DELL POWER SOLUTIONS · FEBRUARY 2008

POWER SOLUTIONS

DELL.COM/PowerSolutions

FEBRUARY 2008 • $4.95

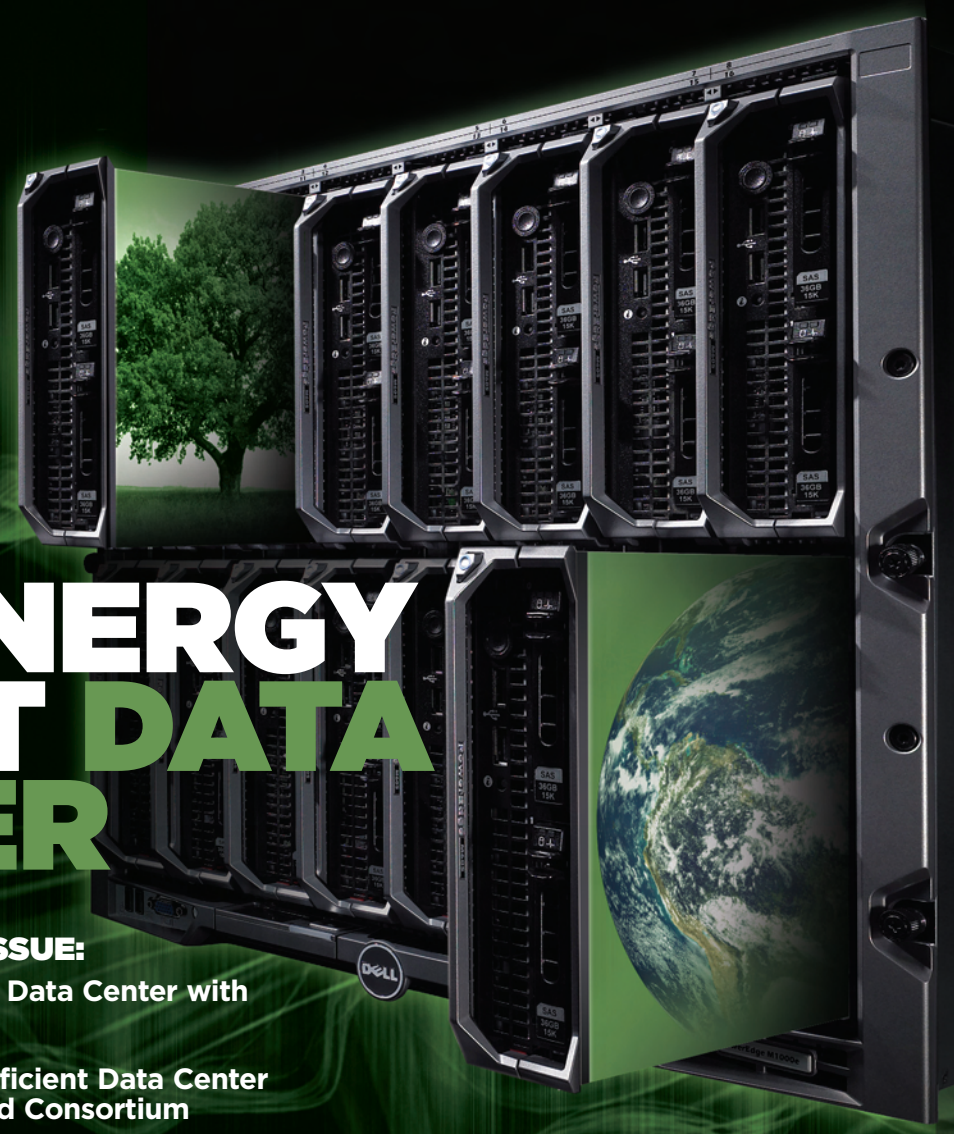

THE ENERGY SMART DATA CENTER

ALSO INSIDE THIS ISSUE:

Unlocking Your Hidden Data Center with Green IT Best Practices

Enabling the Energy-Efficient Data Center Through The Green Grid Consortium

Exploring Snap-in Scalability with the Next-Generation Dell PowerEdge M1000e Modular Blade Server Enclosure

PLUS:

SIMPLIFYING STORAGE WITH DELL EQUALLOGIC ARRAYS

SEE PAGE 36

THE ENERGY SMART DATA CENTER

DPS 2008-02 (FEB 08)

Learn more: intel.com/go/10GbE

Intel® 10 Gigabit XF SR Server Adapter
Now available for Dell PowerEdge* Servers

FEBRUARY 2008

POWER SOLUTIONS

DELL.COM/PowerSolutions

10 COVER STORY

THE ENERGY SMART DATA CENTER

By John Pflueger, Ph.D., and Albert Esser, Ph.D.

Going green can be the secret to significant cost savings as well as aggressive performance growth. Dell offers a comprehensive strategy that includes virtualization and consolidation onto energy-efficient systems, best practices for power and cooling optimization, and expert services that can help businesses achieve immediate benefits.

16 BEST PRACTICES

UNLOCKING YOUR HIDDEN DATA CENTER

Dr. Albert Esser, vice president of data center infrastructure at Dell, shares his perspective on the momentum that is building behind green data center design. Plus: How IT organizations can put unused capacity to work and leverage tactics for fast, flexible growth that helps maximize the bottom line.

SPECIAL SECTION

EQUALLOGIC iSCSI PEER STORAGE

36 **Inside the EqualLogic PS Series iSCSI Storage Arrays**

By John Joseph, Eric Schott, and Kevin Wittmer

Built on a patented peer storage architecture, the EqualLogic PS Series of Internet SCSI (iSCSI) storage arrays offers high performance, reliability, intelligent automation, and seamless virtualization of a single pool of storage to enable simplified enterprise storage deployment, ease of management, and comprehensive data protection.

42 **Flexible Virtualization with EqualLogic PS Series iSCSI Storage Arrays**

By Timothy Sherbak

By integrating server virtualization with the storage virtualization of EqualLogic PS Series iSCSI arrays, any size organization can create a powerful, highly flexible virtualized infrastructure to enhance data center capabilities and benefits.

48 **Business Continuity for SMBs with EqualLogic PS Series iSCSI Storage Arrays**

By Kevin Wittmer

EqualLogic PS Series iSCSI storage arrays enable organizations of all sizes to deploy an enterprise-level storage infrastructure offering comprehensive data and application protection.

JOIN THE REGENERATION

The ReGeneration is a global movement dedicated to sustaining the world's natural environment, with a focus on community dialogue and participation. To read the ReGeneration blog, watch videos, get tips and downloads, or get involved by joining up yourself, visit **ReGeneration.org**. For more about Dell's commitment to improving the environment, visit **DELL.COM/Earth**.

CONTENTS

EDITOR'S COMMENTS

6 **Moving Along the Green IT Continuum**

By Tom Kolnowski

NEWS AND TRENDS

8 **Dell EqualLogic Offerings to Deliver Storage Primed for Server Virtualization**

Everdream Acquisition Extends Reach of Remote Management Services

FEATURE SECTION

GREENING THE DATA CENTER

20 **The Green Grid: Enabling the Energy-Efficient Data Center**

By Christian Belady and John Pflueger, Ph.D.

The Green Grid—an international consortium dedicated to improving energy efficiency in data centers and business computing ecosystems—is developing platform-neutral standards, metrics, measurement methods, processes, and technologies that promise to help conserve energy for sustainable growth.

24 **Increasing Energy Efficiency with Dell/EMC CX3 Storage**

By Rodan Zadeh

Energy-efficient technologies in Dell/EMC CX3 series storage are designed to help enterprises optimize application performance while reducing energy use and controlling ongoing operational costs.

28 **Managing Energy Use with Dell Client Manager from Altiris**

By Todd Mitchell

Dell Client Manager software from Altiris provides a simple, effective way to manage power schemes and other settings for Dell client systems to help reduce power consumption and energy costs.

32 **A Systems-Level Approach to Efficient Data Center Design**

By Fred Stack

Dell-Liebert Energy Smart Solutions can help organizations increase performance while maintaining existing levels of energy use, or provide sufficient cooling for a maximum-performance infrastructure.

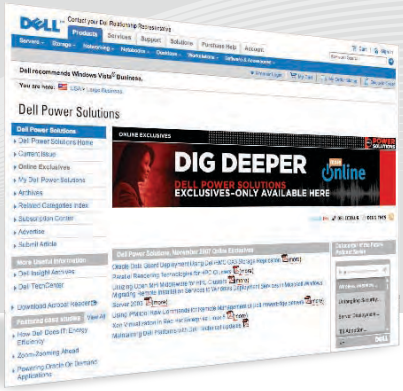

MORE ONLINE

DELL.COM/PowerSolutions

SEE IT HERE FIRST

DELL POWER SOLUTIONS ONLINE EXCLUSIVES

GREENING THE DATA CENTER

Best Practices for Increasing Data Center Energy Efficiency

By Paul Rad, Max Thoene, and Tim Webb

By consolidating systems using virtualization and Dell PowerEdge Energy Smart servers, organizations can retire legacy hardware to help significantly reduce power and cooling requirements.

High-Efficiency Cooling Through Computational Fluid Dynamics

By Paul Rad; Kailash Karki, Ph.D.; and Tim Webb

Predictive modeling based on computational fluid dynamics enables enterprises to configure their data centers for optimal cooling, helping maximize efficiency, reduce costs, and meet both current and future IT requirements.

Evaluating Energy Efficiency in InfiniBand-Based Dell PowerEdge Energy Smart Clusters

By Ramesh Radhakrishnan, Ph.D.; Rizwan Ali; and Vishvesh Sahasrabudhe

By building HPC clusters using the InfiniBand interconnect and Dell PowerEdge Energy Smart servers, organizations can help reduce power consumption without compromising performance.

Building Highly Available HPC Clusters with IBRIX Fusion and the Dell PowerVault MD3000

By Amina Saify, Aziz Gulbeden, and Onur Celebioglu

Qualified high-performance computing (HPC) cluster configurations based on the IBRIX Fusion file system and Dell PowerVault MD3000 storage can provide a robust, cost-effective, high-performance architecture.

Securing iSCSI Storage Networks

By Surendra Bhat, Lokesh Singh, and Santosh Bhadri

Following best practices can help administrators secure their networks when deploying iSCSI in enterprise environments.

Disabling Local Configuration and Remote Virtual KVM in the DRAC 5

By Sriranjan Bose and Abhay Salunke

Dell Remote Access Controller 5 (DRAC 5) firmware version 1.30 enables administrators to easily disable local configuration and remote vKVM (virtual keyboard, video, mouse) to help increase data center flexibility and security.

Smart Card Logon in the DRAC 5

By Sriranjan Bose and Abhay Salunke

DRAC 5 firmware version 1.30 introduces the smart card logon feature, which is designed to provide secure two-factor authentication on Microsoft Windows platforms.

CASE STUDIES

104 **Breaking the Energy Barrier**
Dell PowerEdge Energy Smart servers helped MarketLive cut power consumption by an estimated 50 percent while providing twice the processing power in the same amount of space.

106 **Excellent Online Service 24/7**
Dell PowerEdge servers, Dell/EMC storage, and Dell OptiPlex desktops help online retailer Tesco.com support a 30 percent annual customer expansion rate.

108 **Advancing Medical Treatment Efficiency**
Deploying an IT environment built on Dell PowerEdge servers and VMware virtualization software enabled Taiwan Adventist Hospital to create a highly available IT environment that allows medical staff to focus on patients.

110 **Real-World Education**
By consolidating 90 percent of its physical servers onto Dell PowerEdge blade servers using the VMware virtualization platform, the Ryerson University IT group is poised to react quickly to university needs.

MORE ONLINE

DELL.COM/PowerSolutions

INSIDE DELL IT

Optimizing Call Center Efficiency with Automated Verification and Intelligent Routing

By Glen Curry, Ralph Hilliard, and Carthikeyan Shanmuganathan

The Dell IT group is working to automate parts of customer interaction to enhance both agent productivity and customer experience, and potentially simplify Dell call center activities.

How Dell Simplified IT Management by Designing a Business Service Management Dashboard

By Joshua David, Tom Painter, and Frank Tang

The innovative Dell Business Service Management Dashboard helps enhance IT efficiency, increase uptime, and provide executive-level visibility into the Dell IT environment.

Advanced Web Technologies Help Save Customer and Developer Time at Site Checkout

By Jun Yang

In a recent site upgrade, Dell IT identified Web technologies that help increase flexibility and streamline site development.

Dell EarthWatch Application Offers a Global View of Web Customer Engagement

By Shawn Fielding and Christos Pattichis

EarthWatch combines hourly traffic on the Dell Web site with Google Earth images to provide a near-real-time view into Web customer engagement worldwide.

How Dell Centralized Worldwide Server Management Using Microsoft Operations Manager 2005

By deploying Microsoft Operations Manager 2005, the Dell IT group was able to centralize and consolidate management of its Microsoft systems worldwide.

YOUR INFORMATION, YOUR WAY

Dell Power Solutions articles are also available online at **DELL.COM/PowerSolutions**. Check the *Dell Power Solutions* Web site for our early-release articles, how-to's, implementation studies, and expert tips you won't find anywhere else.

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

DEPARTMENTS

BLADE SERVERS

52 **The Next-Generation Dell PowerEdge M1000e Modular Blade Enclosure**
By Chad Fenner

56 **Exploring the Dell PowerEdge M1000e Network Fabric Architecture**
By John Loffink

60 **Managing Dell PowerEdge M1000e Blade Servers with the Avocent iKVM Switch**
By Stephen M. Hahn and Chad Fenner

62 **Simplified Management with Altiris Deployment Solution for Dell Servers 3.0**
By Eric Szewczyk

FLEXIBLE COMPUTING

70 **Simplifying IT with Dell On-Demand Desktop Streaming**
By Aaron Prince and Bharath Vasudevan

STORAGE

74 **Extending the Advantages of Data Center Consolidation with De-duplication Technology**
By K. E. H. Polanski

78 **Safeguarding Data with Dell PowerVault Data Protection Solutions**
By Sanjeet Singh and Jason Buffington

81 **A Look Inside Microsoft System Center Data Protection Manager 2007**
By Jason Buffington and Sanjeet Singh

84 **Cost-Effective Archiving with the Dell PowerVault RD1000**
By Jeff Boles

SYSTEMS MANAGEMENT

88 **Facilitating Microsoft Windows Vista Migration Through Application Virtualization**
By Coby Gurr

DELL DATA CENTER SOLUTIONS PAGE 66

CREATING A HYPER-EFFICIENT HYPER-SCALE DATA CENTER

By Jimmy Pike, Ty Schmitt, Frank Frankovsky, and Todd Brannon

Hyper-scale computing "clouds" can deliver critical network-centric applications and flexible access to powerful compute and storage resources. The Dell Data Center Solutions team offers a wide range of services to help create integrated solutions designed to minimize acquisition and operating costs, maximize energy efficiency, and enable rapid scalability.

NETWORK AND COMMUNICATIONS

92 **Accelerating Application Traffic with Intel 10 Gigabit Ethernet Server Adapters**
By Jordan Plawner and Travis Vigil

HIGH-PERFORMANCE COMPUTING

96 **Microsoft HPC Solutions Boost Manufacturing Innovation**
By Aditya Krishnan

SECURITY

100 **Securing Virtualized Environments with McAfee IntruShield**
By John Vecchi

SPECIAL SECTION: PARALLEL APPLICATION DEVELOPMENT SOLUTIONS

112 **Extending OpenMP to Clusters**

POWER SOLUTIONS

DELL.COM/PowerSolutions

EDITORIAL STAFF

EDITOR-IN-CHIEF AND PUBLISHER Tom Kolnowski
MANAGING EDITOR Deb McDonald
FEATURES EDITOR Kathryn White
DEPUTY MANAGING EDITOR Liza Graffeo
SENIOR EDITOR Jim Duncan
EDITORIAL ASSISTANT Amy J. Parker
ART DIRECTOR AND COVER DESIGNER David Chan
DESIGNER AND ILLUSTRATOR Cynthia Webb

STAFF WRITERS Romy Bauer, Jeanne Feldkamp, and Julie Jervis

CONTRIBUTING WRITERS Rizwan Ali; Christian Belady; Santosh Bhadri; Surendra Bhat; Jeff Boles; Sriranjan Bose; Todd Brannon; Jason Buffington; Onur Celebioglu; Glen Curry; Joshua David; Albert Esser, Ph.D.; Chad Fenner; Shawn Fielding; Frank Frankovsky; Aziz Gulbeden; Coby Gurr; Stephen M. Hahn; Ralph Hilliard; John Joseph; Kailash Karki, Ph.D.; Aditya Krishnan; John Loffink; Todd Mitchell; Tom Painter; Christos Pattichis; John Pflueger, Ph.D.; Jimmy Pike; Jordan Plawner; K. E. H. Polanski; Aaron Prince; Paul Rad; Ramesh Radhakrishnan, Ph.D.; Vishvesh Sahasrabudhe; Amina Saify; Abhay Salunke; Ty Schmitt; Eric Schott; Carthikeyan Shanmuganathan; Timothy Sherbak; Lokesh Singh; Sanjeet Singh; Fred Stack; Eric Szewczyk; Frank Tang; Max Thoene; Bharath Vasudevan; John Vecchi; Travis Vigil; Tim Webb; Kevin Wittmer; Jun Yang; and Rodan Zadeh

CONTRIBUTING PHOTOGRAPHERS Tony Bolding, Bryan Kuntz, Adran Matte, and Joey Pena

ONLINE DESIGN Joi Chevalier
SPECIAL INSERTS MANAGER Kristin Kreisel

SUBSCRIPTIONS

Subscriptions are complimentary to qualified readers who complete the online subscription form. To sign up as a new subscriber, renew an existing subscription, change your address, or cancel your subscription, access the online Subscription Center forms at **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, listens to its customers and delivers innovative technology and services they trust and value. Uniquely enabled by its direct business model, Dell is a leading global systems and services company and No. 34 on the Fortune 500 list. For more information, visit our Web site at **DELL.COM**.

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

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

February 2008

STOCK NUMBER: DPS 2008-02 (FEB 08)

MOVING ALONG THE GREEN IT CONTINUUM

Tom Kolnowski
Editor-in-Chief and Publisher
tom_kolnowski@DELL.COM
DELL.COM/PowerSolutions
DELL.COM/RSS

Enterprises planning to pursue a green data center infrastructure have a broad range of options to choose from. For some, the answers may lie in building entirely new, green-spec data center shells that are architected from the ground up to be energy efficient, environmentally responsible, and highly sustainable. Approaches for extending the green envelope here are numerous and varied and may include incorporating recycled and renewable building materials in construction, using renewable energy sources, taking steps to safeguard indoor environmental air quality, installing highly efficient room- and rack-level cooling systems, and deploying the latest energy-efficient servers and networked storage systems. Green certification of a new data center through the U.S. Green Building Council's Leadership in Energy and Environmental Design (LEED) program can also be a worthy goal, as can following best practices defined by the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE).

But short of rearranging bricks and mortar, there is much organizations can do within their current data centers to make healthy progress along the green IT continuum. A desire to advance a pragmatic approach to green IT was the impetus for our feature section in this issue, Greening the Data Center. Our cover story, "The Energy Smart Data Center" (starting on page 10), and a companion article, "Best Practices for Unlocking Your Hidden Data Center" (starting on page 16), provide insight and guidance on how organizations can make meaningful power, cooling, and resource utilization changes within the four walls of their existing data centers—changes that can yield significant, positive green IT results and may even defer or help avoid the need for a planned data center expansion. Four more articles in the feature section cover various green IT topics, including an inside look at The Green Grid consortium (page 20) and methods of boosting data center density without increasing energy usage using Dell-Liebert Energy Smart Solutions (page 32). Three Web-exclusive articles on green IT are also available for download at DELL.COM/PowerSolutions.

In addition, we are pleased to announce our new Case Study section (starting on page 104), where we profile companies and organizations worldwide that have enjoyed success simplifying their IT infrastructures with Dell™ solutions. Featured in this issue are e-commerce provider MarketLive (United States), online retailer Tesco.com (Europe and Asia), health care provider Taiwan Adventist Hospital (Taiwan), and Ryerson University (Canada). Organizations interested in participating in the case study program can contact the Dell Worldwide Customer Reference Program team at references@DELL.COM.

DELL #1

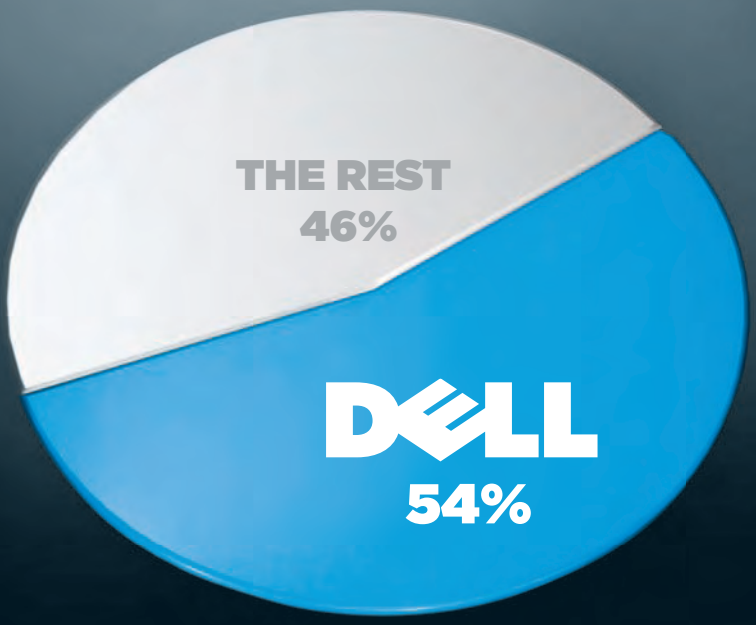

According to Forrester Research, more than half the enterprise companies in North America and Europe rely on Dell for notebook and desktop computers.
Who do you rely on?

–How Enterprise Buyers Rate Their PC Suppliers And What It Means For Future Purchases, Forrester, November 2007.

GET YOUR FREE COPY OF FORRESTER'S REPORT AT
DELL.COM/Numberone

Actual Forrester quote: "Dell is clearly the No. 1 enterprise desktop and laptop supplier." Survey question: "From which vendor did you purchase desktops in the last 12 months?" Base: 565 PC decision-makers at North American and European Enterprises.

DELL EQUALLOGIC OFFERINGS TO DELIVER STORAGE PRIMED FOR SERVER VIRTUALIZATION

Businesses of all sizes are constantly looking for ways to simplify management and lower the cost of enterprise storage. To advance these objectives, Dell has recently inked an agreement to acquire EqualLogic, a leading provider of high-performance storage area network (SAN) solutions optimized for server virtualization.

Built on a patented peer storage architecture, the award-winning[1] Dell™ EqualLogic® Internet SCSI (iSCSI) arrays offer world-class performance, scalability, and reliability on a platform that is easy to install, maintain, and manage. This enables organizations of all sizes to cost-effectively consolidate storage while helping protect data and applications for enhanced business continuity and disaster recovery.

Dell EqualLogic iSCSI arrays are also optimized for virtualization, allowing them to be easily integrated into Dell and other server virtualization solutions. By deploying Dell EqualLogic shared storage, organizations can reap the full benefits of virtual storage pools, enhanced storage management, and high levels of automation while helping streamline the migration to a virtualized IT infrastructure.

"Our customers will be dealing with the largest increase in data we have seen in our history over the next few years," says Michael Dell, chairman and CEO of Dell. "Leading the iSCSI revolution will help Dell accelerate IT simplification and virtualization and will drive the Dell value proposition into more areas of the enterprise storage business."

At press time, Dell announced plans to grow EqualLogic's successful channel-partner programs with current and future Dell EqualLogic offerings, and also make Dell EqualLogic offerings available direct from Dell. For in-depth technical coverage on the Dell EqualLogic arrays, turn to page 36.

[1] The EqualLogic PS3800XV storage array was awarded SAN Product of the Year at the Techworld Awards 2007, by vote of a judging panel made up of IT industry experts and Techworld editors who evaluated the PS3800XV on the criteria of strategy, creativity, innovation, and effectiveness; see www.event-space.com/techworld2007/winners2007.asp?m_pid=0&m_nid=20945.

EVERDREAM ACQUISITION EXTENDS REACH OF REMOTE MANAGEMENT SERVICES

Dell recently completed its acquisition of Everdream Corporation, a leading provider of software-as-a-service (SaaS) solutions for remote-service management. By helping organizations support a highly mobile workforce, these offerings are key to the Dell IT simplification strategy. Dell plans to extend its SaaS offerings to include remote management of globally distributed end-user devices such as desktops and notebooks in addition to servers, storage, and printers.

Built around a powerful, scalable, and cost-effective SaaS platform model, the comprehensive set of Everdream software services is delivered through a single integrated management console—enabling businesses to manage their mobile, distributed workforces over the Web virtually anytime.

"The Everdream acquisition is a significant step in Dell's next-generation service offerings and provides Dell an expanded presence in the rapidly growing SaaS-enabled managed services market," says Steve Schuckenbrock, senior vice president at Dell. "Everdream has developed industry-leading technology that enables customers and partners to easily and affordably manage PCs from anywhere in the world."

Everdream products will continue to be developed for and offered through channel partners, and will also serve as the foundation for future Dell service offerings. For more information on Dell and Everdream, visit www.everdream.com.

THE ENERGY SMART DATA CENTER

Going green can be the secret to significant cost savings as well as aggressive performance growth. Dell offers a comprehensive strategy that includes virtualization and consolidation onto energy-efficient systems, best practices for power and cooling optimization, and expert services that can help businesses achieve immediate benefits.

By John Pflueger, Ph.D.
Albert Esser, Ph.D.

Related Categories:

Dell Data Center Solutions (DCS)
Dell Infrastructure Consulting Services (ICS)
Dell OptiPlex desktops
Dell PowerEdge servers
Dell PowerVault storage
Energy efficiency
The Green Grid
Green IT
Power and cooling
Simplify IT
Virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

Companies go green for a variety of reasons. Thanks to a 56 percent increase in global electricity prices since 2002,[1] green initiatives are often conceived as a response to the rising cost of energy. Some businesses go green simply to improve customer and public perception of their organization. Still others are driven by environmental concerns to minimize the impact of global carbon dioxide emissions.

Whatever the motivation, IT is a key component of green initiatives.[2] IT infrastructure accounts for a disproportionate share of energy relative to head count and operating costs—and these energy issues can inhibit business innovation and growth. Power and cooling capabilities are maxed out in many data centers, and power and cooling expenses frequently outstrip the cost of IT equipment. IT managers are also grappling with operational issues such as difficult-to-manage "hot spots" while trying to meet the needs of an ever-increasing server population.

Conflicts between IT and facilities management groups are often the cause of inefficient practices in the data center. These conflicts typically surface when organizations are planning for upcoming data center changes such as consolidations and expansions, new redundancy requirements, or incremental power requirements. IT managers are typically concerned with the ability of the IT facility—including its space, power, and cooling capacity—to fully support server consolidation, virtualization, or high-performance computing initiatives, and the ability of existing server room and data center cooling systems to support new equipment. For example, in one typical enterprise usage scenario, power distribution and cooling equipment consumed 59 percent of data center power, whereas compute servers accounted for less than 30 percent (see Figure 1).

Existing server room and data center power distribution systems cannot always support future growth—and organizations frequently cannot afford the costs associated with a major business disruption. Meanwhile, facilities management groups have different priorities when planning data center changes. They must consider power usage caps from the utility company as well as energy costs, which often exceed equipment costs. And in many cases, companies are literally running out of space in the data center because existing server racks are full. Moreover, the data center may be running out of circuit breakers. Facilities managers often cannot deploy any additional servers until a new data center is brought online.

SIMPLIFYING THE GREEN DATA CENTER

Because energy consumption is about much more than just the IT equipment, Dell takes a comprehensive approach to greening the data center. Dell engineers start by designing energy-efficient platforms with built-in power management features—and suppliers are encouraged to do the same. Examples include high-efficiency power supplies and optimized thermal design, and can also include silicon changes, component changes, and energy-efficient motherboard design. (For details on other aspects of Dell environmental initiatives, see the "Green, not greenbacks" sidebar in this article.)

Dell also helps organizations optimize the office IT environment through effective client power management (see the "Best practices: Client usage policies that pay off" sidebar in this article). In addition, data center optimization measures such as enhancing system utilization through virtualization and workload management technologies, highly efficient cooling architectures, and energy-efficient storage design can complement these efforts and enhance efficiency.

A positive change in the energy usage of IT equipment can have a direct effect on power and cooling requirements (see Figure 2). For example, in an internal Dell data center usage scenario, a 10 percent improvement at the server level netted about an 8 percent improvement at the facility level. However, efficiency improvements in cooling or power delivery were relatively independent; a 10 percent increase in power delivery efficiency manifested itself as about a 4 percent improvement at the facility level.

Dell Infrastructure Consulting Services experts can help data center managers pinpoint facility and IT improvements that enable efficient operation. For example, simply running the facility at a

[1] "Electricity Prices for Industry," by the Energy Information Administration, U.S. Department of Energy, June 7, 2007, www.eia.doe.gov/emeu/international/elecprii.html.

[2] For an executive perspective on why green data center design is gaining momentum and tactics for immediately putting unused IT capacity to work, see "Best Practices for Unlocking Your Hidden Data Center," in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080198-Esser.pdf; and DELL.COM/HiddenDataCenter.

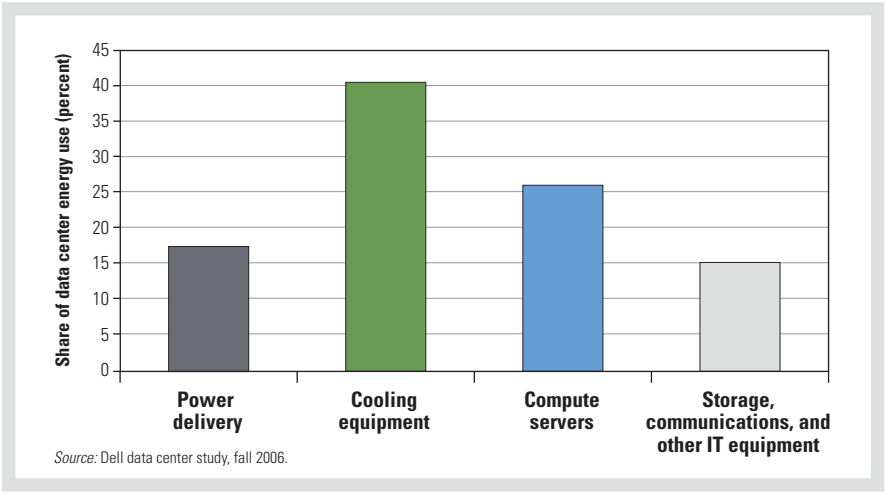

Figure 1. *Power distribution and cooling equipment consumed 59 percent of data center power in a typical enterprise usage scenario*

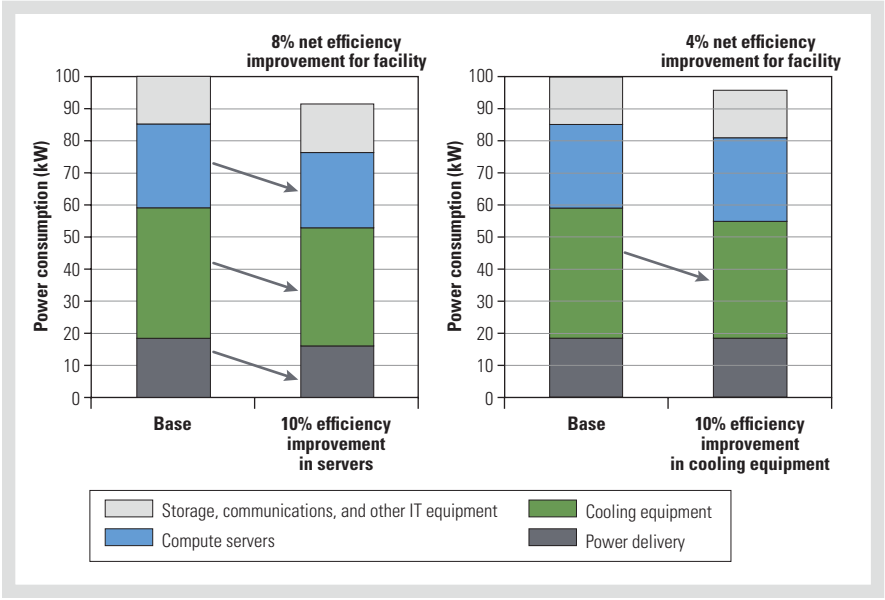

Figure 2. *Increased server efficiency can lead to improvements in overall data center efficiency*

slightly higher temperature than usual can often provide an opportunity to increase efficiencies. Coupled with air-handler energy-efficiency options, savings of 5 percent or more at the facility level are not uncommon. Savings may result from staged compressor operation, slowdown of airflow with variable-frequency drives, increased chiller efficiency, or minimized condensation and re-humidification in the air handler.

ADVANCING EFFICIENCY THROUGH VIRTUALIZATION AND CONSOLIDATION

Through virtualization and consolidation, enterprises can reduce the number of physical servers in the data center while dramatically increasing compute capacity. Virtualized servers can provide more efficient space utilization, more effective energy consumption in the remaining servers, and lower overall power and cooling requirements than traditional non-virtualized servers (see Figure 3).

Virtualization typically allows organizations to consolidate anywhere from 3 to 20 virtual systems onto a single physical platform. Results vary greatly based on the particular applications, user demand, and configuration of both the old and new servers; Dell customers commonly report 10:1 consolidation ratios. Because they are designed to facilitate dense computing environments—with substantial memory to help remove barriers to running

GREEN, NOT GREENBACKS

Contrary to popular belief, greening the data center does not necessarily mean revamping server and storage infrastructures. Dell offers several non-technology options to help organizations reduce environmental impact.

Carbon offset programs such as the Dell "Plant a Tree for Me" and "Plant a Forest for Me" initiatives make it easy to help offset carbon emissions associated with the electricity required to power computer equipment. Through these programs, Dell partners with the Conservation Fund and Carbonfund.org to plant trees in sustainably managed reforestation projects.

By shipping multiple products in a single box with reduced or reusable packaging, Dell can help simplify deployment and reduce environmental impact. When organizations select the "Multipack" option for consolidated delivery, they can minimize on-site packaging waste and recycling costs, streamline on-site logistics, and help lower shipping costs. The Multipack option also helps reduce redundant documentation, cables, and other unnecessary extras—helping both speed deployment and reduce deployment costs.

In addition, Dell offers a variety of asset recovery programs to help ensure proper disposal of retired IT equipment. Besides being environmentally friendly, asset recovery allows organizations to maximize the residual value of systems they no longer use. For more information about Dell environmental initiatives, visit DELL.COM/Earth.

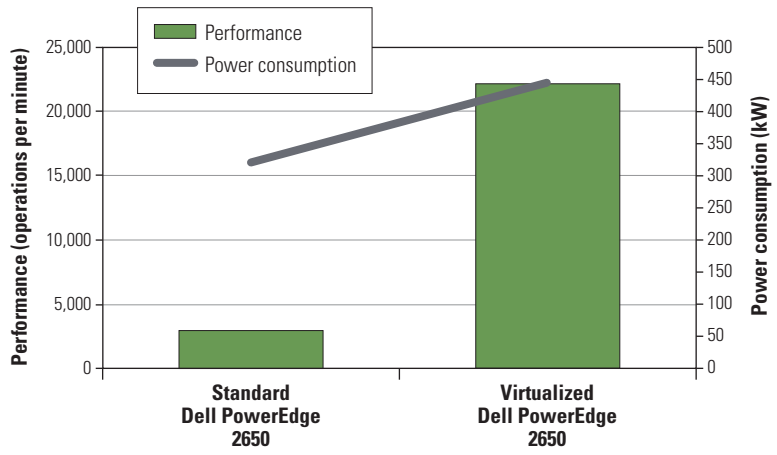

Source: "Competitive Power Savings with VMware Consolidation on the Dell PowerEdge 2950," by Scott Hanson, Dell Enterprise Technology Center, August 2007, DELL.COM/Downloads/Global/Solutions/comppowersavings.pdf.

Figure 3. *Virtualization increases server utilization and helps lower power and cooling requirements*

memory-bound applications—Dell™ PowerEdge™ servers can be the driving force behind a variety of data center efficiency improvements.

MANAGING PERFORMANCE PROACTIVELY

Replacing legacy servers with higher-performance, higher-capacity systems designed with efficient thermal characteristics and power delivery can be an effective way to cut power consumption and costs in the data center. Dell PowerEdge Energy Smart 1950, PowerEdge Energy Smart 2950, and PowerEdge Energy Smart 2970 servers—industry-standard systems specifically configured to optimize energy efficiency—can help deliver dramatic savings by reducing power consumption through features such as low-voltage processors, high-efficiency memory, high-efficiency power supplies, and energy-optimized BIOS settings. Dell studies show that a PowerEdge Energy Smart system configuration can deliver up to 21 percent greater performance per watt over a similarly configured server[3] and enable savings of up to US$200 per server per year in energy expenses.[4]

Today, the power and cooling challenge is focused squarely on servers. However, it is important to proactively manage storage to help prevent potential power and cooling issues. With data volumes growing exponentially, increasing volumes of storage will likely be required in virtually every organization. Right-sizing storage systems, drive mixes, and media to information usage can help achieve considerable power savings. Dell offers expertise and services that can help organizations match application profiles with a range of energy-efficient storage options for enhanced utilization and efficiency.

Additionally, because up to 60 percent of the power consumption in a storage array may be attributed to disk drives, reducing the number of physical drives by using centralized storage—such as the Dell PowerVault™ MD3000i Internet SCSI

BEST PRACTICES: CLIENT USAGE POLICIES THAT PAY OFF

By reaching beyond the data center to implement simple usage policies for power management throughout the enterprise, organizations can achieve significant energy conservation and put dollars back into the bottom line immediately. These common-sense strategies may seem small, but they can make a big difference:

- **Activate power management settings:** Dell Energy Smart power management settings can be turned on before client systems are shipped from the factory, enabling organizations to reduce power consumption by up to 78 percent using thermally optimized Dell OptiPlex™ Energy Smart 755 desktop systems compared with the previous generations of OptiPlex desktops. Similarly, energy-efficient Dell Latitude™ Energy Smart notebooks consume up to 70 percent less energy on average than previous-generation OptiPlex GX620 desktops.*
- **Stagger startups:** Creating a staggered schedule for powering up computers helps avoid overburdening network servers and triggering increased costs for peak energy demands.
- **Turn off desktops at night:** During the approximately 100 hours per week that computers are unused, organizations should switch them off to avoid creating heat and increasing demand for power and cooling throughout the building.

*Based on product specifications and 7 hours of typical office use, 1 hour of maximum-performance use, 1 hour of idleness, and 15 hours of sleep for 264 days a year, with 24 hours of sleep for the remaining 101 days each year. For more information, visit DELL.COM/Energy.

[3] Based on AC power measurements using an Extech 380803 Power Analyzer taken during the peak load of the SPECjbb2005 benchmark test performed by Dell Labs in March 2007. The PowerEdge Energy Smart 2970 configuration consisted of two dual-core AMD Opteron™ 2212 HE processors at 2.0 GHz; eight 1 GB, 667 MHz double data rate 2 (DDR2) error-correcting code (ECC) 1R dual in-line memory modules (DIMMs); and two 73 GB, 10,000 rpm, 2.5-inch Serial Attached SCSI (SAS) drives running the Microsoft® Windows Server® 2003 Enterprise x64 Edition OS, compared with a PowerEdge 2970 standard configuration consisting of two dual-core AMD Opteron 2212 processors at 2.0 GHz; eight 1 GB, 667 MHz DDR2 ECC 1R DIMMs; and two 73 GB, 10,000 rpm, 2.5-inch SAS drives running the Windows Server 2003 Enterprise x64 Edition OS. Actual performance and power consumption will vary based on configuration, usage, and manufacturing variability. For the latest SPECjbb2005 benchmark results, visit www.spec.org/jbb2005.

[4] Based on energy costs obtained from the International Energy Agency using a worldwide average of US$0.13/kWh based on 2003 and 2004 global average cost of energy. Assumes server operation 24 hours a day, 365 days a year. Actual AC power measurements were obtained using an Extech 380803 Power Analyzer taken during the peak load of the SPECjbb2005 benchmark test performed by Dell Labs in November 2006. The PowerEdge Energy Smart 2950 configuration consisted of two dual-core Intel® Xeon® 5148LV processors at 2.33 GHz; four 1 GB, 667 MHz fully buffered DIMMs; and two 73 GB, 10,000 rpm, 2.5-inch SAS drives running Windows Server 2003 Enterprise x64 Edition, compared with a PowerEdge 2950 standard configuration consisting of two dual-core Intel Xeon 5140 processors at 2.33 GHz; four 1 GB, 667 MHz fully buffered DIMMs; and two 73 GB, 10,000 rpm, 3.5-inch SAS drives running Windows Server 2003 Enterprise x64 Edition. Actual performance and power consumption will vary based on configuration, usage, and manufacturing variability. For the latest SPECjbb2005 benchmark results, visit www.spec.org/jbb2005.

(iSCSI)/Fibre Channel storage area network (SAN) or the Dell PowerVault NX1950 unified network attached storage (NAS) and SAN data storage system—can also help lower energy costs. Eliminating boot drives in servers using boot from SAN can yield significant savings for a large fleet of servers. iSCSI helps lower the cost of deploying a SAN by using an Ethernet-based network compared with a Fibre Channel fabric—while single-instance storage, a feature that is part of the Dell NAS family, helps eliminate duplicate files.

SCALING COMPUTE DENSITY COST-EFFECTIVELY

Traditional raised-floor data centers typically have about a 5 kW limit per rack—so, for example, even if 75 percent of the physical space in a rack remains vacant, the 5 kW limit would mean that the rack is operating at its full capacity. Dell-Liebert Energy Smart Solutions are designed to help organizations increase computing density within the same power envelope. By helping reduce fan power, reduce mixing and short-circuiting of air, and eliminate the negative effects of condensation, alternative cooling products can be an energy-conscious alternative to traditional raised-floor cooling.

For example, Liebert XD cooling systems from Emerson Network Power provide a pumped refrigerant infrastructure with cooling modules that can be placed directly above or alongside high-density racks. The Liebert XD system can be up to 30 percent more efficient than traditional perimeter cooling, helping to overcome cooling-related density limitations so that organizations can increase computing capacity in cost-effective increments. In this way, Dell-Liebert Energy Smart Solutions help maximize data center layout and density, allowing organizations to defer the cost of expensive data center leases or build-outs.[5]

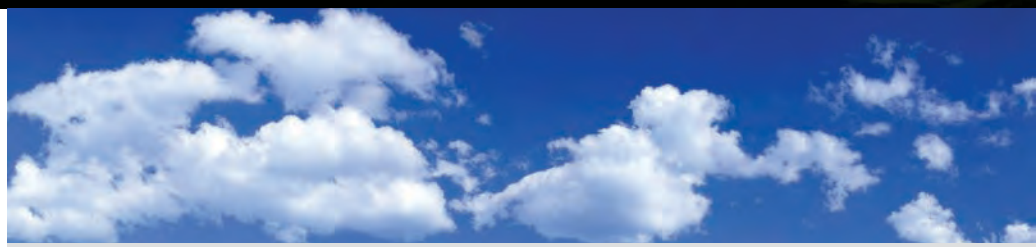

SIMPLICITY AND EFFICIENCY: GREEN IN ACTION

The Dell Cloud Computing™ Solution is expected to cut Ask.com's power bills while supporting ever-higher computational demands.

How can organizations increase their computing capacity with the same power expense? That was the key question for Ask.com, the popular online search engine that helps millions of people find information on the Web every day.

With its advanced ExpertRank ranking algorithm that helps identify the most authoritative sites on the Web—a process that demands many additional calculations that other search engines do not perform—Ask.com requires an enormous amount of processing power and uses thousands of servers to run its search engine. These servers are linked into clusters, or "clouds," that are spread across several data centers. Because the company was committed to reducing its environmental impact, it adopted energy-efficient servers throughout—and then went one step further. The Ask.com IT group decided to locate one of its new data centers in eastern Washington, where it could use renewable hydroelectric power, and developed a novel evaporative cooling system that requires less electricity than a traditional cooling system. But it was still searching for other ways to offset its carbon output.

IDENTIFYING ENERGY SMART BENEFITS WITH EXPERT SERVICES

Dell Energy Smart Data Center Assessment services, offered by the Dell Infrastructure Consulting Services team, are designed to help identify potential problems or opportunities for improvement in data center efficiency. Identifying inefficiencies is the first step toward achieving efficient utilization of power, cooling, space, servers, storage, and communications equipment. Improved power usage practices help mitigate the risk of hardware failure due to overheating and can help reduce energy expenses. Organizations may also be able to consolidate and expand IT capacity without building new facilities or exceeding existing heating and cooling limits.

During a Dell Energy Smart Data Center Assessment, Dell experts conduct a review of the organization's existing power and cooling capacities and practices. They then perform computational fluid dynamics analysis to highlight problem areas.[6] At the completion of the review, an organization receives an assessment report with practical advice for maximizing power and cooling in that specific environment. Virtualization Readiness Assessment services also are available to help customers plan effectively, implement rapidly, and maintain efficiently.

CULTIVATING A STRATEGY FOR SUSTAINABLE GROWTH

Whatever an organization's motive is for going green, Dell can deliver products and

[5] For more information about Dell-Liebert Energy Smart Solutions, see "A Systems-Level Approach to Efficient Data Center Design," by Fred Stack, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080216-Emerson.pdf.

[6] For more information, see "Best Practices for Increasing Data Center Energy Efficiency," by Paul Rad, Max Thoene, and Tim Webb, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080185-Rad.pdf; and "High-Efficiency Cooling Through Computational Fluid Dynamics," by Paul Rad; Kailash Karki, Ph.D.; and Tim Webb, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080187-Rad.pdf.

> **"By partnering with Dell on multiple green initiatives, we can help ensure that our data centers minimize energy use and waste. We are pleased to be using less and giving back more."**
>
> —Chuck Geiger
> Executive vice president of technology and engineering at Ask.com
> December 2007

Dell Data Center Solutions (DCS) approached Ask.com with a proposal for a new strategy: let DCS custom-design a purpose-built rack server—optimized for the Ask.com data centers—that would have very low electricity needs yet deliver high computing power.

Initially, says Chuck Geiger, executive vice president of technology and engineering at Ask.com, the company just wanted straightforward optimizations to help minimize power draw. However, Dell DCS engineers went a step further, designing a server around Ask.com's specific workload requirements and selecting only the components required. To streamline it, Dell DCS removed all unnecessary components, including PCI slots and the secondary power supply, which is a helpful backup when a server is working alone but typically unnecessary in a cloud computing configuration.

"In the system we co-developed with Dell DCS, we were able to right-size the power supply precisely for the parts in our design," says Geiger. "Compared to the servers we were using before, the energy bills for these are 18 percent lower. And with the thousands of machines we use, that 18 percent savings adds up."

Dell and Ask.com continue to optimize the cloud computing solution. For example, the two companies are testing other server components, which potentially could be another 15 percent more energy-efficient and need fewer repairs compared with the current solution.

To help reduce packaging waste, new servers arrive at Ask.com in Dell Multipack bundles. Ask.com also became the first corporate customer to participate in the Dell "Plant a Forest for Me" program. Through partnerships with Dell, the Conservation Fund and Carbonfund.org plant trees in managed reforestation projects to sequester carbon. To help offset the carbon emitted by Ask.com, those organizations plan to plant thousands of trees on behalf of Ask.com.

"While we are always evaluating our server options, this has been a good process for us so far, what with the savings in power, maintenance, and packaging; high performance and reliability; and opportunity to deepen our company's working relationship with Dell," Geiger says. "The Dell DCS team really listened to us. They have worked with us closely to provide the best solution, in a collaborative co-development process. And we are extremely happy with the results."

For more information about Dell DCS Cloud Computing Solutions, visit DELL.COM/CloudComputing.

services to help make green goals a reality. Energy-efficient features are built into Dell products from the desktop to the data center. Dell can also help organizations extend the benefits of efficient equipment through intelligent data center design—enabling them to optimize power and cooling configurations to achieve high computing density without performance degradation.

Dell partnerships with industry leaders and industry consortiums such as The Green Grid; the Climate Savers Computing Initiative; the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE); the Storage Networking Industry Association (SNIA); the Distributed Management Task Force (DMTF); and the Standard Performance Evaluation Corporation (SPEC) create broad-based opportunities for enterprises to address their energy efficiency issues by leveraging cost-effective, standards-based power and cooling equipment.[7]

John Pflueger, Ph.D., manages the Dell technical strategy for facility and system thermals, focusing on ways to improve efficiency. He has 16 years of experience in product development, product marketing, and product management. John has a Ph.D. from the Massachusetts Institute of Technology.

Albert Esser, Ph.D., serves as vice president for data center infrastructure at Dell, where he is responsible for enhancing Dell's enterprise-class IT solutions by sharing customer insights with the company's Server, Storage, Data Center Solutions, and Services teams. Albert has an M.S. and a Ph.D. in Electrical Engineering from the University of Aachen. He holds 12 U.S. patents.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell energy efficiency initiatives:
DELL.COM/Energy

Dell power and cooling solutions:
DELL.COM/HiddenDataCenter

Dell Infrastructure Consulting Services:
DELL.COM/ICS

Dell environmental initiatives:
DELL.COM/Earth

Dell IT simplification:
DELL.COM/Simplify

[7] For more information, see "The Green Grid: Enabling the Energy-Efficient Data Center," by Christian Belady and John Pflueger, Ph.D., in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080199-GreenGrid.pdf.

BEST PRACTICES FOR UNLOCKING YOUR HIDDEN DATA CENTER

Dr. Albert Esser, vice president of data center infrastructure at Dell, shares his perspective on the momentum that is building behind green data center design. Plus: How IT organizations can put unused capacity to work and leverage tactics for fast, flexible growth that helps maximize the bottom line.

Get ready for a major shift in the way data centers operate: according to Gartner, "By 2008, 50 percent of current data centers will have insufficient power and cooling capacity to meet the demands of high-density equipment."[1] Many companies may respond to this challenge by performing a major data center facility upgrade in the next several years.

Concerns about power consumption are behind much of this momentum to upgrade. Operating high-density equipment such as blade servers demands an enormous amount of energy. For example, rack enclosures that accommodate the equivalent of 60–70 1U servers typically require 20–25 kW of power per rack.

Many data centers do not have the power or cooling capacity to extend this level of rack density across the raised-floor area to help meet growing performance demands. Traditional data centers have been provisioned to supply 40–60 W per square foot, but densities of 20–25 kW per rack would require power and cooling capacity of up to 600 W per square foot—a level that typically exceeds acceptable efficiency or cost levels.

In addition to imminent power and cooling problems, growing awareness about the environmental impact of excessive energy usage is driving business leaders to explore ways to make IT infrastructures more efficient than they are now. Today, compute servers typically account for 26 percent of total energy consumption in a typical data center, while a whopping 59 percent of total energy is consumed by power distribution and cooling equipment.[2]

Dr. Albert Esser, vice president of data center infrastructure at Dell, recently spoke with *Dell Power Solutions* about the connections between green IT, data center performance, and Dell's flexible approach to creating an environmentally friendly infrastructure.

