance level for the Dell system update along with the associated component updates (see Figure 2). *Note:* If a component update version is not present on a Dell server, it will not meet the particular system update version compliance level and thus will not be considered in compliance.

[3] For more information about modifying this file, visit www.microsoft.com/downloads/details.aspx?FamilyID=8dfa57f6-291d-4ece-8b07-50bb3ebea2ab.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

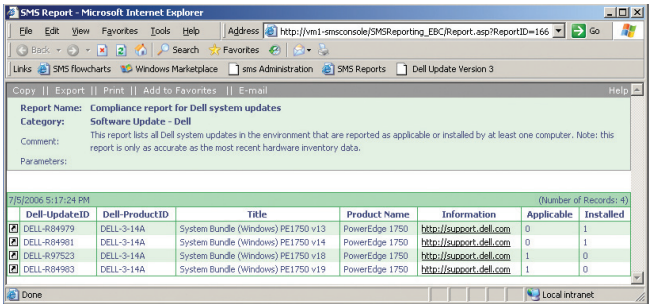

Figure 2. SMS 2003 compliance report for Dell system updates showing available update packages for a specific Dell server model

Identify phase

The Identify phase begins when a hardware update is released. Administrators may want to update Dell servers for a number of reasons, including the following:

- Correcting a security concern
- Correcting a system problem based on a Dell Support recommendation
- Updating a hardware driver or firmware component to a minimum level required by an application, such as Server Administrator
- Gaining access to a new feature or improved performance in a hardware component, such as a RAID controller
- Updating all hardware components (such as BIOS, ESM, or PERC drivers or firmware) for a particular server model as part of periodic maintenance

Activities carried out during the Identify phase include identifying patch types, assigning update severity levels, discovering update sources, identifying relevant updates, and obtaining and verifying updates.

Identifying patch types

The SMS 2003 ITDU introduces different terminology for the tool's critical components: Dell component updates, Dell system updates, Bundle Applicator, and Dell Update Packages.

Dell component updates. A Dell component update is a single hardware update, such as a BIOS, firmware, or driver update, packaged in the update package format with standardized command-line and graphical user interfaces for installation. Administrators can select component updates when running the Distribute Software Updates Wizard (DSUW), as shown in Figure 3.

Administrators should keep in mind that when they have more than one component update configured in a single package, they cannot control the installation order. Best practices recommend deploying only one component update version in a package at a time, which helps provide consistency in deployments.

Dell system updates. A Dell system update is a collection of Dell component updates configured in a single SMS package that has been tested by Dell to bring a specific Dell server model to a known and supported state. All component updates are automatically configured appropriately for deployment.

For example, a system update for a Dell PowerEdge 2600 server might include five PERC component updates and one BIOS component update. Even though the system update includes multiple PERC component updates, only the relevant and required component updates are installed on a server (for example, if a server has a PERC 4, Dual Channel integrated, only that PERC component update is installed). This automatic configuration is a key feature of SMS 2003, because it helps eliminate guesswork and associated operational complexity—administrators do not need to target patches to the appropriate servers.

Note: When configuring a system update in an SMS package, administrators cannot choose the component updates to include in the system update. If certain component updates cannot be deployed to a particular environment, the recommended procedure is to deploy a single component update in an SMS package.

Bundle Applicator. Bundle Applicator (BundleApp.exe) is a subcomponent of the PDK that installs Dell system updates. Using the Catalog.xml Dell Version Control Catalog to obtain the list of included component updates in a system update, Bundle Applicator applies all required component updates in a logical sequence.

Occasionally, a component update included in a system update may require a reboot before the next component update is installed—for example, a BIOS update may require an immediate restart. In this case, Bundle Applicator applies all component updates up to and including the component update that requires the reboot, after which SMS initiates the reboot. After the reboot, the SMS software

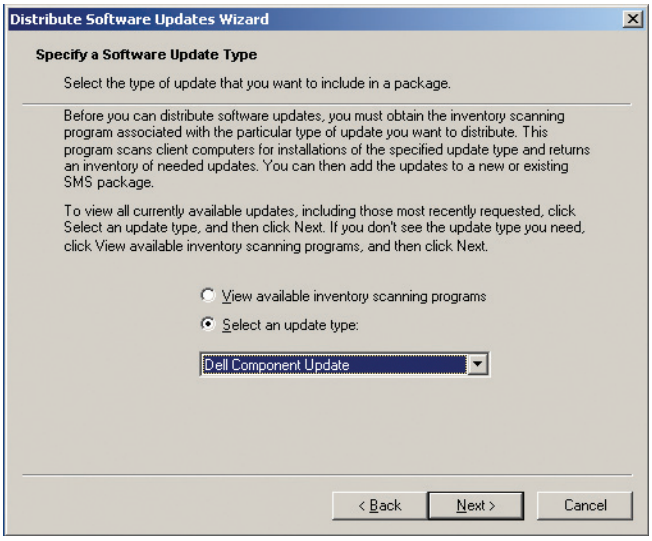

Figure 3. Dell Component Update selection in the Distribute Software Updates Wizard

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

update management feature reruns the advertisement that initiated the system update. This advertisement restarts Bundle Applicator, and Bundle Applicator applies the next component update.

Dell Update Packages. A Dell Update Package is an intelligent installer technology for any updatable element on a Dell server, such as the BIOS, a device driver, or a peripheral's firmware, that can be wrapped. All component updates use this installer technology as well.

Note: Administrators should not normally change the default settings in the DSUW, because these settings have been tested in a supported configuration by both Dell and Microsoft.

The standard Dell Update Package provides several benefits for administrators. First, it allows for a consistent method to update any Dell server system software component in the supported matrix. The package also verifies all prerequisites before the update is applied. In addition, administrators do not need to create additional media, such as a floppy disk, just to perform the update. Administrators can apply an update even if Dell OpenManage software has not been installed, and they can deploy packages remotely using SMS software distribution. Furthermore, administrators can perform an update without shutting down the system (although reboots are still required for some, such as BIOS updates).

Assigning update severity levels

Dell assigns hardware updates a severity value between 1 and 3. Level 3 severity denotes *optional* updates that contain changes affecting only certain configurations (for example, an update to Server Administrator). Level 2 severity denotes *urgent* updates that can help improve system reliability or availability (for example, a BIOS update that adds functionality or resolves a noncritical issue). Level 1 severity denotes *recommended* updates that can help keep system software current and compatible with other system modules such as firmware, BIOS, drivers, and software (for example, a BIOS update that resolves a bug that could result in loss of data, or an update to Server Administrator that addresses a security vulnerability).

Discovering update sources

Dell hardware updates are typically released quarterly. Updates can be discovered through the Dell File Watch e-mail notification service, SMS 2003 Dell update reports, SMS 2003 ITDU Scan component, or SMS 2003 ITDU Synchronization Host component.

Dell File Watch e-mail notification service. To help stay current on hardware updates, administrators can subscribe to the Dell File Watch e-mail notification service (available at support.dell.com/support/notifications/filewatch.aspx). After registering, administrators receive e-mail notifications when Dell releases hardware updates.

SMS 2003 Dell update reports. Many SMS reports can be used to determine applicable Dell hardware updates. For example, administrators can run a report to identify all Dell PowerEdge 6450 servers that fall below the recommended Dell system update baseline.

SMS 2003 ITDU Scan component. The Scan component (ITDUScan.exe) performs ongoing automated scans of servers for installed and missing Dell system and component updates. The Scan component runs Inventory Collector (Invcol.exe) first to extend the WMI class, which then populates the Win32_Extended_Patch_State WMI class with all installed and applicable hardware updates using the Catalog.xml Dell Version Control Catalog. The Dell Version Control Catalog is an important PDK component that identifies applicable and installed updates. An update for a network interface card driver, for example, would not be reported because this hardware component is currently not supported by the Dell Version Control Catalog (Dell anticipates that the content and coverage of this catalog will improve with each release).

The Scan component uses the configuration settings in such files as Scan.ini and Scanconfig.xml to scan Catalog.xml and either hold the hardware update status in WMI for a period of time set by the SMS hardware inventory site or expedite the results to the SMS site server after the scan has completed. Because updates do not report back to the SMS site database until the next SMS site hardware inventory has run, the Scan component package provides an additional program to run with the expedited switch to start the hardware inventory cycle immediately after the scan has completed. In a large SMS environment, administrators must test and carefully evaluate the impact of running a hardware inventory before selecting the expedited program, because it can cause performance degradation on the SMS site server.

The Dell Scan component is downloaded to the *X*:\windows\system32\vpcache*PackageID* directory (where *X* represents the drive where the SMS 2003 ITDU was originally installed and *PackageID* is the package ID) on all Dell servers targeted to receive the SMS 2003 ITDU advertisement.

SMS 2003 ITDU Synchronization Host component. The Synchronization Host component (ITDUSync.exe) periodically downloads the PDK component files from the Dell Web site on a schedule set by the administrator. These PDK files are then sent to the SMS distribution points and used by Dell servers to conduct scans of missing and applicable updates.

Administrators should keep two considerations in mind regarding how the synchronization host obtains its updates and the importance the synchronization host plays in the hardware update life cycle. First, after the setup process completes, the synchronization host advertisement triggers the Synchronization Host component to run on a computer specified during the setup process. The Synchronization Host component obtains the Internet location or the path of a network shared folder from the ITDUDownload.ini file. If, after setup, administrators have changed the original method of obtaining updates, they must edit each of the Source fields in the ITDUDownload.ini file.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

If administrators have full Web access from the site server, then downloading the latest PDK files by using the synchronization host should not be a problem. However, if they have decided to run the synchronization host on a system isolated from the production servers, either in a separate domain or behind a demilitarized zone, they must make additional configurations explained in the SMS 2003 ITDU administrator's guide.[4]

The second consideration is that, if the synchronization host fails for a reason that is not detected before a new hardware update scan is conducted, the latest PDK components may not be used by the Dell servers to conduct scans of installed and applicable updates, which can result in misleading and incorrect data being sent to the SMS site database. Regularly monitoring the synchronization host status is critical, particularly after the release of a new Dell Version Control Catalog. By examining the ITDUSync.log on the SMS site server or running an SMS report such as "Software update health," administrators can stay current on the synchronization host's health.

Update life cycle. Dell typically updates the PDK files quarterly, with the possibility of mid-quarter releases for updates to existing releases and hot fixes. Administrators should run the synchronization host daily to catch any updates or hot fixes relating to the SMS 2003 ITDU released before the new quarter begins. Anything that changes on the Dell update Web site is also typically changed in the Dell Version Control Catalog—that is, if the release ID is updated, the Catalog.xml Dell Version Control Catalog is also updated.

Dell update Web site. Once administrators are notified of a hardware update, they may need to visit the support.dell.com Web site to obtain more information about what the update fixes, enhances, or stabilizes. At this site, they should use the Advanced Search option so that all drivers appear for the Dell server they wish to update.

Every Dell release includes a release title explaining the Dell server model affected, a release date indicating when the update was officially released for use in a production environment, and a category naming the particular hardware component applicable for the update. Dell also provides a criticality rating for the update along with the included fixes or enhancements.

Identifying relevant updates

Identifying relevant hardware updates means screening out those that do not apply to the server models in a particular environment or addressing the issues identified in those that do apply. For example, suppose an environment includes a Dell PowerEdge 2550 server running Microsoft Exchange Server that has experienced slow disk I/O for the past two months. Although this slow I/O has not affected SLAs, it has caused concern because the server has fallen below the past year's baseline. After the Scan component has run on the Dell servers, the tool identifies an applicable RAID controller driver designed to fix a known issue similar to the performance degradation being experienced. After reviewing the information through the DSUW, the administrator determines that the update is indeed relevant for this environment.

After using the Scan component to identify updates, administrators can use the Software Updates node in the SMS Administrator console to determine whether the updates are requested by a Dell server, which is an effective way to identify relevant updates.

Obtaining and verifying updates

Administrators have several ways to obtain Dell Update Package and PDK files. This topic is discussed in more detail in the SMS 2003 ITDU administrator's guide, which is included in the software installer download.

Integrated update management

The four-stage patch management process described in this article helps administrators assess, identify, evaluate and plan, and deploy updates in IT environments. To continue reading about the remaining two phases of the process—the Evaluate and Plan phase and the Deploy phase—see Part 2 of this article at www.dell.com/downloads/global/power/ps3q06-20060249-Microsoft-SOE.pdf. Using this process with the Systems Management Server 2003 Inventory Tool for Dell Update provides a powerful way for administrators to integrate both software and hardware updates into a simplified process that can help keep servers running reliably and efficiently.

Edited with permission from Microsoft Corporation. Copyright © 2005 Microsoft Corporation. All rights reserved.

FOR MORE INFORMATION

"Applying Updates for Dell PowerEdge Servers Using Microsoft Systems Management Server 2003: Part 2." *Dell Power Solutions,* August 2006. www.dell.com/downloads/global/power/ps3q06-20060249-Microsoft-SOE.pdf

SMS 2003 Inventory Tool for Dell Update Version 3 installer download:
www.microsoft.com/smserver/downloads/2003/tools/dellupdates.mspx

[4] Administrators should also review the "Task 4: Deploy the Software Update Inventory Tools" section in the *Systems Management Server 2003 Operations Guide* (available at www.microsoft.com/technet/prodtechnol/sms/sms2003/opsguide/ops_6v3p.mspx), which provides additional guidance on configuring the synchronization host in an unattended mode with an account that has pass-through authentication through the firewall.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

PURE DELL OPENMANAGE NEWSLETTER

Flexible Management for the Scalable Enterprise

PREMIER ISSUE! August 2006

View from the Top

Reducing Data Center Complexity

Systems management headaches continue to hold many enterprises back. Faced with legions of disparate management tools, few IT administrators are able to look at the big picture—let alone fix issues arising from interdependencies. The Dell approach enables IT administrators to monitor industry-standard platforms using best-of-breed management tools from a central console.

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

A few years ago, IT departments were demanding management tools to oversee each application environment. And the industry responded with fervor—so much so that yesterday's solution has actually become one of today's biggest problems.

As each environment has evolved—and with it a mix of hardware, software, storage, and networking platforms—organizations face a mix of old and new, basic and complex, proprietary and obsolete management tools that demand a significant amount of time and resources. Furthermore, these tools typically fail to offer any meaningful insight into the health and status of the overall system, let alone provide a means to fix failures that arise from interdependencies.

Several vendors have developed huge frameworks that attempt to serve as systems management consoles and cover multiple environments, but these tools rarely include what's needed to remedy every situation. In contrast, Dell is advancing standards-based systems management options to simplify operations, improve resource utilization, and scale out cost-effectively. This approach helps enterprises leverage existing IT investments now—and it helps the industry move toward the goal of developing a single repository for management information that combines in one agent all relevant instrumentation about hardware, OS, and applications.

Today, a major Dell innovation in hardware management is helping to quicken progress: development toolkits that allow industry leaders such as Microsoft, Altiris, and Novell to integrate Dell hardware management. As a result, IT administrators running Microsoft® Systems Management Server 2003 and Microsoft Operations Manager 2005 can monitor Dell™ servers and update server BIOS, firmware, and drivers from the same management console they use to monitor and update the applications and OS. This provides a practical way to deploy systems from existing partner tools, so that Dell hardware can be integrated with best-of-breed systems management solutions to help simplify the overall IT environment.

This partnering approach by Dell is also helping the industry move toward the long-term, holistic objective based on the Common Information Model—a Distributed Management Task Force (DMTF) standard that defines how agents are represented as a common set of objects and how they interact. Dell is actively involved in creating standards and profiles—instrumentation and methodology for interacting with hardware—that are currently under consideration by several industry organizations, including the DMTF and the Storage Networking Industry Association (SNIA).

Standards such as the DMTF Web Services for Management (WS-Management) protocol and the SNIA standard data model enable IT administrators to design system elements that can be used to model systems for deployment across the entire enterprise. Industry standards such as these form the cornerstone of the Dell scalable enterprise strategy. In this way, Dell is playing a pivotal role in driving greater levels of functionality, enabling administrators to build linkages into system profiles and create methods that quickly isolate failures and then automate decisions to act upon those failures.

There is a lot to do in the next few years, but here at Dell we continue with our commitment to innovations that help ensure our hardware fits into whatever systems management tool is best for your enterprise environment. ■

INSIDE THIS ISSUE:

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Real World

How Danish Confederation Reduces IT Administration Time

The Confederation of Danish Industries saves up to 12.5 working days in IT administration thanks to Microsoft® Systems Management Server (SMS) 2003 and SMS 2003 Inventory Tool for Dell Update

> *"I am impressed at how we managed to update the BIOS, firmware, and drivers on 49 servers without a single one crashing—the whole update operation ran completely without a hitch."*
>
> *—Michael Frederiksmose*
> *Operational Engineer*
> *Dansk Industri*

Today, hardware and software are essential organizational assets. As more emphasis is placed on IT solutions to improve business processes, it is vital for enterprises to find cost-effective infrastructure management tools that help reduce operational costs.

The Confederation of Danish Industries (CoDI) is Denmark's leading private trade organization, committed to supporting manufacturing and service industries throughout the country. This means advocating policy at the local, national, and international level as well as providing consulting and information services to more than 6,400 member companies.

Managing these activities and information assets depends on a robust IT platform. As CoDI grew over the years, the organization developed an IT infrastructure based on Microsoft technology that today includes more than 50 Dell™ PowerEdge™ 1850 servers and PowerEdge 2850 servers. Manually updating the drivers, firmware, and BIOS took between one and two hours per server, which translated into as many as 12.5 working days to update the server infrastructure. Plus, servers had to be updated in a strict order so the process was prone to human error.

To gain efficiencies, in 2004 CoDI deployed Microsoft Systems Management Server (SMS) 2003, a comprehensive solution for providing software updates cost-effectively across Microsoft environments. At the same time, CoDI implemented the Microsoft SMS 2003 Inventory Tool for Dell Update (ITDU).