WHAT DOES GOING GREEN ENTAIL ON A TACTICAL LEVEL?

There's a common misconception that going green means buying lots of new, more energy-efficient equipment. In fact, in many cases, the easiest way to be more environmentally friendly is to consolidate your infrastructure onto fewer servers.

Most data centers are not even near capacity, and there is an enormous potential savings there—a whole "hidden data center." A lot of data centers run at maybe 15 to 30 percent CPU utilization without realizing it. But there is so much more capacity in those servers. Even running at 60 percent, those servers are often not fully loaded.

[1] "The Data Center Power and Cooling Challenge," by Michael Bell, Gartner, Gartner Symposium/ITxpo 2007, San Francisco, CA, April 22–26, 2007.

[2] Dell data center study, fall 2006. Data is based on one example data center. While energy consumption in typical data centers will likely be close to these estimates, the percentages could be significantly higher or lower in other cases.

HOW DOES VIRTUALIZATION HELP MAKE DATA CENTERS GREENER?

For many companies, the current mentality is to buy servers. Your data center is planned around servers and infrastructure for servers, the number of servers, what kind of servers you have in there. One application per server used to be the norm—and once a system was commissioned, you didn't mess with it. These guidelines served us well in the past. But now it's clear that, in both cases, some efficiencies and productivity are left on the table. Requiring one application per server frequently leads to equipment running at low utilization rates—wasting processor power. Leaving legacy systems in place results in server populations that are not leveraging the best that today's technology has to offer.

What you really need is *application space*. So, say you have a facility that needs to support your enterprise resource planning implementation, your document management applications, and your customer relationship management database. If you deploy each of those workloads on virtual machines (VMs), you don't have to worry about buying servers every time you add something new. You basically just look at how fully your servers are loaded, and if you're running out of space, you just buy capacity.

Dell has a virtualization assessment service that helps organizations consolidate servers and storage with excess capacity, which also lets them turn off outdated equipment or replace aging systems with new ones. You don't have to buy all-new Dell hardware to make your data center greener. It's not a rip-and-replace model. Virtualization features are built into Dell servers. For example, iSCSI storage systems such as the Dell PowerVault MD3000i modular disk array enable seamless VM mobility across physical servers by allowing VMs to be mapped directly to shared storage.

Organizations can typically fit anywhere from 5 to 20 VMs on a single physical platform, depending on the application, user demand, system configuration, and the configuration of the old or replaced servers. A typical ratio is 8:1. Our customers frequently report 10:1 or 12:1 ratios as well.

HOW DO GREEN INITIATIVES AFFECT DATA CENTER PERFORMANCE?

Sometimes people get fooled by the idea that you have to choose between being green and being effective. You don't. The two concepts are related—more effective data centers tend to be greener, and greener data centers tend to be more effective.

It's common sense, really. The best hardware in the industry is the most cost-effective and energy efficient, and that makes for the greenest data center you can have. And it goes the other way too—if your data center is not green, it isn't as efficient or effective as it could be.

City buses are a great analogy. If you're in the bus business, you buy big buses because they are an efficient way to transport lots of people. But the bus isn't efficient if you have only two or three people on it all the time. That is what we do right now with our data centers. Businesses have servers with lots of capacity, but they may use only a fraction

HOW EFFICIENT SERVER UTILIZATION HELPS SAVE ENERGY

When a server is underutilized, it uses a disproportionate amount of power and cooling. But as utilization increases, the additional power and cooling required to meet computational needs is relatively lower.

Two tactics can help improve this situation. First, companies may increase utilization rates, either through consolidation or virtualization. Power throttling and putting servers in sleep mode or turning them off when they are idle also can help conserve energy, cut costs, and boost the bottom line.

Consider this example: at 20 percent utilization, a Dell™ PowerEdge™ 2950 server requires 293 W. At 80 percent utilization, the same server requires just 358 W—only a 22 percent increase, despite the 60 percent jump in utilization (see Figure A).*

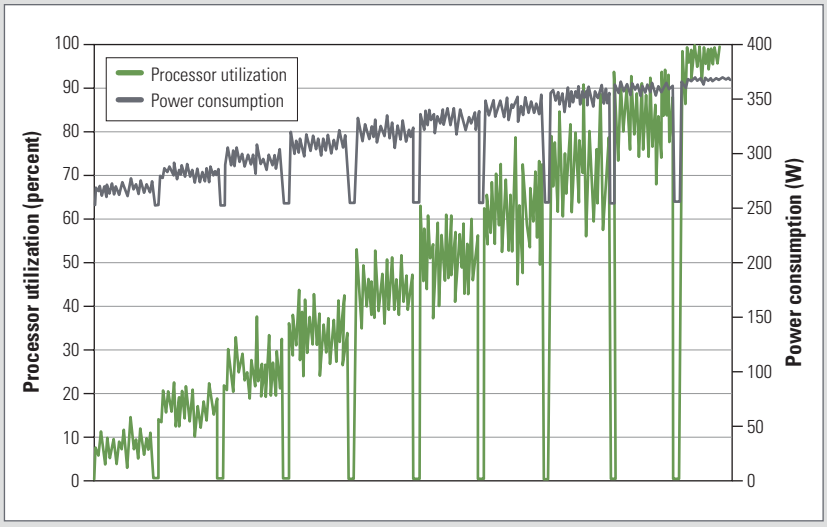

Figure A. *As server utilization increases, the additional power and cooling required to meet computational needs becomes relatively lower*

*Dell Server Performance Analysis lab study, November 2, 2007.

> **"The key thing for CIOs to remember is that they can make their company's IT greener on their own schedule and with their own IT people. Dell's approach is not to come in and take everything over—you choose what you can afford."**

of it. This situation leads to inefficiencies that can make workload consolidation a big deal for green initiatives.

HOW DOES COOLING IMPACT EFFICIENCY IN THE DATA CENTER?

One of the biggest misconceptions about data centers is that colder is better. You don't have to cool your data center like a meat locker. You can actually gain efficiencies in your facility if you allow a higher set point for air and water temperatures.

For example, we've found we can improve airflow efficiency in servers by lowering the flow with variable-frequency drive blowers and allowing a hotter exhaust. That's because any kind of energy that goes into moving air is wasted energy. Pushing air through the system only moves the heat—so fan power is wasted power.

On the air-conditioning side, chiller cooling efficiency improves with increased water temperature. It's almost a linear curve, where for every degree you can raise the temperature of the cooling coil, the more effective and efficient your air-conditioning will be. In many data centers, you can just crank up the temperature on your thermostat five degrees and see up to a 5 percent increase in efficiency. The magnitude of these easy savings is absolutely astounding.

Recently we told a customer that these types of changes would allow four times more servers into a data center than they had anticipated. Not four more servers—but four *times* more servers. And it was a big data center.

WHAT IS DELL'S APPROACH TO GOING GREEN?

Dell can give companies a lot of help on the path to being green, from helping with buying decisions to providing the hardware to making recommendations for improving efficiency. We also have, for one server line, an in-depth study that shows up to 10 to 15 percent better efficiency for Dell servers than for comparable systems from competitors.[3] So just using Dell equipment can help businesses make their data centers greener.

On the operations side, Dell has online resources to help customers with infrastructure questions. We have "Plant a Tree for Me" and "Plant a Forest for Me" programs that give companies a way to offset their data center carbon emissions. We offer consolidated delivery options that can help reduce packaging by shipping multiple products in a single box, and we have an asset recovery program that helps make sure that retired IT equipment is disposed of properly.

At Dell, we believe in simplifying IT—including power and cooling. While the industry has witnessed steadily increasing complexity and cost associated with energy management, we offer simple, cost-effective solutions.

If you want to take your power consumption levels back to what they were five years ago without making drastic changes, Energy Smart Dell PowerEdge 1950, PowerEdge 2950, or PowerEdge 2970 servers are an excellent option. They have about the same power draw as servers used five years ago, but they are designed to be significantly more energy efficient than legacy systems and applications. For example, Dell studies show that a PowerEdge Energy Smart system configuration can deliver up to 21 percent greater performance per watt over a similarly configured server.[4]

Or, in a more comprehensive approach, you can buy the latest and greatest cooling technology. We can hook you up with Liebert and other partners to design a new infrastructure and the tools and services to support it. One way or another, Dell can help make your IT infrastructure as efficient as possible in terms of application utilization—which also translates to low cost of ownership and simplified operations.

The key thing for CIOs to remember is that they can make their company's IT greener on their own schedule and with their own IT people. Dell's approach is not to come in and take everything over—you choose what you can afford. But virtually anything you do with Dell can help reduce the environmental impact of your data center.

Albert Esser, Ph.D., serves as vice president for data center infrastructure at Dell, where he is responsible for enhancing Dell's enterprise-class IT solutions by sharing insights gained from customers with the company's Server, Storage, Data Center Solutions (DCS), and Services teams. Albert has an M.S. and a Ph.D. in Electrical Engineering from the University of Aachen. He holds 12 U.S. patents.

[3] Based on testing conducted by Principled Technologies in December 2006 on comparable configurations of servers from Dell, Hewlett-Packard, and IBM. For detailed results, see "SPECjbb2005 Performance/Watt on Three Servers with Dual-Core Processors," by Principled Technologies, December 2006, DELL.COM/Downloads/Global/Products/Pedge/En/enrg_smrt_perf_brief.pdf.

[4] Dell Datacenter Capacity Planner 2.0 online calculator tool, available at DELL.COM/Energy. For more information, see "The Energy Smart Data Center," by John Pflueger, Ph.D., and Albert Esser, Ph.D., in *Dell Power Solutions,* February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080179-CoverStory.pdf.

"DELL AND EMC...BEST REPUTATION FOR PRODUCT QUALITY"

Nexus Energy Software (April 2007)
Brian Lenane, VP Client Services

Dell/EMC brings you the performance and reliability you need in a storage solution. Our systems are easy to deploy, manage, and upgrade so your IT team can spend less time on maintenance and more time on big ideas.

SIMPLIFY IT BY CALLING (877) 671-3355 OR VISITING DELL.COM/EMC

DELL | EMC²

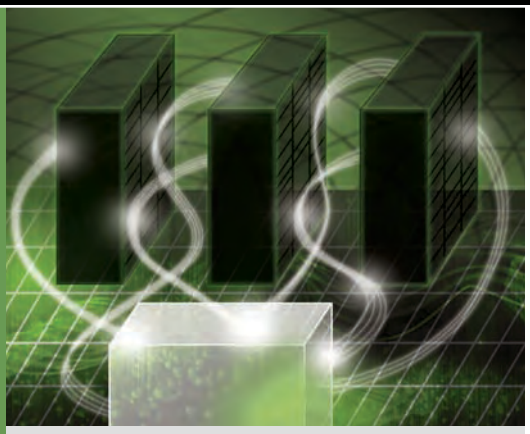

THE GREEN GRID: ENABLING THE ENERGY-EFFICIENT DATA CENTER

By Christian Belady
John Pflueger, Ph.D.

An efficient data center can enhance business performance and help lower costs. The Green Grid—an international consortium dedicated to improving energy efficiency in data centers and business computing ecosystems—is developing platform-neutral standards, metrics, measurement methods, processes, and technologies that promise to help conserve energy for sustainable growth.

Related Categories:
Energy efficiency
The Green Grid
Green IT
Industry standards
Power and cooling

Visit DELL.COM/PowerSolutions for the complete category index.

The emergence of the Internet and the information economy has led organizations to turn to IT for ways to both improve business processes and provide additional value to customers. Now, however, increasing energy costs and decreasing energy availability mean IT organizations must understand the relationship between their activities and the resources they consume.

Servers and data centers used double the amount of energy in 2006 than they did in 2000, according to estimates in a recent report to the U.S. Congress by the Environmental Protection Agency (EPA).[1] Furthermore, about 50 percent of that energy was consumed by the power and cooling infrastructure supporting IT equipment in data centers, the report states. Even under current efficiency trends, the report concludes that by 2011 this consumption could nearly double again to more than 100 billion kWh—a US$7.4 billion annual electricity cost. Figure A in the supplemental online section of this article, available at DELL.COM/PowerSolutions, shows the projected electricity use by IT space type from 2007 to 2011 based on current efficiency trends.

However, cost is not the sole reason organizations are concerned about rising energy consumption. As IT capabilities become increasingly critical to both the enterprise infrastructure and business model, IT departments must be able to meet the growing computing needs of the organization.

UNDERSTANDING COMPETITIVE OPPORTUNITIES TO DRIVE EFFICIENCY

These issues are serious, but they also present opportunities: organizations that effectively manage their cost per application will likely be in a position to competitively differentiate themselves in the marketplace. The first step in the process of understanding and improving energy efficiency is awareness. But while organizations are all too aware of the importance of deploying efficient IT resources, many still lack the tools to understand how to measure the improvements as well as how to determine the real impact of IT energy costs on the business model.

In the past, the total cost of ownership (TCO) for IT equipment was based primarily on acquisition cost. However, to gain an accurate understanding of TCO, organizations must also consider the efficiency of data center operations. Meaningful metrics help provide an objective view of the true cost of data center operations in terms of both energy consumption and

[1] "Report to Congress on Server and Data Center Energy Efficiency: Public Law 109-431," by the U.S. Environmental Protection Agency, August 2, 2007, www.energystar.gov/ia/partners/prod_development/downloads/EPA_Datacenter_Report_Congress_Final1.pdf.

resource utilization. Metrics also help organizations understand the connection between holistic improvements and potential bottom-line savings. This understanding enables these organizations to justify any costs that might be associated with operational changes or energy-efficient IT equipment.

Data center design is also an important aspect of improving efficiency. By adopting industry best practices for design and engineering data centers based on similar design principles used for servers—such as airflow and component configurations—organizations can help minimize hot spots and hot and cold mixing, and optimize space utilization. As a result, effective data center design can play a major role in significantly reducing energy consumption.

New technologies can also play a key role in improving data center efficiency. However, organizations must be able to identify relevant technologies with true potential, understand how best to implement them, and measure the improvement afforded by these technologies.

BENEFITING FROM MULTIDISCIPLINARY EXPERTISE

The Green Grid (www.thegreengrid.org) is an international consortium comprising a wide range of IT companies and professionals, working together to address all of the preceding concerns. The Green Grid objectives include promoting industry awareness, meaningful metrics, design best practices, and emerging technologies. When The Green Grid officially launched in February 2007, its founding members recognized that improving efficiency could not be accomplished by individual experts or companies alone. Today, more than 130 organizations around the world—including software and hardware manufacturers, infrastructure providers, and consulting firms—have joined to spearhead the effort to develop and promote energy efficiency for data centers and information service delivery.

To help achieve these objectives, The Green Grid is taking a broad-reaching approach that includes working out the technical details around measuring and managing for efficiency, and collaborating with other organizations such as the EPA; the U.S. Department of Energy; the American Society of Heating, Refrigerating, and Air-Conditioning Engineers (ASHRAE); and the Distributed Management Task Force (DMTF) to benefit from multidisciplinary expertise. And because energy efficiency is truly a global problem, outreach efforts are also underway in other parts of the world, including Europe and Asia.

In addition to promoting the adoption of energy-efficient standards, processes, measurements, and technologies, The Green Grid also focuses on defining meaningful, user-centric models and developing metrics, standards, measurement methods, processes, and new technologies to help improve performance against the defined metrics. (For more information, see "The Green Grid metrics" in the supplemental online section of this article, available at DELL.COM/PowerSolutions.)

COMPARING KEY DATA CENTER CHARACTERISTICS AND PERFORMANCE SCHEMA

While metrics help organizations measure performance, the next step is to compare and assess improvement. One of The Green Grid's current deliverables is a baseline market study that identifies today's common practices with respect to measuring and managing data center efficiency. With this data, companies can assess where they are relative to the industry. A subsequent deliverable The Green Grid plans to offer in the future is a database that enables organizations to compare energy efficiency performance against a wide range of data center characteristics. Establishing a database enables data identified by metrics to be collected into a tool that can help organizations self-assess on key performance elements, such as data center design or daily operations.

The Green Grid plans on collecting two main sets of data for the database: characteristics of data centers (location, size, server population, types of workload, and so on) and performance schema that revolve around the metrics and their key components. By enabling a wide range of query fields, the database can be designed to help organizations extract compelling and constructive comparisons such as performance expectations in similar climate conditions, in similar-size businesses or market sectors, or under specific types of workload. Sharing this type of knowledge enables organizations to engineer their data centers for optimum efficiency and to base IT investment decisions on measurable improvements geared to drive sound business strategies.

One of the main challenges associated with measuring data center efficiency is that the data currently comes from a variety of equipment in different formats. From spreadsheets to handwritten notes, few organizations can aggregate data for a consistent, reliable view in real time. The Green Grid is chartered to develop standards for how equipment reports data, what sort of data is reported, and the protocols, data transportation fabric, and application formats necessary to realize this vision. These efforts require coordination across many industry organizations.

Improving interoperability in the data center and developing specifications for instrumentation that reports vital information in a standard way can enable organizations to gather information in a meaningful format and analyze results in real time, rather than through the current process, which can take days or even weeks. This approach enhances data center efficiency by reducing the resources required while affording the insights necessary to make smart decisions quickly.

ADVANCING TECHNOLOGIES AND BEST PRACTICES TO ACHIEVE MEASURABLE RESULTS

Another element of The Green Grid strategy is to identify existing, underutilized technologies and emerging technologies that show the potential for data center

efficiency improvements. The Green Grid identifies relevant technologies, provides guidance, and establishes industry enablement plans for their adoption. By collaborating with external agencies and organizations, The Green Grid can help broaden the range of technology options from server and storage architecture to include items related to power distribution and cooling.

The Green Grid membership base, currently over 130 members representing a wide range of industry stakeholders and market segments, enables it to provide an unbiased, consensus-driven perspective on industry questions and challenges around energy efficiency issues. So far, the organization has published recommendations on efficiency metrics, power distribution in data center around AC/DC, and best practices for enhancing operational efficiency. By communicating best practices across the industry now, The Green Grid is creating a baseline for future developments.

The industry has already shown a significant interest in the published metrics, recommendations, and best practices. While many organizations have been unsure how to adopt energy efficiency metrics in the past, today a large number are actively managing their data centers for enhanced efficiency using The Green Grid metrics as a guide. And they are reporting results. For example, Figure 1 shows efficiency improvements for a Microsoft data center measured over a three-year period using one of The Green Grid's data efficiency metrics—Power Usage Effectiveness (PUE). Data is normalized to the first measurement. This Seattle-area data center is eight years old, with about 60,000 square feet of critical space and 5 MW of critical load. Over this three-year period, the data center team reduced the PUE by over 20 percent. Other interesting observations emerge from the graph: for example, there are seasonal peaks during the summer and distinct valleys in the winter. Microsoft tracks PUE daily for all of its data centers, and each new-generation data center is designed to be better than the last. Without the PUE metric, it would be difficult to know whether improvements were being made.

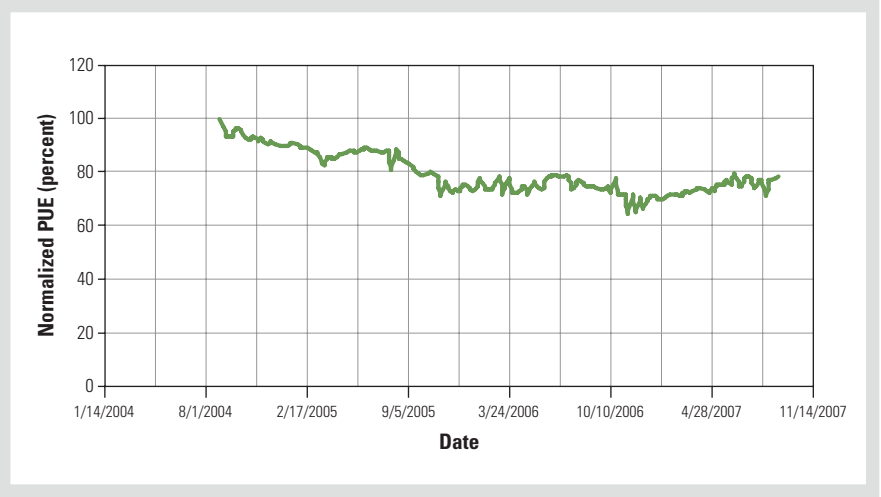

Figure 1. *Example scenario: normalized energy efficiency gains measured by PUE metric*

LIGHTING THE WAY TO ENERGY-SAVVY DATA CENTERS

The Green Grid consortium is also successfully promoting important messages around energy efficiency and providing an industry-wide platform for further discussion around the issues. A key component of this platform is The Green Grid Technical Forum,[2] designed to bring members and other industry stakeholders together to further the advance of energy efficiency in data centers and business computing ecosystems. The forum provides an opportunity for unbiased, critical discussion around energy use and data center efficiency. Industry participants can learn more about the data center metrics, current and emerging techniques for managing data centers for efficiency, and future technical deliverables from the organization. The Green Grid also hosts discussions with leading industry policy makers from around the world, helping build momentum toward addressing important energy issues.

As data centers undergo a period of major growth stimulated by increasing business demand for data processing and storage, the questions around energy efficiency continue to grow as well. Through The Green Grid's efforts, companies should be able to meet these business demands while simultaneously improving efficiency and minimizing TCO. This will become ever more relevant as energy prices rise and organizations recognize the huge competitive advantage that can be gleaned from understanding how to optimize data center efficiency.

Christian Belady is principal power and cooling architect in the Global Foundation Services group at Microsoft. He is also a participant in The Green Grid Technical Committee and the Metrics and Measurement Work Group, and is an alternate on The Green Grid Board of Directors.

John Pflueger, Ph.D., manages the Dell technical strategy for facility and system thermals, focusing on ways to improve the efficiency of computer systems and data centers. He is a participant in The Green Grid Technical Committee and is on The Green Grid Board of Directors.

[2] See www.thegreengrid.org/events/technical_forum.

Network Nirvana

The Secure Choice for Open Convergence

Our integrated, open-architecture approach to understand and manage the priority and security of unified communications doesn't lock you into a particular voice, video or data vendor.

Let us show you how Enterasys Secure Networks™ solutions for switching, routing, wireless and intrusion detection/prevention can automatically discover, classify, prioritize and secure any convergence application without sacrificing performance.

It's VoIP without compromise as we support wired and wireless IP telephony solutions from Avaya, Cisco, Mitel, Nortel, Panasonic, Shoretel and Siemens, while also ensuring compliance with E-911 and CALEA regulations.

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

We Secure *Any* Network

Download our **whitepaper on secure convergence** at www.enterasys.com/convergence to learn more about how we secure and prioritize a phone, camera, printer and computer all running on the same network. **When you're ready, call us at 877-801-7082 or arrange for a live demo** at www.enterasys.com/demo.

enterasys®
Secure Networks™

There is nothing more important than our customers.

INCREASING ENERGY EFFICIENCY WITH DELL/EMC CX3 STORAGE

By Rodan Zadeh

Enterprises must consider multiple variables when designing an efficient storage environment, including performance, power and cooling requirements, and energy costs. Energy-efficient technologies in Dell/EMC CX3 series storage are designed to help enterprises optimize application performance while reducing energy use and controlling ongoing operational costs.

Related Categories:

Best practices

EMC

Energy efficiency

Power and cooling

Visit DELL.COM/PowerSolutions for the complete category index.

Rising energy prices and increases in density and power requirements for enterprise servers and storage have made controlling energy use a key part of data center management. According to the U.S. Department of Energy, the average cost of electricity in the United States increased 36 percent between August 2000 and August 2007, from US$0.0691/kWh to US$0.0968/kWh, with costs typically highest for regions along the densely populated coasts and in Alaska and Hawaii.[1] Enterprises in other parts of the world, including Europe and Asia, may face even greater challenges than those in the United States in terms of energy costs as well as limited power availability and floor space.

The high cost of energy is only one of the factors contributing to rising operational costs. The amount of power and cooling needed for each square foot of data center floor space is also increasing, as racks are filled with increasingly powerful systems designed for the same footprint as their predecessors. And while IT managers have typically welcomed increases in processor speeds, they may now be concerned about the increased power and cooling requirements that future acceleration could bring. An IDC study, for example, has found that both servers and storage can be major contributors to energy demand and to rising power and cooling costs, a trend that will likely continue increasing in the future (see Figure 1).[2]

Enterprises can take a variety of approaches to help address energy use. This article focuses on how data centers built on Dell™ hardware can use the advanced energy-efficient features of Dell/EMC CX3 series storage to help optimize storage performance, reduce energy use, and ultimately reduce ongoing operational costs.

MULTITIERED STORAGE DESIGN

The Dell/EMC CX3 UltraScale™ architecture allows administrators to deploy a wide variety of disk drive types within a single array. Dell/EMC CX3 disk array enclosures support simultaneous use of Serial ATA (SATA), 2 Gbps Fibre Channel, and 4 Gbps Fibre Channel drives to help increase flexibility within each storage array, particularly when combined with virtual logical unit (LUN) technology (see the "Virtual LUN technology" section in this article).

Because most disk drive technologies use a similar electromechanical design, drives with the same rotational speed often have similar power usage. Dell/EMC CX3 1 TB, 7,200 rpm SATA II disk drives can help dramatically

[1] Table 53, "Estimated U.S. Electric Utility Average Revenue per Kilowatthour to Ultimate Consumers by Sector, Census Division, and State, August 2000 and 1999," by the Energy Information Administration, U.S. Department of Energy, in *Electric Power Monthly*, November 2000, tonto.eia.doe.gov/ftproot/electricity/epm/02260011.pdf; and Table 5.6.A, "Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State, August 2007 and 2006," by the Energy Information Administration, U.S. Department of Energy, in *Electric Power Monthly*, November 2007, tonto.eia.doe.gov/ftproot/electricity/epm/02260711.pdf.

[2] "Storage Power Management Is Heating Up: Challenges and Solutions," by IDC, Doc #202573, July 2006.

increase storage system capacity. These high-capacity drives offer a 50 percent increase in capacity per spindle over previous-generation 500 GB drives and can provide comparable performance in large-block, sequential I/O environments—without requiring increased power consumption.

Dell/EMC CX3 10,000 rpm 4 Gbps Fibre Channel drives, available in 146 GB, 300 GB, and 400 GB sizes, are designed to increase storage system capacity through enhancements to linear and area densities and tracks per inch—the key factors affecting overall drive capacity per spindle. The 400 GB drives offer more than twice as much capacity per spindle as previous-generation 146 GB drives while providing equivalent or improved power usage and increased performance per spindle. Meanwhile, Dell/EMC CX3 15,000 rpm 4 Gbps Fibre Channel drives, available in 73 GB, 146 GB, and 300 GB sizes, are designed to increase performance through enhancements to disk operations such as rotational latency and seek rates—the key factors affecting access times.

The flexibility offered by Dell/EMC CX3 drives enables administrators to deploy drives of different capacities and performance levels to meet the needs of different applications within their environment. And because these drives have different power profiles, administrators can move data dynamically between drive types as needed to help control and manage energy use.

	Number of drives	Line current	Typical power consumption*	Annual energy costs**	Typical heat dissipation
7,200 rpm SATA II	15	1.4 A	0.29 kVA	US$726	900 Btu/hour
10,000 rpm 4 Gbps Fibre Channel	15	1.5 A	0.31 kVA	US$794	1,000 Btu/hour
15,000 rpm 4 Gbps Fibre Channel	15	1.6 A	0.32 kVA	US$821	1,100 Btu/hour

*Based on drive specifications. Actual power consumption will vary based on configuration, usage, and manufacturing variability.

**Based on a rate of US$0.1005/kWh, the average U.S. commercial retail price of electricity as given in Table 5.6.A., "Average Retail Price of Electricity to Ultimate Customers by End-Use Sector, by State, August 2007 and 2006," by the Energy Information Administration, U.S. Department of Energy, in *Electric Power Monthly*, November 2007, tonto.eia.doe.gov/ftproot/electricity/epm/02260711.pdf.

Figure 2. *Typical annual energy costs and heat dissipation for different Dell/EMC CX3 drive types*

Figure 2, which compares typical annual energy costs and heat dissipation for various Dell/EMC CX3 disk drives, illustrates typical differences between these drives. (Note that these figures compare only the drive type, not their relative power consumption per gigabyte.)

Based on the differences shown in Figure 2, administrators could set up a tiered structure for different types of data in their environment. Because each storage tier has different information and power requirements throughout its life cycle, administrators should ensure that they have an appropriate configuration when designing a specific storage environment.

FLEXIBLE RAID CONFIGURATION

Choosing the appropriate RAID configuration for particular applications can be key to maximizing performance and data protection while optimizing energy use. For example, although RAID-0 typically uses less power than other RAID types of the same capacity, it does not include safeguards against drive failures. And while RAID-1 provides high levels of availability, it is not well suited for writing large I/O sizes; in addition, because it is designed to duplicate all data, it can also double storage and energy costs.

Dell/EMC CX3 storage is designed to support multiple RAID types within each array and storage group, helping increase flexibility when designing a storage environment that can meet specific performance, data protection, and energy requirements. Figure 3 compares typical annual energy costs and heat dissipation for different RAID configurations in arrays using 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives. Figure 4 compares these same elements for a Dell/EMC CX3-80 configured with ninety 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives with one configured with fifteen of these Fibre Channel drives and seventy-five 750 GB, 7,200 rpm SATA II drives.

As these figures illustrate, using multiple tiers within a given Dell/EMC CX3 array can have a significant impact on storage power and cooling characteristics. Administrators can, for example, move aging data from high-performance Fibre Channel drives to high-capacity SATA II drives in a RAID-5 configuration to help

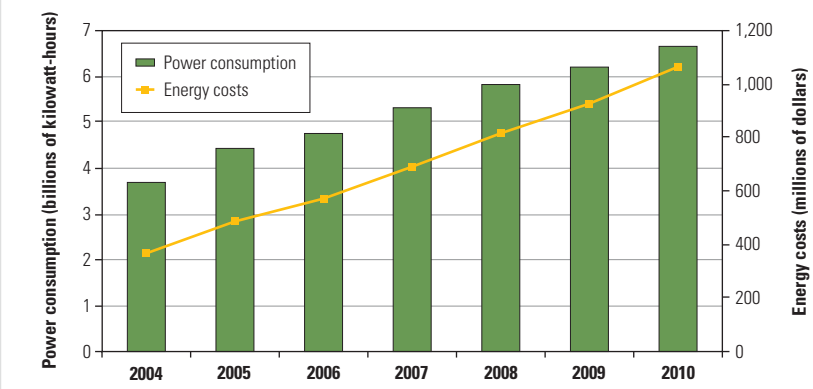

Source: "Storage Power Management Is Heating Up: Challenges and Solutions," by IDC, Doc #202573, July 2006.

Figure 1. *Worldwide power consumption of new disk drive shipments into enterprise solutions, 2004–2010*

reduce power and cooling requirements for that system. As Figure 4 shows, this approach can also help increase overall storage capacity in a given array.

STORAGE CONSOLIDATION AND EFFICIENT UTILIZATION

Consolidating storage by migrating from direct attach storage to a storage area network (SAN) can help significantly increase hardware utilization and energy efficiency while helping simplify management, reduce maintenance costs, and increase availability. Dell/EMC CX3 storage is designed to support the high utilization rates of SAN environments without sacrificing performance. However, administrators should be sure to take both the utilization and performance levels required by their environment into account when planning this type of consolidation to help ensure the design meets their needs.

METALUN TECHNOLOGY

MetaLUN technology is designed to allow administrators to expand the capacity of RAID groups and individual LUNs as needed to help accommodate growing storage requirements, without reducing application performance and availability. This technology allows administrators to deploy storage without needing to build in extra LUN capacity to accommodate future growth. By helping avoid the need to supply power to disks that may not be used right away, metaLUN technology can help increase utilization and reduce operating costs. MetaLUN technology is built into the EMC® FLARE® operating environment used by Dell/EMC CX3 storage, and does not require additional licenses or installation.

	Number of drives	Line current	Typical power consumption*	Annual energy costs**	Typical heat dissipation
RAID-0	4	0.7 A	0.14 kVA	US$348	400 Btu/hour
RAID-3	5	0.7 A	0.15 kVA	US$390	500 Btu/hour
RAID-5	5	0.7 A	0.15 kVA	US$390	500 Btu/hour
RAID-6	6	0.8 A	0.17 kVA	US$434	600 Btu/hour
RAID-10	8	1.0 A	0.20 kVA	US$520	700 Btu/hour

*Based on drive specifications. Actual power consumption will vary based on configuration, usage, and manufacturing variability.
**Based on a rate of US$0.1005/kWh, as described in Figure 2.

Figure 3. *Typical annual energy costs and heat dissipation for different RAID configurations in a Dell/EMC CX3 storage array using 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives*

Figure 5 illustrates the impact metaLUN technology can have on energy costs and heat dissipation. When not using this feature, administrators would typically need to provision all of the capacity during deployment whether they need it or not. As shown in Figure 5, the total power and cooling requirements of a 32 TB RAID-5 configuration are considerably higher than those of a 16 TB RAID-5 configuration. By using metaLUN technology, however, administrators could initially provision the 16 TB configuration, then expand it as their storage requirements increase, helping avoid the operating costs of the 32 TB configuration until they require that level of capacity.

VIRTUAL LUN TECHNOLOGY

Virtual LUN technology enables administrators to seamlessly migrate data between LUNs in an array without requiring application downtime, providing increased flexibility and control over the storage environment. It also enables them to transparently change characteristics such as RAID configuration, number of disk spindles, alignment offset, stripe element size, and LUN size, as well as to change drive types from SATA II to Fibre Channel and vice versa. Like metaLUN technology, virtual LUN technology is built into the EMC FLARE operating environment used by Dell/EMC storage, and does not require additional licenses or installation.

As shown in Figure 6, using this capacity to migrate aging or archival data from a 16-drive RAID-5 Fibre Channel array to a 7-drive RAID-5 SATA II array can significantly reduce annual energy costs while providing a comparable level of storage capacity.

EMC NAVISPHERE QUALITY OF SERVICE MANAGER

The EMC Navisphere® Quality of Service Manager (NQM) software available as an option with Dell/EMC CX3 storage systems enables administrators to monitor

	Number of drives	Total capacity	Line current	Typical power consumption*	Annual energy costs**	Typical heat dissipation
Non-tiered configuration	Ninety 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives	19.32 TB	11.2 A	2.35 kVA	US$5,995	7,700 Btu/hour
Tiered configuration	Fifteen 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives and seventy-five 1 TB, 7,200 rpm SATA II drives	58.24 TB	10.3 A	2.17 kVA	US$5,519	7,100 Btu/hour

*Based on drive specifications. Actual power consumption will vary based on configuration, usage, and manufacturing variability.
**Based on a rate of US$0.1005/kWh, as described in Figure 2.

Figure 4. *Typical annual energy costs and heat dissipation for a Dell/EMC CX3-80 array in non-tiered and tiered configurations*

	Number of drives	Line current	Typical power consumption*	Annual energy costs**	Typical heat dissipation
16 TB LUN	64	6.9 A	1.43 kVA	US$3,634	4,600 Btu/hour
32 TB LUN	128	13.4 A	2.80 kVA	US$7,091	9,100 Btu/hour

*Based on drive specifications. Actual power consumption will vary based on configuration, usage, and manufacturing variability.
**Based on a rate of US$0.1005/kWh, as described in Figure 2.

Figure 5. *Typical annual energy costs and heat dissipation for two RAID-5 LUNs in a Dell/EMC CX3 storage array using 300 GB, 15,000 rpm 4 Gbps Fibre Channel drives*

	Number of drives	Total capacity	Line current	Typical power consumption*	Annual energy costs**	Typical heat dissipation
1 TB, 7,200 rpm SATA II drives	6	4.59 TB	0.7 A	0.16 kVA	US$396	500 Btu/hour
300 GB, 15,000 rpm 4 Gbps Fibre Channel drives	16	4.03 TB	2.0 A	0.41 kVA	US$1,040	1,300 Btu/hour

*Based on drive specifications. Actual power consumption will vary based on configuration, usage, and manufacturing variability.
**Based on a rate of US$0.1005/kWh, as described in Figure 2.

Figure 6. *Typical annual energy costs and heat dissipation for SATA II and Fibre Channel Dell/EMC CX3 arrays*

and manage application performance. Administrators can use the NQM monitoring features to obtain a logical view of system performance, both for the entire storage system and for specific applications, to help them evaluate current service levels and determine what service levels are possible in specific environments. The NQM management features, meanwhile, enable them to set performance targets for high-priority applications and performance limits for low-priority applications to help them make effective use of storage resources. Combining these features with the built-in NQM scheduling capabilities enables administrators to easily control energy use during peak hours to help reduce overall operational costs.

Performance targets and limits are based on one of three key performance characteristics: response time, bandwidth, and throughput. Administrators can give a high-priority application utilizing a Dell/EMC CX3 storage system a target-specific response time, and limit the bandwidth available to a noncritical application to help ensure that the storage system has resources available for critical applications. Administrators can then use the NQM scheduling feature to dynamically adjust these targets based on enterprise requirements, automatically giving priority to specific applications during specific times. For example, they could give the e-mail application a certain performance target during the day, during its peak hours of operation, then create an overnight performance target for the backup application to help ensure that the e-mail data is backed up in a timely way (see Figure 7).

Because energy costs in industrial and commercial environments rise during peak hours, and because active disks consume more energy than inactive disks, managing application access to avoid highly disk-active tasks such as backup-to-disk operations during peak hours can help reduce overall energy costs. In addition, because NQM directly links application policies with specific LUNs, enforcing explicit application policy criteria can reserve active disk use for those LUNs for off-peak times to help reduce energy costs.

ENERGY-EFFICIENT DELL/EMC CX3 STORAGE

Dell/EMC CX3 series storage systems provide an extensive feature set designed to optimize energy efficiency. Combining options such as multitiered storage design, flexible RAID configuration, and SAN consolidation with intelligent management features such as metaLUN technology, virtual LUN technology, and EMC Navisphere Quality of Service Manager can help enterprises optimize storage performance while reducing energy use and maintaining control of overall energy costs.

Rodan Zadeh is consultant marketing manager for midrange storage platforms at EMC.

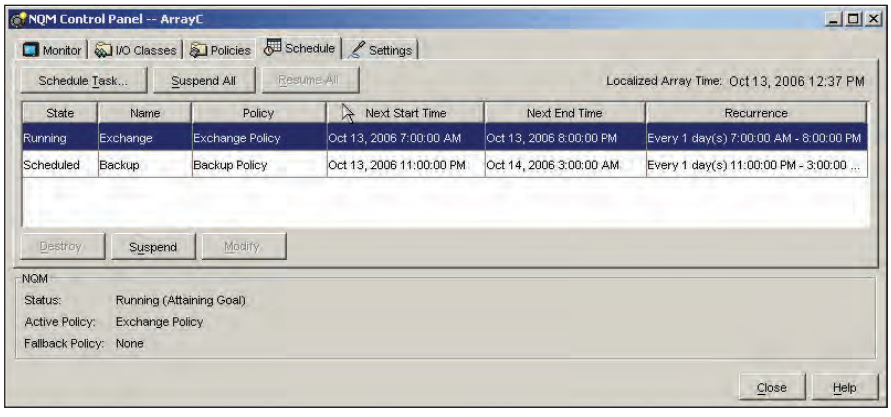

Figure 7. *Scheduled performance policies in EMC Navisphere Quality of Service Manager*

MANAGING ENERGY USE WITH DELL CLIENT MANAGER FROM ALTIRIS

By Todd Mitchell

Configuring and maintaining power settings for client systems can be a critical part of enterprise efforts to manage energy usage. For organizations built on Dell™ desktops, notebooks, and workstations, Dell Client Manager™ software from Altiris provides a simple, effective way to manage power schemes and other settings to help reduce power consumption and energy costs.

Related Categories:

Altiris

Power and cooling

Symantec

Systems management

Visit DELL.COM/PowerSolutions for the complete category index.

For many enterprises, clear strategies for environmental stewardship and efficient energy use have become increasingly important, with goals that may include reducing their overall power consumption, working to extend the useful life of their products, and providing equipment recovery and recycling services. For IT managers, reducing power use for client systems can be a key part of these strategies. Altiris (now part of Symantec) and Dell have partnered to make Dell Client Manager, a flexible, easy-to-use tool to help administrators manage power schemes and other settings for the Dell OptiPlex™ desktops, Dell Latitude™ notebooks, and Dell Precision™ workstations in their environments, particularly when combined with systems using Dell Energy Smart configurations.

Key energy management features of Dell Client Manager include power scheme tasks enabling administrators to collect current power scheme inventories and activate customized power schemes remotely on Dell systems, as well as automated notification policies that can monitor these inventories for noncompliant systems. Administrators can also use Dell Client Manager to configure BIOS settings related to power management and take advantage of Intel® vPro™ technology for secure, scheduled power cycling of remote clients. Using these features to actively manage client energy use can help significantly reduce power consumption and related costs for Dell systems.

Dell Client Manager is available at no additional cost as part of the Dell OpenManage™ suite, and can be downloaded from the Client Management section at DELL.COM/OpenManage.

DELL ENERGY SMART SYSTEMS

The Dell Energy Smart offering, available when purchasing systems from the Dell Web site, combines energy-efficient hardware with processor and OS power management to help reduce power consumption and energy costs. Dell Energy Smart settings help optimize energy savings in part by setting OS power management capabilities to place systems into a low-power mode after 15 minutes of inactivity. Using Energy Smart configurations for Dell desktops, notebooks, and workstations, for example, can help reduce power consumption by up to 78 percent compared with non–Energy Smart Dell system configurations.[1]

[1] Based on product specifications and 7 hours of typical office use, 1 hour of maximum-performance use, 1 hour of idleness, and 15 hours of sleep for 264 days each year, with 24 hours of sleep for the remaining 101 days each year.

> **"The Dell Energy Smart offering combines energy-efficient hardware with processor and OS power management to help reduce power consumption and energy costs."**

Dell also provides an online energy calculator[2] to help organizations identify potential savings from using Energy Smart features with Dell desktops, notebooks, and workstations. For example, this tool shows that a Dell OptiPlex GX620 desktop with an Intel Pentium® D processor and a 17-inch CRT monitor would use approximately 966 kWh per year, while a comparable OptiPlex GX620 with Energy Smart power management and a 17-inch flat panel monitor would use only 368 kWh per year, saving US$59.80 per desktop annually—which, for an enterprise with 500 desktops, translates to savings of US$29,900 each year. Replacing a standard OptiPlex GX620 with an Energy Star–compliant OptiPlex 745 desktop with an Intel Core™2 Duo processor and an external 17-inch flat panel monitor used for four hours each day would bring energy use down to 214 kWh per year, saving US$75.20 per desktop annually—or, for enterprises with 500 desktops, a total of US$37,600 each year.[3]

These projected savings assume, of course, that Energy Smart power configurations are maintained throughout a system's life cycle—which is where Dell Client Manager can help. By enabling administrators to set power configurations across the Dell client systems in their environment, this tool can help keep power management consistent and in line with IT policies. If settings are changed, Dell Client Manager can notify administrators through e-mail or automatic help-desk tickets, or automatically enforce a variety of other responses.

DELL CLIENT MANAGER POWER SCHEMES

Dell Client Manager supports multiple models of Dell OptiPlex desktops, Dell Latitude notebooks, and Dell Precision workstations, and requires the Microsoft® Windows® 2000 OS or later. It uses the Altiris Task Server™ component to define and distribute power schemes across Dell systems. To provide a familiar interface and enable administrators to define power settings quickly and easily, the power scheme task settings are designed to look and function similarly to the Microsoft Windows Power Options Properties window (see Figure 1). When administrators activate a power scheme on a remote system, it also appears on that system's list of available power schemes as "Altiris Power Scheme" (as shown in Figure 1).

Administrators can create a power scheme in Dell Client Manager by selecting View > Tasks in the menu bar, selecting Tasks > Task Management > Client Tasks in the left pane, then right-clicking on the Client Tasks folder and selecting New > Task/Job. In the dialog box, they can then select "Power Scheme Task" and configure the appropriate settings. To distribute a power scheme to remote Dell client systems, administrators can select View > Tasks in the menu bar, select Tasks > Task Management > Power Scheme Task in the left pane, and click "Create Schedule" in the right pane. They can then provide a name, description, and schedule time, and click "Select computers" to select specific systems or

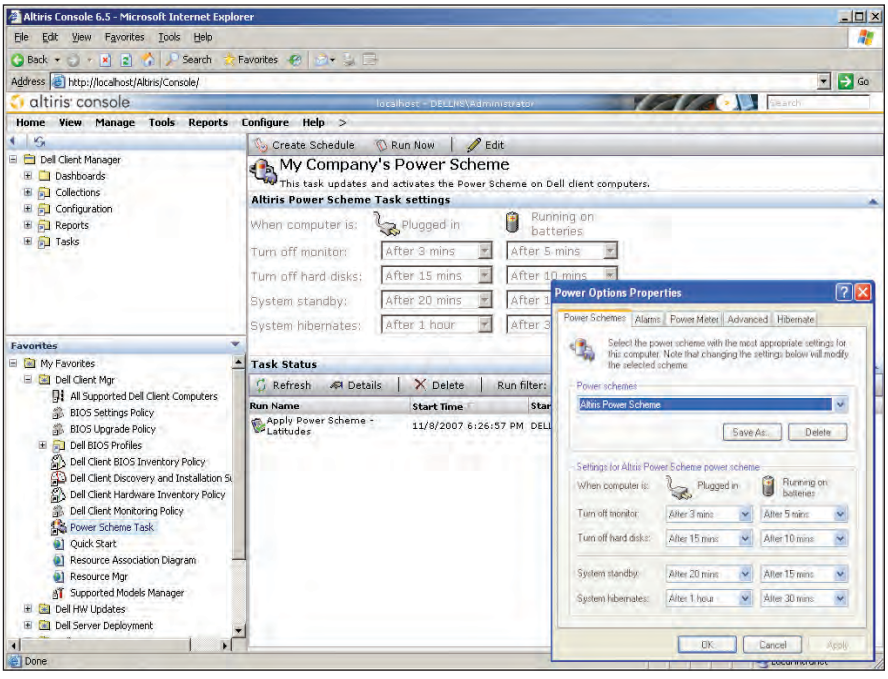

Figure 1. *Dell Client Manager power scheme shown with the Microsoft Windows Power Options Properties window*

[2] Available at DELL.COM/HTML/US/Products/OptiPlex/Calculator/Index.html.

[3] Based on product specifications; 7 hours of typical office use, 1 hour of maximum-performance use, 1 hour of idleness, and 15 hours of sleep for 264 days each year, with 24 hours of sleep for the remaining 101 days each year; and an average energy price of US$0.10/kWh.

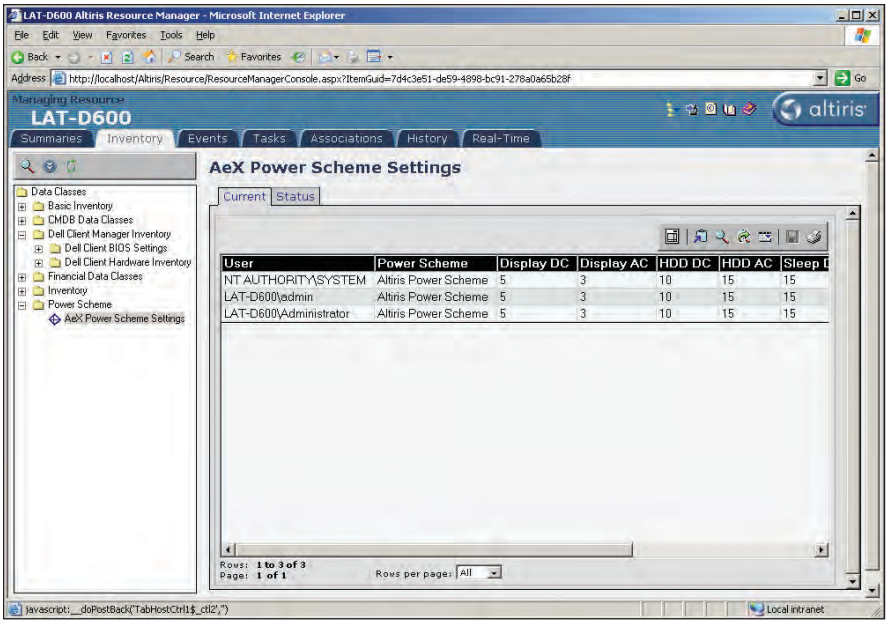

Figure 2. *Power scheme settings displayed in Altiris Resource Manager*

collections to apply the power scheme to. When they are finished, they can either click the OK button to save the task settings or click the Run Now button to distribute the power scheme immediately.

Dell Client Manager can also inventory active power scheme settings for each user on a client system. This inventory is reported back to the Altiris configuration management database and can be used for a variety of reporting and alerting functions. Administrators can also view power scheme settings for specific clients in the Inventory tab of that system's Altiris Resource Manager software (see Figure 2).

Power scheme inventories can be particularly useful when combined with notification policies that check for noncompliant settings (see Figure 3). When a noncompliant setting is detected, these notification policies can trigger a variety of automated actions, including creating a help-desk ticket, sending an e-mail to an administrator, launching a custom script or application, e-mailing a report, and so on. Administrators can create notification policies in Dell Client Manager by selecting View > Tasks in the menu bar, navigating to Assets and Inventory > Inventory > Windows > Notification Policies, then right-clicking on the Notification Policies folder and selecting New > Notification Policy.

BEYOND POWER SCHEMES

In addition to configuring power schemes, organizations can use Dell Client Manager to manage power settings in a number of additional ways. For example, administrators can use Dell Client Manager to set BIOS policies to manage settings related to energy use, including power management settings, wireless functions, Intel SpeedStep® technology, and so on (see Figure 4).

Dell Client Manager also supports the Intel vPro technology available on many Dell systems, which enables administrators to remotely perform secure, scheduled power cycling. Management console policies can help ensure that vPro-based systems on a network are powered down after a certain time and then, optionally, powered back up for normal business hours. For organizations with a large number of systems that are billed on a time-of-use rate, using schedules to stagger when they are powered up can

"Administrators can use Dell Client Manager to set BIOS policies to manage settings related to energy use, including power management settings."

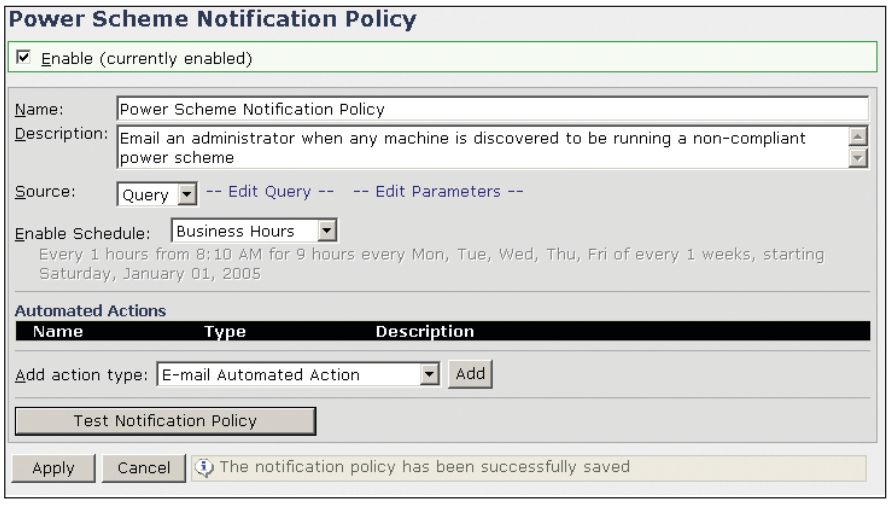

Figure 3. *Dell Client Manager power scheme notification policy*

> **"Dell Client Manager also supports the Intel vPro technology available on many Dell systems, which enables administrators to remotely perform secure, scheduled power cycling."**

reduce peak demand usage to help control energy costs. Administrators can take advantage of the vPro power-on capabilities for advanced management functions as well, including powering up systems to apply OS or antivirus updates and then powering them down (see Figure 5).