Developed by Dell specifically for use by software distribution tools such as SMS, the ITDU automatically updates Dell servers using SMS 2003 by placing a small software pack on each server, which checks the versions of the drivers, firmware, and BIOS installed. SMS 2003 stores this data and then accesses the Dell Web site at regular intervals to determine whether new versions are available. If so, it updates the servers automatically.

"With the ITDU, I can see which servers are not completely updated," says Michael Frederiksmose, operational engineer at Dansk Industri. "This applies whether individual components or all the components are involved." As a result, SMS 2003 and the ITDU have optimized network administration efficiency for Frederiksmose and his team.

"We managed to update the BIOS, firmware, and drivers on our servers without a single one crashing," Frederiksmose explains. "I was especially anxious about updating our Microsoft Exchange Server system and file servers, which run in clusters—but the whole update operation ran without a hitch." And because the servers are now updated automatically in accordance with each system's requirements, the ITDU has helped make the procedure safe by helping to eliminate human error. Using the ITDU, the IT team now updates servers with just a few clicks.

"It is a great asset for us that the ITDU is integrated with Microsoft tools for managing software updates," Frederiksmose says. "The screens are generally the same and the familiarity makes the tool easy to use." The speed and simplicity of automated server updates using SMS 2003 and the ITDU help provide cost-effective systems management for CoDI's growing IT infrastructure, enabling Frederiksmose and his team to spend more time focusing on core business development. As a result, CoDI can deliver information and services cost-effectively to members—helping Danish industry improve its competitive edge. ■

Challenge

Denmark's major trade organization, the Confederation of Danish Industries, needed to improve the time-consuming and error-prone manual process for updating its Dell server environment to help reduce costs and improve efficiency.

Solution

Hardware:

- Dell PowerEdge 1850 servers
- Dell PowerEdge 2850 servers

Software:

- Microsoft Systems Management Server (SMS) 2003
- Microsoft SMS 2003 Inventory Tool for Dell Update
- Microsoft Exchange Server 2003

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

TECH CORNER

By David Weber,
Enterprise Technologist, Dell Inc.

Making Systems Management Tools Work for You

Few solutions exist to manage server deployment, configuration, and updates effectively in a complex environment. Dell has worked with key partners, including Microsoft, to develop toolkits that extend existing software management capabilities to monitor and control Dell™ servers from a common interface. This innovative approach to systems management is enabling business benefits like the following representative examples recently reported from the field.

Problem: How can our IT team proactively monitor the overall health and status of the business-critical solution stack—including managed applications, operating systems, virtualization technology, servers, and data center infrastructure—while they have to grapple with different management tools that do not talk to each other?

Solution: Dell is helping to simplify operations by designing its hardware instrumentation to fit into Microsoft® Operations Manager (MOM) 2005. From a single console, MOM 2005 provides centralized monitoring, diagnostics, and problem solving. By integrating the management of software components and hardware platforms, MOM 2005 enables you to monitor and correlate information from the hardware, to the OS, up through the infrastructure layer, and out to the applications—similar to slicing through the layers of a cake.

For example, an enterprise deployed MOM 2005 with the core OS management pack, the Microsoft Active Directory® directory service management pack, the Dell server hardware management pack, and the Microsoft Exchange Server 2003 management pack. When a field representative called in saying she was unable to access Exchange, the help desk checked the management pack in MOM and saw that Exchange services were up and running. Nothing appeared to be wrong until the operator investigated what was going on at the next level—and noticed that Active Directory replication was not occurring on the server consistently. Without the ability to quickly associate Active Directory with the Exchange services, solving this connection problem could have been a long and frustrating process.

Problem: How can our IT team stay on top of the operations for a global enterprise infrastructure when they are already hard-pressed to keep up with day-to-day monitoring and control activities?

Solution: Enterprises are reporting huge time and resource savings by deploying Dell hardware management packs on MOM 2005. For example, one organization deployed MOM 2005 with the Dell toolkit across its large Microsoft Windows Server® 2003 environment to monitor all file and print servers as well as servers running Active Directory, Exchange Server 2003, and the Microsoft SQL Server™ 2005 database platform.

Before implementing MOM 2005, the enterprise required six people on the IT operations team to monitor Exchange. After implementing MOM 2005 it needed only one—freeing valuable staff resources to extend their monitoring expertise across the enterprise. ■

Benefits

- Reduces the time required to manage the Dell server environment
- Helps ensure that core technology of Dell servers is up-to-date
- Enhances insight into the health of the server environment
- Minimizes the threat of human error during server updates
- Frees IT administrators to focus on core business development

GUIDED TOUR

MOM 2005 Administrator Console

It is easy to get productive as soon as Microsoft Operations Manager (MOM) 2005 is installed. Start with this Guided Tour of the Administrator console's home page, which provides access to the following features and capabilities:

1. The Administrative Setup Tasks section links to wizards for common administrative setup tasks such as installing and uninstalling agents, managing computers, and setting global configurations.
2. The Management Pack Setup Tasks section links to wizards for management pack setup tasks, including importing and exporting management packs; managing operators (their roles, views, and configurations); and managing notification groups. For example, a MOM 2005 administrator can configure Operator console views so that a Microsoft Exchange Server administrator can work only with Exchange systems.
3. The Operations section contains links for launching the Operator console, the Reporting console, and the Web console.

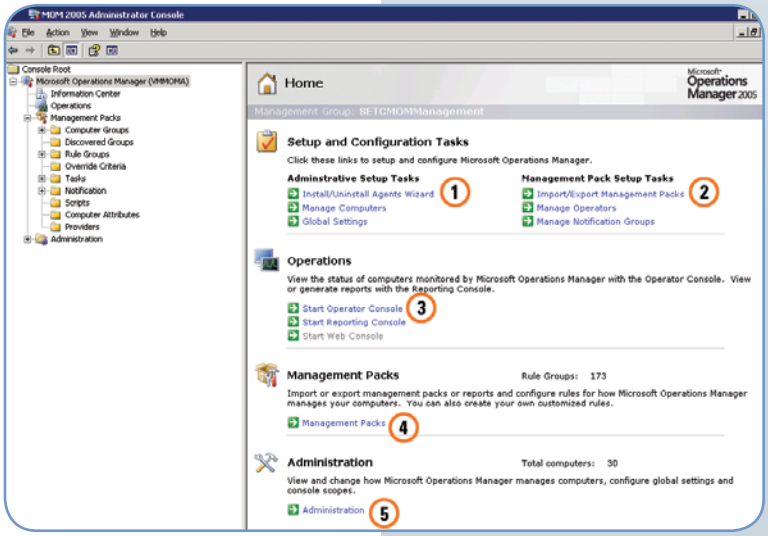

MOM Administrator console's home page

4. The Management Packs link launches the management pack details pane, which accesses advanced configuration tasks and author rules.
5. The Administration link launches the Administration details pane, from which administrators can configure views—including the definition of role-specific views and scopes, where a scope can be a specific computer or computer groups. ■

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

PARTNER INSIGHTS

Moving Closer to the Dynamic Data Center

By Brad Anderson, General Manager, Windows and Enterprise Management Division, Microsoft Corporation

Microsoft's partnership with Dell has taken a giant step forward in helping enterprises achieve tighter software and hardware integration and leading the industry in innovative systems management solutions. Earlier this year, Microsoft announced several exciting developments in the company's Dynamic Systems Initiative.

The Dynamic Systems Initiative enables IT organizations and developers to capture knowledge in the form of models to design self-managing systems and automate ongoing operations. This is an important advance toward reducing costs, simplifying systems management, and freeing staff to focus on core business development.

Microsoft continues to invest heavily in software research and development, working with key partners such as Dell to deliver end-to-end solutions that can be integrated across application development tools, operating systems, applications, hardware, virtualization technologies, and management tools.

Virtualization advances

We believe that virtualization is a major attribute of what we call the dynamic data center—a system that readily complies with business and IT policies, requires little time and effort to maintain and troubleshoot, and provides fast, flexible response to changing business requirements.

Recent progress includes the announcement of the Microsoft® System Center Virtual Machine Manager, a virtual data center management application; Microsoft plans to make the beta version available for download this summer. Together with the intended acquisition of Softricity, a maker of application virtualization and on-demand streaming software, these virtualization management solutions and current System Center products such as Microsoft Operations Manager (MOM) 2005 and Systems Management Server (SMS) 2003 provide a consistent approach to managing physical and virtual environments.

Currently, SMS 2003 is the industry leader in comprehensive change and configuration management solutions. SMS makes it easy to manage, support, and maintain a distributed network of computers. Furthermore, SMS 2003 allows you to deploy updates and patches from the hardware through the OS and out to all the applications that run on the OS across platforms from multiple partners. SMS also enables the deployment of software updates for internally developed applications.

As part of our commitment to streamline systems management, we are working with partners such as Quest (which is also a Dell partner) and Jalasoft to extend Microsoft products to manage and monitor non-Windows platforms. These partners have natively integrated and extended the System Center products to enable management of the Linux,® Sun Solaris, and IBM® AIX operating systems. As a result, network devices can be monitored by MOM 2005 or configured and updated through SMS 2003.

Over 200 partner-developed add-ons or extensions for MOM and SMS are listed on the Microsoft System Center Web site. The feedback from organizations that have implemented these solutions has been hugely positive, affirming the value of our commitment to developing strong management tools that promote cross-platform integration.

Management standards

We believe it is essential to enable enterprise-wide management information to be stored in a single repository. Developing an industry standards–based System Definition Model (SDM) is essential to the Dynamic Systems Initiative. The SDM will provide a common modeling language that enables IT organizations and developers to capture relevant organizational knowledge from distributed systems and create models based on that information. We are working with the industry to make SDM the standard for modeling knowledge in a way that can be used by any intelligent management solution, not just SMS and MOM.

At the same time, we are leveraging industry standards to provide a consistent way to manage virtual and physical environments across heterogeneous platforms. This will be accomplished using standards such as the Distributed Management Task Force Web Services for Management (WS-Management) protocol and the SDM.

Microsoft and Dell are working together to support highly scalable, highly available enterprise applications with seamless cross-platform management tools. Our strong focus on delivering value and innovation through the tight integration of standards-based data center components will continue to streamline the change-management process and help drive business success. ■

> After studying the core drivers of IT complexity, Microsoft concluded that the process of designing, deploying, and updating systems was getting out of control—taking way too much time, costing far too much money, and typically resulting in poor workload balancing across inflexible system resources. In response, the Microsoft Dynamic Systems Initiative focuses on open standards that enable flexibility of choice and centralized systems management.

August 2006

This publication is for informational purposes only, and may contain typographical errors and technical inaccuracies. The content is provided as is, without express or implied warranties of any kind. Dell, the Dell logo, Dell OpenManage, and PowerEdge are trademarks of Dell Inc. IBM is a registered trademark of IBM Corporation. Linux is a registered trademark of Linus Torvalds. Microsoft, Active Directory, SQL Server, and Windows Server are trademarks or registered trademarks of Microsoft Corporation. Other trademarks and trade names may be used in this document to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

The Dell OpenManage newsletter is published quarterly by the Dell Product Group, Dell Inc., Mail Stop RR5-03, One Dell Way, Round Rock, TX 78682, U.S.A. Copyright © 2006 Dell Inc. All rights reserved. Reproduction in any manner whatsoever without prior written permission from Dell is strictly forbidden. Information in this publication is subject to change without notice.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Extending Microsoft Operations Manager 2005 to Linux and UNIX with

Quest Management Xtensions

for MOM

Microsoft® Operations Manager (MOM) 2005 is the industry standard for systems monitoring and management in Microsoft Windows® environments. Quest® Management Xtensions for MOM is designed to natively and seamlessly extend MOM 2005 functionality to Linux- and UNIX-based systems, helping enterprises simplify operations and leverage their existing investment in MOM by integrating management of mixed OS environments from a centralized console.

BY PAUL BARCOE-WALSH

Related Categories:

Enterprise management

Linux

Microsoft Operations Manager (MOM)

Quest Software

Systems management

UNIX

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

Knowledge is power, particularly as it applies to the systems, applications, and services that can be instrumental in meeting bottom-line business goals. Regardless of the industry, the size of the enterprise, or the platform on which vital business resources operate, disruption in the availability and performance of enterprise applications can result in dissatisfied customers, lost revenue, and a competitive disadvantage—ultimately the difference between success and failure in today's fast-paced and highly competitive markets. The more knowledge an IT organization has about the current state, trends, and potential problems of business-critical applications and services, the better chance it has of maintaining the availability and performance that constituents demand.

The ability to monitor the availability and performance of mission-critical applications and remedy problem conditions proactively to avoid system downtime is essential for efficient enterprise operations. Even enterprises that dedicate the majority of their IT investment to Windows resources and Microsoft technologies typically have some investment in Linux- or UNIX-based systems. Organizations running a mixed systems environment often implement operations management on a platform-by-platform basis. Because Windows is a ubiquitous desktop and server platform, Windows-specific tools are firmly entrenched and widely deployed. In comparison, Linux- and UNIX-based systems may not have equally consistent or capable operations management tools to

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

perform the necessary monitoring and prescriptive remedies. For example, when more than one Linux or UNIX OS is deployed, organizations may find themselves using a number of different tools to perform the same task—one tool for the operations management of Windows-based systems, another for Sun Solaris–based systems, another for IBM® AIX–based systems, yet another for Hewlett-Packard HP-UX–based systems, and so on.

As an alternative, organizations may look for one comprehensive operations management tool that works across all platforms. However, true cross-platform operations management can be difficult to achieve and even more difficult to maintain. All too often, cross-platform approaches fail to provide the requisite breadth and depth. Ideally, the comprehensive capabilities of platform-specific approaches would be extended across all other platforms. Given the proliferation of Windows resources, extending best-of-breed Windows-specific operations management to Linux- and UNIX-based systems can be a cost-effective solution. For Windows environments, the industry standard is Microsoft Operations Manager (MOM) 2005. Microsoft has positioned MOM as a key component of its Dynamic Systems Initiative (DSI).

In accordance with DSI goals, MOM 2005 is designed to help IT departments quickly and efficiently determine the root cause of system disruptions and enable quick resolutions.

In accordance with DSI goals, MOM 2005 is designed to help IT departments quickly and efficiently determine the root cause of system disruptions and enable quick resolutions. MOM 2005 delivers open and scalable operations management for Microsoft Windows Server® environments along with comprehensive event management, proactive monitoring and alerting, reporting and trend analysis, and system- and application-specific knowledge. These capabilities can help improve the availability and performance of applications and services.

MOM 2005 in mixed OS environments

Microsoft has designed MOM to interoperate with monitoring tools for non-Windows-based systems. For example, the MOM Connector Framework (MCF) Web service enables administrators to achieve bidirectional alert forwarding and synchronization between MOM and monitoring tools such as IBM Tivoli or Hewlett-Packard OpenView. However, MCF does not address the underlying disadvantages of using separate operations management tools for Windows, Linux, and UNIX platforms.

Quest Software extends the capabilities of MOM 2005 to Linux- and UNIX-based systems with its Quest Management Xtensions for MOM. Quest Management Xtensions for MOM is an alternative to multiple platform-specific operations management tools for heterogeneous enterprise environments, extending the capabilities of MOM 2005 to Linux- and UNIX-based systems. Through the MOM 2005 Software Development Kit (SDK), Quest Management Xtensions for MOM provides a standards-based MOM agent for each supported Linux and UNIX platform, enabling Linux- and UNIX-based systems to look and act like MOM-monitored Windows-based systems. In this way, the monitoring, reporting, and management pack capabilities of MOM are extended to Linux (Red Hat® Enterprise Linux and Novell® SUSE® Enterprise Linux Server) and UNIX (Solaris, AIX, and HP-UX) environments.

This integration is achieved through the implementation of standards unique to each supported platform. MOM 2005 is based on the Windows Management Instrumentation (WMI) standard, which allows MOM 2005 to achieve in-depth, secure, and seamless monitoring of Windows-based systems. Unfortunately, Linux and UNIX platforms have no similar standards that would allow a tool like MOM to achieve the same level of integration and control. Inconsistent standards across Linux and UNIX platforms are typically the underlying reason why most cross-platform applications offer limited capabilities, haphazard platform support, and high overhead—and why they require an infrastructure that can result in another layer of management, expense, and complexity. Quest Management Xtensions for MOM helps overcome these disadvantages through the implementation of standards.

Quest is the primary author and maintainer of an open source Web-Based Enterprise Management (WBEM) project called OpenWBEM—which is functionally equivalent to WMI for Linux and UNIX environments, providing the same level of integration, control, and extensibility for a wide range of Linux- and UNIX-based systems. Because OpenWBEM, through Quest Management Xtensions for MOM, is integrated natively and uniquely on each OS, the supported systems appear in the MOM consoles and generate reports and alerts similarly to MOM-monitored Windows-based systems.

Another advantage of the OpenWBEM implementation in Quest Management Xtensions for MOM is its ability to preserve the unique attributes, capabilities, and personalities of Linux and UNIX environments. Quest Management Xtensions for MOM does not force Linux- and UNIX-based systems into Windows-centric operations management. On the contrary, it extends the capabilities of MOM to the world of Linux and UNIX. A Linux- or UNIX-specific management pack gathers information, generates reports and alerts, and resolves problems on Linux- and UNIX-based systems in the same way that those systems would be handled on their own. Administrators can use a single powerful and familiar tool, MOM 2005, for automated monitoring, alerting, and resolution of Linux and UNIX problems.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

IT Executive Learning Series

By IT Leaders, for IT Leaders

Executive Learning Series Summit - August 30 & 31

Dell invites you to join us in Austin, Texas, on August 30 and 31 for our 2006 IT Executive Learning Series Summit. The Summit is an IT-to-IT event hosted by Dell Information Technology featuring interactive sessions addressing the problems that IT professionals face on a daily basis, along with special keynote speeches from Dell vice president and CIO Susan Sheskey and CTO Kevin Kettler. This complimentary event provides an outstanding opportunity for you to connect with your industry peers and gain insight into several Dell IT best practices.