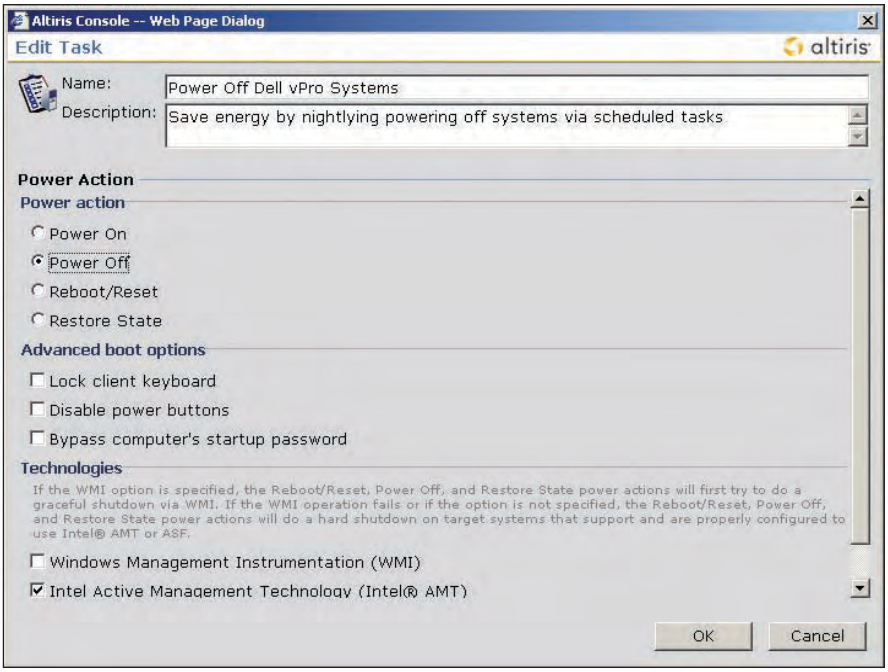

Figure 5. *Dell Client Manager task using Intel vPro for enhanced power management*

FLEXIBLE POWER MANAGEMENT FOR DELL SYSTEMS

Dell Client Manager is designed to help IT administrators actively manage power schemes and related settings for Dell OptiPlex desktops, Dell Latitude notebooks, and Dell Precision workstations. For enterprises working to build and implement environmentally friendly IT policies while simultaneously reducing power consumption and related energy costs, Dell Client Manager can be a key part of their overall strategy for energy efficiency.

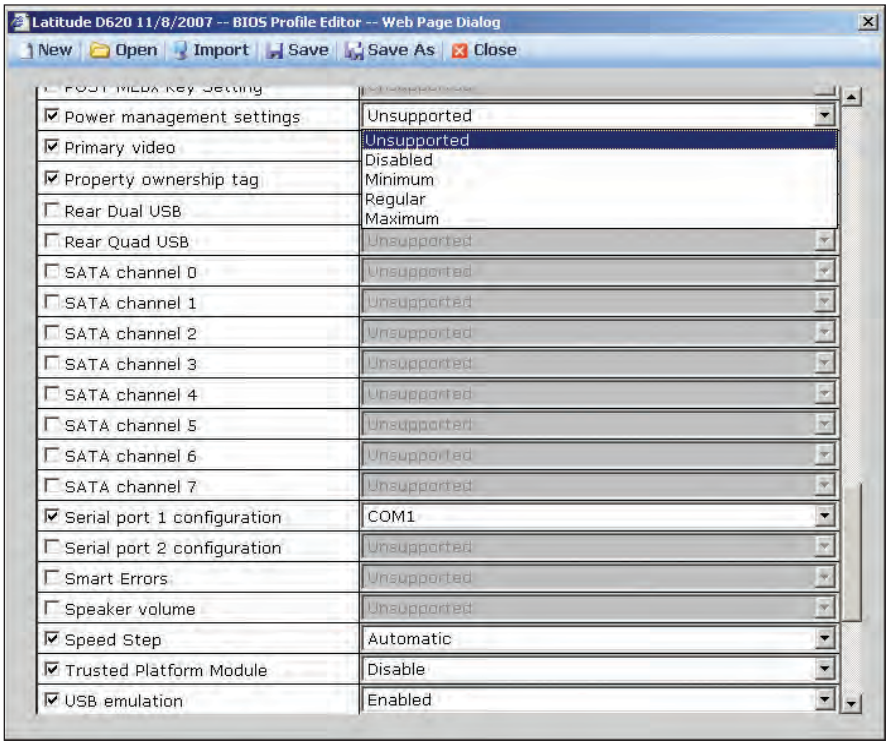

Figure 4. *BIOS policies in Dell Client Manager*

Todd Mitchell is a technical director at Altiris, now part of Symantec. Todd has a bachelor's degree from Brigham Young University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Client Manager:
www.altiris.com/dellclientmanager
DELL.COM/OpenManage

Dell Energy Smart:
DELL.COM/EnergySmart

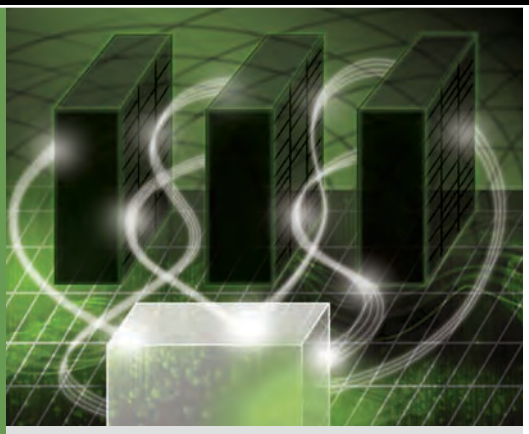

A SYSTEMS-LEVEL APPROACH TO EFFICIENT DATA CENTER DESIGN

By Fred Stack

The increasing need for high-performance, rack-dense servers has strained the limits of many data center cooling systems. Dell-Liebert Energy Smart Solutions can help organizations overcome these limits, enabling them to increase performance while maintaining existing levels of energy use or to provide sufficient cooling for a maximum-performance infrastructure.

Related Categories:

Data center cooling

Dell PowerEdge servers

Emerson Network Power

Liebert Precision Cooling

Power and cooling

Visit DELL.COM/PowerSolutions for the complete category index.

Rising demand for computing capacity and constraints on power and cooling are forcing many organizations to consider building new IT facilities much earlier than they had anticipated. For example, according to a survey by the Data Center Users' Group, 96 percent of data centers are expected to reach capacity by 2011, even though many of those facilities were built in the last 10 years.[1]

Physical space alone is typically not the problem. As server form factors continue to shrink while processing power grows, many facilities have ample physical space for additional servers. What they lack is the power and cooling to support additional devices. Rapidly increasing rack densities have forced data center managers to leave open space in racks or spread racks out across the room to help ensure proper cooling—an inefficient use of space and cooling systems that can limit data center growth.

Dell-Liebert Energy Smart Solutions are designed to help enterprises overcome these limitations. By taking a systems-level approach to data center design that combines efficient Dell™ PowerEdge™ Energy Smart servers with Liebert DS™ and Liebert XD™ precision cooling systems from Emerson Network Power, these solutions enable IT managers to provide sufficient cooling for a maximum-performance infrastructure or to increase overall performance without requiring additional energy use.

DEPLOYING EFFICIENT INFRASTRUCTURE TECHNOLOGIES

Designing an efficient infrastructure requires coordinating technologies at the server level and the room level. At the server level, efficient servers incorporate advanced processor, memory, power, and thermal design technology to help increase performance while reducing power consumption. Complementing these servers with an efficient infrastructure helps enable organizations to both increase data center rack processing power and reduce cooling-related energy costs.

At the room level, Emerson Network Power provides variable-capacity cooling systems, such as the Liebert DS precision cooling system, with Copeland Digital Scroll™ compressors and variable-speed motors that enable highly efficient operation at partial loads. The Liebert DS also employs an intelligent control system that can facilitate coordination between multiple cooling units in a room.

[1] "Emerson Network Power Presents Industry Survey Results That Project 96 Percent of Today's Data Centers Will Run Out of Capacity by 2011," by Emerson Network Power, November 16, 2006, www.liebert.com/information_pages/NewsRelease.aspx?id=2386.

> **"By taking a systems-level approach to data center design, these solutions enable IT managers to provide sufficient cooling for a maximum-performance infrastructure or to increase overall performance without requiring additional energy use."**

Computer room air-conditioning units can typically provide efficient environmental control for rack densities of up to 5 kW per rack. When using densities higher than that level, organizations should add supplemental cooling systems close to the heat source to assist with heat removal. (The room cooling systems are still necessary to control humidity and air quality and assist in balancing temperatures across the room.) By reducing the distance cold air must travel across the room, this approach helps reduce the energy required to move it.

For example, the flexible and efficient Liebert XD precision supplemental cooling system from Emerson Network Power uses a pumped refrigerant infrastructure to support cooling modules placed directly above or alongside high-density racks, supplementing the air coming up through the floor and helping eliminate cooling-related density or capacity limitations. Compared with traditional water-based room cooling systems, Liebert XD systems are designed to reduce energy requirements when operating at both full and half load.

TESTING DELL-LIEBERT ENERGY SMART SOLUTIONS

Advances in server and cooling system technology can have a major impact on data center performance and efficiency when server and infrastructure manufacturers work together to take a true systems-level approach to data center design. This philosophy has led to the creation of Dell-Liebert Energy Smart Solutions.

To demonstrate how these solutions can help address the challenges that IT management face today, in summer 2007 Dell and Emerson Network Power experts simulated a model data center using a baseline environment of 50 racks of Dell PowerEdge 1850 servers, each with two dual-core Intel® Xeon® 7030 processors at 2.8 GHz, 4 GB of double data rate 2 (DDR2) memory, and two 36 GB, 15,000 rpm hard drives. Using the SPECjbb2005 benchmark to measure performance, the test team determined that each server produced a benchmark rating of 46,360 business operations per second (bops) while using 405 W of power, for a total power consumption of 243 kW. The test environment assumed a raised-floor cooling capacity of approximately 5 kW per rack, which limited each rack to 12 servers, for a total of 600 servers in the data center. Cooling was provided by three Liebert Deluxe FH740C computer room air-conditioning units, which used a total of 86.6 kW of power. Overall, this baseline environment had a total benchmark performance rating of 27,816,000 bops and a total power consumption for the servers and cooling units of 329.6 kW, with a total facility power consumption of 471 kW.

Next, the test team upgraded the environment to increase computing capacity without increasing overall power consumption by replacing the 600 PowerEdge 1850 servers with 1,148 PowerEdge Energy Smart 1950 servers, each with two dual-core Intel Xeon 5148LV processors at 2.33 GHz, 4 GB of DDR2 memory, and two 73 GB, 10,000 rpm hard drives (see Figure 1). They also reconfigured the data center from 50 racks of 12 servers each to 28 racks of 41 servers each. Each PowerEdge 1950 produced a SPECjbb2005 rating of 84,312 bops while using 234 W of power, for a total power consumption of 268.6 kW. In addition to upgrading the servers, the test team replaced the three fixed-capacity Liebert Deluxe FH740C units with two variable-capacity Liebert Deluxe FH740C units and one Liebert XD supplemental cooling system using Liebert XDV units on

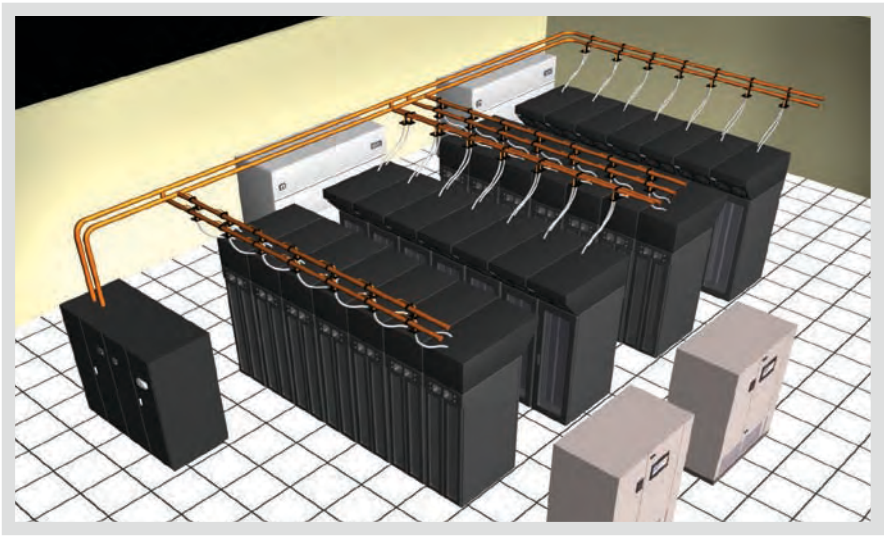

Figure 1. *Dell PowerEdge Energy Smart 1950 servers and energy-efficient Liebert cooling systems provided increased performance without increased power consumption*

top of the racks to support the increased rack density, a setup that used a total of 56.2 kW of power for cooling. Overall, this environment had a total benchmark performance rating of 96,790,176 bops and a total power consumption for the servers and cooling units of 324.6 kW, with a total facility power consumption of 464 kW—a 350 percent increase in performance over the baseline environment while using a comparable amount of power.

The test team next upgraded the environment further to help maximize processing performance by filling all of the available space in the 50 racks, replacing the 600 PowerEdge 1850 servers with 2,100 PowerEdge Energy Smart 1950 servers, for a total power consumption of 491.4 kW (see Figure 2). They also upgraded the cooling systems to handle the increased heat load, replacing the three fixed-capacity Liebert Deluxe FH740C units with two variable-capacity Liebert Deluxe FH740C precision cooling systems and two Liebert XD supplemental cooling systems using Liebert XDV units on top of the racks, a setup that used a total of 75.2 kW of power for cooling. Overall, this environment had a total benchmark performance of 177,055,200 bops and a total power consumption for the servers and cooling units of 566.6 kW, with a total facility power consumption of 852 kW—an 81 percent increase in power consumption, but a 637 percent increase in performance over the baseline environment.

"IT managers do not need to wait for complicated control schemes—which may take years to fully develop—to achieve significant increases in efficiency and performance. By taking a systems-level approach to data center design and combining multiple technologies available today, Dell-Liebert Energy Smart Solutions can help increase both efficiency and performance."

DESIGNING FOR HIGH EFFICIENCY AND PERFORMANCE

IT managers do not need to wait for complicated control schemes—which may take years to fully develop—to achieve significant increases in efficiency and performance. By taking a systems-level approach to data center design and combining multiple technologies available today, Dell-Liebert Energy Smart Solutions can help increase both efficiency and performance and help remove limits to continued data center growth.

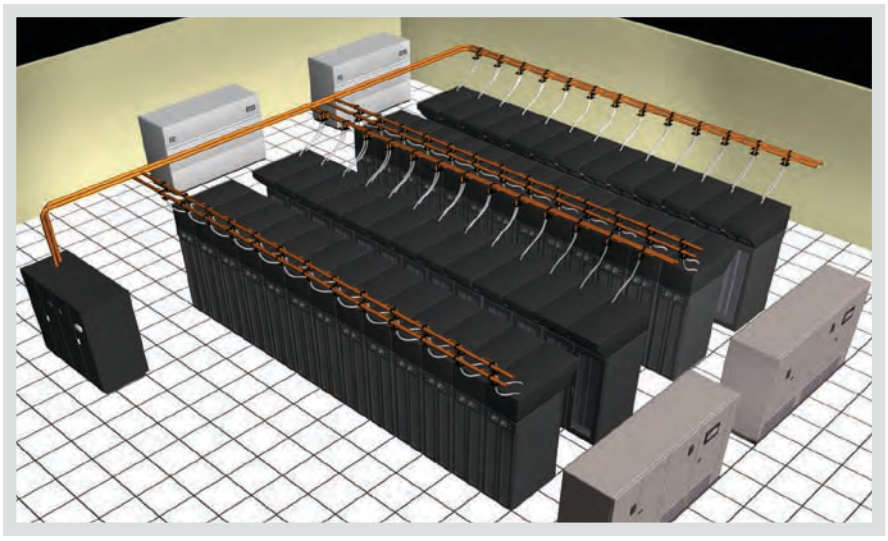

Figure 2. *Liebert cooling systems help remove constraints on rack utilization for maximum data center performance*

Fred Stack is vice president of marketing for Liebert Precision Cooling, part of Emerson Network Power. He is responsible for new product development road maps that reflect evolving market demands and incorporate new technology, and that serve as the foundation for development of new products and application engineering tools. He has previously worked in technical marketing and sales for Rockwell International and several divisions of General Electric. Fred has a B.S. in mechanical engineering from Tulane University and an M.B.A. from Southern Methodist University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Data center solutions from Emerson Network Power:
dell5.liebert.com

For you, it's a problem you didn't see coming.

For your business, it's a customer you won't see coming back.

You can't anticipate every problem. But Emerson Network Power and its Liebert power and cooling technologies can help you create an IT infrastructure that is ready for anything—unplanned outages, unpredictable growth or unexpected technologies.

One example is the **Liebert NX**, a software-scalable UPS that can double in capacity without adding or modifying hardware. Download our white paper, *Powering Change in the Data Center*, and discover what Liebert technologies can do for your operating flexibility, at flexibility.liebert.com.

Liebert flexibility
Just another reason why Emerson Network Power is the global leader in enabling Business-Critical Continuity.™

EMERSON™
Network Power

EMERSON. CONSIDER IT SOLVED.™

INSIDE THE EQUALLOGIC PS SERIES iSCSI STORAGE ARRAYS

By John Joseph
Eric Schott
Kevin Wittmer

Built on a patented peer storage architecture, the EqualLogic® PS Series of Internet SCSI (iSCSI) storage arrays offers high performance, reliability, intelligent automation, and seamless virtualization of a single pool of storage to enable simplified enterprise storage deployment, ease of management, and comprehensive data protection.

Related Categories:

EqualLogic

Internet SCSI (iSCSI)

Storage

Visit DELL.COM/PowerSolutions for the complete category index.

Storage management is at the center of an increasingly complex IT problem. Applications continually demand more from their storage systems—as performance, capacity, availability, backup, and disaster recovery needs intensify, so do expectations of the servers and storage that support them. As a result, IT professionals need a storage solution that integrates a full complement of high-end management capabilities, includes data protection, and operates intelligently. They need a reliable storage system that expands transparently and provides consistent data availability regardless of the server configuration, OS, or application.

The EqualLogic PS Series of Internet SCSI (iSCSI) storage arrays offers a fundamental change in the way enterprises think about purchasing and managing storage. Built on a patented peer storage architecture, these solutions offer enterprise-class performance and reliability, intelligent automation, and seamless virtualization of a single pool of storage for simplified storage management. The PS Series combines an intelligent, automated management framework and a comprehensive set of enterprise data services with a fault-tolerant hardware architecture that supports most major operating systems. It delivers a modular and cost-effective solution that can be deployed in appropriate increments for small and medium businesses, while also being cost-effective for large enterprises requiring high-end capacity and performance.

PEER STORAGE ARCHITECTURE

The EqualLogic PS Series is based on the unique *peer storage architecture*. In this context, *peer* describes the collaboration and equal partnership of a single, simple architecture; arrays are designed to function as peers, working together to share resources, evenly distribute loads, and collaborate to help optimize application performance and provide comprehensive data protection.

The result is an intelligent storage array that can deliver rapid installation, simple management, and seamless expansion. Using patent-pending page-based data mover technology, members in a storage area network (SAN) work together to automatically manage data, load balance across resources, and expand to meet growing storage needs. Because they share this architecture, enterprises can use PS Series arrays as modular building blocks for simple SAN expansion (see Figure 1). This architecture provides the basis for numerous features and capabilities, including peer deployment, control, provisioning, protection, and integration (see Figure 2).

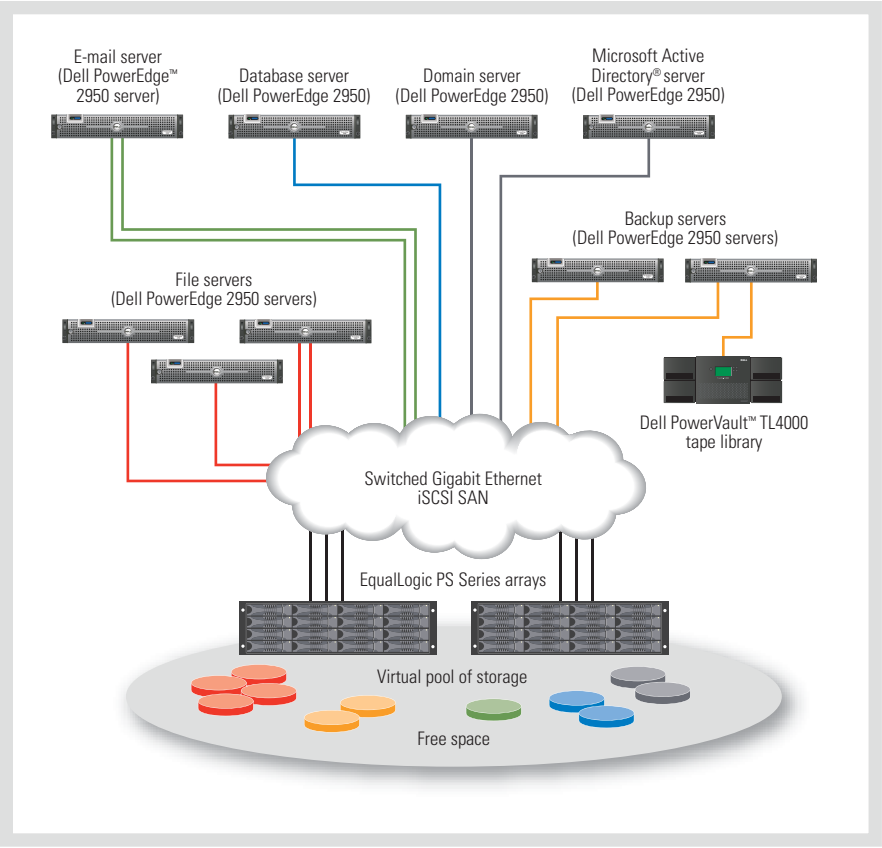

Figure 1. *Grouping multiple EqualLogic PS Series arrays helps increase capacity and performance linearly—without increasing management complexity*

Modular design with enterprise reliability

EqualLogic PS Series arrays have a modular design that allows enterprises to purchase only the storage they need, when they need it—helping prevent both underutilization and over-provisioning. Its peer storage architecture uses industry-standard protocols, disk drives, and network interfaces to provide cost-effective connections and high-performance access to data across heterogeneous environments, including Microsoft® Windows®, Novell® NetWare®, Novell SUSE® Linux® Enterprise Server, Red Hat® Enterprise Linux, Solaris, UNIX®, HP-UX, Mac OS X, IBM® AIX, and VMware® ESX Server environments.

Designed to meet and exceed the rugged requirements of the data center, EqualLogic engineered fault tolerance into the PS Series hardware design. Its components are fully redundant and hot swappable with optional dual controllers, standard dual fan trays, and dual power supplies. The hot-swappable controller module features high-performance dual-core 64-bit processors with a high-speed I/O bus and twin 64-bit double data rate (DDR) channels. Each control module is equipped with 1 GB of high-speed battery-backed DRAM. Each disk drive is interconnected with its own independent, hot-swappable serial channel and connected mechanically with an inertial dampening chassis that helps eliminate drive vibrations. Self-tuning controller caches are battery backed and mirrored across controllers to help protect these components without compromising performance.

EqualLogic PS Series arrays support Serial Attached SCSI (SAS) and Serial ATA (SATA) disk drives with a variety of performance and capacity options. Enterprise-class RAID protection governs hot-swappable disk drives, including RAID-5, RAID-10, and RAID-50 support; hot sparing; automatic rebuilds; accelerated rebuild times; advanced stripe integrity algorithms; online RAID set expansion; geometry transformation; and patent-pending predictive media-error detection and correction. Not only can administrators service the arrays without taking them offline, but the system isolates faults to help prevent cascading failures or loss of protection during service. The result is high levels of data protection and performance even during service procedures. The redundancy of the design helps eliminate single points of failure and provide enterprise-level availability and reliability.

PEER DEPLOYMENT: RAPID INSTALLATION

Peer deployment is an automatic sensing and SAN configuration technology in EqualLogic PS Series arrays designed to eliminate complex and cumbersome manual tasks. It can enable administrators to set up and deploy PS Series arrays within minutes of opening the shipping box.

Once the array has identified the network topology, the peer deployment intelligence conducts a system health check to help ensure that components are fully functional, and automatically builds RAID sets. With a few simple steps and without

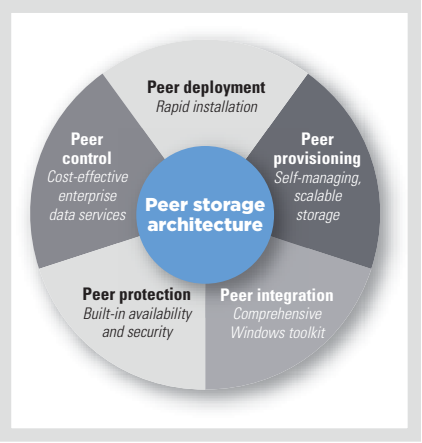

Figure 2. *EqualLogic PS Series enterprise-class features are based on the patented peer storage architecture*

special expertise, administrators can deploy an enterprise-class SAN seamlessly, without downtime. This capability helps eliminate the complex manual configuration of other SANs, allowing administrators to focus on servers and applications.

PEER CONTROL: COST-EFFECTIVE ENTERPRISE DATA SERVICES

Peer control offers virtualized storage management with a single view. EqualLogic PS Series arrays are designed to be self-managing; the system continuously monitors storage resources and can automatically load balance data across controllers, network connections, and disk drives to help deliver optimal performance. Peer control also automates key functions for configuration, management, storage pooling, and data distribution, helping minimize the complexity of storage administration.

The PS Series includes built-in storage features previously available only to top-tier data centers, enabling best-practice storage for organizations of all sizes. In addition to automatic load balancing, the arrays include enterprise software features such as automatic snapshot management, automatic replication, volume cloning, volume management, storage virtualization, thin provisioning, SAN boot capability, role-based administrative management, historical performance trending and reporting, tiering and pooling, and multipath I/O with no additional licensing fees or host-based software to implement.

PS Group Manager

Administrators can manage PS Series arrays in a SAN through a single interface—the PS Group Manager (see Figure 3). This Web browser–based GUI helps eliminate the need for a dedicated management workstation or server and allow administrators to remotely manage virtually any aspect of their EqualLogic iSCSI-based SAN. In addition to using the Web interface, administrators can also manage PS Series arrays using a scriptable command-line interface over Secure Shell (SSH) and Telnet. Built-in monitoring and notifications provide e-mail, syslog support, and comprehensive Simple Network Management Protocol (SNMP) monitoring and traps—all standard features.

Volume snapshots

PS Series arrays provide space-efficient snapshots that support up to 512 snapshots per volume, thousands of snapshots per array, and read-only as well as read/write snapshots. Snapshots can be used for quick recovery and offloading backup operations. PS Series arrays implement safe snapshot recovery in which data is not discarded unintentionally, helping provide the flexibility to implement true enterprise IT storage.

Automatic load balancing

The EqualLogic approach to optimizing performance is simple: maximize the horsepower of all available components in the SAN. Because of the unique peer storage architecture used by EqualLogic PS Series arrays, instead of relying solely on individual components to deliver high performance, peer control leverages all relevant components. It automatically load balances to seamlessly spread data across active storage resources, helping maximize performance by intelligently optimizing the available network connections, cache, controllers, and drives. Through this optimization, each PS Series array can support high transactional workloads and database applications. When multiple arrays are added together, the PS Series quickly scales, with performance growing linearly with each additional array.

Tiering and pooling

The EqualLogic tiering capability enables administrators to prioritize applications within a SAN by placing them on separate storage resources, each optimally configured for the required service level of the application. Administrators can configure separate storage pools within a single SAN to help build an efficient, flexible, easy-to-manage storage environment. Using this "SAN within a SAN," administrators can gain the advantages of consolidation, but can also easily separate workloads as needed: by application, by service level, by disk type, by cost, or even by department within the organization. Because online data movement is built into all PS Series arrays,

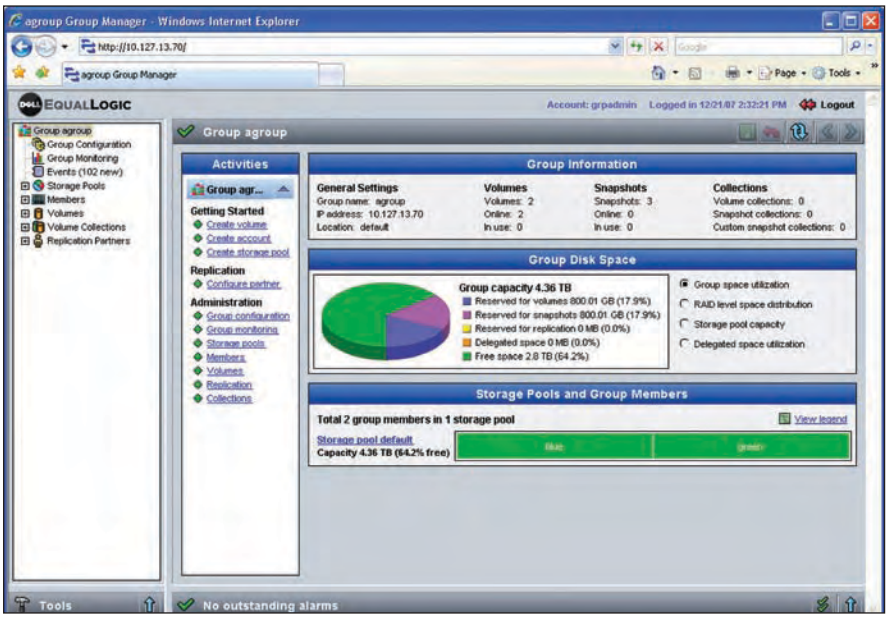

Figure 3. *The EqualLogic PS Group Manager interface helps simplify storage management across PS Series arrays*

EqualLogic PS Series iSCSI storage arrays offer enterprise-class performance and reliability, intelligent automation, and seamless virtualization

administrators can adjust application resources and move data between different pools of storage without downtime or disruption.

PEER PROVISIONING: SELF-MANAGING, SCALABLE STORAGE

In combination with the modular design of PS Series arrays, *peer provisioning* enables administrators to dynamically provision resources to meet application requirements—including not only disk space, but also connectivity, security, performance, and data protection. When application requirements change, the storage configuration can change seamlessly.

The intelligence of peer provisioning enables administrators to quickly and easily expand storage capacity without affecting data availability. They can start with a single storage array—then, when their storage, performance, and network requirements grow, they can add more members to help increase capacity, performance, and network bandwidth automatically and linearly. Expansion is linear, enabling administrators to scale not only disk drives but also controllers, ports, and cache, and performance can potentially increase as the environment grows.

Configuring one or more EqualLogic PS Series arrays as a PS Group also enables administrators to manage the arrays as a single system with a shared pool of storage. EqualLogic peer provisioning automates the key functions needed to configure, manage, and scale storage, helping eliminate much of the complexity of storage administration. Each group member is automatically configured and participates in balancing the load, distributing data, and tracking host access to data, without requiring user intervention—helping keep storage management simple, regardless of scale.

The power of peer provisioning enables a key feature of PS Series arrays: on-demand growth. Because capacity expansion in PS Series arrays is designed to be nondisruptive, administrators can add storage resources while applications remain online. As a result, enterprises do not need to tie up capital by over-provisioning—leaving underutilized storage assets—as IT managers often do to help avoid future downtime. With a PS Series SAN, enterprises can purchase the storage capacity they need today, and plan for future storage growth knowing that when they need to expand, they can do so easily and without interrupting performance. In addition, PS Series arrays are designed to fully integrate with Microsoft Virtual Disk Service (VDS), enabling application-level provisioning by supporting applications.

Thin provisioning

Thin provisioning is an important advanced feature of peer provisioning that enables the automatic addition of physical capacity on demand up to preset limits. Advanced thin provisioning helps make buy-as-you-grow storage management and virtualization seamless for servers and applications. When administrators create a volume, they can size it for the long-term needs of the application without initially allocating the full amount of physical storage. Instead, as the application needs additional storage, capacity is allocated to the volume from a free pool. The EqualLogic implementation of thin provisioning provides enhanced flexibility and safety controls—with proactive, user-defined threshold alarms and controls, administrators can depend on automatic space allocation without worrying about reaching allocation limits or unexpected depletion of physical storage.

Peer provisioning can offer significant economic benefits to enterprises of all sizes, because they can purchase physical storage when they actually need it. These benefits include increased asset utilization, reduced management costs, reduced floor space footprint, reduced power and cooling costs, and smart, efficient capital expenditures.

"The PS Series includes built-in storage features previously available only to top-tier data centers, enabling best-practice storage for organizations of all sizes."

PEER PROTECTION: BUILT-IN AVAILABILITY AND SECURITY

Peer protection starts with a robust design that helps avoid single points of failure and provide enterprise-level availability and reliability, including built-in PS Series features such as application-aware snapshots for quick recovery and remote replication for disaster protection. It encompasses comprehensive system monitoring and high-availability features such as E-mail Home, multipath I/O, and Auto-Replication that help provide comprehensive protection against system failures or outages. These features enable administrators to quickly create an end-to-end solution that can help provide comprehensive protection against multiple types of failure or outage.

Comprehensive system monitoring

EqualLogic system monitoring capabilities provide administrators with a comprehensive view of the health and status of their SAN. Within the PS Group Manager, administrators can view the status of individual SAN components such as drives, power supplies, and controllers, as well as the overall system. PS Series arrays also include the Auto-Stat Disk Monitoring System (ADMS), which proactively scans disk drives in the background to help detect media anomalies and correct them. Combined with automatic sparing and spare utilization, this feature helps enhance protection and ensure optimal disk performance.

E-mail Home

To help ensure the availability of systems and data, PS Series arrays come standard with group event notification methods (e-mail, syslog, and SNMP), but also allow administrators to enable E-mail Home functionality. E-mail Home can automatically contact EqualLogic customer support if a hardware component such as a disk, control module, fan tray, or power supply fails, or if the firmware on a PS Series array is updated. E-mail Home functionality enables EqualLogic to rapidly respond to issues and assist administrators.

Multipath I/O

EqualLogic multipath I/O provides high availability and performance load balancing across multiple network ports (host bus adapters and/or network interface cards) for Windows, Linux, UNIX, and VMware environments. By leveraging the ease of use and cost-effectiveness of Ethernet, multipath I/O helps remove single points of failure between the server and the storage.

Auto-Replication

EqualLogic Auto-Replication remotely replicates data from one PS Group to another over a standard IP network over long distances, helping provide high levels of data protection and disaster tolerance. Auto-Replication offers the advantages of geographic isolation—a critical component in any true disaster recovery plan—without the traditional complexity. Administrators can quickly and easily configure volumes for replication, letting the PS Series arrays manage the underlying hardware resource complexity. A catalog of recovery points is maintained at each disaster recovery site, providing multiple points of recovery from which to choose in the event of a disaster—a choice that helps protect against "cascading" failures such as software viruses.

Making a disaster recovery site operational is simple, even if the primary site already stores terabytes of data. EqualLogic PS Series arrays enable administrators to perform an initial manual sync by copying primary site data to transportable media, physically shipping the media to the remote site, unloading the data, and then starting up the automatic replication. Automatic failover and failback functions enable nondisruptive testing of the disaster recovery deployment in addition to facilitating continuous access to data. Auto-Replication is completely array based, helping free IT administrators from the arduous task of managing host-based software as well as the economic burden of host software licenses.

PEER INTEGRATION: COMPREHENSIVE WINDOWS TOOLKIT

Peer integration provides a comprehensive software toolkit to facilitate the deployment, ongoing management, and protection of EqualLogic SANs in Microsoft Windows environments. As a Microsoft Simple SAN solution, the PS Series is certified to be interoperable and easy to set up and manage for Windows platforms. The EqualLogic Remote Setup Wizard facilitates rapid SAN configuration in these environments, while EqualLogic Auto-Snapshot Manager provides a feature-rich tool to help protect and recover data for Windows applications.

Remote Setup Wizard

Each PS Series array comes with the EqualLogic Remote Setup Wizard, an easy-to-use tool that can transform the way administrators set up their SANs. From a Windows-based system, administrators can have a PS Series array up and running in just minutes. Configuration of multipath I/O between Windows-based servers and an EqualLogic SAN—a multistep operation with the basic Microsoft iSCSI driver—can be vastly simplified with the Remote Setup

"EqualLogic system monitoring capabilities provide administrators with a comprehensive view of the health and status of their SAN."

> **"Peer integration helps simplify the deployment of shared SAN storage in Windows-based application environments, helping eliminate the worries of time-consuming storage management."**

Wizard. This tool helps shorten configuration time and ensure that multipath I/O has been configured properly, helping maximize availability and performance. When administrators need more capacity or performance, they can also use the Remote Setup Wizard to add members to their existing SAN.

Auto-Snapshot Manager for Windows

EqualLogic Auto-Snapshot Manager is a feature-rich application for protecting and recovering Windows application data, including protection for systems using Windows NT File System (NTFS) and Microsoft SQL Server™ databases. Auto-Snapshot Manager, included with all PS Series arrays, is an intuitive, easy-to-use GUI that creates "clean" file system and SQL Server database snapshots, clones, and remote replicas. Integration with Microsoft Volume Shadow Copy Service (VSS) provides hardware-based snapshots with full or differential copies. A scheduler with e-mail notification options is included to help simplify maintenance operations. In the event that a file system or SQL Server database recovery is necessary, Auto-Snapshot Manager provides several quick restore options to help maximize the availability of critical data and applications.

Auto-Snapshot Manager also allows administrators to safely take snapshots of, back up, and restore their Microsoft Windows Server® 2003 OS–based applications with advanced third-party backup software modules by supporting VSS and VDS interfaces. This capability helps eliminate the extensive system configuration and scripting that other systems may require for snapshot-based backups.

Peer integration helps simplify the deployment of shared SAN storage in Windows-based application environments, helping eliminate the worries of time-consuming storage management and free up time to focus on delivering enhanced Windows-based business solutions.

FLEXIBILITY WITH LOW TOTAL COST OF OWNERSHIP

The EqualLogic PS Series represents an advancement in storage economics, from purchase and setup to operation and upgrades. Unlike traditional SANs, the EqualLogic PS Series comes complete with enterprise-class software features and applications at no additional charge—there is no additional software to install or service costs to incur to initiate the data management and protection features for this enterprise-class SAN. EqualLogic arrays scale on demand and online, helping enterprises increase their storage resources without disrupting the application environment or budget. This packaging model, combined with the ease of use and automated intelligence of PS Series arrays, enables EqualLogic systems to provide a high return on investment.

Whether used to consolidate a direct attach storage infrastructure, migrate data from an existing SAN, streamline data protection processes, or just add capacity, the EqualLogic PS Series offers a family of high-performance, self-managing storage arrays designed to meet the requirements of SAN or network attached storage environments for organizations of all sizes. Based on the patented peer storage architecture, the PS Series of storage arrays is designed to be comprehensively interoperable and upgradable, seamlessly scaling arrays without disrupting application or data availability.

John Joseph is vice president of marketing of the Dell™ EqualLogic iSCSI storage business. He has worked in the computer industry for 22 years, for the last 5 years leading the EqualLogic marketing team. John has a B.S. in Mechanical Engineering from Worcester Polytechnic Institute and an M.B.A. from Clark University.

Eric Schott is a senior director of product management for the Dell EqualLogic product family, where he is responsible for strategy and planning for iSCSI storage networking products. Eric has more than 25 years of experience developing enterprise technology products; has extensive expertise in storage, real-time systems, log-based file systems, volume managers, and cluster products; and frequently contributes to articles in leading publications and speaks at key industry events.

Kevin Wittmer is director of product marketing for the Dell EqualLogic product family. He has more than 21 years of experience developing and bringing to market enterprise storage products. He is currently the vice chair for the Storage Networking Industry Association (SNIA) IP Storage Forum board of directors, and previously served on the board of directors for the SATA International Organization (SATA-IO). Kevin has a B.S. in Electrical Engineering from Rensselaer Polytechnic Institute and an M.S.E.E. and M.B.A. from Worcester Polytechnic Institute.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

EqualLogic:
www.equallogic.com

FLEXIBLE VIRTUALIZATION WITH EQUALLOGIC PS SERIES iSCSI STORAGE ARRAYS

By Timothy Sherbak

The advent of Internet SCSI (iSCSI) technology, combined with the latest in server virtualization software, can offer enhanced capabilities and benefits to enterprises of all sizes. This article provides guidance on building a virtualized infrastructure with EqualLogic® PS Series iSCSI storage arrays and VMware® virtualization software.

Related Categories:

EqualLogic

Internet SCSI (iSCSI)

Storage

Virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

Internet SCSI (iSCSI) is a flexible and powerful storage area network (SAN) protocol that can deliver superior capabilities and benefits for enterprises of all sizes. In addition to providing enterprise-class data availability and performance, the iSCSI protocol enables breakthrough virtual storage designs that parallel the advanced designs of server virtualization technologies such as VMware Infrastructure 3.[1]

Enterprises of all sizes are building flexible storage infrastructures using iSCSI and advanced virtualization technologies that let them allocate and shift SAN resources dynamically in response to the demands of their virtualized server environments. This article describes a virtualized infrastructure that applies storage and server virtualization technologies to cost-effectively achieve a flexible, high performance, dynamic IT infrastructure that is simple to manage and scale.

SERVER VIRTUALIZATION AND STORAGE VIRTUALIZATION

Server virtualization is one of the most significant technology trends in IT today. Managing server hardware and software as separate infrastructure components makes it possible to significantly increase data center operating efficiencies. At a high level, server virtualization works by encapsulating guest operating systems into a set of virtual machine (VM) files. The OS and its unique execution variables and associated data sets are maintained in these files.

Given sufficient memory and disk storage, a single industry-standard server such as a Dell™ PowerEdge™ server running virtualization software can support many guest VMs. Applications and data are processed in exactly the same way as in the physical environment, using the same OS facilities for making system requests. The only difference is that the OS is now a virtualized guest running in a virtual system environment.

VM files can be transparently migrated from one server hardware platform to another, giving administrators the ability to allocate and shift physical server resources in response to changes in application workloads. A scalable, highly resilient, flexible server environment is created, helping increase resource utilization and IT flexibility and reduce operating costs.

Storage virtualization is typically defined as a technology that allows discrete storage systems to operate as a single resource. In light of recent advancements in server virtualization, the concept of storage virtualization is being further refined as a way to create an abstraction layer between the storage hardware and logical data volumes.

Given sufficient protocol support (as with iSCSI), virtual storage products are now being designed that allow data volumes to be located and striped across multiple (and diverse) physical storage resources,

[1] For more information on iSCSI and virtualization, see "iSCSI: Changing the Economics of Storage; Part 2—Deploying iSCSI in Virtualized Data Centers," by Matt Baker and Travis Vigil, in *Dell Power Solutions*, August 2007, DELL.COM/Downloads/Global/Power/ps3q07-20070401-Baker.pdf.

including storage systems, RAID groups, disk types, and controllers. This approach not only helps increase performance and scalability, but also allows data volumes to be transparently moved from one set of resources to another without disruption to the operating systems and applications that are utilizing the data. A scalable, highly resilient, flexible storage environment is formed, helping increase storage utilization rates and significantly reduce operating costs.

NETWORKED STORAGE WITH VMWARE INFRASTRUCTURE 3

In a VMware Infrastructure 3 environment, multiple physical servers are networked into a cooperating set of computing resources (see Figure 1). In lieu of a standard Microsoft® Windows® or Linux® OS, each of the x86-based standard servers run a copy of VMware ESX Server, which provides the abstraction layer between the physical servers and guest operating systems. Multiple ESX Server systems are organized into a cluster that can be controlled and coordinated by VMware VirtualCenter, a centralized management system.

Within the virtualized environment, the VM configuration parameters, OS, data sets, and applications are encapsulated in a set of files managed by the VMware Virtual Machine File System (VMFS). VMFS is a cluster file system that coordinates access across ESX Server systems, enabling the migration of live VMs across systems in the ESX Server cluster.

The power of this architecture comes from the ability to run virtually any guest OS, its applications, and its data without modifying the physical servers in the cluster. As the need arises, additional server resources can be nondisruptively added to the ESX Server cluster, and current workloads are load balanced to take advantage of the newly available resources.

Effective VMware virtualization deployments depend on shared access to storage—in other words, a SAN. SANs help ensure that each ESX Server system has ready access to the data sets of other VMs, enabling virtually immediate re-hosting of the VM. This capability helps eliminate the time-consuming necessity of copying VM files, applications, and data from one ESX Server system to another.

Networked storage (iSCSI or Fibre Channel) is required to enable powerful features within VirtualCenter, like VMware VMotion™ technology, which enables the online migration of active VMs without interruption, and Distributed Resource Scheduler, which monitors and automatically migrates VMs to help balance workloads across resources within an ESX Server cluster. VMware High Availability (VMware HA), which also requires networked storage, monitors the functioning of ESX Server systems within the cluster, and in case of a system failure, can re-host and restart VMs affected by the failure on another ESX Server system within the cluster.

In addition to the requirement for networked storage, a virtualized environment heightens the need for high-performance, highly available, resilient storage to help meet the needs of the aggregated workloads. As critical applications, production workloads, and data assets are consolidated into relatively few resources, the need for high performance, nondisruptive scalability, and continuous availability of the storage assets increases. Enterprise-class storage designed for mission-critical deployments is a basic requirement when building a virtualized IT infrastructure.

In particular, purpose-built storage architectures should be deployed that include mirrored memory write caches, fully redundant hot-pluggable components, online hot-spare disks, environmental monitoring, and enterprise-class disk drives with RAID protection. Furthermore, advanced availability features, including storage controller and I/O path failover, are required to help guarantee data access even in the case of component failure.

Consideration should also be given to evaluating the data protection software features enabled within the storage system. Storage consolidation within a SAN enables the consistent application of data protection and disaster recovery, assuming the basic SAN-based tools are available and enabled. Space-efficient, nonintrusive, snapshot-based point-in-time copies as well as efficient array-based replication tools should be considered basic requirements of the storage infrastructure for a virtualized environment. Multiple use cases exist for the application of these tools in virtualized environments, including disaster recovery, online backup and quick recovery of VMs, simple extraction of lost or corrupted data files, rapid VM and data set provisioning using snapshots and clones for production and temporary use, and server-less tape- and disk-based backup and recovery of the organization's data assets.

A thorough evaluation of the storage platform should be made prior to deployment, reflecting not only on the robustness and feature set, but also on the total cost of ownership over its useful life, with particular consideration of the required system growth in terms of both performance and capacity.

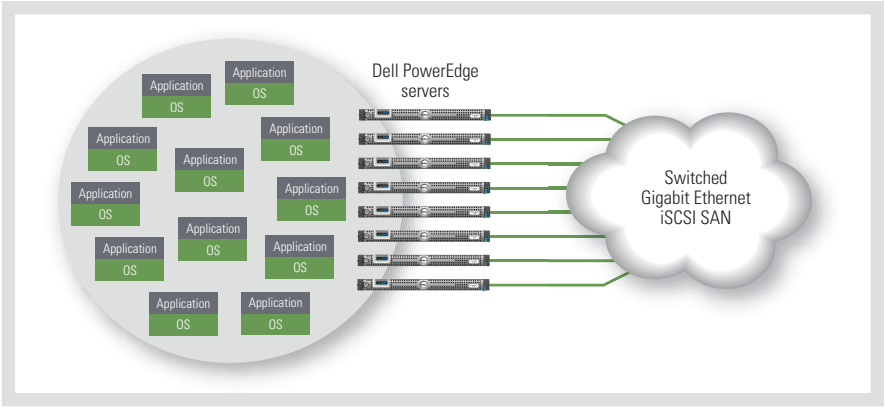

Figure 1. *VMware ESX Server cluster with SAN storage*

iSCSI and VMware Infrastructure 3

With the advent of VMware Infrastructure 3, iSCSI initiators are integrated into the ESX Server kernel, allowing native access to iSCSI storage directly from ESX Server systems. Both hardware initiators and software initiators are available. Network interface card teaming (multiplexing a single logical connection across multiple interfaces) is supported with the software initiator. The hardware initiator includes support for multipathing as well as boot from SAN.

iSCSI support broadens the potential for full-fledged VMware virtualization deployments in small and medium businesses by helping remove the need for a costly and complex Fibre Channel–based SAN deployment. The acquisition, implementation, and operational costs of an iSCSI-based SAN are intrinsically lower than those of a Fibre Channel–based SAN, helping remove an economic barrier to the adoption of server virtualization technologies. Advanced iSCSI-based system architectures can help dramatically change enterprise expectations for large deployments as well. A truly flexible iSCSI-based virtualized environment can help not only reduce costs, but also simplify deployment, provide comprehensive storage management and data protection functionality, and enable seamless VM mobility between hosts.

The addition of iSCSI protocol interfaces to classic storage system designs does not fully exploit the potential of the iSCSI protocol and, in turn, does not adequately solve the storage management challenges faced today by IT managers. Even with the introduction of iSCSI, storage administrators using classic storage systems typically require a rarified level of knowledge for configuring and tuning storage arrays, RAID geometries and data layout considerations, application workload analysis, forced data migration, and complex system upgrades.

ADVANCED VIRTUALIZATION WITH THE EQUALLOGIC PS SERIES

Storage virtualization, a key feature of EqualLogic PS Series arrays, can help simplify provisioning and ongoing management, increase storage utilization, provide virtually unlimited scalability in capacity and performance, and enable online migration of data sets among controllers and storage tiers. By consolidating network-based storage into a simple, flexible, consolidated pool of storage designed to grow in capacity and performance (as in an EqualLogic storage environment), storage virtualization can help provide an exceptional virtualization deployment.

Historically, the adoption of storage virtualization technologies in SAN environments has been limited to enterprise data centers needing a tool for online data migration between heterogeneous storage environments. This add-on design allows for heterogeneous storage and data movement, but does not simplify management of the configuration. Most of these designs have been implemented as external appliances sitting within the SAN. These devices generally add complexity to an already complex environment—introducing additional points of management within the SAN, masking value-added features of the attached arrays, and limiting the performance and scalability of the consolidated SAN storage pool to a single gateway hosting the back-end storage.