Set Your Own Agenda

Sessions are led by Dell IT executives and give a behind-the-scenes look at Dell's IT organization covering a range of topics including data warehousing, systems management, IT security, disaster recovery, and much more. Tours of Dell's Topfer Manufacturing Center and leading-edge Enterprise Command Center will be available. In addition, senior technologists from Dell's Enterprise Product Group will offer special in-depth sessions on the latest trends and product directions in storage and virtualization. Select from a list of sessions to customize an agenda that best fits your interests. Registration is limited, so don't delay in confirming your attendance.

What Your Peers Have Said About Dell IT Events

- "I found the series to be very informative, and more importantly, thought provoking."
- "Substantially more relevant than going to a Gartner or META conference."
- "I appreciated the IT-to-IT discussion as opposed to a sales event."

*Visit **www.dell.com/it** for registration information*

DELL™

Information Technology

Dell and the Dell logo are trademarks of Dell Inc. ©2006 Dell Inc. All rights reserved.

Quest Management Xtensions for MOM allows Linux and UNIX operations data to be gathered and stored in the same repositories and through the same channels as Windows data. Linux and UNIX management packs can be authored in, distributed through, and executed by MOM. However, these operations occur natively on or for the managed Linux- or UNIX-based system.

Quest Management Xtensions for MOM architecture

Quest Management Xtensions for MOM is designed to provide seamless Linux and UNIX operations management from within MOM 2005 using the same consoles, data stores, and reporting mechanisms that MOM uses for Windows-based systems. The product consists of a service installed on the MOM management server and an agent installed on each Linux- or UNIX-based system to be monitored. The service works as a liaison between the agent installed on the monitored system and the MOM management server. All alerts and performance data generated from the agent are passed to MOM using the OpenWBEM standard. Similar to Windows operations management, MOM deploys rules targeted for Linux- or UNIX-based systems through the agent. All rule processing for these systems occurs on the agent as well.

The Quest Management Xtensions for MOM agent runs on each managed Linux- or UNIX-based system to collect and process performance data for significant events. The agent downloads rules from the management server specifying the information the agent should collect, process, and send back to the management server, and then carries out these actions on the managed Linux- or UNIX-based system itself. The agent includes a script server used for running configuration scripts, scripts implementing Common Information Model (CIM) providers, and script responses.

The major components of the Quest Management Xtensions for MOM agent include the following (see Figure 1):

- **Update agent:** Functions as a liaison between the service on the MOM management server and the OpenWBEM CIM Object Manager (CIMOM), forwarding alerts to the service and receiving rules that are then stored in the CIMOM
- **Rule processor:** Applies rules to incoming events from four rule types: time events (specifying when an event occurred), WMI events (reporting server activity and setting alerts that enable administrators to respond quickly to system problems), WMI numeric events (collecting performance data), and a generic log file (parsing log files, handling syslog formatting, and allowing administrators to use scripting for complex parsing)
- **CIM providers:** Produce data that is turned into MOM events by WMI event providers and WMI numeric event providers in two categories: CIM instance providers (supplying CIM-class instances that can be polled periodically, functionally turning WMI numeric events into performance events for data such as load average, available disk space, and available swap space) and CIM indication providers (generating CIM indications that WMI event providers turn into events, such as syslog file events, status changes, Apache access logs, and Apache error logs)

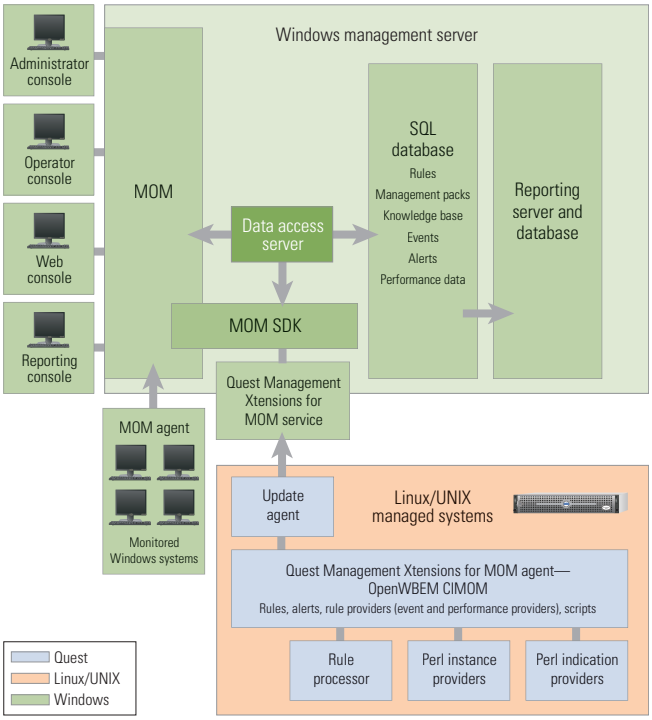

Figure 1. Overview of Quest Management Xtensions enterprise-wide monitoring from an existing MOM infrastructure

Quest Management Xtensions for MOM performance data processing and rules

Quest Management Xtensions for MOM uses the following steps to process Linux and UNIX performance data:

1. A Linux- or UNIX-based system receives rules from the providers, which it then stores in the OpenWBEM CIMOM.
2. A "subscription" is set up for the Linux- or UNIX-based system, which is then monitored by the WMI event provider.
3. The WMI event provider generates an event whenever something new happens to the specific asset being monitored—for example, when an entry is made in a log file.
4. Upon creation of an event, the CIMOM forwards the indication to the rule processor, where the event is processed against the stored rules.
5. Based on the rule, the CIMOM generates an alert or runs a script, which is stored in the CIMOM.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

6. The update agent polls the CIMOM and forwards the alert to the Quest Management Xtensions for MOM service within the MOM management server.
7. The Linux or UNIX alert is then treated and viewed in the same way MOM handles Windows alerts.

Quest Management Xtensions for MOM supports several rule types, including the following:

- **Event rules:** Define how MOM responds to specific events; these include alert-on rules (specifying that an alert has been generated), respond-on rules (generating a response), and consolidation rules (specifying that multiple similar events have been grouped into a single summary event on an agent system)
- **Alert rules:** Allow administrators to specify a response for an alert or for a number of previously defined alerts
- **Performance rules:** Define how MOM processes performance counter data and WMI numeric data; these include measuring rules (collecting numeric values from sources such as WMI, which can include a response) and threshold rules (directing MOM to generate an alert or execute a response when a WMI value or performance counter exceeds a defined threshold)

Management packs in Quest Management Xtensions for MOM

In Windows operations management, the greatest value of MOM comes from its use of management packs. A management pack can include any type of rule or set of rules specific to the performance of a given application or service. The proactive monitoring of systems for thresholds, parameters, and events unique to individual applications or services allows MOM to help IT departments improve system availability and performance.

In most heterogeneous enterprise environments, Linux- and UNIX-based systems run mission-critical services that are just as important to monitor as similar Windows services. Quest Management Xtensions for MOM extends the MOM 2005 management pack capabilities to these systems. The software includes a basic set of Linux and UNIX management packs that monitor performance data and generate events specific to CPUs, memory consumption, swap utilization and activity, storage, and network bandwidth for a wide variety of Linux and UNIX platforms. In addition, application vendors, third parties (including Quest), and administrators can use the MOM management-pack authoring capabilities to create, distribute, and run management packs for a wide range of Linux and UNIX applications and services.

As Quest Management Xtensions for MOM matures and gains widespread implementation, the available library of management packs will likely grow. Quest intends to continue authoring management packs that fit its area of expertise. In addition, partners and other third parties are encouraged to develop, improve, and contribute to management packs under a traditional open source development model. Quest also anticipates that leading Linux and UNIX software vendors and developers will author management packs in much the same way that Microsoft and its partners have created a large number of Windows-specific management packs.

MOM operations management for Linux and UNIX

Operations management allows enterprise IT departments to monitor and quickly resolve disruptions to mission-critical systems and applications. Quest Software has developed Quest Management Xtensions for MOM to extend the operations management MOM 2005 offers Windows environments to Linux and UNIX environments. It uses the same MOM infrastructure, consoles, reporting mechanism, and workflow for Linux- and UNIX-based systems as MOM 2005 does in a Windows-based system. In addition, Quest Management Xtensions for MOM ships with a basic set of management packs and allows Quest, its partners, and other third parties to author and run additional management packs using the same familiar interface and tool that makes these management packs so valuable in Windows environments—a capability that further extends the value of Quest Management Xtensions for MOM in cross-platform environments. This seamless, native, standards-based integration allows organizations to streamline operations management around a single powerful, proven tool for the entire enterprise.

Paul Barcoe-Walsh is a director of product management for systems management and monitoring at Quest Software, where he oversees product direction, strategy, and go-to-market activities for integration with monitoring programs such as Microsoft Operations Manager and Microsoft Systems Management Server. Before joining Quest, he was a key member of the Microsoft Management Business Group's product management team and worked on the beta product and processes for Systems Management Server 2003.

FOR MORE INFORMATION

Quest Management Xtensions for MOM:
www.vintela.com/products/vsm/index.php

Quest Management Xtensions for MOM Webcast:
www.quest.com/dellpower

MOM 2005:
www.microsoft.com/mom

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

LANDesk Server Manager:

Simplified Proactive Management of Dell PowerEdge Servers

LANDesk® Server Manager allows administrators to instantly assess server health and can help keep systems available and running at high performance levels. Its built-in awareness of management attributes specific to Dell™ PowerEdge™ servers and Dell Remote Access Controllers, integration with the Dell OpenManage™ suite, and ability to aggregate Dell-specific patch content can help simplify management and maintenance of Dell server environments.

BY LAURIE NYLUND

Related Categories:

Dell PowerEdge servers

Dell Remote Access Controller (DRAC)

LANDesk

Remote access controllers (RACs)

Remote management

Systems management

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

For both small businesses and large enterprises with complex network infrastructures, server health and availability drive productivity and success, and increased pressure on IT personnel to ensure server uptime means that having the tools to proactively understand, manage, and protect server health has become critical.

The award-winning[1] LANDesk Server Manager not only delivers a comprehensive arsenal of powerful yet simple-to-use tools for centralized server management in heterogeneous environments, but can also help simplify proactive management of Dell server environments. LANDesk Server Manager provides built-in awareness of management attributes specific to Dell PowerEdge servers and Dell Remote Access Controllers (DRACs), integrates with the Dell OpenManage suite, and aggregates Dell-specific patch content for review and download.

LANDesk Server Manager allows administrators to actively monitor Dell hardware and the software running on it for performance, health, operating status, and configuration changes from a centralized management console. Using data from Common Information Model (CIM), Intelligent Platform Management Interface (IPMI), custom application-specific integrated circuits, and other methods, LANDesk Server Manager provides a complete picture of PowerEdge server health and performance. Easy access to troubleshooting tools can help improve efficiency and effectiveness when finding and fixing problems. To enable administrators to proactively manage, protect, update, and maintain PowerEdge servers for optimal performance and availability, LANDesk Server Manager provides centralized change and configuration management tools—including inventory discovery, vulnerability

[1] "eWEEK Labs Picks the Top Products of 2005" by eWEEK Labs in *eWEEK,* December 19, 2005, www.eweek.com/article2/0,1895,1901684,00.asp.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

scanning, patch management, asset management, software distribution, OS imaging, Preboot Execution Environment (PXE) boot, and remote control. Through the LANDesk Server Manager customizable report-generation tools, administrators can easily access and use historical and real-time server data to facilitate trend analysis, capacity planning, and business process management.

Extensive, automated patch management

Staying current on driver and software updates and the latest vulnerabilities and patches, and deploying them to servers, should not be a full-time job. LANDesk Server Manager can help simplify and reduce the time spent on software update and patch management efforts. For example, server-specific patch management tools (see Figure 1) allow administrators to stay updated on patches and vulnerabilities while also enabling them to control what level of automated deployment and remediation makes the most sense for their environment.

To prepare administrators for new and emerging security and performance threats, LANDesk mines content servers and catalogs from trusted industry sources—including Dell—daily for the latest patches and updates. To help verify that the patches install as intended, a team of LANDesk engineers validates every patch that is made accessible from the LANDesk security database. The patch engineering team performs basic conflict-checking and patch-dependency analysis and provides installation notes to help administrators plan remediation efforts.

To provide extensive control over the management and deployment of updates specific to PowerEdge servers (such as firmware, BIOS, and driver patches), LANDesk Server Manager takes advantage of the interfaces exposed by the Dell Partner Development Kit. The vulnerability scanner in LANDesk Server Manager works in concert with the Dell inventory scanner to extract critical information from all Dell servers, such as firmware revision levels, and then matches that information against the latest updates published by Dell. LANDesk Server Manager can then provide information about the latest software updates, vulnerabilities, and patches applicable to PowerEdge servers at any time.

One of the most powerful and time-saving capabilities of LANDesk Server Manager is its ability to aggregate updates specific to the Dell servers distributed throughout an enterprise. This aggregation allows administrators to centrally manage and deploy necessary patches to all servers at once, rather than on a one-to-one basis. It also enables administrators to delay deployment until they have had a chance to sufficiently test the patch in a nonproduction environment. Once ready, administrators can schedule the patch to be deployed in a manner and time frame that best fits their organization's unique needs.

The ability to detect and aggregate patches, centrally manage them, and automate their deployment according to administrators' needs and schedules is not limited to Dell-specific updates. LANDesk Server Manager provides this same capability for all the updates and patches that PowerEdge servers require for the operating systems and software applications running on those servers.

Out-of-band management using IPMI

In spite of the best proactive management efforts, servers can still fail unexpectedly. The health dashboard and customizable health alerts and handlers in LANDesk Server Manager can provide administrators with information about server problems quickly—before end users start contacting administrators. LANDesk Server Manager also enhances administrators' ability to remotely troubleshoot any DRAC-equipped PowerEdge server that fails. Because the inventory manager in LANDesk Server Manager works with Dell OpenManage to gather DRAC-specific information, administrators can use the LANDesk Server Manager Web console to view the events, notices, and warnings recorded in the DRAC logs and DRAC trace logs without ever needing to launch a DRAC session. The console also provides direct access to remotely launch DRAC sessions, giving administrators comprehensive remote access to the DRAC functionality provided by Dell.

The management capabilities for these DRAC-equipped servers increase dramatically when administrators also use the out-of-band IPMI support included in LANDesk Server Manager. Most PowerEdge servers are equipped with more than 100 on-board sensor devices and chips. The deep IPMI support enables LANDesk Server Manager to display data from these sensors and allows administrators to quickly and remotely receive sensor alerts for all DRAC-equipped PowerEdge servers.

For example, from the LANDesk Server Manager console, administrators can use IPMI out-of-band management to check the status of the servers' IPMI-supported sensors for disconnected cables, removed hardware, chassis intrusion, and similar

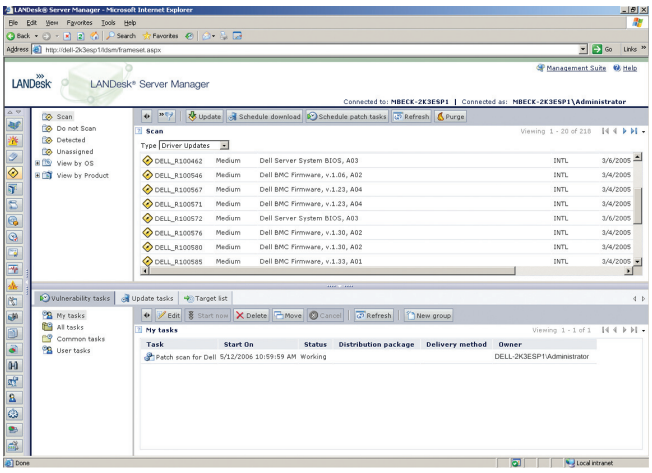

Figure 1. Viewing available Dell updates with LANDesk Server Manager

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

problems. The console provides access to the IPMI system event log, which shows a history of IPMI-related events that occurred on the device, such as the watchdog timer causing a power cycle or a general chassis intrusion. Administrators can also remotely power cycle a server, look at its BIOS settings for problems or conflicts, and make any necessary changes, as well as remotely configure the DRAC or the IPMI-compliant baseboard management controller, including settings for watchdog timers, power configuration, users, passwords, LAN configuration, and Serial Over LAN configuration.

Combining the performance monitoring, vulnerability scanning, change management, OS imaging, and patch management capabilities of LANDesk Server Administrator with its IPMI out-of-band management features enables administrators to proactively manage servers, helps reduce many of the problems that can affect server health and availability, and enhances administrators' ability to quickly bring servers back online.

Integration with Dell OpenManage

Not only does LANDesk Server Manager work with Dell OpenManage to let administrators manage DRAC sessions, but the LANDesk Server Manager Web console also provides access to the server management capabilities offered by Dell OpenManage. By virtue of this interaction, LANDesk Server Manager gathers data on all Dell OpenManage–supported components, presenting through the console a customized view of Dell-specific attributes such as the Dell service tags and Dell OpenManage versions on PowerEdge servers (see Figure 2). It also provides the ability to remotely manage and configure RAID devices.

In addition to the Dell-specific management capabilities provided by Dell OpenManage, the intuitive and easy-to-use LANDesk Server Manager Web console provides access to the comprehensive set of enterprise server management capabilities offered by LANDesk Server Manager, such as performance monitoring, asset management, vulnerability scanning, license monitoring, OS provisioning, and centralized configuration management. Because it is Web based, this console can provide anytime, anyplace remote server management, which can also help simplify and optimize IT efforts and lower total cost of ownership.