Advanced storage virtualization technologies, like those available with PS Series arrays, offer the ability to virtualize storage at the storage device level. This approach helps achieve scalable performance and capacity and reduce overall management by aggregating multiple controllers into a cooperating set of resources—that is, by virtualizing volumes not only across disks within a controller but also across storage systems in a SAN.

EqualLogic PS Series iSCSI storage arrays are an excellent example of such a "scale-out" architectural design. Despite their powerful simplicity, such designs are virtually nonexistent in the Fibre Channel–based SAN world, primarily because of architectural constraints inherent in Fibre Channel network deployments.

For example, devices in a Fibre Channel–based SAN are attached to the network via a World Wide Name (WWN), a physical port address specifically assigned and encoded within the device hardware or firmware. WWNs are assigned for all devices within the SAN, including each physical port within each host server and each physical port within the storage controllers. Furthermore, data paths between the hosts and the storage array are statically set when the host is added to the SAN.

Herein lies the inflexibility of Fibre Channel–based SANs. The topology of the SAN is intrinsically hard coded into the environment, making changes within the infrastructure burdensome, time-consuming, and error prone. Expansion of resources in the host, fabric, or storage layer propagates changes throughout the infrastructure, causing intrusive downtime to the applications and infrastructure.

In contrast, TCP/IP networks support virtual addressing and dynamic routing, whereby paths through the network are not statically defined. With Dynamic Host Configuration Protocol (DHCP), addresses can be dynamically assigned, or through address proxying, physical addresses can be virtualized, making it possible for devices in the network to transparently act on behalf of other devices on the network. If the IP address of the resource being accessed is known from within the network, the entity (such as a SAN initiator) requesting access can find the resource dynamically without prior knowledge of the paths or the resource's physical address.

IP address proxying helps make storage access fully virtualized across multiple EqualLogic PS Series storage controllers. In an EqualLogic group, for example, each member array has 3 active physical Gigabit Ethernet ports; a four-member group has a total of 12 active Gigabit Ethernet ports, providing 12 Gbps of peak bandwidth. Each of these Ethernet ports is assigned an IP address. The members within the group are aware of and coordinate use of these multiple IP addresses and the underlying physical resources. External access to the group by all hosts within the SAN is exclusively addressed

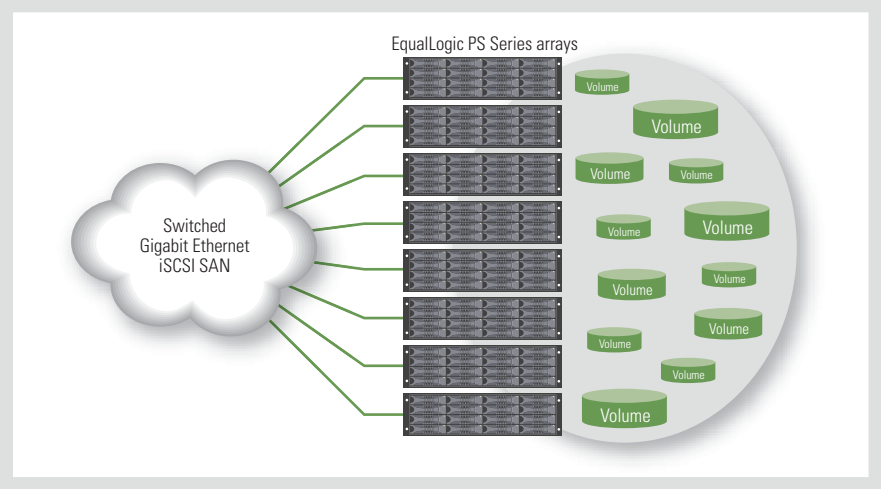

Figure 2. *EqualLogic virtualized iSCSI-based SAN*

through the group IP address, a unique IP address that transcends all of the underlying member IP addresses. The only IP address known by the attached hosts in the SAN infrastructure is the group IP address. I/Os may be serviced by any physical Gigabit Ethernet port.

Shielding host systems within the SAN from being aware of the physical topology of the storage serving its data is vitally important, particularly in a virtualized server environment. By virtualizing the physical I/O ports, a flexible storage utility is created. Pooling storage assets within the SAN is now possible, enabling a rich set of features to help maximize storage utilization and effectiveness, extending beyond the limitations of a single physical end-to-end connection to a single storage controller.

Data is accessed by the host without intimate knowledge of where in the storage utility the data is stored, providing the opportunity for data volumes to span controller boundaries and exploit the combined resources of multiple controllers. Controllers coordinate among themselves to help balance workloads and optimize storage resource utilization.

Changes in storage infrastructure can be achieved seamlessly, without affecting host connectivity or data access. The iSCSI-based infrastructure can be scaled nondisruptively, automatically applying the resources of the additional controllers and disks to help extend performance and capacity. Multiple tiers of storage can be seamlessly integrated in the utility, providing for automatic placement of data within a volume to the appropriate tier based on the volume's access patterns.

A flexible storage architecture is particularly appealing in a virtualized server environment because the scale-out architectures of both server and storage infrastructures complement each other, providing an end-to-end virtualized infrastructure. By deploying both virtualized servers and virtualized storage together, IT managers achieve an infrastructure that is particularly adaptable to changing business requirements by providing nondisruptive workload migration and balancing across all physical resources (both storage and server resources) as well as online resource expansion that can immediately apply to preexisting server and storage workloads without intervention and tuning. For more information on the advantages of this type of environment, see the "Integrating server and storage virtualization" sidebar in this article.

EqualLogic virtualized iSCSI-based SANs

A virtualized iSCSI-based SAN complements a VMware ESX Server farm. The foundation of an EqualLogic virtualized iSCSI-based SAN is the PS Series storage array (see Figure 2). Each array is designed to be fully redundant, containing disks, multiple high-performance network interfaces, redundant controllers with mirrored battery-backed caches, and other advanced features. The disks can be automatically protected with RAID (RAID-5, RAID-10, or RAID-50) and hot spares. Multiple models of arrays exist, leveraging high-performance 10,000 rpm and 15,000 rpm Serial Attached SCSI (SAS) drives or cost-effective, high-density Serial ATA (SATA) II drives.

A PS Series storage group comprises a single array or multiple arrays. A group is a virtualized resource, appearing to ESX Server systems as a single entity that offers network storage access to a single large pool of storage—a SAN composed of a single virtually scalable, high-performance storage system. Unlike a conventional SAN, in which multiple controllers appear as independently managed islands of storage, each group member in an EqualLogic SAN can cooperate with other members to automate resource provisioning and performance optimization.

ESX Server administrators can create data volumes from within the EqualLogic group storage pool. The group exports its volumes as iSCSI targets protected with security measures that include authentication and authorization. Upon connection, ESX Server systems work with EqualLogic volumes as VMFS data stores, which contain multiple VMs and virtual disks. A variation on this is a raw device map in which the SAN-resident volume is mapped directly to a corresponding virtual disk within the ESX Server environment, circumventing the VMFS data store. A third alternative is to employ the native iSCSI software initiator of the individual guest OS to connect directly to the SAN-resident volume. This last option allows administrators to use all array-based functionality and allows seamless integration with high-level applications, particularly backup applications that use Microsoft Volume Shadow Copy Service (VSS) or Virtual Disk Service (VDS) as an underlying protocol.

Automatic load balancing

Volumes are distributed among the PS Series group's member arrays, with data placement and access continually adjusted for optimal performance as resources are added or workloads change. When an array is added as a group member, its disk space is added to the group's storage pool. Volumes can be automatically re-striped and distributed across the members of the storage pool. Controller resources are also dynamically adjusted based on the workloads being generated by the ESX Server farm. Data and network I/O to the group can be automatically load balanced across the group members' resources.

A key feature of PS Series arrays that sets them apart from traditional storage is their dynamic scalability and automated configuration, enabling administrators to flexibly meet changing needs while maintaining availability. As capacity and performance requirements increase, a group can be scaled linearly in both capacity and performance—all while online. New members "learn" configuration and performance information from the group—with no manual intervention. Data and client connection load balancing can occur automatically as the group scales. I/O activity is monitored, and data and network connections are adjusted as needed.

Nondisruptive scalability

The scalability model allows for automated, online expansion across storage dimensions, and the EqualLogic PS Series architecture helps nearly eliminate downtime caused by expanding or managing a storage system. Because capacity can be added so easily, IT managers need to buy only the storage necessary for today's applications, helping ease budget constraints caused by excessive purchases.

Additional advantages of the virtualized SAN environment include the transparent application of storage tiers. EqualLogic storage groups can use multiple storage tiers contained within the storage pool and automatically optimize data placement based on workload. Alternatively, storage tiers may be segmented into independent resource pools to help guarantee specific resources to specific workloads, concurrently maintaining the flexibility to migrate volumes online from one tier to another, transparently to the ESX Server hosts.

Quick, intelligent provisioning

Within the EqualLogic environment, storage is quickly provisioned simply and easily with the click of a button. The key functions needed to configure, manage, and scale storage are automated, helping cut administration tasks for volume and capacity growth. Given the intelligence built into the PS Series software, decisions with respect to RAID type and data layout can be made automatically at the time of provisioning, and optimized as workload patterns for the newly provisioned storage evolve. As new resources become available, the resources can be automatically applied where and when they are needed. By helping eliminate complex tasks and

INTEGRATING SERVER AND STORAGE VIRTUALIZATION

Combining server virtualization with the storage virtualization of EqualLogic PS Series iSCSI arrays can enable organizations to create powerful, highly flexible virtualized environments, offering a variety of advantages:

- **Aggregation of virtualized assets on consolidated hardware:** Operational procedures and best practices can be standardized and consistently applied to both storage and server assets (both physical and virtual), helping provide a resilient infrastructure with high resource utilization, high levels of service, and enhanced protection of information assets.
- **Simple, centralized management:** Increased management efficiencies can be achieved by centralizing management through intuitive, graphical management tools accessible from virtually anywhere on the network, providing a comprehensive view to provision, monitor, and manage the entire virtualized infrastructure.
- **Flexible and quick deployment of virtualized resources:** Organizations can quickly adapt to changing and growing business needs by reducing the time to provision and deploy new applications through quick provisioning methods available for both servers and storage.
- **Online, nondisruptive resource reallocation and expansion:** As workflows and business priorities change, both storage and server resources can be simply reallocated online, with no disruption to operations. In addition, physical resources can be easily expanded online without downtime. Workloads can be automatically rebalanced across these newly available resources without disruption to applications.
- **Common IP network–based infrastructure:** The IT environment is simplified by basing operations on IP networking, including the interconnect for client access, inter-server communication, storage access, and off-site data replication. The organization's inherent IP networking expertise is leveraged, helping reduce training and ongoing management costs.
- **Enterprise-class resiliency:** With redundancy built into the physical server, network, and storage architecture, as well as component failure detection and failover software implemented within each layer of the infrastructure, overall reliability, availability, and service levels are enhanced.
- **Advanced data management and disaster recovery:** A rich set of server and SAN-based data protection tools help ensure the organization's critical assets are protected and immediately recoverable at the local or a remote site.

enabling fast and flexible storage provisioning, PS Series solutions can help dramatically reduce acquisition and ongoing operational costs and make enterprise-class shared block storage practical for organizations of all sizes, from small businesses to large enterprises.

Thin provisioning extends existing EqualLogic provisioning features, helping make the "buy-as-you-grow" storage model of the modular PS Series seamless for servers and applications. Thin provisioning is an important advanced feature that enables the automatic addition of physical capacity on demand up to preset limits. With advanced thin provisioning, buy-as-you-grow storage management and virtualization are made seamless for servers and applications. When a volume is created, it can be sized for the long-term needs of the application without initially allocating the full amount of physical storage. Instead, as the application needs additional storage, capacity is allocated to the volume from a free pool.

EqualLogic thin-provisioning capabilities help provide extensive flexibility and user safety controls. These include the ability to turn thin provisioning on and off as needed for any volume, allowing users to cost-effectively test the most suitable applications and volumes for thin provisioning, with the knowledge that they can return to "normal" provisioning online. EqualLogic implementation of thin provisioning also provides enhanced alerts and controls—with proactive, user-defined threshold alarms and controls, administrators can depend on automatic space allocation without worrying about reaching allocation limits or unexpected depletion of physical storage.

Automated management

The PS Series architecture is designed to simplify storage management in several ways. RAID configuration and hot sparing is automated, and dynamic storage and network I/O load balancing can occur automatically as resources and performance metrics change. No longer must administrators manually map application data to specific physical devices and controllers.

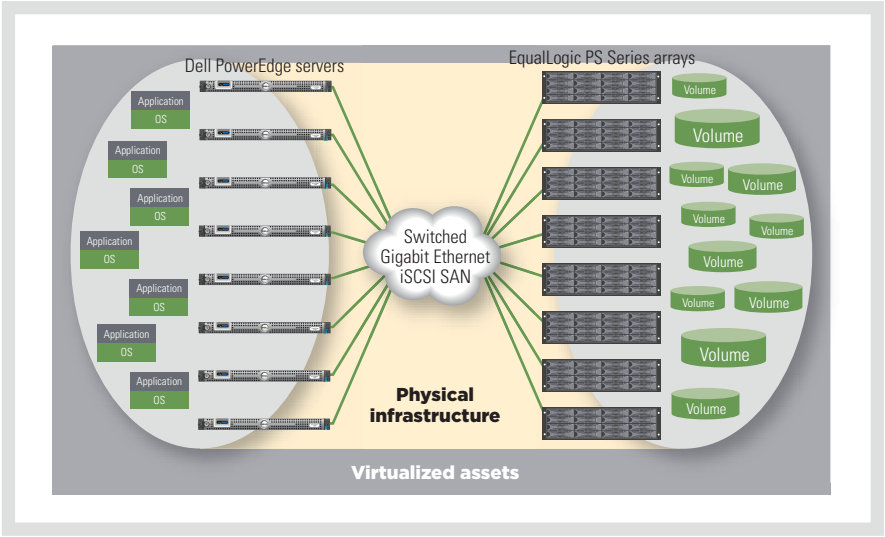

Figure 3. *Virtualized servers combined with virtualized storage*

A primary advantage of this automated virtual storage system is that storage administration remains stable as the group is expanded. All administration is performed at the group level, through the EqualLogic PS Group Manager, an intuitive, single-pane administrative console. Because the group is managed as a single logical system, the ongoing operational costs of storage management can remain fixed even as storage grows.

SIMPLE, COST-EFFECTIVE, DYNAMIC IT ENVIRONMENT

Virtualized iSCSI-based SANs, such as those enabled by EqualLogic PS Series arrays, are changing enterprise expectations of how simple a storage infrastructure can be to deploy, manage, and grow. As Figure 3 illustrates, complementary server and storage virtualization built on iSCSI-based SANs create a fully abstracted pool of physical resources—one that can provide higher performance, scalability, ease of use, and flexibility compared with classic storage array architectures.

EqualLogic PS Series iSCSI arrays are designed to reduce technical complexity and cost barriers imposed by classic SAN architectures without compromising performance, scalability, or resiliency. In enterprise data centers, a virtualized SAN utilizing PS Series arrays offers a cost-effective, comprehensive solution to help reduce both management complexity and total cost of ownership.

Timothy Sherbak is senior manager of virtualization solutions marketing for the Dell EqualLogic product family. He has over 20 years of engineering, sales, and marketing experience in computer systems and storage networking technologies. Timothy has a B.S. in Systems and Information Science from Syracuse University and an M.B.A. from the Massachusetts Institute of Technology Sloan School of Management.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

EqualLogic:
www.equallogic.com

EqualLogic and VMware:
www.equallogic.com/vmware

EqualLogic PS Series demonstrations:
www.equallogic.com/demo

BUSINESS CONTINUITY FOR SMBs WITH EQUALLOGIC PS SERIES iSCSI STORAGE ARRAYS

By Kevin Wittmer

Robust business continuity has historically been out of reach for small and medium businesses (SMBs) because it was often too costly and overly complicated. EqualLogic® PS Series Internet SCSI (iSCSI) storage arrays enable organizations of all sizes to deploy an enterprise-level storage infrastructure offering comprehensive data and application protection.

Related Categories:

EqualLogic

Internet SCSI (iSCSI)

Storage

Visit DELL.COM/PowerSolutions for the complete category index.

Costs and complexity have typically placed robust business continuity (BC) out of reach for small and medium businesses (SMBs). While large companies could afford expensive Fibre Channel storage area networks (SANs), the highly trained staff to manage them, duplicate data centers, channel extenders, and expensive replication software, SMBs were often restricted to making backup tapes and storing them off-site.

Today that has changed, and organizations of all sizes can cost-effectively build an infrastructure that offers robust protection without exorbitant costs or specially trained staff. Advanced functionality that was first developed for mainframe environments has been enhanced over time, making it available to Microsoft® Windows®, Linux®, and UNIX® environments. As a result, company size and platform selection no longer determine functionality level. This is due to multiple technological advances—in particular intelligent storage, Internet SCSI (iSCSI) connectivity, server virtualization, and wide area network (WAN) acceleration.

These advances carry significant advantages for SMBs, because data is just as critical for these organizations as it is for large enterprises—and downtime for SMBs can be much more damaging than it is for large enterprises. This article examines how organizations of all sizes can create a cost-effective infrastructure that delivers comprehensive protection—from simple data protection to disaster recovery (DR) to BC.

THE STAGGERING COST OF DOWNTIME

In today's highly competitive business environment, data center outages can be devastating. Regardless of the cause—hurricane, fire, accident, hacker attack, or some other type of disaster—production downtime is not only costly, but in some cases can be ruinous. With a mobile workforce, global customers wanting to do business around the clock, and continually greater dependence on technology, companies need to not only protect data, but continue business operations virtually uninterrupted. The cost of downtime, depending on the industry, can range from thousands to millions of dollars per hour—costs due not only to DR expenses, but also to lost sales, customer defection, and lack of productivity. Add to that a damaged reputation in the marketplace and diminished shareholder confidence, and the cost of downtime can be staggering.

However, organizations that can continue business operations through outages, large or small, can gain competitive advantage—and sometimes even take market share from competitors.

TAPE: GOOD FOR ARCHIVING, BAD FOR AVAILABILITY

Tape backup has typically been the most common protection solution for SMBs, and remains a good long-term archiving method for most large and small organizations. However, there are several reasons that necessitate that organizations complement their tape solutions with other BC solutions. First, backup windows are shrinking because of the huge growth in data volumes, requirements for long retention and fast access, and generally high reliance on data and technology. Second, because tape backup can be time-consuming, many organizations cannot back up often enough to adequately protect themselves.

Instead, organizations need a continuum of protection schemes that include tape, storage array–based data protection, remote replication for recovery after a failure or disaster, and BC during outages and common IT maintenance procedures.

DATA PROTECTION

The first line of defense is to protect data where it is stored, and that means keeping storage units running as well as providing safety features. EqualLogic PS Series iSCSI arrays are fundamentally built to protect data with enterprise-level availability and reliability. Redundant, hot-swappable components—including disk drives, controllers, network interfaces, power supplies, and cooling fans—mean that component failure need not result in downtime or data loss. The system monitors disk drives in the background to help detect and correct problems before they occur; in addition, should a drive fail, a spare is automatically placed into service and configured. Disks and arrays "learn" the configuration from those already installed, so no administrator intervention is necessary. Controller caches are mirrored and battery backed, and the arrays support multiple RAID levels.

Of course, backup is a basic part of data protection—simply backing up to tape can help protect against complete data loss, although restoring from tape can be time-consuming. Disk-based snapshot functionality is a big improvement—snapshots are created quickly and easily without disrupting operations, and can provide multiple restore points to minimize lost work. Many organizations take snapshots and then run tape backups from them, helping keep data available by virtually eliminating backup windows, speeding restore operations, and minimizing errors.

DISASTER RECOVERY

The focus on DR has been greatly magnified in recent years. Increasing virus and hacker attacks, hurricanes, and electrical brownouts serve as reminders of how vulnerable data is.

The keywords for DR are *copies* and *distance*. For DR, organizations must have more than one data copy, and they must store those copies some distance away from the primary data center on different physical systems—remote replication is the standard for DR. Today, remote replication is available for iSCSI-based SANs, and does not use the specialized equipment that Fibre Channel networks require. iSCSI enables organizations to replicate data across standard IP networks and provide multiple replicas from various points in time. Should a virus corrupt data, administrators can almost immediately roll back to a previous replica to help prevent data loss. Deciding how and where to replicate depends on an organization's needs and available locations. Some organizations may replicate from the primary data center to one remote location; others may replicate the same data to multiple locations. Organizations with branch offices often replicate from each branch to a central DR site, and then back up data from there.

Configuring a DR implementation depends on two important factors that each organization must identify: recovery time objective (RTO) and recovery point objective (RPO). RTO indicates how quickly the organization can restore data—typically minutes or hours, and in some cases days. Some operations and data types may only tolerate a very short RTO, while others can survive longer delays. RPO indicates how much data loss the organization can tolerate, and that determines how often it replicates data—every hour, three times per day, and so on. Many organizations define different RTOs and RPOs across the enterprise—uniformity is not important as long as they can easily and cost-effectively match data types to protection levels.

BUSINESS CONTINUITY

BC differs from data protection and DR because it describes not only a level of protection that helps speed recovery, but also a strategy that helps reduce the importance of recovery speed. For BC, the focus is not on how long it takes to get back in operation—the focus is on *staying in operation* regardless of failure, outage, attack, or corruption. For example, if an organization has a duplicate data center at another location with data copies, it can quickly bring operations online at that location.

Another way to continue business is with built-in storage features that help minimize disruption during standard IT maintenance tasks. Because of the patent-pending EqualLogic page-based volume management, data on PS Series arrays can be automatically moved *while it is in use*. This is an advancement of revolutionary proportions because it enables data movement without downtime. This capability means that organizations can add or move capacity among storage tiers, and automatically balance loads across disks and arrays, without interrupting users—to help IT staff get their job done and enhance performance while business continues.

Common performance improvement tasks such as load balancing across disks and arrays are done automatically by PS Series arrays—helping both avoid interruptions to business and optimize performance. Automatic disk sparing and multipath I/O enable the array to automatically replace a failed disk or manage network throughput without intervention or downtime. These features can add tremendous value during

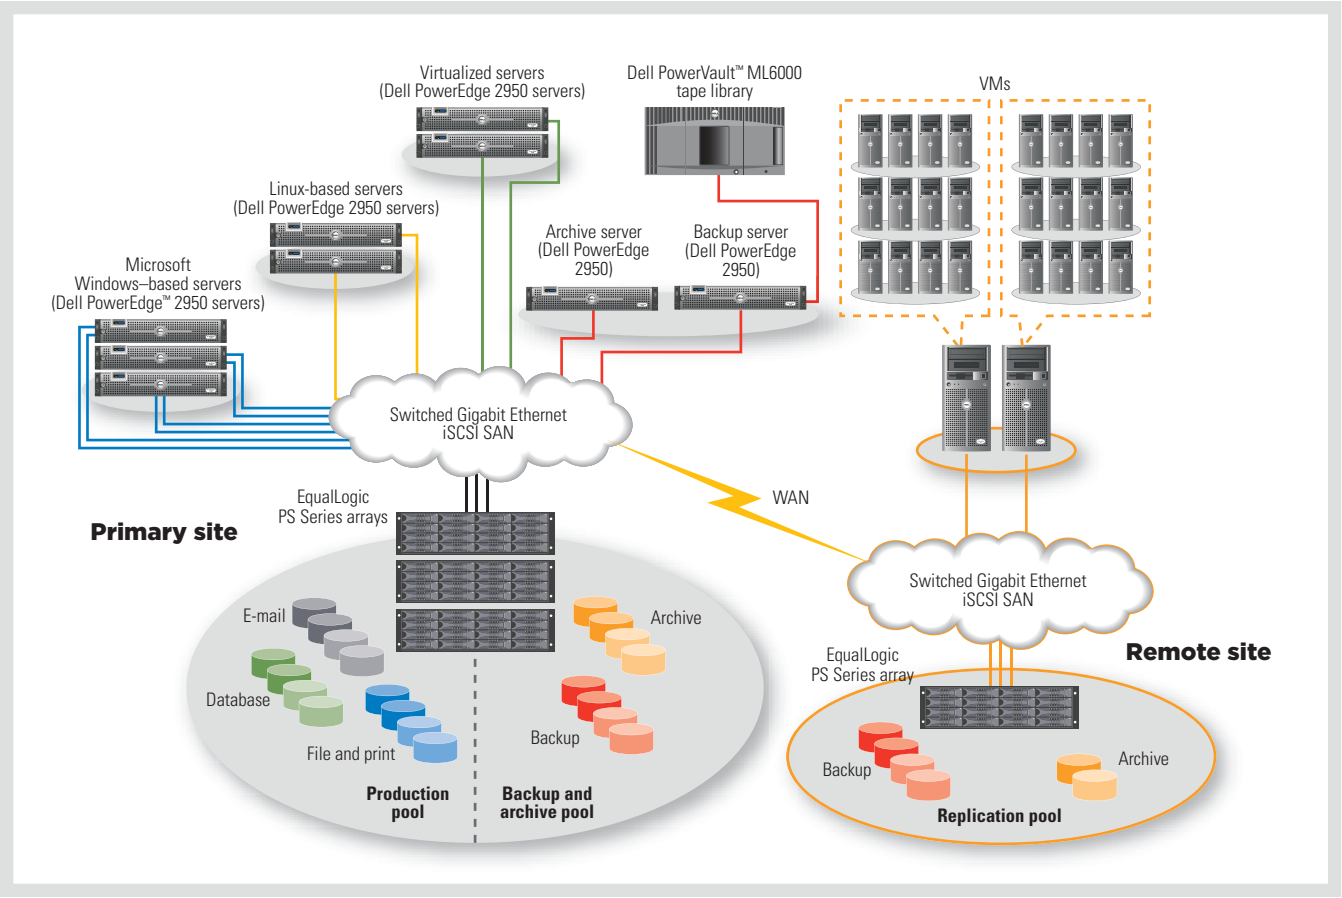

Figure 1. *EqualLogic PS Series iSCSI arrays enable business continuity for organizations of all sizes*

regular business operations, and can actually help keep problems and accidents from becoming disasters.

BUSINESS CONTINUITY STRATEGIES AND TECHNOLOGIES

SMBs need the same comprehensive protection as large organizations—but, in addition, they need it to be easy to use, nondisruptive, and cost-effective.

Data center equipment costs have dropped sufficiently to make *duplicate data centers* possible for many organizations—for example, EqualLogic PS Series iSCSI arrays make remote replication much more cost-effective than it has been in the past because the arrays are not cost-prohibitive and they include base and advanced software functionality (see Figure 1). In addition, while physical server costs have dropped, the development of *server virtualization* has dramatically altered the landscape. Products from companies such as VMware enable multiple virtual machines (VMs) to reside on a single physical host—as a result, organizations can build a second data center without duplicating hardware. For example, an organization's production data center may have 100 physical servers running business applications, with data shared on a SAN. At its secondary data center, it can run those 100 servers as VMs using only 10 physical servers, with each physical server running 10 VMs. This approach enables the business to maintain a duplicate data center cost-effectively to help remain operational following a failure or other outage. Some organizations simply could not have BC without using virtualized servers.

Server virtualization also lets organizations replicate multiple applications to the same target that normally would not work well on the same system. For example, administrators might replicate servers running Microsoft SQL Server™, Microsoft Exchange, and Oracle® application software to different VMs on the same replication target, which typically costs much less than buying individual physical servers for each application. Different versions of the same application can be treated the same way—Microsoft Exchange Server 2003 and Exchange Server 2007 can both be replicated to the same target server, but to different VMs.

Virtualization also enables organizations to share data from a SAN and move VMs between physical servers to enhance application management—and administrators can make these modifications without users even knowing it. For primary or remote data centers, this capability helps provide BC during maintenance and performance optimization.

STORAGE TIERING

Another key to keeping BC cost-effective is to not treat all data the same way. Instead of continuously replicating all data to a remote location on high-speed disks so that any piece of data is almost instantly recoverable, most organizations set up storage tiers in which some data is replicated less often, to higher-capacity (less expensive) disks with slower performance—such as Serial ATA (SATA)—than the organization's most critical data. Some data should simply be archived to tape—available if needed, but at a lower cost than disk storage.

Storage tiers should be linked to the business value of the data involved, the ability to re-create it, and the requirement for speedy access. Organizations can first identify the required RTOs and RPOs for all data and set up tiers of storage and services that help maximize utilization while lowering total cost of ownership. Also, platforms like Microsoft Windows Storage Server 2003 and software like Microsoft Exchange Server 2007 include single-instance storage, enabling organizations to store only a single copy of identical data, such as e-mail attachments that go to multiple mailboxes. This approach can help streamline data movement and minimize bandwidth requirements for DR and BC.

WAN OPTIMIZATION

Bandwidth is expensive, and accounts for a significant portion of the cost of both DR and BC; in addition, bandwidth latency creates challenges for recovery and restore processes. These issues can affect the selection of remote sites, the amount of data that can be replicated, RTOs, and RPOs. However, today organizations like F5, Riverbed, and Citrix offer WAN acceleration and optimization technologies that enable organizations to replicate a large amount of data while increasing efficiency in a cost-effective way. Using various compression, de-duplication, and optimization techniques, these solutions accelerate WAN traffic in ways that can be significantly less costly than purchasing more bandwidth.

CLUSTERING, iSCSI SAN BOOT, AND THIN PROVISIONING

Clustering servers is a way to create BC, because applications can be distributed across multiple nodes for performance optimization as well as failover. This approach is aided by SAN implementations, because data can be shared among servers. Booting servers from the array itself, another standard feature of EqualLogic PS Series arrays, also enables BC; if a server fails, another can be deployed almost immediately using the boot volume on the SAN. This feature also enables central provisioning and management of VMs, as well as implementation of diskless servers such as blades, making a robust infrastructure more cost-effective than it has been in the past.

Thin provisioning, included as a standard feature in PS Series arrays, is another strategy designed to prevent downtime by enabling applications to grow nondisruptively—administrators can add capacity on demand up to preset limits. For systems that do not have online expansion, IT staff can allocate virtual disk capacity up front but not logically provision it. As a result, applications can grow when needed without downtime by allocating increments of actual capacity on demand from a free pool. This strategy does require some diligence, however, because real and perceived capacity limits may differ.

BUSINESS CONTINUITY FOR ORGANIZATIONS OF ALL SIZES

EqualLogic PS Series iSCSI arrays enable organizations of all sizes to create an enterprise-level storage infrastructure that provides data protection, DR, and BC in a cost-effective way. iSCSI connectivity means the SAN uses standard Ethernet instead of a complicated and expensive Fibre Channel network. This approach helps make the array itself much more cost-effective than it would be otherwise, so organizations can apply the savings to create robust DR and BC implementations. iSCSI also helps eliminate the need for IT staff to be specially trained in a new network protocol—the organization can leverage the existing skills of its staff.

Equally important to BC, load balancing and other management tasks are handled *by the array*, not by administrators. Continuously monitoring themselves, PS Series arrays allocate disk space—along with connectivity, security, and performance—dynamically for every application as needed. This proactive management helps prevent downtime and keep business applications running. Remote replication can be done over a WAN without expensive add-on components like channel extenders. Organizations can replicate data between PS Series arrays using Ethernet without buying additional software and licenses for various servers, because the arrays come with the necessary functionality already included: all of the management and protection features described in this article come standard with all PS Series arrays.

Because EqualLogic PS Series arrays include this advanced functionality, upgrading from direct attach storage can provide organizations with data protection, DR, and BC capabilities at the same time. They do not have to think of building a BC infrastructure over time—when they install a PS Series SAN, they get all the capabilities in one cost-effective package.

There is no software to add on to use the enterprise-class data management and protection features of the arrays, and as new capabilities are developed, they are delivered at no additional cost for systems under warranty or service plan. This is a fundamental difference from many storage vendors who charge for add-on software to provide snapshot, replication, and other features. Features such as these help make BC a reality for organizations of all sizes, with pools of storage, server power, and network bandwidth designed to operate like an integrated business utility. A few years ago this kind of environment was typically reserved for the very largest, wealthiest organizations. Today, a robust BC infrastructure is available and cost-effective for organizations of all sizes.

Kevin Wittmer is director of product marketing for the Dell™ EqualLogic product family.

THE NEXT-GENERATION DELL POWEREDGE M1000e MODULAR BLADE ENCLOSURE

By Chad Fenner

The innovative new Dell™ PowerEdge™ M1000e modular blade server enclosure is designed to meet evolving requirements for blade servers in enterprise environments—including comprehensive functionality comparable to standard rack-mounted servers, efficient energy use through Dell Energy Smart technology, upgradable components, and simplified deployment and management.

Related Categories:

Blade servers

Dell PowerEdge blade servers

Dell PowerEdge servers

Visit DELL.COM/PowerSolutions for the complete category index.

The next-generation Dell PowerEdge M1000e modular blade server enclosure begins a new generation of blade technology. Designed by Dell with over 54,000 hours of work and 30 patents solely on the architectural design of the enclosure, the PowerEdge M1000e incorporates the benefits of an enhanced power and cooling design, a flexible high-speed I/O subsystem, and integrated management capabilities designed to simplify administration tasks.

Blade enclosures offer a number of unique advantages over typical rack-mounted servers that help simplify IT data centers, including the following:

- **Increased density:** While only 40 standard 1U rack-mounted servers can fit in 40U of rack space, that same 40U can hold up to 64 blades, enabling administrators to efficiently use available space.
- **Rapid deployment:** Installing a blade enclosure in a rack takes a similar amount of time as installing a rack-mounted server. However, after the enclosure is installed, each blade can take seconds to install, while each rack-mounted server requires starting again from scratch.
- **Reduced power and cooling requirements:** Blades can help drastically reduce the amount of cooling required in a data center. In comparison with 1U servers, the PowerEdge M1000e is expected to save up to 24 percent in power requirements.
- **Reduced cabling:** One PowerEdge M1000e enclosure with 16 blades can help reduce the number of cables by over 93 percent in comparison with a 1U server. This also helps greatly reduce the number of cable failure points, helping simplify management and reduce installation costs.
- **Simplified management:** One (redundant) management card can manage up to 16 servers, while each rack-mounted server needs its own management card.

To meet the performance and efficiency needs of demanding enterprise usage scenarios, Dell designed the new PowerEdge M1000e modular blade enclosure to also provide the following:

- **Outstanding power efficiency:** Already providing revolutionary power efficiency with the PowerEdge 1955 blade server, Dell designed the PowerEdge M1000e to increase power efficiency even further.
- **Enhanced I/O switches:** Before the PowerEdge M1000e, blade I/O switches from major vendors were not upgradable, meaning that upgrades resulted in throwing away the initial switch investment and starting from scratch.
- **Reduced complexity:** PowerEdge M1000e functionality avoids requirements of many blade enclosures for expensive upgrades, complicated

software, difficult deployment, or even clumsy KVM (keyboard, video, mouse) cables off the front of each blade.

- **Comparable functionality to rack-mounted servers:** The PowerEdge M1000e enables the same functionality and features as a 1U server. A blade form factor should not preclude standard enterprise features such as hot-swappable hard drives, quad-core processors, and similar speed and total number of dual in-line memory module (DIMM) slots.

COMPREHENSIVE FUNCTIONALITY

The PowerEdge M1000e enclosure features comprehensive redundancy and a high degree of flexibility. This 10U enclosure features up to 16 blade slots, a completely passive and highly available midplane delivering up to 5.44 Tbps of total I/O bandwidth, plus support for up to six high-speed switches, six power supplies, nine fans, two Chassis Management Controllers (CMCs), and one KVM switch.

Dell designed the blades for the PowerEdge M1000e enclosure to meet stringent enterprise requirements. Each blade comes with comparable features as Dell would design into a normal 1U rack-mounted server, including two Intel® or AMD processors, eight DIMM slots, two hot-swappable hard drives, and two expansion slots for connectivity such as Ethernet, Fibre Channel, or InfiniBand.

While these blades each use only one blade slot, the PowerEdge M1000e can also handle blades of other sizes. These include double-high blades or double-wide blades. These blades are designed to fit into any PowerEdge M1000e slot regardless of size or processor. Because Dell anticipates that the PowerEdge M1000e will last for three generations of servers and five years, this support is key to helping increase the enclosure's flexibility.

EFFICIENT ENERGY USE

Energy efficiency has been a hallmark of Dell blade servers. For example, a study by Principled Technologies comparing the previous-generation Dell PowerEdge 1955 blade server system with the HP BladeSystem c-Class blade server system found that the Dell system used an average of 5–7 percent less power per blade than the HP system (up to 10 blades), in addition to achieving higher performance per watt in every configuration tested.[1]

PowerEdge M Series blade technology is designed to increase this efficiency even further by reducing power consumption at the component level, avoiding wasted energy, providing self-regulation features, and integrating many of the Energy Smart technologies available in other Dell systems as well as additional enhancements. With PowerEdge M Series blade technology, Dell takes energy efficiency to a new level. In comparison with the PowerEdge 1955, PowerEdge M Series blade technology is designed to consume up to 18 percent less power per blade using similar blade configurations.[2] In addition, a study by Principled Technologies comparing PowerEdge M Series blade technology with the HP BladeSystem c-Class and IBM® BladeCenter H (8852) found that PowerEdge M Series blade technology had up to 19 percent lower power consumption and 2 percent higher performance per blade than the HP system, resulting in up to 25 percent higher performance per watt on average with a full set of blades in each enclosure. PowerEdge M Series blade technology also had up to 12 percent lower power consumption and 13 percent higher performance per blade than the IBM system, resulting in up to 28 percent higher performance per watt on average with a full set of blades in each enclosure.[3]

Key Energy Smart components integrated into the PowerEdge M1000e include power supplies, fans, and power monitoring and management tools.

Power supplies

Standard power supplies typically achieve their highest efficiency only when under extremely high loads—which, because they do not typically run under those loads, means that enterprises often do not receive the full benefit of that efficiency. The Dell Energy Smart power supplies in the PowerEdge M1000e enclosure are designed not only to provide up to 91 percent efficiency, but also to achieve that level of efficiency at load levels typical of real-world use (see Figure 1).[4]

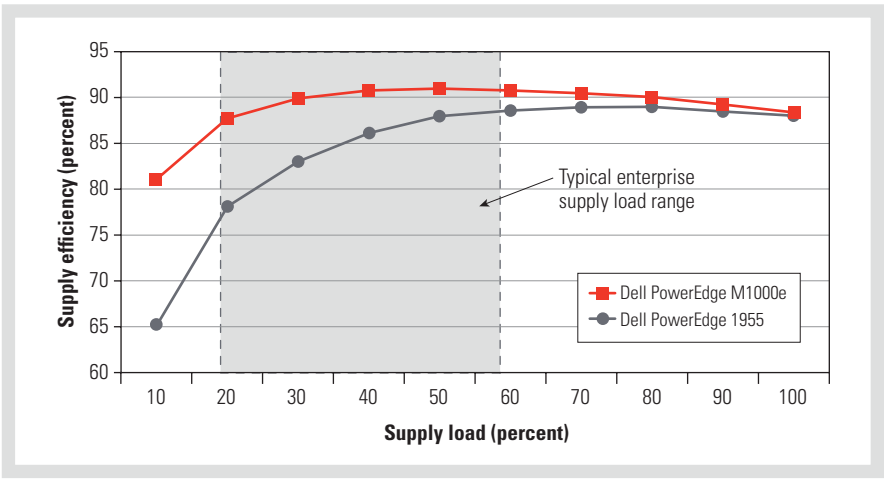

Figure 1. *Power supply efficiency curves for Dell PowerEdge M1000e and PowerEdge 1955 blade servers*

[1] "SPECjbb2005 Performance and Power Consumption on Dell and HP Blade Servers," by Principled Technologies, June 2007, DELL.COM/Downloads/Global/Products/Pedge/En/DellHPBladeServerPerfPwr_report.pdf.

[2] Based on product specifications; actual power consumption will vary based on configuration, usage, and manufacturing variability.

[3] Based on tests performed in December 2007 by Principled Technologies using the SPECjbb2005 benchmark.

[4] Based on product specifications; actual efficiency will vary based on configuration, usage, and manufacturing variability.

The PowerEdge M1000e also provides within the CMC a tool called Dynamic Power Supply Engagement (DPSE) allowing the system to automatically put lightly used power supplies into standby mode, driving up utilization on the other supplies and overall system efficiency—potentially saving up to 150 W per enclosure.[5] Administrators can set power supplies in *n* + *n*, *n* + 1, or *n* + 0 redundancy configurations.

Fans

Based on the fan design of high-performance aircraft, the Dell Energy Smart fans used in the PowerEdge M1000e are designed to cool only to the level needed for each blade and allow air to move efficiently while still maintaining moderate sound levels. The fan control algorithm senses the configuration of each blade and chooses how fast each fan needs to spin to produce the air needed to cool each blade, giving each blade the air it needs while helping avoid wasted air and power through overcooling.

Set up in zones, unused fans can turn off when not needed—for example, when some enclosure slots are empty. The PowerEdge M1000e ships standard with all nine fans at no additional cost.

Power monitoring and management

Administrators can use the PowerEdge M1000e CMC to monitor and manage power use for the whole system and for individual blades. For example, they can set upper limits on enclosure power use to help ensure it adheres to data center requirements, or do so for individual blades to help ensure that high-priority blades consistently receive the level of power they need. Administrators can receive notifications when system power consumption nears these upper limits so they can take appropriate action if needed, and can set the system to automatically throttle the blades if the upper limit is reached. The power management software also provides information on high and low power consumption to help administrators easily track enclosure or blade energy use over time (see Figure 2).

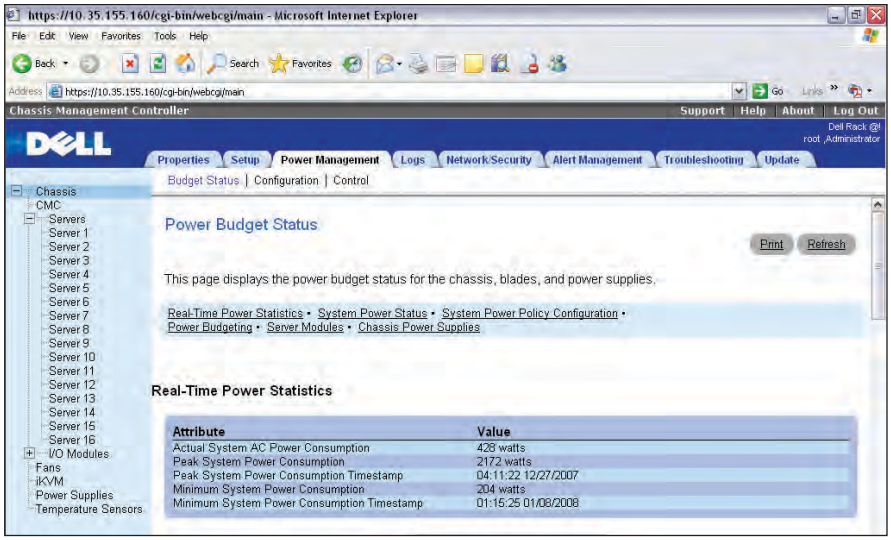

Figure 2. *Dell PowerEdge M1000e chassis management controller display showing high and low energy use*

FLEXI/O SUBSYSTEM

Because Dell anticipates that the PowerEdge M1000e will accommodate three generations of server blades over the next five years, it designed the enclosure to support not only current technologies but also those expected to be coming in the future. For example, to help ensure that the PowerEdge M1000e enclosure continues to meet the needs of leading-edge networks, it is designed to not only support current fabrics such as 10 Gigabit Ethernet, 4 Gbps Fibre Channel, and double data rate (DDR) InfiniBand, but also to eventually support future technologies such as 8 Gbps Fibre Channel and quad data rate (QDR) InfiniBand.[6]

Dell also sought to solve a significant I/O issue facing previous blade switches: no modular upgradability. Other blade switches typically require organizations seeking to move to the next technology to throw away the existing switch and start from scratch. New Dell PowerConnect™ and Brocade switches such as the PowerConnect M6220 can upgrade as the technology changes. For example, the Dell PowerConnect M6220 begins with four Gigabit Ethernet ports and has two optional ports for up to four other ports. Administrators can then begin with all Gigabit Ethernet ports and in the future add 10 Gigabit Ethernet uplinks. This design can help provide investment protection and upgradability over time. For organizations with preexisting rack-based switches, the PowerEdge M1000e contains optional Ethernet or Fibre Channel pass-through modules.

In addition, while previous server blades have typically only supported high throughput on some of their ports, every expansion port on blades fitting into the PowerEdge M1000e is designed for high-throughput fabrics. Each half-height blade comes with two Gigabit Ethernet network interface cards on the motherboard and four expansion ports. Each of the expansion ports can run at full-speed 10 Gigabit Ethernet, 4 Gbps or 8 Gbps Fibre Channel, or DDR or QDR InfiniBand, providing up to 324 Gbps of bidirectional throughput for each blade. This extreme throughput per blade, the new modular switches, and the backplane of the enclosure make up the FlexI/O subsystem and required thousands of hours of Dell research and development.

[5] Based on product specifications; actual power consumption will vary based on configuration, usage, and manufacturing variability.

[6] For more information, see "Exploring the Dell PowerEdge M1000e Network Fabric Architecture," by John Loffink, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070500-Loffink.pdf.

SIMPLIFIED DEPLOYMENT AND MANAGEMENT

Most IT environments use a mix of different server form factors. The PowerEdge M1000e enclosure is designed to fit into existing heterogeneous environments and includes features to help simplify local and remote management, including flexible design features such as a rotating LCD, CMCs for enclosure and blade management, and integration with management software such as the Dell OpenManage™ suite and Altiris Deployment Solution for Dell Servers software.

Flexible design features

Dell designed the PowerEdge M1000e for easy deployment. Simplified deployment has been a focus in previous-generation blade servers as well; for example, a Principled Technologies study comparing the Dell PowerEdge 1955 blade server system with the HP BladeSystem c-Class blade server system found that the Dell system saved more than two hours of deployment time per enclosure.[7] For the PowerEdge M1000e, Dell also offers flexible shipping options that can provide organizations with a comprehensively configured system right out of the box at no additional charge.

Once deployed, the PowerEdge M1000e is also designed for flexibility and ease of use in data center environments. A rotating LCD on the front of the enclosure helps ensure that the display is easy to read regardless of its rack position, which can be particularly useful when preloading IP addresses or checking enclosure status reports. Handles on the front of each blade are designed to help IT staff carry multiple blades at once, while color coding makes it easy to identify which components are hot swappable and which components administrators can replace themselves. And a cable management bar that attaches to the rack can help administrators manage enclosure cables.

The front of the enclosure includes a control panel to provide an easy connection to the Avocent iKVM (integrated KVM) switch on the back of the enclosure. This switch is designed to seamlessly assimilate into Avocent KVM infrastructures and can be managed through the Avocent On-Screen Configuration and Reporting (OSCAR) interface.[8]

Chassis Management Controllers

The CMC is a powerful systems management tool providing comprehensive access to enclosure components, including status, inventory, alerting, and first-boot-device configuration. Administrators can also use it to set security parameters or flexible restrictions for internal components. The PowerEdge M1000e includes one CMC as a standard feature, and can support an optional second CMC for active/passive redundancy to help ensure access to the enclosure should the first CMC fail.

Each server blade includes an integrated Dell Remote Access Controller (iDRAC) that works in conjunction with the enclosure CMC. Similar in functionality to the DRAC 5 available on other Dell PowerEdge servers, the iDRAC is designed to provide management of individual blades. This information is also available through the CMC, but administrators only managing particular blades can do so using the iDRAC even if they do not have CMC access.

Integration with management software

Administrators can use Dell OpenManage server management software to manage blades in the PowerEdge M1000e enclosure just as they do other Dell servers, including monitoring, deploying, and configuring server blades using information provided to the Dell OpenManage software by the CMCs and iDRACs. Dell OpenManage software can also seamlessly integrate with other management consoles from Microsoft, Altiris, and Novell.

In addition, to help simplify deployment, Dell now includes Altiris Deployment Solution and the Altiris Deployment Solution for Dell Servers add-on on an optional DVD with Dell PowerEdge blade servers. This software is designed to simplify and accelerate server deployment and includes rip-and-replace functionality that can help administrators quickly deploy replacement blades with the same configuration as the previous blades.[9]

POWERFUL, VERSATILE BLADE SERVERS

While blade servers can offer multiple advantages over standard rack-mounted servers, evolving enterprise requirements and technological advances have also created an opportunity for further improvement. The innovative new Dell PowerEdge M1000e modular blade enclosure is designed to meet these evolving enterprise requirements, offering functionality comparable to standard rack-mounted servers, efficient energy use through Dell Energy Smart technology, the FlexI/O subsystem with modular components, and a high-speed midplane plus simplified management. By deploying the PowerEdge M1000e in their data centers, enterprises can take advantage of these features to help reduce operational costs and create an environment that can continue growing to meet their needs both now and in the future.

Chad Fenner is a senior product manager for blade servers at Dell. He has a bachelor's degree from Trinity University in San Antonio, Texas.

[7] "Out-of-Box Comparison Between Dell and HP Blade Servers," by Principled Technologies, June 2007, DELL.COM/Downloads/Global/Products/Pedge/En/DellHPBladeServerOOB.pdf.

[8] For more information on this switch, see "Managing Dell PowerEdge M1000e Blade Servers with the Avocent iKVM Switch," by Stephen M. Hahn and Chad Fenner, in *Dell Power Solutions,* February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070529-Avocent.pdf.

[9] For more information, see "Simplified Management with Altiris Deployment Solution for Dell Servers 3.0," by Eric Szewczyk, in *Dell Power Solutions,* February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070310-Altiris.pdf; and "Simplifying Blade Server Management with Altiris Deployment Solution Rip and Replace," by Eric Szewczyk and Chad Fenner, in *Dell Power Solutions,* February 2007, DELL.COM/Downloads/Global/Power/ps1q07-20070274-Altiris.pdf.

EXPLORING THE DELL POWEREDGE M1000e NETWORK FABRIC ARCHITECTURE

By John Loffink

Modular components can be key to keeping pace with fast-moving IT advancements. The architecture of the new Dell™ PowerEdge™ M1000e modular blade server enclosure and its 10th-generation Dell server technology is designed to support both current and future network technologies, helping protect enterprise modular server investments.

Related Categories:

Blade servers

Dell PowerEdge blade servers

Dell PowerEdge servers

Visit DELL.COM/PowerSolutions for the complete category index.

Designed from the ground up to support current and future generations of server, storage, networking, and management technologies, the PowerEdge M1000e modular blade server enclosure includes the headroom for multiple generations of connectivity. As part of its modular architecture, key PowerEdge M1000e components, such as the LAN on Motherboards (LOMs) and mezzanine cards in the server blades and the I/O modules (IOMs) and midplane in the enclosure, are designed to support three high-speed fabrics per server blade: one standard Gigabit Ethernet (GbE) fabric and two additional customizable fabrics providing sufficient bandwidth to support interconnects such as 10 Gigabit Ethernet (10GbE), Fibre Channel, and InfiniBand. The enclosure uses redundant hot-pluggable components throughout to help maximize system availability.

The PowerEdge M1000e enclosure supports up to 16 half-height server modules, each occupying a slot accessible in the front of the enclosure. This article discusses the PowerEdge M1000e fabric architecture when using these half-height blades. Blades in larger form factors can support a proportionally higher level of fabric connectivity.[1]

MODULAR SERVER I/O

To understand the PowerEdge M1000e architecture, it is necessary to first define four key terms: *fabric*, *lane*, *link*, and *port*.

A *fabric* is defined as a method of encoding, transporting, and synchronizing data between multiple devices. Examples of fabrics are GbE, Fibre Channel, or InfiniBand. Fabrics are carried inside the PowerEdge M1000e system between server modules and IOMs through the midplane. They are also carried to the outside world through the physical copper or optical interfaces on the IOMs.

A *lane* is defined as a single fabric data transport path between I/O end devices. In modern high-speed serial interfaces, each lane comprises one transmit and one receive differential pair. In reality, a single lane is four wires in a cable or traces of copper on a printed circuit board: a transmit positive signal, a transmit negative signal, a receive positive signal, and a receive negative signal. Differential pair signaling provides improved noise margin for these high-speed lanes. Various terminologies are used by fabric standards when referring to lanes. PCI Express (PCIe) calls this a lane, InfiniBand calls it a physical lane, and Fibre Channel and Ethernet call it a link.

[1] For more information on the PowerEdge M1000e, see "The Next-Generation Dell PowerEdge M1000e Modular Blade Enclosure," by Chad Fenner, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080206-Fenner.pdf.

> **"Key PowerEdge M1000e components are designed to support three high-speed fabrics per server blade: one standard Gigabit Ethernet fabric and two additional customizable fabrics."**

A *link* is defined here as a collection of multiple fabric lanes used to form a single communication transport path between I/O end devices. Examples are four-lane (x4), eight-lane (x8), and sixteen-lane (x16) PCIe, or four-lane 10GBase-KX4. PCIe, InfiniBand, and Ethernet call this a link. The differentiation has been made here between *lane* and *link* to help prevent confusion over Ethernet's use of the term *link* for both single- and multiple-lane fabric transports. Some fabrics such as Fibre Channel do not define links, as they simply run multiple lanes as individual transports for increased bandwidth. A link as defined here provides synchronization across the multiple lanes, so they effectively act together as a single transport.

A *port* is defined as the physical I/O end interface of a device to a link. A port can have single or multiple lanes of fabric I/O connected to it.

SERVER BLADE I/O ARCHITECTURE

There are three supported high-speed fabrics per PowerEdge M1000e half-height server module, with two flexible fabrics using optional plug-in mezzanine cards on the server. The server blades used in the PowerEdge M1000e enclosure are designed to support multiple network topologies and provide sufficient bandwidth to accommodate future upgrades. I/O fabric integration encompasses LANs, storage area networks (SANs), and Interprocess Communication (IPC) networks. The blade ports connect through the enclosure midplane to the associated IOMs in the back of the enclosure, which then connect to the LAN, SAN, or IPC network.

As shown in Figure 1, the first embedded high-speed fabric, referred to as fabric A, comprises dual GbE LOMs and their associated enclosure IOMs. The LOMs are based on the Broadcom BCM5708 NetXtreme II Ethernet controller and support TCP/IP Offload Engine (TOE) and Internet SCSI (iSCSI) boot. Although the blades currently have a dual GbE LOM configuration, the enclosure midplane design allows future support for up to four GbE LOMs in each half-height blade. The blade LOMs and Ethernet fabric B and C mezzanine cards also support Wake-on-LAN (WOL). SAN boot is supported by iSCSI- and Fibre Channel–enabled cards.

In addition to fabric A, the blades support additional fabrics B and C by installation of two optional dual-port I/O mezzanine cards in each half-height blade. These cards can currently support a wide array of Ethernet (including iSCSI), Fibre Channel, and InfiniBand technologies and use a common form factor to provide flexibility in fabric configuration. Fabric B and C I/O mezzanine cards have identical mechanical, electrical, and management specifications, providing a high level of flexibility and modularity.

The optional mezzanine cards connect to the blade chipsets through eight-lane (x8) PCIe interfaces, providing up to 16 Gbps of bandwidth per mezzanine card with Gen1 PCIe. Both the PCIe fabrics and external fabrics are routed through high-speed, 10 Gbps–capable air dielectric connector pins through the planar and

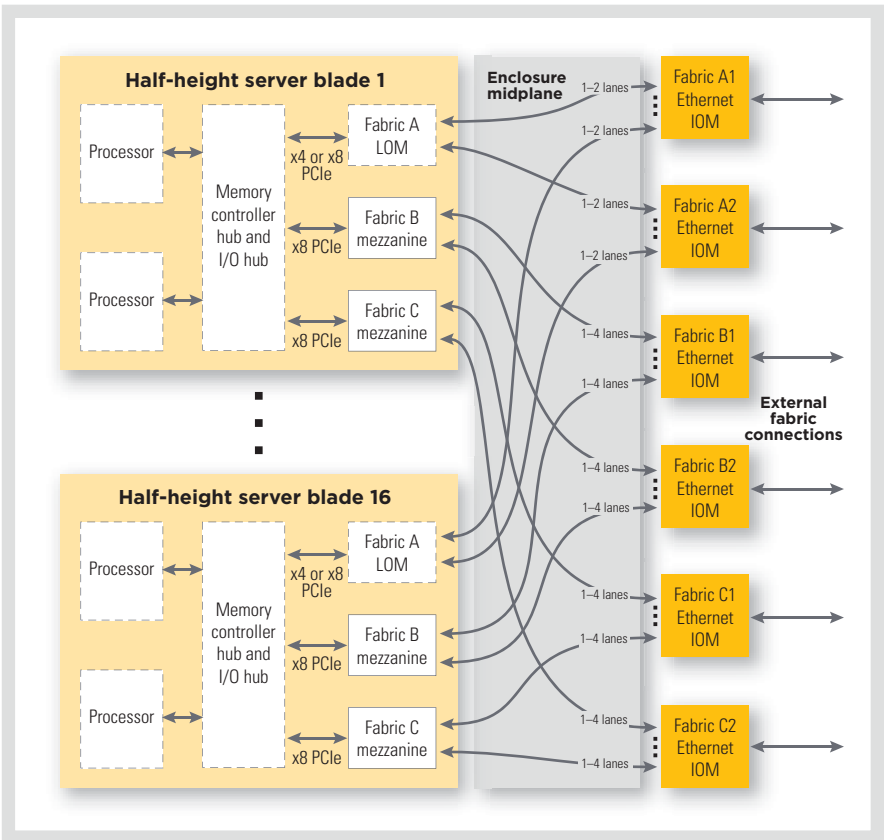

Figure 1. *Dell PowerEdge M1000e high-speed fabric architecture*

midplane. To enhance signal integrity, the signals isolate transmit and receive signals to help minimize crosstalk. Differential pairs are isolated with ground pins, and signal connector columns are staggered—helping minimize signal coupling.

Because of the modular architecture and flexibility to configure mezzanines and IOMs in the system, administrators could inadvertently hot plug a blade with the wrong mezzanine card into the system. To help prevent the accidental activation of mis-configured fabrics on a blade, the PowerEdge M1000e systems management hardware and software include fabric consistency checking. For example, if administrators have configured Fibre Channel IOMs in the fabric C slots, then all blades must have either a Fibre Channel mezzanine card or no mezzanine card in that fabric slot. If, in this case, they hot plugged a blade into the enclosure that had a GbE mezzanine card in its fabric C slot, the system would automatically detect this misconfiguration and alert administrators of the error so they could appropriately reconfigure the blade.

ENCLOSURE I/O ARCHITECTURE

A key capability introduced with the PowerEdge M1000e is comprehensive 10/100/1,000 Mbps Ethernet support when using Ethernet pass-through modules. In the past, pass-through connections were limited to 1,000 Mbps or GbE speeds. The PowerEdge M1000e enables organizations to connect to legacy 10/100 Mbps infrastructures using Ethernet pass-through or switch technology. This feature uses in-band signaling on a 1000Base-KX transport and does not require administrator interaction to function.

The enclosure midplane can support up to four GbE links per blade for fabric A, providing up to 4 Gbps of bandwidth per half-height blade. Fabrics B and C are routed as two sets of four lanes from mezzanine cards on the blades to the IOMs in the back of the enclosure. Supported bandwidth ranges from 1 to 10 Gbps per lane depending on the fabric, or up to 80 Gbps per mezzanine card.

Because each mezzanine card connects to the blade chipset through an eight-lane PCIe link, the system has no throttle points constricting I/O bandwidth. Dell anticipates that when multi-lane 10GBase-KR technology is available, blades will have moved to PCIe 2.0 or better, providing full end-to-end I/O bandwidth from the server blades to the enclosure IOMs.

Midplane

The midplane—a large printed circuit board providing power distribution, fabric connectivity, and systems management infrastructure—is the focal point for all connectivity within the PowerEdge M1000e enclosure, and is designed to provide scalable bandwidth for both current and future generations of servers and infrastructure. As is required for fault-tolerant systems, the PowerEdge M1000e midplane is passive, with no hidden stacking midplanes or interposers with active components. The I/O fabrics and systems management infrastructure are designed to be fully redundant for each hot-pluggable component.

I/O fabrics are routed through 10 Gbps–capable high-speed connectors and dielectric material. The I/O channels have been simulated to 10GBase-KR channel models. Following industry standards, the fabrics internally support a bit error rate of 10^{-12} or better. Dell has made a high level of investment to help ensure scalable bandwidth for current and future generations of servers and infrastructure.

I/O modules

The back of the PowerEdge M1000e enclosure contains systems management, cooling, power, and I/O components. The IOMs are used as pairs, with two fully redundant modules for each blade fabric, and can be pass-through or switch modules. Pass-through modules provide direct one-to-one connectivity from each LOM or mezzanine card port on each blade to the external network. Switches provide an efficient way to consolidate links from the LOM or mezzanine cards

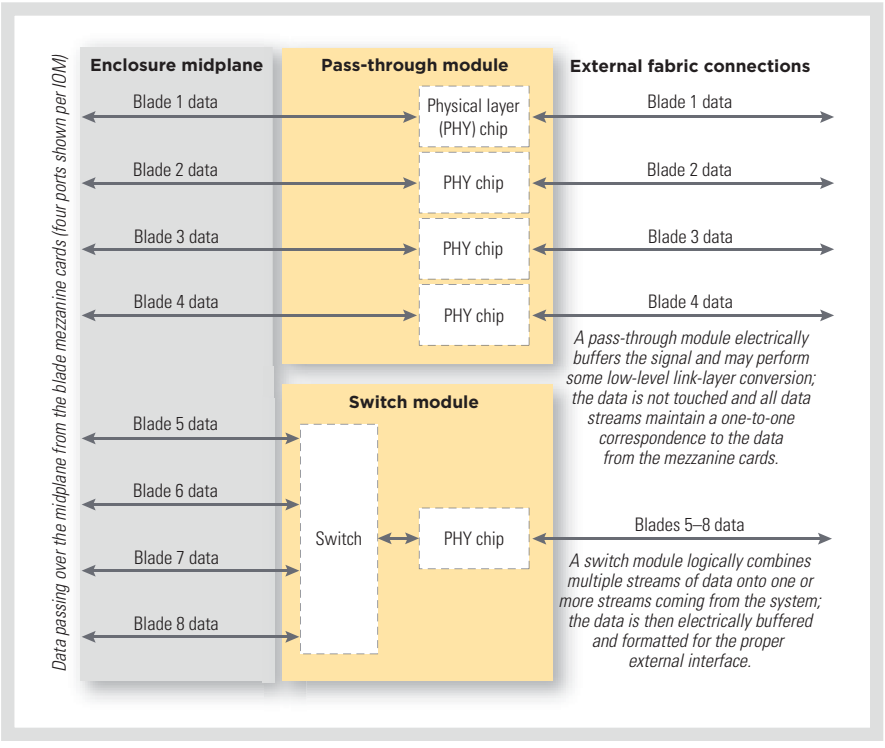

Figure 2. *Pass-through and switch modules as part of the Dell PowerEdge M1000e architecture*

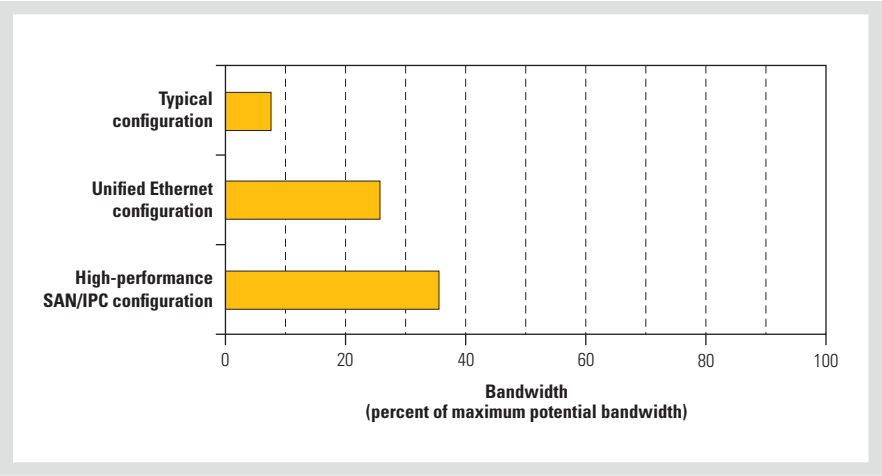

Figure 3. *Bandwidth requirements for three possible Dell PowerEdge M1000e configurations compared with total potential system bandwidth*

on the blades to network uplinks. IOMs are designed to be fully compatible across slots. Figure 2 illustrates these different components as part of the PowerEdge M1000e architecture.

The Dell PowerConnect™ M6220 Ethernet switch and Cisco Catalyst blade switch (3032, 3130S-G, and 3130S-X) for the PowerEdge M1000e enclosure demonstrate the advanced modularity of the IOMs. These Ethernet switch modules introduce sub-module I/O expansion to help maximize flexibility and extension. They provide a single hardware design that allows scaling from a cost-effective GbE configuration to one that utilizes switch stacking to interconnect multiple switches within or between enclosures or to one that adds support for 10GbE uplinks to the core network with both copper and optical interfaces.

Scalable system bandwidth

Assuming a full population of GbE and 10GbE lanes in all three fabrics, the PowerEdge M1000e architecture can support a total potential bandwidth of 5.44 Tbps. As shown in Figure 3, a typical configuration of four lanes of GbE and two lanes of Fibre Channel per server blade uses less than 10 percent of this total potential bandwidth. A unified Ethernet configuration—one that aggregates all network, storage, and IPC-over-Ethernet links by using two lanes of GbE and four lanes of 10GbE per blade—uses only about 25 percent of the total potential bandwidth. Even a high-performance SAN/IPC configuration using two lanes of GbE, two lanes of 8 Gbps Fibre Channel, and two lanes of QDR InfiniBand still uses only about 35 percent of the total potential bandwidth.

The PowerEdge M1000e is also designed to provide comprehensive support for near-, medium-, and long-term I/O infrastructure requirements in a flexible, cost-effective way. An example of this flexibility is the routing of dual two-lane paths for fabric A and dual four-lane paths for fabrics B and C. In the near term, 10GBase-KX4 routing supports 10GbE connectivity. 10GBase-KX4 routing uses all 4 lanes, each running at a 2.5 Gbps data rate to achieve total link data bandwidth of 10 Gbps. Today and in the near-term future, 10GBase-KX4 is the most cost-effective, ubiquitous solution for 10GbE fabrics over midplanes. A fully configured PowerEdge M1000e system supports two dual GbE fabrics and two dual 10GbE fabrics with 10GBase-KX4 routing, supporting the leading edge of unified network topologies. Such a configuration could, for instance, provide redundant 10GbE links per blade for traditional network traffic, and another set of redundant 10GbE links for iSCSI or Fibre Channel Over Ethernet networked storage, and still maintain dual two-lane GbE links for systems management or other low-bandwidth requirements.

MODULAR BLADE SERVER ARCHITECTURE

The architecture of the Dell PowerEdge M1000e modular blade server enclosure and its 10th-generation Dell server technology has been designed specifically with modularity in mind, providing customizable multi-lane fabrics that can support both current and future network technologies and help protect enterprise blade server investments. As part of the Dell focus on simplifying IT, there is no confusing support matrix to follow for PowerEdge M1000e mezzanines or IOMs and no multiplication of mezzanine and IOM form factors and design standards, helping avoid potential compatibility and configuration difficulties. The PowerEdge M1000e supports one mezzanine design standard and one IOM design standard to enable true modularity at the system and subsystem level, helping simplify extension and enhancement now and in the future.

John Loffink is an engineer/strategist in the Dell Server Advanced Engineering Group. He has over 20 years' background in servers, enterprise storage, hardware design, and high-availability computing. John holds a B.S. in Electrical Engineering from the Florida Institute of Technology.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge M1000e:
DELL.COM/Servers

MANAGING DELL POWEREDGE M1000e BLADE SERVERS WITH THE AVOCENT iKVM SWITCH

By Stephen M. Hahn
Chad Fenner

The Avocent® iKVM (integrated keyboard, video, mouse) switch, an analog KVM switch integrated into the Dell™ PowerEdge™ M1000e modular blade enclosure, is designed to simplify administrative tasks by letting IT staff easily view, monitor, and manage server blades and by providing access to the remote management and power control functions of the PowerEdge M1000e Chassis Management Controller.

Related Categories:

Avocent

Blade servers

Visit DELL.COM/PowerSolutions for the complete category index.

In data center environments, simple, flexible server monitoring and management tools can be critical to maintaining efficient operations. The Avocent iKVM (integrated keyboard, video, mouse) switch, an analog KVM switch integrated into the Dell PowerEdge M1000e modular blade enclosure,[1] is designed to let administrators easily view, monitor, and manage each of the 16 server blades in the PowerEdge M1000e and provide access to the Dell Chassis Management Controller (CMC) for management and power control. Key features of the Avocent iKVM switch include local access to each server blade, one-to-one administration for up to 16 administrators accessing each of the 16 blades, integration with the Avocent On-Screen Configuration and Reporting (OSCAR®) graphical user interface for configuration and to work with existing Avocent KVM infrastructures, integration with the Microsoft® Active Directory® directory service, and access to the integrated Dell Remote Access Controller (iDRAC) on each blade.

The Avocent iKVM switch is compatible with standard USB keyboards and pointing devices as well as with VGA monitors that support Display Data Channel (DDC). Its video connections support display resolution ranges from 640 × 480 at 60 Hz up to 1,600 × 1,200 at 75 Hz. When necessary, administrators can update the Avocent iKVM switch firmware through the CMC firmware update utility.

USING THE AVOCENT iKVM SWITCH

The Avocent iKVM switch provides traditional KVM switching between server blades in the Dell PowerEdge M1000e enclosure as well as a scanning feature, which allows administrators to scan through preselected blades while pausing for a specific time at each blade. They can also select a specific blade through the OSCAR main dialog box, as shown in Figure 1. In addition to the blade names (which administrators can set themselves or pull from the CMC automatically during setup) and chassis slot numbers, the OSCAR interface provides status information for each blade: a green circle, for example, indicates that the blade is online and functioning properly, while a red *X* means the blade is offline or not functioning properly. A yellow circle indicates that a blade is online but unavailable—for example, because another user is controlling it remotely. The letters in the right column provide additional information, such as whether a blade is being accessed by a user panel (a green letter) or blocked by a specific user channel (a black letter), and whether the blade is being accessed through the rear panel (the letter *A*) or the front panel (the letter *B*). When administrators

[1] For more information on the PowerEdge M1000e, see "The Next-Generation Dell PowerEdge M1000e Modular Blade Enclosure," by Chad Fenner, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080206-Fenner.pdf.

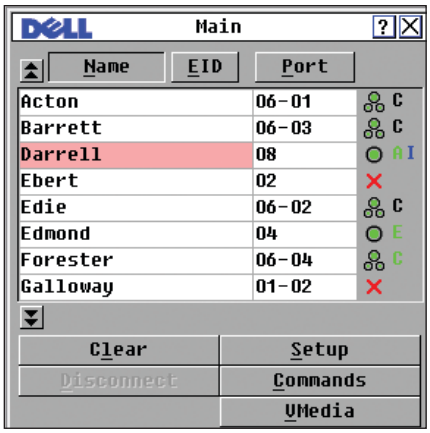

Figure 1. *Avocent OSCAR main dialog box for the Dell PowerEdge M1000e blade server*

select a blade to view and control, the Avocent iKVM switch can automatically reconfigure the keyboard and mouse to the proper settings for that blade.

Administrators also have the option of using hotkey sequences to access server blades without using the OSCAR main dialog box. For example, they can configure the switch so that rather than pressing the Print Screen button twice to bring up the main dialog box and then selecting a blade, they can press the Print Screen button and then the first few characters of a blade name or slot number to display that blade immediately.

The Avocent iKVM switch also enables administrators to send keyboard and mouse actions to multiple server blades simultaneously, helping simplify the process of managing multiple blades at once. They can configure the switch to broadcast keystrokes and mouse movements independently as well, providing additional management flexibility.

The switch provides basic security capabilities to help prevent unauthorized users from controlling the server blades. For example, the OSCAR interface allows administrators to protect the system with a screen saver password: after a defined time period, the screen saver engages and access is prohibited until the password is entered. Administrators can also protect the OSCAR configuration itself with a password. If they ever lose or forget this password, they can clear it through the PowerEdge M1000e CMC.

ACCESSING THE CHASSIS MANAGEMENT CONTROLLER

The CMC is a hot-pluggable systems management hardware and software controller designed to provide remote management and power control for the Dell PowerEdge M1000e blade server. It includes its own processor and memory, and draws power from the modular blade enclosure. Its key features include the following:

- **Support for Microsoft Active Directory authentication:** Directory services such as Microsoft Active Directory maintain a common database of information needed to control network users and assets; organizations using Active Directory can use it to provide access to CMCs. Administrators can centralize CMC user IDs and passwords in Active Directory using the standard schema or an extended schema.
- **Comprehensive monitoring:** CMCs provide access to system information and the status of components such as server blades, power supplies, fans, and temperature sensors.
- **Access to system event logs:** CMCs generate a hardware log of events that occur on the chassis, which includes the severity, time, and a description of the event. Administrators can view, save a text file version of, and clear the hardware log from the CMC interface.
- **Automated alerts:** Administrators can configure CMCs to send e-mail messages or Simple Network Management Protocol (SNMP) traps to alert them of warnings or errors related to temperatures, hardware mis-configurations, power outages, fan speeds, and so on.

In addition to other ways of accessing CMCs, administrators can use the Avocent iKVM switch to establish racadm console connections. They can then use racadm commands to configure CMC properties and perform remote management tasks through a command-line interface. Administrators can display a list of available racadm commands, syntax information, and other instructions by using the command `racadm help`. They can list syntax and options for specific racadm commands by using the command `racadm help` *`command`*. The Dell CMC user guide also provides a list of available racadm commands.

PROVIDING SIMPLE, FLEXIBLE BLADE SERVER MANAGEMENT

The Avocent iKVM switch is designed to provide easy access to each of the 16 server blades and the CMC in Dell PowerEdge M1000e modular blade enclosures. Using the switch with the Avocent OSCAR interface and taking advantage of the racadm console connections can give administrators comprehensive access to and control over the Dell PowerEdge M1000e blade server, helping provide simple, flexible monitoring and management capabilities.

Stephen M. Hahn has been the Dell account manager at Avocent for the last eight years. He has a bachelor's degree from Boise State University.

Chad Fenner is the product marketing manager for blade servers at Dell. He has a bachelor's degree from Trinity University in San Antonio, Texas.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Avocent iKVM switch:
www.avocent.com

Dell PowerEdge M1000e:
DELL.COM/Servers

SIMPLIFIED MANAGEMENT WITH ALTIRIS DEPLOYMENT SOLUTION FOR DELL SERVERS 3.0

By Eric Szewczyk

Altiris® Deployment Solution™ for Dell Servers software offers simple drag-and-drop bare-metal provisioning and ongoing server management of Dell™ PowerEdge™ servers. Version 3.0 of this software introduces multiple features and enhancements designed to further simplify key administrative tasks, and incorporates support for the Dell PowerEdge M1000e modular blade enclosure and PowerEdge M600 series server blades.

Related Categories:

Altiris

Blade servers

Dell PowerEdge blade servers

System deployment

Systems management

Visit DELL.COM/PowerSolutions for the complete category index.

In the past, Altiris Deployment Solution has been available primarily as a download from the Altiris Web site. However, Dell has also begun offering this software on an optional DVD with Dell PowerEdge blade servers. This DVD includes the standard Altiris Deployment Solution software and the Altiris Deployment Solution for Dell Servers add-on for server management, as well as product documentation and a quick-start video to help administrators carry out the initial installation and configuration steps. The Altiris Deployment Solution for Dell Servers add-on is designed to integrate with the Dell OpenManage™ Deployment Toolkit (DTK) to facilitate bare-metal provisioning of Dell hardware, including automated configuration and update deployment for the BIOS, baseboard management controller (BMC), Dell Remote Access Controller (DRAC), and RAID components. The DVD sleeve includes a registration code that administrators can use at the licensing portal on the Altiris Web site to retrieve their software licenses.

Altiris Deployment Solution for Dell Servers 3.0 introduces multiple features and enhancements over previous versions designed to further simplify key administrative tasks. It also incorporates support for the Dell PowerEdge M1000e modular blade enclosure and PowerEdge M600 series server blades, including planned enhancements to its rip-and-replace feature designed to work with updates to the PowerEdge M1000e enclosure.[1]

UPDATED HARDWARE AND ENVIRONMENT SUPPORT

Altiris Deployment Solution for Dell Servers 3.0 with Service Pack 1 (SP1) is bundled with version 2.4 of the DTK, and supports 6th- through Dell servers—including the Dell PowerEdge M1000e enclosure and PowerEdge M600 series blades—and both Microsoft® Windows® Preinstallation Environment (WinPE) and Linux® automation environments.

This release no longer supports the Microsoft MS-DOS® OS as an automation environment for bare-metal provisioning. Administrators can still use MS-DOS for simple capturing and distribution of image-based OS deployments and executing of other pre-OS scripts; however, an MS-DOS preboot environment typically cannot match the speed of a true 32-bit environment. Linux is typically the preferred replacement for MS-DOS, primarily because it is open source; offers more drivers than MS-DOS; supports x86, x86-64, and Intel® IA-32 platforms; and helps accelerate boot-from-network times and server deployments by minimizing overhead, with a Preboot

[1] For more information on the PowerEdge M1000e, see "The Next-Generation Dell PowerEdge M1000e Modular Blade Enclosure," by Chad Fenner, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080206-Fenner.pdf.

Execution Environment (PXE) image file typically only 15 MB in size, compared with the 160 MB typical of WinPE images.

EASY CONFIGURATION WIZARD

Altiris Deployment Solution for Dell Servers 3.0 introduces a configuration wizard designed to simplify scripted OS installations, help administrators install WinPE or Linux preboot environments if they have not already been created, and create independent PXE images optimized for bare-metal deployments (see Figure 1). To help simplify the process, the software only creates bare-metal OS deployment jobs for the OS media support administrators provide when they first install the software. For example, if administrators only add support for the Microsoft Windows Server® 2003 Release 2 (R2) Enterprise Edition OS, only jobs for deploying that OS on bare-metal servers appear in the Altiris Deployment Solution console. The wizard supports a wide array of server platforms, including the Microsoft Windows Server, Red Hat® Enterprise Linux, and Novell® SUSE® Linux Enterprise Server operating systems.

Administrators can also use multiple instances of a single OS for different builds without needing to duplicate the source files. For example, if administrators must create a variation of their standard Windows Server 2003 build, or need to specify a unique Windows product key during their next deployment, they can simply point to the previously copied source directory and specify a custom answer file with the desired changes. If they use images for OS deployment, they can enable this functionality by combining the Job Builder feature (see the "Customizable bare-metal OS deployments" section in this article) and the basic image capturing and distribution tasks.

INCREASED INTEROPERABILITY THROUGH SEPARATE PXE IMAGE FILES

Administrators can now separate PXE images configured and tested to work with Dell hardware from the default PXE image files created when they install the Altiris Deployment Solution for Dell Servers add-on (see Figure 2). A major advantage of this option is increased interoperability with other Altiris Deployment Solution add-ons, enabling administrators to isolate separate drivers and configurations for hardware from different vendors.

CUSTOMIZABLE BARE-METAL OS DEPLOYMENTS

Altiris Deployment Solution for Dell Servers 3.0 introduces the Job Builder feature, which is designed to help administrators customize bare-metal OS deployment jobs by allowing them to create new jobs based on specific requirements (see Figure 3). For example, when using one of the predefined jobs automatically created based on the OS source files copied during installation, administrators may want to set up the RAID configuration based on specific rules or on a captured RAID configuration file rather than using the default dynamic configuration approach. Selecting the Configure RAID check box enables them to create the individual tasks that make up this RAID provisioning job. This job then appears in the Job Builder folder, and administrators can highlight the job, copy its individual tasks, and insert them in place of the existing dynamic RAID provisioning tasks contained in the original predefined job. (For more information on rule-based RAID provisioning, see the "Rule-based RAID provisioning" section in this article.)

The Job Builder feature is also useful if administrators do not copy OS source files when installing Altiris Deployment Solution for Dell Servers, and prefer to use a single image for Dell servers. They can select the check boxes for specific configuration tasks, then add a Distribute Disk Image task using the desired image.[2]

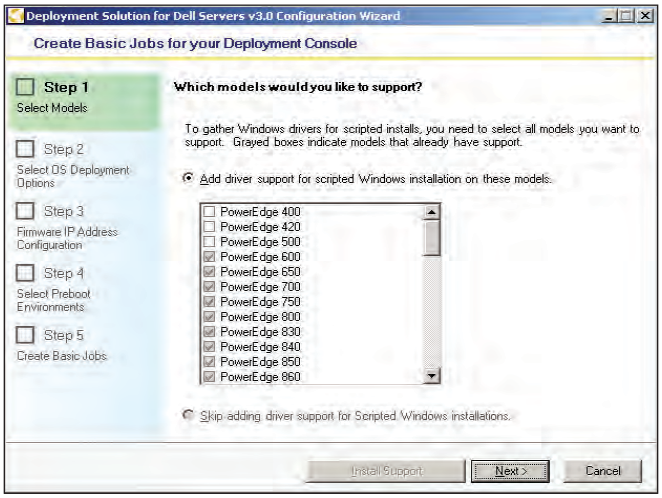

Figure 1. *Altiris Deployment Solution for Dell Servers Configuration wizard*

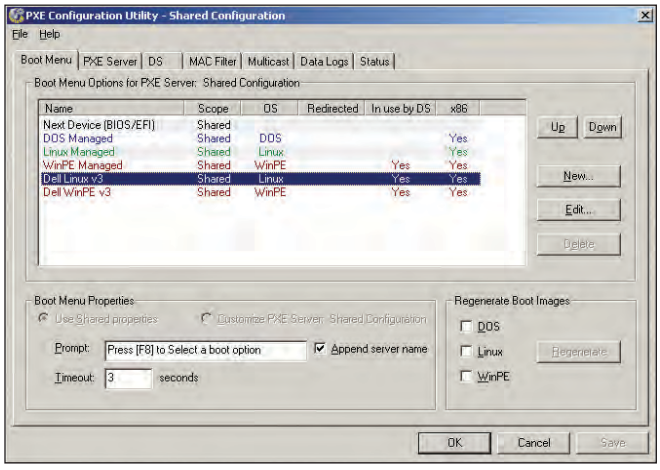

Figure 2. *Altiris Deployment Solution for Dell Servers PXE Configuration Utility showing separate images for Dell servers*

[2] For more information on using hardware-independent images with Dell servers, see "Building a Single Image for Dell PowerEdge Servers Using Altiris Deployment Solution," by Eric Szewczyk, in *Dell Power Solutions*, February 2007, DELL.COM/Downloads/Global/Power/ps1q07-20070193-Altiris.pdf.

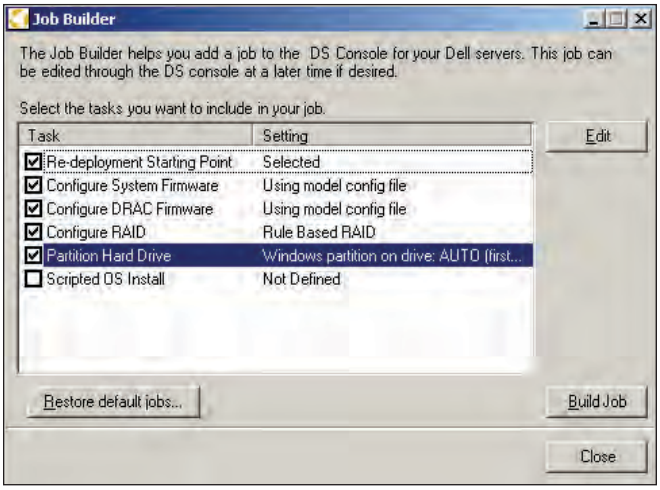

Figure 3. *Altiris Deployment Solution for Dell Servers Job Builder window*

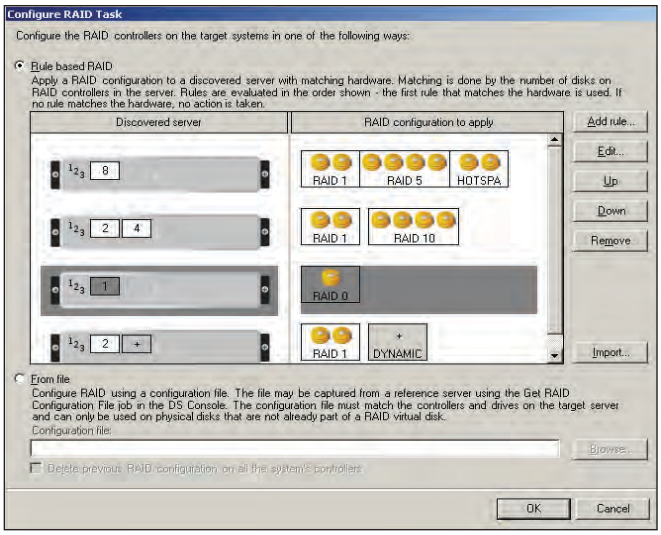

Figure 4. *Altiris Deployment Solution for Dell Servers Configure RAID Task window for rule-based RAID provisioning*

RULE-BASED RAID PROVISIONING

The Dell IT group uses Altiris Deployment Solution for Dell Servers to deploy and manage thousands of servers in its environment, and this group had requested the ability to provision RAID arrays based on a series of rules or from a captured configuration file. Altiris responded by integrating rule-based RAID provisioning into Altiris Deployment Solution for Dell Servers 3.0 (see Figure 4).

Dynamic RAID provisioning is still suitable for server builds where RAID arrays are created based on servers with a single RAID controller and varying numbers of drives. For example, this approach would use RAID-0 for non-blade Dell servers with one drive, RAID-1 for Dell servers with two drives, and RAID-5 for Dell servers with three or more drives. For complex environments in which servers may have more than one RAID controller and administrators need to use other supported RAID configurations such as RAID-10 or RAID-50, or in which they want to create global or dedicated hot spares, rule-based RAID provisioning can significantly simplify the configuration process.

PRE-OS HARDWARE UPDATES

Traditionally, Dell Update Packages have been designed to execute in post-OS Microsoft Windows or Linux environments and apply updates when the server is rebooted. Previous versions of Altiris Deployment Solution for Dell Servers have included predefined jobs enabling administrators to remotely execute these updates by dragging and dropping update jobs to the managed server from within the Altiris Deployment Solution console. In previous releases, administrators could also update the BIOS in a pre-OS environment by extracting an .hdr file from the Dell Update Package and applying it in an MS-DOS preboot environment. By enabling administrators to easily update the BIOS to its most recent version before deploying the OS and placing the server into a production environment, this approach helped reduce the risk of errors or compatibility problems during deployment. However, the DTK no longer supports the MS-DOS preboot automation environment.

Altiris Deployment Solution for Dell Servers 3.0 with SP1 introduces the ability to use Dell Update Packages to update server firmware in Linux preboot automation environments, helping administrators ensure that firmware for BIOS, BMC, DRAC, and RAID components are updated before placing servers in production environments and helping minimize risks to the stability of production systems. Administrators must still apply updates for network interface cards (NICs) and RAID controllers in a post-OS environment, however. To help ensure that Dell firmware and drivers stay up-to-date in production environments, administrators can also take advantage of Altiris Patch Management Solution™ for Dell Servers software, which scans servers for compliance and offers a simple policy-based distribution method to help keep these components updated.[3]

SIMPLIFIED DELL DIAGNOSTIC PARTITION INSTALLATION

In the past, installing the 32 MB Dell diagnostic partition as part of the provisioning process required copying the appropriate upimg.bin file from the Dell OpenManage Server Assistant CD for each Dell server being deployed. Now, however, Dell servers use one common file included by default as part of the installation. Altiris Deployment Solution for Dell Servers 3.0 takes advantage of this change to help reduce the number of onetime post-installation configuration steps required for simple drag-and-drop deployment.

FLEXIBLE OUT-OF-BAND POWER CONTROL

Although past versions of Altiris Deployment Solution for Dell Servers have included power-control features through

[3] For more information on Altiris Patch Management Solution for Dell Servers, see www.altiris.com/products/patchmanagementsolutionfordell.aspx.

Symantec Yellow Books™

Symantec Enterprise Messaging Management for Microsoft® Exchange

A comprehensive approach to effectively managing messaging environments

- Overview of e-mail and instant messaging (IM) security, availability, and resilience concepts
- Best practices for implementing Symantec e-mail and IM management solutions
- Technical information regarding multiple product deployment, configuration sequences, and achieving synergies
- Performance testing powered by Dell™

Available now...

The New Symantec Yellow Book™

Learn how Dell and Symantec are helping organizations protect the integrity and confidentiality of information transmitted in Exchange, while reducing the costs of storage and electronic discovery at the same time.

Highlights of the new Yellow Book:

- Brand new Exchange 2007 performance testing results
- Symantec Mail Security for Exchange 2007
- Symantec Mail Security 8300 gateway appliance, now with Premium Content Control
- Symantec Enterprise Vault 2007
- Performance impacts of SMS and Enterprise Vault in environments of varying sizes

Retail value: $65.00

Obtain a copy by electronic download.
E-mail: mickey_coombs@symantec.com and
kara_bankhead@symantec.com
Subject of e-mail: Dell YB – Protecting Exchange 2007

the DRAC or the Intelligent Platform Management Interface (IPMI)–based BMC, Altiris Deployment Solution for Dell Servers 3.0 enhances the discovery process when searching for these devices, enabling administrators to use these features by right-clicking on the servers in the Altiris Deployment Solution console (see Figure 5). The top half of the right-click menu provides in-band power control through the Altiris Deployment Solution AClient agent, while the bottom half provides out-of-band power control through the DRAC interfaces or IPMI.

ENHANCED BOOT DISK CREATOR FOR DELL HARDWARE

With previous versions of Altiris Deployment Solution for Dell Servers, administrators that did not use PXE in their environment were not able to easily add Dell driver and configuration support to automation boot disks such as CD or USB boot devices. Administrators can now create new configurations from the Boot Disk Creator and easily add Dell drivers and configurations.

SUPPORT FOR DELL POWEREDGE M1000e BLADE SERVER

Altiris Deployment Solution for Dell Servers 3.0 with SP1 supports the 16-blade Dell PowerEdge M1000e enclosure and its PowerEdge M600 series blades, and enables administrators to pre-provision blades by creating virtual bays. Administrators can assign unique characteristics to these bays, including service tags, asset tags, and Media Access Control (MAC) addresses. When a blade PXE boots and reads the primary key lookup, Altiris Deployment Server executes the jobs that have been assigned to that particular bay.

In addition, administrators can pre-provision systems using the ImportComputers55.xls spreadsheet provided with Altiris Deployment Solution. This spreadsheet, located in the .\Samples folder of the Altiris eXpress share, provides a powerful way to assign system

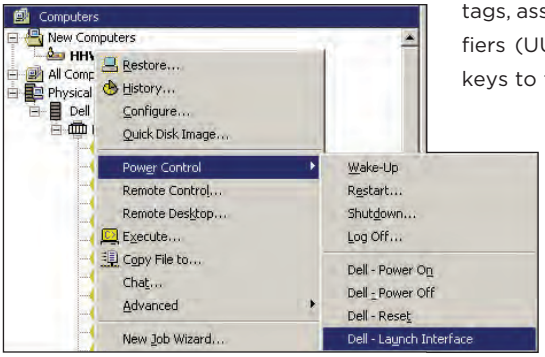

Figure 5. *Altiris Deployment Solution for Dell Servers power control options*

roles and configurations before the systems connect to the network. Once these systems have been pre-provisioned in the Altiris Deployment Solution console, adding servers can be as simple as plugging them in, with Altiris Deployment Solution automatically routing the appropriate builds to each server the first time they are booted on the network.

UPDATED RIP-AND-REPLACE FUNCTIONALITY

Traditionally, blade server rip-and-replace support in Altiris Deployment Solution for Dell Servers has been based on the MAC address of the primary NIC for specific blades. This MAC address is stored in the cache of the Altiris PXE server. When administrators remove a blade from its chassis slot, insert a different blade in that slot, and turn the blade on, the Altiris PXE server first determines that it does not recognize the blade's MAC address and then executes one of four administrator-defined blade change rules.

In the future, Dell plans to add a management feature to the PowerEdge M1000e enclosure based on tying persistent World Wide Names (WWNs) and MAC addresses to individual chassis slots. Once this feature has been added, Altiris Deployment Solution for Dell Servers will no longer be able to use MAC addresses to trigger blade change rules. To accommodate this change, Altiris plans to update Altiris Deployment Solution to use service tags, asset tags, universally unique identifiers (UUIDs), or a combination of these keys to trigger blade change rules.[4]

SIMPLIFIED DEPLOYMENT FOR DELL SERVERS

Altiris Deployment Solution for Dell Servers 3.0, now included with Dell PowerEdge blade servers, introduces multiple features and enhancements over previous versions, including updated hardware and environment support, an easy configuration wizard, features designed to increase administrative flexibility, and support for the Dell PowerEdge M1000e modular blade enclosure and PowerEdge M600 series server blades. Administrators can take advantage of these features to help simplify server deployment in Dell hardware–based data center environments.

Eric Szewczyk is a Dell Alliance technical strategist for Altiris, now part of Symantec. He has a B.A. from the University of Central Oklahoma and is an Altiris Certified Engineer (ACE).

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and Altiris:
DELL.COM/Altiris
www.altiris.com/dell

Dell and Altiris on Juice:
juice.altiris.com/dell

Altiris Deployment Solution for Dell Servers:
www.altiris.com/products/delldeploymentsolution.aspx

Altiris Deployment Solution for Dell Servers resources:
www.altiris.com/delldeploy

[4] For more information on blade server rip-and-replace support in Altiris Deployment Solution, see "Simplifying Blade Server Management with Altiris Deployment Solution Rip and Replace," by Eric Szewczyk and Chad Fenner, in *Dell Power Solutions*, February 2007, DELL.COM/Downloads/Global/Power/ps1q07-20070274-Altiris.pdf.

CREATING A HYPER-EFFICIENT HYPER-SCALE DATA CENTER

By Jimmy Pike
Ty Schmitt
Frank Frankovsky
Todd Brannon

A growing number of organizations today utilize computing "clouds" to deliver many applications familiar to Internet users as well as flexible access to powerful compute and storage resources. Integrated solutions designed to minimize acquisition and operating costs, maximize energy efficiency, and enable rapid scalability can help operators of these and other hyper-scale data centers speed deployment, lower total cost of ownership, and reduce environmental impact.

Related Categories:

Cloud computing

Visit DELL.COM/PowerSolutions for the complete category index.

As computing architectures continue to evolve, server-centric "cloud" and large cluster environments are increasingly being utilized to deliver common applications as well as flexible access to powerful compute and storage resources. A cloud computing environment is one that is designed to cluster and/or load balance computation across thousands, or even tens of thousands, of compute nodes. Users typically connect over the Internet to tap into these resources "in the cloud." Internet search, free e-mail, online media, and multiplayer gaming are representative of this emerging computing paradigm. Cloud computing platforms are also increasingly being deployed by organizations that offer computing resources as well as software and storage hosting services over the Internet. In parallel with this trend, traditional clusters employed across a wide range of industries to perform compute-intensive tasks such as financial risk analysis, geophysical modeling, electronic design analysis, graphics rendering, and medical research continue to grow ever larger.

Cloud and large cluster computing environments have distinctly different design and management requirements from traditional general-purpose environments. For example, in the massively scaled-out architecture of a hyper-scale data center, any inefficiency—such as excess power consumption or poor space utilization—can significantly drive up the cost of computation. As a result, maximizing the efficiency of a compute cloud at the node, rack, and facility level is a key source of business value and an essential design consideration.

Cloud and large cluster computing environments also have different hardware requirements from traditional computing environments. For example, in a hyper-scale data center, it is usually the *application* that provides availability and resiliency, not redundant hardware within the individual nodes. As a result, server-level availability is not required—which significantly changes the configuration and characteristics of individual compute nodes.

UNDERSTANDING HYPER-SCALE DATA CENTER REQUIREMENTS

In a hyper-scale computing environment, efficiency is a key driver of enterprise success. But unlike traditional computing environments, hyper-scale data centers generally have thousands, or even tens of thousands, of compute nodes. In this type of massively scaled-out architecture, even very small inefficiencies can significantly drive up the cost of computation. Also, because the workload requirements of individual nodes in a hyper-scale environment are generally less stringent

than those of nodes in a general-purpose environment, over-provisioning can occur and significantly drive up the cost of computation.

For example, most general-purpose servers and storage come equipped with standard features such as redundant components for availability and reliability that are not needed in a cloud environment. As a result, organizations often pay a "feature tax" on unneeded capabilities. In addition, in most general-purpose servers, the power supply and cooling components are designed to accommodate the power and cooling needs of the maximum configuration of components available on that server. As a result, general-purpose servers often consume more electricity and generate more heat than necessary to support the workload and configuration required of a node in a specific cloud or cluster environment.

Furthermore, because a hyper-scale environment utilizes thousands of servers, the excess heat generated by unnecessary features and over-provisioned power and cooling components typically found in general-purpose servers requires additional cooling that can significantly increase facility-level energy costs. Excess heat and bulk also limit node density, which can increase the amount of space required to deploy a server farm. In the multi-thousand-node environment of a fast-growing compute platform, such space requirements can significantly add to facility costs and slow project schedules.

Data center efficiency considerations are not just limited to individual nodes, however. The design and layout of a data center at the *facility* level is equally important. For example, the layout of racks, power, networking, and heating, ventilation, and air-conditioning (HVAC) systems within a facility is a critical factor contributing to the cost of computation. A facility that is not carefully designed to minimize heat and optimize air flow can generate tremendous excess energy and space costs. It is essential that great care be taken when designing a hyper-scale data center to minimize inefficiency at the node, rack, and facility level.

CO-DEVELOPING DESIGNED-TO-ORDER COMPUTING SOLUTIONS

To help organizations maximize the efficiency and effectiveness of their hyper-scale data centers, Dell recently created the Data Center Solutions (DCS) Division. The DCS team comprises solution consulting, engineering, supply chain, and project management experts with in-depth knowledge regarding the needs of hyper-scale data centers and cloud computing environments. The DCS Division offers a wide range of services, including data center design consulting, designed-to-order hardware, and custom service offerings for rapid deployment and system maintenance.

By employing a close co-development process with organizations, the DCS team can help improve the efficiency of an existing computing platform or assist in designing a new hyper-scale data center from the ground up. The DCS team is organized and scoped to support the needs of these specific computing environments, as shown in Figure 1. By building integrated solutions designed to minimize acquisition and operating costs, maximize energy efficiency, and enable rapid scalability, DCS can help organizations speed deployment, lower total cost of ownership (TCO), and reduce the environmental impact of their compute clouds.

To help organizations maximize the computing power of a hyper-scale data center while minimizing overall costs, the Dell DCS Division offers data center design consulting. In data center design consulting, a DCS team works closely with an organization to help understand its current and future computing requirements, analyze and assess the efficiency and effectiveness of the data center currently in place, and co-design a comprehensive data center solution.

For example, in the assessment phase, the DCS team evaluates the features and specifications of the compute nodes currently in place to help identify areas of inefficiency, such as unneeded features or excess power consumption. The team then evaluates the layout of the facility—including the placement of nodes, racks, networks, and power and cooling systems—to help identify inefficiency in data center layout. The team also applies advanced computational fluid dynamics (CFD) analysis to assess the effectiveness and efficiency of the power and cooling systems.

Following the assessment phase, the DCS team develops a comprehensive data center solution designed to optimize the

	Solutions for typical data center environments	Solutions optimized for cloud computing environments
Deployment size	Any	Typically 1,500 nodes or more
Application	Any	Homogeneous cloud applications, not requiring hardware high availability at the node level
Validation	Broad enterprise interoperability and deep product testing	Point solution only
Support	Comprehensive enterprise hardware, OS, applications, and solutions	Engineer-to-engineer and point solutions only
Service	Numerous packaged options	Site-specific service plan, with parts on-site
Deployment span	Multiple locations and/or staggered over a long time	Limited number of deployment locations with large-scale builds
Design	General-purpose or optimized for widely used solutions	Specific to each organization and optimized for clouds

Figure 1. *Cloud computing platforms are optimized for tightly scoped requirements*

efficiency and effectiveness of the computing environment at all levels of scale. At the node level, for example, the DCS team designs custom-built hardware tailored to meet an organization's specific computing needs, to avoid a "feature tax" for unneeded capabilities. DCS designers also use best thermal practices and state-of-the-art technology to help ensure that custom-designed hardware is as energy-efficient as possible. In particular, DCS designers take advantage of advanced technologies such as the following:

- **High-performance, energy-efficient components:** To help optimize computing power and efficiency in processors, memory, and hard disk drives
- **Right-sized power supplies and fans:** To help meet the specific requirements of each compute node's feature configuration
- **Low-flow fan algorithms:** To help optimize both fan efficiency at the server level and return air temperature to computer room air-conditioning units
- **Increased memory density:** To reduce dual in-line memory module (DIMM) count using technologies such as quad-rank DIMMs to help reduce excess heat and overall power consumption
- **Custom-designed chassis:** To help optimize density and airflow in each node, taking advantage of the absence of components such as redundant power supplies

Furthermore, the DCS team applies best thermal practices to design the layout of the entire facility. For example, using state-of-the-art CFD modeling, the DCS team can design the optimal structure and layout of racks, networks, and power and cooling systems, as shown in Figure 2. In particular, the team can suggest layouts at the component, server, rack, and row level that help optimize power flow, air flow, and heat removal. The team can also size HVAC and power distribution configurations based on a careful assessment of specific power consumption and cooling requirements.

When working with organizations to co-develop custom solutions, DCS system architects leverage industry-leading Dell technology and manufacturing expertise to deliver hardware that can meet specific hyper-scale computing needs with outstanding energy efficiency and functionality. Moreover, because DCS has full access to the outstanding Dell supply chain, manufacturing infrastructure, and global sales and service organization, DCS custom-built hardware can be manufactured and delivered quickly, and offer the same world-class warranty service as general-purpose Dell™ servers and storage.