Informative reports for enhanced decision making

All administrators need intelligent tools to help them be proactive, interpret server data, and enhance overall decision-making capabilities. Through the LANDesk Server Manager console, administrators can generate a wide array of customizable and predefined reports that help facilitate decision-making efforts. By using the LANDesk Server Manager ad hoc query engine, administrators can quickly create a report to provide immediate information on any aspect of a server stored in the application's open standards database (Oracle®, Microsoft® SQL Server™, or Microsoft Data Engine databases). Administrators can even create a report that details the history of changes on any given server to facilitate configuration and change management efforts.

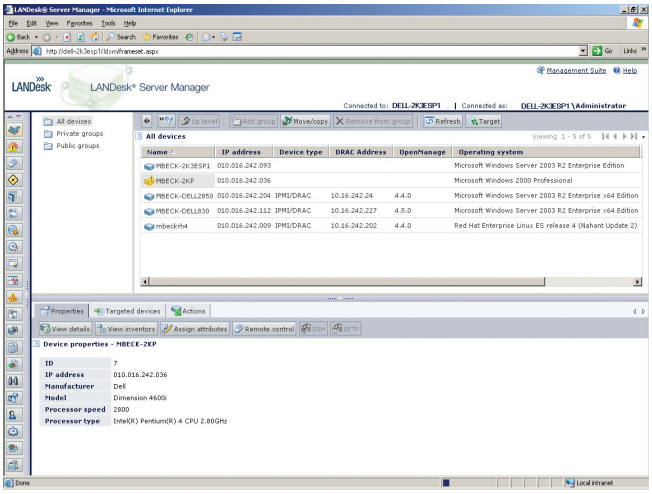

Figure 2. Viewing Dell devices with LANDesk Server Manager

Besides enabling administrators to create their own reports, LANDesk Server Manager includes a variety of predefined reports that address many common systems management situations, such as hardware refresh planning and capacity planning. Included in this extensive array of reports are Dell-specific reports that display information about all PowerEdge servers in an environment, such as firmware revision levels, Dell OpenManage versions, the presence of DRACs, and more.

Interoperability through open standards

The integration of Dell and LANDesk products can help enterprises improve server utilization, provide the data to support informed decisions, and simplify IT management efforts. In addition to creating the active integration between LANDesk Server Manager and Dell OpenManage, Dell and LANDesk work together through their shared partner alliances with industry leaders such as Microsoft, Intel, Novell, LSI Logic, and Avocent.

Dell and LANDesk also actively support open standards and nonproprietary management schemes. Their shared support for industry standards such as IPMI, CIM, Intel® Active Management Technology, Systems Management Architecture for Server Hardware (SMASH), Web Services for Management (WS-Management), and Web-Based Enterprise Management (WBEM), as well as joint participation in the Distributed Management Task Force, help advance the progress of interoperability, reliability, and performance in enterprise server environments.

Laurie Nylund is the director of engineering at LANDesk.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Enterprise RAID Management

Using the Dell OpenManage Deployment Toolkit

The Dell OpenManage™ Deployment Toolkit includes utilities, sample scripts, and other tools for server management in an enterprise environment. This article discusses using and modifying the raidcfg utility and sample script to perform efficient RAID configuration and management for multiple systems.

BY ZAIN KAZIM AND ALAN DAUGHETEE

Related Categories:

Dell OpenManage

Dell PowerEdge servers

RAID

Remote management

Scripting

System deployment

Systems management

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

RAID management in an enterprise environment with a large number of systems can be time-consuming and cumbersome. To gain maximum operational efficiency and stay competitive in the global economy, IT organizations must be able to perform routine systems management tasks like RAID configuration quickly and efficiently. The Dell OpenManage Deployment Toolkit (DTK) can help administrators achieve these goals.

The DTK provides a set of utilities, sample batch scripts, drivers, and configuration files for automated and scripted management of RAID configurations, BIOSs, baseboard management controllers (BMCs), and Dell™ Remote Access Controllers (DRACs) on supported Dell PowerEdge™ systems both efficiently and reliably. It allows administrators to perform tasks like RAID and DRAC configuration simultaneously on multiple systems, thus helping create consistency and standardization across systems management processes. In addition, the sample scripts provided by the DTK are designed to be modified to fit organizations' specific RAID management needs.

Accessing the Dell OpenManage Deployment Toolkit

Unlike other Dell OpenManage products, the DTK is not included on the Dell OpenManage CDs shipped with PowerEdge servers. Instead, it is delivered as a self-extracting zip package for the Microsoft® Windows® Preinstallation Environment (WinPE) and as an ISO image for embedded Linux® operating systems through the Dell support Web site (support.dell.com). Extracting the zip package (in WinPE) or booting the CD created from the ISO image (in Linux) creates a Toolkit folder (/opt/dell in Linux). This folder contains subdirectories with the necessary files to perform deployment tasks. The Tools directory (the bin directory in Linux) carries administrative utilities such as raidcfg and syscfg. These utilities can be used as stand-alone tools for configuring individual components or integrated into scripts for a complete one-to-many scripted mass deployment. The sample batch and shell scripts and configuration files required for a complete system deployment are in the template directory.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

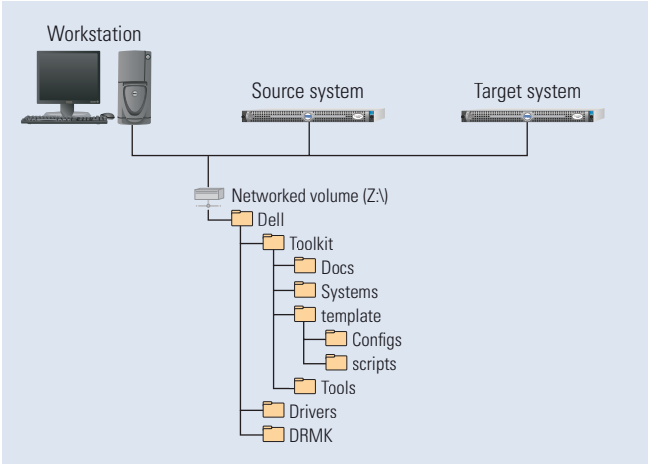

Figure 1. Dell OpenManage Deployment Toolkit directory structure for Windows

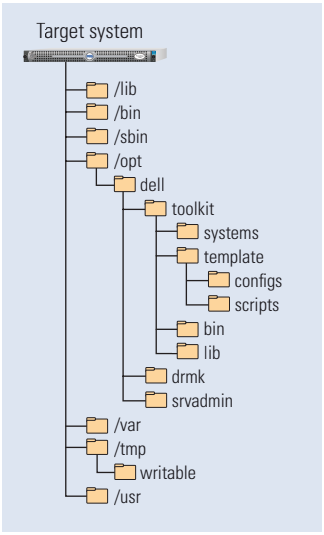

Figure 2. Dell OpenManage Deployment Toolkit directory structure for embedded Linux

The Systems directory contains platform-specific information such as configuration capture files and answer files for specific PowerEdge platforms. For details on other directories and information related to creating WinPE images, administrators should refer to the DTK user's guide. Figures 1 and 2 show the DTK file structure in Windows and embedded Linux environments, respectively.

Using the raidcfg utility and sample script

The DTK includes tools that enable administrators to create RAID configurations and hard drive partitions as well as update the BIOS, BMC, and DRAC firmware versions on supported Dell PowerEdge systems. Among these tools are the raidcfg utility and sample script, which help administrators perform RAID management tasks.

Raidcfg utility

The raidcfg utility is used to configure and report RAID-related information for supported RAID controllers. The utility abstracts any hardware differences in the RAID controllers and allows administrators to use standardized command-line interface (CLI) commands across all supported controllers on supported PowerEdge systems. The raidcfg utility supports various RAID levels (including RAID-0, RAID-1, RAID-5, RAID-10, and RAID-50) and provides switches and parameters to configure the read/write policy, stripe size, failover, and so on. The raidcfg utility performs the following tasks:

- Displays help and usage information
- Reads configuration options from CLI parameters
- Enumerates RAID controllers, physical disks, and virtual disks
- Creates and deletes virtual disks
- Logs activity to a given file name
- Returns specific error codes and messages

Figure 3 shows a representative list of valid options and arguments for the raidcfg utility.[1]

Raidcfg script

The raidcfg sample script (RAIDCFG.BAT in WinPE and raidcfg.sh in Linux) configures RAID controllers detected on the system. No administrator input is required to run this script. The default environment variable values used in this script can be set in the tkenvset script (TKENVSET.BAT in WinPE and tkenvset.sh in Linux). The raidcfg script uses the raidcfg utility to automatically configure the detected controllers on the target system. For example, for the first controller discovered, if the raidcfg utility discovers only one attached hard drive, the script creates a RAID-0 configuration; if two hard drives are discovered, the script creates a RAID-1 configuration; if three or more hard drives are discovered, the script creates a RAID-5 configuration. By default, the code to create RAID-10 and RAID-50 configurations is commented out in the script, but this code can be uncommented to create these configurations for available controllers in the system.

Option	Argument	Description
`-ac` or `action`	`cvd` or `createvdisk`	Creates a virtual disk
`-ad` or `adisk`	`ch:targ`, `ch:targ:lun`, or `ch:targ:enclosure`	Specifies an array disk
`-c` or `controllerid`	Valid controller ID number	Specifies a RAID controller
`-cp` or `cachepolicy`	`d`, `c`	Specifies the virtual disk cache policy
`-fd` or `failoverdrive`	`ch:targ`, `ch:targ:lun`, or `ch:targ:enclosure`	Specifies the location of the failover drive in an array
`-r` or `RAID`	Valid RAID level number	Sets the RAID level for the array (valid values are `0`, `1`, `5`, `10`, and `50`)

Figure 3. Sample list of valid options and arguments for the raidcfg utility

[1] For a complete list of raidcfg commands, refer to the *Dell OpenManage Deployment Toolkit Command Line Interface Reference Guide* at support.dell.com/support/edocs/software/dtk/2.1/CLI/index.htm.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

```
set DT_DRIVE=%systemdrive%
set DT_PATH=%DT_DRIVE%\Dell\Toolkit
:: * Please leave the trailing \ alone for
:: * DT_TOOLS, the code using the environment
:: * variable does not account for this.
set DT_TOOLS=%DT_PATH%\Tools\
set DT_SYSTEMS=%DT_PATH%\Systems
set DT_SCRIPTS=%DT_PATH%\Template\Scripts
set DT_STOP_ON_ERROR=FALSE
set DT_DRMK_PATH=%DT_DRIVE%\Dell\DRMK
```

Figure 4. Tkenvset script settings in WinPE

Modifying the tkenvset and raidcfg scripts for scripted RAID management

The Tools directory contains the executable utilities and sys.ini file. Sys.ini is used by syscfg to determine the system type. The template/scripts directory carries the .BAT or .sh scripts (in WinPE or Linux, respectively) needed for scripted RAID management.

The tkenvset script is where administrators define the execution environment. The default settings can normally remain unless a network share is being used to access the DTK directory structure. In that case, the `DT_DRIVE` variable should be set to the network drive letter (in WinPE) or the mount point of the shared directory (in Linux). When administrators are installing from a CD that contains the boot environment integrated with the DTK directory structure, the utilities reside on the *X*: CD drive; in this case, WinPE sets the `%systemdrive%` environment variable to this drive by default, so the environment file does not need to be changed. Figure 4 shows the tkenvset settings in WinPE.

In Linux, the default DTK directory is /opt/dell/toolkit, and the tkenvset script requires no changes unless a network share is being used. In that case, the `DT_PATH` variable should be set to the mount point of the network share. Figure 5 shows the tkenvset settings in Linux. Administrators should modify the isolinux.cfg file under the isolinux directory of the Linux ISO image prior to burning the contents on the CD. This file is called in the DTK startup phase to mount the share where the DTK contents are located.

The only other change that may be necessary is to the raidcfg script to expand the types of RAID available for use. The raidcfg script probes the available controllers on the system and loops through each one to check the number of available physical disks for use in a RAID container. Based on the number of available disks, the raidcfg script selects a RAID type (for example, RAID-0, RAID-1, or RAID-5). The script has additional RAID settings, such as RAID-10, but certain lines in the script must be uncommented to enable these settings. For example, in Figure 6 the provisions for RAID-10 in the first line have been commented out, but this line could be edited to include RAID-10 by changing it to `%%i in ( 10 5 1 0 ) do`.

```
export DT_PATH=/opt/dell/toolkit
export DT_TOOLS=$DT_PATH/bin
export DT_SYSTEMS=$DT_PATH/systems
export DT_SCRIPTS=$DT_PATH/template/scripts
export DT_DRMK_PATH=/opt/dell/drmk
```

Figure 5. Tkenvset script settings in Linux

The Linux raidcfg script has the same basic mechanism built in, as shown in Figure 7. As in WinPE, uncommenting the first line (by removing the # character at the beginning of the line) and commenting out the second line (by adding a # character at the beginning of the line) allows RAID-10 and RAID-50.

Modifying configuration files for scripted RAID management

After editing the raidcfg script to specify the RAID type, administrators may want to configure the boot OS (WinPE or embedded Linux) to automatically execute the raidcfg script when booting from a network share. For Windows installation, WinPE is the boot environment. After creating a WinPE CD (see the DTK user's guide for details on how to do this), administrators can use the winbom.ini or startnet.cmd file to map the network drive and execute the raidcfg batch file—for example, administrators can add the following lines to the startnet.cmd file:

```
Net use z: \\hostname\sharename /user:valid_username
    password
Z:\raidcfg.bat
```

For Linux deployments with a Dell-provided ISO image, administrators must copy the contents of the ISO image to a holding directory and edit the /isolinux/isolinux.cfg file to perform the same

```
:: for %%i in ( 10 5 1 0 ) do (
for %%i in ( 5 1 0 ) do (
   echo Checking if RAID %%i can be created on controller %CONT_ID% ...
   %DT_TOOLS%\raidcfg.exe controller action=createvdisk controllerid=%CONT_ID%
      adisk=%ARRAY_DISKS% raid=%%i > nul
   if NOT ERRORLEVEL 1 (
      echo RAIDCFG created RAID %%i on controller %CONT_ID%
      goto success
   )
```

Figure 6. Example WinPE raidcfg script displaying available RAID types

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

```
# for r in 50 10 5 1 0;
    for r in 5 1 0;
    do
        RAID_TYPE=$r
        echo Checking if RAID $RAID_TYPE can be created on controller $CONT_ID ...

        $DT_TOOLS/raidcfg controller action=createvdisk controllerid=$CONT_ID adisk=$ARRAY_DISKS
            raid=$RAID_TYPE $RAID50_PARAM > /dev/null
        retval=$?
        if [ $retval -eq 0 ]; then
            RAID_CFG_SUCCESS=1
            break
        fi
        RAID50_PARAM=
    done

    # Could not create a VD
    if [ $RAID_CFG_SUCCESS -ne 1 ]; then
        echo "Error: Could not configure RAID on controller $CONT_ID"
    else
        echo Created RAID $RAID_TYPE on controller $CONT_ID
    fi;
```

Figure 7. Example Linux raidcfg script displaying available RAID types

functions. The isolinux.cfg file has several example uses and need only be edited to include the appropriate names. For example, to configure the Linux image to map a Windows share and execute the raidcfg script on the share, administrators should change the `default` line to `default network-smbfs`, and then modify the line immediately following `network-smbfs` that begins with the `append` keyword. On that line, administrators should edit `share_location` to set the IP address of the target system and the share name, modify `share_script` to be the relative path to the script to be executed (that is, /opt/dell/toolkit/raidcfg.sh), and edit `share_options` to contain `username=`*`username`*`,password=`*`password`*. Next, administrators should save the isolinux.cfg file and create a new CD using the directory the ISO contents were copied to along with the edited isolinux.cfg file. For additional information on creating a CD, refer to the DTK user's guide.

Once the respective files for the WinPE and Linux images are modified, booting the CD on the target platform automatically executes the raidcfg script to configure the desired RAID containers.

Achieving efficient deployment using the Dell OpenManage Deployment Toolkit

The Dell OpenManage Deployment Toolkit enables administrators to deploy multiple Dell PowerEdge servers both efficiently and quickly to help increase server availability and minimize downtime. The DTK is both expandable and scalable, and is designed to integrate easily into existing deployment frameworks. As the need for automated deployment processes in enterprise environments continues to increase, the DTK can help system administrators continue to meet the challenge.

Zain Kazim is a test engineer in the Dell OpenManage Product Test organization. His responsibilities include quality assurance of Dell enterprise products. Zain has a B.S. in Computer Science from Michigan State University.

Alan Daughetee is an engineering technician specialist in the Dell Enterprise System Test organization. He has more than four years of test experience with Dell enterprise products.

FOR MORE INFORMATION

Dell OpenManage Deployment Toolkit User's Guide **and** ***Dell OpenManage Deployment Toolkit Command Line Interface Reference Guide:***
support.dell.com/support/edocs/software/dtk

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Using the Dell OpenManage Deployment Toolkit

with Microsoft Windows Server 2003 Automated Deployment Services

Many applications enable unattended OS installation on multiple systems using a network boot approach. But many of these options lack end-to-end OS deployment functionality. Using the Dell OpenManage™ Deployment Toolkit in conjunction with Microsoft® Windows Server® 2003 Automated Deployment Services can help IT teams deliver an automated installation.

BY GAJANAN MUDALIAR

Related Categories:

Automated Deployment Services (ADS)

Dynamic Host Configuration Protocol (DHCP)

Scripting

System deployment

Systems management

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

The Dell OpenManage Deployment Toolkit (DTK) 2.0 consists of tools and sample scripts used to configure virtual disks, system BIOS information, and Dell™ server components such as the baseboard management controller and Dell Remote Access Controller (DRAC). These commands can be used locally on the system or with remote deployment applications such as Microsoft Windows Server 2003 Automated Deployment Services (ADS) 1.1. Administrators can use the DTK in conjunction with ADS as an end-to-end deployment tool for Dell PowerEdge™ servers. This article discusses setting up the proper environment for automated OS deployment and configuring ADS for such a deployment.