CUSTOMIZING SERVICES FOR RAPID DEPLOYMENT AND HYPER-SCALE SUPPORT

DCS can also help organizations enhance efficiency by helping reduce deployment time and simplify maintenance. To do this, DCS offers a comprehensive range of services tailored to meet the service needs of hyper-scale data centers. Key DCS service offerings include the following:

- **Single point of contact:** Each organization is assigned a dedicated services account manager to help ensure that issues are properly routed and rapidly resolved.
- **Fully integrated racks:** By installing fully integrated racks, DCS can minimize the time required to deploy a hyper-scale data center. Deployment services include integrating the hardware solution into a rack and setting up cabling off-site, delivering it to the data center, unpacking it, connecting it to network uplinks and a power source, inspecting it, powering it up, conducting "burn-in" testing, and removing trash. DCS can deploy the solution into an existing rack in the data center as well.
- **On-site parts:** Dell-owned parts can be inventoried on location to enable rapid systems repair. Service technicians can easily scan parts in and out of inventory, and parts can then be replenished automatically. Alternatively, parts can placed off-site near the organization.

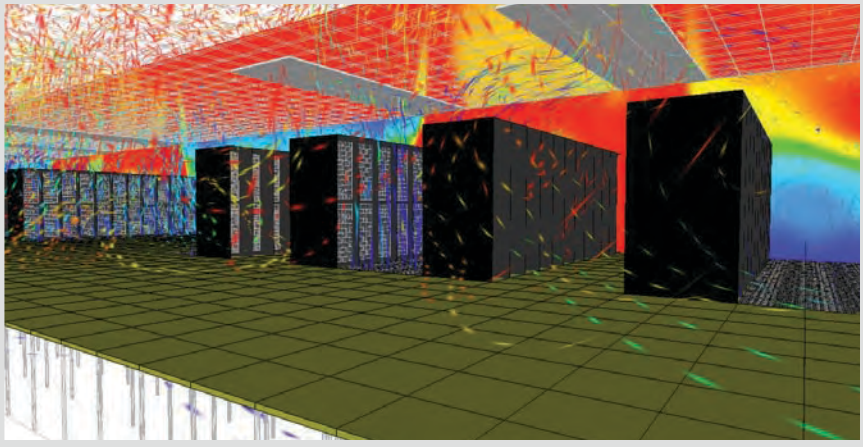

- Optimize layout and airflow paths and minimize air consumption, bypass, and recirculation
- Utilize techniques such as cold or hot air containment based on environment
- Make cable management, raised floor and drop ceiling plenum, and return path to the HVAC system as unrestrictive and efficient as possible
- Use closed loop-feedback to drive HVAC and power distribution units
- Size power distribution based on an accurate understanding of power consumption requirements
- Balance all parameters critical to the operational efficiency of a data center

Figure 2. *A data center facility and nodes are an intrinsically linked ecosystem*

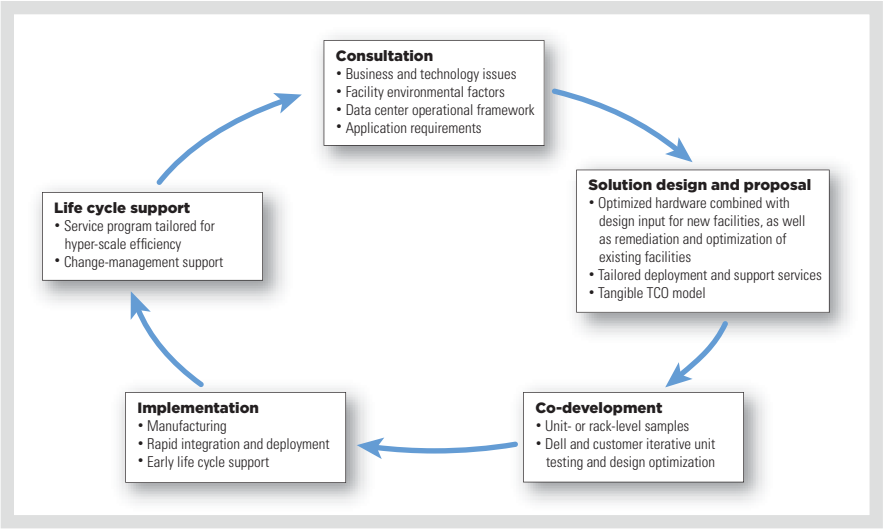

Figure 3. *DCS offers optimized solutions throughout the life cycle of a hyper-scale data center*

- **Node-level exchanges:** Fully tested and burned-in whole units can be located on-site, which allows an impaired system to be replaced immediately while it is repaired outside the rack, helping increase system uptime. Alternatively, whole units may be located off-site and delivered on a scheduled basis at the same time that impaired units are removed for repair at the off-site location.
- **On-site hardware troubleshooting and repair:** To help organizations rapidly troubleshoot and repair problems, DCS can deploy highly skilled personnel on-site either periodically or full-time.
- **Custom maintenance and support services:** To streamline ongoing support, DCS also provides several maintenance and support service options that can be customized to meet an organization's specific needs.

ENHANCING EFFICIENCY AND PERFORMANCE IN AN ENVIRONMENTALLY FRIENDLY WAY

In a hyper-scale computing environment, efficiency is a key source of business value. By offering optimized solutions throughout the life cycle of the data center, DCS can help organizations dramatically improve efficiency and lower the TCO of their hyper-scale data centers, as shown in Figure 3. In particular, by offering data center design consulting, designed-to-order hardware, and customized service offerings, DCS can help organizations rapidly deploy and cost-effectively maintain a hyper-scale data center that is designed for lower acquisition costs, lower power consumption, higher HVAC efficiency, lower network infrastructure costs, and lower deployment and management overhead than a traditional data center.

Moreover, by optimizing for energy efficiency at all levels of scale, Dell also helps ensure that DCS-designed data centers are environmentally friendly. Plus, DCS offers an optional carbon offset program in partnership with Carbonfund.org, which plants trees to offset the estimated "carbon footprint" of a participating organization's data center.

Jimmy Pike is director of system architecture and distinguished engineer in the Dell DCS Division. He is responsible for the strategic system architecture of the enterprise product line, and has more than 25 years of experience in the server industry. He was previously employed by Intel as part of its Enterprise Product Group, where he served as the director of telecom server development, director of server architecture, and director of engineering. He has received numerous patents through extensive work in the area of symmetric multiprocessing for both large and small systems.

Ty Schmitt serves as principal thermal and mechanical architect in the Dell DCS Division. He previously served seven years as the senior manager for the enterprise thermal, acoustical, and mechanical architecture division at Dell, and seven years as an enterprise mechanical architect and lead mechanical engineer responsible for the mechanical development of 10 enterprise and storage platforms. Ty has a bachelor's degree in mechanical engineering from Texas A&M University and holds 37 U.S. patents.

Frank Frankovsky leads the Solution Architect team in the Dell DCS Division and was a key leader in the inception and launch of DCS. Since joining Dell in 1991, he has held leadership roles in Dell enterprise marketing and technical sales organizations. Frank has a B.B.A. in Marketing from Stephen F. Austin State University and completed the Executive Management Program at the Jesse H. Jones Graduate School of Management at Rice University.

Todd Brannon is market development manager for the Dell DCS Division. Todd joined Dell in 1995 as a founding member of the System Performance Analysis lab and has since held a variety of enterprise marketing roles, including leading the launch of the Dell server and storage businesses in Japan. Todd has a B.S. in Electronic Media from Syracuse University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell Cloud Computing™ Solutions:
DELL.COM/CloudComputing

SIMPLIFYING IT WITH DELL ON-DEMAND DESKTOP STREAMING

By Aaron Prince
Bharath Vasudevan

The Dell™ On-Demand Desktop Streaming™ solution with Dell EasyConnect™ technology is designed to stream virtual disks from servers to diskless client systems. Deploying this solution can help enterprise administrators simplify IT, increase efficiency, and enhance security while providing end users with stable, robust client systems.

Related Categories:
Ardence Software
Citrix Systems
Dell OptiPlex desktops
Dell PowerEdge servers
Flexible computing
Simplify IT
Systems management
Visit DELL.COM/PowerSolutions for the complete category index.

For enterprise administrators, managing growing numbers of client systems can prove to be a difficult task. Concerns ranging from physical security to patch management can increase the total cost of ownership of these systems far beyond their initial purchase price. The Dell On-Demand Desktop Streaming solution in conjunction with Dell EasyConnect technology is designed to address these challenges, helping simplify IT, increase efficiency, and enhance security while providing end users with stable, robust client systems that can provide performance levels similar to those of comparably configured systems using traditional hard drives.

UNDERSTANDING DELL ON-DEMAND DESKTOP STREAMING

The Dell On-Demand Desktop Streaming solution utilizes Dell OptiPlex™ client desktops, Dell PowerEdge™ servers, and Citrix® Provisioning Server for Desktops software. In this solution, the OptiPlex clients boot like traditional clients, but receive the desktop session from the server in real time. To help minimize network traffic, the clients only pull the necessary information from the server on demand, rather than the entire image.

By utilizing diskless client systems that stream their virtual disks (known as vDisks) from a networked PowerEdge server in real time, this solution enables administrators to maintain tight control over their environment while helping simplify management and enhance security. On-Demand Desktop Streaming differs from other thin-client solutions in that the OptiPlex clients use their local processor, memory, and graphics resources, enabling them to achieve performance levels similar to those of comparably configured systems using local hard drives.

Administrators can configure client systems to provide a variety of boot options using one of two vDisk modes:

- **Standard mode:** This one-to-many mode streams a read-only vDisk to multiple client systems, helping both increase physical security and decrease the number of images to maintain, with updates to a single vDisk automatically updating multiple clients when the clients reboot.
- **Private mode:** This one-to-one mode creates an individual read/write vDisk for each client, essentially storing the client hard drive securely on a data center server.

Enterprises can deploy this solution in either standalone or highly available configurations. In the standalone configuration, one streaming server handles

booting, licensing, streaming, and caching for all clients. The highly available configuration, in contrast, incorporates up to 4 licensing servers and 10 streaming servers. This configuration allows clients to fail over to a different server if their current server fails, and enables load balancing across multiple servers. To help avoid performance degradation, enterprises should size their On-Demand Desktop Streaming deployment such that it can handle the total user load even following a server failure. Figure 1 illustrates an example of each configuration, with the highly available configuration incorporating two streaming servers and network attached storage (NAS) for client write caching.

SIMPLIFYING IT AND INCREASING EFFICIENCY

When clients boot using standard mode, they cache the writes that would normally be applied to their hard drive to storage on a networked server. This approach allows the vDisk to be read-only and remain unchanged regardless of the number of clients using it. If an OS patch is required, administrators can switch the vDisk to private mode, apply the changes, and redeploy the vDisk in standard mode. The changes are then propagated to all clients using that vDisk the next time they reboot, helping simplify routine system maintenance.

Moving data from individual consumer-class hard drives—which may have a higher failure rate relative to other components—to a redundant array of enterprise-class server drives also helps significantly reduce the risk of data loss from a drive failure. And storing data in a controlled data center environment, where administrators can maintain optimal power and cooling conditions, helps significantly reduce the risk of data loss from other hardware failures as well.

Another advantage of centralized user data is that the client hardware becomes easily replaceable. When hardware failures occur, enterprises should not need a large team of experienced IT staff to reinstall and reconfigure the replacement hardware. Instead, administrators should be able to simply plug in the replacement client and start streaming the same image immediately—helping significantly reduce end-user downtime following a failure. In addition, because standard mode does not allow clients to commit changes to the OS, performance typically does not degrade over time—every time a client boots, it should start with a fresh OS that can perform as it did the first time the client booted. This approach helps significantly reduce problems related to OS stability, performance, or functionality, helping free IT staff to focus on other matters.

ENHANCING SECURITY

Discarding traditional client systems carries the risk of residual, and potentially sensitive, information remaining on the hard drive. Because the Dell On-Demand Desktop Streaming solution does not store data on local client systems, organizations can typically discard these systems as needed without worrying about this type of residual data. When working with extremely sensitive information, administrators can also use the Port Blocker feature of the On-Demand Desktop Streaming software to help prevent the removal of data from the client system by selectively disabling components such as USB ports, CD drives, or floppy disk drives. This additional layer of security helps protect against someone removing data from the IT environment either inadvertently or intentionally.

The Dell On-Demand Desktop Streaming solution also helps control spyware and viruses. In standard mode, rebooting a client system restores the OS

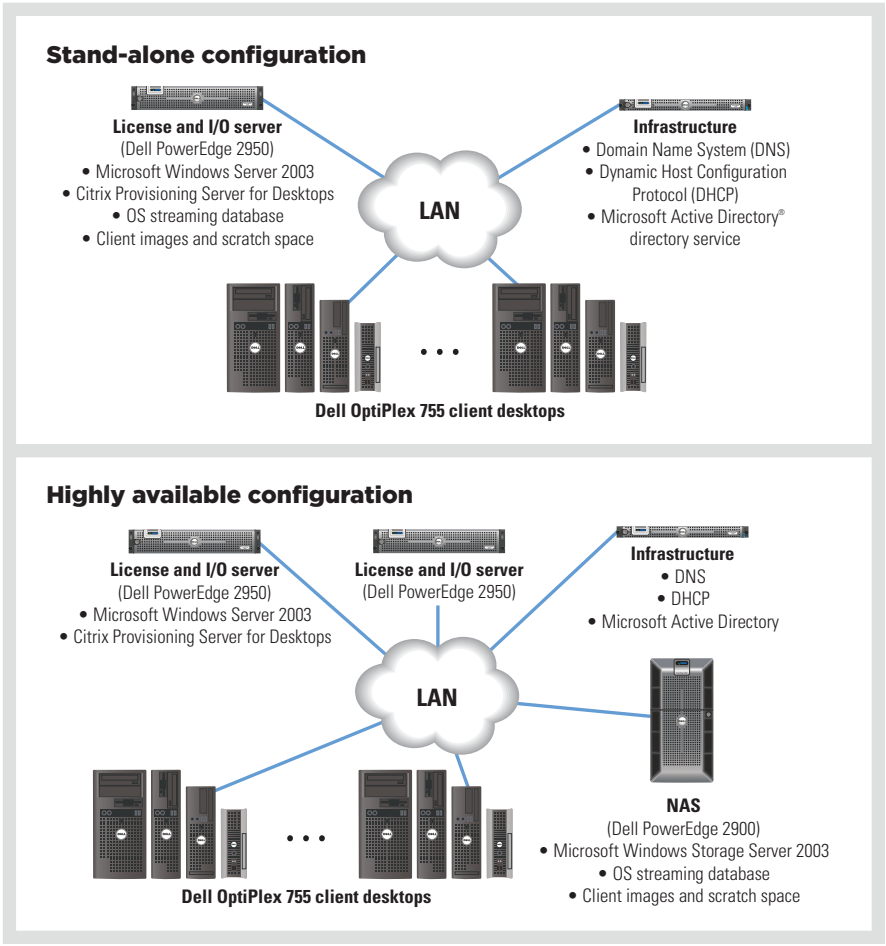

Figure 1. *Stand-alone and highly available configurations of the Dell On-Demand Desktop Streaming solution*

to its original state, helping avoid undesirable software that can compromise security or reduce performance. This approach can allow administrators to quickly and efficiently return infected clients to a known state through a simple reboot.

INTEGRATING DELL EASYCONNECT TECHNOLOGY

Dell EasyConnect technology is designed to optimize desktop systems to help automate client deployment and save time. By embedding key information in the BIOS of client systems, Dell EasyConnect helps provide simplified licensing, easy deployment, and enhanced server management.

Simplified licensing

Dell On-Demand Desktop Streaming solutions with Dell EasyConnect technology include client license information embedded in the BIOS of client systems, avoiding the need for separate client licensing. Embedding these licenses on clients instead of servers also enables administrators to avoid the need to reassign licenses between servers when moving servers or clients. Nor do they need to retrieve lost or corrupt licenses following a server failure, enabling them to easily have clients back up and running once the server is brought back online or replaced. These features help reduce deployment time and effort, and allow administrators to add systems without needing to purchase, install, or activate new licenses.

Easy deployment

EasyConnect also provides an embedded bootstrap in the BIOS of client systems, helping eliminate the need for Preboot Execution Environment (PXE) servers to boot clients over a network. If the organization already uses PXE, this approach helps avoid the need to divide the network to segregate multiple PXE servers. Or, if the organization does not want to use PXE, this solution can fit seamlessly into the environment. Because only clients with Dell EasyConnect technology can boot over the network in this type of configuration, administrators can maintain tight control over network access.

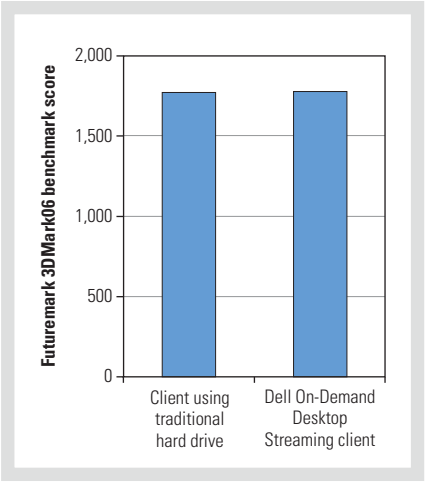

Figure 2. *Performance test results comparing a Dell OptiPlex desktop using a traditional hard drive with one using Dell On-Demand Desktop Streaming*

Enhanced server management

Adding clients to the On-Demand Desktop Streaming database in traditional deployments can be time-consuming, requiring administrators to first gather information such as the Media Access Control (MAC) address of each client and then enter this information into the database. Dell EasyConnect integrates this information into the BIOS of client systems, allowing clients to be automatically added to the database when they first boot, using their Dell service tag as a unique identifier. After clients are added to the database, administrator-specified settings can be automatically applied and the proper vDisk can be automatically associated with the service tag and delivered to the client without administrator intervention. This approach allows administrators to quickly bring many clients online while helping ensure that they are all configured in the same way.

TESTING DELL ON-DEMAND DESKTOP STREAMING

In August 2007, Dell engineers performed tests using the Futuremark 3DMark06 benchmark on Dell OptiPlex desktops—one with a traditional hard drive and one using Dell On-Demand Desktop Streaming simultaneously with 99 other clients using the same vDisk in standard mode on a stand-alone server. The test environment consisted of a Dell PowerEdge 2950 server with one dual-core Intel® Xeon® 5160 processor at 3.0 GHz, 4 GB of RAM, and the Microsoft® Windows Server® 2003 Release 2 (R2) Standard Edition OS connected through a Gigabit Ethernet network to Dell OptiPlex 745 desktops, each with one Intel Core™2 Duo E6400 processor, 1 GB of RAM, an ATI Radeon™ X1300 PRO graphics processor, and the Microsoft Windows® XP Professional Edition OS.

The 3DMark06 benchmark stresses processor, memory, and graphics resources using high-resolution game sequences and assigns a score based on system performance. Figure 2 shows that both systems returned virtually the same score—1,783 for the traditional hard drive configuration, and 1,788 for the streaming configuration—demonstrating that Dell On-Demand Desktop Streaming does not result in a client performance penalty.

BUILDING A FLEXIBLE OS STREAMING ENVIRONMENT

The Dell On-Demand Desktop Streaming solution with Dell EasyConnect technology is designed to simplify IT, increase efficiency, and enhance security for enterprise administrators while providing end users with stable, robust client systems that can provide performance levels similar to those of comparably configured systems using local hard drives. Administrators should carefully evaluate the number of servers and streaming server options when choosing a configuration. Deploying a properly sized and configured environment can help administrators increase IT agility, security, and productivity in their infrastructure.

Aaron Prince is a systems engineer on the Dell End-to-End Solutions team.

Bharath Vasudevan manages the Dell End-to-End Solutions team.

When a server or router fails, your business can come to a stop in a matter of seconds. Now more than ever, you need **anytime, anywhere, comprehensive management** of your data center assets. Avocent DSView® 3 software, in conjunction with Dell™ KVM over IP switches, MergePoint® service processor managers and ACS console servers, lets you **securely view, access, and control every device in your data center** in real-time from a single console. To learn how you can ensure that your data center is always available, even when you're not, visit us at http://dell.avocent.com/availability

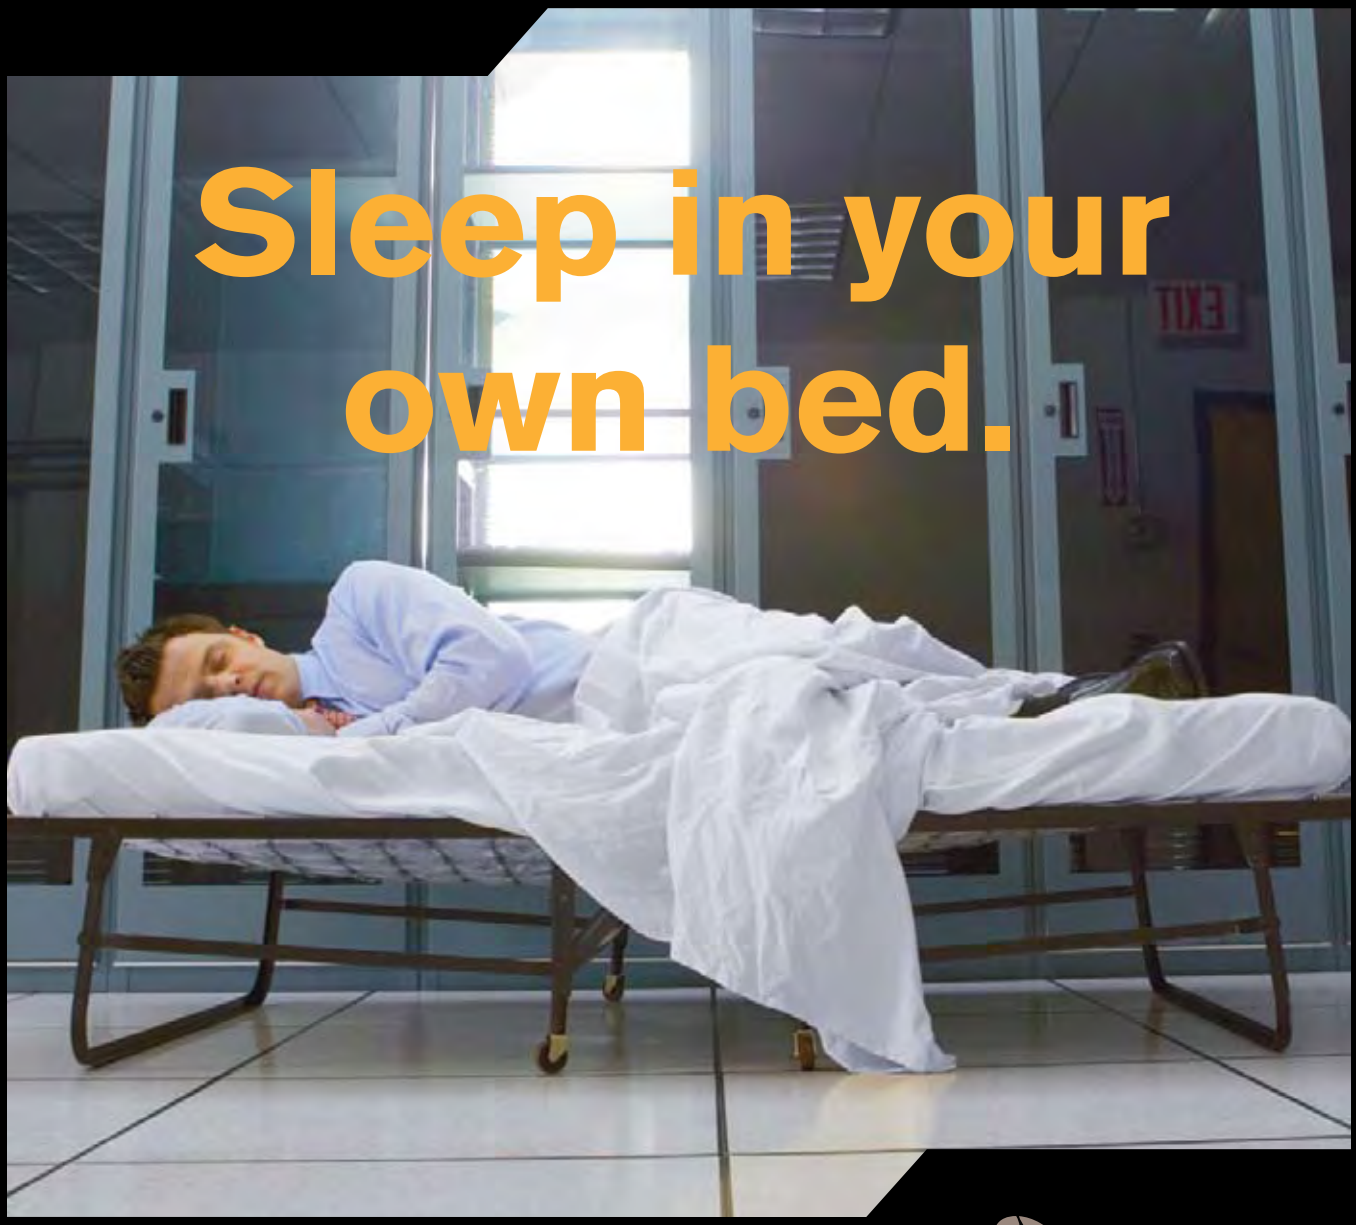

Avocent.
The Power of Being There.

EXTENDING THE ADVANTAGES OF DATA CENTER CONSOLIDATION WITH DE-DUPLICATION TECHNOLOGY

By K. E. H. Polanski

Hardware consolidation has become a key way for enterprises to increase efficiency and reduce ongoing data center costs. By combining this approach with de-duplication technologies such as those in CommVault® Simpana™ 7.0 software, which are designed to remove unnecessary copies of data, they can reclaim storage space and extend consolidation down to the data itself.

Related Categories:

CommVault

Dell PowerVault storage

Server consolidation

Storage

Storage software

Visit DELL.COM/PowerSolutions for the complete category index.

Enterprises must constantly contend with the pressures of limited floor space, energy, and cooling systems. One widespread approach to coping with these pressures is server consolidation through virtualization, which helps reduce the number of physical servers required to support key systems. To complement this type of approach and help further increase return on investment, many organizations are also considering implementing de-duplication technologies, such as those in CommVault Simpana 7.0 software. These technologies are designed to eliminate duplicate copies of data, reclaim storage space, and extend consolidation from the physical infrastructure down to the data itself.

UNDERSTANDING DIFFERENT APPROACHES TO DE-DUPLICATION

De-duplication technologies encompass a variety of approaches. One common type is hardware-based de-duplication, in which the de-duplication functionality is integrated into the storage system itself. This model can work well for individual storage systems, allowing them to store more logical data with less storage space than would be possible without de-duplication; however, it is limited to data within that specific device. Software de-duplication, in contrast, can work with many different types of storage across a data center environment, providing more flexibility than hardware de-duplication.

De-duplication can also work on both the file level and the block level. File-based de-duplication uses hash algorithms to find and remove duplicate files, while block-based de-duplication uses hash algorithms to find and remove duplicates of the blocks that make up the files. By providing reference pointers to a single copy of a file or block, rather than storing multiple copies in different locations, de-duplication helps reduce storage requirements and increase efficiency. By working along two different dimensions, file- and block-based approaches can serve as complements to one another, with the file-based approach removing "horizontal" copies of files and the block-based approach working "vertically" within these files to remove duplicate blocks.

When evaluating de-duplication approaches, organizations should keep in mind that they are not implemented in isolation, but as part of a comprehensive system of data management used in a particular environment. That environment typically includes some combination of backup and recovery, archive, snapshot, and replication technologies in addition to de-duplication. It may also require encryption or compression when data is transmitted across wide area networks or stored on tape for long-term

> **"De-duplication technologies are designed to eliminate duplicate copies of data, reclaim storage space, and extend consolidation from the physical infrastructure down to the data itself."**

off-site retention, and may include service-level agreements for levels of data recovery that the organization must support and enforce.

When taking these factors into account, the differences between file- and block-based de-duplication can be significant. For example, when file-based de-duplication removes copies of files, the image of the file is preserved. When block-based de-duplication removes identical blocks, in contrast, the image is not preserved, and must be rebuilt before files can be restored—an important consideration for organizations where rapid recovery is a high priority. In addition, while file-based de-duplication can integrate easily with encryption and compression processes, doing so with block-based de-duplication typically causes a net expansion of the data footprint.

Both file- and block-based de-duplication can be hardware independent, but file-based approaches are typically software implementations, while block-based approaches are typically hardware implementations. In these forms, file- and block-based de-duplication can work together, with the software enabling de-duplication of files stored on virtually any disk while the hardware provides complementary block-based functionality for specific devices.

USING DE-DUPLICATION WITH BACKUP AND ARCHIVE SYSTEMS

Single-instance store (SIS) capabilities, such as those provided by CommVault Simpana 7.0, are a key method of file-based de-duplication. By comparing the contents of files to locate duplicates regardless of differences in file names, properties, or attributes, SIS helps eliminate duplicate copies of files and attachments found in backups and archives stored on disk. This method helps ensure that exact digital duplicates can be matched and eliminated across data sets, applications, clients, and OS platforms to help maximize compression.

As new backup and archive copies are made, this compression ratio can continuously increase as the number of eliminated duplicate files increases. As shown in Figure 1, SIS is particularly well suited for eliminating duplicate files across backup cycles, eliminating duplicate files across backup and archive copies, and consolidating tape copies of data to SIS disk copies for search and discovery. SIS can apply to both primary backup and archive copies of data as well as to secondary copies created by storage policies. As logical copies of data age and are gradually retired, reference links to the single physical copy can be eliminated until none remain, after which the physical copy can also be removed. Figure 2 lists some key features and advantages of the SIS capabilities of CommVault Simpana 7.0.

IMPLEMENTING DE-DUPLICATION IN THE DATA CENTER

Tape backup, the traditional method of backup and archiving, is still useful and appropriate for many uses, particularly off-site rotation and long-term data preservation. However, combining de-duplication with Serial Attached SCSI (SAS) and Serial ATA (SATA) storage systems enables organizations of all sizes to maximize the advantages of simple, cost-effective disk-based data management.

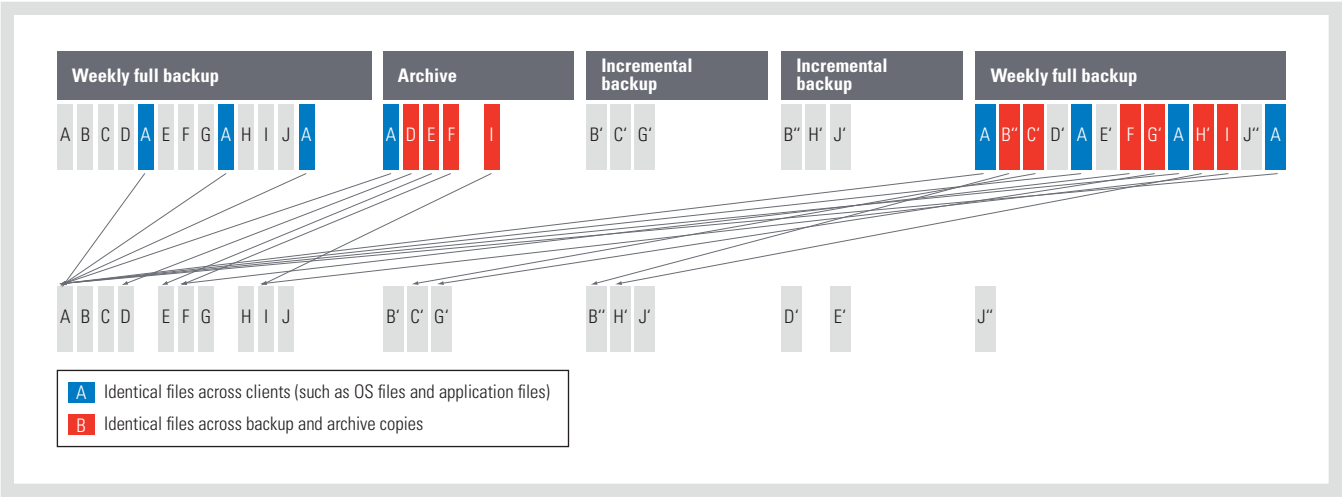

Figure 1. *Elimination of duplicate files across backups and archives through single-instance store de-duplication*

Storage systems such as the Dell™ PowerVault™ MD3000 and PowerVault MD3000i, for example, combine disk-based storage with tape and are well suited for SIS backup and archiving for small organizations. Using SIS with these systems can help organizations keep more data on disk for longer periods than they could without this functionality, and can help accelerate recovery while also preserving data on tape. Organizations can also use the same approach with Dell/EMC CX series arrays, perhaps in combination with virtual or physical tape libraries, to help accommodate data growth while still taking advantage of de-duplication features.

In addition, organizations can combine SIS with other technologies to help further increase storage efficiency. For example, backup systems that enable aggressive "incremental forever" backup cycles without compromising data recovery can dramatically reduce the amount of data duplicated as part of the normal backup process, helping increase the cost-effectiveness of disk-based data retention even without de-duplication. Data archiving systems can also help reduce the amount of data that organizations must back up; by moving old data to a secondary storage location for long-term retention, organizations can avoid the need to repeatedly protect that data across multiple full backup cycles. This type of holistic approach to data management can help maximize compression ratios and reduce the amount of data that must be stored on either disk or tape.

"As enterprises strive to reduce storage costs, particularly those related to backup and archiving, they can take advantage of de-duplication technologies such as those in CommVault Simpana 7.0 to help reclaim storage space, accelerate data recovery, and maximize return on investment."

INCREASING STORAGE EFFICIENCY

De-duplication can be a key part of efforts to consolidate hardware and increase efficiency. As enterprises strive to reduce storage costs, particularly those related to backup and archiving, they can take advantage of de-duplication technologies such as those in CommVault Simpana 7.0 to help reclaim storage space, accelerate data recovery, and maximize return on investment.

Feature	Advantages
Removal of duplicate files from backup and archive copies	By reducing the physical disk space required to store logical data, organizations can reduce storage requirements, increase the amount of data they can hold in short-term retention on disk, and enable rapid recovery and search access.
Support for existing disk hardware	Organizations can typically implement the software without changing their physical infrastructure, easily migrate SIS copies to new disk arrays, and use the software with different types of storage from different vendors.
Support for encryption and compression	Unlike block-based de-duplication, organizations can combine SIS with data compression and encryption to secure and further compress data sent over networks and written to disk.
On-demand use	Administrators can run de-duplication when the initial backup or archive operation occurs or apply it later to existing data copies, including those made with previous versions of CommVault software for searching or staging purposes.
Rapid restore	Unlike block-based de-duplication, SIS does not need to rebuild file images before restoring files, helping accelerate data recovery.
Support for block-based de-duplication	Organizations can write SIS data copies to de-duplication appliances, enabling them to combine SIS file duplicate removal and hardware-based block duplicate removal to help further reduce storage requirements.

Figure 2. *Key advantages of the single-instance store features in CommVault Simpana 7.0*

K. E. H. Polanski is a partner at the KEHP Group, a marketing and public relations firm specializing in storage and data management. She was previously the director of product marketing at CommVault, and has worked in product marketing and management, business development, and channel marketing at EMC, Legato, the Qualix Group, and Octopus Technologies. She has a degree in Computer Science from Augustana College.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

CommVault:
www.commvault.com

Dell and CommVault:
DELL.COM/CommVault

vmware

THE POWER OF VIRTUALIZATION

The Power of Virtualization enables mainframe-class IT infrastructures that are always-on, offer capacity on demand and provide policy-based automation in IT environments.

Dell and VMware jointly integrate and certify VMware Infrastructure 3 Software with Dell PowerEdge Servers, Storage and Infrastructure Consulting Services to deliver virtual infrastructures that simplify and maximize IT investments across enterprises, institutions and agencies–large, medium or small.

THE POWER OF VIRTUALIZATION :: vmware.com/virtuallyspeaking

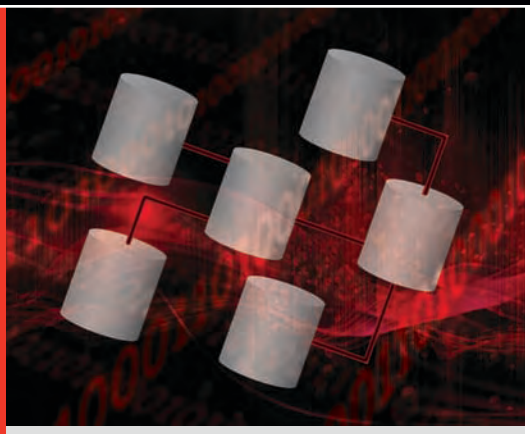

SAFEGUARDING DATA WITH DELL POWERVAULT DATA PROTECTION SOLUTIONS

By Sanjeet Singh
Jason Buffington

Dell and Microsoft have worked together to create unified disk-to-disk-to-tape systems—Dell™ PowerVault™ Data Protection Solutions—that help avoid the drawbacks of traditional tape- and disk-based data protection. Optimized for Microsoft® software–based workloads, these systems can provide versatile, comprehensive protection for environments built with Dell PowerVault servers and Microsoft System Center Data Protection Manager 2007.

Related Categories:
Business continuity
Data security
Dell PowerEdge servers
Dell PowerVault storage
Microsoft System Center
Storage
Visit DELL.COM/PowerSolutions for the complete category index.

Over the past several years, Dell and Microsoft have worked together to bring advanced storage platforms into the mainstream by combining cost-effective, standards-based Dell servers and storage with the Microsoft Windows® Storage Server platform. The combination of Dell hardware and Microsoft software expanded with the introduction of Microsoft Windows Unified Data Storage Server (WUDSS), which offers not only optimized file sharing and storage management features but also block-level Internet SCSI (iSCSI) functionality, helping provide efficient and easy-to-manage storage while avoiding some of the costs and complexities of Fibre Channel networks. Dell and Microsoft continue to bring advanced storage technologies to the mainstream with the introduction of data protection appliances designed for improved backups.

Although nightly tape backups are still a standard practice, many organizations have found that they also need more comprehensive protection than these nightly systems can provide—and have realized that tape is not an ideal medium for many data recovery situations. To help meet the needs of these enterprises, Dell and Microsoft have again teamed up to deliver unified disk-to-disk-to-tape Data Protection Solutions based on Dell PowerVault servers and Microsoft System Center Data Protection Manager (DPM) 2007.[1] Designed for ease of use and simplified deployment, these Dell PowerVault systems are preconfigured with the Microsoft Windows Storage Server OS, management functions, and DPM 2007. They provide a combination of disk-based continuous data protection (CDP) and traditional tape backup specifically designed to protect major Microsoft workloads such as Microsoft Exchange, Microsoft SQL Server™, Microsoft Office SharePoint® Server, and Microsoft Virtual Server software as well as file services in the Microsoft Windows Server® 2003, Windows Server 2008 (code-named "Longhorn"), Windows XP, and Windows Vista® operating systems.

CHALLENGES OF TRADITIONAL DATA PROTECTION APPROACHES

Tape is still an integral component of data protection infrastructure, and often a good choice for whole-server

[1] For more detailed information about DPM 2007 functionality, see "A Look Inside Microsoft System Center Data Protection Manager 2007," by Jason Buffington and Sanjeet Singh, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070485-Singh.pdf.

recoveries and long-term archiving. However, it also has key disadvantages:

- Organizations can lose up to a full day's worth of data depending on when a server fails: if it fails early in the morning, they may lose very little or no data, while if it fails at the end of the day, they may lose a full day's worth of data. Typically, administrators must assume that they will lose half a day's worth of data each time a server fails.
- Data recovery can be time-consuming, often requiring administrators to restore a full backup, then each differential backup (or perhaps just the latest incremental backup). In addition, if the tape fails during this process, they must then take additional time to retry with the previous tape. And if the failed tape was a weekly full backup, they not only must take that additional time, but may have lost additional data as well.
- Tape typically provides poor performance when restoring individual files, folders, databases, and mailboxes—operations that usually constitute a majority of data recovery tasks.

Various disk-to-disk approaches can help overcome these problems, including virtual tape libraries (VTLs), real-time replication (RTR), and CDP.

Virtual tape libraries. VTLs use disk-based storage as simulated tape, which can provide higher performance than real tape. In addition, because traditional tape backup systems see VTLs as high-performance tape devices, they can typically begin utilizing VTLs immediately. VTLs can often serve as a bridge technology suitable for very large enterprises that are already heavily invested in their backup infrastructure. When those enterprises upgrade their backup software, they can also replace their VTLs with fully integrated backup-to-disk technologies.

The downside of this type of system is that the tape backup methodology remains the same, utilizing traditional full, incremental, and differential backups and typically restricting backup frequency to once per day. In addition, the virtual tape still behaves as a large sequential repository, rather than the random access medium that disk actually is.

Real-time replication. While VTLs use disk as a substitute for tape, RTR treats disk as disk, usually by asynchronously replicating data from primary production disks to secondary disks in an application-agnostic way. RTR can capture data changes as they occur—typically with low latencies—and send them to a secondary server, often with the goal of increasing system availability. If the primary server fails, the secondary server has a near-current copy of the data and can easily take over the primary role.

RTR is particularly useful when core enterprise applications do not incorporate high-availability features. But because RTR focuses on maintaining a near-current copy of production data for failover, rather than past iterations for data recovery, it is actually providing high availability rather than true data protection.

Continuous data protection. As the name suggests, CDP is designed to provide continuous protection (through replication, synchronization, or mirroring), but also includes multiple previous restore points. By definition, it typically includes the ability to restore to any previous point in time or operational transaction (depending on workload). Some currently available CDP solutions offer these extremely granular recovery points. For many enterprises, however, choosing between 10:48:23 or 10:48:24 or 10:48:25, and so on, may be a burden more than a useful feature, and the administrator may not even know whether these recovery points are consistent. Instead, using DPM to create continuous, predictable, and consistent recovery points every 15 minutes or every hour can help simplify matters for both end users and IT staff. Moreover, because CDP software often focuses on a single application, enterprises using this approach may need to invest in different software for different applications, increasing complexity and introducing different levels of coverage.

One challenge with all three of these data protection approaches is that they often come from vendors other than traditional tape backup companies. Enterprises often must look at a tape backup system for long-term retention and an advanced disk-to-disk system (such as RTR or CDP) for rapid recovery, and may later discover that the two do not interoperate, and might not even be able to coexist on certain systems.

DATA PROTECTION SOLUTIONS BASED ON DELL POWERVAULT AND DPM 2007

Many enterprises use some form of nightly tape backup, but also recognize the need for more comprehensive data protection than these nightly backups can provide. To help meet this need, Dell and Microsoft have created unified disk-to-disk-to-tape systems based on Dell hardware and DPM 2007 (see Figure 1). These comprehensive

"Many organizations have found that they need more comprehensive protection than nightly systems can provide—and have realized that tape is not an ideal medium for many data recovery situations."

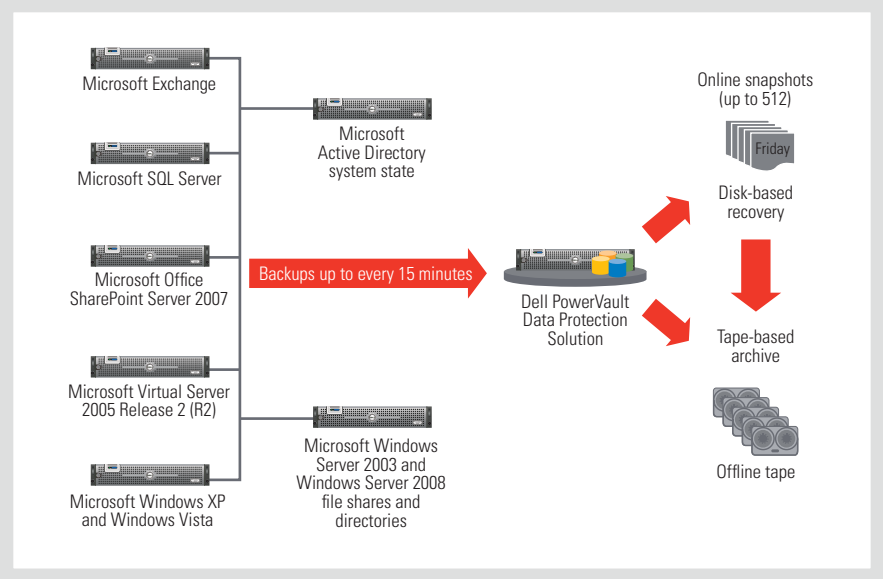

Figure 1. *Comprehensive disk-to-disk-to-tape backup system based on Dell hardware and Microsoft System Center Data Protection Manager 2007 software*

systems are optimized for core Microsoft workloads and can offer multiple advantages:

- **Consistent design and support:** Because the Data Protection Solutions utilize Dell hardware and Microsoft software and are designed expressly for Microsoft workloads, enterprises can take advantage of technical support from the company that created the applications themselves, without relying on a third party.
- **Comprehensive capabilities:** The DPM 2007–based Data Protection Solutions are designed to provide both disk-based CDP and traditional tape backup. And while they may not protect every niche application, they can provide broad support for the primary Microsoft workloads.
- **Cost-effectiveness and flexibility:** By creating multiple systems with different storage capacities based on cost-effective software, Dell and Microsoft have helped reduce the cost barrier to using disk-based protection for enterprises of all sizes.
- **Simplified management:** DPM 2007 is designed to simplify management through intuitive user interfaces and wizards.

The Data Protection Solutions are preconfigured to help simplify deployment. Administrators can deploy them simply by powering them on, joining them to the appropriate Microsoft Active Directory® domain, and installing the DPM 2007 agent on each server running Windows by using the DPM 2007 Administrator Console, deployment software such as Microsoft System Center Configuration Manager 2007, Microsoft Systems Management Server 2003, or even a Group Policy within Microsoft Active Directory.

After deployment, DPM 2007 creates an initial copy and then monitors data as it changes. It then periodically creates snapshots to build multiple disk-based copies designed for rapid and reliable recovery. By storing only the block-level differences between one instance and the next, the DPM 2007 server can efficiently maintain several recovery points on disk that can be supplemented with long-term tape-based archives. When necessary, administrators can simply select a restore point to roll back data to a previous point in time.

UNIFIED DATA PROTECTION FROM DELL AND MICROSOFT

Although still a good choice for some purposes, traditional tape- and disk-based data protection can have disadvantages, including poor performance and a lack of flexibility. Dell PowerVault Data Protection Solutions built with Microsoft System Center Data Protection Manager 2007 combine disk and tape backups to create versatile, easy-to-use data protection systems for Microsoft workloads.

Sanjeet Singh is a global product marketing manager in the Dell Enterprise Storage Group. He has eight years of experience in developing and delivering business-critical technologies, including databases and data protection. Sanjeet has an M.S. in Computer Engineering from Purdue University and an M.B.A. from the University of Texas.

Jason Buffington is the senior technical product manager for Data Protection Manager and other storage solutions at Microsoft. He has more than 16 years of experience in data protection and business continuity, and studied Computer Science at Texas A&M University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Buffington, Jason, and Sanjeet Singh. "A Look Inside Microsoft System Center Data Protection Manager 2007." *Dell Power Solutions*, February 2008. DELL.COM/Downloads/Global/Power/ps1q08-20070485-Singh.pdf

Dell storage:
DELL.COM/Storage

Microsoft System Center Data Protection Manager 2007:
www.microsoft.com/dpm

A LOOK INSIDE MICROSOFT SYSTEM CENTER DATA PROTECTION MANAGER 2007

By Jason Buffington
Sanjeet Singh

Microsoft® System Center Data Protection Manager (DPM) 2007 is designed to combine the advantages of traditional tape backup and disk-based continuous data protection processes in a unified platform. This article delves into the key functionality of DPM 2007 and how it can help streamline and enhance data protection processes.

Related Categories:
Business continuity
Data security
Dell PowerEdge servers
Dell PowerVault storage
Microsoft System Center
Storage
Visit DELL.COM/PowerSolutions for the complete category index.

For years, the basic minimum standard in traditional data protection has been nightly tape backup. The standard approach has changed very little: once a week, production server data is streamed to tape media during nonpeak hours, a process referred to as a full backup. Then, each night, backup agents identify files or other data objects that have changed and send those to tape as well, in the form of either incremental backups (collective changes in the data since the last full backup) or differential backups (changes in the data each day). Incremental backups typically allow faster recoveries than differential backups, while differential backups typically allow faster nightly backup processes than incremental backups.

Although tape hardware technology has advanced, the disadvantages of nightly tape backup have remained the same, including the potential to lose a full day's worth of data, lengthy recovery times, and poor performance when restoring individual files or other data objects—which can constitute a significant percentage of recovery operations. Although disk-based approaches such as virtual tape libraries (VTLs), real-time replication (RTR), and continuous data protection (CDP) are designed to avoid or mitigate some of these problems, they also come with their own disadvantages. VTLs use the same approach as traditional tape backup, simply substituting disk media for tape media. RTR maintains a near-current copy of production data but does not typically provide multiple recovery points, and therefore is more suitable for high availability than data protection. And CDP is often designed specifically for individual applications, requiring enterprises to purchase and deploy multiple systems that provide differing capabilities and may not be compatible with one another.

Microsoft System Center Data Protection Manager (DPM) 2007 is designed to combine the advantages of traditional tape backup and disk-based CDP processes using an integrated disk-to-disk-to-tape system that is optimized for Microsoft workloads such as Microsoft Exchange, Microsoft SQL Server™, Microsoft Office SharePoint® Server, and Microsoft Virtual Server software, as well as file services in the Microsoft Windows Server® 2003, Windows Server 2008 (code-named "Longhorn"), Windows® XP, and Windows Vista® operating systems. Utilizing key DPM 2007 functionality can help administrators streamline and enhance their data protection processes while increasing flexibility and simplifying backup management.[1]

[1] For a more detailed discussion of traditional tape backup; disk-based VTL, RTR, and CDP approaches; and how Dell and Microsoft are working together to create DPM 2007–based data protection systems, see "Safeguarding Data with Dell PowerVault Data Protection Solutions," by Sanjeet Singh and Jason Buffington, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf.

UNDERSTANDING MICROSOFT SYSTEM CENTER DATA PROTECTION MANAGER 2007

DPM 2007 incorporates two key data protection features—express full backups and transaction replication. The express full feature provides block-level synchronization utilizing Microsoft Volume Shadow Copy Service (VSS) along with the built-in VSS writers for supported Microsoft server applications and the VSS writer in the native file systems for Window Server 2003, Windows Server 2008, Windows XP, and Windows Vista. The transaction replication feature uses application-aware agents that can protect Microsoft Exchange and SQL Server data through their transaction logs.

Express full backups

The express full feature creates a bitmask of the physical volumes that are being protected and is designed to work across different data sources. When administrators select a data source to protect, such as a SQL Server database, DPM 2007 identifies the files that make up the database and which disk blocks compose those files. The DPM bitmask table essentially contains a 0 or 1 for each block that composes a file or other data object, with 0 denoting a block that has not changed and 1 denoting a block that has changed. As data changes during the course of a day or week and the blocks are updated, the DPM agent changes the bits for those blocks from 0 to 1.

When a scheduled express full task executes, the DPM agent invokes a VSS snapshot of the data area, including the Windows file system and any source-specific VSS writers (such as Exchange, SQL Server, SharePoint, or Virtual Server). Within each application that uses VSS, the VSS writer helps ensure that the data is consistent and ready to be backed up. In addition, VSS provides the ability to have a "frozen" set of disk blocks (a shadow copy) while the production disk continues to service active I/O. The DPM agent then fetches only the changed blocks from within the shadow copy, while the production processes continue to function.