The DTK tools are located in the G:\Dell\Toolkit\Tools folder, which is governed by the `DT_TOOLS` environment variable, and the scripts are located in the G:\Dell\Toolkit\Template\Scripts folder. The scripts use the tools to automate the end-to-end deployment. These scripts are meant to be modified according to an administrator's needs.

Understanding ADS

ADS comprises the Controller Service, Network Boot Service (NBS), and Image Distribution Service (IDS); the Volume Imaging Tools; and a set of agents. These features help administrators deploy and administer large numbers of servers running a Microsoft Windows® OS. This article focuses on the Controller Service, NBS, and agents.

The Controller Service governs all activity in the ADS environment. NBS, in conjunction with a network Dynamic Host Configuration Protocol (DHCP) server, provides ADS with boot command capability. Preboot Execution Environment (PXE)–enabled devices use DHCP to discover the PXE service, and then use NBS to begin a communication session with ADS. The system on which the OS is to be deployed is represented as a device in the ADS Microsoft Management Console (MMC). Each device is controlled by the Microsoft Windows Preinstallation Environment (WinPE) agent, which is installed locally on the server. The set of agents provides pre- and post-deployment management capabilities.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Setting up the deployment environment

The DHCP server can be configured to be the same server as the ADS server. Deployment of ADS consists of the Controller Service, NBS, and IDS. In a subnetted or virtual LAN (VLAN) environment, the DHCP broadcast scope is usually constrained to the IP subnet or VLAN of the attached network. This scope is the same for PXE broadcasts. In these environments, administrators must configure DHCP forwarding at the routing points to bring the devices back to one DHCP and NBS server. For environments that cannot provide DHCP forwarding in the routers, Microsoft provides a DHCP forwarding service in its server operating systems to perform this function.

Setting up ADS

Prior to the installation of ADS on a server, an instance of Microsoft SQL Server™ software should be running on the server. If SQL Server 2000 is not installed on the server, it should be installed with the Service Pack 3 update. SQL Server 2000 maintains the ADS database. If the deployment is performed for initial testing purposes, then ADS should be deployed in a private VLAN that does not interfere with the production network. If a PXE setup is already deployed on the network, then ADS should not be deployed on the same network, because doing so may cause problems at later stages.

To start the ADS MMC, administrators should enter "ads.msc" in the Run dialog box. The ADS MMC is used as an administrative tool to configure all devices and deploy jobs to individual devices.

Creating the ADS repository

The Deployment Agent Builder service, which is a component of NBS, uses a copy of OS files from the WinPE CD to create Deployment Agent images. The Deployment Agent Builder service uses the files in this directory to populate the Deployment Agent Builder repository, *systemdrive*:\Program Files\Microsoft ADS\nbs\repository\Windows. The following commands can be used to respectively list, add, and delete the repository:

```
adsrepository -l
adsrepository -a repository_name path
adsrepository -d repository_name
```

Adding a device

A device can be added and manipulated using either a command-line interface (CLI) or a graphical user interface (GUI), but adding a device from the MMC GUI is simpler and more convenient than using the CLI. To add a device, administrators should retrieve the server network interface card's Media Access Control (MAC) address from the server's BIOS setup. This MAC address should be manually updated in the inventory of the ADS database.

After a system is added as a device in the ADS MMC, its status is displayed as "Disconnected" and the control is displayed as "No." To deploy a job to the server, administrators can right-click on the particular device and select "Take Control" from the drop-down menu. The other way to add a device to the ADS repository from a CLI is by issuing the following commands:

```
adsdevice -a device_name /adminmac 12_digit_mac_
   address
adsdevice -l
```

Device variables can be added to the device to use them in a job. After adding a device in the ADS repository, administrators should assign a default template for the device to run.

Using ADS and the DTK for an unattended deployment of Windows Server 2003

The following components can help administrators with unattended OS deployments:

- **DHCP server:** Server that assigns IP addresses to the clients
- **ADS server:** Server that works as a centralized deployment server
- **Dell PowerEdge clients:** Servers on which the OS needs to be deployed; these servers should be configured to boot through the network
- **Sequence editor:** Tool provided by the ADS server; the ADS server issues jobs to the clients that consist of individual tasks sequenced together and stored as XML files on the ADS server
- **Network share:** Scripts and OS dumps that act as a centralized share

For an unattended OS deployment on a bare-metal system, administrators should perform the following tasks:

1. Configure the system-specific configuration in the system BIOS.
2. Configure the DRAC component.
3. Configure a single virtual disk on the first RAID controller (by default, the integrated Dell PowerEdge Expandable RAID Controller).
4. Create a partition (for example, a Dell utility partition) and map the system drive to the system partition.
5. Install the OS on that valid partition.

Carrying out these steps using the DTK can be done as follows:

1. Use the syscap.bat script to capture the system configuration and the sysrep.bat script to replicate the configuration, provided that a configuration file with appropriate settings is passed.
2. Use the raccap.bat and racrep.bat scripts to do the same for the DRAC component.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

3. Use the raidcfg.bat script to configure the RAID controller. After this configuration, reboot the system—otherwise WinPE cannot recognize the created virtual disk.
4. After the reboot, use the partcfg.bat script to create the Dell utility partition and format the data partition. Make sure that the utility partition bin file (upinit.bin) is placed in the path set in %DT_SYSTEM%\%DT_PLATFORM%\upinit.bin. For example, for a PowerEdge 2800 server, the path would be G:\Dell\Toolkit\Systems\PE2800\upinit.bin.

Steps 1 and 2 are optional, and specific to the vendor deployment environment. These steps can be performed using the sample scripts (syscap.bat, sysrep.bat, raccap.bat, and racrep.bat) provided with the DTK. These scripts can be deployed as an ADS job from the ADS server.

Creating a single ADS job

Administrators can use any XML editor to create an ADS job, but ADS does provide an XML editor for this purpose. To open the sequence editor, administrators should run seqed.exe from the Run dialog box. Figure 1 shows the sequence editor console with the available commands displayed.

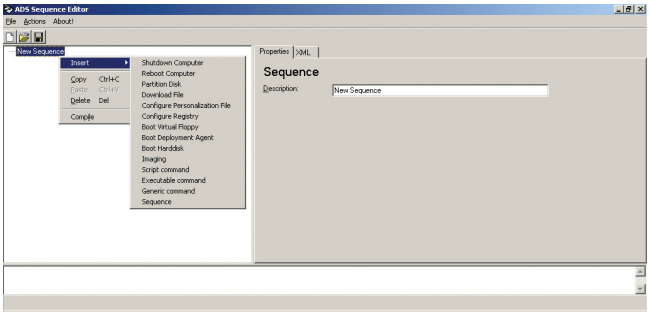

Figure 1. ADS sequence editor console

During the deployment life cycle, administrators may frequently use the script command job and the executable command job. They can modify the DTK sample scripts as needed and use the script command job to map each script to a single script command job.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

The easiest PC recovery tool on the market

Works even when Windows® won't boot

Recover Pro™ 6 can completely restore your PC immediately after a major operating system failure. Because your data and the recovery mechanism are stored in a secure area of your hard drive, Recover Pro 6 is immune to operating system failures. No recovery CDs, network access, or IT involvement are required. Simply press one button and you're back in business in minutes. Recover Pro 6 complements your anti-virus software to help ensure your business continuity even in the event of a virus attack or operating system failure. To order Recover Pro 6 contact Dell at: 1.800.WWW.Dell (800.999.3355) SKU#AX-0044XX6T-E-BOX or www.phoenix.com/en/Products

phoenix technologies

Secure from the START

To create a single script command job, the following conditions are necessary:

- The script file to be mapped to the script command in the ADS sequence editor should be in a common location or a network share.
- This network share should be mapped using the `net use` command as a local drive in the client computer—for example:

```
net use drive_letter: \\ip or hostname\share_name
    password /user:username
```

  This command can be included in the i386\system32\startnet.cmd file to automate this step. The command is executed whenever the client boots to WinPE.
- The *drive_letter* value should be reflected in the `DT_DRIVE` environment variable in the tkenvset.bat file (for example, throughout this article the `DT_DRIVE` is G:).
- When the DTK files are copied to the network share, the hierarchy of the directories should be maintained. This hierarchy plays an important role in the environment variable assignment in the tkenvset.bat file, which stores all the variable values.

To create an ADS sequence script, administrators can take the following steps:

1. From the ADS sequence editor, create a new sequence.
2. Give an appropriate name to the new sequence.
3. Insert a script command in the sequence.
4. Provide the full path of the script file in the Command field.
5. If necessary, provide a parameter to pass to the script in the Parameter field (this step typically is not required).
6. Run the raidcfg.bat script, provided in the G:\Dell\Toolkit\Template\Scripts folder.
7. Run a reboot command on the client so that the OS can recognize the virtual disk.
8. To execute the partcfg.bat script, set the upinit.bin file to the appropriate directory path. For this script to function, two modifications are necessary. First, create a platform-specific folder for the upinit.bin file (for example, in a PowerEdge 2800 server, the file should be in G:\Dell\Toolkit\Systems\PE2800). Second, set the appropriate path for the unzip.exe utility (typically G:\Dell\Toolkit\Tools) in the partcfg.bat file.
9. Pass the parameter for the type of OS to be installed into the Wininst.bat script. For Windows Server 2003, this parameter is `NET`; for Windows 2000, it is `W2K`.
10. Place the Unattend.txt file in the G:\Dell\Toolkit\Systems\PE2800 folder.

For more information about subsequent steps of an unattended installation, see technet2.microsoft.com/WindowsServer/en/Library/ef4f2523-df0a-4a54-a02b-8bb4dc4461b91033.mspx.

Avoiding deployment pitfalls

Administrators should take note of the following deployment guidelines:

- Scripts deployed from the ADS server cannot set an environment on the client server. The environment variables should be set in the scripts and used internally in the scripts themselves.
- Parameters cannot be passed dynamically from one script to another during runtime.
- The upinit.bin file varies for each server. This file is available on the Dell OpenManage Server Assistant CD.
- If the virtual CD is enabled, the WinPE OS may assign the first created virtual disk as Disk1 and not Disk0; in this case, the `DT_HD` variable should be changed to 1, because by default its value is 0.
- The `NET` and `W2K` parameters passed in the Wininst.bat file are case sensitive.
- If the Wininst.bat script fails in execution, administrators should clear the system partition before reexecuting the script. To do so, they can rerun first the partcfg.bat script and then the updated Wininst.bat script.

Combining tools for automated deployment

The Dell OpenManage Deployment Toolkit contains a useful set of tools for configuring Dell PowerEdge servers and server components. These tools, used in conjunction with Microsoft Windows Server 2003 Automated Deployment Services, can help administrators automate and simplify end-to-end OS deployment for large numbers of servers.

Gajanan Mudaliar is an engineering analyst on the Tape Storage Engineering team in the Dell Product Group. His primary area of expertise is testing and troubleshooting tape hardware and automation devices. Prior to joining Dell, Gajanan had seven years of IT experience in technical support and system administration for enterprise storage installations. He has a B.E. in Electronics and Telecommunication from Shivaji University in Kolhapur, India.

FOR MORE INFORMATION

Windows 2000 default Unattend.txt file:
www.microsoft.com/resources/documentation/Windows/2000/server/reskit/en-us/deploy/dghn_ans_meuv.asp

ADS Administrator's Guide:
www.microsoft.com/resources/documentation/WindowsServ/2003/all/ads/en-us/admin_top.asp

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Best Practices: Measuring the Success of Enterprise Testing

To determine whether an enterprise solution is ready to deploy in production, enterprise IT organizations must perform substantial testing before deployment. Testing requires extensive planning and numerous resources, but enterprises should also determine whether the investment in testing is paying off.

BY CYNTHIA LOVIN AND TONY YAPTANGCO

Related Categories:

Enterprise management

Enterprise testing

Performance

Planning

Scalable enterprise

System deployment

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

Determining the success of a product development and test effort is a matter of ensuring the organization has delivered a product that meets its customers' quality expectations.[1] Of course, an organization cannot wait for its customers to provide feedback as to whether it has achieved this goal, especially if it has not. Rather, an organization must measure several dimensions of the development and testing to determine whether it has been successful in meeting quality goals.

Developing a measurement strategy

A measurement strategy needs to address four key aspects:

- **Process:** Metrics for the effectiveness and efficiency of the test process, which address the consistency of test efforts either by comparing current efforts with past performance or by measuring the consistency of test activities for different teams within an organization
- **Quality:** Metrics that address the quality of the entity being tested, typically addressing both the volume and the severity of defects
- **Schedule:** Metrics that assess the overall schedule adherence to the test plan or that compare the duration per test phase against past programs
- **Cost:** Metrics that measure investments in a particular test project

A good approach to measurement is provided by the Six Sigma framework, which describes four families of measures. These families are known as the PFQT measures, which detail the productivity (process), financials (cost), quality, and timeliness (schedule) of a project.[2] An effective measurement strategy includes metrics that address these four areas.

Building a measurement infrastructure

Prior to implementing a measurement strategy, organizations must create a measurement infrastructure to track the number and type of defects reported during the development and test process as well as the number of resources allocated to test activities. This infrastructure also should set the goals and standards for the development and test organizations.

[1] This article is part of a series of articles examining best practices in enterprise testing. For previous articles, visit www.dell.com/powersolutions.

[2] For more information, see "Data Management Plans Can Reduce Project Cycle Times" by David Wetzel, *iSixSigma,* March 9, 2005, www.isixsigma.com/library/content/c050309b.asp.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Defect-tracking system

A defect-tracking system can be as simple as a local database or as sophisticated as one of the many defect-tracking tools available today. A key characteristic of a defect-tracking system is the ability to categorize defects by problem area, state, and severity. The problem area typically can be associated with the development team responsible for the component. The state of a defect is its current phase in the defect life cycle. Typically defect phases include *created*, *fixed*, and *closed*, but they can also include other states such as *assigned* (assigned to a specific developer for resolution) or *rejected* (determined not to be a valid defect). For severity, defect-tracking systems may use a scheme of three to five defect types. For example:

- **High:** These are referred to as severity 1 defects in many defect-tracking tools. These critical defects cause serious problems at customer sites and must be fixed prior to the product being shipped.
- **Medium:** Known as severity 2 defects in many defect-tracking tools, these defects can cause some part of a product feature set to deviate from requirements. In some instances, these defects are specific to a configuration or to a specific set of user actions.
- **Low:** These are referred to as severity 3 defects in many defect-tracking tools. These defects generally do not affect the functionality of the product but may give a poor perception of the rigor of the development and test process—for example, misspelled words or inconsistencies between the user interface and the user guide. Many organizations choose not to fix all of these defects, although a high volume of these types of defects can cause the perception of a poor-quality product.

Metric	PFQT family	Description
Defect management	Quality	Trend of defects found and fixed per week; indicator of defect volume and required capacity to resolve outstanding defects
Defect arrival rate	Quality	Pareto distribution of incoming defects; may indicate complexity of defects to be fixed and possible feature hot spots
Defect incoming/ closure gap	Quality	Indicates how much change is going to occur on the product; may be used to support decisions about increasing or decreasing test coverage
Defect aging	Quality	Pareto distribution of outstanding defects by severity classified by < 7 days, < 14 days, < 30 days, and > 30 days; another measure of organizational capacity and defect complexity
Test execution progress	Timeliness	Test cases executed compared to plan; pass rate can be correlated to number of defects to be found and focus areas for regression testing

Figure 1. Common operational metrics

Resource-tracking system

A resource-tracking system is necessary to measure the number of hours spent on test activities. Useful metrics include the number of total hours spent on the project (for planning, execution, and reporting), the number of manual or automated test hours, and the number of hours spent during each test phase.

Goals and standards

Goals and standards are the benchmarks in a measurement system. These may be defined as a percentage improvement over historical programs, or they may be based on an industry source or a best practice from another organization doing similar work.

Identifying key measurements

A measurement system needs to provide metrics in three categories: operational, project assessment, and business.

Operational metrics

Operational metrics are designed to measure the project progress during the test execution phase—whether that phase is unit, product, or system testing. Organizations may use these metrics to assess the stability of the product at a given point in time, release readiness, and adherence to test entrance and exit criteria. Operational metrics are typically analyzed weekly during the test cycle to enable the team to make adjustments to the plan as required. Figure 1 describes common operational metrics.

Project assessment metrics

A key objective of any organization is to execute the test cycle for a given project according to the test plan. Project assessment metrics are designed to provide the organization with feedback on the efficiency of test planning, execution, and reporting efforts. These metrics can help increase the accuracy of planning efforts as well as ensure that comparable resources are being utilized for similar programs across the organization. Project assessment metrics should be produced shortly after the end of the test project, and may be useful in helping test and development teams drive improvements. Figure 2 describes common project assessment metrics.