Rather than using approaches such as RTR (which is usually application agnostic and may not ensure write-order integrity) or CDP (which is often designed for a specific application), DPM uses the underlying VSS capabilities in Windows file systems and Microsoft applications to provide data-consistent protection capabilities across different data sources. After invoking the VSS snapshot, DPM can transmit the changed blocks while the production file system continues to operate. When DPM has synchronized all the marked blocks with the DPM server, the snapshot is released.

Microsoft terms these *express full backups*. Administrators should note that this option is not the same as a traditional full tape backup. When the backup process is complete, DPM does have a near-exact copy of the production data as it was when the process began—effectively a full backup. However, unlike a traditional full tape backup, DPM tracks and sends only changed blocks, ignoring unchanged blocks and stagnant files. The advantage of this approach is significantly reduced bandwidth requirements, which helps make DPM particularly useful for centralized backup of data from branch offices and other sources across a wide area network (WAN).

The result is a network-optimized synchronization of production data to the DPM server. DPM can then facilitate snapshots of the replicated area to provide up to 512 recovery points by storing only the block-level changes between each iteration. DPM express full backups would typically occur as a nightly scheduled task, but can occur as often as every half hour or as seldom as once per week.

Transaction replication

Transaction replication uses DPM agents to secure any application transaction logs that have closed since the last time window (as often as every 15 minutes) and transmit them to the DPM server. This feature sounds similar to log shipping, in which the secondary server applies logs automatically and repeatedly to maintain a near-current copy of the data, with the goal of increasing availability. However, because DPM is designed to increase data protection rather than availability, the DPM server retains each set of transaction logs as a differential backup. And just as a traditional tape backup might use its latest full backup plus some set of differentials to restore data to a previous point in time, DPM can use the live replica data plus the differentials to restore a production server to any previous point in time that the DPM server has transaction logs for.

INTEGRATING REPLICATION FEATURES

Combining the express full and transaction replication features can provide significant administrative flexibility. For example, because a DPM 2007 server can maintain up to 512 VSS snapshots, performing one express full backup each week in addition to synchronizing transactions every 15 minutes could result in an Exchange or SQL Server system having more than 344,000 total recovery

"DPM 2007 is designed to combine the advantages of traditional tape backup and disk-based CDP processes using an integrated disk-to-disk-to-tape system optimized for Microsoft workloads."

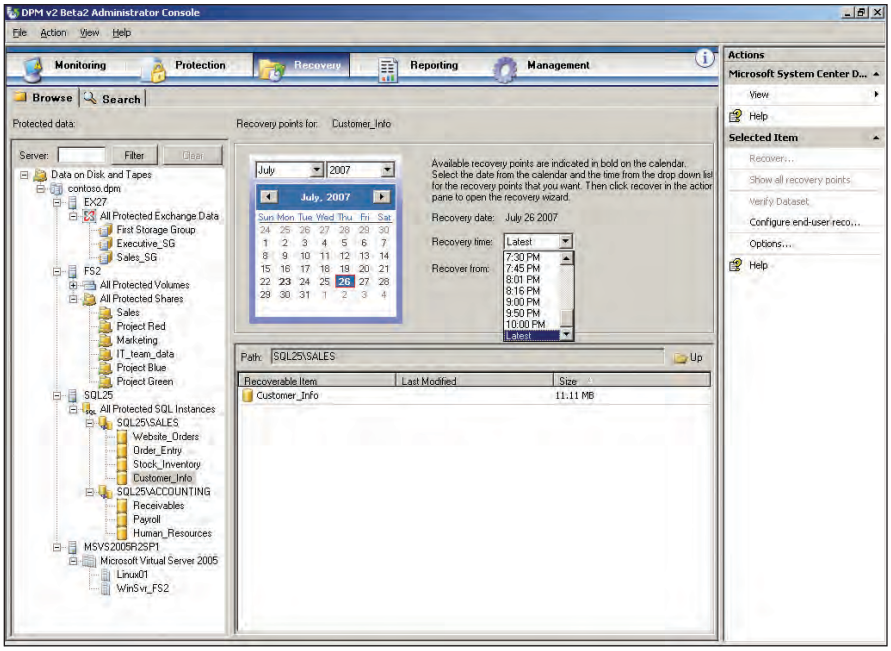

Figure 1. *Simplified recovery in the Microsoft System Center Data Protection Manager 2007 Administrator Console*

points. Performing one express full backup each day in addition to capturing transaction replication differentials every 15 minutes could result in nearly 50,000 recovery points, and potentially allow administrators to recover data to any 15-minute point in time over the past year and a half. Realistically, organizations would normally have 30–90 days of data on disk-based backup, including 15-minute transactional differentials, and then offload this data to tape for long-term retention.

Although this potentially large number of recovery points may seem imposing, the DPM Administrator Console is designed to simplify data recovery (see Figure 1). For example, if a database server's hard drive fails at 11:58, the DPM server could potentially recover the data to 11:45, 11:30, 11:15, and so on. However, the DPM Administrator Console also allows administrators to simply choose "Latest" for the latest available recovery point. DPM could then first recover the server to 11:45. Then—provided administrators have followed the best practice of maintaining the databases on one volume and the transaction logs on a different volume—it can automatically play forward any surviving transaction logs from 11:46 up through 11:58. When this process is complete, the database could be within a single transaction of its state when the hard drive failed.

Using the DPM 2007 express full and transaction replication features together offers multiple advantages. Instead of needing one tool for nightly tape backup and then one or more "better than nightly" disk-to-disk technologies, DPM enables a single platform to first protect data on disk and then offload that data to tape. This approach helps simplify management by using a single agent on the production server, a unified replication engine from production server to protection server, and a consistent interface regardless of media type. In addition, rather than requiring different application-specific agents and protection features, DPM provides a unified system to protect Microsoft Exchange, SQL Server, SharePoint, and Virtual Server data—meaning that if administrators repurpose a SQL Server system to run Virtual Server, they would not have to waste a SQL Server–specific backup agent or purchase an additional Virtual Server–specific backup agent.

PROTECTING CRITICAL DATA FOR KEY MICROSOFT WORKLOADS

Dell and Microsoft are working together to create unified disk-to-disk-to-tape Data Protection Solutions based on Dell™ PowerVault™ hardware and Microsoft System Center Data Protection Manager 2007 software. Designed to be cost-effective and easy to deploy and manage, these systems can provide flexible, simplified data protection in environments built on Dell hardware and Microsoft software.

Jason Buffington is the senior technical product manager for Data Protection Manager and other storage solutions at Microsoft. He has more than 16 years of experience in data protection and business continuity, and studied Computer Science at Texas A&M University.

Sanjeet Singh is a global product marketing manager in the Dell Enterprise Storage Group. He has eight years of experience in developing and delivering business-critical technologies, including databases and data protection. Sanjeet has an M.S. in Computer Engineering from Purdue University and an M.B.A. from the University of Texas.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Singh, Sanjeet, and Jason Buffington. "Safeguarding Data with Dell PowerVault Data Protection Solutions." *Dell Power Solutions*, February 2008. DELL.COM/Downloads/Global/Power/ps1q08-20070484-Singh.pdf

Dell storage:
DELL.COM/Storage

Microsoft System Center Data Protection Manager 2007:
www.microsoft.com/dpm

COST-EFFECTIVE ARCHIVING WITH THE DELL POWERVAULT RD1000

By Jeff Boles

The Dell™ PowerVault™ RD1000 removable disk drive provides a flexible, cost-effective way for small and medium businesses to meet data growth and long-term archiving requirements—designed to provide both higher performance and simpler management than traditional tape media while being more portable and durable than traditional external hard drives.

Related Categories:

Backup

Backup, recovery, and archiving (BURA)

Dell PowerVault storage

Storage

Visit DELL.COM/PowerSolutions for the complete category index.

Small and medium businesses (SMBs) are often challenged by tremendous data growth and seemingly ever-increasing demands for data retention. IT managers, for example, face a litany of regulations and requirements—such as those of the Sarbanes-Oxley Act, Federal Information Security Act, and Health Insurance Portability and Accountability Act (HIPAA)—that prescribe how they must archive and maintain aging data. Facing these rapidly increasing requirements, these organizations need storage that combines high performance, long-term data integrity, and ease of use in a single system. Unfortunately, traditional low-end tape and hard drive solutions typically force them to compromise on at least one of these needs.

Advances in removable disk drive technology have now made high-performance, reliable, simple-to-use storage significantly easier to obtain than it has been in the past. The Dell PowerVault RD1000 removable disk drive can provide multiple advantages for SMBs seeking to meet the challenges of data growth and retention requirements by delivering high performance, simple management, optimal portability, and exceptional durability in a single cost-effective storage system.[1]

TRADITIONAL ARCHIVING AND THE DELL POWERVAULT RD1000

For SMBs, traditional approaches to long-term data retention can carry a variety of disadvantages. For example, many archiving applications continuously add and delete data from long-term storage media such as tape, as new data is captured and old data ages beyond its specified retention period. Over time, this process can lead to large amounts of non-contiguous storage or gaps in utilization, resulting in poor performance, fragmentation, and decreased capacity—and making it difficult for administrators to recover unused space and maintain the required levels of performance.

Built in a 3.5-inch form factor using industry-standard 2.5-inch removable hard drives, the Dell PowerVault RD1000 incorporates a hardened case designed to tolerate a wide range of physical forces (such as dropping and crushing) and environmental conditions (such as high heat and humidity). This removable cartridge can help meet the demands of data growth and retention while avoiding some of the disadvantages of traditional low-end tape and hard drive backups. It combines the advantages of

[1] For an introduction to the PowerVault RD1000, see "Introducing the Dell PowerVault RD1000: A Portable Disk-based Replacement for Traditional Low-End Tape Backup," by Curt Krempin, in *Dell Power Solutions*, February 2007, DELL.COM/Downloads/Global/Power/ps1q07-20070207-RD1000.pdf.

> **"The Dell PowerVault RD1000 removable disk drive can provide multiple advantages for SMBs seeking to meet the challenges of data growth and retention requirements."**

both tape and disk media, is easy to use, has few parts that can fail or wear out, and can handle many cycles of removal and insertion. By combining removability and high-performance disk, the PowerVault RD1000 provides high sequential and random I/O performance, simple implementation and management, scalability, ongoing capacity optimization, and long-term integrity.

The ruggedized removable drive cartridges used by the PowerVault RD1000 can be accessed through the PowerVault RD1000 dock, which is designed for simplicity, contains no read/write mechanisms, and utilizes a minimal number of components to help maximize its life span. The Serial ATA (SATA) II removable drive cartridges—available in 80 GB, 120 GB, 160 GB, and 300 GB capacities—are designed not only for easy off-site storage, but also for capacity scaling similar to tape by allowing enterprises to build and add to a pool of rotated cartridges. This removability also helps minimize power consumption compared with always-on near-line disk systems; for more information, see the "Power conservation with the Dell PowerVault RD1000" sidebar in this article.

KEY ADVANTAGES OF THE DELL POWERVAULT RD1000

Both traditional SMB tape and hard drive solutions for long-term storage require compromises. In comparison with these solutions, the PowerVault RD1000 can provide higher performance, equivalent or better removability and scalability, better cost-effectiveness, and higher durability.

High performance. Dynamic access to traditional tape backups can be extremely slow, and even basic backup performance with tape can suffer when backup clients are unable to sustain consistent throughput rates to help avoid tape stops and starts. PowerVault RD1000 disk cartridges are designed to deliver high levels of performance for both sequential and random I/O—including data transfer speeds of up to 45 MB/sec—that are well suited for the dynamic nature of near-line archival storage and retrieval.

Removability and scalability. The PowerVault RD1000 is designed to exceed the performance of traditional low-end tape while matching its advantages in removability and scalability. The removable cartridges provide SMBs with quickly accessible near-line archival storage and protected backup storage in a single system that can be easily expanded simply by adding cartridges. Administrators can also easily archive and protect data offline without performing complex data migrations or managing multiple types of media.

Cost-effectiveness. Over time, the cost of replacing worn tape can far exceed the initial cost of the tape system itself. PowerVault RD1000 disk cartridges are designed to handle far more read/write cycles than typical tape media can, and typically should not require replacement over a normal archiving life span—helping significantly reduce the total cost of ownership for data retention.

POWER CONSERVATION WITH THE DELL POWERVAULT RD1000

Energy efficiency has become increasingly important to help control power and cooling costs in enterprise IT environments, and backup and archiving systems are no exception. The Dell PowerVault RD1000 removable disk drive is designed to balance high performance and durable archival storage with efficient power conservation. For example, in addition to enabling administrators to remove the cartridges for long-term archiving and storage, these cartridges are designed to rapidly spin down during periods of inactivity even when plugged in. This feature helps reduce not only the power consumption that other disk-based storage devices typical require to continue running, but also the daily wear on drive components, helping extend the system's life span and further reduce total cost of ownership.

—Curt Krempin
Product marketing manager,
Dell Enterprise Product Group

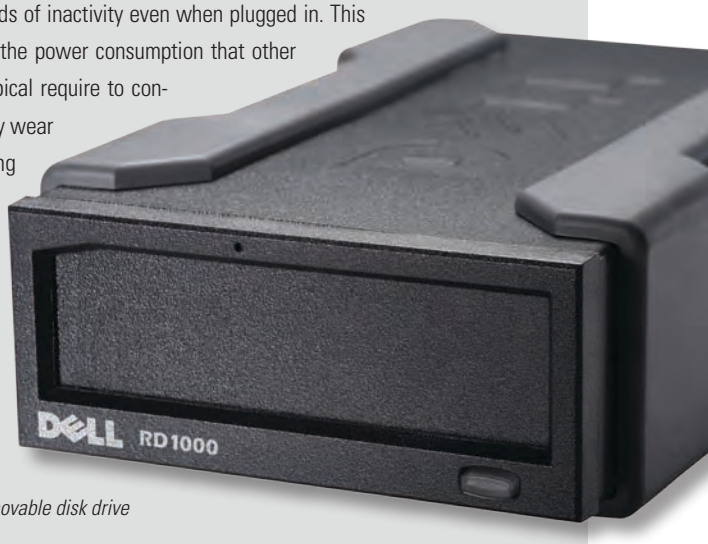

Dell PowerVault RD1000 removable disk drive

> **"The Dell PowerVault RD1000 offers high-performance, durable, easy-to-use disk-based storage in a single system. Its versatility and long life expectancy make it well suited to help SMBs meet the challenges of data growth and long-term data retention in a flexible, cost-effective way."**

Durability. Durability is critical to long-term storage media, and PowerVault RD1000 storage is designed to outlast typical enterprise data retention requirements. Best practices for digital preservation recommend that enterprises regularly migrate media to reduce the risk of data loss from both media failure and access mechanism obsolescence. The PowerVault RD1000, however, helps avoid these risks by providing removable backup and archival storage using durable, highly protected disk cartridges that integrate both the access mechanism and the media itself. Key features include the following:

- Special-purpose 2.5-inch drive cartridges that secure drive heads by ramp loading them at the end of activity, helping remove a source of accidental media and head collisions that can damage traditional drives
- Inactive-state spin down to help avoid wear on drive components
- Casing and enclosure designed for high insertion rates
- Ruggedized shell designed to withstand a three-foot drop onto concrete as well as prolonged exposure to heat, dust, electrostatic discharge, and other environmental contaminants

To help validate these product specifications, in July 2007 Percept Technology Labs, an independent testing and consulting company, tested the PowerVault RD1000 to evaluate how long its cartridges might typically last during real-world use. These tests were designed to provide rough estimates of the expected service life of the drives under specific storage and handling conditions.

The Percept team tested a total of 80 PowerVault RD1000 cartridges. After writing data to each drive to completely utilize its capacity, verifying the data, and establishing a baseline for performance, they separated the cartridges into five groups. Each group was placed into its own environmental chamber set at high temperature and humidity levels, with temperatures in different chambers ranging from 140°F to 176°F and humidity levels ranging from 10 percent to 85 percent. At specified intervals of between 336 and 1,000 hours, the test team removed the cartridges and checked the data for error. Failed drives—defined as those returning an unrecoverable read error when reading any data on the disk—were removed from the test, with the rest returned to the environment chamber for another monitoring interval. This process was repeated until the total duration for each chamber was reached, a period of between 2,000 and 4,000 hours.

The Percept team next plotted the test results in ReliaSoft ALTA, a software package designed to calculate predicted life spans based on data from accelerated conditions such as those used in the environmental chambers. The results indicated that a PowerVault RD1000 cartridge stored in moderately controlled conditions (78°F with 95 percent noncondensing relative humidity) can provide a potential life expectancy of 30 years or more. In stringently controlled conditions (68°F with 30 percent relative humidity), the results indicated that this potential life expectancy can increase to 100 years or more—well over the typical requirements for long-term enterprise data storage.

HIGH-PERFORMANCE, DURABLE, COST-EFFECTIVE ARCHIVING

The Dell PowerVault RD1000 is designed to provide significant advantages over both traditional low-end tape and traditional external hard drives, offering high-performance, durable, easy-to-use disk-based storage in a single system. Its versatility and long life expectancy make it well suited to help SMBs meet the challenges of data growth and long-term data retention in a flexible, cost-effective way.

Jeff Boles is an analyst and the director of technology validation services at the Taneja Group, an analysis and consulting organization focused on the storage and server industry. He has more than 20 years of IT experience in data protection, storage systems, networking, and servers, as well as in product development, validation, and strategy.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerVault RD1000:
DELL.COM/RD1000

NO GUESSWORK

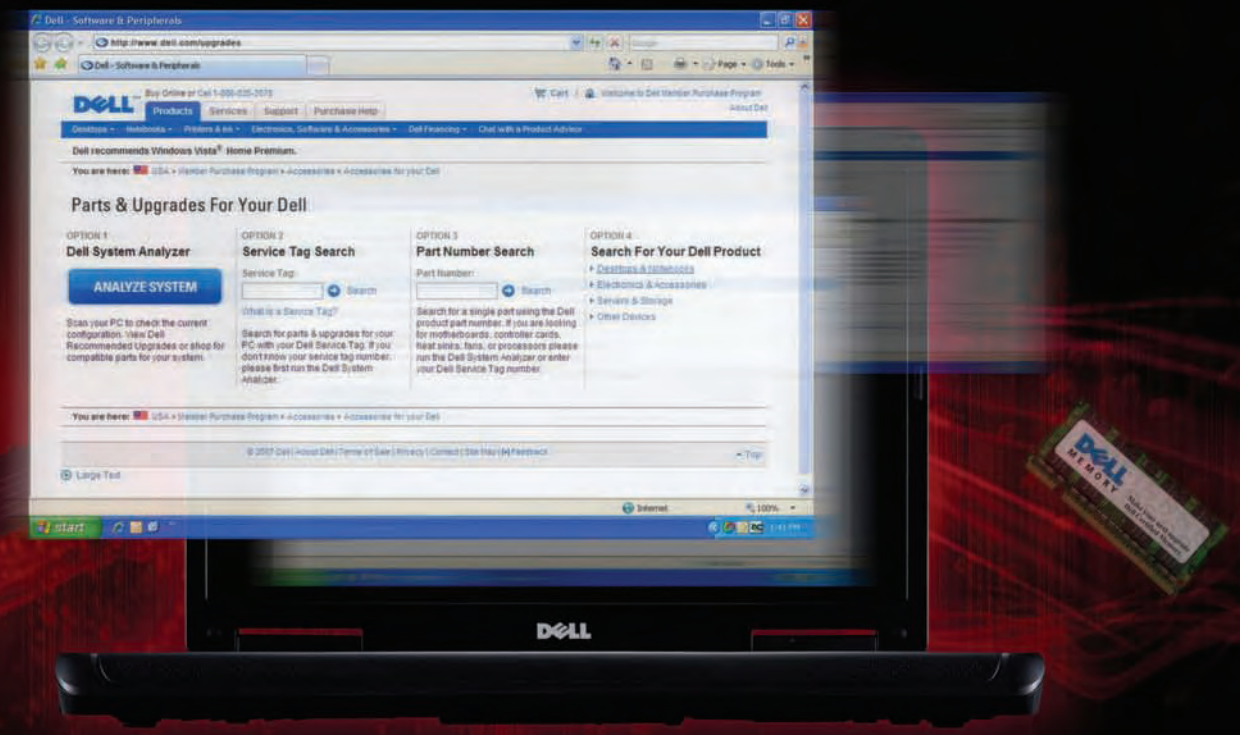

AN EXCITING, INNOVATIVE ONLINE SOURCE FOR COMPATIBLE PARTS AND UPGRADES.

Memory. Hard drives. Power cords. Dell has made it easier than ever for small businesses to keep their systems updated and optimized. Introducing the **"Parts For Your Dell"** website – an easy-to-use resource for parts and upgrades that apply specifically to your Dell systems. Here, our online tool scans your PC and gives you the recommendations you need to help find compatible replacement parts and upgrades. The only thing you won't find is – guesswork.

SIMPLIFY FINDING PARTS AND UPGRADES AT DELL.COM/Upgrades

FACILITATING MICROSOFT WINDOWS VISTA MIGRATION THROUGH APPLICATION VIRTUALIZATION

By Coby Gurr

Many enterprises are considering migrating to the Microsoft® Windows Vista® OS, but have postponed doing so because of concerns about application compatibility. By enabling the creation of self-contained virtualized applications, LANDesk® Application Virtualization can help minimize risk, save time, increase user satisfaction, and reduce costs for support and regression testing during this type of migration.

Related Categories:

LANDesk

Systems management

Virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

As enterprises plan their next hardware refresh of Dell™ desktops and notebooks, they may be strongly considering whether to deploy the Microsoft Windows Vista OS or continue using a previous version such as the Microsoft Windows® XP OS. Although some have already migrated to Windows Vista, a significant percentage of IT professionals are currently still only considering or planning this migration, primarily to take advantage of the latest security enhancements.

The slow adoption rate results from multiple factors, but one of the most critical is uncertainty regarding application compatibility: enterprises cannot seriously consider a migration to Windows Vista if their critical applications cannot run in that OS. To help address this problem, they might try to convince their software vendors to port their legacy applications to Windows Vista. This approach might meet with some success, but it typically requires a significant waiting period.

Enterprises might consider implementing OS virtualization to help solve the problem—running two different operating systems on each computer, with Windows Vista in one virtual machine and the legacy OS with its legacy applications in another. However, by splitting resources between two operating systems, this approach can double both the required hardware resources and administrative tasks, forcing administrators to maintain and patch two different operating systems on every computer instead of just one.

For many organizations, neither of these approaches provides an acceptable resolution to the application incompatibilities that can plague OS migrations such as a move to Windows Vista. However, application virtualization through software such as LANDesk Application Virtualization can help solve this problem without requiring application porting or OS virtualization. This approach can help eliminate application incompatibility concerns, greatly simplify application deployment time and effort, and enhance enterprise security.

UNDERSTANDING APPLICATION VIRTUALIZATION

The term *virtualization* can mean different things to different people. Many administrators think of virtualization in terms of *OS virtualization*, such as a VMware® virtualization platform running on Dell PowerEdge™ servers that allows each server to host multiple operating systems. Others may think of virtualization in terms of *hardware virtualization*, such as Intel® Virtualization Technology or AMD Virtualization™ technology in processors designed to

> **"Application virtualization helps solve the problem of application incompatibility in Windows Vista—and other operating systems—by adding a layer of isolation between the OS and the application."**

increase the efficiencies and performance of OS virtualization environments.

OS virtualization and hardware virtualization focus primarily on enabling physical servers and computers to effectively host multiple virtual OS environments. *Application virtualization* has a very different focus: helping eliminate application incompatibilities. It can accomplish this goal by freeing an application from dependencies on physical hardware or the host OS, enabling it to execute in its own isolated virtual environment (see Figure 1). While initial application virtualization efforts focused primarily on eliminating application conflicts, current solutions such as LANDesk Application Virtualization can do much more than that.

MINIMIZING OS MIGRATION RISKS AND CHALLENGES

A key advantage of application virtualization is its ability to minimize OS migration risks and challenges. For enterprises that have been seriously considering a move to Windows Vista, but have been hesitating because of concerns about application compatibility, application virtualization can provide a solution.

Application virtualization can help solve the problem of application incompatibility in Windows Vista—and other operating systems—by adding a layer of isolation between the OS and the application. Different application virtualization solutions achieve this layer of isolation in different ways. The most effective and easiest to manage are typically those that package an application into a virtualized version of itself without requiring the installation of special drivers or extra infrastructure on the client or on a back-end server.

Application virtualization solutions such as LANDesk Application Virtualization can create self-contained executables that include their own small user-mode OS kernel, similar to a runtime environment. The executable for this virtualized application typically includes a small virtual registry, a small virtual file system, and other components the application needs to run, allowing it to function independently in isolation from the host OS.

This level of OS independence and isolation also enables the virtualized application to dynamically remap its virtual file system and registry locations for simple, rapid migration from one OS to another. For example, this functionality would allow a single virtualized version of 2007 Microsoft Office to run on Windows 2000, Windows XP, or Windows Vista without requiring changes to the executable file or the host.

To help achieve the level of application isolation necessary to resolve application incompatibilities between operating systems, the goal should be to take advantage of an application virtualization solution that allows enterprises to seamlessly run a self-contained version of their virtualized application across multiple OS versions without requiring any drivers, reboots, administrator access, additional infrastructure, or special modifications to the OS environment.

SIMPLIFYING APPLICATION DEPLOYMENT

In addition to facilitating OS migration, application virtualization solutions such as LANDesk Application Virtualization can provide significant advantages for application deployment—a task that can require

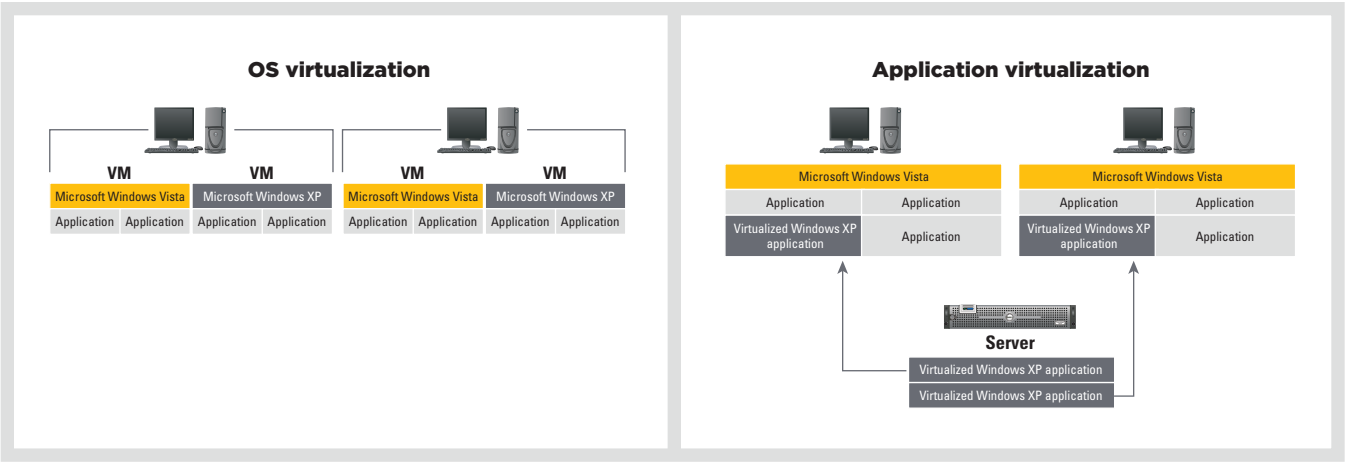

Figure 1. *Retaining legacy applications using OS virtualization compared with application virtualization*

"Application virtualization gives administrators a way to create virtualized applications that help minimize risk, save time, increase user satisfaction, and reduce costs for support and regression testing."

an enormous investment in time and administrative resources. Traditional application deployment carries the risk of interfering with other applications or with the overall environment, potentially reducing user productivity and satisfaction and requiring administrators to field numerous complaints and, if necessary, work to roll back the environment to its pre-deployment state and then perform regression tests to fix problems created by the new applications.

Application virtualization gives administrators a way to create virtualized applications that help minimize risk, save time, increase user satisfaction, and reduce costs for support and regression testing. For example, many application virtualization solutions allow administrators to make applications available or unavailable on demand, helping eliminate the need to carry out a long installation process. If an application experiences a problem, administrators can also easily reset it to its original state without having to completely reinstall it.

Different application virtualization solutions employ different methods to help simplify application deployment, but one method is typically the simplest to execute and manage while also requiring the least amount of infrastructure overhead: packaging everything that an application needs into a single self-contained executable. Instead of installing the application, administrators can simply copy this executable to a hard disk, USB drive, CD, or network share and run it from there. This approach allows them to take a large application like 2007 Microsoft Office—which can contain thousands of registry keys and other files—and compress it into a single, manageable executable that does not require installation.

When the virtualized application executable runs, it transparently streams the compressed blocks of data required to run the application to the end user's desktop. It can automatically start the necessary virtual services, such as loading dynamic-link library (DLL) and data files and accessing registry entries. This streaming process enables applications such as Microsoft Word to load and run extremely quickly—often even faster than if they ran as a traditional local installation.

This type of virtualized application, running as a self-contained unit in isolation from the OS and other applications, helps minimize the risk of application conflicts or other problems that may negatively affect the environment. Solutions such as LANDesk Application Virtualization can package the application so that it does not modify registry keys in the host OS or proliferate DLL files throughout the file system, which can create havoc with the system or other applications. They can also enable two versions of the same application to run on the same computer without conflicts, even where both virtualized and non-virtualized versions exist.

Virtualized applications that run in this way often employ a *sandboxing* technique. When using sandboxing, in addition to running in its own self-contained and isolated virtual environment, the solution's virtual file system helps further protect the host from potential conflicts by redirecting changes made by or to the application to an isolated per-user, per-application sandbox directory. For example, a sandbox directory for a virtualized version of a Web browser might contain bookmarks and user preferences, as well as automated updates to the browser application itself. Sandboxing these application components helps prevent them from affecting or conflicting with other applications or aspects of the host OS. As a result, this approach helps greatly reduce the regression testing that administrators must perform following the release of a patch or a new version of an application.

Another advantage of sandboxing is that it can enable administrators to avoid repackaging and redeploying a virtualized application every time an update arrives. Instead, a simple update to the application sandbox updates the application itself. In addition, administrators can easily roll back a virtualized application to its original state on a per-user basis by deleting the user's sandbox for that application.

ENHANCING ENTERPRISE SECURITY

Solutions such as LANDesk Application Virtualization, which can run applications exclusively in user mode without kernel-mode code or device drivers, help make it virtually impossible for a virtualized application to crash a system, while also helping prevent kernel-mode calls from violating local security policies. This

"Solutions such as LANDesk Application Virtualization help make it virtually impossible for a virtualized application to crash a system."

approach also means that virtualized applications can run without requiring administrative privileges. As a result, these types of virtualized applications can run in user mode on locked-down computers without requiring the installation of device drivers.

In addition, administrators can further increase overall system stability by using sandboxing to limit and contain user-specific and application-specific changes. Combining these features allows administrators to easily maintain secure, clean, stable user desktop environments.

AVOIDING PROPRIETARY INFRASTRUCTURE REQUIREMENTS

While application virtualization can offer multiple advantages for application isolation, deployment, stability, and management, not all solutions are equal. Some may fail to provide complete application isolation, allowing potential application conflicts. Others may have limited OS support that does not include support for Windows Vista. Many can add significant complexity to an environment by requiring preinstalled proprietary software infrastructures on the client and server. To add a layer of isolation between the OS and application, some may require a proprietary client on the host computer, device drivers, or a back-end server, which adds to the required infrastructure and can significantly increase cost and complexity.

In these examples, administrators end up managing and maintaining multiple clients and servers for each of their desktop environments. In addition, when their application virtualization vendor updates its virtualization client, if any one of their applications need that updated client, they typically must repackage and redeploy all of their applications to use that updated client, because many solutions do not support using more than one client on a computer at a time.

On the other hand, self-contained virtualized applications such as those created with LANDesk Application Virtualization do not require a proprietary client, special drivers, or back-end servers. These self-contained executables also have a zero-execution footprint, allowing the virtualized applications to run from a network share, USB drive, or CD without requiring preinstalled software on the client computer or additional infrastructure. Administrators can even deploy virtualized applications to reliable high-speed network subnets located at branch offices, allowing users at those offices to stream the compressed applications on demand over the network directly into their computer's memory, without requiring local installation or caching. Avoiding the need for front-end clients, drivers, or back-end server components helps significantly reduce infrastructure costs and simplify deployment and management.

> **"LANDesk Application Virtualization sets the standard for application virtualization by delivering application isolation without the costs and complexities of additional front-end and back-end proprietary infrastructures."**

SETTING THE STANDARD FOR APPLICATION VIRTUALIZATION

As the demand for seamless interoperation of multiple applications continues to grow, so does the need for robust application virtualization solutions. The advantages of seamless integration and application isolation provided by application virtualization can help administrators minimize OS migration risks, simplify application deployment and change management, enhance management of and control over software access, enable rapid deployment of software framework technologies, reduce multi-application regression testing costs, and boost IT responsiveness when changing and updating applications.

LANDesk Application Virtualization sets the standard for application virtualization by delivering application isolation without the costs and complexities of additional front-end and back-end proprietary infrastructures. It can package both simple and complex applications into virtual environments within a single self-contained executable, then transparently merge a virtual system environment with a real system environment to provide true application isolation, providing enterprises with the tools they need to seamlessly run multiple applications without the hassles of application deployment or the worries of application incompatibilities. LANDesk Application Virtualization provides a clientless application virtualization solution that can enable organizations to easily create, maintain, and control secure, clean, stable user desktop environments throughout their enterprise.

Coby Gurr is a business line manager for systems management at LANDesk.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

LANDesk Application Virtualization:
www.landesk.com/products/virtualization/index.aspx

ACCELERATING APPLICATION TRAFFIC WITH INTEL 10 GIGABIT ETHERNET SERVER ADAPTERS

By Jordan Plawner
Travis Vigil

By using Intel® 10 Gigabit Ethernet server adapters in Dell™ PowerEdge™ servers and unified storage systems such as the Dell PowerVault™ NX1950, organizations can achieve the performance required for high-bandwidth applications and meet the demand for increased storage capacity while moving toward network convergence on a simplified Ethernet infrastructure.

Related Categories:
Dell PowerEdge servers
Dell PowerVault storage
Ethernet
Intel
Intel networking
Internet SCSI (iSCSI)
Network interface card (NIC)

Visit DELL.COM/PowerSolutions for the complete category index.

The growing demand for high-bandwidth, real-time applications by enterprises and consumers can require IT departments to make significant changes in their IT infrastructures. The increasing use of radio-frequency identification (RFID) technologies, decision support software, data mining, virtualized applications, and multimedia content is already pushing internal networks to their limits. At the same time, the rising adoption of Internet software distribution, high-volume data transactions, software as a service, and communications applications is creating additional infrastructure burdens. These trends are likely to accelerate. Internet communications applications, for example, are being used with increasing frequency among both home and business users: a recent survey found that three out of four U.S. Internet users streamed video online in May 2007.[1] Organizations need to find new ways to deliver more data to more people in more locations than they can right now.

Supporting high-bandwidth, real-time applications requires IT groups to add not only bandwidth to their infrastructure, but also storage capacity. Organizations are already facing the challenges of scaling storage capacity to accommodate the growing volume of health care records, financial data, and customer information required for data mining, customer relations, and e-commerce. The creation, distribution, and archiving of high-resolution content, including the media that people are increasingly streaming from the Internet, will likely require enterprises to scale capacity further. One market research firm recently predicted an eightfold increase in new storage capacity shipments per year between 2005 and 2012.[2]

As organizations begin to modify their infrastructures to accommodate these trends, they will likely need to do so without significantly adding to the cost or complexity of IT. Enterprises will also likely need to increase storage capacity without major expenditures in hardware acquisition or management, and to prepare their networks for increased throughput without adding complexity through the management of multiple networks.

SUPPORTING HIGH-BANDWIDTH, REAL-TIME APPLICATIONS

By combining the capabilities of a network attached storage (NAS) system and Internet SCSI (iSCSI), unified storage systems such as the Dell PowerVault NX1950 can provide organizations with cost-effective ways to scale storage capacity and address the challenges of supporting high-bandwidth, real-time

[1] "3 Out of 4 U.S. Internet Users Streamed Video Online in May," by comScore, July 17, 2007, www.comscore.com/press/release.asp?press=1529.
[2] "2007 Digital Storage in Entertainment Creation and Distribution," by Coughlin Associates, available at www.tomcoughlin.com/techpapers.htm.

applications. Unified systems can help deliver the performance required for real-time applications while also helping simplify storage administration and reduce the cost and complexity of adding storage capacity.

Unified storage systems can help optimize information processing for real-time applications by removing processing burdens from the host server. The PowerVault NX1950, for example, stores both file and application data in a single system. While client systems can access file data using Common Internet File System (CIFS) and Network File System (NFS) standards, servers can access application data through the iSCSI protocol. With iSCSI technology, blocks of information are transmitted over an IP network between the server and storage system. Application workloads are managed by host servers, and storage workloads are handled by the PowerVault NX1950. This division of labor enables servers to devote more resources to real-time applications than they can with other types of storage.

A unified storage system can also help simplify storage administration and reduce costs. While Fibre Channel storage systems such as Dell/EMC CX series storage area networks (SANs) offer the performance, reliability, and functionality required by many enterprise functions, the cost of acquiring and maintaining Fibre Channel storage places these systems beyond the reach of some organizations. A unified storage system that employs iSCSI can offer a cost-effective, simplified alternative to a Fibre Channel system. Administrators of iSCSI systems can use standard, cost-effective Ethernet switches, routers, cables, and network interface cards (NICs). Consequently, the overall cost of acquiring and maintaining an iSCSI system can be less than that of a Fibre Channel system with comparable capacity. Because many organizations already manage an Ethernet network, moving from a combination of Fibre Channel and Ethernet networks to a single Ethernet network can also reduce the complexity of network administration.

At the same time, the use of iSCSI storage can bring added advantages over Fibre Channel systems. For example, the use of Ethernet and TCP/IP technology by iSCSI storage can help extend a system's geographic flexibility beyond the limits of a Fibre Channel system. Consequently, organizations can incorporate iSCSI storage into storage centralization and disaster recovery strategies. iSCSI systems can also help organizations realize the benefits of using shared storage with virtualization without incurring the cost or complexity of deploying a Fibre Channel SAN.

In addition, a unified storage system that uses iSCSI can provide administrators with a cost-effective way to capitalize on boot capabilities offered by SANs. By maintaining OS images on a shared storage system, administrators can streamline server provisioning, easily deploy upgrades and patches, and accelerate disaster recovery. A unified system such as the PowerVault NX1950 can provide these advantages without the infrastructure costs of a Fibre Channel–based SAN.

BUILDING TOMORROW'S INFRASTRUCTURE

While many enterprises find that the Gigabit Ethernet bandwidth used by typical unified systems is sufficient for running current applications, other organizations are adopting servers with Gigabit Ethernet LAN on Motherboard chips in conjunction with multi-port adapter cards to take advantage of multi-gigabit speeds. Some of today's streaming media, data backup, and decision support applications require that increased throughput.

The introduction of Intel 10 Gigabit Ethernet (10GbE) server adapters enables organizations to use unified storage to meet the needs of today's throughput-intensive applications while helping build a foundation for future growth. By using a variety of technologies to distribute packets across multiple processor cores, these adapters provide more than additional bandwidth. They also help accelerate processing for real-time applications.

OPTIMIZING PERFORMANCE WITH MULTI-CORE PROCESSORS

To provide the level of performance required for high-bandwidth, real-time applications, an infrastructure needs to deliver not only the necessary bandwidth capacity but also the means to accelerate the processing of network data packets. A network packet undergoes considerable processing between the time it arrives at the physical controller to the moment its payload is delivered to the application. On many of today's systems, this processing is performed by a single

INTEL ADAPTER TECHNICAL SPECIFICATIONS

Intel 10 Gigabit Ethernet (10GbE) adapters provide comprehensive support for multiple IEEE standards. Three types of fiber-optic adapter are currently available:

- Single-port SR (short reach: up to 300 m)
- Dual-port SR
- Single-port LR (long reach: up to 10 km)

Intel 10 Gigabit XF SR dual-port server adapter

Key technical specifications of Intel 10GbE adapters include the following:

- 32 transmit queues and 64 receive queues per port
- 16 virtual machine device queues per port
- 512 KB receive buffer (divisible into eight individual packet buffers)
- 320 KB transmit buffer (also divisible into eight individual packet buffers)

processor. But multiple I/O requests to a single processor can create a bottleneck while leaving its other resources available for I/O processing.

This problem cannot be solved simply by adding bandwidth. In fact, adding network controllers only amplifies the problem: additional packets increase both processor congestion and latency. The solution lies in the ability to distribute data packet processing across multiple cores working in parallel. By capitalizing on the processing capabilities of these cores, Intel server adapters distribute the load and help lower the processing time for individual network packets.

Intel 10GbE adapters are specifically designed for multi-core systems and the latest OS features. They can use up to 32 separate hardware NIC queues for transmission and up to 64 queues for receiving—all of which can be mapped to 16 separate processor cores.

The ability to direct streams to different processor cores can provide important advantages in virtualized server environments. By using Intel 10GbE adapters, virtual machines (VMs) hosted by hypervisors that emulate network controllers can rely on a dedicated network stream handled by a single processor core. When multiple VMs are in use, they can share the controller ports while taking advantage of their own privately processed packet stream. This approach enables significant enhancements in the performance of virtualized applications.

As illustrated in Figure 1, Intel 10GbE adapters use multiple technologies that work together to distribute packets across multiple processor cores and help accelerate processing:

- **Multiple descriptor queues:** To spread the workload across multiple processor cores, network traffic streams are divided into queues through receive-side scaling (RSS), filtering based on the Media Access Control (MAC) address, or the use of virtual LAN tags. The packet queues can be accessed by driver threads running on different processor cores so that multiple cores can process network packets in parallel. The Intel 10GbE adapters' 32 transmit queues and 64 receive queues can be mapped to up to 16 processor cores. For servers equipped with several multi-core processors, multiple descriptor queues can facilitate powerful load-balancing functionality.
- **Receive-side scaling:** To determine which queue to use for incoming packets, network adapters residing on systems using the Microsoft® Windows Server® 2003 or Windows Vista® operating systems can use RSS, which directs packets to different queues without the need for reordering. (On systems using the Linux® OS, this technology is known as scalable I/O). RSS is intelligent in its distribution of packet processing, and is also programmable. Consequently, network controllers with multiple queues can efficiently direct multiple TCP/IP streams to different processor cores for handling.
- **Message-Signaled Interrupts Extended (MSI-X):** As part of the PCI Express (PCIe) standard, MSI-X technology facilitates efficient communication between queues and specific processor cores, enabling the network adapter hardware to direct an interrupt to the designated core when network packets are placed in the processor queue. In this way, MSI-X provides an important enhancement to standard MSI, which could pass interrupts only to a single processor core. Intel 10GbE adapters give each queue its own set of MSI-X controllable interrupt vectors to help provide efficient packet management and fine-tune the processor load.
- **Intel Virtual Machine Device Queues (VMDq):** VMDq provides multiple hardware queues and offload features to help reduce the software overhead associated with sharing a single network controller between multiple VMs. Previously, a network switch emulated by the virtualization software sorted and routed packets individually to VMs. That process typically introduced significant delays in the network packet processing. With VMDq, individual hardware queues are associated with the simulated network interfaces of running VMs. The network controller itself performs the routing of received packets, helping substantially reduce overhead. VMDq is also used on outbound VM packets to provide transmit fairness and to help avoid a single VM blocking access to the controller.

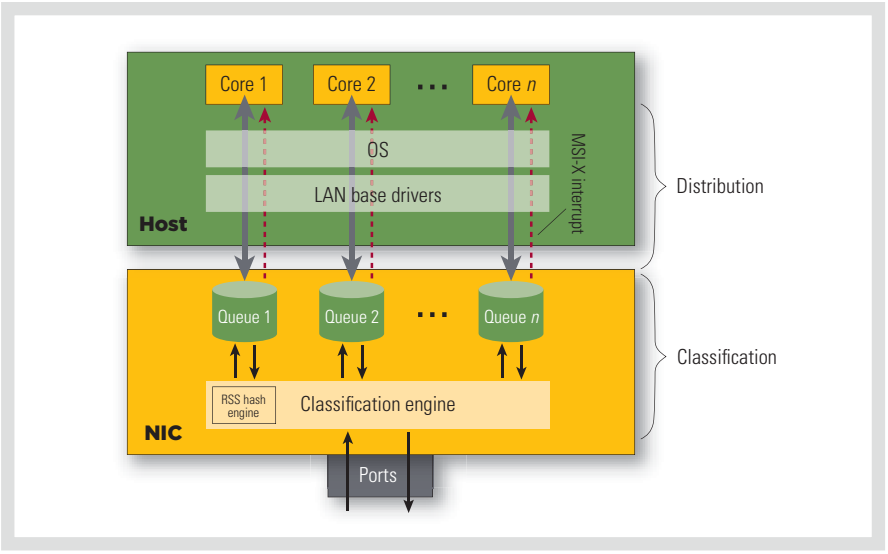

Figure 1. *Intel 10 Gigabit Ethernet adapters sort, group, and direct network packets across multiple processor cores to help reduce network bottlenecks*

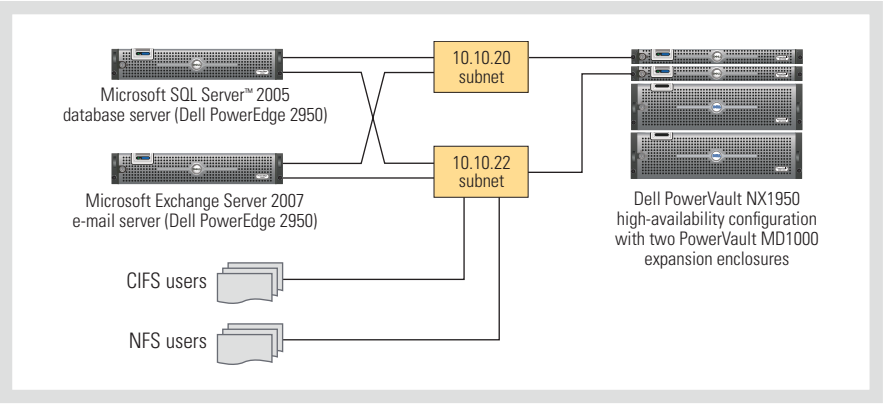

Figure 2. *Dell PowerEdge 2950 servers connected to Dell PowerVault NX1950 storage using Intel 10 Gigabit Ethernet adapters help deliver a cost-effective file-sharing solution that can support high-bandwidth enterprise applications*

- **Intel I/O Acceleration Technology:** Intel 10GbE adapters also include Intel I/O Acceleration Technology, which helps optimize bandwidth by redirecting header processing and speeding up memory access to packet components.

For more information on the technical aspects of the Intel 10GbE adapter hardware, see the "Intel adapter technical specifications" sidebar in this article.

EXPANDING THE DELL 10 GIGABIT ETHERNET ECOSYSTEM

Intel 10GbE adapters are validated for the Dell PowerVault NX1950 storage system and for Dell PowerEdge servers, while Dell PowerConnect™ 6200 series switches provide the connectivity for a Dell 10GbE infrastructure.

Dell PowerVault NX1950 storage system

The Dell PowerVault NX1950 network storage system offers a range of advantages for organizations that need a cost-effective way to address the growing demand for storage capacity and to support high-bandwidth, real-time applications (see Figure 2). While the unified design helps optimize the processing power of host servers, the simple, cost-effective scalability of the PowerVault NX1950 can help address rapidly rising demand for storage capacity. The integrated PowerVault MD3000 modular disk storage array holds up to 15 Serial Attached SCSI (SAS) drives for up to 4.5 TB of storage capacity. Administrators can also easily add up to two PowerVault MD1000 external storage expansion arrays, for a total of 45 drives and up to 13.5 TB of total storage capacity in a single 3U enclosure. In many cases, these SAS drives are less expensive than comparable Fibre Channel disks.

The PowerVault NX1950 also helps simplify storage deployment and management. While setup wizards help administrators configure the system quickly, the included Microsoft Windows® Unified Data Storage Server 2003 OS provides an integrated console in a familiar management environment for administrators to conduct essential tasks.

Dell PowerEdge servers

Intel 10GbE adapters are validated and available as an option on ninth-generation Dell PowerEdge servers. By offering a selection of multi-core processors, including dual- and quad-core Intel Xeon® processors on several models, PowerEdge servers can provide the performance required for real-time processes. The Intel adapters add cost-effective iSCSI connectivity while helping optimize the network processing performance of multi-core processors to deliver the performance required for high-bandwidth applications.

Dell PowerConnect 6200 series switches

Advanced Dell PowerConnect 6200 series switches can incorporate 24 or 48 ports of 10/100/1,000 Mbps Ethernet in a 1U form factor, enabling administrators to connect up to 576 servers or clients and up to nearly 3 TB of capacity in a single stack. These switches also support up to four 10GbE uplinks for direct connectivity to 10GbE routers, enterprise backbones, and data centers. The modular design enables administrators to upgrade to advanced stacking or 10GbE when needed.

PREPARING FOR THE FUTURE

Unified storage systems such as the Dell PowerVault NX1950 offer organizations a cost-effective, scalable, easy-to-manage solution for the growing demand for storage capacity, while Intel 10GbE server adapters help increase the viability of iSCSI storage even further for enterprise data centers. With increased bandwidth and technologies designed to accelerate network processing by capitalizing on the capabilities of multi-core processors, Intel 10GbE server adapters can help prepare organizations for today's and tomorrow's high-bandwidth, real-time application requirements.

Jordan Plawner is a product planner and technologist in the Intel LAN Access Division.

Travis Vigil is a product marketing strategist for Dell iSCSI and Dell PowerConnect solutions.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Intel 10GbE server adapters:
www.intel.com/go/adapters

Dell PowerEdge servers:
DELL.COM/PowerEdge

MICROSOFT HPC SOLUTIONS BOOST MANUFACTURING INNOVATION

By Aditya Krishnan

In manufacturing concerns, high-performance computing (HPC) has evolved from a cost center to a strategic asset. The Microsoft® HPC solution stack facilitates seamless access to distributed computing resources—enhancing end-user productivity and performance to accelerate global product innovation cycles.

Related Categories:

Computer-aided engineering (CAE)

High-performance computing (HPC)

Microsoft

Visit DELL.COM/PowerSolutions for the complete category index.

Global manufacturing companies are facing intense competition on many levels. As new markets emerge worldwide and competitive pressure mounts to meet accelerating time-to-market goals, these businesses must reevaluate their entire operational infrastructure. Many organizations are attempting to simplify their supply chains to lower development and manufacturing costs while continuing to deliver high-quality, innovative products in ever-shorter design cycles. All of these factors combine to ratchet up requirements for a high-performance computing (HPC) infrastructure that can support an agile manufacturing organization.