Business metrics

Ongoing improvement is a requirement for any organization. Business metrics address whether an organization is executing tests

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Metric	PFQT family	Description
Mean time to close	Productivity	Average time required to close a defect after it has been discovered
Test plan effectiveness	Productivity	Number of defects discovered per 100 hours of test case execution
Invalid defect rate	Productivity	Number of invalid defects expressed as a percentage of all defects; typical reasons a defect would be considered invalid are "working as designed," "already reported," "invalid configuration," and "cannot duplicate"
Adherence to resource plan	Financials	Variance by percentage in overall labor costs as compared with the test plan
Defect incoming rate by test phase	Quality	Percentage of defects found by test phase (unit test, product test, and system test); a best practice is discovery of 60% of defects during unit test, 35% during product test, and 5% during system test
Defect closure rate by test phase	Quality	Percentage of defects found by test phase (unit test, product test, and system test); defect closure should not lag significantly behind defect discovery
Volatility index	Timeliness	Number of scope changes introduced during the test cycle, including an impact analysis
Schedule adherence	Timeliness	Variance of calendar days from beginning of unit testing to end of system testing as compared with the test plan

Figure 2. Common project assessment metrics

Metric	PFQT family	Description
Defect discovery accuracy	Productivity	Number of defects discovered per month for all projects by percentage of valid versus invalid
Test plan effectiveness	Productivity	Number of defects discovered per 100 hours of test case execution as compared to historically similar programs
Test case execution profile	Productivity	Number of overall tests executed, percentage of manual versus automated tests, and defect discovery by each type
Adherence to resource plan	Financials	Variance by percentage in overall labor costs as compared with historically similar programs
Adherence to capital plan	Financials	Variance by percentage in overall capital costs as compared with historically similar programs
Field escapes	Quality	Number of customer support escalations after field release, including pareto distribution of reasons for escape
Defect incoming rate by test phase	Quality	Comparison of the incoming rate of defects by test phase as compared with historically similar programs
Defect closure rate by test phase	Quality	Comparison of the closure rate of defects by test phase as compared with historically similar programs
Schedule adherence	Timeliness	Variance of calendar days from beginning of unit testing to end of system testing as compared with historically similar programs

Figure 3. Common business metrics

with consistency and improving in its test efforts. Business metrics compare test planning, execution, and reporting efforts with historical efforts of a similar type. They are also used to assess consistency of test activities within the same team, among different teams with the same mission, and among various regions of a global organization. Business metrics are typically reviewed monthly or quarterly. Figure 3 describes common business metrics.

Gauging the success of test efforts

The primary goal of measuring the development and test process is to help ensure that enterprises are meeting the quality expectations of customers. Quality measures indicate whether a product is performing to the specifications defined by the marketing team and implemented by the product development group.

No one measure can be used to determine the quality level of a product. It may be tempting to continually introduce more and more metrics to assist in the management team's analysis, but a succinct list of measures can be used effectively. Of course, customers will always provide the final verdict as to whether an enterprise has met its quality goals and whether its measurement system needs tuning.

Measurement systems address key aspects of product development and testing, including consistency, effectiveness, and efficiency. A well-defined measurement system can help organizations deliver high-quality products and provide a basis for continual process improvement.

Cynthia Lovin is a senior consultant test engineer in the Dell Product Group Global Test Department. She has Six Sigma Green Belt certification and a B.A. in Business Administration from The University of Texas at Austin.

Tony Yaptangco is the director of the system test group within the Dell Product Group Global Test Department. Tony has a B.S. in Computer Science from San Diego State University and an M.S. in Engineering Management from National Technological University.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Advanced Configuration Options in

Microsoft SQL Server 2005

The ability to effectively manage application and database performance is critical for IT organizations. The Microsoft® SQL Server™ 2005 database platform includes an array of configuration settings that can help improve resource utilization and optimize system performance. This article discusses some of the advanced configuration options in SQL Server 2005 and their associated best practices and limitations.

BY ANANDA SANKARAN AND SHABANA M.

Related Categories:

Characterization

Database

Microsoft SQL Server 2005

Microsoft Windows

Performance

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

Microsoft SQL Server 2005 provides advanced configuration options that can significantly affect system performance. These instance-wide options—which can be set through SQL Server Management Studio or the sp_configure stored procedure—control various aspects of database operation and performance. With the appropriate settings for these options, enterprise IT organizations can help ensure efficient resource utilization and high performance for their SQL Server 2005 environments.

Setting the recovery interval parameter

When an instance of Microsoft SQL Server starts, all transactions at the time the instance was previously stopped need to be rolled back or forward based on their commit status. The recovery interval option specifies the maximum time (in minutes) for performing recovery operations and controls the frequency of checkpoints issued in each database. During each checkpoint, SQL Server flushes all log information and modified data pages from the buffer cache to the disk. A checkpoint is issued whenever the number of data modifications made after the last checkpoint reaches a certain limit estimated by SQL Server (based on the ability of SQL Server to perform recovery operations within the specified recovery interval) or whenever the log becomes 70 percent full. The default value of this parameter is 0 (indicating automatic configuration), which usually means a recovery time of less than one minute and a checkpoint approximately every minute for active databases. This parameter does not take into account the time required for undoing long-running transactions that span several checkpoints.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Server	Dell™ PowerEdge™ 2800 server with two Intel® Xeon® processors at 3.6 GHz, 2 MB L2 cache, and 8 GB of RAM
Storage	Dell/EMC CX300 Fibre Channel storage system with one disk array enclosure
RAID configuration	Two 10-spindle RAID-10 volumes for database files and one 4-spindle RAID-10 volume for the log, with 15,000 rpm Fibre Channel disks
Storage interconnect	Brocade SilkWorm 4100 Fibre Channel switch and QLogic QLE2360 Fibre Channel host bus adapter
Network interconnect	Dell PowerConnect™ 5324 switch
Client servers	Eight Dell PowerEdge 750 servers
OS	Microsoft Windows Server 2003 Enterprise x64 Edition with Service Pack 1
Application	SQL Server 2005 Enterprise Edition (x64)

Figure 1. Test server configuration

In December 2005, a team of Dell engineers conducted stress tests to measure the impact of changing the recovery interval parameter on database performance. The server configuration used for the tests is described in Figure 1. A generic transaction-based workload from a set of client servers was applied on a SQL Server 2005 deployment at recovery interval values of 0, 5, 10, and 15. OS and database performance counters were collected during a steady period when the load was constant and the same for each test. Figures 2 and 3 show the variation of the significant counter values. In Figure 2, the checkpoint pages per second metric reflects the number of checkpoints issued by the server, and disk transfers per second reflects the total number of reads and writes issued to the logical disks. As the recovery interval was increased from 0 to 15, the number of checkpoints issued decreased, and hence the disk transfers decreased as well. Figure 3 shows a corresponding decrease in disk-transfer response time, suggesting an improvement in I/O performance.

Normally, the default value of 0 is the recommended recovery interval setting. However, in certain deployments system performance may be affected because checkpoints are issued too frequently. In such cases, administrators can increase the value in small increments, which reduces checkpoint frequency and can help improve system I/O performance. This improvement is, however, obtained at the expense of database recovery time during a restart, and the recovery interval should never be changed to a large value that increases recovery time to unacceptable levels. Administrators should thoroughly assess the trade-off between performance and recoverability before changing this parameter from the default value. In failover clustering configurations, the recovery interval should be left at the default value, because large values may lead to long failover times.

Configuring thread scheduling

The SQL Server 2005 affinity mask options assign specific threads to processors, reducing both thread migration across processors and context switching. SQL Server 2005 includes two affinity mask options: affinity mask (also known as CPU affinity mask) and affinity I/O mask. In most cases, the OS default affinity provides the optimal performance. However, the affinity mask option can be useful for resource allocation when multiple instances of SQL Server are hosted on the same system.

Two other configuration settings related to thread scheduling are the max worker threads and lightweight pooling options. Each SQL Server instance maintains a pool of threads or fibers for its use. A fiber is a lightweight thread that requires fewer resources than a thread and can switch context in user mode. The max worker threads option controls the maximum size of a pool of threads, and

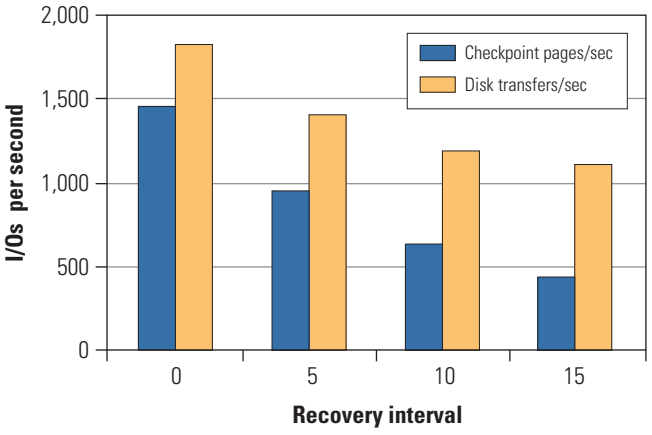

Figure 2. I/O variation with different recovery intervals

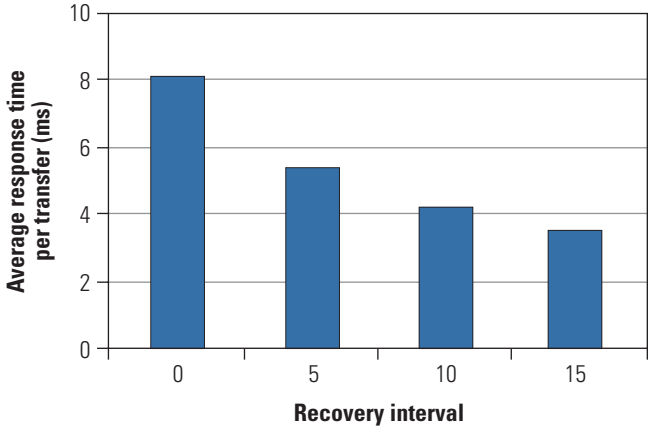

Figure 3. I/O response-time variation with different recovery intervals

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

the lightweight pooling option controls whether an instance of SQL Server uses threads or fibers.

Enabling large-memory support

SQL Server 2005 can provide support for up to 1 TB of RAM with the appropriate 64-bit edition of both SQL Server 2005 and the Microsoft Windows Server® 2003 OS. On 32-bit Microsoft Windows® operating systems, SQL Server 2005 provides large-memory support using Address Windowing Extensions (AWE). AWE is a set of extensions to the Windows management functions that allows applications to acquire physical memory directly and map views of physical memory to their virtual address spaces, enabling physical memory usage beyond 4 GB and up to 64 GB. Each Windows process has its own virtual address space, the size of which depends on the architecture (32-bit or 64-bit). A 32-bit process can map only up to 4 GB; only 2 GB of this space (3 GB with the `/3GB` boot option) is made available to the application, and the rest is reserved for the OS.

SQL Server can directly acquire up to 64 GB of physical memory or the maximum supported by the OS as nonpaged memory using AWE (see Figure 4). AWE support was included with SQL Server 2000 on 32-bit platforms, but SQL Server 2005 enhances the AWE support by dynamically allocating and managing AWE-mapped memory to balance overall system memory needs when deployed on Windows Server 2003. The following should be taken into account when enabling AWE with SQL Server 2005:

- AWE support is provided only with SQL Server 2005 Enterprise Edition and Developer Edition.
- A maximum of 64 GB of physical memory is supported through AWE; if 3 GB of virtual address space is configured for SQL Server (using the `/3GB` switch), then the overall physical memory support is limited to 16 GB.
- Physical Address Extension (PAE) support for accessing physical memory greater than 4 GB is provided only on Windows 2000 Advanced Server and Datacenter Server and on Windows Server 2003 Enterprise Edition (32-bit) and Datacenter Edition (32-bit).
- Only database pages can reside in the physical memory allocated through AWE. Memory allocated through AWE cannot be used for supporting additional database users, threads, databases, queries, and other objects that permanently reside in the virtual address space.
- SQL Server 2005 Analysis Services cannot take advantage of AWE-mapped memory.

AWE requires the database engine to run under a Windows account with "Lock pages in memory" permission. Administrators can further control memory allocation using the min server memory and max server memory options, which enforce upper and lower limits on SQL Server memory utilization. The SQL Server 2005 AWE mechanism differs when used with Windows 2000 Server as compared with Windows Server 2003. On Windows 2000 Server, AWE memory allocation is static—locked during startup and not released until shutdown. If memory allocation is not controlled using the max server memory setting, all available physical memory except 128 MB is used. On Windows Server 2003, AWE memory is allocated dynamically and returned to the system as required to balance the overall system needs. The max server memory option must be set to guarantee additional memory for other system applications or SQL Server instances. In failover clustering scenarios, all server nodes must have adequate physical memory to handle the memory requirements of all SQL Server instances for both before and after failover load conditions.

AWE is not necessary on 64-bit versions of SQL Server 2005 because virtual address space is not limited to 4 GB. However, the Windows account under which the database engine is running should still have "Lock pages in memory" permission. This setting enables SQL Server to hold on to the allocated memory and prevents the OS from paging the allocated memory, which can help provide performance improvements. The dynamic memory management mechanism releases locked memory appropriately if a system-wide resource shortage occurs. A 64-bit edition of SQL Server is better suited than a 32-bit edition for supporting heavy analytical processing workloads because of the large amount of

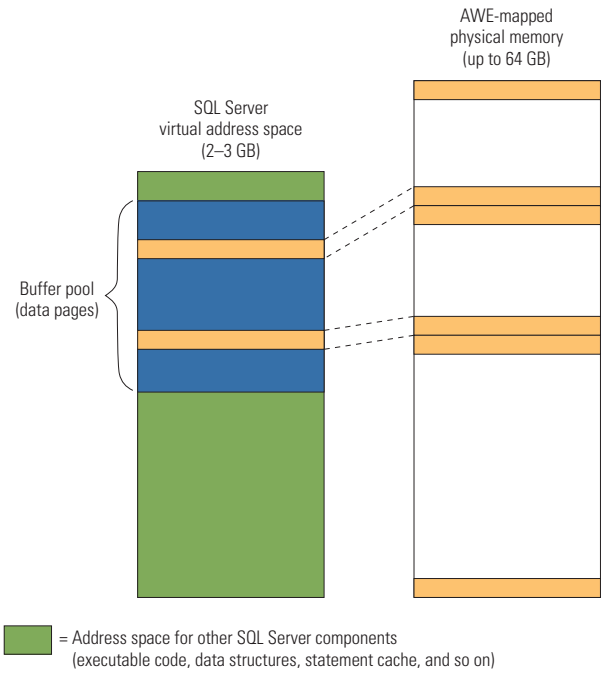

Figure 4. AWE mapping of physical memory into SQL Server virtual space

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

virtual address space available. On a 32-bit edition, only a limited amount of address space in the buffer pool can be used for operations like sorting, hashing, and grouping, which can have a significant impact on analytical processing workloads.

Earlier releases of SQL Server had a limited amount of memory available at startup. SQL Server 2005 introduces support for hot-add memory on Windows Server 2003 Enterprise Edition and Datacenter Edition, enabling usage of physical memory added to the system without restarting. Hot-add memory is supported on all 64-bit versions of SQL Server 2005; it is supported on 32-bit versions only if AWE is enabled. The hot-add memory feature requires specific hardware support.

SQL Server 2005 introduces support for hot-add memory on Windows Server 2003 Enterprise Edition and Datacenter Edition, enabling usage of physical memory added to the system without restarting.

Configuring the temporary database: tempdb

The tempdb database is a globally available system database for holding temporary objects within a SQL Server instance. This database is re-created every time SQL Server is started and does not persist across restarts. Only one tempdb database is created per SQL Server instance. Temporary objects that can be held in tempdb include the following:

- **User objects:** Temporary tables, stored procedures, variables, and cursors
- **Internal database engine objects:** Temporary or intermediate results during query processing, such as sorting, spooling, and hashing
- **Row version information:** Snapshot isolation or online index operations

The size of the tempdb database can be crucial for system performance—if the database is not sized appropriately to meet workload needs, the system can become too occupied trying to auto-grow the database. SQL Server 2005 makes extensive use of tempdb and requires more disk space than earlier SQL Server versions. This change is because of the tempdb space requirements of certain features introduced in SQL Server 2005 and enhancements to existing features. SQL Server 2005 is designed to improve tempdb performance by caching temporary tables and table variables where applicable and minimally logging tempdb operations. Page allocation mechanisms have also been enhanced, which can help improve performance.

Administrators should consider application workload and which SQL Server features are being used when determining an appropriate tempdb size. The physical placement of the tempdb database files is also important for performance: best practices recommend that they be placed on logical volumes separate from those that host other databases or applications in the system. Other best practices for optimizing tempdb performance include using RAID striping and multiple data volumes, configuring the Simple Recovery model[1] for the database, allocating an appropriate initial size to accommodate typical workloads, and setting suitable auto-growth values for the database files for certain unplanned peak conditions.

Configuring efficient, high-performance SQL Server 2005 environments

Microsoft SQL Server 2005 includes configuration options that can help manage system utilization effectively. Administrators should thoroughly understand and assess these options before applying them to particular environments. The default settings for these options should satisfy the requirements of most deployments, but in scenarios where significant control of SQL Server resource usage is necessary, administrators can modify these options to suit their deployment needs.

Ananda Sankaran is a systems engineer consultant in the High-Availability Cluster Development Group at Dell. His current interests related to high-availability clustering include storage systems, application performance, business continuity, and cluster management. Ananda has a master's degree in Computer Science from Texas A&M University.

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

FOR MORE INFORMATION

Dell and SQL Server 2005:
www.dell.com/sql

SQL Server 2005 documentation:
msdn2.microsoft.com/en-us/library/ms203721.aspx

[1] For more information about the Simple Recovery model, visit msdn.microsoft.com/library/default.asp?url=/library/en-us/adminsql/ad_bkprst_60s9.asp.

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Configuring Linux to Enable Multipath I/O

Storage is an essential data center component, and storage area networks can provide an excellent way to help ensure high availability and load balancing over multiple redundant data paths. To take advantage of these benefits in Linux® OS environments, enterprise IT organizations can use applications to set up multipath I/O configurations.

BY TESFAMARIAM MICHAEL, REZWANUL KABIR, JOSHUA GILES, AND JOHN HULL

Related Categories:

Fibre Channel switches

Storage

Storage area network (SAN)

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

In the data center environment, to minimize downtime and service disruptions, IT departments must avoid single points of failure in any highly available system. For storage area networks (SANs), administrators can set up multiple redundant data paths (multipaths) between servers and storage systems to help avoid interruptions in data flow should a hardware failure occur.

To manage a multipath I/O configuration, administrators should ensure that the server OS supports multipath I/O and is configured properly to access data from the storage system and fail over to secondary data paths when necessary. For Linux operating systems, two multipath I/O applications are available: device mapper multipath and EMC® PowerPath® software. This article provides an overview of each application and highlights the advantages and disadvantages of each.

Understanding the basics of multipath I/O

A typical highly available SAN configuration may include a Dell™ PowerEdge™ server containing several host bus adapters (HBAs), two Fibre Channel switches, and a Dell/EMC CX series storage array, as shown in Figure 1; a cluster configuration would include multiple PowerEdge servers. As the figure shows, multiple data paths are configured between the server and the storage system to provide the necessary redundancy. In such a configuration—for example, a PowerEdge server running Red Hat® Enterprise Linux 4—a logical unit (LUN) on the CX storage that is assigned to the server is detected as many times as there are paths available. When an HBA driver loads, the SCSI midlayer initiates a scan of its bus and detects all assigned storage LUNs through every available path. Accordingly, that many SCSI disk devices are registered by the OS. In Figure 1, there are four paths to the storage system, so a LUN assigned to the attached server is detected four times, and four SCSI disk devices are detected by the HBA driver and registered with the server OS.

Despite the benefits of redundant paths, there are challenges to consider. These challenges include identifying a particular device for I/O and managing multiple devices of the same physical device.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

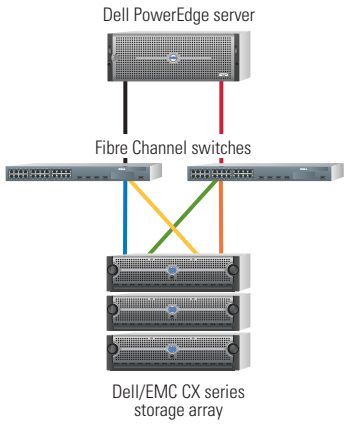

Figure 1. Basic highly available SAN configuration

Different storage systems manage paths to a particular LUN in different ways. Some systems provide an *isotropic* or *symmetric* view of the paths, where all paths are treated as equal. In these cases, all paths are active, and I/Os can be directed to any of them. Other storage systems, such as Dell/EMC CX series Fibre Channel systems, implement *asymmetric* arrays. In this case, paths to the same LUN are divided into active/passive groups, limiting the number of accessible devices at any given time by half.

Active/passive cluster formations allow only one storage processor at a time to be actively performing I/Os to its assigned LUNs. The processors in Dell/EMC storage systems are grouped as storage processor A (SPA) and storage processor B (SPB). A particular port in these systems is associated with only one of these storage processors, and a LUN can be owned by only one of these processors at any given time. Default LUN ownership is specified during its creation. When a failure occurs, the ownership of a LUN can be changed to the other storage processor; this process is known as LUN trespassing. LUN trespassing is achieved by sending a device-dependent trespass command to the storage system. All asymmetric arrays require special hardware handlers to implement this mechanism to either fail over or fail back.

From the server side, a device is referred to as active if its path can be traced to a port owned by the storage processor that owns the LUN. A passive device's path to the LUN can be traced to a port not owned by the storage processor that owns the LUN. To differentiate active devices from passive devices, administrators can issue the command `sfdisk -l /dev/sdX`, where `/dev/sdX` is the SCSI disk device. Devices that do not return an I/O error are active, and those that do return an I/O error are passive. Under the current implementation of multipath I/O in Dell/EMC storage systems, Linux can use only active devices for I/O; passive devices can be used only when the LUN trespasses to the other storage processor. Once a LUN trespasses, the passive devices become active and vice versa. I/Os that were issued before the trespass but did not complete, along with all future I/Os, are then redirected to the devices that just became active.

When the Linux multipath application tools are used, a physical LUN on the storage system that is registered multiple times to the server OS—including active and passive devices—is bound into a single device, providing applications on the server side with a single point to perform I/Os. For instance, for the server shown in Figure 2, if the internal drives are combined in one logical drive in a RAID configuration, Linux would register it as /dev/sda. A LUN on the storage system that is assigned to the attached server would be detected four times by the HBA driver because there are four paths to it. These four devices are registered as /dev/sdb, /dev/sdc, /dev/sdd, and /dev/sde. In this case, because the LUN is owned by SPA, /dev/sdb and /dev/sdd are active and the other two devices are passive. Thus, only /dev/sdb and /dev/sdd can be used for I/O. Both /dev/sdc and /dev/sde become usable only if the LUN trespasses to SPB.

On the server, because the LUN has two active devices associated with it, administrators could mistakenly try to use them as two different devices by mounting these two active devices separately, which can allow data corruption or loss. For instance, if /dev/sdb is used for some data (after partitioning, creating a file system, and mounting it), and then /dev/sdd is later formatted with a file system treating it as a free device, all data will be wiped out in the LUN.

To help avoid this type of situation, administrators should use Linux multipath I/O applications. A multipath application provides a single of point of access by binding the four devices into a single

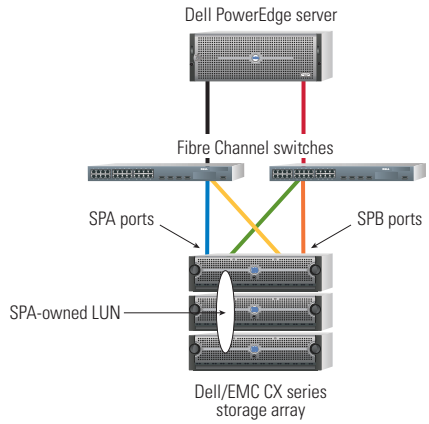

Block device	DM multipath device	Data path	Mode
/dev/sdb	/dev/dm1	Black, blue	Active
/dev/sdc		Black, yellow	Passive
/dev/sdd		Red, green	Active
/dev/sde		Red, orange	Passive

Figure 2. Basic highly available SAN configuration with a LUN owned by SPA

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

device, which can be partitioned, formatted, and mounted. This single device can then be used to distribute I/Os onto all the underlying active devices using a given set of algorithms.

Using device mapper for multipath I/O

Native Linux multipath I/O support has been added to the Linux 2.6 kernel tree with the release of 2.6.13, and has been backported into Red Hat Enterprise Linux 4 in Update 2 and into Novell® SUSE® Linux Enterprise Server 9 in Service Pack 2. It relies on device mapper (DM), a tool for mapping block devices that provides logical volume management, software RAID, and multipath functionality. Combining DM with the multipath user-space application can help create a native Linux multipath I/O configuration.

The overall architecture for DM multipath support in Linux is flexible and modular. DM multipath has a convenient plug-in design that allows administrators to enhance functionality by plugging in a module that achieves the desired result. For example, the DM multipath module has two hooks built into it: path selector and hw handler. The path selector hook is used to determine how I/Os should be distributed among various available paths, and the hw handler hook is used to take hardware-specific actions (for example, LUN trespassing).

Because of this modular architecture, administrators can implement a path selection algorithm (currently only a round-robin algorithm is supported) and register it with the path selector hook to use that particular algorithm to select paths. Similarly, administrators can implement a hardware-specific handler (for example, dm_emc) and register it with the hw handler hook of the DM multipath module to allow hardware-specific actions. For example, Dell/EMC CX series systems require the dm_emc handler to perform LUN trespassing for failover or failback.

In addition to these packages, some DM kernel modules, such as dm_multipath, dm_round_robin, and dm_emc, are also required. DM includes a user-space configuration tool (dmsetup) and a library (libdevmapper). DM multipath support also includes a multipath configuration file (multipath.conf), an init script (multipathd), udev rules, a device map creation tool from partitions (kpartx), and a multipath executable binary, among others. Udev is a recent Linux user-space application that manages devices (/dev/ directory) dynamically.

When DM multipath starts, it retrieves the universally unique identifier (UUID) of all the block devices in /proc/partitions (except those excluded in its configuration file) by issuing the `scsi_id -eg -s /block/sd`*X* command. It then groups all the block devices with the same UUID and creates a single device for them in /dev/mapper/. When this device is created, it can be partitioned with `fdisk` or `parted`. The partitions can be registered in /dev/mapper/ using `kpartx`, formatted with a file system, and mounted for usage.

DM multipath uses round-robin algorithms to balance I/Os across all active paths. If it experiences a failure when performing I/Os on the active devices because of a path disconnection, the DM kernel module (dm_emc in the case of Dell/EMC CX series systems) issues a trespass command (`switch-over`) to the system to switch over ownership of the LUN. Until the LUN trespasses successfully, all I/Os are queued. Once the trespass is successful, the passive devices become active and the active devices become passive, and DM multipath shifts I/Os (including those queued) to the new active devices.

Setting up the multipath configuration

To set up a multipath I/O configuration, administrators must first gather the UUIDs of the block devices. As mentioned earlier, the `scsi_id` command can be used to obtain the UUID of a block device. The default device naming can be changed by specifying aliases to UUIDs. These aliases as well as other settings for the multipath I/O configuration are set in the configuration file. The following steps describe how to configure systems for multipath I/O; these steps use the sample configuration shown in Figure 2.

The block devices have the same UUID, because they are all devices for one physical LUN. Administrators can issue the following commands to obtain the UUID for the four block devices in Figure 2:

```
scsi_id -g -s /block/sdb
scsi_id -g -s /block/sdc
scsi_id -g -s /block/sdd
scsi_id -g -s /block/sde
```

The output of all four commands is the same: a long hexadecimal number. A multipath configuration file has four sections: `devnode_blacklist`, `defaults`, `multipaths`, and `devices`. Visit *Dell Power Solutions* online at www.dell.com/powersolutions to see the sample multipath configuration file referred to in this article.

The `devnode_blacklist` section lists devices to be excluded from the multipath, which thus will not be probed for UUIDs. In the sample file online, all IDE devices (/dev/hd[a-z]) are excluded. When DM multipath starts, it will not issue any commands to these devices.

The `defaults` section assigns the default values to the specified multipath parameters. In the sample multipath configuration file online, these parameters include `multipath_tool`, which passes any argument to the `multipath` command; `polling_interval`, which dictates how often the devices should be pinged; and `default_selector`, which specifies the algorithms. Note that `default_hwhandler` should be set to `1 emc` to load the dm_emc module and issue all the necessary commands, including the trespass command, to the Dell/EMC CX series systems.

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

```
fdisk /dev/mapper/dm1
kpartx -l /dev/mapper/dm1        # lists all partitions on this device
kpartx -a /dev/mapper/dm1        # adds all partitions on this device in /dev/mapper/
```

Figure 3. Commands to create and add partitions

```
mke2fs -j /dev/mapper/dm1p1      # creates a file system
mkdir /data
mount /dev/mapper/dm1p1  /data   # mounts device on /data
df -h /data                      # displays device properties
```

Figure 4. Commands to create a file system on partitions and mount a device

The `multipaths` section embeds as many entries of the multipath stanza as there are available LUNs assigned to the server. The internal multipath stanza specifies the UUID (or `wwid`, as shown in the sample figure online) and the alias to the LUN. This figure includes only one multipath entry, which sets the `wwid`, `alias`, and `path_checker` (to check the path regularly) variables. It is important that these variables are set. For the `wwid` value shown in the sample multipath configuration file online, an alias device dm1 is created in /dev/mapper/ when DM multipath starts.

The `devices` section, similar to the `multipaths` section, also embeds the device stanza. In an environment with multiple SAN storage systems, several device entries are necessary. This internal stanza shows vendor-specific SAN settings.

After setting the multipath configuration file, administrators can start DM multipath by issuing the `multipath` command. To generate detailed return messages, they can issue the command as `multipath -v3 -ll`. This command displays useful information such as the size of the LUN, the alias, a list of active and passive devices, and other settings. It also displays the alias devices created (/dev/mapper/dm1 in the sample multipath configuration file online).

Next, administrators should create a partition on /dev/mapper/dm1 using `fdisk` or `parted` and add the partitions to /dev/mapper/ using the commands shown in Figure 3.

Finally, administrators should create the file system on the partitions (/dev/mapper/dm1p1) and mount the device using the commands shown in Figure 4.

The LUN can then be accessed using the /data mounting point, and data can be read from and written to it. To verify that the LUN can be used as expected, administrators should perform some I/O activities by copying files to /data. At the same time, they should confirm the I/O activity by issuing the command `iostat -d 1`.

Using EMC PowerPath for multipath I/O

EMC PowerPath provides similar functionality as DM multipath, but also includes features not present in DM multipath, such as a variety of algorithms (including round-robin), the ability to set priority for its devices, and the ability to report current configurations.

PowerPath for Linux is packaged in the Red Hat Package Manager (RPM™) format. EMC releases new or updated versions regularly. Usually a particular version supports a specific Linux release, such as Red Hat Enterprise Linux or SUSE Linux. The package can be downloaded from the EMC Web site at www.emc.com.

Once downloaded to the server, the package can be installed using the `rpm -ivh` command. For example, if EMCpower.LINUX-4.4.0-337.rhel.i386.rpm is downloaded for use on a 32-bit Intel architecture (IA-32) system running Red Hat Enterprise Linux 4, administrators can install this package by issuing the following command (the majority of the package's files are copied to /etc/opt/emcpower):

```
rpm -ivh EMCpower.LINUX-4.4.0-337.rhel.i386.rpm
```

PowerPath includes powermt, a powerful management utility for its devices. Its man page (man powermt) provides specific information about the utility. Among other features, powermt allows administrators to display the current settings; set priority, policy (algorithms), and mode; remove a particular HBA or device; and restore a removed HBA or device.

In addition, PowerPath comes with its own init script and can be started and stopped from the command line. When stopping PowerPath, administrators should be sure that there is no I/O activity—that is, PowerPath should not be in use by any application. After installation, administrators can start PowerPath by issuing the command `service PowerPath start`.

As with DM multipath, once started PowerPath gathers the UUIDs of the block devices and bundles the devices with the same UUID into a single device, /dev/emcpower*X*. However, it does not use a configuration file. As PowerPath identifies the LUNs, it enumerates them as /dev/emcpowera, /dev/emcpowerb, and so on. Because PowerPath relies on how the HBA driver has detected the LUN and created the block devices, and does not

Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

use an administrator-supplied configuration file, its enumeration of LUNs can vary from one node to the next in clustered servers. For the configuration in Figure 2, /dev/sdb, /dev/sdc, /dev/sdd, and /dev/sde are all bundled to /dev/emcpowera. This device can be partitioned, formatted with a file system, and mounted using the commands shown in Figure 5.

```
fdisk /dev/emcpowera            # partitions the device
mke2fs -j /dev/emcpowera1       # formats with ext3 file system
mkdir /data
mount /dev/emcpowera1 /data     # mounts the partition
df -h
```

Figure 5. Commands to partition, format, and mount a device

To stop PowerPath, administrators should first confirm that all of the /dev/emcpower*X* devices are not in use (that is, they must stop all I/Os to the devices and unmount them). They can then issue the command `service PowerPath stop`.

Comparing DM multipath with EMC PowerPath

When considering which multipath application to deploy, IT departments must take into consideration the features, level of manageability, and type of support. Given that DM multipath is relatively new, PowerPath is much more feature rich. For example, DM multipath provides only round-robin algorithms, but PowerPath provides nine different policies, including round-robin, adaptive, and basic failover. PowerPath also supports dynamic load balancing, automatic path failover, and online recovery. In addition, PowerPath allows administrators to set different priority levels for its devices, benefiting applications that use the devices with higher priorities.

Although DM multipath is relatively new compared with PowerPath, it has solid backing in the Linux community and is expected to develop into an even stronger alternative to PowerPath in the future.

For manageability, PowerPath has an advantage in heterogeneous OS environments, because it is supported on the Microsoft® Windows®, Linux, UNIX®, and Novell NetWare® operating systems. DM multipath is available only on Linux and has relatively immature management support. However, DM multipath does allow for a consistent mapping of devices to LUNs in a cluster environment, which PowerPath does not.

Support for PowerPath is limited to the specific Linux operating systems supported by EMC, which typically includes only Red Hat Enterprise Linux and SUSE Linux. Because PowerPath is proprietary software, administrators must be running both a supported OS and a supported kernel to have PowerPath support, which can be inconvenient because of the large number of Linux distributions unsupported by PowerPath. Also, when new kernels are released by Linux vendors, there may be a lag between the kernel release and PowerPath support for that release. DM multipath does not have these limitations, because of its GNU General Public License and inclusion with most Linux distributions. Any new kernel released by a Linux vendor includes DM multipath support by default.

Choosing the appropriate multipath I/O application

Linux device mapper multipath and EMC PowerPath both provide viable and robust multipath I/O capability for Linux operating systems on Dell PowerEdge servers and Dell/EMC storage systems. Choosing the appropriate application depends on the specific data center environment and the necessary features and support. Although DM multipath is relatively new compared with PowerPath, it has solid backing in the Linux community and is expected to develop into an even stronger alternative to PowerPath in the future.

Tesfamariam Michael is a software engineer in the Dell Database and Application Engineering Department of the Dell Product Group. Tesfamariam has a B.S. in Electrical Engineering from the Georgia Institute of Technology, and a B.S. in Mathematics and an M.S. in Computer Science from Clark Atlanta University.

Rezwanul Kabir is a systems engineer in the Dell Linux Development Group. He has a B.S. in Computer Science and Engineering from Bangladesh University of Engineering and Technology and an M.S. in Computer Science from New Mexico State University.

Joshua Giles is a software engineer at Red Hat. His interests include operating systems, grammars, automata-based programming, and support vector machines (machine learning). Joshua has a B.S. in Computer Science from the New Mexico Institute of Mining and Technology.

John Hull is the manager of the Linux OS Development team at Dell. He has a B.S. in Mechanical Engineering from the University of Pennsylvania and an M.S. in Mechanical Engineering from the Massachusetts Institute of Technology.