The Microsoft HPC platform is designed to address critical factors that support manufacturing success or failure. In close partnership with Dell and advanced manufacturing and computer-aided engineering (CAE) software providers, Microsoft is delivering on a consolidated solution stack that makes the power of parallel processing seamlessly accessible to analysts, designers, engineers, and other end users while simplifying platform deployment and management for IT architects and administrators.

IDENTIFYING SUCCESS FACTORS FOR MANUFACTURERS

Success for manufacturers in a competitive business environment depends on multiple factors, including increasing operational performance, ensuring regulatory compliance, streamlining product innovation cycles, and accelerating what Microsoft calls *time to insight*.

Increasing operational performance. Manufacturers must continually increase operational performance and optimize processes along multiple dimensions, working to reduce design and manufacturing costs while increasing product quality and rapidly bringing new products to market. While companies can use several strategies to help reduce overall costs—ranging from Six Sigma to lean manufacturing—deploying an advanced HPC software infrastructure in an open, scalable environment can be critical to success.

Streamlining product innovation cycles. Meeting end-user demands and creating competitive differentiation requires companies to innovate ceaselessly and efficiently in new product development. Introducing new, high-quality products and accelerating time-to-market cycles while controlling costs can be key to maintaining a thriving business.

Accelerating time to insight. Intense competition requires manufacturing organizations to bring continually improved, high-quality products to market, with compressed development times and reduced costs. Microsoft refers to the competitive advantage of accelerated design cycles as time to insight—the important

element of evaluating and making critical decisions with maximum efficiency.

Ensuring regulatory compliance. Complying with environmental and safety regulations as well as auditing policies can represent a major cost—one that that becomes more difficult to manage each year. As regulations become increasingly stringent, organizations must embed compliance processes at every stage of design and manufacturing. State-of-the-art software and HPC operating environments can support this requirement by simplifying integration.

> **"Microsoft HPC solutions enable analysts, designers, engineers, and other end users to achieve maximum performance and high productivity while simplifying deployment and management for IT architects and administrators."**

MEETING THE CHALLENGES OF COMPUTER-AIDED SIMULATION

High-end computer-aided simulation, or virtual prototyping, can help global manufacturers meet many innovation-related challenges. By incorporating simulation at every stage of the engineering design process, these organizations can bring high-quality, competitive products to market quickly and efficiently. However, implementing an effective solution requires overcoming a number of additional challenges, to meet the needs of both end users relying on HPC infrastructures—including analysts, designers, and engineers—and the administrators who deploy and manage them.

For end users, for example, lengthy simulation runtimes during product design and evaluation cycles can inhibit decision making, delaying time to insight and ultimately time to market. Similarly, limited hardware and software resources restrict exploration of multiple design options: although large simulation models can yield more accurate predictions than small models, they can be incredibly difficult to run on complex HPC solution stacks assembled from a variety of sources. In addition, the lack of availability and scalability for many HPC applications can prevent end users from easily transferring design and evaluation processes from workstations to clusters, creating significant barriers to a seamless experience.

IT architects and administrators, meanwhile, face their own set of challenges. For example, integrating HPC clusters with existing IT infrastructures can be difficult, both for initial deployment and subsequent management. IT groups may perceive implementing HPC clusters as financially risky, time-consuming, and technically challenging, and may not want to devote the time and resources required to research and size the systems nor allocate the budgets required for new solution stacks in the manufacturing environment. And running HPC infrastructures as silos—separate from the organization's overall Microsoft Windows® OS–based environment—can increase the time, cost, and required skills for HPC infrastructure management.

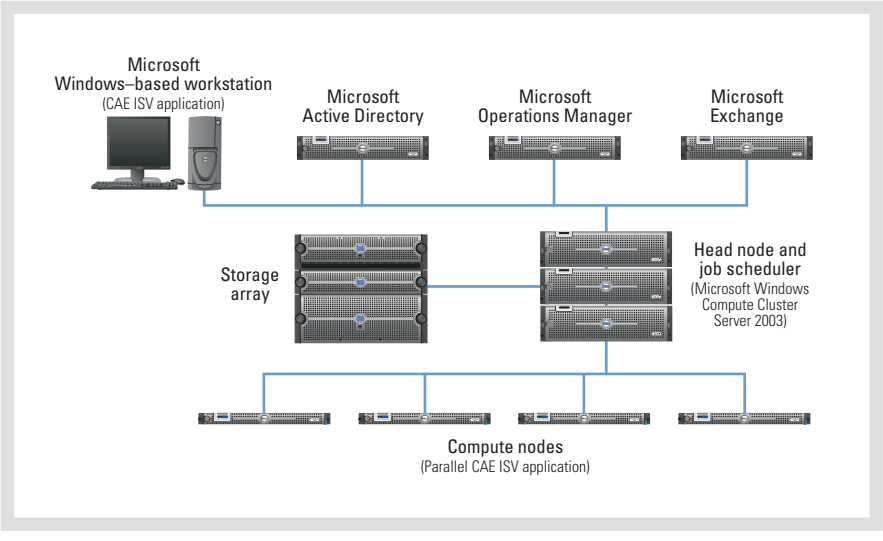

Figure 1. *Integrated Microsoft HPC solutions can help maximize performance and simplify management*

ACCELERATING INNOVATION WITH MICROSOFT HPC SOLUTIONS

Microsoft HPC solutions have been developed with all of these challenges in mind, enabling analysts, designers, engineers, and other end users to achieve maximum performance and high productivity while simplifying deployment and management for IT architects and administrators (see Figure 1). Working closely with the development and integration teams at Dell, leading independent software vendors (ISVs), and other third-party solution providers, Microsoft has designed a flexible HPC platform that enables companies to procure, deploy, and manage HPC cluster infrastructures that run strategic third-party CAE applications and in-house applications on Windows-based HPC clusters, while taking advantage of

existing investments in their overall Windows environment. By using clusters based on Microsoft Windows Compute Cluster Server 2003, organizations can use the applications with a very similar end-user experience, while helping ensure that users can harness the power of HPC clusters to run fast, detailed simulations. For more information on Windows Compute Cluster Server 2003, see the "Microsoft Windows Compute Cluster Server 2003" sidebar in this article.

For analysts, designers, engineers, and other end users, tight desktop integration creates a seamless workflow when scaling from their workstations to clusters. Among other advantages, this integration enables them to create and submit jobs directly from workstations or from within CAE applications, and to monitor the status of jobs running on an HPC cluster while utilizing user-profile contexts that already exist in the Microsoft Active Directory® directory service.

For IT architects and administrators, basing their HPC clusters on Microsoft Windows Compute Cluster 2003 enables them to take advantage of their familiarity with the Microsoft Windows Server® 2003 OS to help significantly simplify cluster deployment and management. Because security is a paramount consideration, the Microsoft HPC platform helps ensure that jobs run in the context of each user's Active Directory profile and are integrated right down to the Microsoft Message Passing Interface (MS MPI) stack.

MICROSOFT WINDOWS COMPUTE CLUSTER SERVER 2003

Microsoft Windows Compute Cluster Server 2003 provides a platform for highly demanding high-performance computing (HPC) environments—one designed for simplicity, ease of deployment and use, high performance, and high reliability. By helping reduce the cost and complexity barriers of typical HPC implementations and offering world-class Microsoft support, this HPC platform can help increase performance and productivity for both end users and IT architects and administrators.

Windows Compute Cluster Server 2003 has been designed with a keen understanding of the performance requirements of HPC environments. It harnesses the power of off-the-shelf x86-64 systems, with the ease of use of Microsoft Windows environments, while helping ensure seamless integration with Microsoft Active Directory to help provide a secure, cost-effective, high-productivity HPC solution.

BUILDING INTEROPERABLE, INTEGRATED HPC CLUSTERS

Microsoft is continually expanding and strengthening close working relationships with leading third-party CAE software vendors. Many ISV applications have already been ported for use in Microsoft Windows Compute Cluster Server 2003 environments, and Microsoft anticipates many more to come in the future. Microsoft developers, working in conjunction with leading application developers, are creating CAE solution stacks with a focus on enabling a seamless experience for users moving from desktops and departmental workstations to HPC clusters.

The Microsoft HPC platform is designed to deliver the power of commercial off-the-shelf x86-64 supercomputing power to manufacturing and engineering workgroups. By helping simplify the deployment and management of HPC clusters as well as integration with existing infrastructure and tools, this interoperable platform can help accelerate both time to insight and time to market. Once implemented, a thoroughly tested, tightly integrated solution stack from Microsoft can help reduce the turnaround time for design cycles, simplify systems management, and reduce total cost of ownership—a key success metric for HPC clusters in competitive business environments.

Aditya Krishnan is a technical product manager at Microsoft and a member of the Microsoft High-Performance Computing team. He has a bachelor's degree in Computer Science and Engineering from the University of Madras and a master's degree in Computer Science with a focus on large scale-out computing from Stanford University.

"Once implemented, a thoroughly tested, tightly integrated solution stack from Microsoft can help reduce the turnaround time for design cycles, simplify systems management, and reduce total cost of ownership—a key success metric for HPC clusters in competitive business environments."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Microsoft HPC solutions for manufacturing:
www.microsoft.com/hpc/manufacturing

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

ATI FIREGL GRAPHICS

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

AMD
Smarter Choice

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

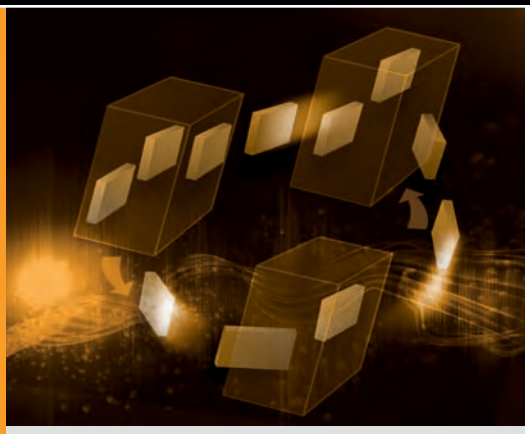

SECURING VIRTUALIZED ENVIRONMENTS WITH McAFEE INTRUSHIELD

By John Vecchi

Virtualization technology can offer significant advantages in enterprise data centers, but can also carry its own specific security risks. McAfee® IntruShield® network security and intrusion prevention system appliances are designed to provide comprehensive defenses against a wide array of external and internal threats in virtualized environments.

Related Categories:

McAfee

Security

Virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

Virtualization technology has become a key tool in enterprise data centers. By enabling administrators to run multiple virtual machines (VMs) on a single physical server, it can offer tremendous advantages for both data centers and the overall enterprise, including simplified management, efficient utilization of hardware resources, and the flexibility to support specialized requirements for testing, support, and specialty applications.

As the use of virtualization technology increases, so does the need for comprehensive security strategies to help protect virtualized environments. Tools such as McAfee IntruShield network security and intrusion prevention system (IPS) appliances can help enterprises secure both the data center as a whole and virtualized environments in particular against a wide array of external and internal threats.

UNDERSTANDING SECURITY IN VIRTUALIZED ENVIRONMENTS

Along with the significant advantages of virtualized environments come multiple security risks and challenges. All too often, however, security has been an afterthought for early adopters of virtualization technology. Best practices recommend implementing a formal security and information protection strategy that addresses the specific needs of virtualized environments.

Virtualized and non-virtualized environments share many of the same security challenges, but some are specific to virtualization (see Figure 1). Traditional computing platform threat vectors—including malware, worms, spyware, Trojan horses, and other attacks targeting software vulnerabilities—are still a concern, and the risk of propagating infections increases even further when enterprises do not take measures to help ensure the integrity of VMs. Setting up more than half a dozen VMs on a single physical server is like setting up a new data center, and as with any data center, its assets need protection.

One of the greatest vulnerabilities in a virtualized environment is the hypervisor, the platform that enables VMs with different operating systems to run on the physical system. It offers a single point of attack and can be vulnerable to *hyperjacking* attacks designed to take control of all VMs under its management. In addition, virtualized environments are typically more transient than non-virtualized environments. For example, servers may be programmed to go online and offline at unscheduled times for various reasons, such as load balancing. Those servers are potentially vulnerable to denial-of-service (DoS) attacks or other attacks (since they may have been offline during the latest patch update) that can cripple entire virtualized server farms.

PROTECTING VIRTUALIZED ENVIRONMENTS WITH McAFEE INTRUSHIELD

To help lock down a virtualized environment, it is critical to have a comprehensive approach to security—on the server, on the desktop, and on the network. McAfee offers security and compliance solutions for both virtualized and non-virtualized environments that are scalable, centrally managed, and comprehensive, spanning all three data center elements.

McAfee IntruShield network security platforms are IPS appliances designed for the first line of defense in virtualized environments—the network layer. These multidimensional, multi-vector appliances offer integrated protection through an easy-to-use, application-specific integrated circuit (ASIC)–based platform designed to provide broad physical and virtual asset protection, maximized business availability, and reduced security costs.

The IntruShield architecture integrates patented signature, behavioral anomaly, and DoS detection on a single virtualized appliance (see Figure 2). Its built-in IPS technology is designed to provide proactive, highly accurate protection against a wide range of network threats and attacks, including zero-day attacks, cyber attacks, and malware; spyware, phishing, and other unwanted programs; DoS, distributed DoS (DDoS), and SYN flood attacks; encrypted attacks, worms, Trojan horses, and evasions; instant messaging and peer-to-peer applications; voice over IP (VoIP) threats and vulnerabilities; and threats and vulnerabilities specific to VMware® virtualized environments.

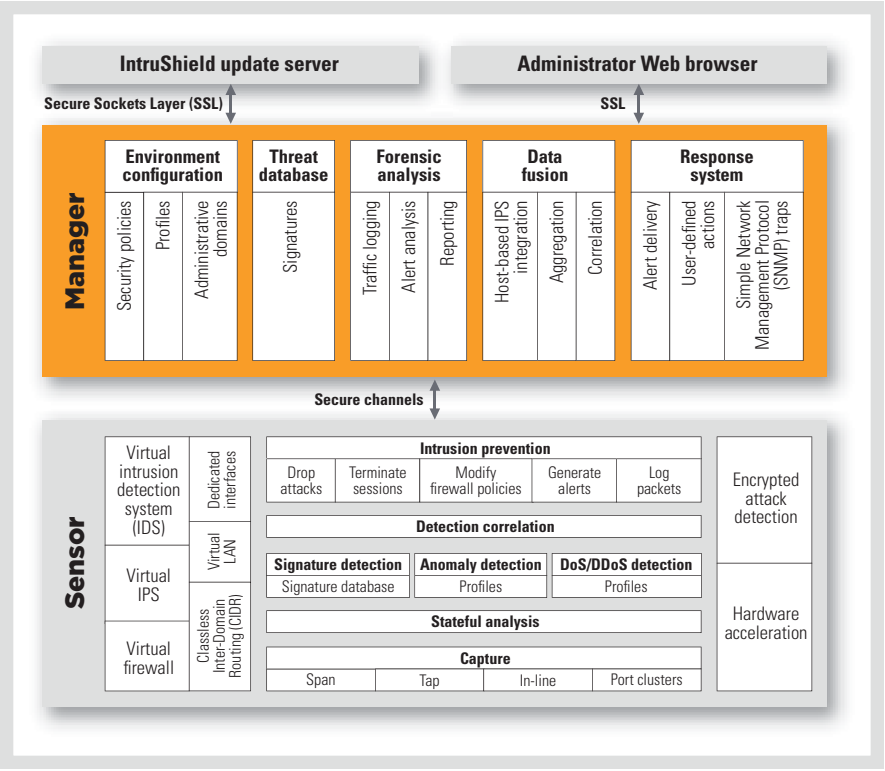

Figure 2. *McAfee IntruShield architecture*

The IntruShield network security platform also consolidates additional security features in a single integrated console, including a virtual IPS, a virtual internal firewall, built-in physical and virtual host quarantine, and protocol-based dynamic rate limiting. And, to further increase detection and prevention accuracy, the IntruShield architecture employs a combination of threshold-based and patented self-learning, profile-based detection techniques. For more information on how IntruShield security features can help defend against a variety of attacks, see the "McAfee IntruShield combats external and internal threats" sidebar in this article.

Vulnerabilities common to virtualized and non-virtualized environments	Vulnerabilities specific to virtualized environments
■ Unsecured accounts, such as those with no password or no password expiration ■ Unnecessary services ■ Backdoors such as those created by malware, worms, spyware, Trojan horses, and other attacks ■ Mis-configured Network BIOS shares or FTP servers ■ Un-patched software allowing attacks such as buffer overruns and SQL injection	■ Offline, noncompliant, or under-protected VMs ■ Vulnerabilities in the virtualization platform or hypervisor ■ Propagation and activation of infected virtualized images ■ Hyperjacking, in which a single attack can provide simultaneous access to multiple VMs and hypervisor rootkits

Figure 1. *Potential security vulnerabilities in both virtualized and non-virtualized environments*

DEPLOYING MULTIFACETED ENTERPRISE SECURITY

Comprehensive network security for virtualized environments requires implementing multiple components while maintaining network performance and flexibility. These components include protection for virtualized and non-virtualized resources and dynamic network devices, proactive protection for un-patched VMs, isolation of infected hosts, management of and limits on traffic and bandwidth, discovery of active hosts, and simple deployment and use.

The high-performance McAfee IntruShield network security platform treats virtualized environments holistically as part of the network infrastructure, just like other physical devices. It is designed to protect these environments

McAFEE INTRUSHIELD COMBATS EXTERNAL AND INTERNAL THREATS

McAfee IntruShield is designed to provide comprehensive protection against multiple types of attacks on network security in virtualized environments, including both external and internal threats.

For example, a rogue access point or laptop, when given network access, could allow hackers to access a virtual machine (VM) on the network, launch attacks, steal sensitive data, or gain other unauthorized access. McAfee IntruShield and McAfee Network Access Control solutions work together to help prevent the rogue access point and hackers from penetrating the VM. IntruShield uses "black-hole" filtering to block the attacker's traffic before it reaches the intended recipients, log the packets for later analysis, and alert administrators that intruders are attempting to breach the system.

In terms of internal threats, IntruShield could also help prevent a disgruntled employee from breaking into virtualized servers to steal credit card records, intending to sell the data or embarrass the company. IntruShield separates the network into virtual local networks and enforces firewall functionality using access control lists, which restrict internal user traffic to authorized subnets or to individual IP addresses. If an employee attempts to maliciously subvert network security measures, IntruShield can block the attacks before they reach sensitive systems. Once it has detected the attacks, IntruShield can automatically quarantine the malicious system and quickly alert administrators.

by addressing four primary aspects of security risk management (SRM): management, containment, threat prevention, and compliance and control (see Figure 3).

Management

IntruShield was designed to support the dynamic, flexible nature of virtualized environments in which administrators can launch, move, and remove virtual systems without physically reconfiguring or rearranging hardware. The IntruShield built-in virtual IPS capability—which supports up to 1,000 virtual systems on a single appliance—integrates seamlessly with virtualized environments, enabling administrators to make network changes without physically unplugging and plugging in LAN cables. IntruShield also helps protect network segments as new VMs are brought online on different physical platforms, and can proactively apply and manage VM patches the same way it does physical devices. While administrators create, test, and install patches, IntruShield helps protect vulnerable VMs with an in-line IPS.

The IntruShield virtual IPS functionality also allows administrators to easily create custom IPS and firewall security policies for virtualized environments. By adjusting different parameters, they can match security policies to specific VMs. In addition, close integration with McAfee ePolicy Orchestrator® software provides administrators with comprehensive system visibility for intelligent, centralized management and priority-based decision making.

Containment

When IntruShield detects that VMs or physical servers have become infected, violated security policies, or could propagate threats to other network devices, it can quickly quarantine those systems to help prevent further damage. IntruShield isolates quarantined systems from the rest of the network infrastructure and notifies administrators so that they can resolve the problem, restore the systems, and bring them back online as quickly as possible.

Threat prevention

By helping block unwanted traffic at the network source—before it enters the virtualized environment—IntruShield can effectively secure both VMs and physical servers from attacks. When strategically placed in-line at different points on the network, it helps safeguard the entire virtualized data center (see Figure 4).

IntruShield is preconfigured with a recommended policy designed to provide accurate, proactive blocking for hundreds of attacks. It can also boost threat protection for virtualized environments by providing multiple vulnerability-based signatures for specific protection against potential exploits in the VMware virtualization platform. McAfee Avert® Labs continually adds to these signature sets as new vulnerabilities or threats arise to help comprehensively protect virtualized environments from the latest known and unknown exploits. And the IntruShield behavioral anomaly learning engine—which scans and remembers typical VM behaviors—helps detect unusual behavior so that IntruShield can quickly alert administrators to potential security problems.

The dynamic nature of virtualized environments, which allows administrators to quickly deploy new VMs, is one of their key advantages, but it also brings potential risk.

Security aspect	Advantages
Management	■ Increased flexibility with minimal physical network changes ■ Customized security policies to meet specific VM requirements
Containment	■ Containment of infected VMs ■ Blocking of outbound malicious traffic
Threat prevention	■ Comprehensive security for virtualized and non-virtualized environments ■ Integration with McAfee Foundstone, ePolicy Orchestrator, and Remediation Manager software for simplified threat detection and remediation ■ Default blocking and vulnerability protection for VMware virtualization platforms
Compliance and control	■ Pervasive security for active and inactive hosts ■ Compliance and governance rule sets, including network isolation

Figure 3. *Key aspects of McAfee IntruShield security for virtualized environments*

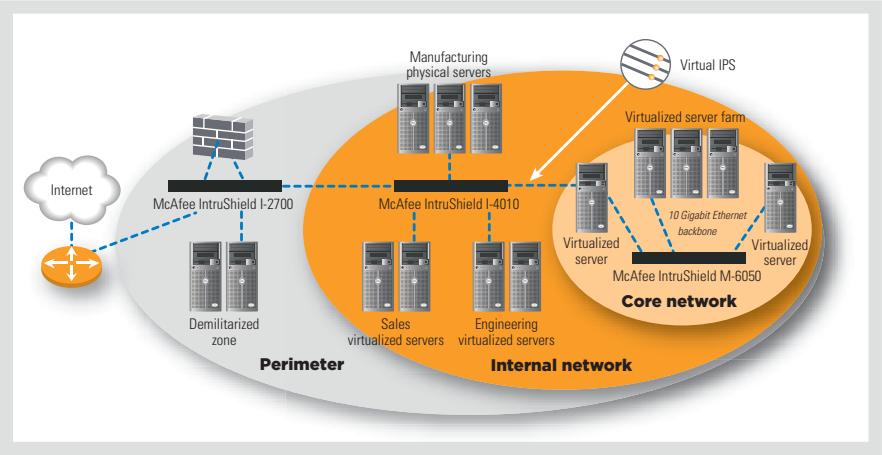

Figure 4. *Strategic placement of McAfee IntruShield appliances to help safeguard the entire virtualized data center*

New VMs could be vulnerable to certain attacks that administrators may not be aware of. IntruShield integration with the McAfee Foundstone® vulnerability management solution enables administrators to launch real-time scans of both VMs and physical servers to help determine the potential level of vulnerability, after which the system vulnerability status is updated in the Foundstone risk database.

If a system was offline the last time a vulnerability scan took place, however, the Foundstone database may be out-of-date. In this case, if the system begins to propagate infected or malicious traffic once it is back online, IntruShield can detect this traffic, initiate a block and/or quarantine, and alert administrators to the problem. Administrators can then quickly launch a real-time Foundstone scan on the system. After the scan is completed and potential vulnerabilities have been discovered and presented in the IntruShield console, the system profile is updated in the Foundstone database. Administrators can then apply the appropriate patches to help fix the vulnerability. The closed-loop process provided by McAfee SRM integration also enables them to automatically initiate help-desk remediation tickets through McAfee ePolicy Orchestrator and McAfee Remediation Manager. In the meantime, IntruShield visibility and proactive network protection capabilities help reduce the urgency of needing to deploy system patches immediately.

Compliance and control

IntruShield provides comprehensive network and system security visibility and control. Administrators can discover and see all active hosts on the network regardless of uncontrolled usage. For example, if an enterprise has a common operating environment deployed as its standard policy, administrators can monitor new VMs to help ensure that they comply with that policy, detect and protect traffic coming from the noncompliant VM, and contain and quarantine the noncompliant VM to help protect the rest of the systems on the network.

IntruShield also includes *rate limiting*, a key security and efficiency control that helps prevent low-priority VM network traffic from saturating host resources. Rate limiting helps ensure that each VM residing on a particular host has sufficient resources to continue operating at a high performance level. Granular, dynamic rate-limiting capabilities allow administrators to easily manage bandwidth by application, protocol type, or port allocation.

Also helpful to compliance and control is the IntruShield virtual internal firewall. This firewall extends the virtual IPS architecture to the internal firewall, enabling administrators to internally deploy Layer 3 and Layer 4 access control lists (ACLs) in a granular, virtual way. As such, this feature enables them to easily enforce different virtual firewall policies for a range of IP addresses on a port, right down to individual hosts. By applying ACL rules that restrict access to certain resources from specific parts of the corporate network, administrators can also use IntruShield to enforce a generic security policy for internal compliance.

To help truly secure virtualized environments, network isolation is a necessity. The ability to isolate a virtualized environment from the network or other systems is a primary purpose of a network IPS such as IntruShield. In fact, network isolation is a key component to many compliance rule sets as well, such as payment card industry regulations. And as an in-line security device, IntruShield can help organizations meet a host of governance policies.

CREATING SECURE VIRTUALIZED ENVIRONMENTS

As organizations embrace the advances of virtualization, implementing security strategies for this technology becomes increasingly important. McAfee IntruShield network security and IPS appliances can provide comprehensive protection to help both defend the data center as a whole and meet the specific security challenges of virtualized environments.

John Vecchi is the director of product marketing for network security solutions at McAfee, responsible for product marketing activities for the McAfee portfolio of network security solutions, including the award-winning IntruShield network IPS product line. He has a B.A. in International Business from the University of St. Thomas.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

McAfee IntruShield:
www.mcafee.com/intrushield

BREAKING THE ENERGY BARRIER

When MarketLive found that limited power and rack space in its data centers had become a barrier to growth, Dell™ PowerEdge™ Energy Smart servers helped the company cut power consumption by an estimated 50 percent while providing twice the processing power in the same amount of space.

CHALLENGE

MarketLive needed to overcome power issues in leased data centers, enabling the company to continue expanding its infrastructure and keep up with a rapidly growing business volume.

SOLUTION

Standardizing on Dell PowerEdge Energy Smart servers helps conserve power and reduce space requirements. The company is exploring server virtualization to further help conserve space and lower costs in the future.

BENEFITS

- The server makeover is projected to cut energy consumption in the data center in half.
- The Dell servers enable the company to double the number of servers per rack and obtain a 50 percent increase in the return on its leased space investment.
- MarketLive now has ample room to accommodate new business.
- With server deployment time reduced by 30 percent, the IT staff can focus on the needs of the company.

Related Categories:
Case study, MarketLive, power and cooling, power management, server consolidation, virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

The MarketLive 5 e-commerce application is one of the most successful software products of its kind, powering the online sales of many well-known companies. The application, which includes a suite of sales, promotion, and management tools, provides retail marketing and merchandising capabilities for customer Web sites. The ability of the MarketLive software to easily integrate with other software products has helped fuel the company's double-digit growth.

MarketLive also offers full hosting, management, and monitoring services. To deliver these services, MarketLive operates nearly 400 servers in two data centers. With that much hardware, maximizing rack density is essential. In 2006, as MarketLive grew beyond 150 customers and the number of servers mushroomed to support them, the limits of data center power availability became a barrier to growth. "We had packed so many servers into our leased space that all of a sudden we started blowing circuit breakers, which brought down servers," says Derek Irving, manager of the company's Hosting and Managed Services group. "As in many data centers, there was no additional power available. If we couldn't add servers, we couldn't support the growth of our business."

The group decided to expand to a second data center, and turned to Dell for advice on reducing server space and power requirements. "We know and trust our Dell representatives and their products," says Irving. "With other manufacturers, we have seen server failure rates as high as 30 percent out of the box. But we have never had any problems with Dell, and its servers are extremely reliable. We knew Dell would give us balanced advice on saving energy while preserving reliability."

MARKETLIVE EXPANDS WHILE SIMPLIFYING OPERATIONS

Irving asked Dell to help make the new MarketLive data center as energy efficient as possible, and asked for an assessment of energy-efficient Intel® and AMD

processors. "The Dell Infrastructure Consulting Services staff was very familiar with these processors, and explained that the Dell engineering team had created a power-optimized server line—called Energy Smart—with components designed to minimize energy consumption," Irving says. "Not only does the processor consume less energy, but so does the rest of the server."

To help improve the data center's energy efficiency while meeting its demanding performance requirements, Irving standardized on Dell PowerEdge Energy Smart 1950 servers with two dual-core Intel Xeon® processors. "The PowerEdge 1950 gives us a lot of processing power in a compact, energy-efficient 1U server," he says.

Dell Custom Factory Integration services staff also offered to help reduce installation time by loading the OS on the servers before shipping them from the factory. Irving estimates that this service saves about 30 percent in labor costs per server. "When we receive new Dell servers, we take them directly from the boxes and put them in the racks without having to install the hard drives or the OS, which is something most other suppliers require."

The Dell servers also provide MarketLive with peace of mind. "I know Dell is factory-testing the servers with the OS installed. So when I receive the server, I am confident it is ready to go," Irving says. "I don't question whether it is going to work. All we have to worry about is running our business."

DELL POWEREDGE SERVERS HELP INCREASE EFFICIENCY

Results from the PowerEdge Energy Smart servers have been dramatic. "Since we started using the Dell servers, we estimate that we have reduced our power consumption in the data centers by 50 percent," says Irving. "That is even better than we had hoped."

> **"Since we started using the Dell servers, we estimate that we have reduced our power consumption in the data centers by 50 percent."**
>
> —Derek Irving
> Manager of the Hosting and Managed Services group at MarketLive
> November 2007

The rack density of the Dell servers also enabled MarketLive to reduce its leased space requirements by 50 percent. "We had reached the limits of the data center's power supply and cooling capabilities with just 10 servers per rack," says Irving. "Now, with the Dell servers, we can comfortably double that. We are providing twice as much computing capacity from our leased space investment, and still have plenty of room to grow."

The move to PowerEdge Energy Smart servers benefits MarketLive customers as well. "The availability problems associated with being near the power limits of the data center affected up to 20 of our customers at a time, depending on how many users the server handled. Having redundant systems did not help," Irving explains. "E-commerce customers like ours need to be online 24/7, so losing a server can cause a significant slowdown for their business. Staying within the data center power constraints is essential to meeting our service-level agreements and helping our customers succeed."

DELL SERVICES HELPS MARKETLIVE OPTIMIZE ITS INFRASTRUCTURE

MarketLive is now working with the Dell Infrastructure Consulting Services team to optimize its infrastructure in other ways. For example, Dell Virtualization Consulting Services engineers are helping the MarketLive IT team plan for server virtualization. By consolidating hosting and management services on virtualized servers, MarketLive can reduce complexity, streamline infrastructure administration, and help lower costs further while enhancing scalability for its customers. "Dell takes a holistic view of our infrastructure, helping us with both server and storage virtualization," adds Irving. "Dell listens to what we are trying to do and collaborates with us to come up with the best solution." In addition, MarketLive uses Dell Enterprise Gold Support services for replacement parts and troubleshooting, and for training as the MarketLive team learns to use the Dell OpenManage™ server management suite.

Looking back, Irving reflects, "I am very proud of what we have accomplished. Since deploying the Dell servers, we have reduced power consumption and increased availability and scalability. Not only is it great for our company's bottom line, but I feel good about taking an ecologically sound approach. When you can do the right thing and do well for your business, everyone wins."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell Energy Smart:
DELL.COM/EnergySmart

EXCELLENT ONLINE SERVICE 24/7

Dell™ PowerEdge™ servers, Dell/EMC storage, and Dell OptiPlex™ desktops help online retailer Tesco.com support a 30 percent annual customer expansion rate.

CHALLENGE

Tesco.com needed an IT provider with the technology and support services to meet a significant rate of growth in its business. In particular, it wanted an organization with the expertise to maximize the uptime of its Web site, which serves 250,000 customers per week and continues to grow steadily.

SOLUTION

Tesco.com worked with Dell to create an IT infrastructure consisting of Dell PowerEdge servers, Dell/EMC storage area networks, and Dell OptiPlex desktops to help ensure an excellent online shopping experience for its customers.

BENEFITS

- High levels of performance and scalability support a 30 percent annual customer expansion rate.
- Combination of powerful servers with dual-core processors in 64-bit environment helps provide a cutting-edge infrastructure.
- Reliable infrastructure helps maintain uptime to deliver service excellence.

Related Categories:
Case study, database, server consolidation, Tesco.com

Visit DELL.COM/PowerSolutions for the complete category index.

The Internet has provided a significant opportunity for retailers to boost their services and business. However, for online brands such as Tesco.com, it is crucial that consumers receive the same excellent level of service they get in-store. As a result, online brands need reliable IT solutions that can support their promises and help drive growth in a rapidly evolving market.

Since its launch in 2000, Tesco.com has aimed to be first for customers in online retailing—with offerings that range from electronics and books to broadband and financial services. Today, with the company's focus on creating a robust IT infrastructure, it has won awards for innovative processes that support online grocery shopping.

"IT has always been fundamental to our business," says Chris Howell, head of IT operations and infrastructure at Tesco.com. "As our customer numbers grow rapidly, we have continued to focus on the fact that improving the reliability and availability of our systems is crucial for creating the ultimate customer shopping experience."

DELL SERVICE PROVIDES KEY SUPPORT

Aware that online shoppers need the same level of service as in-store visitors—minimum waiting times and high availability of produce—the company looked to

> **"Dell offers help without our having to ask. We see a real willingness on the part of Dell to understand our business and our customers."**
>
> —Chris Howell
> Head of IT operations and infrastructure at Tesco.com
> October 2007

build a relationship with a trusted IT provider. A key requirement was that the IT infrastructure help ensure maximum Web site uptime with no more than three hours of downtime over a 12-month period. "We wanted a provider with robust systems that would help ensure our customers had a great online shopping experience 24 hours a day, seven days a week," says Howell. "In addition, we needed an organization with solutions that would optimize the performance of our existing Microsoft software."

Key to the success of the relationship between Dell and Tesco.com is the Dell commitment to service. For example, Dell consultants meet with the IT team at Tesco.com each month and carry out most of their communication in person. Dell, like Tesco.com, understands how service excellence leads to success.

Says Howell, "Dell offers help without our having to ask. We see a real willingness on the part of Dell to understand our business and our customers. Dell gives us the opportunity to use the best technology solutions and gain the right level of support to help ensure our Web site performs well.

"When we need new products," he adds, "the lead time is short—only a few weeks. And if I call Dell and say I need something sooner, I know it will arrive. The Dell team does everything it can to help us, and its solutions provide the great performance needed to run our systems."

DELL HARDWARE SUPPORTS CUSTOMER EXPANSION

Since Tesco.com began working with Dell, it has rolled out a range of client and data center solutions. For example, Tesco.com employees now use Dell OptiPlex desktops running the Microsoft® Windows Vista® OS. The advanced features in Windows Vista help developers create and improve the user experience at the Tesco.com Web site, and help simplify systems integration.

"When we need new products, the lead time is short—only a few weeks. And if I call Dell and say I need something sooner, I know it will arrive. The Dell team does everything it can to help us, and its solutions provide the great performance needed to run our systems."

—Chris Howell
Head of IT operations and infrastructure at Tesco.com
October 2007

At the data center level, Dell PowerEdge servers and Dell/EMC storage area networks are supporting an annual customer expansion rate of 30 percent. One data center has upgraded its Dell deployment, and two new data centers are supporting Dell systems running the 64-bit Microsoft Windows Server® 2003 OS.

IT support personnel control the server infrastructure using the Dell OpenManage™ systems management suite. Dell OpenManage helps the IT team effectively manage its entire environment from one convenient user interface. "Having a Dell infrastructure helps simplify adding new systems to our data centers," says Howell. "It allows us to grow so we can meet the needs of our customers. And when we need new products, we can respond quickly."

Most importantly, the powerful 64-bit Dell platform allows Tesco.com to run leading-edge applications from Microsoft and enhance the functionality of its Web site. The improved performance and reliability of this infrastructure has helped the retailer transform its customer experience, helping minimize customer wait times and promoting fast, easy online shopping.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell Services:
DELL.COM/Services

ADVANCING MEDICAL TREATMENT EFFICIENCY

By deploying an IT environment built on Dell™ PowerEdge™ servers and VMware® virtualization software, Taiwan Adventist Hospital created a highly available IT environment that enables medical staff to focus on patients.

CHALLENGE

Taiwan Adventist Hospital needed to update its IT infrastructure to help ensure the efficient delivery of high-quality medical treatment. The hospital also required a remote backup system to overcome sudden power shutdowns or system errors in Taiwan's disaster-prone location.

SOLUTION

Dell PowerEdge enterprise servers enabled the hospital to deploy a virtualized infrastructure based on VMware software and to consolidate the hospital's major systems.

BENEFITS

- Comprehensive Dell solutions helped create a highly available IT environment where hospital staff can focus on high-quality medical care.
- Benefiting from a single point of contact at Dell and integrated solutions, Taiwan Adventist Hospital can look forward to significant cost savings in IT procurement and equipment maintenance.

Related Categories:
Case study, Dell PowerEdge servers, server consolidation, Taiwan Adventist Hospital, virtualization, VMware

Visit DELL.COM/PowerSolutions for the complete category index.

Efficiency, quality, and cost-effectiveness are essential for hospitals and enterprises alike, but maximizing these elements can be a struggle for today's hospitals. To help ensure the quick delivery of high-quality medical treatment, seamless integration of networks and IT is critical.

"During daily medical routines, the nurses usually record patient information such as body temperature, blood pressure, and prescriptions after the doctors' diagnosis," explains Meiken Hayashi, director of the Information Centre at Taiwan Adventist Hospital. "This is not only time-consuming, but also prone to human error."

This situation has highlighted a growing need for IT systems that help improve medical staff efficiencies. "Let medical staff focus on their jobs and leave the rest to technology," says Hayashi.

Key to improving hospital efficiency was consolidating hundreds of different information systems and simplifying data entry procedures. Medical staff could then input patient information easily into various systems across numerous departments such as the pharmacy, the medical records office, the cashier, and the registration counter. This approach could in turn streamline operations and help minimize human error.

Since the implementation of its electronic medical record (EMR) system in 2002, Taiwan Adventist Hospital has continued to deploy electronic dispensary, data backup, and network remote backup systems. Dell solutions have helped cut the average diagnosis time by more than 50 percent, reducing average patient consultation times from 60–90 minutes to only 30 minutes. And deploying VMware virtualization software on Dell PowerEdge servers—enabling the hospital to host 11 virtual machines on a single physical server—has helped improve performance, reduce storage requirements, and lower procurement costs.

DELL OFFERS HIGH-PERFORMANCE, COST-EFFECTIVE IT

"Our staff relies heavily on our IT infrastructure for daily work, thus making system stability our foremost concern," says Hayashi. "That is why we have chosen Dell to be our IT provider."

Because all medical records such as X-ray film and prescriptions are stored in the hospital's EMR system, any downtime can seriously impede patient treatments.

"When we were considering the EMR system in 2002," adds Hayashi, "we evaluated the various offerings against criteria provided by our doctors, namely, space savings, speed, and system stability."

According to Hayashi, the hospital also selected Dell for its product performance, cost-effectiveness, and brand reputation. While attending medical conferences overseas, Taiwan Adventist Hospital doctors had heard many accolades for Dell systems, which provide the hardware platform for many distinguished hospitals around the world. "Dell products are stable. Since our implementation in 2002, we have enjoyed smooth and glitch-free operations," adds Hayashi.

"The hospital requires powerful management software to monitor and maintain our hundreds of software systems. High-performance Dell systems have proven to meet our needs. We have benefited from the Dell recommendations and seamless integration with Dell partner solutions to achieve outstanding IT management."

—Meiken Hayashi
Director of the Information Centre at Taiwan Adventist Hospital
October 2007

PERSONAL SERVICE HELPS DELL MEET HOSPITAL NEEDS

Throughout his interactions with the Dell Taiwan sales team, Hayashi has been most impressed with the company's ability to provide a single point of contact. By engaging directly with Taiwan Adventist Hospital, the Dell Taiwan sales team has gained a deep understanding of the hospital's practical needs, enabling Dell to recommend a solution that streamlines the hospital's IT processes.

When comparing Dell with other vendors, Hayashi commended the Dell build-to-order model, which not only helped minimize procurement costs but also allowed easy customization to meet the hospital's unique needs. "A competent vendor must always be ready to offer multiple options and a benefits analysis," notes Hayashi. The hospital appreciated Dell efforts in proposing numerous options to address specific requirements.

Taiwan Adventist Hospital has deployed Dell PowerEdge servers, Dell OptiPlex™ desktops, Dell Latitude™ and Inspiron™ notebooks, Dell/EMC storage, and VMware Infrastructure 3 virtualization software. This reliable and comprehensive solution has allowed Taiwan Adventist Hospital to assign a single IT staff member out of its team of 12 to manage the hardware, freeing other staff members to focus on software development.

Hayashi comments, "The hospital requires powerful management software to monitor and maintain our hundreds of software systems. High-performance Dell systems have proven to meet our needs. We have benefited from Dell recommendations and seamless integration with Dell partner solutions to achieve outstanding IT management."

DELL PLATFORMS SUPPORT WORLD-CLASS MEDICAL SERVICE

"IT should serve two purposes—solution and prevention," says Hayashi. Citing the example of Typhoon Nari in 2001, Hayashi says that even though the hospital was not affected by the flooding caused by the typhoon, it still decided to set up a backup system. Since then, the hospital has also implemented data and remote network backups, and plans to deploy a comprehensive remote backup system in the future.

Dell solutions have helped ensure business continuity and dependable medical treatment for patients at Taiwan Adventist Hospital, even in the face of disaster. Looking ahead, the hospital plans to continue leveraging Dell expertise to help achieve its goal of full-scale electronic system deployment.

"Dell products are stable. Since our implementation in 2002, we have enjoyed smooth and glitch-free operations."

—Meiken Hayashi
Director of the Information Centre at Taiwan Adventist Hospital
October 2007

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell health care solutions:
DELL.COM/HealthCare

REAL-WORLD EDUCATION

By consolidating 90 percent of its physical servers onto Dell™ PowerEdge™ blade servers using the VMware® virtualization platform, the Ryerson University IT group is poised to react quickly to university needs.

CHALLENGE

The ever-increasing demand for IT services at the university put significant stress on IT staff and data center capacity. Ryerson University wanted to speed server provisioning; reduce space, power, and cooling needs; and implement rapid disaster recovery.

SOLUTION

Dell Infrastructure Consulting Services led a project to identify underutilized servers and consolidate onto Dell PowerEdge blade servers running VMware Infrastructure 3 and mirrored Dell/EMC storage area networks.

BENEFITS

- Consolidating 130 servers onto 13 physical hosts frees data center space for future expansion.
- Flexible management tools help cut server provisioning time from weeks to minutes.
- Centralized storage enables simplified backup and mirroring, helping ensure continuity of school operations.

Related Categories:
Backup, recovery, and archiving (BURA); case study; Ryerson University; server consolidation; virtualization

Visit DELL.COM/PowerSolutions for the complete category index.

Ryerson University was founded in 1948 to train a growing workforce to meet the demands of a booming postwar economy. Today, the university leads Canada in career-focused education. The university offers undergraduate and graduate degrees in professional programs as well as the arts and sciences, and includes 51 undergraduate programs and 27 graduate programs.

The university IT group must support the large number of computer applications Ryerson incorporates into its broad curriculum. However, its one-application-per-server deployment methodology had resulted in over 150 Dell PowerEdge servers across two increasingly full data centers. "Our data centers, located in downtown Toronto, could not be easily or cost-effectively expanded," says Eran Frank, manager of technical support at Ryerson. "In addition, data center complexity had grown to the point where server and storage administration were time-consuming. Managing server-attached storage was tedious, and backups were complicated. But our biggest problem was the five weeks it took to provision new servers. We just were not as agile as we wanted to be." In addition, because many servers supported only a handful of users, underutilization was common.

The solution was clear to Frank. "Server consolidation using virtualization was the answer," he says. "That would not only free space in the data centers, but would also help increase IT staff productivity, reduce backup complexity, and speed server provisioning. To make sure we followed best industry practices, I needed a partner that had extensive experience with virtualization."

VIRTUALIZED ENVIRONMENT HELPS OPTIMIZE RYERSON IT

A longtime Dell customer, Frank turned to Dell for implementation help. The first step was to characterize the workloads and identify underutilized components. The Dell Infrastructure Consulting Services (ICS) team performed a virtualization readiness assessment that included analysis of processor and disk utilization. "The results were shocking," says Frank. "Many of our servers were averaging less than 2 percent processor utilization. Dell consultants also found that we were using only 22 percent of our disk capacity."

The Dell consultants recommended consolidating 130 physical servers onto 13 virtualized servers managed by VMware Infrastructure 3 software. As a foundation

for the virtualized environment, Frank chose Dell PowerEdge 1955 blade servers with two quad-core Intel® Xeon® processors each, which provide both the high performance to host 10 virtual machines (VMs) each and the small form factor to free space for growth. Using blade servers could also help lower total cost of ownership by reducing server power consumption and overall management and maintenance costs.

To help maximize the advantages of server consolidation, the Dell consultants also recommended moving from direct attach storage to a storage area network (SAN). This centralized storage enables the Ryerson IT team to move VMs between servers as needed without disrupting end users, perform efficient backup and archive operations using VMware Consolidated Backup, and set up a mirrored SAN array at a second data center using EMC® MirrorView™ software so they can quickly recover from a disaster. Frank selected the Dell/EMC CX3-20 storage array because of its built-in high-availability features and scalability. "The Dell/EMC SAN array has redundant components such as dual power supplies and dual Fibre Channel interfaces," says Frank. "Also, the array can be easily expanded with very little additional investment. Right now we are at 10 TB, but we can scale to 83 TB in the same enclosure."

MIGRATION HELPS FREE SPACE AND REDUCE ENERGY USE

Dell Services consultants assisted the Ryerson IT staff as they performed the migrations and provided training on how to manage the streamlined environment. After the Dell ICS team demonstrated the conversion process on the first 10 servers, the Ryerson IT group consolidated the rest of its 130 servers onto the Dell blade servers. "Each blade is configured for about 10 VMs," says Frank, "which maintains a conservative 60 percent utilization rate and leaves plenty of headroom for peaks in use or to accommodate additional VMs."

> **"At our expected growth rate, we could pay back the entire cost of the Dell virtualization project within a year or so. Add the other sources of cost savings, and I must say that this project has delivered outstanding financial benefits to our school."**
>
> —Eran Frank
> Manager of technical support at Ryerson University
> November 2007

The migration helped free data center space for further expansion. "Virtualization on Dell blade servers helped us consolidate from five racks of 2U servers to one rack of Dell blade servers and the SAN," says Frank. "That reduced the amount of floor space devoted to servers and storage by 80 percent, and really positions us to accommodate the ever-increasing IT demands of the university."

The Dell blade server design and virtualization also helped reduce power consumption in the data center by 80 percent. "Energy savings was not the catalyst for consolidation," says Frank. "But, thanks to the reduction in servers and the energy-efficient design of Dell blade servers, it is a great ancillary benefit."

SERVER VIRTUALIZATION PROVIDES QUICK RETURNS

IT responsiveness has been transformed by virtualization, enabling the Ryerson IT staff to respond quickly to requests for new applications or increased processing power. In the past, provisioning a physical server could take as long as five weeks, delaying the introduction of new services. "Using VMware software on Dell servers, we can now provision a VM in 15 minutes," says Frank. "Also, we do not have to buy new equipment every time we add a new application."

According to Frank, the school expects to save more than US$500,000 over three years from increased IT staff productivity, reduced energy costs, rapid provisioning, and increased scalability. Eliminating the need to buy a new physical server for every new application is expected to more than pay for the entire virtualization project in less than a year. "When we provision virtual servers instead of physical servers, I estimate that we save US$5,000 for every new application we add," says Frank. "At our expected growth rate, we could pay back the entire cost of the Dell virtualization project within a year or so. Add the other sources of cost savings, and I must say that this project has delivered outstanding financial benefits to our school."

Ryerson University educators are now on a course to provide their trademark hands-on education experience for years to come. Adds Frank, "It is very satisfying to know that IT is so instrumental in furthering the university's educational goals."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell Virtualization Services:
DELL.COM/VirtualizationServices

Extending OpenMP to Clusters

OpenMP is a well-known parallel programming paradigm for shared-memory multiprocessors. In the past, OpenMP has been confined to Symmetric Multi-Processing (SMP) machines and teamed with Message Passing Interface (MPI) technology to make use of multiple SMP systems. A new system, Cluster OpenMP, is an implementation of OpenMP that can make use of multiple SMP machines without resorting to MPI. This advance has the advantage of eliminating the need to write explicit messaging code, as well as not mixing programming paradigms. The shared memory in Cluster OpenMP is maintained across all machines through a distributed shared-memory subsystem. Cluster OpenMP is based on the relaxed memory consistency of OpenMP, allowing shared variables to be made consistent only when absolutely necessary.

Performance Considerations for Cluster OpenMP

Some memory operations are much more expensive than others. To achieve good performance with Cluster OpenMP, the number of accesses to unprotected pages must be as high as possible, relative to the number of accesses to protected pages. This means that once a page is brought up-to-date on a given node, a large number of accesses should be made to it before the next synchronization. In order to accomplish this, a program should have as little synchronization as possible, and re-use the data on a given page as much as possible. This translates to avoiding fine-grained synchronization, such as atomic constructs or locks, and having high data locality.

The OpenMP memory model allows individual reads and writes to memory to be done in any order, as long as the synchronization operations (flushes) are done in a strict order–the same order in which they appear in the original user's program. The lack of ordering between reads and writes to memory makes their concurrent execution possible, but all flushes in a program must be serialized, adding overhead to the program.

Latency to L1	**1-2 cycles**
Latency to L2	**5-7 cycles**
Latency to L3	**12-21 cycles**
Latency to memory	**180-225 cycles**
Gigabit Ethernet latency to remote node	**~ 28000 cycles**
InfiniBand* latency to remote node	**~ 23000 cycles**

Figure 1. Itanium® processor latency to cache and memory compared with messaging latency to remote nodes. Cluster OpenMP currently runs on Itanium®-based platforms and on systems based on processors that support Intel® EM64T running Linux*.

Figure 1 shows the number of processor cycles required for an access to different levels of cache and the latency to access a value in the memory of a remote node. This shows that access to the memory of a remote node is approximately 100 times slower than access to the local memory, and thousands of times slower than access to a value in cache. This comparison should drive home the point that local, rather than remote, memory should be used as much as possible.

Cluster OpenMP provides shared memory across a cluster for an OpenMP program. It takes advantage of the relaxed memory model of OpenMP to optimize the memory accesses in an OpenMP program.

Cluster OpenMP does not perform well for all types of programs, but programs with certain characteristics can achieve reasonably good performance on a cluster, compared with attainable performance on a hardware shared memory machine.

This white paper can be downloaded in its entirety from http://assets.devx.com/goparallel/19403.pdf and shows performance results for a set of applications, showing that most were able to achieve greater than 70 percent of the performance of OpenMP programs run on a shared memory machine.