FOR MORE INFORMATION

EMC PowerPath:
software.emc.com/products/software_az/powerpath.htm

 Reprinted from *Dell Power Solutions,* August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Using Intel® Multi-Port Server Adapters to Optimize Virtual Infrastructures

By Robert Hays and Eric Pan

Using virtualization to consolidate physical servers, storage, and networks; improve resource utilization; and streamline business continuity processes is an increasingly popular strategy undertaken by organizations struggling to contain IT costs. To optimize the benefits of virtualization, multi-port server adapters can be an essential hardware component that helps ensure availability, reliability, and performance within a virtual infrastructure.

Recommended LAN Configuration for VMware ESX Server* 3

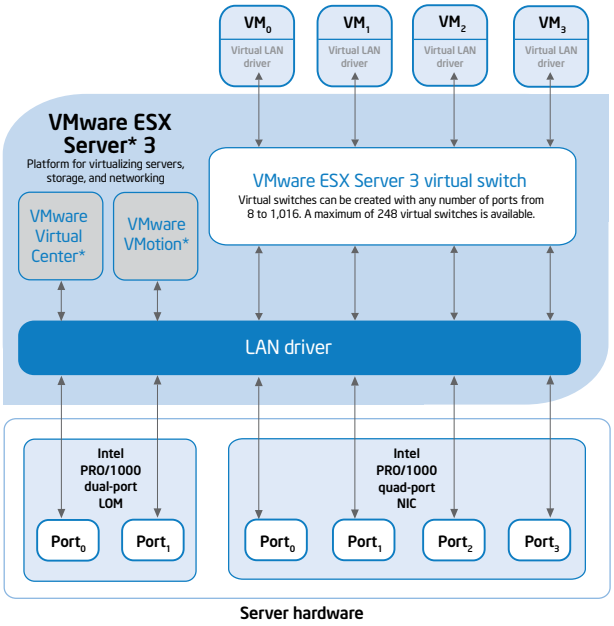

In an era where the continuing challenge is to do more with less, IT departments are casting a critical eye on server sprawl and improving business continuity. Responding to company growth and rapidly changing business needs, many organizations add substantial numbers of new servers—which can increase hardware costs, systems management complexity, administrative overhead and site costs. Compounding this problem is underutilization of server resources. Improving server utilization from the typical 10 to 15 percent up to 80 or 85 percent[1] can have dramatic effects on total cost of ownership—helping to reduce both equipment and administrative costs.

The requirement to lower total cost of ownership for IT assets is driving organizations to reduce the number of physical servers, switches, and storage devices by implementing virtual infrastructures. This approach can help reduce IT costs while increasing flexibility and responsiveness, streamlining business continuity, and optimizing performance and reliability of IT infrastructures. The flexibility to pool resources enables distributed services to optimize and manage IT environments so they can respond immediately to changing business needs and cost-effectively scale the IT capacity of the enterprise.

[1] In multiple-workload tests, a virtualization platform comprising Dell PowerEdge* 2850 servers configured with single-core and dual-core Intel® Xeon® processors DP, Intel® PRO/1000 Multi Port Server Adapters, and VMware ESX Server* software achieved 85 percent CPU utilization. For more information, see "VMware ESX Server Performance on Dual-Core Dell PowerEdge 2850 Servers" by Todd Muirhead and Dave Jaffe, Ph.D., in *Dell Power Solutions*, May 2006, www.dell.com/downloads/global/power/ps2q06-20060236-Muirhead.pdf.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Case Study
Intel® PRO Multi-Port Adapters
Design and Manufacturing

Virtualization Helps Woodward Governor Company Consolidate Servers and Improve IT Responsiveness

Virtual Computing Delivers Cost Savings

Challenges

- Increase IT efficiency while controlling costs
- Serve end users better by improving performance, reliability, and IT responsiveness
- Scale quickly and flexibly to meet business demands

Solutions

- Dell PowerEdge* 2850 servers (single-core and dual-core Intel® Xeon® processors)
- Intel® PRO/1000 Dual Port Network Adapters
- Intel® PRO 1000MT Quad Port Network Adapters
- VMware ESX Server* virtualization software
- VMware VirtualCenter* management software
- VMware VMotion* technology

Benefits

- Reduced the number of new physical server deployments by 50 percent using virtual servers
- Kept IT costs to a minimum every year for the past three years
- Accelerated provisioning and simplified resource allocation using Intel® multi-port adapters and VMware–based virtual servers compared to a deployment comprising all physical servers
- Enhanced network reliability and optimized performance for high-bandwidth applications using Intel multi-port adapters and VMware virtualization software

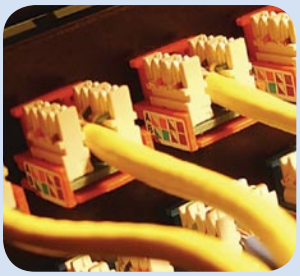

As one of the world's largest independent designers, manufacturers, and service providers of energy control systems for industrial and aircraft engines, Woodward Governor Company is adept at helping customers gain efficiency and high performance. Unfortunately, a few years ago the IT department at Woodward found it difficult to do the same for its internal customers.

As business needs grew and changed, the IT organization, which supports roughly 3,500 global end users, began searching for a way to optimize computing resources and current investments while providing the bandwidth and reliability needed for business-critical applications. When Woodward discovered that server virtualization could help meet these challenges, the company began consolidating servers and never looked back.

Since implementing VMware virtualization software, Woodward has reduced the need for separate physical servers and reigned in server sprawl. Of nearly 100 potential server replacement or deployment

continued from previous page

Providing the foundation: Virtualization software

VMware Infrastructure* 3 is a feature-rich suite designed to deliver a production-ready virtual infrastructure that empowers a responsive data center. This product suite offers tremendous benefits to IT environments and enhances efficiency, increases availability, and provides dynamic workload management. At the core of VMware Infrastructure 3 is VMware ESX Server* 3—a platform for virtualizing servers, storage, and networking.

VMware ESX Server 3 software runs on industry-standard Intel® Xeon® processor–based platforms, providing the virtualization layer that allows multiple virtual machines (VMs) to coexist in isolation from one another on a single physical server as well as virtualizing storage and networking. ESX Server 3 software typically enables organizations to run from one to eight VMs per processor on two-processor and four-processor systems such as the Dell PowerEdge* 1950 server, Dell PowerEdge 2950 server, and Dell PowerEdge 6850 server as well as the Dell PowerEdge 1955 blade server.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

opportunities, 60 percent have been on virtual machines. New workloads are first evaluated as potential virtualization candidates before any hardware is purchased, and as existing servers age and are replaced, the IT department continues to pursue consolidation opportunities to help improve IT efficiency and system performance.

Virtual environment

To enable virtualization, Woodward is running VMware ESX Server 2.5 and VMware VirtualCenter 1.3 on 10 Dell PowerEdge 2850 servers, all configured with a 2 Gbps Fibre Channel connection to a Dell/EMC CX300 SAN (see figure at right). Each server has two Gigabit Ethernet connections plus an Intel PRO/1000 quad-port server adapter. One port on the quad-port adapter is dedicated for a connection to an iSCSI-based network, while the other three ports are connected to the data center LAN switches. Some of the adapter ports are dedicated to a specific virtual machine (VM) that may have high bandwidth needs while other adapter ports are bonded together to provide a dual-gigabyte trunk for VMs.

Woodward has more than 80 VMs spread across 10 Dell PowerEdge 2850 servers. Some of the virtualized servers were running outdated operating system versions on aging hardware that was becoming unreliable and a challenge to support. The legacy servers were moved into Woodward's virtual environment, where they are now running in VMs on leading-edge hardware that would not necessarily be supported by the legacy operating systems.

Other examples of enterprise applications running in VMs on Dell PowerEdge 2850 servers with single- and dual-core Intel Xeon processors include database servers, software update servers and systems management servers.

Role of Intel multi-port adapters

When the company began implementing server virtualization, it deployed both Intel PRO dual- and quad-port adapters extensively to support the new virtual computing infrastructure. Woodward purchases Intel PRO adapters specifically because Intel offers full support for VMware software and demonstrates consistently high throughput.

Intel PRO multi-port network adapters are now a key component in Woodward's toolkit. Moving from dual-port to quad-port NICs has provided Woodward with enhanced server virtualization capabilities, enabling improved bandwidth as well as increased flexibility and scalability.

For example, Intel PRO network adapters enable the IT department to optimize network reliability and performance for high-bandwidth applications, and to rapidly deploy new applications and services. As a result, the IT department can easily allocate server resources to support variable workloads. Having multiple ports available allows Woodward's IT staff to leverage features such as network segmentation, adapter teaming for high throughput, load balancing, and fault tolerance.

Using Intel PRO multi-port Gigabit Ethernet network adapters and VMware virtualization software running on Dell PowerEdge servers, Woodward's IT department has been able to maximize resource utilization and minimize IT costs every year for the past three years—which indicates a very real return on investment.

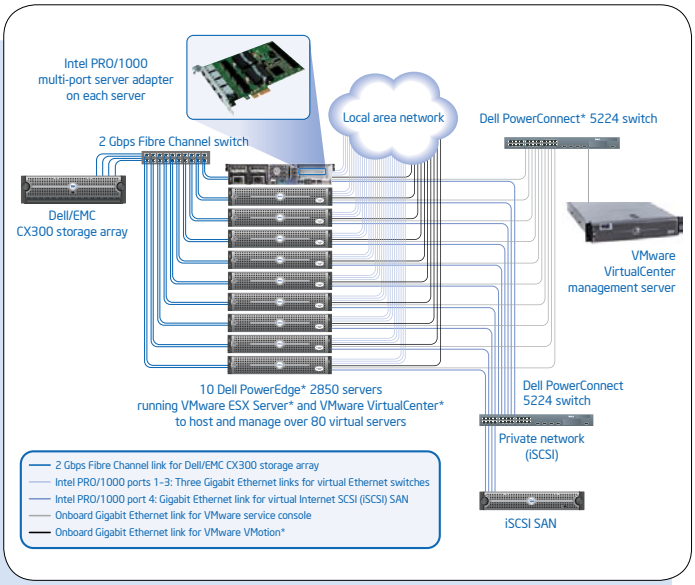

GOT A GREAT STORY ABOUT INTEL® NETWORK SERVER ADAPTERS?

Tell us about it at www.morph2cool.com, and if we complete a case study on your implementation, you'll get a pair of Oakley® Thump sunglasses!

Additional components of the VMware Infrastructure 3 suite feature distributed services designed to increase high availability and minimize downtime and system recovery times.

Enabling technology: Multi-port server adapters

To achieve the benefits of server virtualization, administrators must ensure that the appropriate number and type of network connections are available on physical servers. Best practices for VMware ESX Server 3 recommend a minimum of three, and typically four or more, network ports—which can be a challenge if physical servers are constrained by an insufficient number of PCI slots. With VMware ESX Server 3, two network ports, which are provided as standard LAN on Motherboard (LOM) connections, are dedicated for VMware VirtualCenter* management software and VMotion* technology (see figure on first page). Additional network ports, via add-in server adapters, are required to support guest operating system application networking and/or networked storage data. In choosing the right server adapter for virtualized servers, IT administrators

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

Understanding the Benefits of PCI Express*

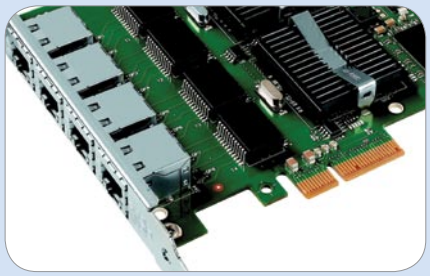

PCI Express[3] is designed to provide faster Ethernet network performance for desktop and server slots than that of PCI and PCI Extended (PCI-X). PCI Express derives a performance advantage from serial bus architecture that supports dedicated, bidirectional I/O at 2.5 GHz, compared to the 133 MHz parallel bus supported by PCI-X.

Designed for PCI Express, the latest family of Intel® PRO Ethernet multi-port server adapters enables organizations to realize maximum benefit from server virtualization and consolidation by taking advantage of the higher bandwidth and scalability of PCI Express.

Dell PowerEdge* servers can support multiple PCI Express slots to accommodate multiple Gigabit Ethernet adapters for network performance–enhancing techniques such as network segmentation and network adapter teaming. In high-density, low-profile rack servers, slots can be fully utilized by selecting dual- or quad-port Gigabit Ethernet adapters.

Intel PRO 1000 Network Server Adapters that support PCI Express include the following:

- Intel PRO/1000 PT (U3867) or PF (GF668) Server Adapters for PCI Express* slots: High-performance, self-configuring 10/100/1,000 Mbps connection
- Intel PRO/1000 PT Dual Port Server Adapter (X3959): 2 Gbps copper server connections in a single PCI Express slot
- Intel PRO/1000 PT Quad Port Server Adapter (EXPI9404PT, EXPI9404PTBLK): 4 Gbps copper server connections in a single PCI Express slot
- Intel PRO/1000 PF Dual Port Server Adapter (EXPI9402PF): 2 Gbps fiber server connections

continued from previous page

should consider factors such as the number of guest operating systems, application network bandwidth requirements, cable type, PCI slot availability, and redundancy/failover protection requirements.

Multi-port server adapters such as Intel® PRO/1000 PCI-X and PCI Express* Dual and Quad Port Server Adapters provide up to four ports in a single network card. These adapters, available for both copper and fiber networks, conserve valuable PCI slots in servers while helping to eliminate network bottlenecks in connection-dense virtual computing environments.

Reliability is critical in a consolidated infrastructure

On a server running multiple VMs, a port failure may result in a costly business interruption. Intel PRO multi-port Gigabit Ethernet[2] adapters are designed to prevent network downtime. Intel® Advanced Network Services (ANS) software allows administrators to maximize uptime with redundant, teamed connections. Intel multi-port adapters can be configured to automatically switch to a secondary link when a server's primary link fails, and features such as switch fault tolerance and test switch configuration help administrators verify compatibility and enhance uptime.

Realizing the benefits of virtual infrastructures

Virtual infrastructures empower organizations to create resource pools from servers, storage and networks so that IT resources can be matched and scaled to specific enterprise requirements—significantly enhancing IT efficiency and responsiveness. In this way, a virtual infrastructure consisting of VMware Infrastructure 3 software, Intel Xeon processor–based Dell PowerEdge servers, and Intel PRO multi-port server adapters can help IT managers simplify operations, improve resource utilization, and scale cost-effectively.

Robert Hays is a marketing manager in the Intel Digital Enterprise Group, where he leads the LAN Access Division's product marketing team. Rob has spent many years at Intel defining and bringing innovative networking products to market for client and server platforms. He has a bachelor's degree in Computer Engineering from Georgia Tech.

Eric Pan is a senior alliance marketing manager at VMware, where he leads marketing-related initiatives between VMware and Dell. He has a B.S. in Management Information Systems and an M.B.A. in Finance.

[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.
[3] For more information about PCI Express*, see the white paper "PCI Express Ethernet Networking" by Intel Corporation, September 2005, www.intel.com/network/connectivity/products/whitepapers/mesh_pcie_whitepaper.pdf.

For More Information

Henderson, Bill, and Tracy D. Edwards. "Using Intel Multi-Port Server Adapters to Enable Virtual Infrastructure in the Data Center." *Dell Power Solutions*, May 2005. www.dell.com/downloads/global/power/ps2q05-20050131-Intel.pdf.

Intel Network Solutions:
www.intel.com/network/connectivity/index.htm

Virtualization solutions from Intel and VMware:
www.intel.com/network/connectivity/solutions/virtualization.htm

Gigabit solutions from Intel and Dell:
www.intel.com/go/dellgig10

VMware virtual infrastructure software:
www.vmware.com/dell

Dell and VMware:
www.vmware.com/virtuallyspeaking
www.dell.com/vmware

Intel, the Intel logo, and Intel Xeon 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.
Copyright ® 2006 Intel Corporation All rights reserved.

Reprinted from *Dell Power Solutions*, August 2006. Copyright © 2006 Dell Inc. All rights reserved.

PURELY UNCOMPLICATED

NEW DELL POWEREDGE SERVERS. BORN TO REDUCE COMPLEXITY.

Behold the new Dell PowerEdge family of servers. Fewer system images help you save time during software updates, and built-in LCD displays help turn datacenter chaos into a thing of the past. For a 360° online view of this pure leap forward, visit www.dell.com/poweredge. Form combined with function. That's the direct path to success. That's pure Dell.

PURE DELL

www.dell.com/poweredge | 1.866.217.3055

Dell cannot be responsible for errors in typography or photography. Dell, the Dell logo and PowerEdge are trademarks of Dell Inc. © 2006 Dell Inc. All rights reserved.

Oracle Database 10*g*

#1 On Windows

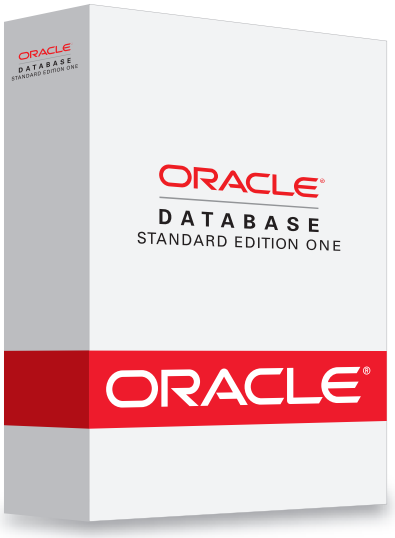

Starts at $149 per user

Oracle Database 10*g*—
The World's #1 Database. Now For Small Business.

ORACLE®

oracle.com/start
keyword: #1onWindows
or call 1.800.633.0675

Terms, restrictions, and limitations apply. Standard Edition One is available with Named User Plus licensing at $149 per user with a minimum of five users or $4995 per processor. Licensing of Oracle Standard Edition One is permitted only on servers that have a maximum capacity of 2 CPUs per server. For more information, visit oracle.com/standardedition

Copyright © 2006, Oracle. All rights reserved. Oracle, JD Edwards and PeopleSoft are registered trademarks of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners.