DELL POWER SOLUTIONS · JUNE 2009

POWER SOLUTIONS

DELL.COM/PowerSolutions

JUNE 2009 • $4.95

EFFICIENT ENTERPRISE

BOOSTING YOUR DATA CENTER IQ

FEATURE SECTION

Data Center Workhorses: New 11th-Generation Dell PowerEdge Servers and Blades

iSCSI for the Enterprise: New Dell EqualLogic PS6000 Series Storage

Computing Pods: Large-Scale Data Center Building Blocks

MEETING ECONOMIC CHALLENGES THROUGH VIRTUALIZATION
SEE PAGE 56

MANAGING ENDPOINT SECURITY FOR MOBILE CLIENTS
SEE PAGE 63

THE EFFICIENT ENTERPRISE

STK NO: DPS-200906-US

POWER SOLUTIONS

JUNE 2009

DELL.COM/PowerSolutions

THE EFFICIENT ENTERPRISE: BOOSTING YOUR DATA CENTER IQ

By Jeanne Feldkamp, Debra McDonald, and Tom Kolnowski

Economic pressures are transforming virtually every area of business, and the data center is no exception. New-generation Dell™ servers, storage, and management infrastructure offer outstanding enterprise efficiency together with exceptional virtualization, scalability, and support—helping IT executives to cut costs without compromising strategic organizational goals.

DATA CENTER WORKHORSES: NEW DELL POWEREDGE RACK AND BLADE SERVERS

By Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner

New 11th-generation Dell PowerEdge™ rack and blade servers are designed from the ground up for high performance, energy efficiency, and simplified management—offering a host of advantages in enterprise IT environments to help optimize resources and reduce total cost of ownership.

ENTERPRISE STORAGE PAGE 50

BOOSTING SAN PERFORMANCE WITH DELL EQUALLOGIC PS6000 ARRAYS

By Dylan Locsin, Ujjwal Rajbhandari, and Wendy Chen

Solid-state drives (SSDs) offer a new tier of high-performance, highly reliable Dell EqualLogic™ PS Series storage. The Dell EqualLogic PS6000S array provides cost-effective SSD-based storage that can easily integrate into multitiered storage area network (SAN) environments—providing a simplified, scalable way to help meet the performance demands of targeted applications.

CLIENT SECURITY PAGE 63

UNDERSTANDING THE BASICS OF DEPLOYING AND MANAGING SECURE ENDPOINTS

By David Schweighofer

With the ranks of mobile and remote workers expanding every day, cost-effective endpoint security is a growing business concern that enterprise IT organizations avoid at their peril. Dell advances enterprise-wide protection with a balanced approach that focuses equally on safeguarding data and preventing unauthorized access.

CONTENTS

DEPARTMENTS

VIRTUALIZATION

CLIENT SECURITY AND VIRTUALIZATION

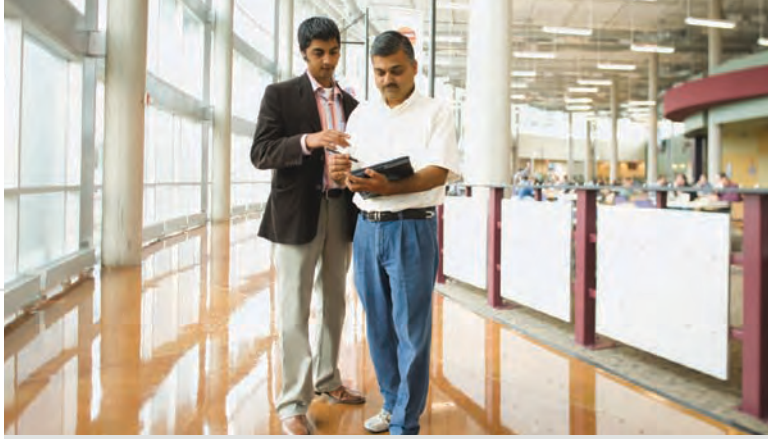

CASE STUDIES

VIRTUALIZATION PAGE 56

MEETING ECONOMIC CHALLENGES THROUGH ACCELERATED VIRTUALIZATION

By Steve Burrows and Ron Oglesby

In difficult economic times, many enterprises are seeking ways to streamline IT budgets without affecting operations. Evaluated and refined across thousands of real-world deployments, Dell virtualization solutions and services offer a comprehensive approach designed to quickly reduce total cost of ownership and maximize return on investment.

SEE IT HERE FIRST

DELL POWER SOLUTIONS ONLINE EXCLUSIVES

MORE ONLINE

DELL.COM/PowerSolutions

Really Remote Computing
From oil fields to glaciers to the middle of the ocean, ruggedized technology is breaking ground in some startling new frontiers.

Right Here, Right Now
Location-based services have the potential to transform the way your company builds relationships with customers. But a personalized, actionable marketing strategy is only as effective as the business intelligence systems driving it.

The Top 5 Mistakes Companies Make When Trimming Their IT Budget
Take some of the sting out of your shrinking IT budget by avoiding gaffes that could come back to haunt you later.

VMware ESXi Systems Management on Dell PowerEdge Servers
By Balasubramanian Chandrasekaran, Sudhir Shetty, Viswanathan Balakrishnan, Alok Sinha, and Charu Chaubal, Ph.D.
By integrating VMware® ESXi with Dell PowerEdge servers and using tools such as Dell OpenManage™ Server Administrator to monitor, inventory, and configure these systems, IT organizations can create scalable, highly manageable virtualized environments.

Evaluating iSCSI SANs: Dell PowerVault MD3000i Outperforms HP StorageWorks 2000i
In Principled Technologies tests designed to evaluate Internet SCSI (iSCSI) SAN performance for online transaction processing and Microsoft® Exchange Server 2007, the Dell PowerVault™ MD3000i array outperformed the HP StorageWorks 2000i array across all tested metrics.

BLADE SERVERS PAGE 94

INTEGRATING DELL POWEREDGE M1000e BLADE ENCLOSURES WITH DELL EQUALLOGIC SANs

By Tony Ansley

Different approaches to integrating Dell PowerEdge M1000e modular blade enclosures with Dell EqualLogic PS Series iSCSI SAN arrays offer specific trade-offs in cost, performance, and complexity. By understanding the available options and following best practices for deployment, administrators can implement a consolidated, highly scalable infrastructure that can meet their specific needs.

INDEX TO FEATURED DELL CUSTOMERS

INDEX TO ADVERTISERS

POWER SOLUTIONS
DELL.COM/PowerSolutions

June 2009

MORE METTLE FOR THE NOMAD

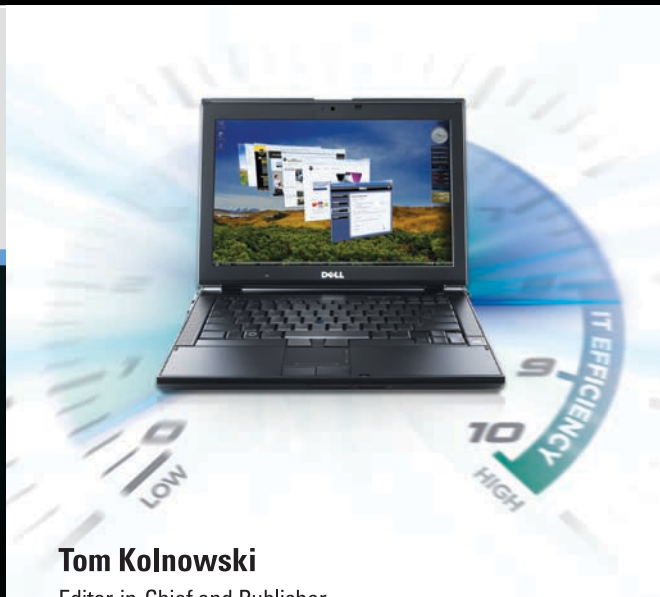

Tom Kolnowski
Editor-in-Chief and Publisher
tom_kolnowski@DELL.COM
DELL.COM/PowerSolutions
powersolutionsdigital.dell.com

My corporate office space and its hardwired, tethered hardware (phone, flat-panel display, and dock) often remain unused for days at a time—having long been supplanted by a smartphone, laptop, and working-where-it-makes-sense approach to chasing editorial deadlines. Recently, the Dell™ Latitude™ D620 laptop that had been at the core of my nomadic existence reached the end of its service life. What follows are observations from a workaday life with my latest-generation Latitude E6400, amassed as we zeroed in on the press deadline for this issue:

- **Not-so-heavy metal:** Having held the very light, bare-metal castings of a preproduction Latitude E-Family chassis last August when we visited with Ken Musgrave, head of the Dell worldwide industrial design and usability team, I was eager to get a production unit in hand. As I opened its lid for the first time, the Latitude E6400 did not disappoint—the cast magnesium alloy chassis and LCD back panel, joined with solid metal hinges, give the distinct impression that the laptop is hewn from a contiguous ingot of metal. Yet it is lighter and consumes less space than the laptop it replaced.
- **An illuminating experience:** As a Microsoft® Word warrior, I incessantly beat on my keyboards. But the Latitude E6400 keyboard has an overall feeling of solidity, with no detectable bounce. I was certain I would have no use for the optional backlighting, but a multi-hour editing session from a dimly lit coffee shop made me a believer. The backlit keys are much brighter than expected, go dim after 15 seconds or so of inactivity to save power, and are easily toggled on or off.
- **Applications en masse:** Nonstop editorial workflow means running the full suite of Microsoft Office applications as well as heavy-footprint mainstays from the Adobe® Creative Suite: the Acrobat®, InDesign®, and Photoshop® applications. With the Microsoft Windows Vista® Enterprise OS loaded on the Latitude E6400, these applications seem to launch faster and behave better than they did on my last laptop. And, as a confirmed addict of the Windows Vista Flip 3D feature, I find that cascading through a dozen or more open windows is decidedly smooth and fast.

Migrating to the Latitude E6400 has convinced me that the latest generation of mobility gear is a significant advance over what was standard fare even a few short years ago, further blurring the lines between workplace and work location, networking seamlessly and ubiquitously—and always on precisely when and where I need to be. For much more information on the digital nomad existence, visit www.digitalnomads.com.

ENTERPRISE-CLASS FILE SHARING

Dell™ NX4 breaks down traditional barriers allowing users to share files between Windows®, Linux®, and UNIX® environments

LEARN MORE AT DELL.COM/NX4

THE EFFICIENT ENTERPRISE:

BOOSTING YOUR DATA CENTER IQ

Economic pressures are transforming virtually every area of business, and the data center is no exception. New-generation Dell™ servers, storage, and management infrastructure offer outstanding enterprise efficiency together with exceptional virtualization, scalability, and support—helping IT executives to cut costs without compromising strategic organizational goals.

By Jeanne Feldkamp
Debra McDonald
Tom Kolnowski

Related Categories:

Dell EqualLogic storage
Dell PowerEdge blade servers
Dell PowerEdge rack servers
Efficient Enterprise
Power and cooling
Systems management
Virtualization

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

Weathering a tough economic climate means putting a hard stop on wasteful spending. Data center complexity further complicates budget requirements, driving CIOs to scrutinize the IT infrastructure and devise innovative ways to get the job done while consuming as few resources as possible.

As energy costs rise, so does the cost of running a data center. Nonetheless, information is the lifeblood of virtually every business—and data center capacity must continue adjusting and expanding to keep pace, even when power and cooling capabilities may have reached their limit.

At times like this, it is no surprise when the finance organization calls for belt-tightening tactics that have worked before: simply delay noncritical expenditures. Unfortunately, short-term budgetary measures intended to produce immediate spending relief may undercut long-term organizational goals.[1] It is essential for IT executives to work proactively with business strategists to advance overall enterprise efficiency. Rather than paring back, many organizations are forging ahead with strategic IT infrastructure deployments designed to increase data center intelligence while reducing operating expenses—effectively boosting their data center IQ. Measures that can help achieve this desired state include the following:

- **Enhanced energy efficiency:** Much of the electricity consumed in a conventional data center goes toward cooling. Administrators can ease power and cooling requirements by implementing energy management best practices as well as consolidating servers and storage. In addition, transitioning to equipment that incorporates the latest advancements in energy-saving technology can help significantly reduce operating expenses. Recent strides in rack server efficiency may fit the bill for energy savings, but with the onset of the latest generation of full- and half-height blade servers, this may be the time to increase the population of blades in the data center mix.
- **Increased server and storage virtualization:** By facilitating server and storage consolidation and the efficient use of computing resources, virtualized environments enable considerable cost savings compared with all-physical configurations. Moreover, while server virtualization is essential to laying the foundation for scalable, highly automated IT infrastructures, organizations may also find this an opportune time to begin progressing along the storage virtualization front by integrating the latest in Internet SCSI (iSCSI) storage arrays.
- **Smart systems management:** In the past, a variety of tools were required to manage heterogeneous systems. Today, administrators can access comprehensive functionality from a single centralized console designed to manage the entire data center—helping IT departments reduce the administrative burden and increase productivity.
- **Intensified focus on smooth integration and scalability:** Organizations are going back to basics with particular emphasis on productivity, efficiency, and value. To support these objectives, data center infrastructures must be flexible enough to change as new systems are added and to grow with the needs of the organization.

INTRODUCING THE NEW GENERATION

Dell servers, storage, and management systems can help IT organizations meet today's economic challenges while supporting fast, seamless growth. New 11th-generation Dell PowerEdge™ servers feature a high-performance, energy-efficient system architecture that can be optimized to support demanding virtualized environments. A range of rack, tower, and blade servers designed for easy customization provide flexibility through innovative design and usability enhancements, helping simplify management of complex data center infrastructures while helping reduce total cost of ownership.[2]

[1] For more information on budget pitfalls to avoid, see "The Top 5 Mistakes Companies Make When Trimming Their IT Budget," in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090357-Budgets.pdf.

[2] For more information on new 11th-generation PowerEdge servers, see "Data Center Workhorses: New Dell PowerEdge Rack and Blade Servers," by Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090246-Nguyen.pdf.

PowerEdge servers also take advantage of the Intel® Xeon® processor 5500 series architecture, which offers outstanding flexibility and bandwidth including up to three memory channels per processor. Best-practices Double Data Rate 3 (DDR3) configurations enable organizations to optimize performance with settings for targeted enterprise applications and virtualized environments.

Building on the capabilities of Dell servers and storage, Dell OpenManage™ systems management tools offer enhanced operations and standards-based commands designed to integrate with existing systems. The new Dell Management Console Powered by Altiris™ from Symantec™—based on the modular Symantec Management Platform framework—offers a holistic view of the data center and unified management of enterprise-wide IT assets through a single console and configuration management database.[4]

Because the Dell Management Console is device agnostic, it is designed to manage a wide range of third-party hardware devices and operating systems using a common data source and a single view into the IT infrastructure. These centralized systems management capabilities also allow administrators to track storage resources, monitor performance, and perform updates quickly and cost-effectively. In addition, the Dell Unified Server Configurator (USC) enabled by the Lifecycle Controller in 11th-generation PowerEdge servers facilitates OS deployment with built-in driver installations, firmware updates, hardware configuration, and diagnostics. Because the USC is available even when the OS is not, it allows added flexibility in provisioning and customizing systems to suit enterprise requirements.

CASE IN POINT

"As a utility company, PG&E is committed to conserving energy. With rightsized power supplies, physical designs that maximize airflow, and the ability for us to tune BIOS settings to enhance energy efficiency, the new generation of Dell PowerEdge servers will help us deliver outstanding service to our customers while controlling our own energy usage."

—Tony Villa
Server hardware product manager at PG&E
March 2009

Efficiency-enhancing features are built into Dell storage as well. The new Dell EqualLogic™ PS6000 series of iSCSI storage area network (SAN) arrays provides an exceptional storage foundation for enterprises. Through their virtualized scale-out architecture and rich software features, these arrays are designed to deliver self-optimized performance, virtualized server integration, integrated data protection for key business applications, consolidated management, and flexible deployment without compromising data availability.[3] The simplicity of management integrated into the EqualLogic iSCSI SAN line, its ease of integration into broader IT infrastructures, and a comprehensive set of storage management features help drive down SAN total cost of ownership.

MAXIMIZING PERFORMANCE AND ENERGY EFFICIENCY

New 11th-generation Dell PowerEdge servers are built from the ground up with energy-efficient components and energy management features designed to minimize system power consumption. For example, these servers include power supply units that are rightsized for system requirements, policy-driven power and thermal management, and highly efficient, standards-based Energy Smart components. Advanced Dell thermal control helps to deliver optimal performance while minimizing system and fan power consumption. Together, these enhancements help maximize energy efficiency and reduce operating expense without compromising performance.

All together, these technology advances make it possible for IT executives to take a giant leap forward in enterprise efficiency. "Energy efficiency is key to our business model—both from a cost and an environmental perspective," says Patrick Pulvermueller, managing director at hosting services provider Host Europe. "The Dell PowerEdge R710 will allow us to double capacity and computing power without increasing our energy costs and carbon footprint. For us, that's a winning formula."

Tony Villa, server hardware product manager at Pacific Gas and Electric (PG&E), has a similar perspective. "As a utility company, PG&E is committed to conserving energy. With rightsized power supplies, physical designs that maximize airflow, and the ability for us to tune BIOS settings to enhance energy efficiency, the new generation of Dell PowerEdge servers will help us deliver outstanding service to our customers while controlling our own energy usage."

[3] For more information on virtualized storage capabilities, see "Boosting SAN Performance with Dell EqualLogic PS6000S Solid-State Drive Arrays," by Dylan Locsin, Ujjwal Rajbhandari, and Wendy Chen, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090236-Locsin.pdf.

[4] For more information on the Dell Management Console, see "Systems Management Simplified," by Pascal Nicolas, Jeanne Feldkamp, and Tom Kolnowski, in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090191-CoverStory.pdf.

PowerEdge servers further boost efficiency by helping to increase the amount of work each processor can perform in a given time period. In each 11th-generation PowerEdge model, processors with the Intel Xeon processor 5500 series architecture are designed to adapt to software in real time—allowing systems to process more tasks simultaneously than they could otherwise. Intel Turbo Boost Technology enables PowerEdge servers to enhance performance during peak usage periods.

Field tests vouch for energy-efficient performance characteristics. "When we benchmarked the Dell PowerEdge M710 blade servers with Intel Xeon 5500 processors, we saw an increase in power efficiency of at least 25 percent," says Helge Meinhard, head of server and storage procurement at European Organization for Nuclear Research (CERN) Central IT. "Additionally, we enabled simultaneous multi-threading, and saw power efficiency go up even further. We could make a potential 50 percent overall improvement."

Intel Intelligent Power Capability further enhances efficiency by putting PowerEdge servers into low power states when demand decreases, helping to reduce operating costs and energy use. "The Dell PowerEdge R610 server and PowerEdge M610 blade server will help us reduce overall power consumption and get the most out of what we do use," says Gary Jung, manager for the Scientific Cluster Support Group in the IT Division at University of California, Berkeley (UC Berkeley) Laboratory.

CASE IN POINT

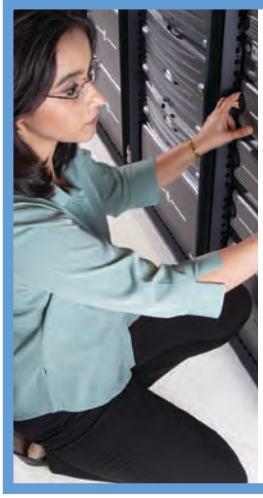

"When we benchmarked the Dell PowerEdge M710 blade servers with Intel Xeon 5500 processors, we saw an increase in power efficiency of at least 25 percent. Additionally, we enabled simultaneous multi-threading, and saw power efficiency go up even further. We could make a potential 50 percent overall improvement."

—Helge Meinhard
Head of server and storage procurement for CERN Central IT
March 2009

BENEFITING FROM VIRTUALIZATION ENHANCEMENTS

A range of enhancements available in the new Dell PowerEdge servers help organizations make the most of high-availability virtualized data center configurations. By helping reduce the complexity of the IT environment while adding to the flexibility and scalability of compute options, virtualization can give enterprises a competitive edge.

Featuring the Intel Xeon processor 5500 series architecture, embedded hypervisors, and expanded memory and I/O, 11th-generation PowerEdge servers are designed to deliver outstanding system performance and support a large number of virtual machines per server. "With greater memory capacity than the previous-generation hardware, the Dell PowerEdge R710 server will enable us to increase the number of virtual machines on each physical server by approximately 60 percent," says Greg Barton, senior analyst at engineering, construction, and operations firm CH2M HILL. "As a result, we can retire more servers and reduce IT costs."

Dell blade servers are also enabling a substantial boost in virtualization efficiency. "VMware virtualization software operates seamlessly on the Dell PowerEdge M610 blade server. Previously we needed five blades to run 100 virtual servers in our environment, but we'd only need three M610s to do the same work," notes Gauthier Catteau, systems and network engineer at Ministère de l'Education Nationale (Meduc).

Optional factory-integrated virtualization capabilities allow administrators to specify custom configurations at the point of purchase to deliver customized solutions out of the box—helping to further simplify deployment and implementation of virtualized infrastructures. For Alex Rodriguez, director of systems engineering at data center services provider Expedient Communications, these capabilities translate into extremely fast time to payback. "We plan to move the entire Expedient virtualization server farm to Dell PowerEdge R710 servers as quickly as we can roll them out," he says. "The additional memory, I/O

CASE IN POINT

"VMware virtualization software operates seamlessly on the Dell PowerEdge M610 blade server. Previously we needed five blades to run 100 virtual servers in our environment, but we'd only need three M610s to do the same work."

—Gauthier Catteau
Systems and network engineer at Ministère de l'Education Nationale (Meduc)
March 2009

bandwidth, and processor speed deliver so much additional capacity that the new servers will pay for themselves almost as soon as I put them in."

SIMPLIFYING INTEGRATION AND MANAGEMENT

In today's economy, ease of integration and management are essential considerations. To that end, 11th-generation Dell PowerEdge servers offer world-class system and image commonality. Enabling consolidation onto a reduced number of highly functional hardware platforms helps IT departments further simplify integration and management.

In addition to consistent hardware design, PowerEdge servers provide a variety of features designed to streamline systems administration. For example, the USC delivers "instant on" integrated manageability through a single access point. This tool offers a one-stop shop for deploying operating systems with built-in driver installations, firmware updates, hardware configuration, and issue diagnoses. Remote management capabilities can then extend these features to administrators working outside the data center. "We have a large, heterogeneous environment," says Host Europe's Pulvermueller, "but with iDRAC, we'll be able to manage all elements, regardless of where we are. Remote management means peace of mind, maximum uptime, and reduced costs."

The Dell Management Console helps simplify IT and create stability by shrinking infrastructure management to a single console. Because it offers centralized access to management tools and a common data source for managing the entire infrastructure, this platform enhances productivity and frees administrators to focus on strategic priorities. "With Dell Management Console, we will be able to consolidate our three separate IT management consoles into a single unified interface," says Kevin Jones, project manager at international transportation company CSX. "That alone will save us quite a bit of time."

Built on the Symantec Management Platform, the Dell Management Console has an easily extensible, modular foundation designed to provide secure access to a comprehensive range of systems management functionality, from basic hardware management to advanced features such as asset and security management enabled through the purchase of optional plug-ins to Dell and partner tools. "In the past, the IT group had to log on to each individual server to check its status and manage routine tasks," says Aaron C. Duncan, network administrator at Chicago Public Radio. "By consolidating server management into a single tool, with a single user-friendly interface, the Dell Management Console will help us dramatically reduce the time we spend on server management."

The 11th-generation PowerEdge server family is designed to deliver simplicity of operations and innovative, administrator-friendly design features. Expedient's Rodriguez says his organization benefits daily from this simplicity: "At Expedient, we reconfigure servers in response to customer requests every day, and we can change components on the Dell PowerEdge R610 server in just a few minutes, which means that we have more time to spend improving customer service."

Each model in the new lineup of PowerEdge servers takes advantage of Dell system image commonality. This commonality enhances productivity because administrators who learn one system also understand how to manage other PowerEdge models and next-generation Dell servers. Logical component layout and power supply placement also help ensure straightforward installation and redeployment. According to Andrey Paramonov, senior system administrator/storage administrator at e-mail hosting company Intermedia, "The physical design and management capabilities of the Dell PowerEdge M610 blade server enable us to deploy a new server up to six times faster than a standard stand-alone server."

Rip-and-replace processes are not a typical option for enterprises looking to increase data center efficiency. Because they are based on industry standards and validated with Dell partners, 11th-generation PowerEdge servers help ensure seamless integration into existing infrastructures. "We're confident that the Dell PowerEdge R610 servers and PowerEdge M610 blade servers will fit easily into our environment," says UC Berkeley Laboratory's Jung. "We manage over 1,400 production servers, and our experience has been that the Dell systems take noticeably less effort to deploy and maintain than those from other vendors."

Rodriguez agrees that Dell systems integrate easily into the existing data center. "The component density, durability, and easy access features of the Dell PowerEdge R610 servers save us time during both deployment and maintenance,"

CASE IN POINT

"We plan to move the entire Expedient virtualization server farm to Dell PowerEdge R710 servers as quickly as we can roll them out. The additional memory, I/O bandwidth, and processor speed deliver so much additional capacity that the new servers will pay for themselves almost as soon as I put them in."

—Alex Rodriguez
Director of systems engineering at Expedient Communications
March 2009

CASE IN POINT

"The physical design and management capabilities of the Dell PowerEdge M610 blade server enable us to deploy a new server up to six times faster than a standard stand-alone server."

—Andrey Paramonov
Senior system administrator/ storage administrator at Intermedia
March 2009

he says. "We can have a new PowerEdge server installed and ready to load applications in less than an hour."

SCALING SEAMLESSLY WITHIN THE SAME DATA CENTER FOOTPRINT

Data centers must be able to scale effectively while using data center real estate efficiently. Scalability starts at the hardware level while working within the given power resources. That is why 11th-generation Dell PowerEdge servers—and PowerEdge M-Series blade servers in particular—are designed for scalability from the inside out. As application needs increase, PowerEdge M610 and PowerEdge M710 blade servers are designed to scale up to 128 cores and 1,536 GB of memory per 10U PowerEdge M1000e modular blade enclosure. Administrators can also effectively scale I/O application bandwidth with end-to-end 10 Gigabit Ethernet or Fibre Channel connectivity.

"The Intel Xeon 5500 processors utilize hyper-threading. This gives us eight cores per CPU, and you can add another core as needed. This scalability makes the Dell PowerEdge R610 a very attractive option," says Folkert de Gans, senior systems and network administrator at housing company Woonzorg.

Moreover, by integrating Dell server, storage, and network infrastructure elements together with advanced virtualization and software tools, administrators can significantly advance data center automation. Today, large-scale data center building blocks called computing pods are designed to aggregate server, storage, and network nodes into a simple, cooperating set of computing resources. Computing pods also incorporate a comprehensive management infrastructure that can be tailored to particular organizational requirements to enhance enterprise efficiency and enable fast, flexible service response to changing business conditions.[5]

To further advance enterprise efficiency and support business growth, Dell offers a comprehensive range of services, from consulting engagements all the way to outsourcing management of IT tasks to Dell. Through a modular approach that allows enterprises to modify their systems management services as requirements change, the Dell Services team can help IT leaders streamline data center operations and customize support plans to help meet evolving enterprise needs.

CAPITALIZING ON LONG-TERM EFFICIENCY GAINS

For many executives, the response to difficult economic times, tight budgets, and escalating energy costs is a bet-your-business decision. While IT organizations are being called upon to make short-term service cuts, many business visionaries recognize that the biggest long-term savings may come from strategic capital investments that advance data center efficiency and intelligence—particularly cost-efficient deployments of blade servers and iSCSI storage. In addition, adhering to energy management best practices, transitioning to energy-saving systems, and integrating virtualization technology helps simplify management and boost data center automation significantly.

Besides helping to survive today's economic storm, these strategies are fundamental for continued growth. By increasing overall enterprise efficiency and scalability, IT executives can keep pace with ever-changing business requirements while continuing to innovate essential systems and services that give their organizations the competitive edge.

Jeanne Feldkamp is a business and technology writer based in San Francisco and a frequent contributor to *Dell Power Solutions* Magazine.

Debra McDonald is the managing editor of *Dell Power Solutions* Magazine.

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

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell EqualLogic PS Series:
DELL.COM/EqualLogic

Dell Services:
DELL.COM/Services

[5] For more information on computing pods and Dell Business-Ready Configurations, see "Computing Pods: Large-Scale Building Blocks for Intelligent, Automated Data Center Deployments," by Timothy Sherbak and Chris Auger, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090238-Sherbak.pdf.

DATA CENTER WORKHORSES: NEW DELL POWEREDGE RACK AND BLADE SERVERS

By Edward Yee
Indrani Paul
Robert Tung
Truc Nguyen
Chad Fenner

New 11th-generation Dell™ PowerEdge™ rack and blade servers are designed from the ground up for high performance, energy efficiency, and simplified management—offering a host of advantages in enterprise IT environments to help optimize resources and reduce total cost of ownership.

Related Categories:

Blade servers
Dell PowerEdge blade servers
Dell PowerEdge rack servers
Dell PowerEdge servers
Dell PowerEdge tower servers
Energy efficiency
Green IT
Intel
Power and cooling
Servers

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

In enterprise IT environments, success often depends not only on maximizing performance, but also on controlling costs by reducing power consumption, cooling requirements, and administrative complexity. To help meet these needs, the new 11th-generation Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 rack-mountable servers and PowerEdge M610 and PowerEdge M710 blade servers are designed from the ground up for high performance, energy efficiency, and simplified management. These general-purpose, two-socket servers are designed for use in a wide range of environments, including mainstream enterprises, medium-sized data centers, remote offices, and growing small organizations. Their flexible design provides a variety of options for internal storage and I/O expansion and incorporates the latest advances in virtualization, performance, power and cooling, systems management, and usability. These technologies are packed into highly available rack-mountable chassis that can fit into standard Dell and third-party 1,000 mm depth racks.

Key enhancements in 11th-generation PowerEdge servers range from the new dual- and quad-core Intel® Xeon® processor 5500 series with QuickPath Interconnect (QPI) technology, high-speed PCI Express (PCIe) 2.0 I/O interconnects, and embedded Gigabit Ethernet network controllers to flexible chassis options, energy-efficient power supplies, and intelligent cooling. (For more details on enhancements specific to PowerEdge M610 and PowerEdge M710 blade servers, see the "Dell PowerEdge M-Series modular evolution" sidebar in this article.) In addition, a breakthrough systems management design incorporates advanced power management, system monitoring, hardware configuration, deployment, and updates. These enhancements can provide a variety of benefits in enterprise IT environments—helping simplify management, control power and cooling requirements, optimize hardware resources, and reduce total cost of ownership.

NEXT-GENERATION SYSTEM ARCHITECTURE

The 11th-generation PowerEdge server architecture is designed for high performance and optimized energy efficiency. Key components of PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers include the processors and integrated memory controller, I/O chipset, and embedded network controllers.

Processors and integrated memory controller

Optimized for performance with the power efficiencies of a low-power micro-architecture and based on 45 nm process technology, the Intel Xeon processor 5500 series incorporates the new Intel QuickPath integrated memory controller and point-to-point link interface using QPI technology. PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers

New 11th-generation Dell PowerEdge servers help simplify management, control power and cooling requirements, optimize hardware resources, and reduce total cost of ownership

support up to two of these dual- or quad-core processors, including support for 60 W, 80 W, and 95 W models. These powerful, efficient multi-core processors help maximize performance and performance per watt for data center infrastructures and dense deployments.

The Intel Xeon processor 5500 series features two Intel QPI links capable of up to 6.4 GT/sec, up to 8 MB of shared cache, and the Intel QuickPath integrated memory controller. It also supports several advanced Intel technologies, including Execute Disable Bit functionality, the Intel 64 architecture for flexibility in 32- and 64-bit applications and operating systems, Enhanced Intel SpeedStep® Technology, Intel Virtualization Technology, and simultaneous multi-threading. Selected processors also support Intel Turbo Boost Technology, an OS-controlled operation that can automatically allow the processor to run faster than the marked frequency if the processor is operating below power, temperature, and specified current limits. The Intel Xeon processor 5500 series also supports deeper C-states than previous-generation processors to enhance power conservation.

The integrated memory controllers support Double Data Rate 3 (DDR3) technology, which is designed to provide a high-performance memory interface capable of low-latency response and high throughput.[1]

I/O chipset

PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers use the Intel 5500 and Intel 5520 I/O hub (IOH) chipsets and I/O controller hub 9 (ICH9) to connect the I/O devices and processors. The PowerEdge R610 server uses the Intel 5500 IOH 24D chipset, while PowerEdge R710 and PowerEdge T610 servers use the Intel 5520 IOH 36D chipset. These chipsets support the Intel Xeon processor 5500 series, QPI, DDR3 memory, and PCIe 2.0. The primary components of the IOH are two full-width QuickPath Interconnect links (one to each processor), 24 or 36 lanes of PCIe 2.0, an x4 Direct Media Interface (DMI), and an integrated I/O advanced programmable interrupt controller (APIC).

The QPI architecture consists of serial point-to-point interconnects for the processors and IOH. The system has a total of three QPI links: one connecting the processors to each other and two connecting the processors to the IOH. The QPI link is designed to support up to 6.4 GT/sec, depending on the processor.

PCIe provides the serial point-to-point interconnect for I/O devices. PCIe 2.0 is designed to double the signaling bit rate of PCIe 1.0, from 2.5 Gbps to 5 Gbps. PCIe 2.0 ports are backward-compatible with the PCIe 1.0 transfer rate.

"The 11th-generation PowerEdge server architecture is designed for high performance and optimized energy efficiency."

[1] For more information on DDR3 memory in 11th-generation PowerEdge servers, see "Optimizing DDR3 Memory Settings in New 11th-Generation Dell PowerEdge Servers," by Paul Benson, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20080414-Benson.pdf.

Dell PowerEdge M610 and PowerEdge M710 blade servers (left) offer 11th-generation performance and efficiency advantages in a rack-dense blade form factor

THE DELL POWEREDGE M-SERIES MODULAR EVOLUTION

The launch of 11th-generation Dell PowerEdge servers includes two new blade servers: the PowerEdge M610 and PowerEdge M710. These new blades can work side by side with any previous PowerEdge M-Series blades without restrictions on power, cooling, or management within the PowerEdge M1000e modular blade enclosure.

New 11th-generation PowerEdge servers are designed to utilize similar features and technologies across the entire line to give organizations the choice of server form factor best fitting their data centers. Organizations typically begin considering blade servers when needing more than six servers and when trying to overcome challenges related to power and cooling, cable management, rapid deployment, or server space. In comparison with Dell rack-mount servers, Dell blade servers typically use approximately 30 percent less power per server, require up to 94 percent less cabling, and have 60 percent higher density, all while allowing much faster physical deployment.

Both the PowerEdge M610 and PowerEdge M710 are designed with the same server architecture and features as the 11th-generation rack and tower servers. Physically, the major difference between the blades and the other form factors is that the blades do not handle standard PCI cards and typically hold fewer hard drives, all while using 208/240 V power for the chassis. Although limitations exist in a few areas when achieving high server density, both blades were designed from scratch to help solve two major enterprise challenges: power and cooling, and server consolidation through virtualization.

The PowerEdge M610 is the direct successor to the previous two-socket PowerEdge M600 blade server, and includes enhancements such as new processors and chipsets, 50 percent more dual in-line memory module (DIMM) slots, an internal Secure Digital (SD) card for embedded virtualization, and enhanced server management. It uses one slot in a PowerEdge M1000e enclosure, and 16 can fit into this enclosure in total. The PowerEdge M610 is well suited for standard applications requiring high processor performance in a small space, high-performance computing clusters, virtualization, and other typical applications. The PowerEdge M610 is also designed to be more energy-efficient than the previous-generation PowerEdge M600; combined with the PowerEdge M1000e already leading the industry in power consumption per blade,* Dell expects the PowerEdge M610 to be a compelling new blade server.

The PowerEdge M710 is a two-socket blade server that uses two slots in the PowerEdge M1000e enclosure and incorporates the same features as the PowerEdge M610, but with 18 DIMM slots and twice the I/O connectivity. Designed for virtualization, databases, or other applications needing large amounts of RAM or I/O throughput, the PowerEdge M710 allows organizations that have previously used a four-socket server to use a two-socket server while helping eliminate the typical two-socket server bottlenecks of limited RAM and I/O. Because many applications charge per socket for licensing, for example, using a two-socket server can potentially help limit application costs. The PowerEdge M710 is also designed to be highly power efficient.

Both the PowerEdge M610 and PowerEdge M710 can fit into any new or existing PowerEdge M1000e enclosure. For existing PowerEdge M1000e enclosures, administrators must upgrade their Chassis Management Controller (CMC) firmware (an upgrade that does not affect servers), after which the new servers can work alongside any other M-Series blades. Both servers can handle high-speed, fully redundant I/O throughput such as end-to-end 8 Gbps Fibre Channel, 10 Gigabit Ethernet, and 40 Gbps quad data rate InfiniBand when paired with the integrated switches for each of those technologies. Both servers also support FlexAddress software, which maintains the I/O connections when switching out blades and can help reduce downtime in many environments.

The PowerEdge M610 and PowerEdge M710 are powerful new entries into the Dell M-Series portfolio. For more information, visit DELL.COM/Blades.

*For a competitive power comparison, see "SPECjbb2005 Performance and Power Consumption on Dell, HP, and IBM Blade Servers," by Principled Technologies, December 2007, DELL.COM/Downloads/Global/Products/PEdge/En/pe_blades_specjbb2005.pdf.

	PowerEdge R610	PowerEdge R710	PowerEdge T610
2.5-inch backplane	Six slots	Eight slots	Eight slots
3.5-inch backplane	N/A	Four slots (with optional tape drive) or six slots (without optional tape drive)	Eight slots

Figure 1. *Backplane types and number of slots in Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers*

Embedded network controllers

PowerEdge R610 and PowerEdge R710 servers provide embedded network interface support through two dual-port Broadcom BCM5709C Gigabit Ethernet network controllers; the PowerEdge T610 includes one of these controllers. An x4 PCIe 1.0 link connects the controller to the IOH chipset. The ports support TCP/IP Offload Engine (TOE), IPv4, and IPv6, with Internet SCSI (iSCSI) Offload Engine (iSOE) supported with an optional hardware key. Network Controller Sideband Interface (NC-SI) support enables the embedded network interfaces to be configured for management communication or to share the management traffic with the Integrated Dell Remote Access Card 6 (iDRAC 6) if one is installed in the system.

FLEXIBLE STORAGE CONFIGURATION

PowerEdge R710 and PowerEdge T610 servers have two chassis options that support either 2.5- or 3.5-inch hard drives; the PowerEdge R610 supports only 2.5-inch drives, which enables the server to include six hard drive bays for internal storage expansion. These chassis options support multiple hard drive configurations and a choice of integrated storage controllers: either the Dell Serial Attached SCSI (SAS) 6/iR or the Dell PowerEdge Expandable RAID Controller (PERC) 6/i. The integrated storage controller can be installed only in a dedicated PCIe storage slot.

The hard drives can be SAS, Serial ATA (SATA), or SATA solid-state drives designed for speeds of up to 3 Gbps. The system storage infrastructure is also designed to support speeds of up to 6 Gbps with appropriate cables, enabling administrators to convert the hard drive configuration to use 6 Gbps drives and the next generation of Dell storage controllers when they become available.

The SAS 6/iR is a 3 Gbps SAS controller that incorporates two four-channel SAS connections to SAS and SATA hard drives. It supports RAID levels 0 and 1. The PERC 6/i is a hardware RAID controller with an LSI 1078 RAID-on-a-chip (ROC) processor, a PCIe host interface designed to support up to 2.5 GT/sec, and 256 MB of error-correcting code (ECC) DDR2 memory at 667 MHz. A battery backup unit enables memory contents to be maintained for up to 24 hours if the system loses power. This controller supports RAID levels 0, 1, 5, 6, 10, 50, and 60.

Backplane

The backplane has two x4 mini-SAS connectors for cable connections to the storage controller as well as a power connector to connect to the system board (in PowerEdge R610 and PowerEdge R710 servers) or to the power distribution board (in PowerEdge T610 servers). Mini-SAS connectors are high-density, low-profile connectors designed to support data rates of up to 6 Gbps. Each hard drive slot has two LED indicators per drive slot: one for power and status and one for activity. Different servers can use different backplanes with different numbers of slots, depending on the system and chassis (see Figure 1).

Internal hard drives

PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers support SAS, SATA, and SATA solid-state hard drives. Mixing SAS and SATA drives in one server requires installing SAS drives as a pair in drive slots 0 and 1; the remaining drives must be SATA drives. Chassis options that support 3.5-inch drives can take 15,000 rpm SAS drives and 5,400 rpm or 7,200 rpm SATA drives; those that support 2.5-inch drives can take 10,000 rpm or 15,000 rpm SAS drives, 7,200 rpm SATA drives, and solid-state drives. Chassis options that support 3.5-inch drives also allow mixing 2.5- and 3.5-inch drives, enabling administrators to use a 3.5-inch hard drive carrier adapter to install a pair of 2.5-inch, 10,000 rpm SAS drives in drive slots 0 and 1. The remaining hard drives must be 3.5-inch drives and must all be either SAS or SATA drives.

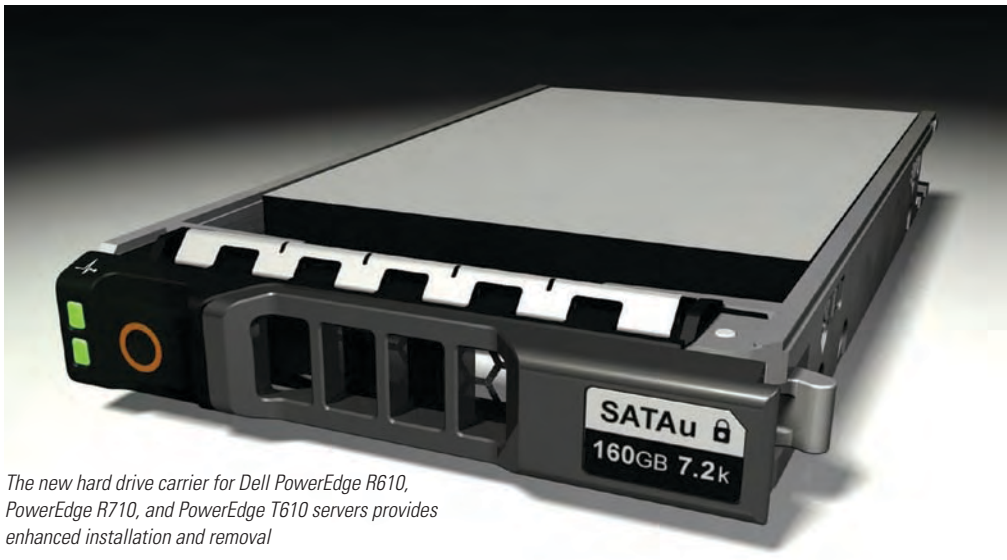

The new hard drive carrier for Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers provides enhanced installation and removal

The new hard drive carrier is designed to match the new industrial design of the servers. A release button replaces the previous-generation release latch for enhanced drive installation and removal. Administrators should install hard drive carrier blanks in any empty drive slots to help maintain system cooling and to provide shielding against electromagnetic interference. Each hard drive carrier has two LED indicators visible from the front of the system, which function similarly to LEDs on previous-generation servers.

PowerEdge R610 and PowerEdge R710 servers also support a diskless configuration with no storage controller installed in the system. This configuration still includes a backplane, enabling administrators to later upgrade their systems with a storage controller and hard drives. Administrators can configure the servers to boot from an external storage controller in a Fibre Channel or iSCSI storage area network, or from an embedded VMware® ESXi hypervisor on the internal Secure Digital (SD) module.

Optical drive

An optional optical drive can connect to the system board through the SATA interface. PowerEdge R610 and PowerEdge R710 servers can support internal slim-line DVD and DVD±RW drives; the PowerEdge T610 server can support an optional 5.25-inch optical drive. Parallel ATA (PATA), or IDE, optical drives are not supported. If an optical drive is not ordered with the server, a blank is installed in its place.

Tape drives

The PowerEdge R610 server does not support internal tape drives. The PowerEdge R710 server can support 3.5-inch half-height internal tape drives in systems with the 3.5-inch, four-slot backplane configuration or 2.5-inch, eight-slot backplane configuration. The PowerEdge T610 server supports 5.25-inch half-height and full-height tape drives. Because the full-height tape drive occupies two drive bays, an optical drive cannot be installed if a full-height tape drive is used in this system. Internal SATA tape drives connect directly to the SATA connector on the system board, internal SCSI tape drives connect to the LSI 2032 PCIe SCSI adapter card, and internal SAS tape drives connect to the SAS 5/iR PCIe adapter card. Blanks are installed in any empty drive bays.

HIGH-SPEED EXPANSION SLOTS

The PCIe slots in PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers are connected to the Intel 5500 or Intel 5520 IOH. In PowerEdge R610 and PowerEdge R710 servers, the slots are on the expansion riser board, while in the PowerEdge T610 server, they are on the system board. The slots are PCIe 2.0 capable, designed to support speeds of up to 5 Gbps; with 8B/10B encoding, the effective data transfer rate is typically 4 Gbps.

	PowerEdge R610	PowerEdge R710	PowerEdge T610
x4 PCIe	None	Slot 1 (full height, full length) and slot 2 (three-fourths height, full length)	Slots 1, 4, and 5 (full height, half length)
x8 PCIe	Slots 1 and 2 (full height, half length)	Slots 3 and 4 (full height, full length)	Slots 2 and 3 (full height, full length)
x16 PCIe	None	Optional slot 3 (full height, full length)	None

Note: Full-length cards are 9.5 inches long, except for slot 1 in the PowerEdge R710 server and slots 2 and 3 in the PowerEdge T610 server, which can accept cards up to 12.2 inches long. Half-length cards are 6.6 inches long or less.

Figure 2. *Available PCIe expansion slots in Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers*

The PowerEdge R710 also supports an optional expansion board that combines slots 3 and 4 into an x16 PCIe slot, with power limited to 25 W. Administrators can use this slot to connect to external general-purpose graphics processing unit (GPGPU) devices.

Figure 2 lists the available expansion slots for each server. In addition to these expansion slots, each server includes one x4 PCIe 1.0 slot dedicated to the integrated storage controller.

COMPREHENSIVE SYSTEMS MANAGEMENT

The iDRAC 6 in PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers supports traditional hardware monitoring and power control previously handled by the baseboard management controller (BMC). It introduces several components to enhance systems management: the iDRAC 6 Express persistent storage device, an optional iDRAC 6 Enterprise card, and an optional unmanaged persistent storage device. A system control panel and LCD provide easy access to key system information and configuration options.

iDRAC 6 Express

The iDRAC 6 Express is an integrated persistent storage device serving as a managed system services repository. The Dell Unified Server Configurator in the iDRAC 6 Express can help administrators perform system deployment, management, updating, servicing, and diagnostics, helping reduce the number of tools required for media-based provisioning.[2]

Coupled with Dell OpenManage™ software, the iDRAC 6 Express can provide a comprehensive systems management solution. Dell OpenManage 6.0.1 introduces multiple enhancements and advanced features focused on simplifying systems management without compromising functionality, including features designed for system configuration,

[2] For more information, see "Simplify Management with the Dell Unified Server Configurator Enabled by the Lifecycle Controller," by Shelli Allgood, Anand Narayanan, Hai Phung, and Pritesh Prabhu, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090226-Phung.pdf.

deployment, and change management; hardware configuration and updating; backups; third-party toolkits; power and virtualization management; and standard instrumentation for interoperability with solutions from a variety of vendors.

iDRAC 6 Enterprise

The optional iDRAC 6 Enterprise card provides access to advanced features. The card contains an RJ-45 management 10/100 Mbps Ethernet connector and an externally accessible virtual flash media card slot. Its management features include remote access management through the racadm command-line interface, Intelligent Platform Management Interface (IPMI), Web server, Secure Shell (SSH), Telnet, and Serial Over LAN; graphical console redirection through remote vKVM (virtual keyboard, video, mouse); remote virtual floppy, CD, and disk (vMedia); and virtual flash with an SD card in the external SD slot. The RJ-45 connector provides out-of-band management and can be configured as a dedicated network interface port or as a shared interface with the embedded network controllers to support failover.

Unmanaged persistent storage device

The unmanaged persistent storage device consists of two ports: one internal USB port and one internal SD module port. The internal USB port is for an optional USB drive and is located inside the chassis. Administrators can use this drive for storage of custom logs or scratch pads for portable administrator-defined information (although the drives are not hot-pluggable), custom boot and preboot operating systems to help simplify deployment or to support diskless environments, and so on.

The internal SD card slot is dedicated for an embedded hypervisor located on the internal SD module. The internal SD card can contain a bootable OS image for virtualized platforms.

Control panel and LCD

The system control panel is located on the front of the system chassis to provide easy access to switches, the display, and I/O interfaces. Located on the control panel are two USB 2.0 connectors, a standard 15-pin VGA connector (not available in the PowerEdge T610), an Advanced Configuration and Power Interface (ACPI)–compliant power button with integrated green power LED controlled through the iDRAC firmware, and a 128-by-20-pixel LCD. To help reflect the true ambient temperature, a temperature sensor is located on the control panel board.

The system LCD is a graphical display controlled by the iDRAC firmware; both the iDRAC firmware and BIOS can send error codes and messages to this display. The LCD also has two navigation buttons, a select button, and a system ID button, which administrators can use to set up the embedded network controller IP address, Media Access Control (MAC) address, system power information, system tag, and so on to identify the system.

The system BIOS and iDRAC Configuration Utility can enter a secure mode that locks the LCD navigation and select buttons as well as the system power and non-maskable interrupt (NMI) buttons. In this mode, for example, the power button cannot be used to turn off the system.

POWERFUL, VERSATILE SECURITY

In addition to providing a system BIOS setup password and system bezel lock, PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers provide enhanced security through a Trusted Platform Module (TPM). The TPM is a microcontroller designed to protect confidential information against external software attacks or physical theft—including generating and storing keys, protecting and authenticating passwords, and creating and storing digital certificates. The TPM is affixed to the motherboard and is compliant with the Trusted Computing Group TPM 1.2 specification. Administrators can enable and activate it through a BIOS setup option.

In addition to protected storage, the TPM also provides digital signature and attestation functionality. The attestation functionality allows the digital signature and protected storage features to function only if the server is in a known good software state. Administrators can also use the TPM to store the encryption keys for the Microsoft® BitLocker™ hard drive encryption feature in the Microsoft Windows Server® 2008 OS.

ENERGY-EFFICIENT POWER SUPPLIES

PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers are designed for energy efficiency. The power supply subsystem supports up to two high-efficiency AC/DC power supply units (PSUs) in a 1 + 1 redundant configuration.

PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers introduce the ability to customize server power capacity. The base system consists of one Dell Energy Smart PSU designed to optimize efficiency for typical configurations. Energy Smart PSUs can enable administrators to replace legacy servers without needing to upgrade their data center power systems. Because most system configurations can use the Energy Smart PSUs, these PSUs also help avoid the need to sacrifice performance or features.

During the boot process, the system determines whether the PSUs can generate enough power for the installed hardware. If they cannot, the LCD panel and screen display a warning at the end of the power-on self-test (POST) before booting the OS. The system then runs using the current configuration after reducing the speed of components such as the processor and memory.

Optional High Output PSUs are available for use with high-end configurations that exceed the power usage supported by the Energy Smart PSUs. These High Output PSUs use the same form factor as the Energy Smart PSUs; because no system modifications are necessary to switch

	PowerEdge R710 and PowerEdge T610		PowerEdge R610	
	Energy Smart PSU	High Output PSU	Energy Smart PSU	High Output PSU
Dimensions (length × width × height)	206.4 mm × 67.5 mm × 76.5 mm		249 mm × 65.5 mm × 38.2 mm	
Status indicators	One two-color LED			
Integrated fans	One 60 mm fan		None	
Fixed input plug	IEC C14			
AC cord rating	15 A at 120 VAC, 10 A at 240 VAC			
Input voltage	90–264 VAC			
Auto-ranging	Yes			
Line frequency	47–63 Hz			
Maximum inrush current	55 A per PSU for 10 ms or less			
Hot-swap and hot-add capability	Yes			
Output power	570 W	870 W	502 W	717 W
Maximum heat dissipation	1,944.9 BTU/hour	2,968.6 BTU/hour	1,712.9 BTU/hour	2,446.5 BTU/hour
Efficiency	86.9%–90.5% at 115 VAC, 88%–92% at 230 VAC	85%–89% at 115 VAC, 87%–90% at 230 VAC	86.5%–90% at 115 VAC, 88%–92% at 230 VAC	85%–88.5% at 115 VAC, 86.5%–90.5% at 230 VAC

Note: All figures based on product specifications; results will vary based on configuration, usage, and manufacturing variability. Dimensions do not include the power supply handle or ejection tab.

Figure 3. *Power supply specifications for Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers*

between these types, administrators that upgrade their servers to a high-end configuration can simply replace the PSUs on their own. The servers can detect a PSU mismatch in which both an Energy Smart PSU and a High Output PSU are installed, resulting in a nonredundant configuration. If a mismatch is detected during the boot process, the LCD panel and screen display a warning at the end of the POST. The system then continues to boot after an administrator has acknowledged the mismatch. If a mismatch occurs when hot swapping or hot adding a PSU, the LCD panel displays a message and the status LED on the PSU shows alternating green and yellow colors.

Figure 3 details the specifications for both types of PSUs in each server model. Both types can offer significant efficiency enhancements over previous-generation PSUs, especially at low loads, and both types are 80 PLUS certified and compliant with Climate Savers Computing requirements.[3]

INTELLIGENT COOLING

PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers are designed for intelligent and efficient cooling, with system chassis that offer increased outflow venting and pulse-width modulation (PWM) fans. Fans in the PowerEdge R710 server are hot swappable. In PowerEdge R610 and PowerEdge R710 servers, a dual-processor system requires all system fans, while a single-processor system can be operated with one fewer fan. The PowerEdge T610 server uses two fans mounted in a cooling shroud for the processor and memory, and can support two additional fans to add redundancy. PSUs in PowerEdge R710 and PowerEdge T610 servers also include an integrated fan. A cooling shroud installed over the processors and memory directs airflow over those components and allows for cover-off operation for short periods of time.

The system fans in PowerEdge R710 and PowerEdge T610 servers have blind-mate connectors and are mounted in a bracket and in a shroud, respectively; in the PowerEdge R710, administrators can insert and remove individual fans without having to remove the bracket. For ease of access to the system board, the entire fan bracket can be removed without having to remove the individual fans. Figure 4 lists the system fan configurations for each server model.

Fan speed is controlled by the iDRAC firmware, which can intelligently determine the fan speed based on the system configuration and ambient temperature. New power management features offer different fan profiles, allowing administrators to operate the fans in minimum-power or maximum-performance modes. The minimum-power

[3] For more information on energy-efficient features in PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers, see "Optimizing the Data Center: New Dell Servers and the Dell Energy Smart Architecture," by Daniel Bounds, John Jenne, and Robert Hormuth, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090247-Bounds.pdf. For more information on 80 PLUS and Climate Savers Computing, visit www.80plus.org and www.climatesaverscomputing.org.

mode allows the fans to operate at reduced speeds, with reduced power consumption and noise, while still maintaining sufficient airflow for system cooling.

ENHANCED CHASSIS DESIGN

PowerEdge R610, PowerEdge R710, and PowerEdge T610 server chassis designs introduce several enhancements, including the following:

- Updated industrial design including a new LCD, bezel, and hard drive carrier
- Tool-less rack latches
- Pull-out tray for the Dell Service Tag and other labels in PowerEdge R610 and PowerEdge R710 servers
- Support for persistent storage (internal USB drive, internal SD module, and external virtual flash media slot)
- Updated PSU removal process
- Updated cable management arm
- Support for multiple types of rails, including static, sliding, and two-post
- Two PowerEdge R710 and PowerEdge T610 chassis options for 2.5- and 3.5-inch hard drives

The PowerEdge R610 has a 1U rack-mount form factor and supports only 2.5-inch hard drives. The PowerEdge R710 has a 2U rack-mount form factor. The PowerEdge T610 uses a tower-optimized chassis that can be converted to fit a rack-mount 5U form factor.

NEXT-GENERATION SERVER DESIGN

Optimized for high performance and energy efficiency, the new 11th-generation Dell PowerEdge servers are designed to take advantage of the latest advances in performance, power and cooling, and systems management. With versatile rack-mountable chassis options and support for multiple types of peripheral devices, these servers are designed to support virtually any type of data center application and can help simplify management, control power and cooling requirements, optimize hardware resources, and reduce total cost of ownership for organizations of all sizes.

New 11th-generation Dell PowerEdge servers feature an enhanced chassis design for flexible data center deployment

Edward Yee is a development engineer at Dell. He has a B.S. in Electrical Engineering from the United States Air Force Academy and an M.S. in Electrical and Computer Engineering from the Georgia Institute of Technology.

Indrani Paul is a systems engineer on the Dell server development team. She has a B.S. in Electronics and Electrical Communication Engineering from the Indian Institute of Technology, Kharagpur, and an M.S. in Electrical Engineering from the Georgia Institute of Technology.

Robert Tung is a systems engineer on the Dell server development team. He has a B.S. in Electrical Engineering from Texas A&M University.

Truc Nguyen is a senior systems engineer on the Dell server development team. He has a bachelor's degree in Electrical Engineering from Rensselaer Polytechnic Institute and a master's degree in Electrical Engineering from Syracuse University.

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.

	PowerEdge R610	PowerEdge R710	PowerEdge T610
Single-processor configuration	Fans 1–5	Fans 1–4	Fans 3 and 4 (with optional fans 1 and 2 for redundant cooling)
Dual-processor configuration	Fans 1–6	Fans 1–5	Fans 3 and 4 (with optional fans 1 and 2 for redundant cooling)

Figure 4. *System fan configurations for Dell PowerEdge R610, PowerEdge R710, and PowerEdge T610 servers*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell PowerEdge blade servers:
DELL.COM/Blades

OPTIMIZING THE DATA CENTER: NEW DELL SERVERS AND THE DELL ENERGY SMART ARCHITECTURE

By Daniel Bounds
John Jenne
Robert Hormuth

The Dell™ Energy Smart architecture, including new 11th-generation Dell PowerEdge™ servers, provides a comprehensive, system-wide approach to balancing performance with energy efficiency. By focusing on four key tenets—design, measurement, control, and reporting—this architecture can support scalable, energy-efficient infrastructures that can help optimize performance per watt and reduce total cost of ownership.

Related Categories:

Blade servers
Dell PowerEdge blade servers
Dell PowerEdge rack servers
Dell PowerEdge servers
Dell PowerEdge tower servers
Energy efficiency
Green IT
Intel
Power and cooling
Servers

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

Limitations on space, power, and cooling capacity combined with rising energy costs present enormous challenges for IT environments. Even as IT departments deploy high-performance, high-density servers to maximize compute capacity without expanding their server footprint, they must also find ways to effectively cool these power-intensive systems—all while controlling operational costs and contending with increasingly restricted IT budgets.

Overcoming power and cooling problems should not require reworking the entire IT infrastructure, purchasing and deploying entirely new technologies, or relying on expensive consultants to guide the time-consuming process of redesigning or rebuilding a data center. Instead, by following best practices to help maximize overall data center productivity—taking advantage of key metrics to evaluate the infrastructure, optimizing data center design for energy efficiency, implementing virtualization to help maximize hardware utilization, and following aggressive hardware refresh strategies—organizations can optimize their resources to help maximize computing power, minimize energy use, and simplify systems management.[1]

A key part of the Dell approach to overcoming the challenges of data center power and cooling is dynamically balancing actual work performed against power consumption, maximizing measures such as performance per watt, capacity per watt, and throughput per watt. These strategies enable existing data centers to support increased computing power while reducing space, energy use, and cooling requirements—unleashing the true potential of the data center.

Dell Energy Smart technologies span a comprehensive range of Dell hardware, software, and services to help organizations implement this approach throughout the data center. Integrated into new 11th-generation Dell PowerEdge servers, these technologies are designed to dynamically manage system performance, power, and thermals at the platform level to help optimize performance per watt and reduce total cost of ownership as part of an overall strategy for energy efficiency.

DESIGNING FOR ENERGY EFFICIENCY

The Dell Energy Smart architecture in 11th-generation Dell PowerEdge servers is built on four core

[1] For more information on best practices for evaluating and enhancing data center efficiency, see "Compute More, Consume Less: Smart Policies Unleash Data Center Productivity," in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090176-Esser.pdf.

tenets: *design*, *measurement*, *control*, and *reporting*. This architecture is primarily based on an out-of-band implementation that supports real-time monitoring and control, and uses extensive platform characterization and advanced features such as Intel® Dynamic Power Technology to help optimize platform-level trade-offs between performance and power consumption.

Design

The holistic, system-wide approach of Dell Energy Smart technologies starts with following efficient design principles across electrical, mechanical, and thermal systems and combining them with intelligent component selection. Electrical contributions to power efficiency include highly optimized power supply units (PSUs), highly efficient voltage regulators (which include static and dynamic phase shedding capabilities), and low-loss printed circuit board layouts and connectors.

Mechanically, Energy Smart technologies are designed for optimized venting and airflow through the platform, helping minimize fan use. Thermal technologies are designed for optimized fan zones, low-power fans, and adaptive fan-speed algorithms. Configuration options for low-power components such as processors, memory, and hard drives help further enhance energy efficiency.

Measurement

Measurement is an integral part of the Dell Energy Smart architecture. Real-time, highly accurate measurements of performance, power consumption, and thermals feed into the Energy Smart control algorithms, enabling the system to make intelligent decisions to help optimize performance per watt. Comprehensive temperature measurements of the ambient air, processors, memory, PSUs, and other areas help optimize fan speeds to the environment and server workload. In addition to these wide-ranging monitoring capabilities, Energy Smart technologies are designed to provide leading-edge monitoring accuracy.

Control

The Dell Energy Smart architecture incorporates firmware running on a high-performance baseboard management controller (BMC) embedded in the system to provide an intelligent, centralized control mechanism. The BMC utilizes Energy Smart measurement technologies to analyze component and platform trade-offs to help determine performance-per-watt optimizations. Power capping, for example, provides a key way for administrators to control energy use by enabling them to define maximum thresholds for power consumption in a given system. Similarly, the Energy Smart thermal algorithm may limit memory throughput rather than allow fan speeds to reach the upper power regions of the fan response curve.

Reporting

The Dell Energy Smart architecture provides a variety of reporting metrics through in-band and out-of-band management interfaces. The Dell Management Console Powered by Altiris™ from Symantec™ provides in-band reporting, while the standard Web Services for Management (WS-Management) protocol provides out-of-band reporting. The architecture is designed to support real-time collaboration between the platform and OS, enabling compatible operating systems to use performance, power, and thermal constraints in their algorithms to make intelligent decisions and help meet scheduling requirements and service-level agreements. Administrators can also view and graph key metrics through the Integrated Dell Remote Access Controller (iDRAC) in 11th-generation Dell PowerEdge servers.

UNDERSTANDING DELL ENERGY SMART TECHNOLOGIES

New 11th-generation Dell PowerEdge servers feature a variety of energy-tuned technologies designed to increase performance while reducing power consumption. Key features include efficient Energy Smart PSUs, Energy Smart system design, the Dell Active Power Controller (DAPC), Energy Smart power management, and high-efficiency processors and memory.

Energy Smart power supplies

Overprovisioned, inefficient server PSUs that draw more power than necessary are common in enterprise data centers. Dell Energy Smart PSUs are engineered for high efficiency and rightsized for typical environments. Design enhancements enable these PSUs to provide higher efficiency than the PSUs of comparable previous-generation Dell PowerEdge servers (see Figures 1 and 2). For 11th-generation PowerEdge M610 and PowerEdge M710 blade servers, dynamically provisioned PSUs help further avoid unnecessary overhead. Energy Smart PSUs can help reduce

"A key part of the Dell approach to overcoming the challenges of data center power and cooling is dynamically balancing actual work performed against power consumption, maximizing measures such as performance per watt, capacity per watt, and throughput per watt."

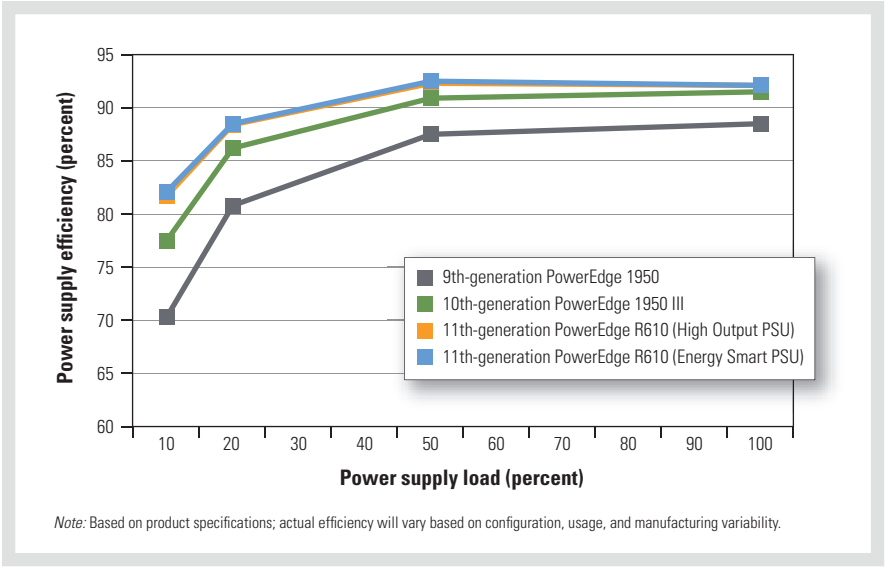

Note: Based on product specifications; actual efficiency will vary based on configuration, usage, and manufacturing variability.

Figure 1. *Power supply efficiency for comparable 1U Dell PowerEdge servers at 230 VAC*

system power consumption by up to 8.4 percent compared with previous-generation Dell server PSUs.[2]

The majority of typical 11th-generation Dell server configurations can be supported by Energy Smart PSUs, helping to immediately lower energy requirements at the server level. For organizations with higher-than-normal performance requirements, optional High Output PSUs are also available for 11th-generation PowerEdge servers. In either case, rather than being locked into a single PSU for a given server, IT departments can choose the PSU that meets their needs. Both Energy Smart and High Output PSUs are 80 PLUS certified, meet Climate Savers Computing standards, and are expected to meet the U.S. government's Energy Star program requirements.[3]

Energy Smart system design

The breakthrough system-level Dell Energy Smart design provides a simplified solution to help overcome inefficiency at the core of the data center: the systems themselves. Incorporating multiple engineering enhancements over previous generations—including high-efficiency voltage regulators, increased venting and airflow, low-flow fan algorithms, and enhanced resource management—this design can help lower overall power consumption and optimize performance per watt in 11th-generation Dell PowerEdge servers. As a whole, Energy Smart system design enhancements can help reduce system power consumption by up to 16 percent compared with previous-generation PowerEdge servers.[4]

Dell Active Power Controller

The DAPC is an OS-independent processor power manager designed to maximize performance per watt starting from the moment the server is powered up. By monitoring the processor and lowering system-level power draw at times of low utilization, it can help reduce system power consumption by up to 15.7 percent compared with OS power management alone.[5]

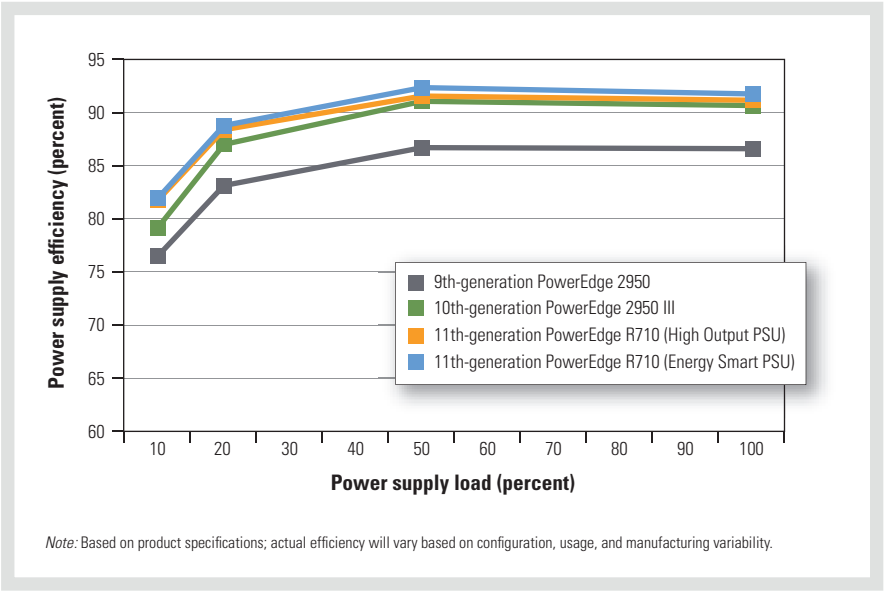

Note: Based on product specifications; actual efficiency will vary based on configuration, usage, and manufacturing variability.

Figure 2. *Power supply efficiency for comparable 2U Dell PowerEdge servers at 230 VAC*

[2] Based on AC power measurements taken by Dell Labs in November 2008 using an Extech 380803 Power Analyzer and comparing total AC load results for a Dell PowerEdge 1950 III PSU with those of a Dell PowerEdge R610 Energy Smart PSU using redundant configurations and a fixed total DC load of 75 W. Actual performance and power consumption will vary based on configuration, usage, and manufacturing variability.

[3] For more information, visit www.80plus.org, www.climatesaverscomputing.org, and www.energystar.gov.

[4] Based on AC power measurements taken by Dell Labs in July and November 2008 using an Extech 380803 Power Analyzer and comparing a Dell PowerEdge 1950 III with a Dell PowerEdge R610. The stated energy savings includes contributions by voltage regulation and fan improvements, respectively. Measurements were taken at 25°C ambient temperature and at 20 percent system load and idle, respectively. The PowerEdge 1950 III was configured with two quad-core Intel Xeon 5160 processors at 3.00 GHz, eight 1 GB fully buffered dual in-line memory modules (DIMMs) at 667 MHz, and two 73 GB, 10,000 rpm, 3.5-inch Serial Attached SCSI (SAS) drives, and ran the Microsoft® Windows Server® 2003 Enterprise x64 Edition OS; the PowerEdge R610 was configured with two quad-core Intel Xeon X5570 processors at 2.93 GHz, twelve 4 GB DDR3 DIMMs, and six 10,000 rpm, 2.5-inch SAS drives, and ran Windows Server 2008 Enterprise Edition. Actual performance and power consumption will vary based on configuration, usage, and manufacturing variability.

[5] Based on testing performed by Dell Labs in February 2009 using an industry-standard performance/power benchmark comparing the Dell PowerEdge R710 running Microsoft Windows Server 2008 with SP1 with DAPC enabled versus Microsoft Windows® Balanced power management enabled.

> **"The Dell Energy Smart architecture in 11th-generation Dell PowerEdge servers can provide significant improvements in performance per watt and total cost of ownership compared with previous-generation PowerEdge servers."**

Energy Smart power management

Incorporating a significant increase in features compared with 10th-generation Dell PowerEdge servers, intelligent Dell Energy Smart management in 11th-generation PowerEdge servers includes features such as power capping, advanced power policies, power scheduling, and device disablement. Power capping allows administrators to define specific power limits to aid power management, after which system-level controls can monitor actual system energy use to help maintain power consumption below this threshold. By enabling the intelligent management of overall system power usage, the Energy Smart architecture can help significantly enhance data center efficiency.

High-efficiency processors and memory

New 11th-generation Dell PowerEdge servers incorporate advanced processors and Double Data Rate 3 (DDR3) memory technologies to help maximize performance and efficiency. The multi-core Intel Xeon® processor 5500 series is designed to provide up to 2.25 times the performance of the existing Intel Xeon processor 5400 series,[6] and can support Intel SpeedStep® Technology, Intel Virtualization Technology, simultaneous multi-threading, and deep C-states for enhanced power conservation. DDR3 memory helps provide a high-performance interface capable of low-latency response and high throughput.[7]

OPTIMIZING PERFORMANCE PER WATT AND TOTAL COST OF OWNERSHIP

The Dell Energy Smart architecture in 11th-generation Dell PowerEdge servers can provide significant improvements in performance per watt and total cost of ownership compared with previous-generation PowerEdge servers. For example, in March 2009, Dell engineers used an industry-standard performance/power benchmark to evaluate the performance and power consumption of a PowerEdge R710 server with two quad-core Intel Xeon X5570 processors. Compared with tests performed by Dell engineers in February 2008 on a previous-generation PowerEdge 2950 III server with quad-core Intel Xeon E5440 processors, the PowerEdge R710 used less power at comparable performance levels and was able to reach higher maximum performance levels (see Figure 3)—thereby offering significantly higher performance per watt than the previous-generation server.

Dell's own IT infrastructure has benefited significantly from a comprehensive approach to energy efficiency that includes pervasive virtualization, system consolidation, hardware refresh, containment strategies, and Energy

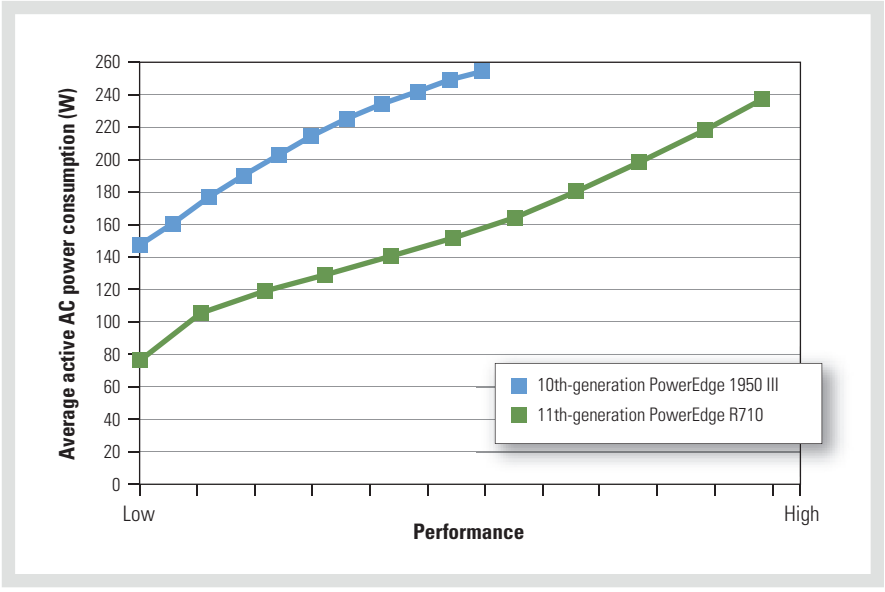

Figure 3. *Performance and power consumption results for a Dell PowerEdge 2950 III server and PowerEdge R710 server*

[6] Based on Intel internal measurements in February 2009; for more information, see "Intel Xeon Processor 5500 Series: An Intelligent Approach to IT Challenges," by Intel, www.intel.com/assets/pdf/prodbrief/321579.pdf. Performance tests and ratings are measured using specific computer systems and/or components and reflect the approximate performance of Intel products as measured by those tests. Any difference in system hardware or software design or configuration may affect actual performance. Buyers should consult other sources of information to evaluate the performance of systems or components they are considering purchasing. For more information on performance tests and on the performance of Intel products, visit www.intel.com/performance/resources/limits.htm.

[7] For more information on processors and other enhancements in 11th-generation PowerEdge servers, see "Data Center Workhorses: New Dell PowerEdge Rack and Blade Servers," by Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090246-Nguyen.pdf. For more information on DDR3 memory in 11th-generation PowerEdge servers, see "Optimizing DDR3 Memory Settings in New 11th-Generation Dell PowerEdge Servers," by Paul Benson, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20080414-Benson.pdf.

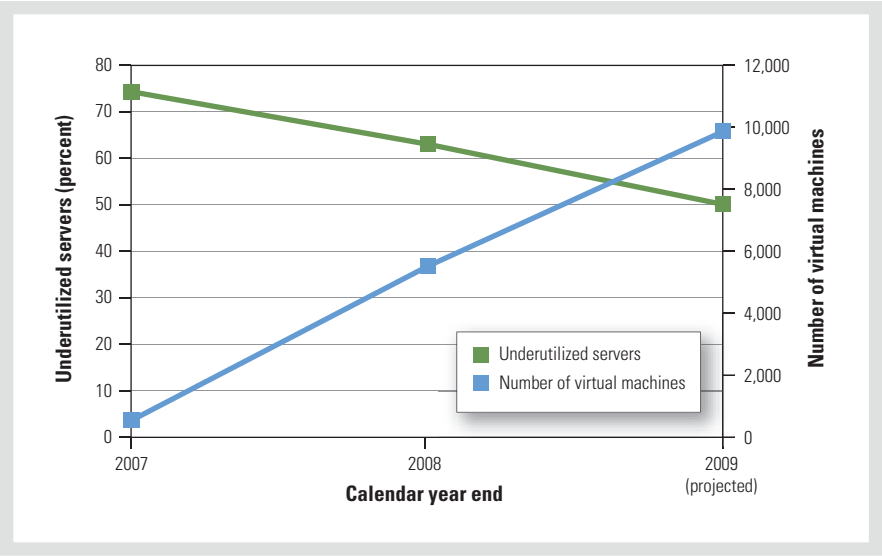

Figure 4. *Total server underutilization and number of virtual machines in the Dell IT environment*

Smart technologies. Figure 4, for example, shows the total percentage of underutilized servers in the Dell IT environment and the number of virtual machines (VMs) being supported by physical hosts at the end of 2007 and 2008, along with projections for the end of 2009. As this figure shows, underutilization has dropped significantly as the number of VMs has grown, helping maximize use of available resources. Dell estimates that its energy-efficiency strategies had saved a total of US$38 million in energy and other costs by the end of 2008, and expects that total to increase even further by the end of 2009. By optimizing its environment for efficiency, in fact, the company was able to avoid the need to build a new data center—a major cost savings.

Upgrading to 11th-generation PowerEdge servers with Energy Smart technology can help lower total cost of ownership in a variety of ways. Organizations can take advantage of the online Dell Server Power and Space Savings Calculator, available at DELL.COM/Switch, to help them evaluate different upgrade scenarios and potential cost savings in their own environments. For example, a single standard 42U server rack with a maximum power consumption of 5 kW can support up to 16 Dell PowerEdge R710 servers, compared with only 12 HP ProLiant DL380 G5 servers. In a data center with 100 servers, this higher rack density means that the Dell servers would require three fewer racks than the HP servers, with a corresponding reduction in power consumption in addition to the advantages of Energy Smart technology. In this scenario, the Dell tool estimates that deploying Dell PowerEdge R710 servers would save up to US$29,947 in energy costs over one year, or up to US$89,841 over three years, compared with the HP ProLiant DL380 G5 servers—a substantial savings.[8]

CREATING SCALABLE, ENERGY-EFFICIENT INFRASTRUCTURES

The Dell Energy Smart architecture provides a comprehensive, system-wide approach to balancing performance and energy use in new 11th-generation Dell PowerEdge servers—capabilities that are critical to simplifying IT and meeting the power and cooling challenges of enterprise data centers. As performance requirements continue to increase, this optimized architecture can help administrators create scalable, energy-efficient infrastructures to support ongoing growth.

Daniel Bounds is a solutions manager for the Dell Enterprise Product Group focused on power and cooling. Daniel has previously held positions with the Dell Global Relationship Marketing Group and Hewlett-Packard. He has a B.A. and an M.B.A. from the University of Texas at Austin.

John Jenne is a system lead in the Dell Enterprise Product Group focused on mainstream rack servers. John has previously held positions with Compaq, MaXXan, and Newisys. He has a B.S. in Computer Engineering from Clemson University and an M.S.E.E. from the University of Houston, and currently holds 15 patents.

Robert Hormuth is the server architecture manager in the Dell Office of the CTO. Robert has previously held positions with Intel and National Instruments. He has a B.S. in Computer Engineering from the University of Texas at Austin.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell Energy Smart:
DELL.COM/EnergySmart

Dell PowerEdge servers:
DELL.COM/PowerEdge

[8] Based on 100 Dell PowerEdge R710 servers configured with Intel Xeon E5504 processors, 24 GB of RAM, and two hard drives, compared with 100 HP ProLiant DL380 G5 servers configured with quad-core Intel Xeon E5410 processors, 24 GB of RAM, and two hard drives. Savings assumes 100 percent system load, power consumption of 5 kW per rack, and an energy price of US$0.10/kWh.

Windfall

Citrix XenDesktop™ lets you deliver desktops as a service from a single image that can be centralized and 100% controlled in the datacenter. Delivering personalized desktops from a single image lets you cut desktop support and lifecycle costs for better TCO. You get bullet-proof security with complete control over desktop data. And you eliminate desktop complexity and replace it with a flexible solution that will enhance IT agility. Visit www.citrix.com to find out how you can deliver desktops as a service and make the management of your users more efficient.

Shred the notion of complex desktops.

Centralize desktop management and let ‘er rip

CITRIX

Deliver Desktops. Not confusion.

Visit www.citrix.com today.

CITRIX®

OPTIMIZING DDR3 MEMORY SETTINGS IN NEW 11TH-GENERATION DELL POWEREDGE SERVERS

By Paul Benson

Dell™ PowerEdge™ R610 and PowerEdge R710 rack servers, PowerEdge T610 tower servers, and PowerEdge M610 and PowerEdge M710 blade servers feature the Double Data Rate 3 (DDR3) memory architecture, which enhances flexibility by enabling up to three memory channels per processor. By understanding best-practices DDR3 configurations, administrators can optimize settings for targeted enterprise applications and virtualized environments.

Related Categories:

Blade servers

Dell PowerEdge blade servers

Dell PowerEdge rack servers

Dell PowerEdge servers

Dell PowerEdge tower servers

Memory

Servers

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

The new 11th-generation Dell PowerEdge server family takes advantage of the Intel® Xeon® processor 5500 series, which utilizes Double Data Rate 3 (DDR3) memory. By providing the flexibility to configure up to three memory channels per processor, the DDR3 architecture allows administrators to optimize memory configurations to suit specific requirements for a wide range of enterprise usage scenarios.

DETERMINING MEMORY REQUIREMENTS

With its flexible design, memory can be configured to meet the demands of a variety of usage scenarios. *High-performance mode* may be appropriate for environments in which memory speed is a top priority—such as servers running enterprise applications and supporting a virtualization infrastructure. For example, if a Web server is primarily delivering user-specific content, then new pages must be created frequently. This scenario typically calls for a high-performance memory configuration because the pages are likely not cached. Application commands must be processed as efficiently as possible, but high-capacity memory is not required because there is little reason to hold on to the result in memory past the time it is displayed to users or stored on disk.

Balanced-performance mode is applicable in a wide variety of business scenarios. Designed for situations that require the best of both worlds—good performance and midrange capacity—this configuration allows for up to twice the memory capacity available in high-performance mode.

High-capacity mode is designed to provide up to three times the memory capacity available in high-performance mode, although the trade-off is low memory speed. High-capacity configurations may be suitable for Web servers handling pages that do not vary from one user to the next. Large-capacity systems are also generally appropriate for virtualized environments and for organizations that have large databases that they want to pull into memory to help avoid increased disk access times.

Mirror mode facilitates balanced performance while providing for the correction of memory errors. It provides a memory backup, so only half of the installed memory is reported to the OS. A data integrity configuration may be appropriate, for example, for 24/7 transaction processing scenarios in which an enterprise cannot afford to take its servers offline for anything other than scheduled maintenance.

DDR3 memory can also be configured for substantial power savings when heating and cooling usage in the data center is a critical concern. Unbuffered dual in-line memory module (UDIMM) technology is typically best suited for power-conscious organizations, as it is designed to consume 1 W less power per memory module than registered DIMM (RDIMM) technology.

SELECTING DDR3 OPTIONS FOR TARGET USAGE SCENARIOS

Each processor in the new 11th-generation Dell PowerEdge server family includes three separate memory controller hubs (MCHs) within the processor package. This design avoids the requirement to transfer memory transactions between the processor and an external device. However, the inclusion of three MCHs per processor introduces considerations about the appropriate way to populate and configure a system based on specific requirements for reliability, availability, and serviceability (RAS) features and speed (see the "DDR3 configuration in 11th-generation Dell PowerEdge servers" sidebar in this article).

The integrated memory controllers in the Intel Xeon processor 5500 series used in Dell PowerEdge R610 and PowerEdge R710 rack servers, PowerEdge T610 and PowerEdge T710[1] tower servers, and PowerEdge M610 and PowerEdge M710 blade servers support DDR3 technology, which is designed to provide a high-performance memory interface capable of low-latency response and high throughput. DDR3 memory is also designed to enable higher bandwidth with lower power usage than DDR2 memory.

New 11th-generation PowerEdge servers support both error-correcting code (ECC) DDR3 RDIMMs and ECC DDR3 UDIMMs at bus speeds of 800 MHz, 1,066 MHz, and 1,333 MHz. PowerEdge R610, PowerEdge T610, and PowerEdge M610 servers have 12 memory sockets split into two sets of 6 sockets, with one set for each processor; PowerEdge R710, PowerEdge T710, and PowerEdge M710 servers have 18 memory sockets split into two sets of 9 sockets, also with one set for each processor. The memory interface supports memory demand and patrol scrubbing, single-bit error correction, and multi-bit error detection.

Offering capacities up to 8 GB, RDIMMs are suitable for a large amount of memory, extensive RAS features, and maximum expandability. RDIMMs enhance expandability because they allow for three DIMMs per channel, instead of two in UDIMMs. In addition, RDIMMs offer address parity, a RAS feature designed to halt the system if an incorrect address is detected. However, unless business requirements call for very high memory capacity, UDIMM ECC is typically a cost-effective alternative to a comparable RDIMM configuration.

The memory subsystem can run in one of three modes:

- **Memory optimized:** This mode uses all three memory channels in a processor, and populates all three channels identically. Although it enables the most DIMM population flexibility and system memory capacity of the three memory subsystem modes, this mode does not support memory mirroring, nor does it support single-device data-correction (SDDC) for x8-based DIMMs. However, it does support SDDC for x4-based DIMMs.
- **Advanced ECC:** This mode joins two controllers into a lockstep mode, thus creating a 128-bit data path. The advantage of this configuration is that it allows SDDC to work for both x4- and x8-based memory devices and supports SDDC features in both types. Memory modules must be identical in size, speed, and technology in corresponding slots.
- **Mirror:** This mode populates two channels per processor identically; the third channel is left unused. The two channels operate as mirrors of each other, with

DDR3 CONFIGURATION IN 11TH-GENERATION DELL POWEREDGE SERVERS

Each processor in 11th-generation Dell PowerEdge servers has three memory controllers, and up to three dual in-line memory modules (DIMMs) can be attached to each controller (see Figure A for an example of the processor and memory layout in a Dell PowerEdge R710 rack server). In models with 12 DIMM slots, each controller has two DIMM slots assigned to it. Both Dell and Intel strongly advise configuring systems in a balanced configuration, in which each channel is similarly populated. However, administrators should keep in mind that if the system has only a single processor, any memory installed in the slots assigned for the second processor cannot be accessed.

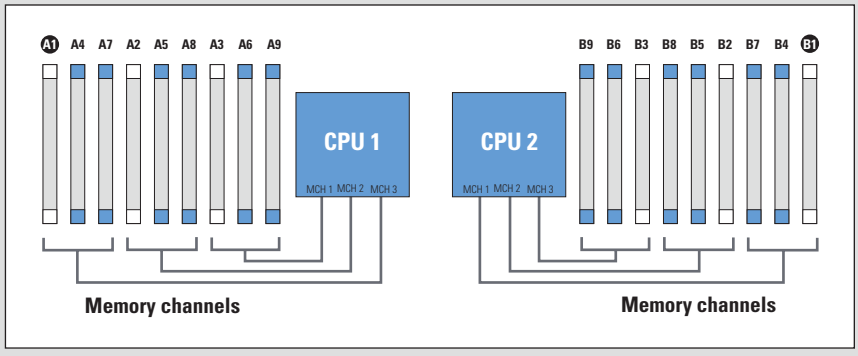

Figure A. *Dell PowerEdge R710 processor and memory layout supporting up to three DIMMs per channel*

[1] The PowerEdge T710 server was not yet released at press time; availability is planned for June 30, 2009.

writes going to both channels and reads alternating between the two channels. This mode is appropriate for typical enterprise applications. If excessive correctable errors occur or an uncorrectable multi-bit error occurs, the problematic channel is disabled and forthcoming memory accesses are routed through the functional channel. The user is informed of this event to service the system. Because the memory is mirrored, half of the installed memory is actually seen and reported by the OS.

This range of memory modes offers several potential advantages. Organizations can specify a single base server model to match specific applications or purposes, which helps simplify the server provisioning process. In addition, IT administrators have the flexibility to modify the server configuration for different applications or purposes at a later time—enhancing its long-term cost-effectiveness.

CONFIGURING DDR3 MEMORY FOR EACH USAGE SCENARIO

While specific requirements vary for individual organizations, understanding best-practices configurations for various DDR3 memory options in typical usage scenarios helps IT administrators achieve optimal performance and energy efficiency in targeted enterprise applications and virtualized environments.

As part of the memory configuration process, selecting a processor that supports required memory speed can be a key consideration. In high-performance configurations at the time of release, DDR3 memory at 1,333 MHz across all three channels per processor allows maximum bandwidth. Frequency limitations mean that this speed can support only one DIMM per channel, and thus a maximum of three DIMMs per processor. Dual-rank 4 GB RDIMMs are designed to provide peak capacity at 1,333 MHz and enable a system capacity of 12 GB per processor.[2]

System memory speed (the speed at which the memory is actually running) is set by the BIOS depending on processor capability, DIMM types used, and the number of DIMMs populated per channel. To run at maximum frequency, administrators should configure no more than one DIMM per memory channel. To help avoid errors in the memory subsystem, the system lowers the memory clock speed as DIMMs are added to each channel. Therefore, this approach supports the least amount of memory but at the highest possible speed among the configurations described in this article.

In a balanced-performance configuration, DIMMs should be installed across all three channels—but DIMMs at 1,066 MHz may be used and can be populated with two DIMMs per channel. Administrators may enable a capacity of 48 GB per processor by installing dual-rank 8 GB RDIMMs at 1,066 MHz.

As of the March 30, 2009, release date, 8 GB RDIMMs offer the highest capacity available for 11th-generation Dell PowerEdge servers. By using dual-rank 8 GB RDIMMs at 1,066 MHz, administrators may enable a capacity of 72 GB per processor for high-capacity configurations. In this case, the system is designed to automatically down-clock itself to 800 MHz because each memory channel will have three DIMMs installed.

For enterprises that require SDDC support for x8-based memory or mirroring in the data integrity configuration, only two channels from each processor should be populated. The memory channel farthest from the processor should not have any DIMMs installed—and it is vital that the DIMMs across the channels match each other.

Power-conscious IT departments can trade low capacity for reduced power consumption by using UDIMMS, which are designed to use less power than RDIMMs. However, UDIMMs are currently limited to a maximum 2 GB capacity and cannot support three DIMMs per channel. They also do not support address parity. Memory frequency is also a component in power consumption—for example, a 1,066 MHz configuration is specified to use less power than one set up as 1,333 MHz.

FINE-TUNING DDR3 FOR SPECIFIC MEMORY REQUIREMENTS

New 11th-generation Dell PowerEdge servers are designed to take advantage of the DDR3 memory architecture utilized by the Intel Xeon processor 5500 series. DDR3 is designed to support a range of performance requirements for enterprise applications and virtualized environments by enabling administrators to configure up to three memory channels per processor. Following best-practices guidelines for various usage scenarios helps administrators configure DDR3 memory options for optimal performance, capacity, data integrity, and energy efficiency.

Paul Benson was part of the 11th-generation Dell PowerEdge development team at Dell, with memory being one of his focus areas. He has a B.E. in Electrical Engineering, Computer Science, and Mathematics and an M.S. in Electrical Engineering from Vanderbilt University, as well as an M.B.A. from the University of North Carolina.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell PowerEdge servers:
DELL.COM/PowerEdge

[2] Dell plans to continue looking into enhancing systems to support higher frequencies with additional memory.

Now, upgrade your server room without overhauling it

APC integrated cooling future-proofs your IT room without breaking the bank

Is your server room a barrier to adopting new technologies?

Consolidation, virtualization, network convergence, blade servers—these new technologies improve efficiency, cut costs, and allow you to "do more with less." But they also bring high-density power, cooling, and management challenges that server rooms were never designed to handle. You're relying on guesswork, depending on building air conditioning, or improvising remedies. So, how can you increase the level of reliability and control in your server room without spending a fortune?

Introducing the APC® by Schneider Electric total server room solution

Now you can get power, cooling, monitoring, and management components that easily deploy together as a complete, integrated solution. Everything has been pre-engineered to work together and integrate seamlessly with your existing equipment. Just slide this proven, plug-and-play solution into most existing spaces—there's no need for confusing cooling configurations or expensive mechanical re-engineering. The modular, "pay as you grow" design lets you be 100% confident that your server room will keep pace with ever-changing demands.

Future-proof your server room easily, cost-effectively

APC takes the hassle out of configuring server rooms. Self-contained InRow® cooling units, high-density NetShelter® enclosures, and the APC rack air containment system combine to create a proper IT ecosystem in almost any surrounding. Rack-level monitoring sensors, intelligent controls built into the cooling unit, and integrated management software provide complete remote control and unprecedented visibility into the entire system. Simply add power protection (like undisputed best-in-class Smart-UPS® or Symmetra® units) and you have a total solution for today, tomorrow, and beyond.

The integrated, cooled, managed server room

1. **Cooling** Effective and energy-efficient InRow cooling units handle high-density heat at its source. Unique variable speed fans automatically adjust to meet changing heat loads.
2. **Power** Energy-efficient, ultra-reliable Smart-UPS and Symmetra UPSs offer scalable runtime. Rack-mount power distribution units (PDU) ensure that a wide variety of devices get plugged in and powered.
3. **Environmental Monitoring & Management** PoE-enabled temperature sensors let you keep an eye on conditions at the rack level. Metered PDUs report on aggregate power draw and tell you which racks have available capacity. Centralized software gives you real-time, data-driven insight into the entire system from anywhere on the network.
4. **Enclosures** Vendor-neutral NetShelter SX rack design handles high-density airflow and power needs.

Learn how to reduce cooling expenses with our FREE Cooling Efficiency kit

Visit www.apc.com/promo Key Code j299w
Call 888-289-APCC x9769 • Fax 401-788-2797

Available through Dell's Software and Peripherals and online at Dell.com

APC® by Schneider Electric

e-mail: esupport@apc.com • 132 Fairgrounds Road, West Kingston, RI 02892 USA • 998-2028 *Full details are available online.

SIMPLIFY MANAGEMENT WITH THE DELL UNIFIED SERVER CONFIGURATOR ENABLED BY THE LIFECYCLE CONTROLLER

By Shelli Allgood
Anand Narayanan
Hai Phung
Pritesh Prabhu

The new Dell™ Unified Server Configurator enabled by the Lifecycle Controller incorporates a comprehensive set of systems management features that operate in a pre-OS environment. Embedded directly in new 11th-generation Dell PowerEdge™ servers, this next-generation tool does not require media such as CDs or DVDs and offers a simplified, flexible way to perform tasks ranging from firmware updates to OS deployment to diagnostics.

Related Categories:

Dell OpenManage
Dell PowerEdge servers
Dell Unified Server Configurator
Diagnostics
Firmware updates
Integrated Dell Remote Access Controller (iDRAC)
OS deployment
Systems management

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

Systems management is a key part of the IT administrator's job, and encompasses tasks such as installing operating systems, updating firmware to be functional and adhere to enterprise policies, configuring hardware, and maintaining IT infrastructures. Over the years, Dell has provided many tools designed to assist administrators with these tasks, including tools that shipped with multiple CDs (such as Dell OpenManage™ Server Assistant and the Dell OpenManage Deployment Toolkit) and a single CD carrying all the software tools (such as the Dell Systems Build and Update Utility). Although these tools helped simplify and streamline management of Dell systems, they also required administrators to save the CDs so the tools could be reinstalled if needed, which added clutter and storage requirements as organizations accumulated multiple versions of these tools.

The new Dell Unified Server Configurator (USC) enabled by the Lifecycle Controller is designed to overcome these challenges. Based on the Integrated Dell Remote Access Controller (iDRAC) embedded in new 11th-generation Dell PowerEdge servers, it provides a single place to perform updates, hardware configuration, OS deployment, and system diagnostics—one that functions independently of both media and platform OS. Because the USC is available even when the OS is not, it helps add flexibility when provisioning servers and customizing them to meet specific requirements. In addition, because the tool is integrated with and embedded in the server, formatting the disk or reinstalling the OS does not remove the tool, helping save the time and costs associated with reinstalling system tools and thereby helping increase administrator productivity.

USC COMPONENTS AND ARCHITECTURE

New 11th-generation Dell PowerEdge servers introduce Embedded Management, which is designed to dramatically reduce the time required for multiple common management tasks. It comprises several interdependent components, including the new Lifecycle Controller and USC interface.

The Lifecycle Controller is a flash chip embedded in the system itself that provides persistent storage for systems management components—including, among others, the system BIOS, firmware, drivers, and Dell OpenManage tools. Administrators can access these components through the USC interface, which is also stored on the Lifecycle Controller.

The USC offers a single access point and interface for performing comprehensive platform updates, hardware configuration, OS deployments, and diagnostics (see Figure 1). Administrators can launch the application by booting the server and pressing F10.

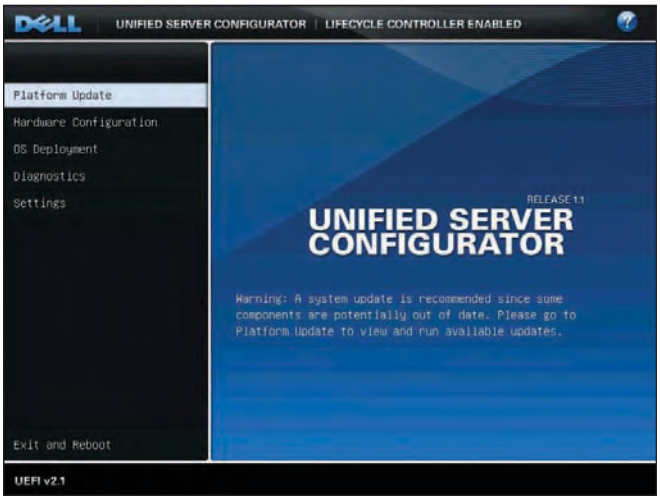

Figure 1. *Main menu functions in the Dell Unified Server Configurator*

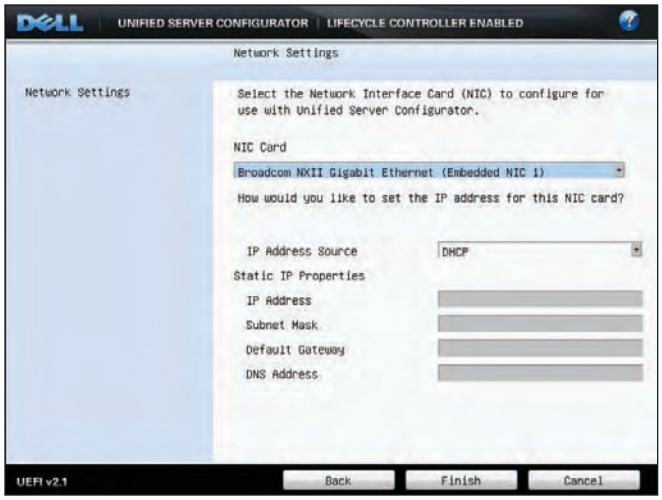

Figure 2. *Network settings in the Dell Unified Server Configurator*

Because the USC is managed by the iDRAC subsystem and runs in a Unified Extensible Firmware Interface (UEFI) environment, running the USC does not require booting the OS.

The UEFI specification was created to replace BIOS, but has also evolved to provide additional functionality that can support powerful applications that run without an OS. This specification includes a Trusted Computing Group protocol that supports Trusted Platform Module (TPM) measurements and secure updates and that the USC is designed to leverage. UEFI—and, therefore, the USC—also has a well-developed network stack and a strong point-and-click graphical environment designed for ease of use, helping make the USC a powerful, simplified pre-OS configuration tool.

The USC application includes built-in support for both Dynamic Host Configuration Protocol (DHCP) and Domain Name System (DNS), and enables administrators to configure network cards to acquire an IP address through DHCP or use static IP properties (see Figure 2). In addition, the application interface supports both text-based and SOCKS 4 proxies, enabling administrators to connect to the network through proxy servers to enable secure downloads to the server.

The USC also includes a task scheduler enabling administrators to deploy multiple updates in batches. This batch processor—called the System Services Manager (SSM)—is automatically invoked when multiple updates are requested. The intelligence in the SSM is designed to defer reboots to the last possible point to help reduce the number of reboots required to perform updates.

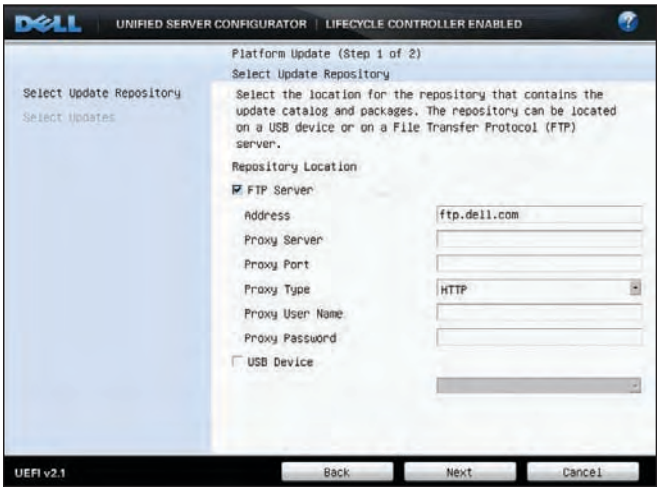

Figure 3. *Update repository settings in the Dell Unified Server Configurator*

The SSM also provides a way to apply OS-based Dell Update Packages (DUPs) in a synchronized way. When administrators apply a DUP through the OS, the updates are scheduled securely on the SSM for execution on next boot.

To provide assistance to administrators learning to use the USC, each USC screen includes context-sensitive help. Administrators can access the relevant help for a specific screen by clicking on the "?" button in the upper-right corner.

INTELLIGENT, EFFICIENT PLATFORM UPDATES

The USC platform update feature enables administrators to maintain the system firmware, diagnostics, BIOS, and drivers at the most current level, which can help maximize system performance and avoid unnecessary downtime. Operating in the UEFI environment, independently of an installed OS, helps increase the security and convenience of this process.

Administrators can perform platform updates from an FTP or local USB repository (see Figure 3); using a USB repository can help improve update efficiency during periods of network congestion. After administrators have selected a repository, the USC can automatically detect the necessary updates and provide the option of manually selecting specific components to update

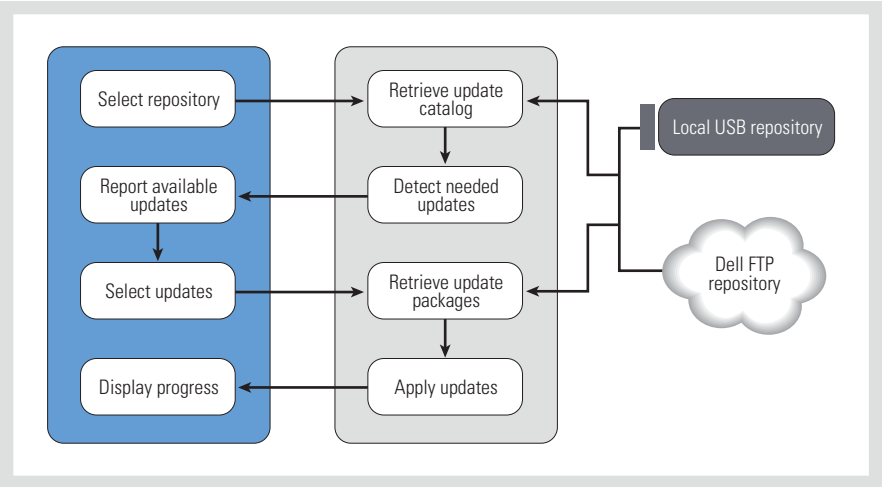

Figure 4. *Platform update process using the Dell Unified Server Configurator*

or deploying all identified updates. Figure 4 illustrates the update process.

Key platform update features

The USC introduces multiple platform update features not previously available in Dell systems management tools, and which are designed to simplify, automate, standardize, and secure the update process for IT administrators.

Simplified, media-less operation. The update process operates in the UEFI environment independently from an installed OS, and works on all supported servers. The update functions are designed to be immediately available after booting the server and pressing F10, and the update process requires no external media such as CDs or DVDs—helping eliminate the time required to find or burn a disc, or to deal with problems such as an unreadable disc.

Efficient, automated inventory and update process. The inventory process starts immediately after the server boots to the USC or after an update process has completed. It collects and records the current level of the system and subsystem firmware into the system life cycle log (LCL), which contains a comprehensive history of the server, including the initial factory loading configuration.

Utilizing the data from the LCL, the USC can automatically detect the necessary and available updates from the designated repository. It then displays the data in a table format, showing both the current level and the available level (see Figure 5). By default, the update process automatically upgrades all identified components, but administrators can also manually select specific components to upgrade.

Increased standardization, consistency, and synchronization. Updates are applied to the system from the SSM batch processor. After the updates are selected from the comparison report table, the updates are registered into the system services information block (SSIB) task list, and the SSM processes and executes the task from this list. Using the SSM and the SSIB helps increase standardization, consistency, and synchronization throughout the update process.

Intelligent reboots to help reduce downtime. The purpose of system updates is typically to keep the system as current as possible, helping avoid unnecessary downtime from system failures. However, the update process itself can cause significant downtime because of the need to reboot before an update can take effect. Applying multiple updates may require several reboots, which can cause considerable disruption.

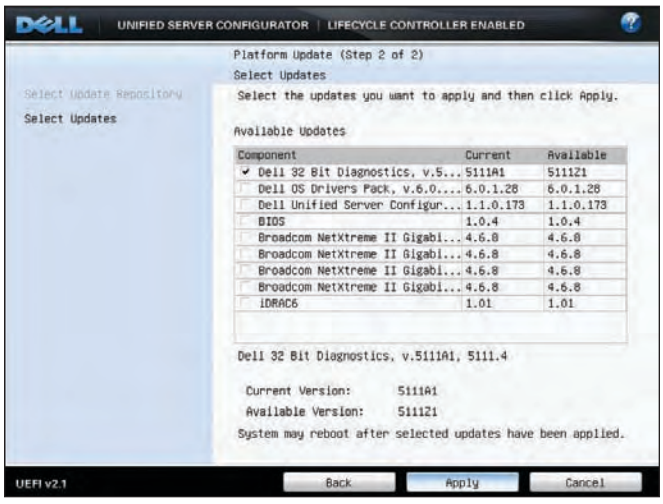

Figure 5. *Available updates detected by the Dell Unified Server Configurator*

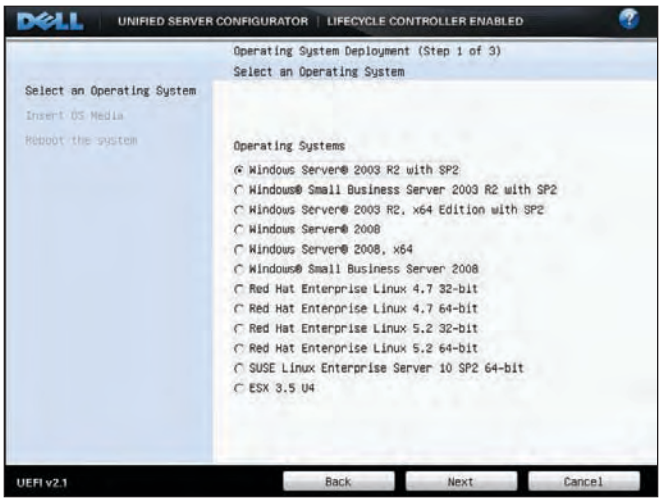

Figure 6. *Supported OS installations using the Dell Unified Server Configurator*

The USC update process is designed to minimize the number of reboots by using the update catalog combined with the parameters in the SSIB task list. Unless an update requires an immediate reboot, the process combines reboots, potentially enabling the server to reboot only once after the updates are completed. In addition, running the update process in the UEFI environment helps reduce boot time by avoiding the need to load a full OS.

Enhanced security. The update process is designed to ensure that the system and data are well protected, and utilizes multiple layers of security validation. For example, external data is stored temporarily in nonvolatile storage, with the data then validated using the data signature to help prevent the intrusion of contaminated data from external sources. Only validated data is staged in the USC and applied to the appropriate components. In addition, USC partitions open as read-only and only authorized programs are allowed to open them as read/write, which helps control access.

FLEXIBLE OS DEPLOYMENT

Previous Dell OS deployment tools such as Dell OpenManage Server Assistant or the Dell Systems Build and Update Utility were media based, requiring administrators to have the appropriate Dell CD or DVD on hand for the server being provisioned. If the CD or DVD was not available, or was not the correct version, administrators would need to download the appropriate image and burn it onto a disc. In addition, the OS drivers present on the CD or DVD could be outdated by the time the server was ready for deployment, requiring administrators to run post-OS DUPs to update their drivers.

The USC OS Deployment wizard provides a flexible, simplified way to install supported Microsoft® Windows® and Linux® operating systems, and does not require separate media-based deployment tools (see Figures 6 and 7; supported operating systems may vary depending on the specific server). If the server has a RAID controller, administrators can also optionally launch the USC RAID Configuration wizard during the deployment process to configure a virtual disk as the boot device.

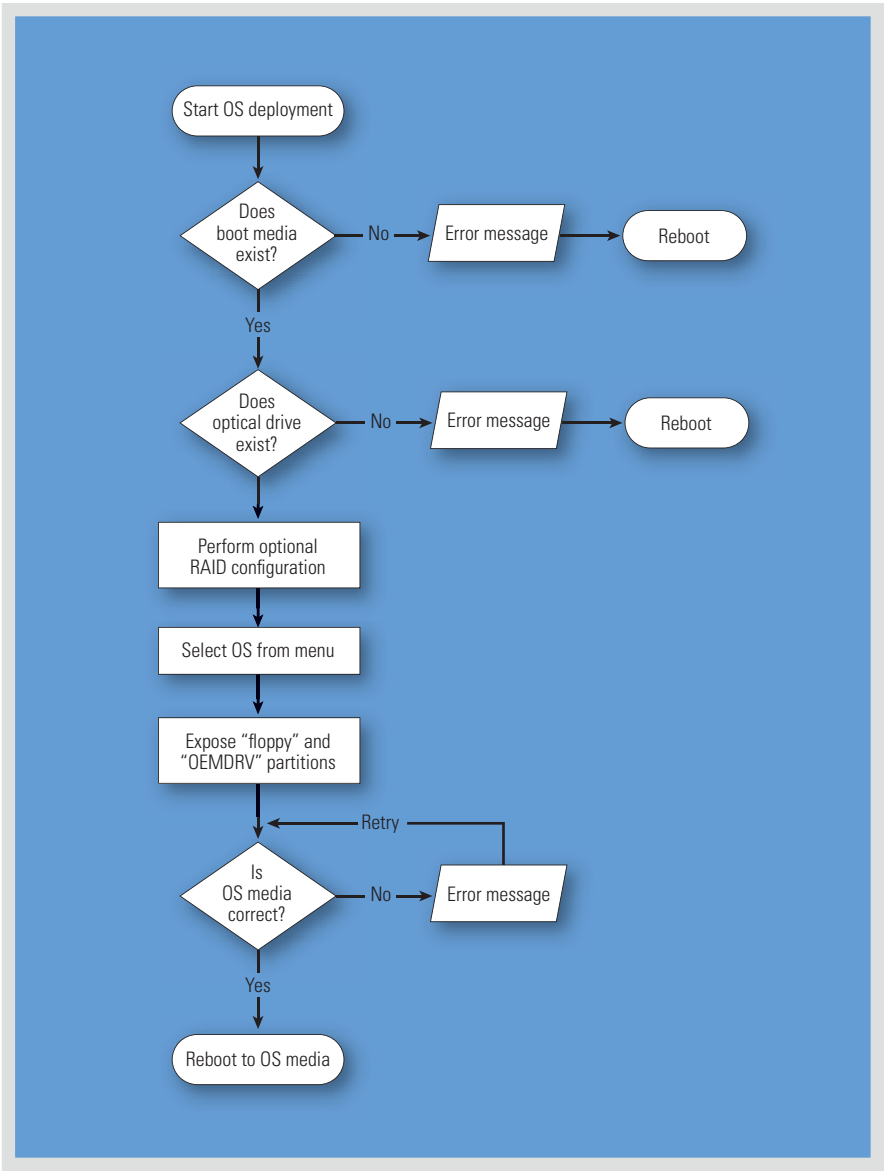

Figure 7. *OS deployment process using the Dell Unified Server Configurator*

After administrators have selected the OS to be deployed, the USC guides them through the process, including extracting the drivers necessary for OS installation and copying them to a staging directory. On Microsoft Windows operating systems, these extracted drivers are installed during the OS installation. On Linux operating systems, administrators must manually install the extracted drivers after the OS installation is complete. After this step is complete and administrators have inserted the OS installation media, the USC verifies the contents and proceeds with the installation.

Key OS deployment features

The USC OS deployment feature is designed to be more flexible and easier to use than previous Dell systems management tools, including support for native OS installation, updatable drivers, one-to-many driver installation, and virtual disk creation.

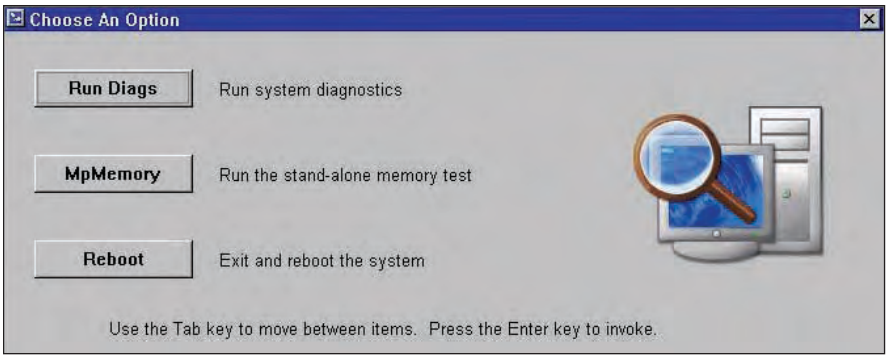

Figure 8. *Diagnostic functions in the Dell Unified Server Configurator*

Native OS installation. The USC uses the native installers provided by the OS installation media. This approach helps make the installation process flexible and extensible to future installer versions released for supported operating systems.

Updatable drivers. The drivers carried by the USC can be updated using the platform update feature. To help ensure that the USC has the most current available drivers, best practices recommend running the USC Platform Update wizard before OS deployment. Doing so helps avoid the need to manually check driver versions, and helps avoid inadvertently installing outdated drivers and then needing to run a post-OS DUP to update them.

One-to-many driver installation. The drivers carried by the USC are exposed on a USB drive, and can be shared across a network and used to provision multiple servers. This feature helps greatly reduce the overhead of setting up multiple servers and enables one-to-many OS deployment. Administrators can easily copy the drivers and plug them into the OS provisioning framework.

Simplified virtual disk creation. For servers that include a RAID controller, administrators can optionally launch the USC RAID Configuration wizard during the OS deployment process. This wizard first displays a list of its current controllers and the virtual disks in the system. If the necessary virtual disks are already in place, administrators can exit the wizard and continue with the OS deployment. Otherwise, the wizard guides them through the process of configuring a new virtual disk as a boot device, including selecting a RAID controller, clearing or ignoring an existing foreign configuration, selecting RAID type and physical disks, and optionally defining advanced parameters such as stripe element size, read and write policy, and hot spare disk assignment.

EMBEDDED DIAGNOSTICS

Diagnostics enable administrators to detect system hardware and memory problems to help reduce the expenses associated with the repair and building of systems (see Figure 8). Because the USC incorporates diagnostics directly in the embedded USC partition, these features ship with the system from the factory and are not removed when administrators install an OS or reformat a drive. In addition, embedding diagnostics in the USC partition enables administrators to easily update the diagnostics using the USC platform update feature, without requiring them to reinstall the utility partition at each update. The platform update feature can automatically perform an inventory and display a comparison report table if a new version is available.

INTEGRATED SYSTEMS MANAGEMENT

The new Dell Unified Server Configurator enabled by the Lifecycle Controller is designed to provide quick and easy access to system life cycle management capabilities in new 11th-generation Dell PowerEdge servers. Because it is embedded and integrated into the server itself, this tool does not require separate media such as Dell OpenManage CDs or DVDs and operates independently of the server OS—helping increase flexibility and save the time and costs associated with tasks such as reinstalling system tools after formatting a disk or reinstalling an OS. This powerful, simplified tool can help administrators rapidly perform systems management tasks and ultimately help increase their productivity.

Shelli Allgood is a lead software engineer in the Dell Systems Management Software Group. Shelli has a B.S. in Computer Science from Baylor University.

Anand Narayanan is a software lead for the Dell Unified Server Configurator in the Dell Product Group. Anand has a degree in Computer Science from Texas A&M University.

Hai Phung is a software engineer in the Dell Systems Management Software Group. Hai has a degree in Computer Science from the University of Minnesota.

Pritesh Prabhu is a software engineer consultant in the Dell Systems Management Software Group. Pritesh has an M.S. in Electrical and Computer Engineering from the University of Texas at Austin.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell systems management:
DELL.COM/OpenManage

COMPUTING PODS: LARGE-SCALE BUILDING BLOCKS FOR INTELLIGENT, AUTOMATED DATA CENTER DEPLOYMENTS

By Timothy Sherbak
Chris Auger

Using advanced virtualization and software tools together with latest-generation server, storage, and network infrastructure, IT organizations can build large-scale computing pods designed to aggregate server, storage, and network nodes into a simple, cooperating set of computing resources. This approach enhances enterprise efficiency by streamlining operations while facilitating fast, flexible response to increasingly complex service-level requirements.

Related Categories:

Blade servers
Dell EqualLogic storage
Dell PowerEdge blade servers
Storage
Virtualization

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

IT infrastructures tend to become ever more complex as administrators support growth with incremental hardware expansion, which in turn calls for increased real estate, power, cooling, and maintenance. Faced with shrinking budgets, many organizations end up dedicating a majority of their IT spending to simply maintaining data center resources—placing serious constraints on their ability to advance new initiatives.

To reverse the tide of complexity caused by piecemeal resource expansion, IT executives are looking for ways to make strategic capital investments geared to reduce overall operating expenses. In particular, transitioning to a virtualized environment can advance the efficiency of the overall IT infrastructure by enabling self-managing data center resources to scale quickly and flexibly to meet specific organizational needs.

CREATING BUILDING BLOCKS FOR AN INTELLIGENT, AUTOMATED DATA CENTER

Many of the elements required to build a highly scalable, automated data center are now available. Strategic integration and coordination of key hardware elements such as standards-based servers, storage, switches, and processors are making these data center advances possible. The required software elements are also available today, including mature operating systems, hypervisors, and management solutions.

Another key requirement of an automated data center is policy-based automation utilizing real-time intelligence. Feedback loops are needed to help recognize problems and optimize the core features of each layer of the infrastructure. Multipathing built into some of today's server and storage infrastructures is one example: if server and storage communication fails on one network path, the infrastructure has the intelligence to try different paths. Other examples include automatic load balancing of workloads across multiple physical servers and storage devices, and migration of workloads to help save energy by minimizing the number of actively running servers.

Many platforms are now optimized for virtualization. To advance data center flexibility, intelligence, and automation, IT managers must be able to integrate virtualization across multiple layers of the infrastructure to create virtualized server, storage, and network resource pools. A key requirement is that specific workloads not be bound to a specific physical resource.

Keeping workloads independent of specific hardware devices allows the flexibility and adaptability needed to migrate workloads, reallocate

resources, and expand resources on the fly. These capabilities help simplify proactive management scenarios to support real-world business processes. For example, IT managers may dynamically reallocate data center resources when end-of-month and financial reports are due, or quickly make resources available for a new project that has to be up and running in 30 days.

Such flexibility and adaptability allows organizations to build as they grow rather than over-provision infrastructure up front based on anticipated needs. An intelligent, self-managing data center that integrates virtualization across multiple layers of the infrastructure is also designed to avoid prolonged service outages because of infrastructure failures, or disruptive system upgrades and manual performance tuning in response to unforeseen performance and growth needs.

THINKING ABOUT IT INFRASTRUCTURE IN NEW WAYS

With many elements of the intelligent, automated data center now available, IT executives are starting to think about the IT infrastructure in innovative ways. They are using existing data center components to build architectures around the concept of a *computing pod*. The pod can be defined as a self-contained building block having the following attributes:

- An optimized power and cooling footprint to facilitate scalability
- Integrated compute, storage, and network nodes of significant size and processing power
- Simplified data center deployment and management to help reduce operational costs

IT managers traditionally think of servers, storage area network (SAN) arrays, and switches as core building blocks. Now they can consolidate these elements into a single large-scale building block that becomes the basic unit of data center deployment (see Figure 1).

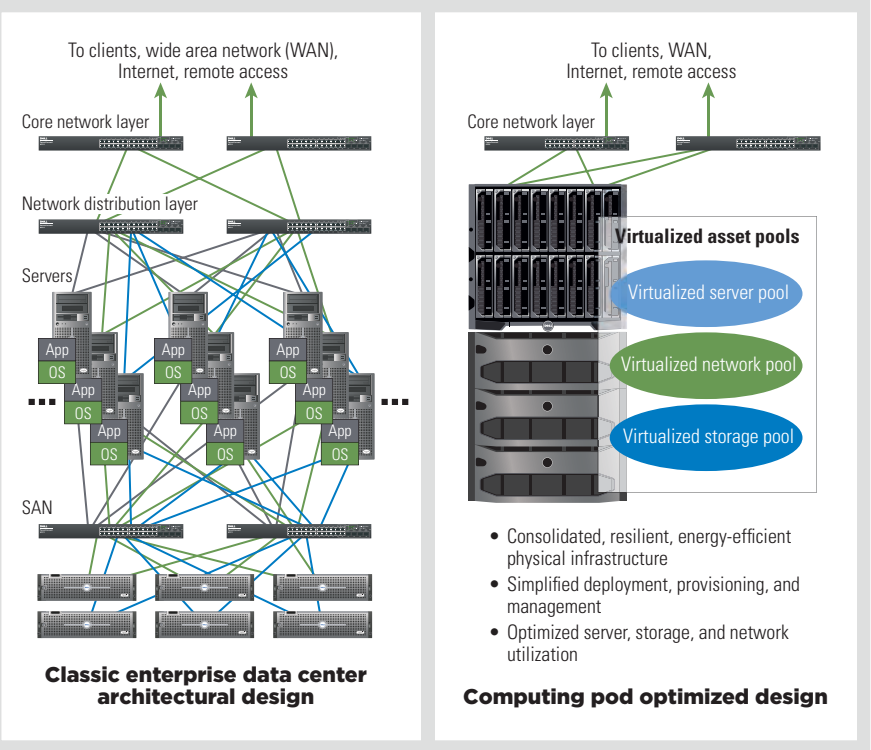

Figure 1. *Comparing a traditional data center architecture with a computing pod that consolidates server, storage, and network resources*

DESIGNING A COMPUTING POD ARCHITECTURE

To design a computing pod architecture, IT organizations can take advantage of advanced virtualization and software tools together with the latest generation of server, storage, and network infrastructure. Dell recommends that each pod include the Dell™ PowerEdge™ M1000e modular blade enclosure with PowerEdge blade servers, SAN arrays, hypervisor, management module, and switches with stacking capability.

Dell PowerEdge enclosures and blade servers

The Dell PowerEdge M1000e enclosure enables administrators to slide in blade servers and switches to expand or upgrade infrastructure as needed. Expansion is simple and internal to the chassis, allowing nondisruptive growth. An important element of the PowerEdge M1000e enclosure is the Dell FlexAddress feature, which effectively virtualizes the physical addresses of ports within the infrastructure. Newly swapped-in components can automatically inherit the attributes of the prior components. This feature enables administrators to modify or change physical network infrastructure without an impact on the production environment.

Organizations using PowerEdge M1000e enclosures can also optimize pod design using the Dell M1000e Rack and Cable Advisor and Dell Virtualization Advisor online tools, available at DELL.COM/RackAdvisor and DELL.COM/Virtualization. The M1000e Rack Advisor guides administrators through the configuration process, including the use of power distribution units and enclosure stacking, while the Virtualization Advisor provides information to help size blade and storage infrastructure according to the aggregate workloads to be hosted. PowerEdge M610 and PowerEdge M710 blade servers are designed to reduce space and power requirements, which in turn can contribute

to reduced total cost of ownership and environmental impact.[1]

Virtualized server and storage platforms

Virtualized server platforms are essential to the computing pod model, enabling administrators to create a flexible pool of server resources that can be allocated dynamically. If one physical server fails, virtual machines can automatically fail over to a different physical system to help ensure continuous availability. IT administrators can also quickly and easily start up virtual machines at a disaster recovery site. In addition, virtualization allows the consolidation of multiple servers onto a reduced number of physical platforms, enabling the computing pod to provide significant processing power without sprawling hardware.

Virtualizing storage extends the benefits of high availability, scalability, and optimization of resources. Internet SCSI (iSCSI)–based virtualized storage arrays such as Dell EqualLogic™ PS6000 series virtualized iSCSI SANs are recommended because they use the same Ethernet technology as other pod elements.[2] The EqualLogic SAN also contains built-in intelligence—such as automatic load balancing, simplified storage provisioning, thin provisioning, replication, and multipathing—that is designed to integrate with pod environments utilizing hypervisor technologies.

IT organizations have a variety of choices for mature virtualization software, including the VMware® vSphere™, Microsoft® Hyper-V™, and Citrix® XenServer™ platforms.

Network stacking and virtual blade switching

Enhanced scalability and simplified management are key objectives of computing pod design. Server stacking and virtual blade switching help deliver these capabilities. Stacking allows IT managers to incrementally grow network infrastructure by adding bandwidth and I/O ports in a seamless way. Cisco switches that include virtual blade switch (VBS) technology, such as the Cisco Catalyst Blade Switch 3100 series switches, enable IT administrators to interconnect multiple switches so that they interact and behave as a single logical switch, managed together as a virtualized pool of network resources.

Stacking and virtual blade switching also enable the computing pod model by creating a self-contained network distribution layer.[3] With its single virtual switch, the blade enclosure becomes the basic computing pod infrastructure unit. The network distribution layer resides entirely within the enclosure, helping avoid the need for traffic among the server and

DELL BUSINESS-READY CONFIGURATIONS: BUILDING BLOCKS FOR GROWTH

As data centers continue to grow in complexity, many traditional approaches to incremental expansion may not be able to keep up with rapid changes in business conditions. IT organizations can enhance business response by taking advantage of advanced software tools and infrastructure such as the following:

- Dell PowerEdge M1000e modular blade enclosure
- Dell PowerEdge M610 and PowerEdge M710 blade servers with the Intel® Xeon® processor 5500 series
- Dell EqualLogic PS6000 series virtualized Internet SCSI (iSCSI) storage area network (SAN)
- Cisco Catalyst Blade Switch 3100 series
- VMware vSphere, Microsoft Hyper-V, and Citrix XenServer virtualization platforms
- Dell Management Console

By capitalizing on large-scale building blocks that help simplify and automate IT operations, the computing pod model helps enterprises create innovative, self-contained data centers that can scale quickly, flexibly, and cost-effectively.

Dell Business-Ready Configurations provide a detailed reference architecture for deploying and implementing a computing pod architecture using Dell blade server and iSCSI storage environments. Based on extensive engineering work in architectural design and certification, the architecture is designed for rapid deployment into production environments, helping to eliminate much of the costly and time-consuming trial and error often encountered during complex infrastructure design and implementation. The solution is optimally configured to run virtualized workloads and is designed to provide redundancy without single points of failure, as well as scalability and manageability. It includes the network architectures, storage configurations, and best practices necessary for deploying and configuring the solution.*

*For more information, see "Business-Ready Configurations for Dell PowerEdge Blade Servers, Dell EqualLogic Storage, and VMware Infrastructure," by Dell Virtualization Solutions Engineering, May 2009; and "Networking Best Practices for VMware Infrastructure 3 on Dell PowerEdge Blade Servers," by Dell Virtualization Solutions Engineering, April 2009, both available at DELL.COM/Virtualization.

[1] For more information on PowerEdge M610 and PowerEdge M710 blade servers, see "Data Center Workhorses: New Dell PowerEdge Rack and Blade Servers," by Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090246-Nguyen.pdf.

[2] For more information on Dell EqualLogic PS6000 series arrays, see "Boosting SAN Performance with Dell EqualLogic PS6000S Solid-State Drive Arrays," by Dylan Locsin, Ujjwal Rajbhandari, and Wendy Chen, in *Dell Power Solutions*, June 2009, DELL.COM/ Downloads/Global/Power/ps2q09-20090236-Locsin.pdf.

[3] Large enterprise data centers, which can potentially house thousands of physical servers, may use a three-tier network architecture. The servers may be connected together through an "access layer" network of switches, which connect to the distribution layer network (which, in turn, connects to the core network layer). In this case, the same conceptual argument for a pod architecture applies. In the pod architecture, server communication and data transfers between servers across the access layer network are consolidated on the pod's backplane-based network infrastructure, isolating it from other parts of the network.

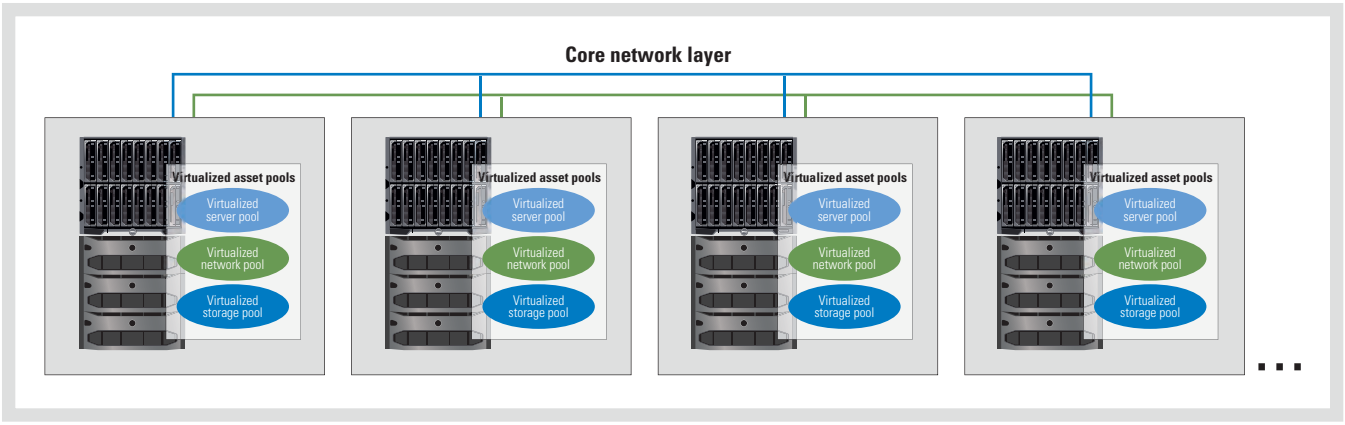

Figure 2. *Scaling enterprise data center architectures with multiple computing pods*

storage nodes to go out onto the data center core network layer. Servers, switches, and the management module are contained within the blade server enclosure environment (see the "Dell Business-Ready Configurations: Building blocks for growth" sidebar in this article).

ENHANCING DATA CENTER CAPABILITIES WITH COMPUTING PODS

Computing pods are designed to offer IT organizations significantly enhanced data center capabilities, including efficient aggregation and networking, increased space savings, unified maintenance and training, and simplified infrastructure design and deployment.

Redundant, high-bandwidth connectivity internal to blade enclosures within the pod helps speed communication and reduce congestion on the core data center network layer. For example, suppose an administrator needs to place a blade server into standby mode and move the virtualized servers to a different blade server in the same stack. None of that traffic actually leaves the backplane or affects the core data center network layer. Even large workloads can move among blade servers without disrupting productivity.

Likewise, access to storage is isolated from the data center's core network infrastructure. The iSCSI SAN networking infrastructure can be integrated directly into the blade enclosure's Ethernet switch infrastructure, either as a virtual LAN or as an independent network. With Dell EqualLogic PS Series arrays, storage nodes are "stacked" as well—that is, multiple cooperating arrays are integrated as a logical pool of storage to help optimize I/O and storage utilization. In this way, administrators can use each pod to aggregate dozens or even hundreds of server and storage workloads in a minimum amount of space.

Data protection procedures may also be consolidated and centralized within the computing pod. Administrators begin by designating one or two blade servers per pod as media servers for backup and recovery with direct connectivity to tape and/or disk resources dedicated to backups. Data backup traffic is limited to the pod's computing, storage, and network infrastructure without affecting external data center resources.

Another blade server within each pod can be designated as a configuration server for pod management. Using the Dell Management Console,[4] hypervisor management tools, and other management applications, this designation enables administrators to manage the production servers in the pod as a unit, from configuration and reimaging to patch management and antivirus updates. As a result, IT staff members need not spend time managing individual servers, which helps reduce IT support costs substantially.

Organizations can maintain a single logical entity—the pod—instead of performing time-consuming, one-to-one server management, storage management, and network management. Managing the complete computing pod environment as a unit not only helps optimize servers and switches, but also helps simplify backup and patch management. IT staff can be efficiently trained to manage the standard pod and its components, using logically consolidated tools and procedures rather than multiple disparate systems.

Computing pods also help simplify infrastructure design and deployment. Designing a data center using the computing pod as a basic building block helps simplify the design target. Organizations using pods are able to reduce the time and cost requirements of data center planning, deployment, and implementation.

DEPLOYING COMPUTING PODS FOR FLEXIBLE, SCALABLE DATA CENTERS

The building block approach to computing pods enables outstanding flexibility and scalability. Blade server, storage, and network nodes within the pod are designed to be easily scalable, and the

[4] For more information on the Dell Management Console, see "Systems Management Simplified," by Pascal Nicolas, Jeanne Feldkamp, and Tom Kolnowski, in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090191-CoverStory.pdf.

pod itself becomes a building block for scaling out the data center.

Organizations of different sizes have distinct opportunities to benefit from the computing pod model. For example, midsize organizations can start with a single pod as a convenient, all-inclusive computing infrastructure in a single rack enclosure and add nodes to the pod or additional pods as requirements grow. In large enterprise environments, IT managers may deploy numerous pods to help meet comprehensive data center infrastructure requirements, assigning different pods for different usage scenarios. For example, pods may be designated for production, development, and Web presence—all sharing a common environment (see Figure 2).

The computing pod model also facilitates virtual desktop infrastructures by allowing IT administrators to easily add a pod to virtualize and aggregate desktops. End users can remotely access the computing pod server infrastructure and run their personal computer environments there.

In addition, a growing number of organizations are utilizing computing clouds to deliver flexible access to compute and storage resources. The cloud computing approach enables organizations to leverage the scalability and manageability of pods to build large data center infrastructures with a pool of flexible resources. The pod design helps minimize acquisition and operating costs, maximize energy efficiency, and enable rapid scalability.

INTEGRATING INFRASTRUCTURE ELEMENTS INTO EFFICIENT, AUTOMATED DATA CENTERS

The computing pod is an exciting model for IT infrastructure that helps significantly advance enterprise efficiency. By elevating data center intelligence through advanced virtualization and software tools in combination with latest-generation server, storage, and network infrastructure, organizations can create highly scalable, flexible, and self-managing data centers. By consolidating data center functionality into a basic unit that also serves as a building block for large-scale data center deployments, computing pods help reduce IT infrastructure complexity, simplify management, and increase cost-efficiency.

Moreover, computing pods are designed to provide a comprehensive, integrated management infrastructure that can be customized to specific organizational requirements. They help reduce and consolidate network traffic by isolating it to the specific resources requiring that traffic. And the large-scale building blocks provide for a flexible computing infrastructure that can quickly and easily adapt to evolving needs and requirements.

Timothy Sherbak is senior manager of virtualization solutions marketing in the Dell Large Enterprise Business Unit, and has over 20 years of experience in computer systems and storage networking technologies. He has a B.S. in Systems and Information Science from Syracuse University and an M.S. from the Massachusetts Institute of Technology's Sloan School of Management.

Chris Auger is manager of enterprise technology in the Dell Public Business Unit, leading a team of enterprise technologists helping government, education, and health care organizations architect world-class data center solutions. He has a B.S. in Mathematics from the University of Texas at Austin.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell Efficient Enterprise Portfolio:
DELL.COM/Switch

Dell M1000e Rack and Cable Advisor:
DELL.COM/RackAdvisor

Dell Virtualization Advisor:
DELL.COM/Virtualization

Dell storage:
DELL.COM/Storage

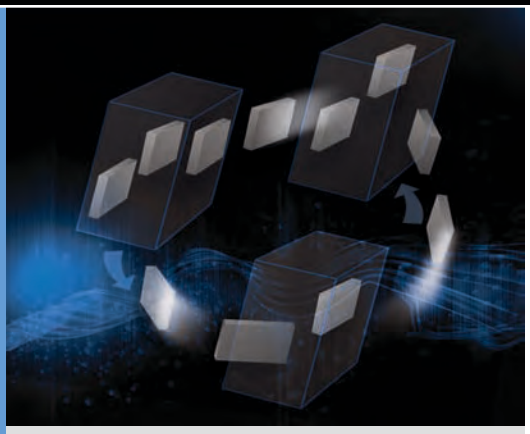

SIMPLIFY MIGRATION OF LEGACY SERVERS TO VIRTUAL MACHINES WITH PLATESPIN MIGRATE

by David Schmidt
Joseph Rispoli

Virtualizing existing workloads offers an excellent way to consolidate underutilized legacy hardware, helping to increase flexibility and reduce costs without compromising performance. PlateSpin® Migrate for Dell can help simplify and accelerate physical-to-virtual migrations onto powerful 11th-generation Dell™ PowerEdge™ servers.

Related Categories:

Blade servers
Dell EqualLogic storage
Dell PowerEdge servers
Microsoft SQL Server
Microsoft Windows Server 2003
PlateSpin
Red Hat Enterprise Linux
Servers
Virtualization
VMware

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

Virtualization has become a mainstream component of enterprise data centers, helping provide consolidation savings, workload migration capabilities, reduced deployment times, and a variety of other benefits. A majority of these benefits are forward-looking: they tend to focus on newly deployed environments, not on legacy server installations. However, because data center ecosystems are increasingly optimized for virtualization, IT teams often seek to virtualize existing installations as well as new deployments. Consequently, migration tools have emerged as a key method of moving legacy OS images from physical servers to virtual machines (VMs).

To help ease the migration of legacy hardware to powerful virtualized environments, Dell now offers PlateSpin Migrate workload portability software from Novell as part of its factory software options. A five-pack license for the physical-to-virtual (P2V) migration capability can be purchased during the configuration of most Dell PowerEdge servers and delivered as part of the server documentation. Administrators can then download the software directly from the PlateSpin Web site. Dell offers support for this software as well as Dell ProSupport services at additional cost.

To successfully perform P2V migrations, however, administrators must next turn to planning and execution, including identifying appropriate candidate servers with consideration for the OS version, specific workload type, and acceptable downtime. This article focuses on migrating an example environment of common legacy servers to virtualized 11th-generation Dell PowerEdge blade servers, outlines factors that administrators should take into account when planning a migration, and recommends best practices for preparing and carrying out the migration.

UNDERSTANDING THE PLATESPIN MIGRATION PROCESS

PlateSpin Migrate is part of the PlateSpin Portability Suite, which uses a client/server model. The client software can be installed on the server host or on a separate Microsoft® Windows® OS–based workstation. The server component requires a Microsoft Windows 2000 Server or Microsoft Windows Server® 2003 OS–based host.

P2V migrations with PlateSpin Migrate use one of two methods: Take Control or Live Transfer. The Take Control method takes the source workload offline during the transfer, while the Live Transfer method is designed to minimize workload downtime. Before performing a Take Control migration, the PlateSpin Portability Suite server temporarily modifies the boot files on the source server so that it can restart into a pre-execution environment. Servers running a Windows OS reboot into a Windows Preinstallation Environment (WinPE), while those running a Linux® OS reboot into a RAM disk. After

rebooting into this environment, the source server's original boot files are restored and network communication with the PlateSpin Portability Suite server is initiated. When the PlateSpin Portability Suite server receives communication from the target hypervisor, the source server's files begin copying to the target VM. The PlateSpin Portability Suite server communicates with the target VMware® ESX host and automatically creates the corresponding source workload. PlateSpin Migrate allows administrators to reconfigure target VM components such as processor, memory, storage, and so on. PlateSpin Migrate also contains a database to help ensure that the appropriate drivers for the target hypervisor are injected into the OS image during the migration. After the migration is complete, the source server may be configured to shut down, while the target VM remains online. Linux-based servers must be migrated using the Take Control method.

PlateSpin Migrate can perform file-based Live Transfer migrations on most Windows workloads. With this method, the PlateSpin Portability Suite server installs a controller on the source server. After network communication is established between the source server, target VM, and PlateSpin Portability Suite server, the controller begins transferring the source server's files to the target VM. When the first transfer pass is complete, the controller checks the source server for modified or recently added files and resends the remaining data. When migrating Microsoft Exchange or Microsoft SQL Server® workloads, the appropriate services are stopped for the final file transfer. After the transfer is complete, the source server shuts down and the target VM remains online. Administrators can fine-tune network and file transfer settings if they have a high-latency network between the source and target.[1]

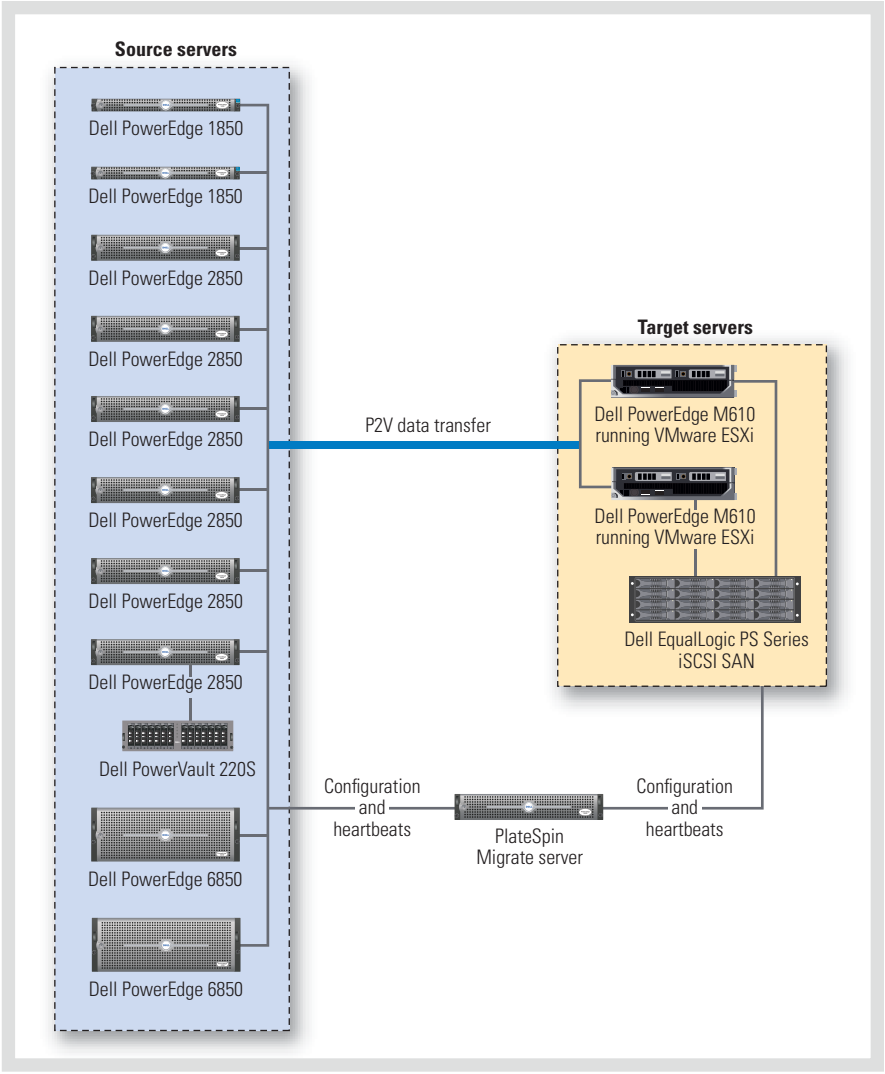

Figure 1. *Example migration environment using PlateSpin Migrate for Dell*

CREATING AN EXAMPLE MIGRATION ENVIRONMENT

A typical characteristic of physical server sprawl is the presence of multiple underutilized servers, each running a single workload. As a case study, in May 2009 Dell engineers and PlateSpin technicians performed an example migration on a set of 10 such legacy servers: 2 Dell PowerEdge 1850 servers, 6 PowerEdge 2850 servers, and 2 PowerEdge 6850 servers. A Dell PowerVault™ 220S disk array provided additional SCSI storage for one of the PowerEdge 2850 servers, and all servers connected to a low-latency Gigabit Ethernet LAN.

The goal of the case study was to use PlateSpin Migrate for Dell to migrate and consolidate the legacy hardware workloads to a target environment consisting of a PowerEdge M1000e modular blade enclosure with two 11th-generation PowerEdge M610 blade servers running an embedded VMware ESXi 3.5 hypervisor. A Dell EqualLogic™ PS Series Internet SCSI (iSCSI) storage area network (SAN) array—in this case an EqualLogic PS5000XV array—provided shared storage.

Figure 1 illustrates the example migration environment; Figure 2 shows the OS, function, transfer method, amount of data transferred, and migration time for each source server. In this environment, the PlateSpin Migrate server established the initial configuration and monitored

[1] For more information, see the "Fine-Tuning Data Transfer Performance" section of the PlateSpin Portability Suite user guide, available at www.platespin.com/downloads/dellp2v.

	OS	Function	Transfer method	Data transferred	Migration time (hours:minutes)*
PowerEdge 1850	Microsoft Windows 2000 Server	Domain controller and DNS server	Take Control	2.89 GB	0:28
PowerEdge 1850	Microsoft Windows 2000 Server	Print server	Live Transfer	2.82 GB	0:27
PowerEdge 2850	Red Hat Enterprise Linux 3	Apache Web server	Take Control	10.02 GB	0:33
PowerEdge 2850	Red Hat Enterprise Linux 3	Apache Web server	Take Control	10.02 GB	0:35
PowerEdge 2850	Microsoft Windows 2000 Server	Common Internet File System (CIFS) file server	Live Transfer	2.77 GB	0:20
PowerEdge 2850	Microsoft Windows 2000 Server	Proxy server	Live Transfer	2.77 GB	0:22
PowerEdge 2850	Microsoft Windows Server 2003 Web Edition	FTP server on Windows Internet Information Services (IIS)	Live Transfer	3.32 GB	0:19
PowerEdge 2850	Microsoft Windows Server 2003 Standard Edition	Microsoft Exchange server with 180 GB database	Live Transfer	185.16 GB	6:23
PowerEdge 6850	Microsoft Windows Server 2003 Enterprise Edition	Microsoft SQL Server database server with 10 GB database	Live Transfer	14.16 GB	0:46
PowerEdge 6580	Microsoft Windows Server 2003 Enterprise Edition	Microsoft SQL Server database server with 50 GB database	Live Transfer	55.69 GB	1:44

*Migration times are rounded to the nearest minute.

Figure 2. *Source servers, data transfer, and migration times during the example migration*

heartbeats from the source and target servers, with the payload data transfer communicated directly from source to target. PlateSpin Migrate limits administrators to 25 simultaneous conversions and 5 conversions per target hypervisor host. For this case study, to help maximize network bandwidth for each migration, the migration team performed only one migration per destination hypervisor at a time.[2]

FOLLOWING RECOMMENDED BEST PRACTICES

Following best practices before, during, and after the conversion process can help ensure a successful migration.

Verify compatibility. Before performing a P2V migration, administrators should use the PlateSpin Analyzer program (part of the PlateSpin Portability Suite) to validate each source workload. This program checks devices on Windows-based source servers for compatibility and flags applications and services that may require special attention during the migration process. The Workload Service Control allows administrators to specify services that should be shut down during the migration, which may include Microsoft Exchange and Dell OpenManage™ software.

Verify licensing and credentials. Administrators should ensure that the PlateSpin Portability Suite has a sufficient number of licenses and the correct type of licenses to perform all the P2V migrations. The PlateSpin Portability Suite must have appropriate authentication credentials for each source server, and all servers must have certain network services running, such as Windows Management Instrumentation (WMI) for Windows-based servers and sshd for Linux servers. Both the source and target servers require administrator or root credentials.

Maintain the source server for critical workloads. When administrators use the Live Transfer approach to help minimize downtime, but must test the target VM before it begins running the production workload, they can use the PlateSpin Server Sync™ feature to keep the source workload active after the migration. After they have completed testing on the VM, they can then initiate one final file transfer pass, after which the VM begins running the production workload.

Understand the needs of specific workloads. Domain controllers can be transferred using the Take Control or Live Transfer method; using the Live Transfer method, however, requires following specific PlateSpin best practices.[3] In the example migration environment, the migration team used the Take Control method to migrate the PowerEdge 1850 domain controller and Domain Name System (DNS) server.

[2] Both PowerEdge 1850 servers were configured with two Intel® Xeon® "Nocona" processors at 3.6 GHz, 3 GB of RAM, one 10,000 rpm, 147 GB hard drive, and two Intel PRO/1000 MT LAN on Motherboards (LOMs); the four PowerEdge 2850 servers running Red Hat Enterprise Linux or Windows 2000 Server were configured with two Intel Xeon "Nocona" processors at 3.4 GHz, 3.5 GB of RAM, one 10,000 rpm, 300 GB hard drive, and two Intel PRO/1000 MT LOMs; the PowerEdge 2850 server running Windows Server 2003 Web Edition was configured with two Intel Xeon "Nocona" processors at 3.6 GHz, 4 GB of RAM, one 10,000 rpm, 300 GB hard drive, and two Intel PRO/1000 MT LOMs; the PowerEdge 2850 server running Windows Server 2003 Standard Edition was configured with two Intel Xeon "Nocona" processors at 3.6 GHz, 4 GB of RAM, one 10,000 rpm, 146 GB hard drive, one 15,000 rpm, 73 GB hard drive, and two Intel PRO/1000 MT LOMs; both PowerEdge 6850 servers were configured with four dual-core Intel Xeon 7140M processors at 3.4 GHz, 16 GB of RAM, one 10,000 rpm, 34 GB hard drive, one 15,000 rpm, 34 GB hard drive, and two Broadcom BCM5704 LOMs; and the PowerVault 220S array was configured with twelve 10,000 rpm, 73 GB hard drives. The two PowerEdge M610 blade servers were configured with two quad-core Intel Xeon E5502 processors at 1.86 GHz, 12 GB of RAM, no hard drives, and two Broadcom BCM5709 LOMs, and the EqualLogic PS5000XV array was configured with sixteen 15,000 rpm, 450 GB hard drives.

[3] For more information, visit support.platespin.com/kb2/article.aspx?id=20501 (login required).

The PowerEdge 2850 Web servers in the example migration environment are running Apache on the Red Hat® Enterprise Linux 3 OS. Because this is a Linux OS, PlateSpin Migrate supports only the Take Control method. Before beginning the migration, administrators should ensure that the file systems to be copied are ext2, ext3, and/or ReiserFS; other file systems are currently unsupported.

For the PowerEdge 2850 Exchange server, administrators would typically use the file-based Live Transfer feature, which supports Windows Server 2003 Standard Edition. They would need to disable Exchange services during the migration, and should ensure that the original settings are restored to the target VM after the conversion is complete by selecting the appropriate check box.

Similarly, administrators would also typically use the file-based Live Transfer feature for the PowerEdge 6850 SQL Server workloads, because Live Transfer supports Windows Server 2003 Enterprise Edition. While the database files are copied from the source server to the target VM, the database service must be stopped; in the Transfer Services window, they should ensure that the SQL Server service is detected and scheduled to be stopped, and that the appropriate check box indicates that database files will be copied. If the database files are extremely large, resulting in an unacceptable level of downtime for the database service, administrators should consider purchasing a license for PlateSpin Protect so that they can perform a block-based Live Transfer migration. Purchasing a license for PlateSpin Protect Enterprise would also enable them to use Microsoft Volume Shadow Copy Service (VSS) to allow Exchange and SQL Server services to continue running during the migration; after the files have been copied, file blocks that have changed since the migration began are then copied at the end without requiring downtime. In this case study, the database files were copied with the OS files by PlateSpin Migrate using Live Transfer.

Choose an appropriate mode. In general, the advanced mode is preferable to the wizard mode because it offers additional options, such as installing VMware Tools in the VM and running custom scripts after the migration. Using advanced mode can help administrators avoid the need to manually perform redundant tasks and accelerate migrations, and allows them to save or schedule a job for execution at a later time.

Disable any unnecessary services. Administrators should typically disable any service that is not to be used in the target environment, such as legacy versions of Dell OpenManage services.[4]

Ensure that the target environment is appropriately sized. During the planning stages of a migration, administrators should ensure that the target virtualized environment has sufficient capacity to meet the workload requirements. For small numbers of servers such as those in the example case study environment, all VMs may be able to run in a one- or two-node virtualized pool. For larger environments, administrators should typically use additional tools such as PlateSpin Recon to help them appropriately size the final configuration.

PlateSpin Recon provides planning and analysis capabilities for large migration environments, collecting hardware, software, and services inventory in addition to workload and utilization information. By assessing the current requirements of the physical environment, it can make detailed sizing recommendations for the target virtualized environment, including specific information for the server models used in the configuration. Dell works closely with PlateSpin to help ensure that the software includes the latest information on Dell servers. In addition, the Dell Global Infrastructure Consulting Services team offers Virtualization Readiness Assessments using PlateSpin Recon to provide a comprehensive, customized analysis of an existing data center configuration.

SIMPLIFYING THE MOVE TO VIRTUALIZATION

PlateSpin Migrate for Dell provides an efficient process for migrating legacy hardware to powerful virtualized environments, one that helps avoid the need for administrators to plan for days of server migrations and incur costly downtimes. In the example case study, most migrations completed in less than an hour, and even the most complex migration completed within a single workday. By offering PlateSpin Migrate software with 11th-generation Dell PowerEdge servers as well as support and services to help with planning and implementation, Dell provides a comprehensive, simplified way to carry out these migrations in enterprise data centers.

David Schmidt is an engineering manager for the Dell Virtualization Solutions Engineering Group.

Joseph Rispoli is an engineer in the Dell Virtualization Solutions Engineering Group.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

PlateSpin:
www.platespin.com

Dell and PlateSpin:
DELL.COM/PlateSpin

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell virtualization solutions:
DELL.COM/Virtualization

More information:
migratetips@platespin.com

[4] For more information, see the "Best Practices (X2P)" section of the PlateSpin Portability Suite user guide, available at www.platespin.com/downloads/dellp2v.

BUILDING THE DELL SUPERCOMPUTER: NEW DELL SERVERS, INTEL CLUSTER READY, AND CLUSTERCORP ROCKS+

By Garima Kochhar
Jacob Liberman
Brock A. Taylor
Jason Bishop

To help simplify the deployment and management of high-performance computing (HPC) clusters, Dell, Intel, and Clustercorp have collaborated to offer scalable, cost-effective HPC solutions based on 11th-generation Dell™ PowerEdge™ servers, Intel® Cluster Ready configurations, and Clustercorp® Rocks+™ cluster management software.

Related Categories:

Clustercorp
Clustering
Dell PowerEdge servers
High-performance computing (HPC)
Intel
Intel Cluster Ready

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

Designing and deploying a high-performance computing (HPC) cluster can be difficult for even the most experienced IT staff. To help simplify cluster acquisition and management, Dell, Intel, and Clustercorp have collaborated to offer an integrated and tested HPC cluster platform based on 11th-generation Dell PowerEdge servers, Intel Cluster Ready configurations, and Clustercorp Rocks+ cluster management software—the commercial version of the award-winning Rocks® software.[1] This collaboration enables Dell HPC clusters to take on the feel of a traditional supercomputer, and enables them to be delivered to the data center ready to run registered applications right out of the box.

DEFINING THE DELL HPC CLUSTER

The clusters resulting from the collaboration between Dell, Intel, and Clustercorp are based on new 11th-generation Dell PowerEdge servers, the Intel Cluster Ready program, and Rocks+ cluster management software.

New 11th-generation Dell PowerEdge servers

New 11th-generation Dell PowerEdge servers feature performance, energy-efficiency, and systems management enhancements that make them well suited for HPC environments. These servers utilize the quad-core Intel Xeon® processor 5500 series, which are frequency and voltage independent to enable granular control over server performance and power consumption. These processors also incorporate a three-level cache hierarchy with a dedicated level 2 (L2) cache to help reduce replacement-driven cache misses, and a shared L3 cache to help minimize communication between sockets. A three-channel Double Data Rate 3 (DDR3) memory controller integrated directly on each processor helps reduce the latency of local memory accesses. The processor sockets are linked through dedicated communication channels based on Intel QuickPath Interconnect technology and accelerate inter-processor communication for remote memory accesses. When combined, these architectural enhancements help eliminate the memory bottleneck associated with legacy frontside bus platforms. These servers are based on the Intel 5500 and Intel 5520 chipsets, which support registered and unbuffered DDR3 dual in-line memory modules (DIMMs) as well as PCI Express (PCIe) 2.0 expansion cards.

New 11th-generation PowerEdge servers also include additional updates and enhancements such as the Integrated Dell Remote Access Controller 6 (iDRAC 6), which is compliant with the Intelligent Platform Management Interface (IPMI) 2.0 specification. The iDRAC 6 includes an integrated video

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

	PowerEdge R610	PowerEdge R710	PowerEdge M610	PowerEdge M710
Form factor	1U rack server	2U rack server	Half-height blade server	Full-height blade server
Total number of DIMM slots	12	18	12	18
Maximum memory capacity	96 GB	144 GB	96 GB	144 GB
Number of PCIe 2.0 slots	Two x8 slots	Two x4 slots and either two x8 slots or one x16 slot	Two x8 mezzanine card slots	One x4 slot and three x8 mezzanine card slots
Number of embedded network interfaces	Four	Four	Two	Four
Maximum number of Serial Attached SCSI (SAS) or Serial ATA (SATA) hard drives	Six 2.5-inch drives	Eight 2.5-inch drives or six 3.5-inch drives	Two 2.5-inch drives	Four 2.5-inch drives
Systems management	iDRAC 6, iDRAC 6 Express, and optional iDRAC 6 Enterprise		iDRAC 6, iDRAC 6 Express, and iDRAC 6 Enterprise	

Figure 1. *Key specifications of 11th-generation Dell PowerEdge rack and blade servers*

controller and provides systems management features such as console redirection, remote power control, chassis hardware alert generation, and remote media boot support. The iDRAC Express persistent storage device provides the ability to automatically pull system firmware updates from an external repository, configure the BIOS, and run diagnostics, and includes an on-board repository with critical OS drivers. Figure 1 outlines the key specifications of 11th-generation PowerEdge rack and blade servers.[2]

Intel Cluster Ready program

Developed in conjunction with hardware and software vendors such as Dell and Clustercorp, the Intel Cluster Ready program focuses on providing cluster solutions that can run HPC applications right out of the box and do not require special cluster expertise to deploy and maintain.[3] Certified clusters provide a standard interface to the application level, and HPC software vendors in turn register applications that execute on top of that interface—helping simplify the purchasing, deployment, and management of HPC clusters.

Each certified cluster includes the Intel Cluster Checker tool, which provides a simple, methodical way to verify the cluster's operational health. Using a modular framework, Intel Cluster Checker proceeds from basic functionality tests to application-level benchmarks, providing comprehensive feedback on whether the cluster is working properly. This tool helps remove the guesswork from maintaining the system, ultimately contributing to reduced systems management costs. As a result of the collaboration between Dell, Intel, and Clustercorp, specific Dell hardware configurations are Intel Cluster Ready certified, and certification tools such as Intel Cluster Checker have been integrated into the Rocks+ deployment framework.

Rocks+ cluster management software

Rocks+ provides a robust, comprehensive software stack for HPC clusters. It is designed to provide all the necessary software components for building, deploying, and managing a cluster, including the Red Hat® Enterprise Linux® OS,[4] cluster middleware, tightly integrated compilers, applications, and communication libraries. Rocks+ 5 introduces multiple features and enhancements, including the Rocks command-line interface, support for cluster virtualization, updated OS distribution based on Red Hat Enterprise Linux 5, and the Dell Roll, which was designed specifically for 11th-generation Dell PowerEdge servers.

DEPLOYING THE CLUSTER WITH ROCKS+ROLLS

The Dell deployment team installs a cluster by combining a carefully selected set of best-of-breed components, which might include low-latency interconnect networks from QLogic or Mellanox, workload management software such as Cluster Resources' Moab Cluster Suite or Platform Load Sharing Facility (LSF), and open source packages such as the OpenFabrics Enterprise Distribution (OFED). These components are managed and delivered by Rocks+ through modular software components called Rolls. After a cluster has been deployed with Rocks+, the Intel Cluster Checker tool confirms that the cluster is Intel Cluster Ready compliant and functioning correctly. Cluster installation and verification can take place at the deployment site or at the factory through

[2] For more information on the architecture and features of 11th-generation PowerEdge servers, see "Data Center Workhorses: New Dell PowerEdge Rack and Blade Servers," by Edward Yee, Indrani Paul, Robert Tung, Truc Nguyen, and Chad Fenner, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090246-Nguyen.pdf; and "Optimizing DDR3 Memory Settings in New 11th-Generation Dell PowerEdge Servers," by Paul Benson, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20080414-Benson.pdf.

[3] For more information on the Intel Cluster Ready program, visit www.intel.com/go/cluster.

[4] Requires separate licenses from Red Hat.

Dell Server and Storage Rack and Stack Services, in which the clusters are racked, cabled, installed, and verified at Dell. Organizations can then connect remotely to benchmark the cluster before it is delivered to the final site.[5]

The engineering and testing effort behind Dell clusters leaves behind the do-it-yourself, on-the-fly approach from the early days of cluster computing and replaces it with end-to-end services throughout the cluster life cycle. Although this approach is reminiscent of traditional supercomputing, Dell clusters remain both cost-effective and scalable.

Dell Roll

The Dell Roll uses the Rocks+ framework to help further simplify cluster deployment and management by automatically configuring cluster hardware and software that is specific to 11th-generation Dell PowerEdge servers. The Dell Roll installs Dell-approved device drivers across the cluster, applies necessary OS fixes and kernel parameters, and modifies BIOS settings for functionality and performance. The Dell Roll can also be used to deploy and configure Dell systems management software and the iDRAC IPMI management interfaces.

Intel Cluster Ready Roll

The Intel Cluster Ready Roll incorporates the components required by the Intel Cluster Ready specification and automates the setup process for the Intel Cluster Checker tool and many of its parameters, helping significantly simplify first-run certification, health checks, and other administrative tasks. The result is a highly scalable approach enabling Dell to replicate Intel Cluster Ready–compliant solutions with a defined set of hardware and software components. The performance characteristics of the hardware and software components are captured as part of the solution design process so that a conforming cluster can be readily checked for performance regression. The Intel Cluster Ready Roll additionally incorporates best practices—such as optimized configuration of Gigabit Ethernet interrupt coalescing for Message Passing Interface (MPI) traffic—and includes the latest versions of Intel runtime components such as the Math Kernel Library (MKL) with optimizations for the Intel Xeon processor 5500 series architecture.

InfiniBand Rolls

Rocks+ can also incorporate additional Rolls that differentiate it from the base Rocks cluster package. The Mellanox OFED Roll was designed by Clustercorp as a core component of Rocks+. It supports OFED 1.4, which includes drivers and firmware for Mellanox InfiniBand host channel adapters (HCAs), fabric management tools, and multiple MPI environments. It also includes comprehensive support for Intel compilers when paired with the Intel Developer Roll from Clustercorp, and helps remove the guesswork and manual steps to create a validated InfiniBand environment.

Clustercorp also offers the QLogic Roll for organizations using QLogic hardware. This Roll—designed by QLogic and tested by Clustercorp—adds support for QLogic OFED+ for TrueScale or InfiniScale HCAs. Integration of QLogic drivers helps provide a standardized deployment strategy that allows organizations to quickly move through the deployment and acceptance phases of cluster setup.

Moab and LSF Rolls

The Moab Roll adds the Moab Cluster Suite to Rocks+, providing an automated configuration for powerful workload and resource management that integrates scheduling, management, monitoring, and reporting for cluster workloads. Platform Computing also worked closely with Clustercorp to create the LSF Roll, which adds Platform LSF to Rocks+ and combines intelligent workload management with flexible, simplified cluster management.

SIMPLIFYING HPC CLUSTERS

Inspired by the pioneers in traditional supercomputing, when HPC systems were expected to be complete solutions rather than piecemeal projects, Dell, Intel, and Clustercorp have worked together to create HPC clusters designed to be simple, comprehensive, and reliable, helping avoid the inherent complexity of managing disparate components. By combining 11th-generation Dell PowerEdge servers, Intel Cluster Ready configurations, and Clustercorp Rocks+, a cluster can be assembled, customized, and verified on the factory floor—helping make it ready for production use right out of the box and enabling administrators and researchers to focus on running their applications rather than on complex cluster tuning and management.

Garima Kochhar is a systems engineer in the Scalable Systems Group at Dell.

Jacob Liberman is a development engineer in the Scalable Systems Group at Dell.

Brock A. Taylor is a staff engineer in the Software and Services Group at Intel.

Jason Bishop is vice president of engineering at Clustercorp.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell HPC solutions:
DELL.COM/HPC

Intel Cluster Ready:
www.intel.com/go/cluster

Clustercorp Rocks+:
www.clustercorp.com

[5] For more information, visit DELL.COM/ServerStorageDeployment.

SIMPLIFY YOUR STORAGE

TO IMPROVE OVERALL IT EFFICIENCY

INTRODUCING THE EQUALLOGIC™ PS6000 SERIES!

THE NEWEST ENTERPRISE-CLASS STORAGE SOLUTION FROM DELL.

Your storage solution should be simple, from initial purchase to scaling out your SAN. With a single, intuitive GUI for multiple group management, plus all-inclusive features like Smart Copy for Microsoft® Hyper-V™* and support for VMware® vStorage, the EqualLogic PS6000 Series iSCSI SAN streamlines and simplifies management while protecting your data. What's more, its patented "peer" architecture allows for on-demand scaling, enabling a true pay-as-you-grow iSCSI SAN solution to help improve overall IT efficiency.

SIMPLIFY YOUR STORAGE AT DELL.COM/PS6000

DELL

BOOSTING SAN PERFORMANCE WITH DELL EQUALLOGIC PS6000S SOLID-STATE DRIVE ARRAYS

By Dylan Locsin
Ujjwal Rajbhandari
Wendy Chen

Solid-state drives (SSDs) now offer a new tier of high-performance, highly reliable Dell™ EqualLogic™ PS Series storage for performance-hungry enterprise applications. The Dell EqualLogic PS6000S array provides cost-effective SSD-based storage that can easily integrate into multi-tiered storage area network (SAN) environments—providing a simplified, scalable way to help meet the performance demands of targeted applications.

Related Categories:

Dell EqualLogic storage

Storage

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

As online transaction processing (OLTP) and other key enterprise applications take on increasing workloads, administrators are pressed to find ways of handling the performance demands on their storage systems. Traditional mechanical hard disk drives (HDDs) have built-in performance characteristics such as seek time and rotational latency that can make it difficult to keep up with escalating needs from transaction-intensive applications.

To deliver adequate performance, many IT administrators rely on work-arounds such as wide striping data on a large number of high-speed disks or short stroking data only on the outer portion of a spinning disk platter. Although these tactics may get the job done, they often leave large amounts of storage underutilized, which can lead to increased complexity and costs from over-purchasing capacity and extra power and cooling for the additional disks.

A new option now available for Dell EqualLogic PS Series storage is solid-state drive (SSD) technology. SSDs offer exceptional performance—enabling significantly faster random read/write response time compared with traditional mechanical HDDs—along with enhanced reliability, energy efficiency, and space efficiency. Historically, SSD technology has often been cost prohibitive, but it can now offer an effective option for targeted application workloads that require top-flight storage performance.

To help organizations take advantage of SSD technology, the Dell EqualLogic PS6000S storage array provides outstanding performance in a cost-effective SSD-based array that can easily integrate into multitiered EqualLogic Internet SCSI (iSCSI) storage area networks (SANs). The EqualLogic PS6000S array is part of the EqualLogic PS6000 series of arrays, which support a range of Serial ATA (SATA), Serial Attached SCSI (SAS), and SSD options that administrators can combine into a high-performance, cost-effective storage platform that is scalable, reliable, and easy to manage.

UNDERSTANDING THE BENEFITS AND CHALLENGES OF SSD TECHNOLOGY

SSD technology can be an optimal solution for administrators who need to provide high-performance storage for specific application workloads such as OLTP database applications. SSDs offer several key advantages over traditional mechanical storage devices, including the following:

- **High performance:** Unlike mechanical HDDs, SSDs have no moving parts. As a result, they have no rotational latency and minimal seek time, and

can deliver significantly faster random read/write response time compared with mechanical HDDs.

- **Low power consumption:** Because SSDs have no moving parts and no motor, they characteristically consume less power and generate less heat than mechanical HDDs. SSDs also usually draw power only during actual read and write operations, consuming little or no power when idle.
- **Enhanced reliability:** Because SSDs have no moving parts, they are less prone to damage from shock or vibration than mechanical HDDs. Consequently, SSDs characteristically are designed to offer higher mean time between failure (MTBF) than mechanical HDDs, making them suitable for rugged environments such as military operation zones and industries in which field deployment can subject the hardware to dust, vibration, and other harsh conditions. The MTBF specified for the 50 GB SSD offered in the Dell EqualLogic PS6000S array, for example, is 2,000,000 hours.
- **Enhanced cost-effectiveness:** Although SSDs typically have a higher cost-per-gigabyte ratio than mechanical HDDs, when used for targeted high-performance application data hotspots—such as tempdb and index tables in Microsoft® SQL Server® or Oracle® database applications—they can turn out to be a cost-effective alternative for achieving comparable or enhanced performance levels to wide striping or short stroking mechanical HDDs. Also, because of their inherent seek and rotational delays, adding HDDs cannot reduce latency to necessary levels for certain workloads.
- **Small form factor:** SSDs also have a much smaller form factor and typically weigh less than mechanical HDDs, helping to reduce some of the physical burdens on data center rack and floor infrastructure.

Despite the many benefits of SSD technology, SSD deployments do present some challenges. For example, SSDs currently offer lower capacity than mechanical HDDs, and that capacity comes at a relatively high cost per gigabyte. Consequently, they are typically practical only in targeted circumstances. Also, it is essential that administrators follow best practices for optimal deployment when integrating SSDs into existing SANs, to help avoid the potential for diminished performance within the overall storage infrastructure.

For example, when SSDs are deployed in a traditional frame-based, dual-controller architecture, limitations in controller capacity or expandability may restrict SSD performance. Because SSDs in this type of deployment are served by the same controllers that support non-SSD devices, the increase in I/Os per second (IOPS) from the SSDs can potentially overwhelm the I/O capacity of the controllers as drives are added. The resulting bottleneck or controller resource contention can diminish the performance of non-SSD devices and contribute to a reduction in the overall performance, effectiveness, and scalability of the entire SAN.

To avoid such pitfalls, IT organizations must carefully consider how incorporating SSDs into their SAN arrays can affect their storage architecture, and take into account the needs of their other application storage. One work-around is to deploy multiple islands of arrays to accommodate SSDs and HDDs separately. However, that approach can add cost and complexity to the storage infrastructure by increasing points of management and negating some of the consolidation advantages that moving to a shared SAN provides.

Instead, SSD arrays can be particularly effective when deployed as part of a tiered environment that delivers targeted storage performance and capacity according to differing workload requirements. For example, an optimized tiered storage environment might combine SATA drive arrays for high-capacity, moderate-performance workloads (tier 2 storage); SAS drive arrays for moderate-capacity, high-performance workloads (tier 1 storage); and SSD arrays for extremely high-performance workloads with limited capacity (tier 0 storage).

FOLLOWING BEST PRACTICES FOR SSD DEPLOYMENTS

Although SSDs can offer dramatic performance advantages, their cost and capacity limitations make them most effective when deployed as a targeted solution in particular use cases. To determine whether a particular scenario is well suited for SSD deployment, administrators must consider several factors, including capacity, performance requirements, and block size. For example, SSDs can be a key solution for workloads with high performance requirements, including low latency and high IOPS, combined with limited capacity.

Compared with traditional HDDs, SSDs are also particularly well suited to situations in which the transfer block size is small, such as OLTP applications. In these cases, SSDs can greatly mitigate the seek and rotational latency delays inherent in mechanical HDDs, and can sometimes be the only viable option in situations where HDD strategies such as wide striping and short stroking cannot satisfactorily eliminate latency or increase transaction throughput. Applications with data sets that have very high read-only or high read/write ratio workloads, such as tempdb and index tables, have the potential to benefit tremendously from SSDs. In contrast, performance benefits provided by SSDs may be less compelling for online analytical processing (OLAP) or data warehousing types of applications that have large-block sequential reads and large storage requirements.

SSDs are also a good option for high-IOPS, low-latency scenarios such as virtual desktop infrastructures in which end-user desktop images are moved to a centralized server and virtualized. Because virtual desktop infrastructure deployments typically generate a high volume of server-based storage transactions, the storage I/O rate is often very high, while at the same time response time must be kept very low to facilitate end-user productivity.

Additionally, some virtual desktop scenarios, such as the VMware® View approach, utilize a linked-clone architecture in which a master desktop image is designed to be accessible by all users. The master desktop image is stored on the SAN, and user-specific data is stored separately and linked to the master copy. Because the master copy requires high IOPS and low latency, this kind of virtual desktop deployment can also be well suited for an SSD-based solution.

INTEGRATING SSD ARRAYS IN A MULTITIER STORAGE ENVIRONMENT

The Dell EqualLogic PS6000S SSD array can be combined with added or existing EqualLogic SAS and SATA arrays to form

BENCHMARKING DELL EQUALLOGIC PS6000S SSD ARRAYS

Solid-state drives (SSDs) have the potential to deliver outstanding performance relative to traditional mechanical hard disk drives (HDDs), especially for targeted high-performance workloads. The Dell EqualLogic PS6000S SSD arrays are priced similarly to midrange EqualLogic PS6000E Serial ATA (SATA) arrays, and in certain usage scenarios, the EqualLogic PS6000S SSD array can deliver higher performance at a lower cost than a traditional high-performance HDD array.

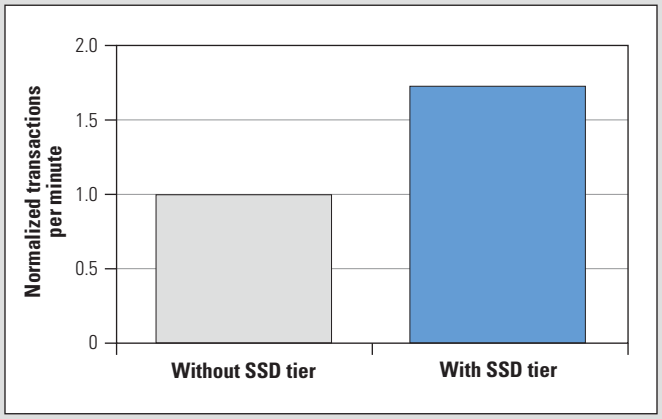

Figure A. *Comparing transactions per minute with and without an SSD tier*

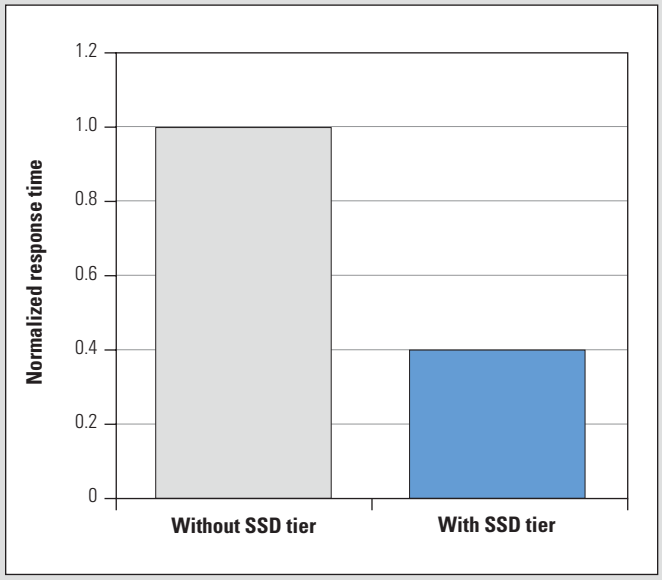

Figure B. *Comparing response times with and without an SSD tier*

To evaluate the performance of the EqualLogic PS6000S SSD array relative to traditional HDDs, in March 2009 Dell's Oracle Solutions team utilized the Benchmark Factory for Databases simulation tool from Quest Software to benchmark the performance of two EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) deployments, one with a tier of SSD storage and one without SSD storage. The test environment configuration included two Dell PowerEdge™ M710 blade servers, each with two quad-core Intel® Xeon® processors at 2.67 GHz, 24 GB of RAM, and four dual-port Broadcom NetXtreme II 5709 Gigabit Ethernet network interface card ports for iSCSI traffic. The external storage consisted of two EqualLogic PS6000XV arrays and one EqualLogic PS6000S array in RAID-10 configurations with two spare drives in each member, using 15,000 rpm Serial Attached SCSI (SAS) drives in the EqualLogic PS6000XV arrays and 50 GB SSDs in the EqualLogic PS6000S array. The network interconnect included two stacked 48-port Dell PowerConnect™ 6248 copper Gigabit Ethernet Layer 3 switches for the SAN. In addition to the 64-bit Oracle Real Application Clusters (RAC) 11.10.7 Enterprise Edition workload, the OS and device drivers included the Microsoft Windows Server® 2003 Release 2 (R2) Enterprise x64 Edition OS with Service Pack 2 (SP2), Microsoft iSCSI Software Initiator 2.0.8, and the EqualLogic Multipath I/O device-specific module.

One deployment used only the two EqualLogic PS6000XV arrays with SAS drives, and the other deployment used one EqualLogic PS6000XV array with SAS drives and one EqualLogic PS6000S array with SSDs, with the read-intensive data sets isolated on the SSDs. The database configuration for both deployments was a two-node Oracle RAC cluster with a database schema size of 130 GB.

The test team calculated both the transactions per minute (TPM) and average transaction response time of the two deployments, each performing the same Oracle online transaction processing (OLTP) workload. In the test configuration, the SAN that included the SSDs demonstrated superior throughput and response time compared with the SAN that did not include the SSDs. Compared with the baseline system configured entirely with mechanical SAS drives, the SAN that included the SSDs delivered up to a 75 percent increase in TPM or up to a 60 percent reduction in average transaction response time (see Figures A and B). Of course, actual performance will vary based on configuration and usage patterns.

The excellent I/O throughput and extremely fast response time demonstrated in these test results indicate that the EqualLogic PS6000S SSD array is suitable for high-performance, low-latency workloads such as OLTP database applications.

a tiered storage environment (see Figure 1). This approach enables IT organizations to provide a customized balance of cost-effective capacity and performance with exceptional throughput and response time (see the "Benchmarking Dell EqualLogic PS6000S SSD arrays" sidebar in this article). In addition, the modular Dell EqualLogic peer storage architecture is designed for seamless integration of multiple EqualLogic arrays, regardless of which hardware generation they belong to or which type of drive they contain.

The EqualLogic storage architecture can overcome the obstacle of controller resource contention from SSDs that may affect traditional storage designs. Because each EqualLogic storage array in a SAN—including the EqualLogic PS6000S array—comes with its own set of controllers with cache memory and network links to service the I/O needs of its disks, SAN performance can continue to scale even as arrays with any drive type are added. By segmenting a pool of SSD arrays, administrators can help ensure that their high-priority workloads have a dedicated set of controllers. Conversely, important but less demanding workloads from other drives also retain their own sets of controller resources. In this way, EqualLogic arrays can take advantage of the performance gains made possible by SSDs while avoiding performance degradation that may result from contention for controller or network port resources, or from the increased complexity and management overhead that may be incurred when administrators deploy separate SAN islands.

Ease of integration can be another benefit of including SSD technology in the EqualLogic storage architecture. EqualLogic PS6000S arrays, like other EqualLogic PS Series arrays, are designed to be rapidly deployed in an EqualLogic SAN without incurring downtime.

INNOVATING STORAGE INFRASTRUCTURE WITH SSD ARRAYS

SSDs have emerged as an excellent high-performance storage option for certain enterprise application workloads. As part of the Dell EqualLogic PS Series of virtualized storage arrays, the EqualLogic PS6000S SSD array can help organizations simply, scalably, and cost-effectively meet escalating application performance demands with enhanced reliability and energy efficiency.

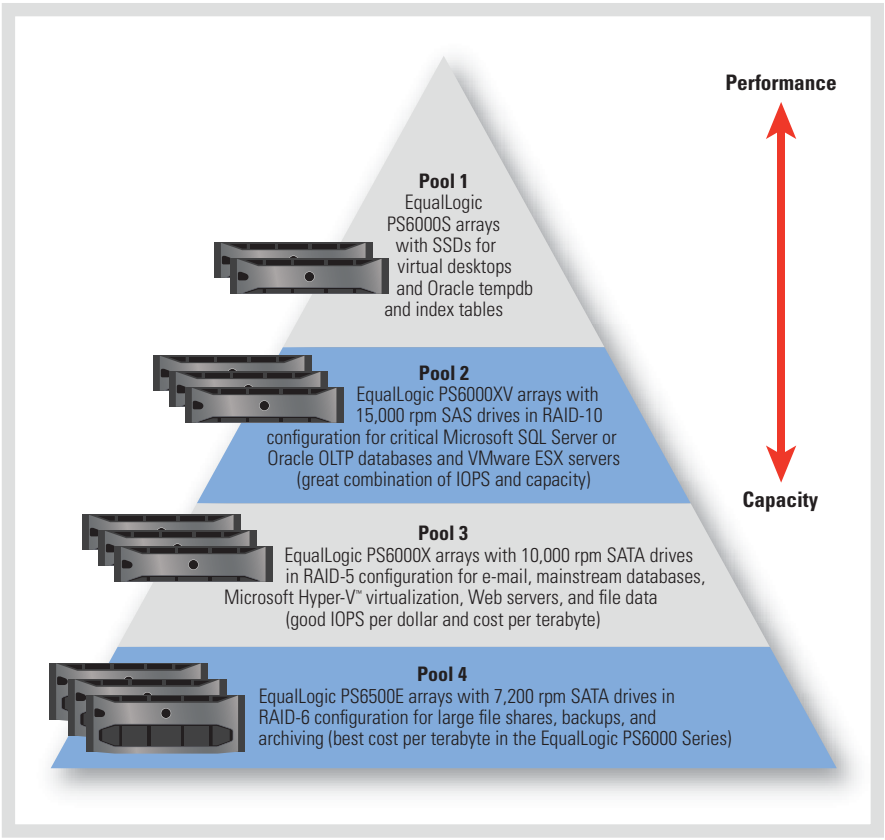

Figure 1. *Matching Dell EqualLogic PS6000 series storage options to workload performance needs*

Dylan Locsin is a product marketing consultant for Dell storage solutions. He has nine years of experience in marketing and communications for storage, networking, and enterprise software technologies for companies including Dell, EqualLogic, NSI Software, Onaro, and NetScout Systems. He has a B.A. from Tufts University.

Ujjwal Rajbhandari is a product marketing consultant for Dell storage solutions. He has a B.E. in Electrical Engineering from the Indian Institute of Technology, Roorkee, and an M.S. in Electrical Engineering from Texas A&M University.

Wendy Chen is a systems engineer on the Oracle Solutions team of Enterprise Solutions Engineering, Dell Product Group. Her primary focus includes Oracle RAC, performance characteristics, disaster recovery, and data protection solutions, and she has previously been an Oracle database administrator. She has an M.S. in Computer Science from Brigham Young University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell EqualLogic PS Series:
DELL.COM/EqualLogic
DELL.COM/PSSeries

NEW TECHNOLOGIES SPEED THE MOVE TO 10 GIGABIT ETHERNET DATA CENTER CONNECTIVITY

By Sunil Ahluwalia

To handle the demands of increasingly bandwidth-hungry applications, IT managers are turning to 10 Gigabit Ethernet technology for virtualized servers and networked storage platforms. New multi-core processor architecture and network interface card technologies are designed to meet rapidly evolving requirements for a truly dynamic data center.

Related Categories:

Data center technology

Dell PowerEdge servers

Ethernet

Fibre Channel

Intel

Internet SCSI (iSCSI)

Virtualization

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

Cost-conscious IT departments have adopted data center virtualization as a way to help increase efficiency and reduce total cost of ownership. As virtualized server and storage deployments increase, infrastructure components are evolving to keep pace. These collective enhancements promise unmatched data center efficiencies, but they pose new challenges too. Systems designed for virtualized environments, such as the Dell™ PowerEdge™ R710 server based on the multi-core Intel® Xeon® processor 5500 series architecture, can provide the compute power required to support a large number of virtual machines (VMs). At the same time, the aggregation of VMs and applications on a server raises the demand for I/O capacity.

Server consolidation can also intensify network storage requirements to support the data needs of VMs, including backup and live migration. By allowing cost-effective Ethernet infrastructure to be used as a storage area network (SAN) fabric, the Internet SCSI (iSCSI) standard has fueled rapid adoption of iSCSI SANs such as Dell EqualLogic™ PS Series arrays. Increasing demand for SAN connections can drive the need for increased I/O.

Escalating I/O demand is driving the data center framework toward 10 Gigabit Ethernet (10GbE) connectivity, with servers, network devices, and storage arrays designed to scale performance dynamically in the face of constantly changing compute needs. The growth in 10GbE implementations is being fueled by increasingly cost-effective per-port capability. Meanwhile, the emergence of the Fibre Channel over Ethernet (FCoE) standard is expected to quicken the pace of 10GbE deployment by fostering a unified data center network.

Many elements of the dynamic 10GbE data center are already in place. The generation of servers based on the Intel Xeon processor 5500 series architecture enables the processing power and adaptability needed to support growing VM deployment without increasing power or space requirements. Intel 10GbE server adapters are designed to take advantage of the capabilities of these platforms and provide the necessary I/O bandwidth.

PROCESSOR ARCHITECTURE: ENHANCING I/O SCALABILITY

The Intel Xeon processor 5500 series combines multiple processors with architectural features that allow for a high level of 10GbE scalability—including specifications designed to enable faster memory, faster interconnect architecture, and a faster PCI Express (PCIe) bus than previous Intel processor architecture generations (see Figure 1).

An integrated memory controller with Double Data Rate 3 (DDR3) memory is designed to enable peak

memory bandwidth of up to 32 GB/sec per socket, which is significantly higher than the memory bandwidth that was specified on previous platforms for the entire system. Another key difference is that the frontside bus has been replaced by an Intel QuickPath Interconnect (QPI) link that provides dedicated, high-speed communication between processors.

To support high-bandwidth I/O and the ability to scale across multi-port 10GbE, the new processor architecture also uses a second-generation PCIe 2.0 bus from the I/O hub to the network interface card (NIC). The PCIe 2.0 bus is designed to provide twice the transfer rate of first-generation PCIe.

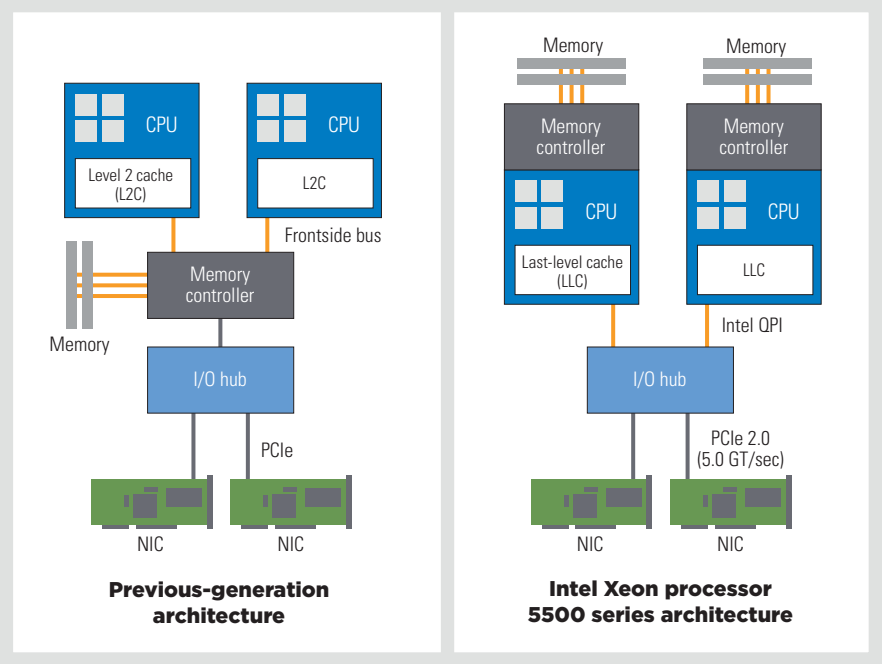

Figure 1. *I/O enhancements of the previous-generation architecture in the Intel Xeon processor 5500 series architecture*

NETWORK CONTROLLER: OPTIMIZING 10GbE FOR I/O VIRTUALIZATION

At the network interface, Intel 10GbE server adapters include Intel Virtualization Technology for Connectivity (VT-c) to enhance I/O performance in virtualized environments. Intel VT-c includes hardware optimizations that help reduce I/O bottlenecks and increase server performance. The Intel VT-c suite consists of Intel Virtual Machine Device Queues (VMDq) and Intel Virtual Machine Direct Connect (VMDc).

VMDq enhances data processing by offloading the network traffic sorting and queuing functionality from the VM monitor in the hypervisor to the Ethernet controller. VMDc is designed to provide direct connectivity to VMs that enables near-native performance and VM scalability.

Intel 10GbE server adapters support iSCSI and network attached storage. Intel 10GbE server adapters on Dell PowerEdge servers can accelerate iSCSI traffic by implementing key stateless offloads such as TCP segmentation offload (TSO). They also support iSCSI initiators from Microsoft®, Linux®, and VMware® platforms and provide a robust iSCSI remote boot implementation. In the next generation of Intel 10GbE Ethernet, planned performance-enhancing features include FCoE offloads and receive-side coalescing (RSC).

CONSOLIDATION: TRANSFORMING THE DATA CENTER

A fundamental shift is taking place in enterprise data centers as IT managers leverage virtualization to consolidate applications and data onto fewer physical servers and storage devices compared with traditional non-virtualized environments. Dell PowerEdge servers with powerful multi-core processors and virtualized iSCSI SANs such as the Dell EqualLogic PS Series arrays help facilitate this transformation to a flexible and cost-effective data center. As enterprise growth requires these systems to continue aggregating VMs and storage, 10GbE connectivity enables the I/O capacity to support increased performance demands and bandwidth-hungry applications.

Together with an ecosystem of powerful applications and today's cost-effective 10GbE networking hardware, innovations in the Intel Xeon processor 5500 series architecture and Intel 10GbE server adapters are helping drive the advance toward an emerging data center model that dynamically adapts to changing business needs.

Sunil Ahluwalia is a senior product line manager responsible for 10GbE products and marketing storage and security technologies for the data center at Intel. He has a bachelor's degree in Computer Science and Engineering from the Institute of Technology in Varanasi, India, and an M.B.A. from the Charles H. Lundquist College of Business at the University of Oregon.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Intel 10GbE adapters:
www.intel.com/go/10gbe

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell EqualLogic PS Series:
DELL.COM/EqualLogic
DELL.COM/PSSeries

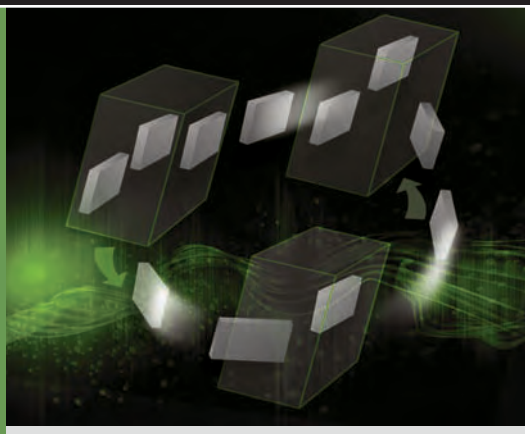

MEETING ECONOMIC CHALLENGES THROUGH ACCELERATED VIRTUALIZATION

By Steve Burrows
Ron Oglesby

In difficult economic times, many enterprises are seeking ways to streamline IT budgets without affecting operations. Evaluated and refined across thousands of real-world deployments, Dell™ virtualization solutions and services offer a comprehensive approach designed to quickly reduce total cost of ownership and maximize return on investment.

Related Categories:

Virtualization

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

The economic downturn has left CIOs under major pressure to limit IT spending. Many enterprises have already reduced expenditure and are focusing on cost optimization and increasing output from existing resources, and the challenge for CIOs is to cut IT operating expenses while maximizing productivity to support continued expansion.

Identifying and addressing hidden IT costs—without cutting deeply into investment capacity—can help organizations adjust to lean times and position themselves for ongoing success. Estimates based on worldwide Dell Global Infrastructure Consulting customer field experience indicate that there are a number of hidden costs with servers:

- **Power and cooling for underutilized servers:** A typical server at low utilization levels can use 60–70 percent of its total power requirements even when idle, and a typical IT environment may have 80–90 percent of its server capacity underutilized in some instances. This inefficient use of resources can typically result in 20–50 percent higher spending on power and cooling than is necessary for the work performed.
- **Data center facilities:** The data center itself requires large-scale investments in building and infrastructure.
- **Systems management:** A lack of standardization across a large number of servers can lead to high maintenance and support charges, and a lack of automated systems management and streamlined administration can lead to lost staff efficiency.

By helping optimize utilization of existing physical resources and reduce the number of physical servers that must be powered, cooled, and managed, virtualization offers a key way to reduce server total cost of ownership (TCO) and increase IT agility. Many enterprises have proceeded cautiously with virtualization to help ensure that they understand the impact it could have on their operations as well as the advantages it can provide. Consequently, even those organizations that have implemented virtualization often are not realizing its full benefits.

Evaluated and refined across thousands of real-world deployments, Dell virtualization solutions and services offer a comprehensive approach designed to reduce TCO and maximize return on investment (ROI). By taking advantage of this experience, organizations can accelerate their adoption of virtualization technology and achieve significant savings in both capital and operational expenditures within the first year.

RETHINKING THE DATA CENTER

According to Gartner, "Virtualization is having a huge impact on infrastructure architectures, processes and operations, people and skills, and the business. It is the most important and impactful trend in infrastructure and operations through 2012—changing how you manage, how/what you buy, how you deploy, how you plan, how you charge."[1]

In short, virtualization encompasses more than simply implementing technology; people and processes are equally important. Addressing all of these aspects in an integrated way helps create the visibility and accountability to accelerate adoption and ROI from virtualization initiatives.

One often overlooked yet important component of a virtualized infrastructure is centralized storage. According to Gartner, "Many respondents [to Gartner's survey] believe that storage technologies, such as thin provisioning, are critical to successful server virtualization, and all agree that server virtualization implementations fundamentally change how storage is backed up and managed."[2] Organizations must consider the impact their virtualization initiatives could have on their current storage and backup environments, including not only the storage capacity required for the virtual machines (VMs) themselves, but also how business continuity and disaster recovery will be affected, how tiering strategies may change based on data use, how high-availability features and flexible resource utilization could affect storage, how data deduplication could help optimize capacity requirements, and how to avoid weak data protection processes that could increase risk.

When properly implemented, virtualization can bring a host of advantages in enterprise environments, such as efficient utilization of storage capacity and resources (see Figure 1). To effectively accelerate the benefits of virtualization, however, organizations must address storage concurrently with the overall virtualization initiative.

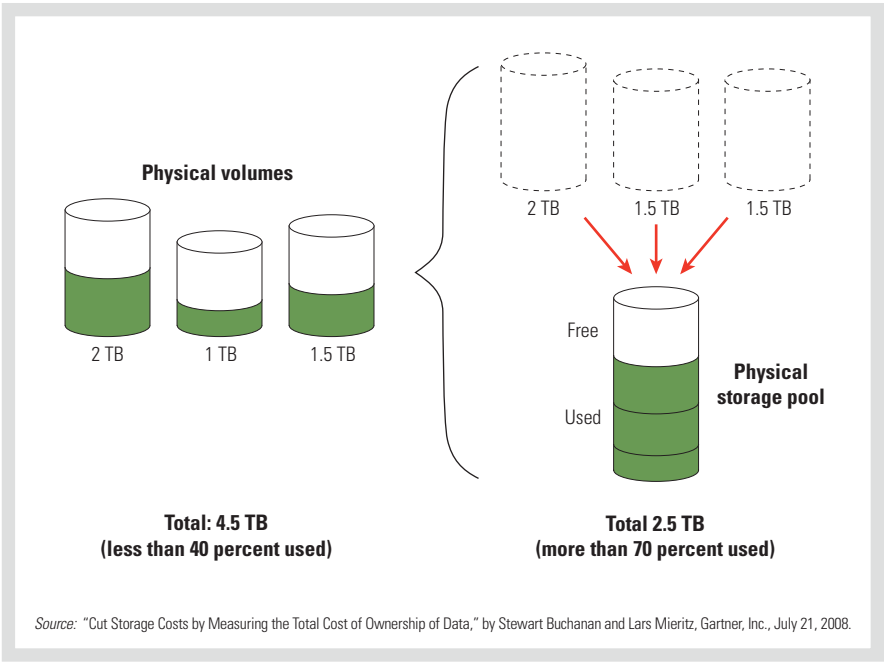

Figure 1. *Improved storage use as a result of thin provisioning virtual volumes*

MAXIMIZING RETURN ON INVESTMENT

Although the organizational approach to virtualization should be strategic, it is often tactical: after reading multiple virtualization success stories, enterprise IT leaders want to move quickly to realize similar cost and agility benefits. Historically, however, these types of technology decisions have tended to be opportunistic, seizing an immediate opportunity without reference to a comprehensive plan.

Virtualization can indeed deliver significant advantages, and the tactical approach can enable enterprises to realize those benefits. However, this approach may not be enough to maximize those benefits. A holistic view of virtualization can help organizations develop a comprehensive strategy to accelerate both implementation and results.

Figure 2 illustrates how this type of strategy can help increase ROI. In this figure, P1 represents the ROI associated with a typical path virtualization solution, and P2 represents the ROI associated with a comprehensive, customized process. The P1 ROI is initially higher than the P2 ROI because this approach does not require an up-front investment in a strategic plan. However, it also plateaus over time, while the P2 ROI continues to increase. In financial terms, the strategic approach generates a positive net present value, which translates into realizing larger savings more quickly than the typical approach.

As an example, consider the following enterprise usage scenario based on aggregated Dell field experience with multiple large companies. In this example, a large company faces multiple IT challenges, including centralizing and streamlining operations through a shared infrastructure, consolidating its data centers to help reduce operational costs, and overcoming organizational resistance to virtualization. This company has 130 locations of various

[1] "Virtualization Changes Virtually Everything," by Thomas Bittman, Gartner, Inc., Gartner Infrastructure Operations and Management Summit, June 23–25, 2008.

[2] "Market Trends: Storage Professional Services, North America, 2008," by Adam W. Couture, Gartner, Inc., August 21, 2008.

sizes, over 15,000 x86-based servers, an overall server utilization level of approximately 10 percent, and a 50:1 ratio of servers to system administrators. Multiple business units are pursuing different optimization strategies, and the company as a whole lacks defined policies or processes for server virtualization, consolidation, or decommissioning.

Initially, this company might seek to optimize its x86-based infrastructure through rational consolidation, decommissioning, and virtualization, including establishing a "VM first" policy for handling new requests. At first the optimizations proceed relatively slowly; however, by establishing an end-to-end process that involves all stakeholders, the company is able to move from tens to hundreds of optimizations per month in the first year. This process is supported by a dedicated optimization organization devoted to achieving the plan, and both the stakeholders and the optimization team have the visibility and accountability, through a project portal, to achieve the plan. Dell participates in this initiative by supporting the company's program management office and providing process efficiency consulting, a template for the project portal, and Dell resources to supplement existing staff and help achieve the desired number of optimizations per month.

The results of the initiative in this scenario could include consolidating the 130 locations to 4 major data centers, reducing the total number of physical servers by 40 percent, increasing server utilization levels to 60 percent, and achieving a 200:1 ratio of servers to system administrators—ultimately helping reduce costs by up to US$30 million in the first year. In the second year, the optimization team could seek to further accelerate the plan—including scaling its optimization adoption rate to 700 servers per month—to attain cost savings of up to US$40 million. In the future, this example company might then plan to expand the optimization scope to include its 5,000 UNIX® OS–based servers.

Although this scenario represents a large infrastructure, the principles can be scaled down and applied to midsize businesses with 250 or more physical servers. The approach itself can be customized to help meet the specific needs of individual organizations—for example, for organizations that prefer to manage rather than execute the server virtualization process, Dell can provide an "optimization factory" for that purpose.

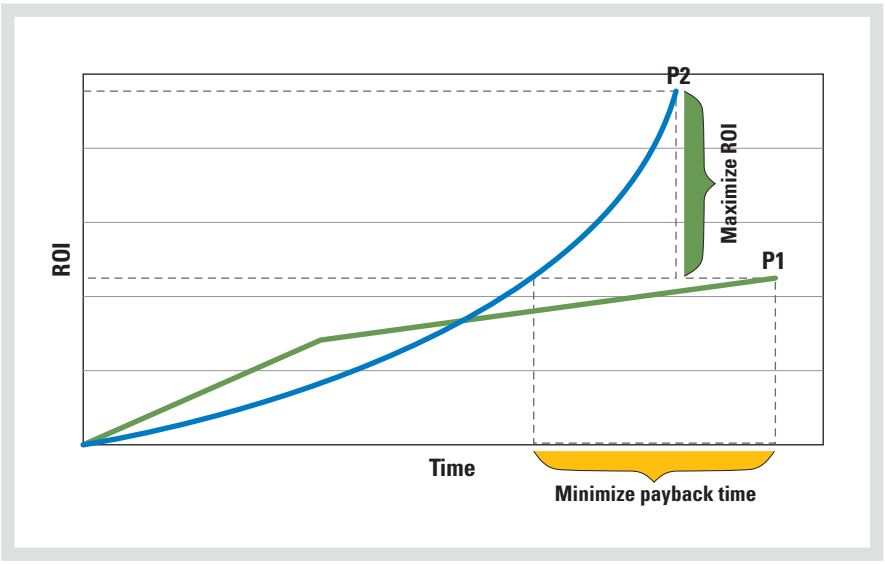

Figure 2. *How a strategic approach to virtualization can help maximize return on investment*

OVERCOMING DEPLOYMENT ROADBLOCKS

After deciding to implement a virtualization initiative, organizations should take three specific steps to help ensure that they can carry out the project quickly and accelerate the resulting ROI.

Step 1: Stop IT spending on new physical servers

Although this step often means stopping projects and/or server life-cycle replacement, server virtualization can enable organizations to reduce their budgets by creating a "VM first" policy stating that all systems being replaced or purchased new will be deployed as VMs. Under this policy, systems are deployed as physical hardware only when there are specific technical reasons to do so. This change in the environment helps continue accelerating the virtualization adoption rate, can help reduce hardware spending by about 50 percent (and in some cases by up to 60–70 percent), and can help reduce power and cooling costs by up to 90 percent.

Step 2: Identify and plan to remove existing depreciated servers

Even without purchasing new servers, organizations are still paying to operate existing servers on a daily basis: existing x86-based servers, for example, can cost between US$700 and US$900 annually in power and cooling alone, not even including additional costs if these servers are hosted in an outsourced or rented facility. By focusing on accelerating the consolidation of these servers to VMs, organizations can reduce monthly costs by approximately US$60 per month for each server. Migrating 200 servers to VMs would then reduce these costs by approximately US$12,000 per month.

This first step in this process is to identify depreciated servers in the environment. In most cases these servers are four or five years old and configured with processors at 2 GHz or slower, which can easily be hosted as VMs. Administrators should prioritize these servers and plan to migrate them as quickly as possible.

Step 3: Identify virtualization roadblocks

With a virtualization policy in place and a list of targeted servers for migration, organizations should next determine how rapidly they can migrate to the new environment. Speed is key to success, because the faster servers can be migrated, the sooner the project can begin returning savings. From an operational perspective, the effect is cumulative: every month a workload is running on a physical server rather than a VM means additional costs that cannot be recouped.

There are four major roadblocks that can slow down migrations, and organizations should address each in a way that can help reduce the cost of the migration while still maintaining a quality process:

- **Untouchable servers:** Too often, servers host applications that are considered untouchable. In today's financial environment, however, companies cannot afford this luxury. Obtaining executive-level sponsorship of the project can help push past political barriers to virtualization.
- **Resource constraints:** IT staffs are typically built to handle maintenance and minor projects, so a major infrastructure shift can overburden the staff and limit the number of migrations that can be accomplished. In this case, organizations could shift resources from other projects to focus on virtualization (especially given that other projects are unlikely to offer the cost savings virtualization can provide) or outsource the migration to a third party. Outsourcing can be expensive, depending on the IT staff's level of expertise and internal organizational processes—but if this approach can keep the cost per migration lower than the operational savings, it can offer a net savings.
- **Inadequate storage and backup infrastructure:** An inadequate storage and backup infrastructure can slow virtualization initiatives and reduce ROI. Organizations should be sure to evaluate this infrastructure early in the process and upgrade if necessary.
- **Time-consuming internal processes:** Migrations are often slowed by an organization's own internal processes, which can be cumbersome, time-consuming, and taxing on staff resources. To help avoid this problem, organizations should determine where automation can help, then determine which processes do not apply to the virtualization migration and remove those processes. To help further accelerate the migration and ROI, they should streamline procedures and processes that do apply to the virtualized environment and enable them to be performed in batches.

BUILDING A STREAMLINED IT INFRASTRUCTURE

Virtualization offers a way for enterprises to reduce server TCO and achieve rapid ROI while still meeting or exceeding IT service-level agreements. Although the natural reaction to difficult economic times is to slow down projects or stop all spending, waiting on a virtualization deployment means spending more in the long run simply to maintain and manage an inefficient existing infrastructure. Investing in a streamlined, aggressively accelerated virtualization initiative and overcoming roadblocks to deployment can help increase flexibility, simplify management, and quickly meet reduced budget goals for IT expenditure.

Steve Burrows is a practice executive in Dell Services focused on virtualization, operational effectiveness, and infrastructure optimization strategies.

Ron Oglesby is a practice executive in Dell Services focused on virtualization and core technologies around the virtual server ecosystem.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell virtualization solutions:
DELL.COM/Virtualization

Dell Virtualization Practice:
DELL.COM/VirtualizationConsulting

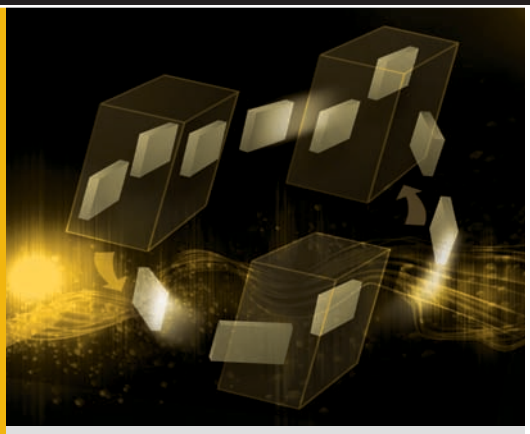

PROTECTING MICROSOFT HYPER-V VIRTUALIZED ENVIRONMENTS USING SYMANTEC BACKUP EXEC

By Ed Casmer

As virtualization becomes increasingly widespread, organizations need backup tools that can effectively protect critical data in virtualized environments. Symantec™ Backup Exec™ 12.5 software, part of the Dell™ PowerVault™ DL2000 – Powered by Symantec Backup Exec, offers simplified, flexible, and robust backup and recovery for environments based on Microsoft® Hyper-V™ virtualization.

Related Categories:

Backup

Backup, recovery, and archiving (BURA)

Data consolidation and management

Dell PowerVault storage

Microsoft Hyper-V

Symantec

Systems management

Virtualization

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

Server virtualization can offer a number of compelling benefits in enterprise data centers, including increased resource utilization, reduced hardware costs, and operational efficiencies. Despite these advantages, however, virtualized environments still raise many of the same concerns as traditional environments—and, critically, the importance of data remains the same whether it resides on a virtual machine (VM) or a non-virtualized server. Given the tools and other processes introduced by virtualization, how can organizations effectively protect the data on virtualized servers and recover it quickly when necessary?

Administrators using traditional backup and recovery tools not designed for virtualization may be struggling with legacy approaches and architectures that add unnecessary time and complexity to implementing data protection strategies. Symantec Backup Exec 12.5, part of the Dell PowerVault DL2000 – Powered by Symantec Backup Exec, provides industry-leading capabilities designed to protect virtualized environments, including those based on Microsoft Hyper-V technology.[1] By taking advantage of these capabilities and following best practices for using Backup Exec with Hyper-V, administrators can implement simplified, flexible, and robust backup and recovery in their virtualized environments.

DATA PROTECTION IN VIRTUALIZED ENVIRONMENTS

The Microsoft Hyper-V hypervisor is a server role integrated directly into Microsoft Windows Server® 2008 x64 Edition operating systems. Like other hypervisors, it enables administrators to consolidate hardware by running multiple VMs on a single physical host server. Its rich feature set is designed to help organizations take advantage of all aspects of server virtualization.

Hyper-V has quickly become a compelling choice for organizations that have standardized on Microsoft platforms. However, virtualization does come with its own disadvantages—for example, the failure of a Hyper-V host server adversely affects all of the VMs on that host, which might include applications such as Microsoft Exchange, Microsoft SQL Server®, and Microsoft Office SharePoint® Server software installed

[1] For more information on the Dell PowerVault DL2000 – Powered by Symantec Backup Exec, see "Simplified Data Protection with Disk-Based Backup from Dell and Symantec," by Sanjeet Singh and Charles Butler, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20080444-Symantec-M.pdf. For more information on using Symantec Backup Exec to protect VMware virtualized environments, see "Best Practices for Protecting VMware Infrastructure 3 with Symantec Backup Exec," in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090153-Symantec-V.pdf.

on those VMs. Therefore, a failed Hyper-V server could potentially reduce productivity for several hours or even days, across multiple departments, while administrators struggle to recover the virtualized environment and individual VMs.

Experienced administrators looking to protect their Hyper-V environment readily understand the frustration and time involved in forcing backup technologies not optimized for virtualization to fit the requirements of virtualized environments. Organizations using tools that are not designed to be virtualization ready must contend with several limitations and time-consuming manual steps, such as installing backup agents in each VM, restoring entire VMs just to recover a single file, maintaining separate backups for system- and file-level recovery, and taking VMs offline during backup operations.

SYMANTEC AGENT FOR MICROSOFT VIRTUAL SERVERS

The Symantec Backup Exec 12.5 Agent for Microsoft Virtual Servers (AMVS) takes advantage of Microsoft Volume Shadow Copy Service (VSS), which is integrated into both Hyper-V and Microsoft Virtual Server 2005 Release 2 (R2) with Service Pack 1 (SP1). In addition to backing up offline VMs, the AMVS can use VSS to back up live running VMs without requiring downtime. Figure 1 outlines key advantages of the AMVS in Hyper-V environments.

The AMVS does not require a backup agent to reside in each VM; instead, it can be activated with a single license key for each Hyper-V host. In certain circumstances, such as running Microsoft Exchange or SQL Server in a VM, running the AMVS in the Hyper-V parent partition does not eliminate the advantages of running a specific application agent within that VM. In addition, an application being VSS compliant does not remove the requirements of regular log truncation, database maintenance, consistency checks, and other critical tasks to help maintain peak performance for those applications. Specific application agent backups can also provide a much more granular recovery mechanism than the AMVS.

During setup, Backup Exec guides administrators through the process of identifying the necessary groups and VMs to enable rapid recovery (see Figure 2). Backup Exec automatically selects the VMs and components such as .vhd, .vmc, .bin, .xml, and .vsv files for backup.

When recovering a VM, administrators can use the Backup Exec console to restore either the entire VM or specific .vhd files. They can also use built-in Backup Exec Granular Recovery Technology (GRT) to recover individual files and folders from within a .vhd file without running a separate backup. Administrators can restore backups to their original locations or to different locations—including alternate Hyper-V hosts or directory locations—and can use the original VM name or a different name.

SYMANTEC BACKUP EXEC AND MICROSOFT HYPER-V

To help ensure successful backup and recovery, administrators should adhere to several specific best practices when protecting Microsoft Hyper-V environments using Symantec Backup Exec 12.5:

- Ensure that the Hyper-V integration services are installed on the host server.
- Ensure that sufficient disk space exists on the media server for all .vhd files that will be copied directly to it for off-host backup.
- Because performance is typically dictated by the slowest component on the backup data path, consider using the

Integration with Microsoft VSS	When protecting an entire Microsoft Windows® OS–based VM, the AMVS can also help protect applications through VSS—enabling the entire VM and application to be recovered simultaneously.
Support for Microsoft Windows Server 2008 Server Core installations using Hyper-V	The AMVS is designed to be fully compatible with Windows Server 2008 Server Core installations using Hyper-V.
Integration with Symantec Backup Exec	The AMVS can integrate with Symantec Backup Exec to automatically discover Microsoft virtualized environments and non-virtualized environments, helping provide seamless, comprehensive data protection.
Agentless VM backup	Protecting VMs using the AMVS does not require installing an agent in each VM.
Simplified licensing	A single AMVS license allows organizations to protect an unlimited number of Microsoft Windows and Linux® OS–based VMs on a host server running Microsoft Windows Server 2008 Hyper-V or Microsoft Virtual Server 2005 R2 with SP1.
Built-in Symantec Backup Exec GRT	Administrators can use built-in GRT to restore individual files and folders from inside Windows-based VMs without restoring the entire VM.
Flexible recovery options	Administrators can use the AMVS to restore VMs to different locations, including specifying a different VM name and virtual network for the restored VM.
Disk-to-disk or disk-to-tape backup and recovery	Flexible backup and recovery capabilities support backup to disk for fast recovery or backup to tape for long-term storage.

Figure 1. *Key advantages of the Symantec Backup Exec 12.5 Agent for Microsoft Virtual Servers*

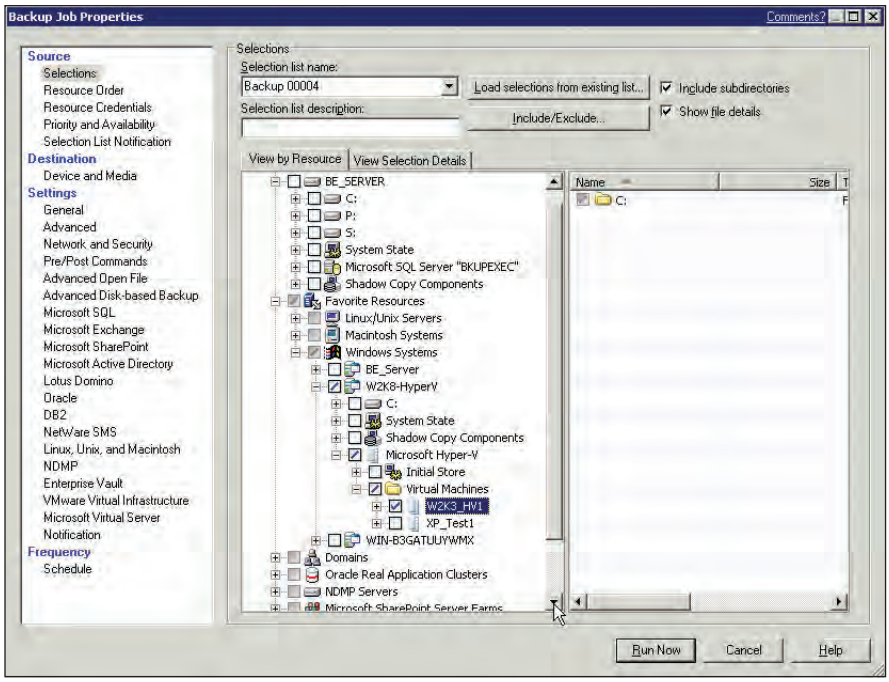

Figure 2. *Virtual machine discovery and selection in Symantec Backup Exec 12.5*

host parent partition as the Backup Exec media server (to avoid the need to move backup data from the host server to a separate media server), using disk-to-disk-to-tape backup (to enable rapid backup and recovery to disk devices while still using tape for long-term storage), and staggering VM backups (to avoid performance impact on the host server, where backup I/O occurs, and related impact on other VMs). In addition, avoid placing Backup Exec backup-to-disk folders on the same physical drive as the source .vhd files being protected.

- Keep in mind that VSS-enabled copy backups do not support application- or database-level full, incremental, or differential backups; do not truncate application log files; and do not provide granular application recovery. These VSS-enabled backups should not replace traditional application- or database-level backups through Backup Exec application or database agents.
- Schedule backups during times of relatively low I/O activity on the VM and avoid large numbers of simultaneous backups to help minimize problems related to snapshots.
- To help accelerate the completion of the snapshot process, ensure that the path from the data store to the snapshot mount point is as fast as possible by configuring the mount point over as many dedicated spindles as possible.

When recovering files in Hyper-V environments, administrators should also keep the following in mind:

- Hyper-V infrastructures do not include a default, built-in mechanism for recovering files directly into a VM. To perform this type of recovery, administrators can install the Backup Exec Remote Agent for Windows Servers inside the target VM (for restore purposes only). The Backup Exec media server can then communicate directly with the individual VMs to transfer recovered files. Alternatively, administrators can perform a client recovery to the host server, then access the restored files and transfer them to the VM directly from the host.
- Granular recovery from within a .vhd file on a tape-based backup is supported, but requires administrators to temporarily stage the .vhd file to a disk and then remove it; this type of recovery typically works best when restoring from a disk-based backup. When performing a granular recovery from tape, administrators should ensure sufficient disk space exists on the temporary staging location.

ROBUST PROTECTION FOR MICROSOFT HYPER-V

Symantec Backup Exec 12.5, part of the Dell PowerVault DL2000 – Powered by Symantec Backup Exec, is designed specifically to meet the data protection needs of virtualized environments, and taking advantage of features such as the Symantec Backup Exec AMVS can help provide simplified, robust backup and recovery for Microsoft Hyper-V virtualized environments. And because Backup Exec incorporates comprehensive support for not only Hyper-V, but also VMware® virtualization as well as non-virtualized environments, it offers a centralized, highly flexible way to help protect heterogeneous environments—combining single-pass and granular recovery advantages in a single simplified backup and recovery tool.

Ed Casmer is a technical strategist in the Partner Alliance Group at Symantec dedicated to the Dell Alliance.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Symantec Backup Exec:
www.backupexec.com
dell.symantec.com/backup-exec

Dell PowerVault DL2000:
DELL.COM/DL2000

SECURITY SIMPLIFIED: UNDERSTANDING THE BASICS OF DEPLOYING AND MANAGING SECURE ENDPOINTS

By David Schweighofer

With the ranks of mobile and remote workers expanding every day, cost-effective endpoint security is a growing business concern that enterprise IT organizations avoid at their peril. Dell advances enterprise-wide protection with a balanced approach that focuses equally on safeguarding data and preventing unauthorized access.

Related Categories:

Dell Latitude laptops

Security

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

In a world where digital information is the key to business and a significant proportion of workers are mobile, data loss is both a worst-case scenario and a daily occurrence. Regardless of whether it is accidental or the result of a security breach, data loss puts organizations at risk of losing market share, damaging shareholder confidence, and incurring compliance fines. Even a single day of downtime can cost thousands of dollars in lost worker productivity, IT administrator time, and data recovery costs.

For many organizations, the risk of business process failures threatens to compromise data security even more than the possibility of an external attack. Securing data on mobile systems can be complex for enterprise IT administrators, which increases the chance that security procedures will be followed improperly or overlooked. Security also can be burdensome and frustrating for end users. Complying with security policies typically requires multiple steps, settings, and actions that can become a barrier to implementation—leading to devastating business consequences.

For these reasons, organizations must protect their IT assets with endpoint security solutions that are simple to deploy and use. By providing a broad choice of world-class security offerings along with deployment and configuration services, Dell can help enterprises secure key data as well as reclaim wasted time and redirect it toward business growth.

UNDERSTANDING REQUIREMENTS FOR ENDPOINT SECURITY

To effectively meet today's endpoint security challenges, enterprise IT departments need a comprehensive approach that addresses three main requirements (see Figure 1):

- **Fast deployment:** Built-in security features such as always-on hard drive encryption help prepare systems for secure operations on delivery.
- **Enhanced protection:** Stringent security technologies for user authentication and access help safeguard precious data and ensure regulation compliance.
- **Smart prevention:** Multiple layers of security help protect valuable assets.

SIMPLIFYING SECURE DEPLOYMENTS

Dell helps simplify mobile security through platforms that are easy to use and manage, including a comprehensive suite of data security options that enable organizations to deploy secure systems direct from the factory. Of course, preventing unauthorized access is equally important. Best practices reach beyond system login procedures to involve facilities

management and human resources processes as well.

> "Dell helps simplify mobile security through platforms that are easy to use and manage."

Encryption provides a first line of defense for data protection

Dell™ systems support several types of encryption. File and folder encryption helps secure data with a digital signature, whether information is in transit over the network or at rest on an internal system—helping minimize the possibility that it will be altered or accessed by unauthorized users either internally or externally. It also helps to secure enterprise data by converting files and text into "cipher text" that can be decoded to its original form only with a valid password or encryption key.

Hardware encryption goes one step further to help ensure data security. Full-disk-encryption hard drives are accessible only through a password, and are designed to work without requiring user intervention or compromising system performance. When unlocked, however, the user interface is easy to use, simple, and transparent.

Dell laptops offer full-disk-encryption hard drives. Dell recently added a 7,200 rpm hard drive option to enhance performance while maintaining a high level of security. Dell also offers an encrypted solid-state drive (SSD) in a range of mobile devices, which combines the durability and reliability of SSDs with strong data protection capabilities.

Protecting sensitive data is a necessary first step to securing enterprise information—but it is not the whole story. By adding multiple access controls, Dell enables organizations to strengthen their authentication processes by implementing multi-factor authentication, which helps ensure that only authorized users get access to valuable information assets. Dell offers a range of authentication technologies, including biometric readers, smart card readers, and contactless smart cards to help support a consistent, cost-effective security strategy for simplifying processes.

Integrated fingerprint readers provide strong authentication

As a result of the strong authentication technology defined by Federal Information Processing Standard (FIPS) 201, a variety of U.S. government employees—including transportation workers and emergency first responders—are using FIPS 201 fingerprint identification technologies as the foundation for biometrically enabled identity credentialing programs.

The Dell Latitude™ E6500 laptop is designed to include an optional FIPS 201–certified fingerprint sensor and Personal Identity Verification (PIV)–compliant smart card reader. Latitude laptops feature the UPEK TouchChip (TCS1) silicon fingerprint sensor, which is designed to meet the demanding FBI fingerprint image quality requirements, fit easily into a tight space, and meet the low-power requirements for laptops.

Many Dell business laptops can be ordered with swipe fingerprint readers. Dell Latitude E-Family laptops and Dell Precision™ mobile workstations offer optional fingerprint readers.

Contactless smart card readers enhance security for mobile employees

Contactless smart cards can be used in various facilities management scenarios—for example, regulating access to company buildings and paying for food at the cafeteria. Dell Latitude E-Family Mainstream and Ultra-Portable laptops offer a radio-frequency identification (RFID)–based contactless smart card reader option. Administrators can enhance security for mobile employees by combining something employees own (a smart card) with something they know (their password) on an identified platform—the Trusted Platform Module (TPM) or Dell ControlVault™ platform.

Embedded smart card readers help make it easy and cost-effective to deploy smart cards for security. Latitude E-Family laptops avoid the requirement for peripheral readers, which occupy a USB port that may be needed for other peripheral devices. Latitude E-Family laptops are the first laptops to offer an embedded contactless smart card reader that supports multiple types of cards.

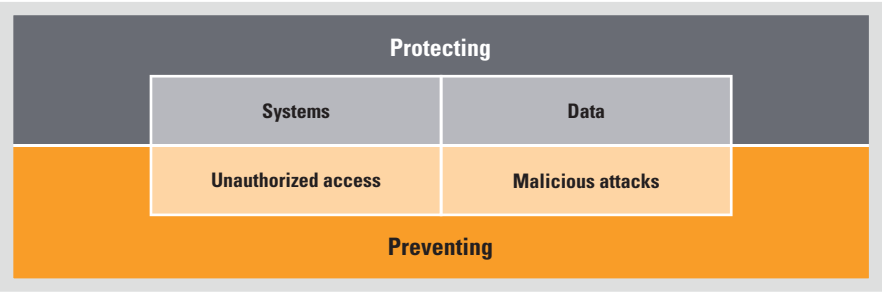

Figure 1. *A comprehensive approach to endpoint security helps protect data and prevent unauthorized access*

Trusted Platform Module software enables secure access to networks

For strong authentication, Dell business clients are often equipped with TPMs. This versatile chip—a microcontroller located on the motherboard of Dell laptops—helps to authenticate a system in an IT infrastructure and stores user credentials such as passwords, digital certificates, and cryptographic keys.

In addition to storage, the chip can securely generate or limit the use of keys for signing and verification as well as encryption and decryption. Capabilities of the TPM include remote attestation, which creates an unalterable summary of the hardware, boot, and host OS configuration to enable a third party to verify that software has not been tampered with. Other capabilities include sealing encrypted data so that it can be decrypted only in the same state.

Dell ControlVault technology dedicates processing power to security

For an additional level of security, Dell ControlVault technology provides a dedicated processor for authentication and certificate storage.[1] Although TPM and ControlVault technologies both store keys and have similar benefits, the ControlVault approach offers additional features designed to improve security. For example, the ControlVault chip can store and execute code using a secure processor—helping protect it from malware attack vectors that typically target RAM or hard drives.

While TPM uses a 160-bit password, ControlVault also supports use of personal authentication methods (such as fingerprint readers, smart cards, and contactless smart cards) to access credentials. The ControlVault chip is designed to store all credential types, allowing a single point of migration and supporting a broad variety of cryptographic algorithms (including Suite B and native error checking and correction). In addition, it supports standard and contactless smart cards, fingerprint readers, and RSA SecurID tokens.

ControlVault helps to protect secure operations by isolating them from the OS environment and memory. Many applications execute their secure operations on the host x86 processor, which exposes it to sniffing of interim values and modification of the final result. Even if the encryption key itself is concealed with another key, the encryption key is still on the hard drive and out in the open. And even if the key is hidden among other data, hackers have programs that can search the hard drive and quickly locate the key.

By sealing off code execution, ControlVault can help protect against these threats. All processing and storage of critical data takes place on a processing and memory chip, which creates a protective boundary. Access to the keys is strictly controlled by an authorization scheme that is designed to prevent any application from accessing the keys without satisfying the authentication requirements set up by the owner or IT manager of a particular ControlVault-protected boundary. In addition, a small memory footprint helps ensure that ControlVault incurs minimal impact on overall system performance.

ControlVault also serves to control access to reference templates. To verify authentication, a reference template—which is created and stored at time of enrollment—must be accessed. Applications usually store this template on the hard drive, thus exposing it to modification, extraction, and copying. ControlVault helps eliminate these threats by allowing applications to store templates inside the ControlVault-protected boundary. Template access is then controlled by an authorization scheme designed to prevent any application from accessing the keys without satisfying the authentication requirements set up by the owner or IT manager of a particular ControlVault-protected boundary.

Dell ControlPoint helps simplify endpoint security management

Authentication is critical to endpoint security—but IT departments must enforce strict rules to strengthen the authentication, which can create complexity for end users. To help simplify using and managing security features, Dell ControlPoint software includes a Security Manager module (see Figure 2).

The ControlPoint platform gathers hardware and security settings within a single intuitive user interface, helping avoid the need to search through multiple control panels for a specific setting. By providing a standardized user interface to access a broad selection of security capabilities, ControlPoint software extends the available authentication offerings and helps simplify management of these options. In addition, the software helps to facilitate implementation and management of multi-factor authentication across multiple devices, including biometrics, smart cards, and contactless smart cards. This combination, when used with preboot authentication and hard drive encryption, provides enterprises with an extra layer of security.

RSA SecurID certification offers strong, cost-effective authentication

The RSA SecurID algorithm is embedded within the Dell ControlVault hardened firmware chip for storage and processing of credentials. RSA SecurID software token seeds are stored within ControlVault, outside the usual attack vector of malicious applications. A one-time RSA password is generated within

> **"ControlVault helps to protect secure operations by isolating them from the OS environment and memory."**

[1] Dell Latitude models E4200, E4300, E6400, E6400 ATG, E6400 XFR, and E6500 as well as Dell Precision models M2400, M4400, and M6400 offer Dell ControlVault and Dell ControlPoint technologies. Other Dell business laptops offer limited versions of these solutions.

Figure 2. *The Dell ControlPoint Security Manager module offers single-view management of security settings, features, and authentication*

the ControlVault chip. Mobile users can conveniently launch the software token from their laptops through Dell ControlPoint Security Manager.

By embedding the RSA SecurID software token within the Dell Latitude ControlVault firmware, the laptop is designed to offer the security of a hardware token combined with the cost-effectiveness and convenience of a software token. This approach helps avoid the need for administrators to replace lost tokens while affording mobile users the convenience of a consolidated device. By linking the two-factor authentication method directly to the laptop, desktop, or workstation, organizations help to ensure that employees are accessing enterprise information only from company computers.

The RSA SecurID software token can be easily licensed and provisioned by organizations with a deployment of Latitude E-Family laptops. Out-of-the-box interoperability with the RSA SecurID Token 4.0 for Windows Desktops allows users or IT administrators simply to install the RSA SecurID Desktop 4.0 application, which is designed to automatically register ControlVault as a token storage device.

ControlPoint Security Manager allows IT professionals to easily access and manage not only RSA SecurID tokens, but also user identification, fingerprint readers, and smart card security technology. The software enables limited access to network resources with a two-factor authenticator linked directly to an enterprise system—meaning that mobile users can use their own mobile device as an authenticator, rather than require an additional dedicated device for two-factor authentication. Integrated two-factor authentication can control access to hundreds of applications, including many from leading virtual private network vendors.

MEETING THE CHALLENGES OF ENDPOINT SECURITY

As the ranks of mobile employees grow, organizations must give resolute focus to protecting IT assets with endpoint security solutions. By understanding how to deploy and use a comprehensive range of endpoint security options, IT organizations can advance enterprise-wide security while enhancing productivity.

Dell endpoint security options start with encryption as a first line of defense for protecting data. In addition, Dell offers FIPS 201–certified fingerprint sensors, PIV-compliant smart card readers, and embedded smart card readers. An embedded contactless smart card in Dell Latitude E-Family laptops supports multiple types of cards. TPM chips offer a security option that helps provide strong authentication, and Dell ControlVault technology offers a dedicated processor for authentication and certificate storage. The RSA SecurID algorithm is embedded within the ControlVault firmware to store and process credentials, and Dell ControlPoint software includes a Security Manager module to help simplify endpoint security management.

Designed to simplify data protection and prevent unauthorized access, the Dell endpoint security options discussed in this article help IT organizations deploy and manage security to facilitate simple and secure data access throughout the enterprise.

David Schweighofer is the worldwide outbound marketing manager for Dell Latitude laptops on the Product Group Marketing team. He has a degree in Marketing and Organization as well as an M.B.A. from the HEC Lausanne business school.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell security solutions:
DELL.COM/Security

Data Protection in a Mobile World:
dell.awakit-webcasts.com/index/webcastsList/series_id/4

Data centers move as fast as business when physical and virtual workloads work as one.

With PlateSpin® workload management solutions from Novell®, what business wants right now, your data center can deliver right now. A single suite of products centrally monitors, manages and optimizes physical and virtual servers for you, automatically shifting workloads to the right server at the right time. Improve server utilization, reduce costs and make your data center more agile so you can respond to business demands in real time. Let us make IT work as one for you.

www.novell.com/workload

Novell®
Making IT Work As One™

HOW TO OPTIMIZE VIRTUAL DESKTOP INFRASTRUCTURE DEPLOYMENTS WITH F5 BIG-IP

By Alan Murphy

Centralized management helps improve efficiency, but delivering desktop images over a network can incur cost, security, and latency trade-offs. Application delivery platforms from F5 Networks such as F5® BIG-IP® devices connected to virtualized Dell™ PowerEdge™ servers help organizations deploy a virtual desktop infrastructure designed to simplify management and reduce costs.

Related Categories:

F5 Networks

Flexible computing

Microsoft

Virtualization

VMware

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

Deploying and managing desktop computers, especially remote desktop computers, can be costly and time-consuming. As a result, many organizations are transitioning to a virtual desktop infrastructure for their client environments to help centralize management and reduce costs. In a virtual desktop infrastructure, desktop images are offloaded to data center servers and distributed to clients over the network using these virtual environments, enabling centralized management of desktop images without compromising the end-user experience. By avoiding the burden of deploying individual desktop images, a virtual desktop infrastructure can help dramatically increase manageability while contributing to reduced total cost of ownership across the enterprise desktop environment.

But first, IT decision makers must resolve concerns about reduced performance and availability, which are common concerns of virtual desktop infrastructure deployments that can frustrate end users and lead to diminished productivity. In particular, the intensified network and server demands of a virtual desktop infrastructure may increase application latency and reduce application availability, especially over wide area network (WAN) and other remote connections. As a result, concerns about application performance, especially in large-scale environments, may prevent some organizations from achieving the potential benefits of virtual desktop infrastructure technology.

F5 Networks helps organizations prepare their IT environments for a successful virtual desktop infrastructure deployment by offering a range of application delivery solutions. When combined with scalable virtual machine (VM) host platforms such as Dell PowerEdge servers, this application delivery system helps organizations dramatically enhance the performance and availability of their enterprise virtual desktop infrastructure deployments, enabling them to take advantage of the benefits of reduced costs, simplified management, and enhanced security while still offering a robust and familiar end-user experience.

ENHANCING DESKTOP SECURITY AND MANAGEABILITY

Traditional enterprise desktop deployments can be costly to manage and difficult to secure. For example, in a typical non-virtualized environment, desktop computers must be managed individually. As a result, tasks such as image deployment and ongoing administration of application updates and security patches can be costly and cumbersome. Because IT

organizations often do not have consistent access to or control over end-user desktops, maintaining security and ensuring regulatory compliance can be difficult as well. For example, in a traditional nonvirtualized desktop environment, failure to prevent users from adding unauthorized applications or other content to their desktops can result in compromised enterprise security and violations of regulatory legislation such as the Health Insurance Portability and Accountability Act (HIPAA).

By enabling centralized control of desktop images, a virtual desktop infrastructure is designed to eliminate many of the problems associated with desktop management without compromising a robust, familiar user experience. For example, because desktop images are stored centrally rather than on individual desktops, administrators can easily deploy images, updates, and patches without needing to manually service individual end-user devices. Centralized management also allows IT organizations to easily identify and remove unauthorized applications or content, which helps simplify enterprise security and regulatory compliance tasks.

Desktop virtualization technology can be deployed in several ways. Typically, the entire desktop image is offloaded to centralized servers and virtualized. In this case, application and desktop processing happens on the servers with specialized platform software, such as VMware® View, delivered over the network, and the proprietary client application performs only graphics processing and other user interface functions. Alternatively, a virtual desktop infrastructure may be deployed over the Web. In this case, centrally stored desktops are accessed through a Web browser. Another option is to virtualize individual applications rather than the entire desktop, which allows the desktop image to remain on the end-user device while certain individual applications are stored on centralized servers and run remotely.

OPTIMIZING VIRTUAL DESKTOP INFRASTRUCTURE DEPLOYMENTS

Although the move to a virtual desktop infrastructure can deliver tremendous benefits, IT organizations must address the potential risks of reduced application performance and availability that can frustrate end users and compromise worker productivity. Because a virtual desktop infrastructure can send large amounts of data over WAN and other remote connections, limitations in network bandwidth and performance are key considerations for organizations considering this type of infrastructure.

Network-induced problems are highly pronounced when data is transferred outside an organization firewall, over congested links, or through access from an autonomous network. Also, load on VM servers running virtualized desktops can negatively impact the performance and availability of those virtual desktop images because processing desktop graphical user interfaces is I/O intensive. Overall, concern about potential performance degradation, especially in large-scale environments, is one of the main reasons why IT organizations may hesitate to adopt a virtual desktop infrastructure.

To help enterprises meet the additional network and server requirements imposed by a virtual desktop infrastructure, F5 Networks offers easy-to-deploy solutions designed to mitigate limitations of WAN connectivity and deliver scalability and performance for VM servers. F5 BIG-IP Local Traffic Manager™ (LTM) systems provide load balancing and traffic management functions that enhance the performance and availability of VMs when using a virtual desktop platform such as VMware View (see Figure 1). Key capabilities of BIG-IP LTM include the following:

- **Load balancing:** BIG-IP LTM is designed to load balance traffic and workloads across VMs to help maximize availability and performance. For example, in a VMware View virtual desktop infrastructure deployment, BIG-IP LTM can monitor the health of VMware View connection servers and balance the workload based on the status of individual connections. BIG-IP LTM can also make load-balancing decisions based on response times.
- **Client connection persistence:** BIG-IP LTM manages client connection persistence based on unique sessions, which can be a more robust approach than persistence based on the typical source IP address. For example, in a VMware View virtual desktop infrastructure deployment, if users are accessing virtual desktops through a Web proxy, BIG-IP LTM can distribute and maintain the connections among

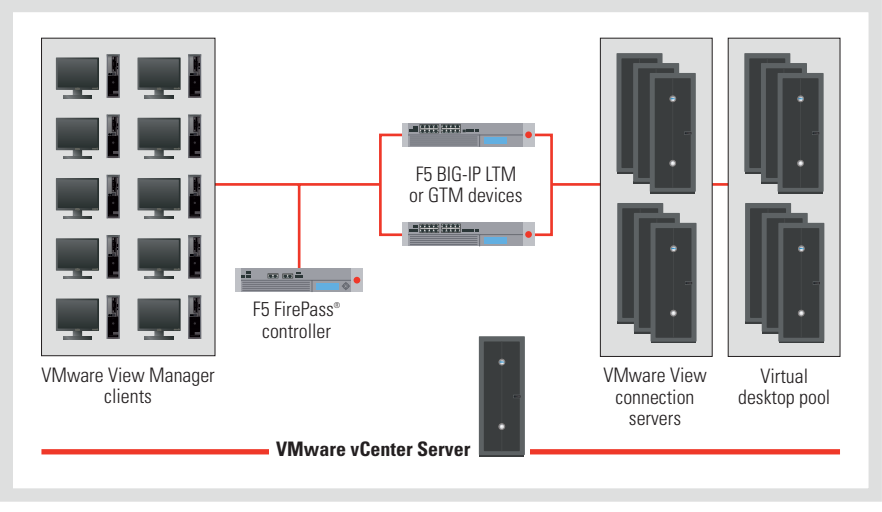

Figure 1. *Traffic management services and load balancing in an F5 and VMware View environment*

> "Many organizations are transitioning to a virtual desktop infrastructure for their client environments to help centralize management and reduce costs."

the VMware View connection servers rather than sending the following connections from that proxy server to a single connection or security server.

- **Secure Sockets Layer (SSL) processing:** SSL processing can place a large burden on VM servers, especially during logon procedures. BIG-IP LTM can offload SSL processing from VM servers, helping free the servers to optimize delivery of virtual desktop infrastructure functionality.
- **Compression:** BIG-IP LTM can provide asymmetric compression to the client for traffic such as JavaScript and HTML, and can offer symmetric compression in architectures with symmetric local traffic managers, such as a BIG-IP LTM device in the data center and a BIG-IP LTM device in the remote office.
- **Data deduplication:** BIG-IP LTM with the BIG-IP WebAccelerator™ module can perform data deduplication in asymmetric deployment architectures, helping to decrease the overall amount of data that needs to traverse the WAN and enhance user experience.
- **Rapid deployment template:** BIG-IP LTM includes an application template specific to VMware View, which provides administrators with a wizard-like interface to rapidly configure and deploy BIG-IP LTM with VMware View. This template helps simplify the deployment process, minimize the risk of configuration errors, and reduce the administrative costs of deployment.
- **User-level service-level agreement (SLA):** Different categories or groups of users may have different service levels in terms of importance and performance; for example, users in an engineering group may have a higher priority for performance than call center employees. With BIG-IP LTM, user connections from the engineering group can be routed to a dedicated pool of desktop VMs, whereas call center users would be routed to general, lower-performing desktop VMs. This approach helps ensure organizations can meet different SLAs for different groups of users.

In addition, F5 Networks offers BIG-IP Global Traffic Manager™ (GTM) systems, which enhance application performance and availability through intelligent routing between data centers. BIG-IP GTM is designed to route incoming virtual desktop infrastructure traffic to the highest-performing data center depending on the location, link conditions, and data center conditions for end users. Intelligent routing helps ensure that clients can access their virtual desktops irrespective of the availability of any single data center.

F5 application delivery devices are designed to work with a range of virtual desktop infrastructure architectures, including VMware View and Microsoft® Virtual Desktop Infrastructure (VDI) technologies, along with application virtualization solutions such as VMware ThinApp and Microsoft Application Virtualization (App-V). F5 platforms are also designed to deliver excellent application performance at all levels of scale, from small deployments to large-scale deployments, and to work seamlessly with VM server infrastructures based on Dell PowerEdge servers. Additionally, F5 Networks offers application templates that help streamline and simplify deployment for a range of virtual desktop infrastructures, including those based on VMware View and VMware ThinApp.

PREPARING THE IT ENVIRONMENT FOR A VIRTUAL DESKTOP INFRASTRUCTURE

Virtual desktop infrastructures have the potential to deliver tremendous benefits to enterprise IT organizations. By centralizing control of enterprise desktops, a virtual desktop infrastructure helps organizations dramatically increase manageability and security and decrease total cost of ownership. Application delivery solutions from F5 Networks help organizations prepare their network and storage infrastructures for the increased load introduced by a virtual desktop infrastructure, helping ensure that end users experience the performance and availability they have come to expect from their desktops.

Alan Murphy is the technical marketing manager for management and virtualization solutions at F5 Networks. Alan provides strategic counsel on the F5 product suite to customers, prospects, and analysts through evaluation and analysis of current technical market trends and the competitive virtualization and security landscape. He has over 15 years of experience in IT, including UNIX® systems administration, systems and data security, and incident response, and holds GIAC Certified Incident Handler (GCIH) and Red Hat Certified Engineer® (RHCE®) certifications.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

F5 Networks:
www.f5.com

F5 virtualization solutions:
www.f5.com/solutions/virtualization

VMware View Reference Architecture Kit:
www.vmware.com/resources/techresources/1084

Get real about virtualization

Virtualizing servers is a great way to cut costs and improve efficiency. But you can't just deploy the technology and call it good. Is your infrastructure ready for the realities of virtualization?

www.f5.com/dell

f5. IT agility. Your way.

AUTOMATING ASSET MANAGEMENT IN DELL ENVIRONMENTS WITH KBOX FROM KACE

By Lubos Parobek

The KBOX™ asset management system from KACE can track IT assets throughout their life cycles without the cost and complexity of traditional software solutions. The KBOX appliance-based approach helps organizations to reduce costs and increase efficiencies in managing leases, warranties, and license agreements through a quick-to-deploy and easy-to-use appliance.

Related Categories:

Asset management

KACE

Systems management

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

A consolidated and centralized approach to asset management is crucial for effectively managing the IT environment. Understanding what hardware and software are installed across an organization is the most basic step to effective asset management, but often one of the most difficult to complete accurately without a major investment in time and resources. The additional challenges of managing warranties and contracts, software compliance, secure environments, and efficient service desks overwhelm many organizations.

An asset management system that enables a comprehensive view into all resources and their assigned dispositions provides an initial baseline for implementing best practices. Keeping inventory current then enables organizations to realize a wide range of operational and business efficiencies.

In turn, an asset management system that provides a detailed database of hardware and software inventory can help identify usage and potential efficiency improvements. In addition to desktops and servers, asset management should extend to printers, network equipment, and other high-value assets along with pertinent financial and contractual data such as vendor names and service and maintenance contract information, helping maintain a comprehensive view of asset information in a single data repository.

BEST PRACTICES FOR ASSET ASSIGNMENT AND TRACKING

KBOX from KACE uses an appliance-based architecture to help simplify ongoing management of distributed systems, enabling IT organizations to increase efficiencies across a number of disciplines including asset management, configuration management, service desks, and systems deployment. The KBOX Asset Management Module focuses on automating non-computer and computer inventory with asset management processes from deployment to retirement, including asset data auditing, tracking, compliance, and reconciliation. KBOX is designed to identify IT assets and keep them updated daily across an organization, utilizing a light agent for discovery and an efficient messaging protocol to communicate with a centralized configuration database. This approach enables frequent inventory updates without interrupting end users, helping to ensure accuracy and informed decision making.

Administrators can define software asset groups and perform daily tracking of associated licenses. Tracking for software compliance can be automated by defining filters to include current and future versions of applications across a vendor's product line. Administrators can automatically monitor the number of installed licenses, even as new versions of an application become available (see Figure 1)—essential

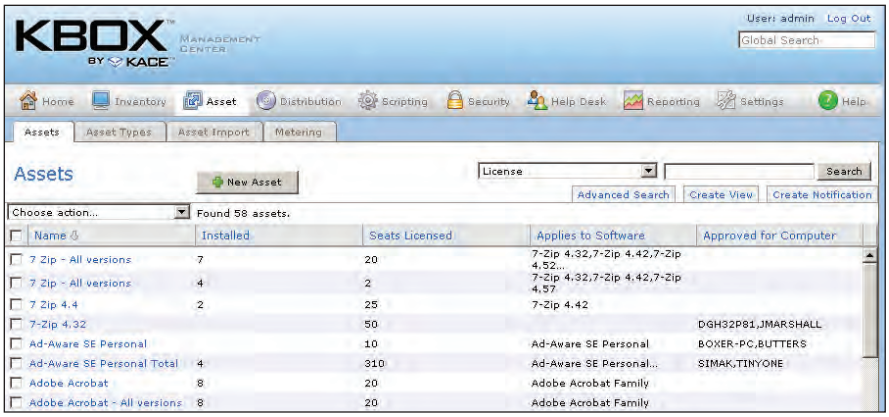

Figure 1. *Tracking software license compliance in KBOX from KACE*

features for managing software license agreements and negotiating ongoing volume purchases.

KBOX is designed to integrate tightly with service tag information for Dell™ hardware to help ease the tracking of warranty information and original configuration details (see the "PCC cuts costs and enhances Dell management with KACE" sidebar in this article). Administrators can easily create new asset types and match Dell equipment with associated service contracts and software licenses. E-mail notifications and alerts can be configured and tied to contract expirations and renewals, allowing the system to be tailored to help meet the asset management and compliance needs of any organization.

KBOX asset management also enhances security by quickly identifying missing hardware or unapproved software or systems found on the network. Providing a comprehensive audit trail of asset changes, including software and hardware, helps increase visibility into unauthorized usage or unexpected changes, thereby advancing security. In addition, the integrated functionality of the KBOX Help Desk Module enhances the efficiency of service desk personnel, providing them with detailed system configuration information to help identify and resolve problems quickly.

BREAKTHROUGH EFFICIENCY TO HELP SIMPLIFY IT

Under challenging economic conditions, effective asset management can help organizations make informed decisions, deliver fast problem resolution, manage service contracts, enforce secure configuration policies, and implement proactive license compliance. The comprehensive KBOX asset management system helps eliminate the complexity and cost of traditional enterprise software solutions, enabling organizations to automate asset management operations for breakthrough efficiency. As a result, enhanced agility helps enterprises accelerate IT response to meet increasing user expectations and ever-changing business requirements.

Lubos Parobek is vice president of product marketing and management at KACE, where he has been intimately involved with the challenges IT organizations face when deploying and managing servers, desktops, and laptops. Lubos previously served more than 10 years in hardware and software product management and marketing, holding a variety of positions at 3Com, AvantGo, and Sybase.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

KACE:
www.kace.com

PCC CUTS COSTS AND ENHANCES DELL MANAGEMENT WITH KACE

Portland Community College (PCC), one of the nation's largest such institutions, supports nearly 5,000 desktop and laptop systems across multiple campuses. Roughly 95 percent are PCs, and of those, the majority are Dell computers. PCC has deployed the KBOX systems management appliance to keep track of Dell system assets, including hardware locations, installed software, where systems are in their service cycles, utilization rates, and more. By leveraging native Dell service tag links from KBOX, PCC can view maintenance history, tie the computers back to purchase order information, and so on—all from the KBOX management console. When combined with additional asset data from Dell, provided through spreadsheets imported into KBOX, PCC administrators can view Dell systems purchased in the last seven to eight years, where they were deployed, when they were replaced, and a full life cycle history.

This comprehensive knowledge helps PCC manage its large desktop and server environment, and has become an essential tool for anticipating the computing needs of different departments. Utilization rates gathered by KBOX help PCC administrators determine, by location or user, which systems require replacement or upgrade. This information also helps the PCC team to get a head start on the lease renewal process and to determine organizational needs when migrating systems off lease. The software inventory process helps further streamline operations—enabling PCC to roll out new software upgrades, including updated utilities and virus protection.

"The KBOX allows us to proactively manage our costs and know everything we need to know about our Dell systems," says Michael Heurer, technology solution services customer support manager at PCC. "Not only are we able to give better support to our users, but with this information, we are able to look at the best opportunities to make purchases and plan for our future."

AIRLINE SOARS WITH BLADES

Jakarta-based airline PT Garuda Indonesia reduces energy use by 40 percent and space requirements by 50 percent with Dell™ PowerEdge™ blade servers.

CHALLENGE

PT Garuda Indonesia wanted to boost the efficiency of its business operations by developing its IT infrastructure further and running key front-end applications on high-performance servers.

SOLUTION

The airline deployed Dell PowerEdge M600 blade servers paired with Dell PowerEdge M1000e modular blade enclosures. The Dell blade servers have helped the company reduce energy consumption, optimize data center space, and provide a reliable, scalable IT infrastructure for its applications.

BENEFITS

- Dell blade servers equipped with quad-core Intel® Xeon® processors helped cut energy use by 40 percent and reduce the data center footprint by 50 percent.
- Dell consultants helped optimize deployment of the server solution, which was completed in three months.
- Dell ProSupport provides custom IT support and services to help meet Garuda Indonesia's requirements.

Related Categories:
Blade servers, case study, Dell OpenManage, Dell OptiPlex desktops, Dell PowerEdge blade servers, Dell ProSupport Services, Dell/EMC storage, PT Garuda Indonesia

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

Jakarta-based PT Garuda Indonesia is Indonesia's flag carrier. Founded in 1949, the national airline flies to 21 domestic destinations and 19 cities in Asia, the southwest Pacific, and the Middle East, while transporting 12 million passengers annually.

Recently, Garuda Indonesia has been challenged by competing international airlines and budget carriers as well as changing economic conditions. In 2003, the Association of Southeast Asian Nations (ASEAN) endorsed Open Sky, a policy targeted for implementation in 2015. Open Sky is designed to promote airline industry competition and to give all airlines from ASEAN the scope to compete on intra-ASEAN routes, which have traditionally been monopolized by local airlines. These factors have placed pressure on Garuda Indonesia to maximize its efficiency and effectiveness; as a result, the company increasingly relies on its IT infrastructure.

IMPROVING APPLICATION PERFORMANCE

Garuda Indonesia runs many applications that require secure, reliable, high-performance servers, such as an enterprise resource planning (ERP) system, business accounting applications, an automated reservation system, the Garuda portal, and Garuda Online Booking—an Internet booking and payment application that integrates other applications and the Internet booking and payment gateway.

Improving application performance on Garuda Indonesia's automated reservation system for travel agents and the airline's international offices is critical. "Each passenger must be served in less than three seconds, regardless of whether the system is accessed from Jakarta, Surabaya, Tokyo, or Riyadh. We need a solid IT infrastructure in place to run our reservation system," says M. Ismed Arifin, vice president of information system solutions for Garuda Indonesia.

While the airline's back-end systems still run on mainframe systems, Garuda Indonesia is focusing on middleware and front-end applications that will run on Intel platform–based servers. The company has a mixed-application environment that includes SAP® Business Suite and a number of Microsoft® products. To boost the effectiveness of the airline's business operations, new servers were needed to run key front-end applications. Garuda Indonesia held an open tender for the new server solution, inviting top-tier vendors to compete in an online auction. From

> **"As business expands, we will have to scale our IT infrastructure to meet new demands, and we know Dell will be there for us."**
>
> —M. Ismed Arifin
> Vice president of information system solutions
> for Garuda Indonesia
> March 2009

this process, Dell was selected as the supplier, besting the competition because of the performance and features of its proposed servers, outstanding overall value, and supporting services.

Garuda Indonesia was pleased to continue its relationship with Dell, having already used Dell hardware for over 15 years. When the company first implemented SAP Business Suite, it had 600 Dell OptiPlex™ desktop clients; today, that number has grown to over 1,000. Dell solution architects assessed Garuda Indonesia's requirements and recommended a suitable hardware solution. The airline deployed a series of Dell PowerEdge M600 blade servers with the latest quad-core Intel Xeon processors, housed in PowerEdge M1000e modular blade enclosures. The blades connect to a set of Dell/EMC CX3-20c storage area network arrays, enabling Garuda Indonesia to replace its existing tower servers and host its business applications. This Dell solution was deployed within three months.

DEPLOYING EFFICIENT BLADE SERVERS

The Dell blade servers have enabled Garuda Indonesia to substantially reduce its energy use. The PowerEdge M1000e enclosure takes advantage of thermal design efficiencies that optimize airflow to help cool the chassis and enable higher performance in a lower power envelope. The PowerEdge M600 blades in the enclosure are based on the Intel Xeon processor 5400 series and the Intel Core™ microarchitecture, helping Garuda Indonesia maximize data center performance and density and improve energy performance. Second-generation Intel quad-core technology helps simplify the process of condensing the airline's applications for efficient operation at a reduced total cost of ownership. "We found that our blade solution reduced energy consumption by as much as 40 percent," says Ismed.

Although many servers are running in Garuda Indonesia's data center, the company has entrusted the Dell blades with running the most critical front-end applications, including Microsoft Exchange Server 2007. Used by 2,000 employees across Garuda Indonesia's global operations, the Exchange platform plays a key role in collaboration and communication. Garuda Indonesia is also currently developing a Business Intelligence application from SAP that will run on blade servers and provide airline management with the latest data on the company's performance.

The Dell blades come with tools such as the Dell OpenManage™ suite that help simplify the deployment and management of the airline's server solution, including alerts enabling the company to act quickly to resolve problems. Deploying the blades into Garuda Indonesia's data center was simple, with 16 blades housed in the PowerEdge M1000e enclosure. The PowerEdge M1000e is designed to support future blade technologies regardless of processor or chipset architecture, helping to secure Garuda Indonesia's investment. "The compact and modular design of the blades has allowed us to get the most out of our data center space, reducing our data center footprint by 50 percent," says Ismed.

SCALING TO MEET EXPANDING BUSINESS NEEDS

Garuda Indonesia has received positive feedback from users of applications that run on the blades and from IT staff benefiting from improved ease of management. The company is confident in Dell's enterprise capability and wants to explore ways that Dell can help keep the airline at the forefront of technology, such as virtualization. Plus, Garuda Indonesia is confident of the committed support services from Dell ProSupport. "Based on our experience with Dell, we are assured that Dell will be committed to invest their time to provide all the consulting and support services we might need," says Ismed.

Looking ahead, Garuda Indonesia wants to further enhance its operations with new customer-facing applications such as the Garuda Frequent Flyer loyalty program, and with the delivery of new fleets for long-haul flights to Europe. "As business expands, we will have to scale our IT infrastructure to meet new demands, and we know Dell will be there for us," concludes Ismed.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades

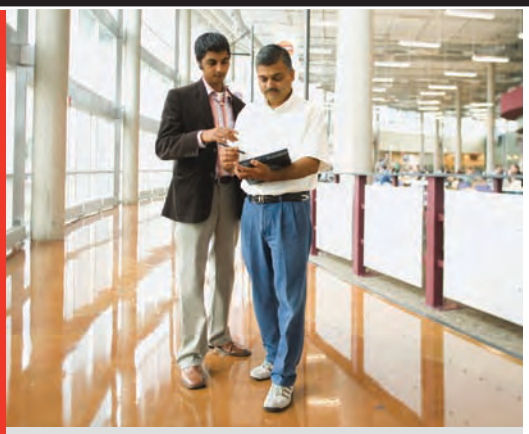

GETTING THE MESSAGE ACROSS

Taylor's University College in Malaysia unifies a decentralized messaging environment with Dell™ PowerEdge™ servers and a Dell/EMC storage array—simplifying e-mail for end users, improving efficiency by 30 percent, and reducing support calls by 20 percent.

CHALLENGE

The messaging infrastructure at Taylor's University College impeded resource sharing and staff mobility, offering little support to accommodate an expanding user base.

SOLUTION

Dell Global Infrastructure Consulting Services helped the college deploy clustered Dell PowerEdge 2950 servers, a Dell/EMC storage area network (SAN) array, and Microsoft® Exchange Server 2003 to centralize and manage its messaging infrastructure.

BENEFITS

- Consolidating to a single Microsoft Active Directory® domain helps improve efficiency by 30 percent.
- Consolidation enables increased performance and throughput.
- Clustered servers provide uninterrupted data flow.
- Centralized SAN enhances e-mail recovery.

Related Categories:
Case study, Dell Infrastructure Consulting Services (ICS), Dell PowerEdge servers, Dell/EMC storage, Microsoft Exchange Server 2003, Taylor's University College

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

Taylor's University College is one of Malaysia's leading educational institutions and the world's largest center for preuniversity studies. Established in 1969 and spread across five campuses, Taylor's offers a comprehensive range of tertiary education courses—from preuniversity classes to diploma and degree programs—and has successfully prepared more than 50,000 graduates for higher education. For this large, dispersed community of staff and students, e-mail messaging is vital to the smooth operation of services on all campuses, but the existing infrastructure had proven to be incapable of supporting Taylor's intended expansion.

UNIFYING THE MESSAGING PLATFORM

The IBM® Lotus® Notes platform, the college's primary messaging software, did not integrate well with Microsoft Active Directory domain servers, making it cumbersome to create new e-mail accounts and manage users. "In the past, when we added one user to the Active Directory domain, we had to create another account for the user separately in Lotus Notes. This is why we had to keep individual servers, one at each campus location, to cater to localized user needs," says Bernard Tee, senior manager in the ICT Development Department at Taylor's. In this messaging environment, end users were often disconnected from their e-mail accounts once they left their specific campuses—so if staff members travelled to another campus, they could not check their e-mail or access e-mail archives that were stored on a server on another campus.

The ICT Development Department recognized the constraints of the current e-mail platform and decided that an integrated solution was required to overcome these limitations. A new campus is scheduled to open this year, so Taylor's sought a platform that could scale easily to handle an increase in staff and students, centralize e-mail storage among campuses, enable simplified user management, and operate with minimal support. However, accommodating the expansion with the existing solution would require costly additional Lotus Notes licenses.

Taylor's wanted to consolidate four separate Active Directory domains into a single domain and migrate its existing Lotus Notes e-mail system to Microsoft Exchange Server 2003—a plan that would eventually serve 2,000 users, with enough space to support up to 5,000 users as enrollments increased. Exchange can integrate seamlessly

with Active Directory to facilitate fast and simple activation of accounts, empower users to self-manage basic e-mail requirements, and offer Taylor's numerous options to enhance its system with thousands of integrated applications.

The college asked various system integrators and vendors, including Dell, to submit proposals. "We explained our problems and asked the vendors for potential solutions. We then short-listed several vendors to submit their final proposals," says Tee. At the close of the tender, the project went to Dell. Having worked with Dell since 1997 as a preferred IT hardware supplier, Taylor's was assured that the partnership would be a good fit. "We had extensive discussions with Dell over our concerns and they offered a comprehensive proof-of-concept solution, which addressed all our needs," says Benjamin Lye, operations manager in the ICT Development Department at Taylor's.

IMPROVING EFFICIENCY AND SIMPLIFYING MANAGEMENT

To help meet the needs of the college's growing user base and future expansion, Dell recommended a clustered configuration based on Dell PowerEdge 2950 servers with Intel® Xeon® processors, a Dell/EMC CX3-20 storage area network (SAN) array, and Dell Fibre Channel switches. Some servers were clustered to run Exchange at the college's Subang Jaya campus, while the rest were used to consolidate the Active Directory user database into a single domain. Dell Global Infrastructure Consulting Services completed the solution in six months, from design to deployment.

The migration's impact was felt almost immediately. "The majority of the user base wasted little time getting up to speed and familiarizing themselves with the new messaging platform," says Lye. With a single Active Directory domain and centralized storage using the SAN array, staff could now conveniently access and retrieve archived e-mails wherever they were. "Basically, the new platform allowed our users to move freely between campuses and share resources easily without any fear of being disconnected from vital communication."

> **"We had extensive discussions with Dell over our concerns and they offered a comprehensive proof-of-concept solution, which addressed all our needs."**
>
> —Benjamin Lye
> Operations manager in the ICT Development Department at Taylor's University College
> February 2009

The cluster topology was designed to provide data integrity by helping ensure redundancy if the Exchange database failed, and to offer headroom to balance a growing Active Directory environment. "By unifying the Active Directories under a single domain, our efficiency improved by about 30 percent," notes Lye. "Clustering ensures we have the throughput and scale at the data center to accommodate a large Active Directory and simplify user management without duplicating the work scope across multiple servers and domains."

REDUCING SUPPORT REQUIREMENTS

Integrating Microsoft Exchange and consolidating the Active Directory environment into a single domain on Dell servers has simplified messaging operations throughout Taylor's. "We've received very positive feedback from users since the implementation," says Lye. Tee adds, "Compared to the previous solution, we receive 20 percent less support queries, and troubleshooting is easier because more support staff are familiar with the Exchange Server. Overall, there's a shorter learning curve using the Exchange Server."

Dell also trained support staff and some users. "Both the software and hardware, such as the Fibre Channel switch and SAN array, were completely new technologies to our people. If not for Dell's expert recommendations and knowledge, Taylor's would not have had the opportunity to test these technologies. Due to Dell consultants' expertise, they were able to train us to use, manage, and administer the solution," notes Lye.

Taylor's cites the excellent working relationship with Dell as key to the project's success. Tee remarks, "When we were looking for a vendor, the service/support relationship took precedence over the solution design and capabilities." He adds, "The Dell team provided excellent support throughout the project. They understood our requirements and delivered what we needed, and in the process, fostered a close working relationship that we've come to rely on."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell solutions for Microsoft Exchange:
DELL.COM/Exchange

KEEPING TRIM

Virtualized Dell™ PowerEdge™ servers and Dell/EMC storage help German automotive textile leader Johann Borgers streamline IT management and reduce maintenance costs by 25 percent.

CHALLENGE

When the Johann Borgers data center environment could no longer reliably support key administrative functions, the IT team needed to replace it with a manageable, scalable, and highly available storage and server solution.

SOLUTION

Deploying virtualized Dell PowerEdge 2900 servers with a Dell/EMC CX3-10c storage area network simplified IT management by enabling the company to shift to thin client computing, allowing employees to add and update software centrally, and reducing routine maintenance time.

BENEFITS

- Simplified infrastructure helps reduce configuration time for new clients from three hours to 30 minutes and helps reduce routine IT maintenance by 25 percent.
- Streamlined management of end-user PCs helps save €24,000 annually.
- Individually upgraded components help maximize the technology life cycle.

Related Categories:
Case study, Citrix Presentation Server, Dell Infrastructure Consulting Services (ICS), Dell PowerEdge servers, Dell/EMC storage, Microsoft Windows Server 2003, virtualization, VMware

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

Johann Borgers is the leading provider of acoustically efficient textile components for the automotive industry, with more than 4,700 employees worldwide. The company's core administrative activities are centralized at its headquarters in Bocholt, Germany, where office-based employees need constant access to key financial and performance data as well as enterprise resource planning (ERP) systems such as inventory management, accounting, and payroll. However, the seven end-of-life servers in the data center could no longer provide the reliability, availability, or performance the company required. Plus, employees accessed aging ERP tools on thick-client PCs, increasing the administrative burden for the IT team.

"We wanted to replace our old servers with a new storage area network [SAN] and virtualized server environment at our headquarters," says Stefan Schmeinck, IT administrator at Johann Borgers. "At the same time, we wanted to move away from our thick-client architecture to streamline software management and make critical applications available to employees centrally." By consolidating storage and key applications, Johann Borgers sought to simplify its IT infrastructure to enable significant administrative cost and time savings, increase scalability, and reduce the need to invest in additional hardware and software to support future applications and services.

BUILDING A SIMPLIFIED INFRASTRUCTURE

Johann Borgers solicited proposals from several solution providers, but finally chose Dell for the project. "Dell offered us the best solution in terms of price and quality," says Schmeinck. "It was also ahead of the others in terms of professional competence, and we were convinced it was the right partner for us from the beginning."

Johann Borgers deployed three Dell PowerEdge 2900 servers with two quad-core Intel® Xeon® processors as well as a Dell/EMC CX3-10c SAN. Installing VMware® ESX 3.5 virtualization software on the PowerEdge servers enabled the IT team to create 10 virtual machines (VMs) running the Microsoft® Windows Server® 2003 Release 2 (R2) OS and Citrix Presentation Server™ 4.5 software. By running Citrix software on these VMs, the IT team can offer its base of 120 users anytime, anywhere access to the Infor ERP application and to Microsoft Office applications. Although remote access functionality is currently limited to headquarters staff, ultimately all

Johann Borgers sites will have real-time access to critical applications and data at company headquarters.

Dell Infrastructure Consulting Services worked with Dell partner CEMA to design the SAN infrastructure. A Dell Deployment Services team delivered and implemented the necessary hardware on schedule; CEMA installed the operating systems and software and deployed the VMware and Citrix systems. "Everything went according to plan," says Schmeinck. "We wanted to complete the project by the end of March, and that's exactly what Dell and CEMA delivered."

The Johann Borgers IT staff received in-depth SAN training as part of the company's Dell ProSupport for IT agreement. The agreement in place at Johann Borgers also provides 24-hour support from Dell service experts, 365 days a year—and guarantees a four-hour, on-site response time if there is a service issue. Plus, the agreement ensures that new or replacement hardware is shipped and delivered rapidly if necessary.

CENTRALIZING IT MANAGEMENT

Transferring management of the local client desktop to the data center offers a single hardware and software management point from which the IT team can manage applications running on VMs and deploy new applications and upgrades without needing to update individual servers—helping speed the delivery of IT services to employees, save significant time on system maintenance, and simplify troubleshooting and manipulation of virtual resources. "Using the Dell SAN and server solution, we have been able to reduce the total cost of running machines by €200 per PC a year," says Schmeinck. "With our existing base of 120 users, which is the tip of the iceberg, this already represents savings of €24,000 annually."

Johann Borgers has also reduced client device setup time from three hours to 30 minutes. "By simplifying our storage and server environment with Dell, we have dramatically reduced the administrative workload of the IT team and the cost of running each PC, delivering total time and cost savings of around 25 percent," says Schmeinck.

> **"We are very satisfied customers. From the earliest days of this project, it was clear that Dell is far ahead of the competition, and the final results we have achieved show this to be true."**
>
> —Stefan Schmeinck
> IT administrator at Johann Borgers
> February 2009

DEPLOYING GREEN IT AND ENHANCING AVAILABILITY

The solution has delivered environmental benefits by reducing the organization's carbon footprint. Virtualization has increased server utilization and decreased the number of physical servers at the Johann Borgers headquarters from seven to three without compromising performance or capacity—enabling the company to reduce running costs, cooling costs, and maintenance costs while making its IT operations far greener than before.

Replacing the end-of-life servers has dramatically improved the availability of critical data and the performance of key IT services. "Our previous infrastructure could no longer provide the scalability and availability we need to support our business," says Schmeinck. "By contrast, the Dell equipment gives us the agility we need to deploy the next generation of business applications rapidly, with no need to purchase additional hardware." In addition, Schmeinck describes the Dell technology as "fail proof," because virtualization keeps key services constantly available even if hardware fails. "If we experience a fault, the VMware diverts processing and storage for key services to virtual machines running on another hardware device. This means we have no downtime whatsoever and can repair faults without impacting our end users."

Server virtualization also improves the allocation of server resources and reduces the need to purchase additional hardware to support new applications and services. "If required, we can expand the solution by adding new hard disks to the SAN or memory modules to the servers," says Schmeinck. "This system will meet our requirements for at least the next five years with very little additional investment."

Based on its experience with this project, Johann Borgers now plans to work with Dell on simplifying its departmental servers and storage. "We are very satisfied customers," says Schmeinck. "From the earliest days of this project, it was clear that Dell is far ahead of the competition, and the final results we have achieved show this to be true."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell virtualization solutions:
DELL.COM/Virtualization

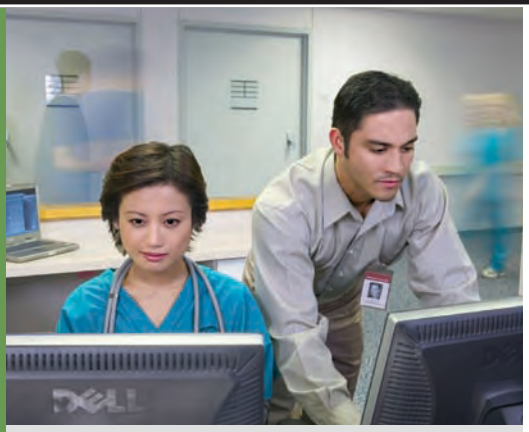

A HEALTHY APPROACH TO SERVER SPRAWL

By migrating to Dell™ PowerEdge™ blade servers, BlueCross BlueShield of Western New York shrinks its data center footprint by 75 percent and improves efficiency by over 50 percent.

CHALLENGE

Health care insurer BlueCross BlueShield of Western New York sought to consolidate its disparate claims-processing server environment to reduce administrative complexity and speed the deployment of new applications and services.

SOLUTION

Dell helped the organization migrate more than 100 servers to Dell PowerEdge M600 blade servers in four PowerEdge M1000e modular blade enclosures without disrupting the business or end users.

BENEFITS

- Shared resources and remote management capabilities help increase IT efficiency by an estimated 60 percent.
- Dell PowerEdge enclosures and blade servers save an estimated 20 percent in hardware costs over 1U and 2U server architecture.
- Dell PowerEdge blade server consolidates racks by 75 percent, helping cut power costs by US$6,000 per year initially and as much as US$40,000 per year.

Related Categories:
Blade servers, BlueCross BlueShield of Western New York, case study, Dell Infrastructure Consulting Services (ICS), Dell PowerEdge blade servers, flexible computing, virtualization, VMware

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

Ideally, the business side of health care should be a secondary concern. Health care insurance provider BlueCross BlueShield of Western New York (BCBS) works hard to help both doctors and patients focus on wellness instead of where to send the bill. BCBS supports more than 800,000 members with simple, straightforward health coverage; easy-to-understand plans; and hassle-free access to information.

BCBS relies on a deep IT infrastructure for everything from claims processing to provider payments and internal operations. However, keeping the BCBS environment running effectively and efficiently had become increasingly expensive and time-consuming. With hundreds of servers running an array of applications on a variety of hardware and operating systems, the BCBS IT team realized that the company had reached an inflection point. The team began by consolidating the IT environment through server virtualization, reducing the server footprint to just over 100 physical servers and 20 virtualized hosts. But they knew that more had to be done to lower costs and ease the management burden on the IT staff.

CONSOLIDATING TO POWERFUL, EFFICIENT BLADE SERVERS

The BCBS IT team calculated that moving from rack servers to blade servers would simultaneously streamline administration and maintenance while preserving the ability to provide dedicated servers when necessary, reducing energy costs, and promoting green technologies. When the team compared solutions from several blade vendors, Dell PowerEdge blade solutions stood out as the clear winner. "The Dell PowerEdge blades and enclosures gave us the best combination of price, performance, and configuration capacity," explains Charles Kibby, technical purchasing coordinator for BCBS. The Dell PowerEdge M1000e modular blade enclosure enabled the team to pack 16 blades into a chassis, compared with a limit of 8 blades per chassis from other vendors. And each PowerEdge blade server could support as many as 20 virtual machines, helping increase efficiency.

The company's prior relationship with Dell also made the decision easy. "We've found that working with Dell allows us to get localized attention when we need it, from the initial planning and configuration all the way through the installation process," says Tim Frank, special projects/IT infrastructure for BCBS.

The team migrated to 64 Dell PowerEdge M600 blade servers with quad-core Intel® Xeon® processors, housed in four PowerEdge M1000e enclosures. BCBS allocated three enclosures to application servers, additional virtualized hosts, and a lab manager environment. The remaining enclosure was dedicated to the next phase in the company's virtualization strategy: VMware® Virtual Desktop Infrastructure (VDI). "VDI running on Dell PowerEdge blades offers us substantial cost savings by replacing standard desktops with remote terminals, with 45 virtual desktops hosted on a single blade," explains Brian Chapman, manager of network systems for BCBS. "All the processing, storage, and most memory resides on a blade instead of the local desktop, extending the average life of the client system from three years to five years or longer. Plus, absolutely no data can be stored on the local terminals, so we're better securing patient data for our customers."

The company's IT team has worked side by side with Dell engineers throughout the entire process. With two months of planning and instruction as well as help with the installation and initial hardware training, Dell Global Infrastructure Consulting Services enabled BCBS staff to integrate the blade project into their normal workload without disrupting the business or end users.

SIMPLIFYING SERVER DEPLOYMENT AND MANAGEMENT

The Dell blade servers led to immediate hardware cost savings for BCBS. "With the performance and scalability of the Dell chassis and blades, we can replace our servers every five years instead of every three years, which we expect to cut our hardware expenditures by 20 percent," says Kibby.

Kibby also credits the enclosure's shared resources and remote management capabilities with easing the administrative complexity of BCBS's environment and accelerating deployment. "With the Dell PowerEdge M1000e blade enclosures and the PowerEdge M600 blade servers, we see a 60 percent increase in efficiency, which lowers the cost of doing business," he explains. "Our IT staff doesn't have to constantly set up or move around cables inside our data center, worry about power management for individual boxes, or other day-to-day functions that you need with a standard 2U or 4U architecture. Instead, they can remotely conduct much of the day-to-day maintenance without having to leave their desks."

> **"With the Dell PowerEdge M1000e blade enclosures and the PowerEdge M600 blade servers, we see a 60 percent increase in efficiency, which lowers the cost of doing business."**
>
> —Charles Kibby
> Technical purchasing coordinator
> for BlueCross BlueShield of Western New York
> April 2009

By shortening the time to production, IT is helping end users quickly realize the benefits of new applications and services. "Previously, providing a server took a day and a half just to get ready for an application to be loaded—unboxing, racking, cabling, loading the operating system, getting the patches ready, and more," says Frank. "With the PowerEdge M600 blades, it will be as little as four hours versus as much as three days to prepare a new server."

OPTIMIZING THE DATA CENTER

The Dell blade servers have helped BCBS take a major step toward realizing a greener data center. "Our data center footprint was reduced from four racks of 1U and 2U servers to the equivalent of one fully configured rack," says Kibby. "Using the Dell PowerEdge M1000e enclosure and the PowerEdge M600 blades, we expect to see an energy savings of US$6,000 per year initially and as much as US$40,000 per year as expiring physical servers are converted to blades. We'll also reduce our space requirements, which will tie in with our LEED Silver Certification."

Ultimately, the BCBS IT team feels that the benefits of migrating to Dell blade servers reach far beyond the data center. "Dell has helped us increase our efficiencies across the board, which accelerates claims processing as well as payment to the health providers we work with," says Frank. "Better yet, the cost savings can be passed on to our end customers." Kibby agrees: "The importance of our Dell enclosure and blades is that they allow us to provide seamless claims processing when supporting patients around the country. That means our customers can be confident we're supporting them as they focus on their health and that of their families."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades

SECURITY THROUGH MATURITY: A BRIEF ON SECURING iSCSI NETWORKS

By Robert Winter
Dan McConnell

Internet SCSI (iSCSI) offers a scalable, simplified, cost-effective way to implement storage area networks (SANs) using standard Ethernet components. By understanding the core elements of network security, key differences between iSCSI and Fibre Channel, and best practices for securing iSCSI networks, administrators can implement robust, secure iSCSI SANs.

Related Categories:

Best practices
Fibre Channel
Internet SCSI (iSCSI)
Network security
Security
Storage
Storage area network (SAN)

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

Internet SCSI (iSCSI) can provide a number of advantages in enterprise storage area network (SAN) infrastructures—offering similar scalability, availability, and manageability as Fibre Channel in a cost-effective way, and enabling IT administrators to work with standard, familiar Ethernet components rather than specialized Fibre Channel equipment. iSCSI-capable Dell™ EqualLogic™ PS Series, Dell PowerVault™, and Dell/EMC storage systems can provide a simplified, scalable way to implement iSCSI SANs.

The underlying technologies—the iSCSI protocol, Ethernet, and TCP/IP transport—help provide a well-known, mature, and cost-effective security infrastructure for iSCSI SANs. Understanding the core elements of effective security strategies, the differences between iSCSI and Fibre Channel security, and best practices for iSCSI security can help administrators create robust, secure iSCSI SAN deployments.

CORE ELEMENTS OF EFFECTIVE SECURITY

Storage-based security can be grouped into two distinct areas, referred to as data at rest (DAR) and data in flight (DIF). Both DAR and DIF have security strategies that include elements of authentication, authorization, and data coherence. Because DAR refers to data contained within a physical device and DIF refers to data moving between devices, not all of these elements need be present at the same time, and each may have differing levels of implementation in a given environment.

Data-at-rest and data-in-flight security

DAR provides security for a physical storage device by encrypting data and providing secure access to data on the device—helping prevent unauthorized data access even if the device is removed from its secure infrastructure. DIF provides network-oriented security, including encryption of packet, or *protocol data unit* (PDU), contents to help ensure payload confidentiality; PDU hashing to help ensure delivery integrity; and secure access to network entities.

DAR security may not be required if physical access to storage devices is strictly controlled, while DIF security may not be required for a SAN that is isolated from other networks. DIF security is recommended if the transported data supports public or shared access, such as remote storage replication over a public network for backup or disaster recovery. An effective security strategy should evaluate the needs of the specific environment and provide guidance on implementation.

Authentication, authorization, and data coherence

In iSCSI environments, *authentication* refers to the login phase authentication of the initiator, or the mutual authentication of the initiator and target. Several

cryptographic techniques are possible for iSCSI authentication: Kerberos V5 (KRB5), Simple Public-Key Generic Security Services Application Programming Interface (GSS-API) Mechanism 1 (SPKM1), Simple Public-Key GSS-API Mechanism 2 (SPKM2), Secure Remote Password (SRP), and Challenge Handshake Authentication Protocol (CHAP). CHAP is by far the most common method and is required for iSCSI devices to implement. It uses one-way, three-phase authentication for targets authenticating initiators and for initiators authenticating targets, which is generally sufficient for the majority of storage environments.

Authorization defines allowed actions. Access control, implemented in most types of networks using access control lists (ACLs), provides authorization of communication channel access based on information contained in PDUs, such as IP address, iSCSI Qualified Name (IQN), virtual LAN (VLAN), and other parameters (also called *n*-tuples). Authorization can be as simple as permitting access to resources from well-known *n*-tuples unique to an iSCSI session, or may include additional layers of authorization that further restrict user access to not only devices but also services.

Data coherence preserves data confidentiality and integrity. Confidentiality is commonly termed encryption, and helps ensure that the data payload is transmitted securely. Integrity is commonly termed hashing, and helps ensure that information received has not been modified. iSCSI data confidentiality provides encryption of information over an iSCSI frame. iSCSI data integrity may be provided at five levels: the Ethernet, TCP, IP, iSCSI, and IP Security (IPsec) checksums. In Ethernet, the checksum is also called the frame check sequence (FCS) and is a 32-bit cyclic redundancy check (CRC). The iSCSI checksum is also called the iSCSI digest, and may cover multiple iSCSI PDUs.

COMPARISON OF iSCSI AND FIBRE CHANNEL SECURITY

iSCSI provides a more mature security solution set than Fibre Channel. iSCSI security today is managed at the endpoints (initiator and target) by way of robust authentication and authorization techniques.

Fibre Channel security is managed in the network and the endpoints, which makes it more difficult to manage than iSCSI and creates additional challenges for security management. The vast majority of critical communications today occur over TCP/IP and not over Fibre Channel. TCP/IP-based security has thus by necessity become increasingly mature over decades of development. Because iSCSI is defined and transported over TCP/IP, it can make comprehensive use of the deep, robust pool of available security technologies and strategies that enable solutions across the authorization, authentication, and data coherence security elements. Fibre Channel, in contrast, has only recently proposed a solution for data coherence: Fibre Channel Security Protocol (FC-SP). Although this proposal borrows from many of the base technologies that IPsec uses, it is not yet deployable and does not therefore have the necessary maturity of the iSCSI solution set. It will likely take years of trial and error before appropriate operational security strategies are developed across Fibre Channel to cover all key security elements. Figure 1 shows the iSCSI and Fibre Channel protocol layers, and Figure 2 summarizes the basic components of iSCSI and Fibre Channel security.

Fibre Channel, which is not a generally well-known or widely used protocol, offers a policy of *security through obscurity*. iSCSI, which runs over well-established protocols, provides *security through maturity* using robust security protocols and techniques. iSCSI therefore can also continue to evolve by taking advantage of the work being done on both Ethernet and TCP/IP standards.

Passive, active, and operational attacks

Attacks on SANs, such as iSCSI or Fibre Channel SANs, typically fall into one of three categories: passive, active, and operational. A *passive attack* is based on snooping of clear-text information in the packet stream, with *n*-tuples comprising IP addresses, Media Access Control (MAC) addresses, World Wide Names (WWNs), IQNs, VLANs, TCP port numbers, and other parameters. Attacks can occur using individual management interfaces on switches or through mirrored switch ports, which can access data on other ports. SANs can also be vulnerable to a passive attack if the attackers have physical access to the fabric cabling.

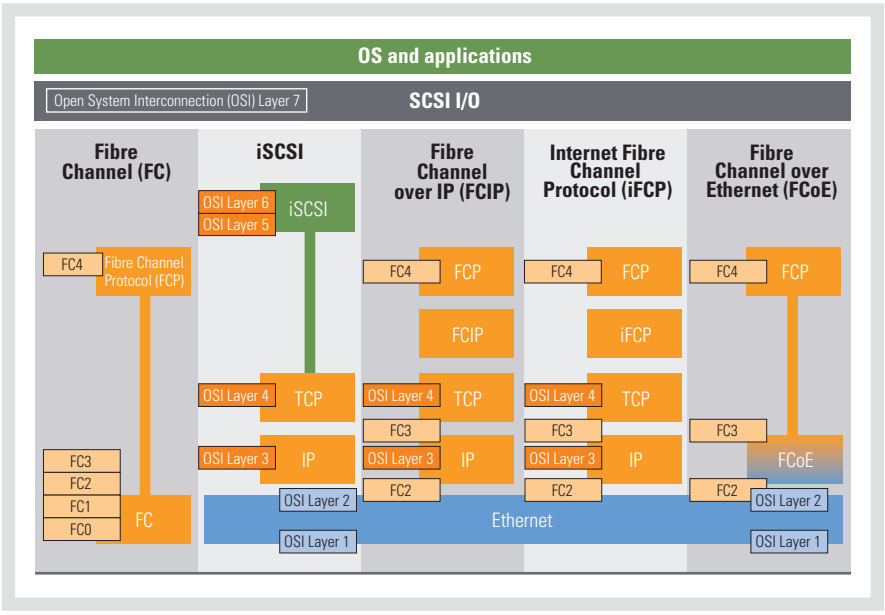

Figure 1. *iSCSI and Fibre Channel protocol layers*

In an *active attack*, those attempting to obtain access typically terminate or intercept a communication endpoint (either an initiator or target) with equipment that they control, and then misrepresent that equipment as the original endpoint. This type of attack can occur only if the attackers have access to the physical or logical network and can thereby intercept the packet streams. Another form of active attack can flood packets to iSCSI or Fibre Channel endpoints, preventing timely access to the network by valid traffic.

An *operational attack* misuses a management portal on an infrastructure component, which might be an initiator, target, or switch. Ethernet and Fibre Channel switches may host an HTTP service, a Simple Network Management Protocol (SNMP) agent, or a Common Information Model (CIM) agent, and these management portals therefore require protection. Another type of attack is a break-in to the server. The underlying storage protocol cannot protect that server's data further, but it is still important that this infected server not be able to access resources that do not belong to it.

Domain isolation and IPsec

Domain isolation can help easily resolve many of the vulnerabilities described in the preceding section. This isolation can take two forms: physical and logical. Physical isolation means that the SAN is entirely closed and has no connection to the rest of the network; this approach can be somewhat impractical, but is a common implementation for Fibre Channel SANs. Logical segmentation can be implemented in iSCSI using VLANs, which are commonly used to isolate management portals across different infrastructures and can also be used to protect access to data types within the same SAN. Fibre Channel supports the similar concept of zoning as well as virtual SANs (VSANs), although VSANs were only recently introduced.

IPsec can provide highly flexible DIF security for iSCSI. Because Fibre Channel is not IP based, it must use non-IPsec methods to create secure tunnels between Fibre Channel SANs or perform translation between IP and Fibre Channel networks. IPsec is designed to be completely transparent to iSCSI protocol operation and provides mechanisms for implementing strong authentication, authorization, and data coherence. It uses two protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). AH is designed to provide authentication and integrity of IP packets, but does not provide confidentiality for the payload; ESP is designed to provide integrity and confidentiality for the payload as well as a form of authentication.

Each of these protocols can operate in either tunnel mode or transport mode (see Figure 3). Tunnel mode is typically used to help ensure the confidentiality of the original IP header and payload between endpoints (such as gateways) connecting two domains, often when data must leave the secure confines of a LAN or wide area network (WAN) and travel between hosts over a public network such as the Internet. Transport mode can provide end-to-end confidentiality but does not encrypt the original IP header, and can thus allow the packet to participate in normal routing operations for the domain.

	iSCSI	Fibre Channel
Authentication	CHAP	CHAP, FC-SP, Fibre Channel Authentication Protocol (FCAP)
Authorization	IQNs, ACLs	WWNs, ACLs
Confidentiality	IPsec ESP and AH protocols in tunnel or transport mode	FC-SP (based on some aspects of IPsec)
Integrity	Ethernet CRC (within LAN), TCP checksum (end to end), IP checksum (end to end), iSCSI digest (end to end), IPsec ESP and AH protocols in tunnel or transport mode	CRC (end to end)
Segmentation	VLANs, IP subnets, Internet Storage Name Service (iSNS)	VSANs, Fibre Channel name services

Figure 2. *Comparison of iSCSI and Fibre Channel security components*

BEST PRACTICES FOR iSCSI SECURITY

Following best practices appropriate for specific deployments or usage models can help organizations secure their iSCSI SANs as part of a comprehensive enterprise security strategy. This section outlines a few of the recommended practices that administrators should consider.

Isolate the iSCSI domain. Because iSCSI SANs can be switched within Ethernet LANs or routed through IP subnets, iSCSI can exist in mixed-traffic networks. Enterprise data centers typically segment different kinds of network traffic by either department or application type. Providing similar segmentation for iSCSI storage traffic through domain isolation—using either physical isolation or logical segmentation—can help secure the iSCSI SAN from other traffic.

Consolidate and secure management portals. Networks often include many management consoles, with switches, hosts, and routers having their own management agents or applications. Best practices recommend limiting the number of management portals and restricting access to them so that only one console has overall control of a particular iSCSI SAN domain. This primary management portal should be accessed using robust security pathways, such as virtual private networks (VPNs), especially if the access is external to the network.

Disable unused switch and router features. Ethernet switches may include features that are unnecessary in an iSCSI

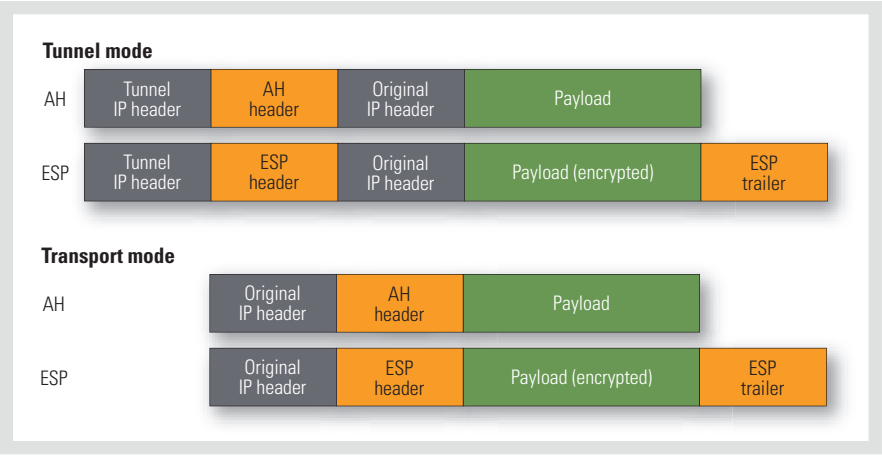

Figure 3. *IPsec headers in tunnel mode or transport mode*

SAN. For example, switches may include up to four management portals—SNMP agent, HTTP server, Telnet command-line interface (CLI) access, and serial CLI access—but obviously not all of these portals are required. Based on the best practice of using a single primary management console, administrators should lock down these portals so that they can access only the primary management console's IP address or configured VLAN. They should also typically turn off port mirroring and other features that may be unnecessary, such as Layer 3 routing on switches used only as Layer 2 devices.

Use access control lists for authorization. ACLs, a common feature of Ethernet switches, match iSCSI PDUs based on a defined *n*-tuple. ACLs can limit access to particular switch ports based on administrator-defined restrictions, with *n*-tuples that do not represent membership in the iSCSI SAN then being rejected by the switch. ACLs provide a powerful way to control network access to services, and best practices recommend using them as a network resource access authorization technique.

Use CHAP when possible. CHAP is used by VPNs and the Microsoft® Windows® OS, and is commonly supported by iSCSI devices. CHAP can have the target authenticate the initiator but also enables targets and initiators to authenticate each other. iSCSI SAN deployments should generally employ the strongest authentication method available, using CHAP when possible for entity authentication.

Use IPsec when appropriate. Although technical and operational concerns may limit its deployment today, when practical, administrators should use IPsec in its minimal required configuration—ESP in tunnel mode between security domains—for data coherence. Typical deployments focus on environments in which data leaves the local data center security domain, and generally employ a security or VPN appliance for using IPsec between data centers.

Define a practical and secure key management strategy. Organizations should define practical and secure key management strategies that fit their needs and resources. Keys for CHAP, DAR, and DIF encryption require careful management—lost keys, for example, may make data inaccessible. Organizations typically develop their own key management strategies and guard them closely, because exposing these strategies can represent a security threat. Best practices recommend that key management strategies include the following minimum elements:

- **Enterprise-wide policy:** The policy should be well-known, enforceable, and consistent across the enterprise. Separate internal groups should not develop their own policies.
- **Secure key creation:** Enabling secure key creation can be difficult, but using purpose-built hardware based on Federal Information Processing Standard (FIPS) 140-2 Level 3 encryption can help simplify this task.
- **Key revocation:** Key revocation policies should cover key aging, key reuse, and re-encryption of old data with new keys.
- **Key escrow:** Keys should be maintained securely between employers and employees such that employees do not have complete ownership or knowledge of the keys. A secure third party might be used to implement this mechanism.
- **Key distribution:** Key distribution may be as simple as having a dedicated organization within an enterprise that manually distributes and manages keys, or as complex as a full Public Key Infrastructure (PKI) implementation.

ROBUST, SECURE iSCSI DEPLOYMENTS

Strong security does not necessarily mean expensive security. By taking advantage of the maturity of Ethernet and TCP/IP technologies that serve as the basis for iSCSI networking and following key iSCSI best practices, administrators can cost-effectively implement robust, secure iSCSI SANs as part of a comprehensive enterprise security strategy.

Robert Winter is a network technologist in the Dell Office of the CTO.

Dan McConnell is a product marketing strategist in the Dell Enterprise Storage Group.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell iSCSI storage solutions:
DELL.COM/iSCSI

ADVANCED FILE SHARING AND MANAGEMENT IN THE DELL NX4

By Scott Sinclair

The Dell™ NX4 network attached storage (NAS) device goes beyond simple file servers to provide advanced features ranging from point-in-time snapshots and replication to deduplication, file-level retention, and tiered storage access. Combined with support for both Internet SCSI (iSCSI) and Fibre Channel, this system can help meet enterprise-class requirements, simplify data management, and reduce costs.

Related Categories:

Data consolidation and management

Dell NX4

Network attached storage (NAS)

Storage

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

The Dell NX4 network attached storage (NAS) device is designed to provide flexible, enterprise-class file storage for Microsoft® Windows®, Linux®, and UNIX® environments. Advanced EMC® features such as point-in-time snapshots, replication, and deduplication enable the Dell NX4 to provide functionality well beyond simple file servers to help save administrative time and reduce costs, while support for both Internet SCSI (iSCSI) and Fibre Channel helps maximize flexibility in multi-protocol environments. File and block storage can be consolidated into a single Dell NX4 system to help simplify management.

The Dell NX4 uses the Data Access in Real Time (DART) OS for multi-protocol network file and block access. DART supports concurrent use of the Common Internet File System (CIFS) and Network File System (NFS) protocols and supports sophisticated locking and access-control mechanisms, enabling seamless file sharing across Windows-, Linux-, and UNIX-based client systems without compromising data integrity and without requiring performance-reducing emulation.

Providing capabilities designed for organizations that require functionality well beyond that of a simple file server, the Dell NX4 helps share files efficiently across both CIFS and NFS systems, protect critical file data with high availability, save wasted disk space where data growth has led to a proliferation of duplicate files, comply with data retention regulations, and implement efficient backup and recovery processes.

HIGH-AVAILABILITY ARCHITECTURE

Because file data can be as vital as application data, the Dell NX4 is designed to provide high availability without compromise—enabling organizations to continue operating at the same performance and service levels even in the event of a failure. The primary/standby architecture supports automatic failover and is designed to eliminate single points of failure from the network to the disk drive. The system is designed for high levels of fault tolerance, while the DART and EMC FLARE® operating environments offer advanced fault detection and isolation capabilities. Administrators can take advantage of hot-pluggable and standby components for transparent high availability and rapid recovery.

POINT-IN-TIME SNAPSHOTS

EMC Celerra® SnapSure™ software enables administrators to create logical point-in-time read-only or read/write copies of file systems and iSCSI logical units (LUNs) to support online backups and quick recovery of deleted files. End users can quickly and easily restore previous versions of a file without administrator involvement. SnapSure also helps save disk space and time by

supporting multiple snapshot versions of a file system or iSCSI LUN.

SnapSure is designed for excellent read performance. It operates on the copy-on-first-write principle: when a block within a production file system is modified, SnapSure saves a copy containing the block's original contents to a separate volume, referred to as the SavVol. Subsequent changes to the same block are not copied to this volume. SnapSure reads the original blocks in the SavVol and the unchanged blocks remaining in the production file system to provide a complete point-in-time image called a checkpoint.

Checkpoints can be read-only or read/write, and can serve as a direct data source for applications that require point-in-time data. Administrators can use checkpoints to restore a production file system or part of a file system to a previous state. Additionally, they can use writable checkpoints to create a test environment for databases and files where, after testing, changes can be applied nondestructively to the production file system. SnapSure also supports online checkpoint access, helping eliminate the need for administrator involvement when a client system needs to list, view, or copy a point-in-time file or directory in a read-only checkpoint or use it to restore information from a previous point in time.

ASYNCHRONOUS REPLICATION

As an option to the Dell NX4 NAS device, EMC Celerra Replicator™ software provides a powerful, simplified tool for asynchronous replication (see Figure 1). This tool uses standard IP-based local area networks (LANs) and wide area networks (WANs) to maintain consistent read-only or read/write replicas across sites, helping simplify configuration and management and enabling deployment of remote replication using familiar IP networking—which can result in significant savings in both staff time and budget. Administrators can easily define policies such as recovery point objectives (RPOs), and the software includes an adaptive scheduler that can determine the size and frequency of updates necessary to meet a given RPO, taking into account available bandwidth, data load, and data transfer concurrency.

Organizations can take advantage of Celerra Replicator in a number of ways:

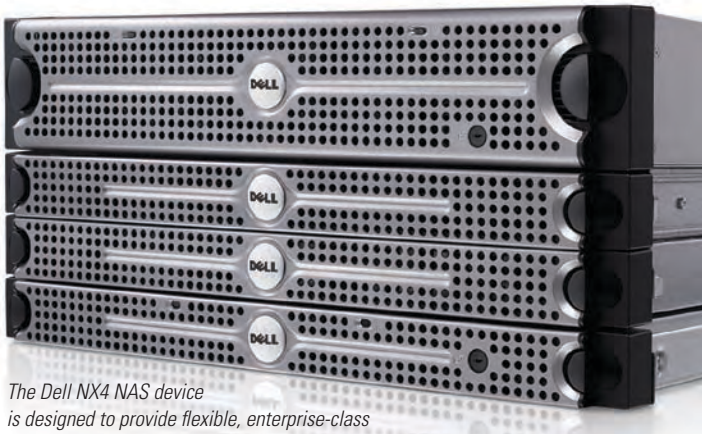

The Dell NX4 NAS device is designed to provide flexible, enterprise-class file storage in heterogeneous mixed-protocol environments

- **Disaster recovery:** Administrators can replicate a duplicate copy of production data to a remote site, where it can be brought online with little downtime in case of a disaster.
- **Content distribution:** One-to-many replication can push content to remote sites—for example, when new engineering or software builds need to be distributed to multiple locations.
- **Backup:** Performing backups with a copy of the production data helps avoid the need to take the production applications offline. The backup can occur locally or at the remote location.
- **Decision support:** File systems and iSCSI LUNs can be replicated to make a copy of a database to be used for data mining and decision support without affecting production applications.
- **Software testing:** Before upgrading software on production systems, administrators can create a duplicate copy of the data and test the upgrade. Writable checkpoints allow administrators to test software with a modifiable copy of the production data.
- **Data center migrations:** When relocating to a new data center, administrators can copy data to the new system and force a failover, helping ensure that the migration is carried out without data loss.

Celerra Replicator can also integrate with Celerra SnapSure and VMware® Site Recovery Manager (SRM) to help simplify

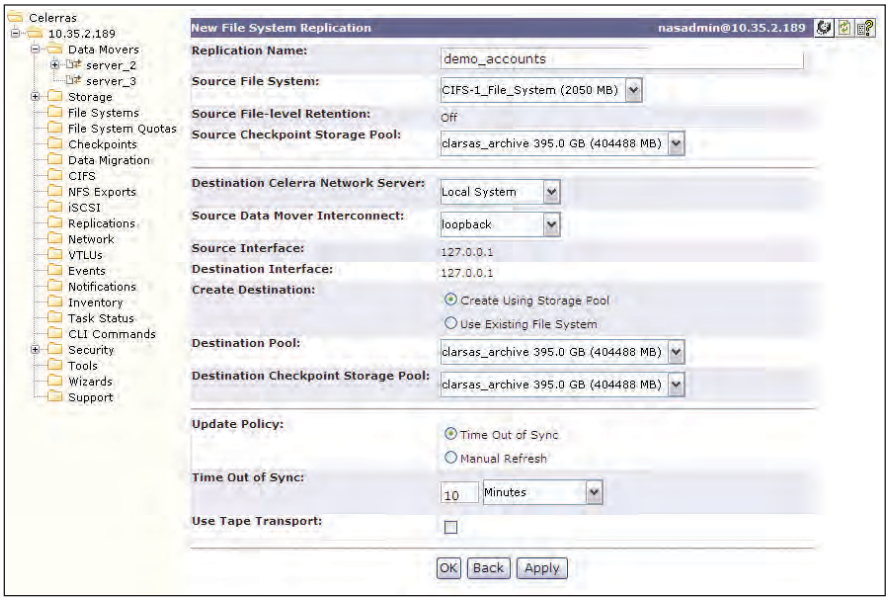

Figure 1. *Celerra Replicator provides a powerful, simplified tool for asynchronous replication*

the testing and execution of disaster recovery processes.

VERSATILE BACKUP AND RESTORE

The Dell NX4 offers several options for backup and restore functionality to help meet the needs of common backup environments and maximize flexibility:

- **Network backups:** Network backups entail mounting CIFS or NFS file systems across a network and backing up to the backup server.
- **Network Data Management Protocol (NDMP) backups:** NDMP backups transfer data to a local backup device and use the LAN only for control information. This approach is typically employed for high-capacity backups and in environments requiring multiprotocol support.
- **Storage area network (SAN) backups:** SAN backups use the LAN for control information and do not involve the server in the data path from the SAN backup server.
- **NDMP volume-based backups (NVBs):** NVBs are an NDMP block-based backup option that can offer significant performance advantages for file systems with large numbers of small files (less than 16 MB). Using NVB in these environments can help significantly reduce backup windows.

DATA DEDUPLICATION AND COMPRESSION

The Dell NX4 was designed to help organizations effectively handle the proliferation of unstructured data, combining file-level deduplication and compression to help intelligently reduce storage requirements with minimal performance impact. A built-in policy engine works in the background, transparently monitoring file activity and file attributes to intelligently identify candidate files. Files that meet certain criteria, such as low access frequency, are both compressed and deduplicated so that only a single instance of those files are retained. Administrators can define specific filters to avoid processing files for which deduplication could result in an undesirable performance impact or only minimal storage savings.

FILE-LEVEL RETENTION

The Celerra file-level retention (FLR) feature in the Dell NX4 enables administrators to protect files from modification or deletion until a specific retention date, helping organizations comply with enterprise, industry, or government requirements. FLR enables administrators, using CIFS or NFS operations, to meet write-once, read-many (WORM) requirements by creating a permanent, unalterable set of files and directories and offering an audit trail that logs user activity in that file system. The Dell NX4 offers this feature as a U.S. Securities and Exchange Commission (SEC) Rule 17a-4(f)-compliant option and as a noncompliant option.

ANTIVIRUS PROTECTION

The optional Celerra Event Enabler (CEE) framework contains the Celerra Anti-Virus Agent (CAVA), which provides an antivirus solution for client systems using Celerra Network Server software. The agent uses CIFS in Microsoft Windows Server® 2003 or Windows 2000 Server domains, and uses third-party antivirus software from companies such as Symantec, McAfee, Computer Associates, Trend Micro, and Sophos to help identify and eliminate known viruses before they infect files on the storage system.

COST-EFFECTIVE TIERED STORAGE ACCESS

The Dell NX4 features the open Celerra FileMover application programming interface (API) for tiered storage access. Administrators can migrate information stored on the Dell NX4 to secondary storage as well as to purpose-built archiving solutions with comprehensive transparency to users and third-party applications. The API requires a third-party policy engine to specify migration policies.

After a migration, file metadata remains on the Dell NX4 in an offline stub file. To a client system or application, the migrated files appear as though they have not moved, enabling the Dell NX4 to remain the primary interface for clients and applications. When file access is requested, the Dell NX4 automatically and transparently retrieves the file from the secondary storage. Using this approach to store critical data on the highly available Dell NX4 and infrequently used data on cost-effective secondary storage can help reduce costs for backup hardware, software, and tape media; increase efficiency through accelerated backups, consistent protection, and increased storage utilization; and enhance service levels through increased performance, high availability, and rapid recovery.

FLEXIBLE, ENTERPRISE-CLASS FILE STORAGE

The Dell NX4 offers a variety of advanced features that can help organizations increase storage efficiency and reduce total cost of ownership. These advanced features, combined with the flexibility of file management for Microsoft Windows, Linux, and UNIX environments and support for both iSCSI and Fibre Channel, make the Dell NX4 well suited to help meet enterprise-class requirements while also helping to simplify management and reduce costs.

Scott Sinclair is a product manager and strategist for Dell storage solutions.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Full white paper:
DELL.COM/NX4FileFeatureWP

Dell NX4:
DELL.COM/NX4

STREAMLINING STORAGE WITH COMMVAULT DEDUPLICATION IN THE DELL POWERVAULT DL2000

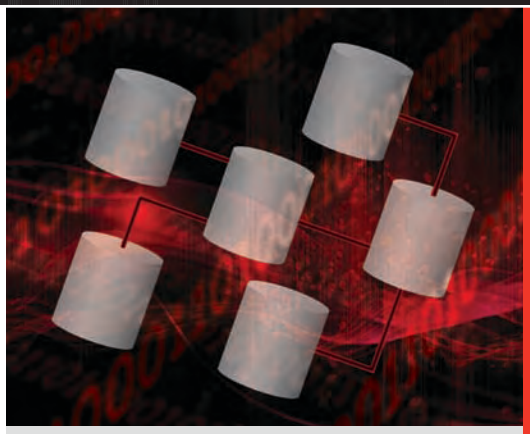

By Jeff Echols

Redundant data is often a major contributor to rising storage costs. The deduplication capabilities of CommVault® Simpana® 8.0 software in the Dell™ PowerVault™ DL2000 – Powered by CommVault are designed to eliminate unnecessary copies of backup and archive data—providing an integrated, simplified way to streamline storage throughout the enterprise.

Related Categories:

CommVault

Data consolidation and management

Dell PowerVault storage

Storage

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

Redundancy is a major contributor to sustained data growth in enterprise data centers. Regular replication, backup, and archiving using traditional tools, for example, often requires organizations to deploy significantly more storage than they actually need simply to maintain unnecessary copies of the same data across many different storage systems. The systems necessary to store this redundant data represent not only an up-front investment in hardware, deployment, and configuration, but also ongoing operational costs for power, cooling, and management.

Deduplication technology, which is designed to eliminate duplicate copies of data, offers an important method of combating uncontrolled data growth. Dell and CommVault have worked together to create a backup and archive deduplication solution specifically designed to address the varied needs of enterprise environments: the Dell PowerVault DL2000 – Powered by CommVault. Combining high-performance Dell hardware with innovative CommVault Simpana 8.0 software, the PowerVault DL2000 – Powered by CommVault provides an integrated, simplified way to streamline storage throughout the enterprise.

STREAMLINING STORAGE REQUIREMENTS AND COSTS

At the core of the Dell PowerVault DL2000 – Powered by CommVault is Simpana 8.0, which integrates backup and recovery, archiving, replication, storage resource management (SRM), and search modules into a single centrally managed platform. Simpana 8.0 also incorporates advanced deduplication technology into this platform, delivering one of the industry's first block-level deduplication solutions for backup and archive data on disk and tape.

Traditional backup and archive architectures repeatedly back up redundant data, consuming valuable disk space with multiple copies of the same blocks. Simpana 8.0, in contrast, can write compressed, deduplicated, and even encrypted blocks to disk, enabling organizations to maximize use of existing storage hardware and helping reduce the amount of data sent over the network (see Figure 1).

After backup and archive data has been deduplicated and stored on disk, administrators can replicate this data to another disk target for disaster recovery or schedule a synthetic full backup copy to tape. Simpana deduplication can not only help reduce data stored on disk, but can also deduplicate that

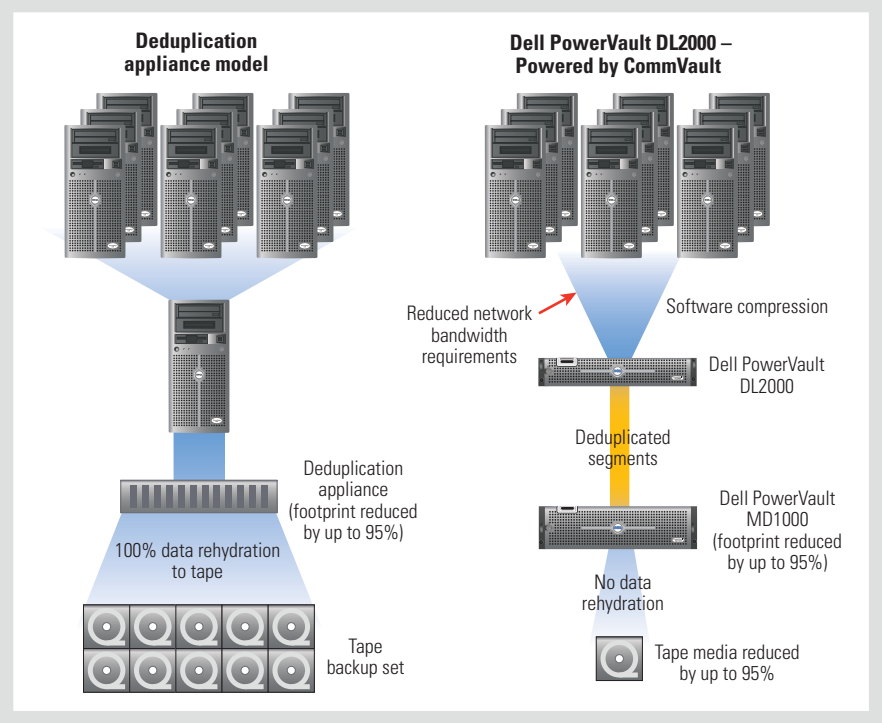

Figure 1. *CommVault Simpana supports block-level deduplication of backup and archive data on disk and tape*

same backup or archive data to tape—potentially saving up to 95 percent on off-site tape storage costs.[1] Disk-to-tape deduplication also avoids the time-consuming process of rehydrating data—expanding it back to its original size—before transferring it to tape, helping dramatically reduce tape media costs.

BACKING UP VIRTUALIZED ENVIRONMENTS

CommVault Simpana 8.0 also introduces the Virtual Infrastructure Agent for VMware® and Microsoft® Hyper-V™ virtualization platforms. In VMware environments, this agent is deployed on a proxy server instead of the virtual machine (VM), and is designed to protect a virtually unlimited number of VMs on an ESX server. This agent can also integrate seamlessly with VMware Consolidated Backup (VCB) to back up VM data over a storage area network (SAN) without requiring scripting (see Figure 2).

The Virtual Infrastructure Agent also integrates CommVault image-level backup technology, enabling it to complete single-pass VMware virtual disk (.vmdk) block-differential backups while maintaining the ability to restore individual files to the VM. Rather than backing up the entire .vmdk file, only changed blocks are backed up, helping save network bandwidth and disk space. This module is enhanced when combined with Simpana block deduplication, because only the changed .vmdk blocks are deduplicated as they are written to disk. Figure 3, for example, outlines how this capability could help significantly reduce disk costs in an example environment with five 100 GB VMs.

REDUCING BACKUP FOOTPRINTS FOR REMOTE OFFICES

Remote office servers can also benefit from a global deduplication strategy based on CommVault Simpana 8.0. For example, reducing backup footprints before they are sent across a wide area network (WAN) to a centralized location can help significantly simplify backup and recovery processes. Offices with short recovery time objectives can deduplicate backups to a local Dell PowerVault DL2000 system while writing the same deduplicated backup stream—with encryption, if necessary—to a local data center for disaster recovery (see Figure 4). The ability to manage the environment from a central console helps provide a scalable, simplified solution. After deduplicated data has been sent to the central site, administrators then have the flexibility to write deduplicated,

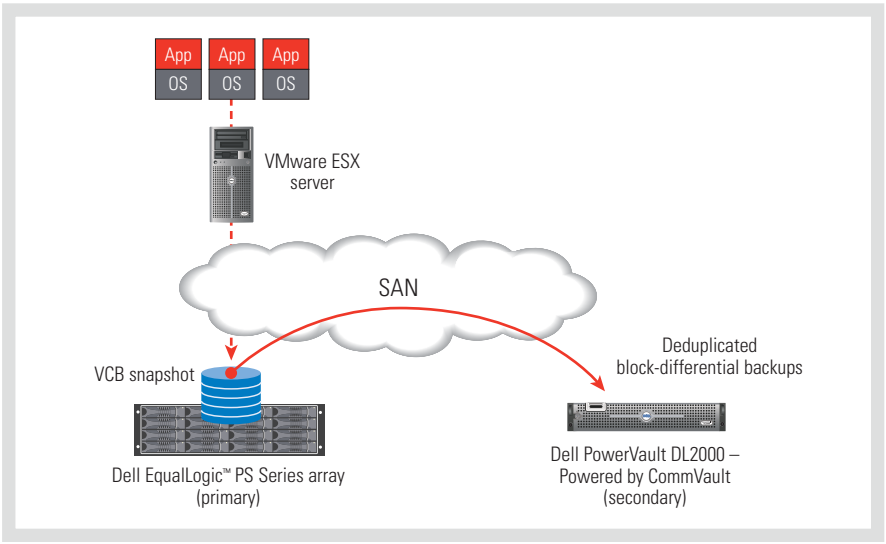

Figure 2. *The CommVault Simpana Virtual Infrastructure Agent helps provide seamless backups in virtualized environments*

[1] Based on daily backups over 20 days with a daily change rate of 5 percent and 50 percent data compression. Actual data reduction will vary based on configuration and usage.

encrypted data to tape before sending it off-site for disaster recovery.

DEPLOYING INTEGRATED DATA MANAGEMENT

The Dell PowerVault DL2000 – Powered by CommVault provides an integrated, comprehensive disk deduplication solution designed for simplified deployment, configuration, and management in environments of all sizes. In addition to disk and tape deduplication backup, this system also offers host-based replication as well as file, e-mail, network attached storage (NAS), and Microsoft Office SharePoint® archiving options. Backup and recovery agents are available for Microsoft Windows®, Novell® NetWare®, Linux®, and UNIX® file systems, as well as for Microsoft Active Directory®, Microsoft Exchange, Microsoft SQL Server®, Microsoft Office SharePoint, Microsoft Office Communications Server, Novell GroupWise®, IBM® Lotus® Notes, Oracle®, and MySQL software.

As administrators struggle with increasingly restricted IT budgets, they need tools that can deliver cost-effective, simplified data management. The deduplication capabilities provided by CommVault Simpana 8.0 in the PowerVault DL2000 – Powered by CommVault can help overcome these challenges, offering an efficient way to consolidate data, streamline storage requirements, and control costs.

Jeff Echols is a director of business development at CommVault. He has a B.S. in Mechanical Engineering and an M.B.A. from the University of Texas at Austin.

	Traditional backup	CommVault Simpana
Number of VMs	5 (100 GB each)	
Full backup size	500 GB	
Full backup size after deduplication	N/A	200 GB
Daily incremental backup size	500 GB	15 GB
Amount of data moved and stored after six days	3,000 GB	275 GB
Compression ratio	1:1	11:1
Disk cost	US$9,000	US$825

Note: Data requirements assume 60 percent deduplication of .vmdk data after one pass, 3 percent change rate for incremental backups, and no deduplication or block-differential backups of .vmdk data in the traditional backup system. Cost calculations assume disk costs of US$3,000 per terabyte.

Figure 3. *Deduplication can help significantly reduce storage hardware costs*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and CommVault:
DELL.COM/CommVault

Dell PowerVault DL2000:
DELL.COM/DL2000

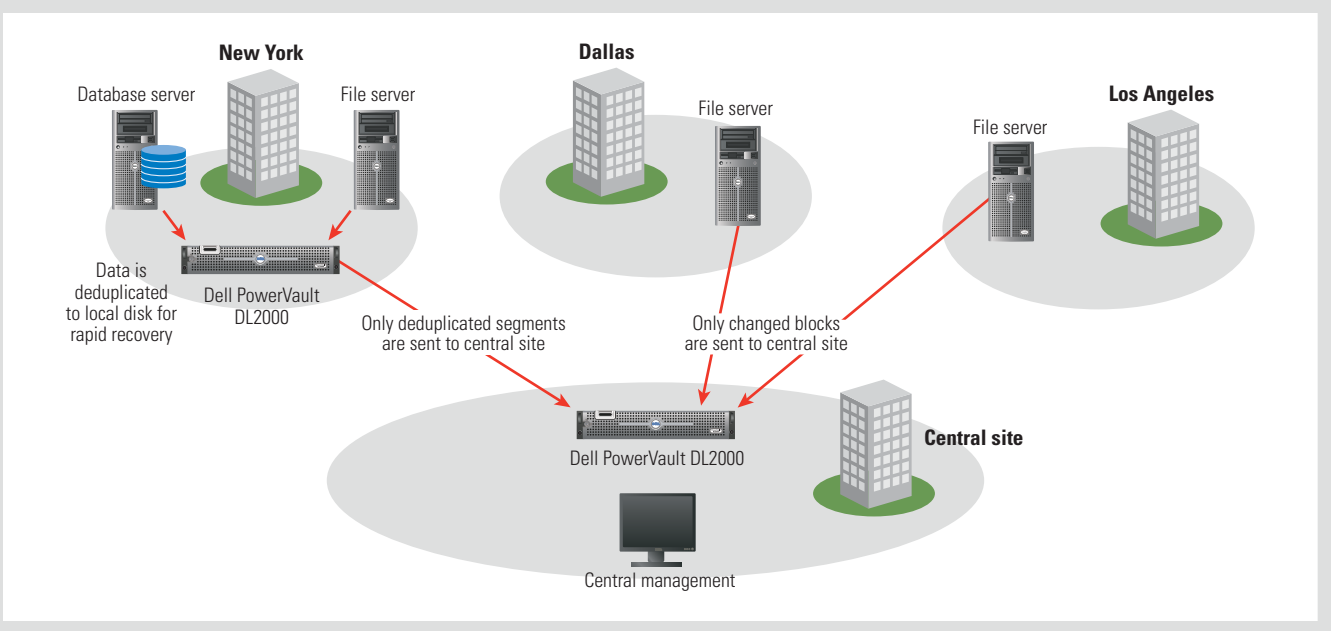

Figure 4. *CommVault Simpana deduplication can support scalable, simplified backups for remote offices*

OPTIMIZING NEARLINE STORAGE IN A 2.5-INCH ENVIRONMENT USING SEAGATE CONSTELLATION DRIVES

By Barbara Craig
Tom McCaffrey

Combining the Dell™ PowerVault™ MD1120 storage array with energy-efficient 2.5-inch Seagate® Constellation™ hard drives can provide the capacity, performance, and reliability to support cost-effective nearline storage as part of a tiered environment.

Related Categories:

Dell PowerVault storage
Energy efficiency
Performance
Power and cooling
Seagate
Storage

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

Keeping up with space constraints and the rising costs of energy are two of the greatest challenges facing IT departments today—and overhauling an entire infrastructure to overcome these challenges can be both costly and difficult. By controlling power consumption and related heat output at the system level, however, organizations can deploy their hardware resources for optimum power and cooling efficiency to help control ongoing costs.

The modular, rack-dense Dell PowerVault MD1120 direct attach storage (DAS) array is designed for efficiency and scalability—holding up to 24 low-power, 2.5-inch hard drives in a compact 2U enclosure for up to 12 TB of total storage capacity. A single dual-port Dell PowerEdge™ Expandable RAID Controller (PERC) 6/E can connect to up to six PowerVault MD1120 enclosures, enabling use of up to 144 total drives and helping simplify management. And now, the PowerVault MD1120 can also support energy-efficient Seagate Constellation Serial Attached SCSI (SAS) hard drives. These industry-first 2.5-inch, 7,200 rpm, 500 GB nearline SAS drives are designed from the ground up specifically for enterprise-class nearline storage, providing the capacity, performance, and reliability to support high-density tiered storage applications.

DESIGNING FOR TIERED STORAGE ENVIRONMENTS

The rotational vibration tolerance of Seagate Constellation drives makes them well suited for tiered storage environments, enabling them to maintain performance in spindle-dense enclosures such as the Dell PowerVault MD1120 even under the most challenging intermix and fan speed conditions. In these environments, administrators assign different categories of data to different types of media to help optimize total storage costs. In terms of hard drives, tier 1 storage would typically use 10,000 or 15,000 rpm drives to support highly transactional applications, while tier 2 storage would typically use high-capacity 7,200 rpm drives to store archived records and data required for regulatory compliance.

Nearline 7,200 rpm, 500 GB Seagate Constellation SAS drives are designed to seamlessly intermix with 10,000 rpm, 300 GB Seagate Savvio® SAS drives and 15,000 rpm, 146 GB Seagate Savvio SAS drives. Because applications requiring a combination of capacity, speed, and data integrity can also be classified as tier 2, Constellation drives are purpose-built not only for high capacity, but also for high performance and data integrity at the system level. Enhanced command queuing can accommodate up to 16 host connections to help eliminate single points of failure, while dual-port capabilities help provide redundancy for failover and data recovery. SAS technology helps reduce data corruption by using end-to-end initiator-target nexus checking for data moving to and from the drive and provides advanced error-reporting features. SAS drives can also help ensure operational continuity by offering

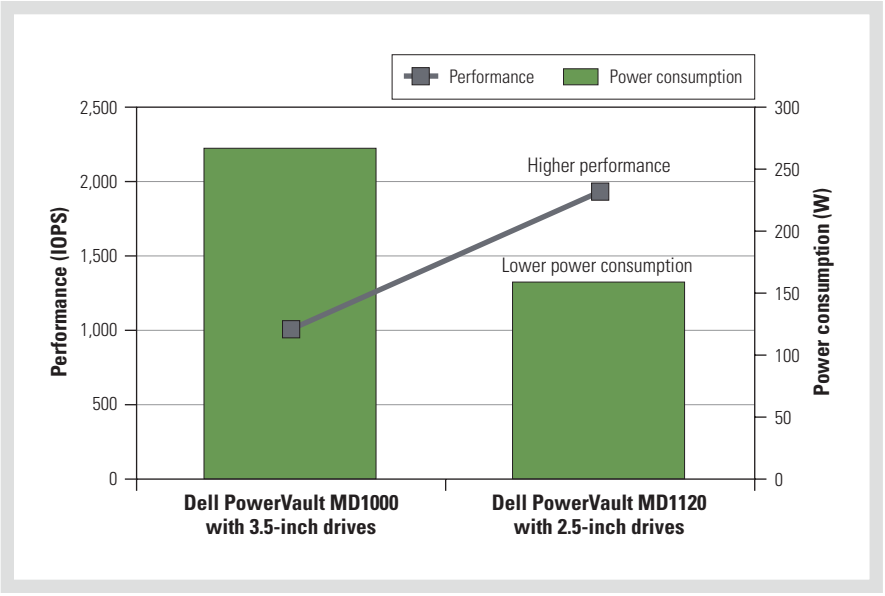

Figure 1. *SPC-1C performance and power consumption for Dell PowerVault MD1000 and PowerVault MD1120 DAS arrays*

seamless compatibility with existing SCSI software and middleware.

OPTIMIZING PERFORMANCE AND ENERGY USE

The 2.5-inch drive slots in the Dell PowerVault MD1120 DAS array can provide higher levels of performance density than typical arrays with 3.5-inch drive slots. To demonstrate the performance and efficiency of 2.5-inch Seagate Constellation drives in the PowerVault MD1120, in January and February 2009 Seagate engineers performed tests using the industry-standard Storage Performance Council (SPC) SPC-1C benchmark. The SPC-1C workload focuses on complex random-I/O performance and is a generally accepted representation of real-world conditions.

The tests compared a 2U PowerVault MD1120 enclosure containing its maximum configuration of twenty-four 2.5-inch, 7,200 rpm, 500 GB Seagate Constellation SAS drives with a 3U PowerVault MD1000 enclosure containing its maximum configuration of fifteen 3.5-inch, 7,200 rpm, 1 TB Seagate Barracuda® ES.2 drives. Each enclosure was connected to a Dell PERC 6/E in a Dell PowerEdge 2950 server configured with two Intel® Xeon® E5405 processors at 2.00 GHz, 2 GB of RAM, and the Microsoft® Windows Server® 2003 Standard Edition OS with Service Pack 2 (SP2). In these tests, the PowerVault MD1120 achieved a peak sustained maximum performance of 1,925 I/Os per second (IOPS). Compared with the PowerVault MD1000, the PowerVault MD1120 provided 93 percent higher performance overall, or 2.9 times the performance per unit of rack space—a major advantage in performance density.

In addition to offering higher performance, performance density, and space efficiency than a PowerVault MD1000 with 3.5-inch drives, a PowerVault MD1120 with 2.5-inch drives is also designed to consume significantly less power. The 2.5-inch, 7,200 rpm Constellation drives are designed to consume less than 3 W of power when idle, meaning that the 24 drive slots in an idle PowerVault MD1120 are designed to consume less power than a single 75 W light bulb. When tested by Seagate engineers using the draft SPC random workload power measurement method, the PowerVault MD1120 with these 2.5-inch, 500 GB SAS drives consumed 140 W when idle and 158 W while under the SPC-1C workload—40 percent less than the PowerVault MD1000 with 3.5-inch, 1 TB SAS drives while providing significantly higher levels of performance (see Figure 1).[1]

DEPLOYING EFFICIENT, COST-EFFECTIVE STORAGE

Tiered storage offers a key strategy to help control hardware costs and optimize storage investments for different types of data. By deploying the Dell PowerVault MD1120 DAS array with 2.5-inch Seagate Constellation hard drives, organizations can combine enterprise-class reliability, nearline performance, and scalability in a rack-dense, energy-efficient storage system.

Barbara Craig is a senior product marketing manager at Seagate. She attended the State University of New York, Pepperdine University, and Vanderbilt University.

Tom McCaffrey is a senior product manager at Seagate. He has a bachelor's degree from the University of Minnesota and a master's degree in Technology Management from the University of St. Thomas in Saint Paul, Minnesota.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Seagate Technology:
www.seagate.com

Dell PowerVault MD1120:
DELL.COM/MD1120

[1] For more information on the SPC-1C benchmark, as well as separate, SPC-audited test results using a configuration with twelve 2.5-inch Constellation drives in a PowerVault MD1120 enclosure (submission identifier C00009), visit www.storageperformance.org/results/benchmark_results_spc1c.

INTEGRATING DELL POWEREDGE M1000e BLADE ENCLOSURES WITH DELL EQUALLOGIC SANs

By Tony Ansley

Different approaches to integrating Dell™ PowerEdge™ M1000e modular blade enclosures with Dell EqualLogic™ PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays offer specific trade-offs in cost, performance, and complexity. By understanding the available options and following best practices for deployment, administrators can implement a consolidated, highly scalable infrastructure that can meet their specific needs.

Related Categories:

Blade servers

Cisco Systems

Dell EqualLogic storage

Dell PowerConnect switches

Dell PowerEdge blade servers

Storage

Storage area network (SAN)

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

Recently, many enterprises have begun adopting blade servers as part of an overall data center strategy focusing on consolidation, virtualization, and green IT initiatives. As part of this type of strategy, these enterprises may also consider parallel deployment of consolidated storage area network (SAN) arrays.

The Dell PowerEdge M1000e modular blade enclosure offers a number of advantages over typical rack-mounted servers, including increased density, rapid deployment, reduced power and cooling requirements, reduced cabling, simplified management, and an advanced integrated I/O module architecture. Dell EqualLogic PS Series Internet SCSI (iSCSI) SAN arrays, meanwhile, can also offer compelling benefits that make them well suited to support blade server deployments. These arrays combine an intelligent, automated management framework and a comprehensive set of enterprise data center services with a fault-tolerant hardware architecture designed to support all major non-mainframe operating systems. These arrays can provide a modular, cost-effective SAN solution designed for scalability, enabling small and midsize organizations to deploy storage in appropriate increments while still helping meet the needs of large enterprises requiring large capacities and high performance.

Successfully integrating a PowerEdge M1000e enclosure with an EqualLogic SAN requires advance planning, because as with other blade solution implementations, the available Ethernet I/O modules in the PowerEdge M1000e—the Cisco Catalyst Blade Switch (CBS) 3130G, Cisco CBS 3130X, and Dell PowerConnect™ M6220 switch—have a limited number of external "fan out" ports. These Ethernet I/O modules connect to networks outside the enclosure through four to eight Gigabit Ethernet ports (or up to four optional 10 Gigabit Ethernet ports) per switch. Because each EqualLogic array requires multiple Ethernet ports in an Ethernet-based SAN infrastructure, and the PowerEdge M1000e I/O modules provide limited numbers of external ports, administrators may need to use one or more external switches to connect PowerEdge M1000e enclosures to the EqualLogic SAN.

Administrators can take any of several approaches to combining PowerEdge M1000e enclosures with an EqualLogic SAN, each with its own trade-offs in cost, performance, complexity, and scalability. This article outlines three possible architectures and discusses the advantages and disadvantages of each architecture.

STRATEGY 1: MULTIPLE TIERS

The first option consists of a multitiered Ethernet switch infrastructure. In this design, the EqualLogic SAN arrays

and a set of external switches are on one tier, and the PowerEdge M1000e enclosures—with their integrated I/O modules and blade servers—are on a separate tier. These tiers are then connected using standard inter-switch linking technologies and link aggregation, allowing the processing of storage requests and data transfer between the SAN arrays and the blade servers. Figure 1 illustrates this architecture. For clarity, this figure shows only connections to primary array controllers; administrators should connect ports from the second array controller in each array to the storage tier switches in a similar fashion.

This design uses both the available stacking ports and the uplinking functionality of the I/O modules, along with external switches such as the Cisco Catalyst 3750-E or Dell PowerConnect 6248, to create a highly scalable SAN infrastructure. Stacking ports are used to link switches within each tier. To provide redundancy and bandwidth scalability, each tier is then connected by joining the I/O modules in the host tier to the external switches in the storage tier with uplinking ports configured into link aggregation groups.

This design strategy offers several advantages. In terms of scalability, the architecture enables administrators to easily add PowerEdge M1000e enclosures to the host tier by inserting the I/O modules for the new enclosures into the stacking loop without taking the EqualLogic SAN infrastructure offline. For example, because Cisco CBS 3130 switches allow up to eight switches to be combined in a single stack, up to eight PowerEdge M1000e enclosures can be integrated into the SAN. Similarly, using stackable Catalyst 3750-E switches in the storage tier enables administrators to add arrays to the existing switches (if ports are available) or add extra switches to the stack loop.

Attaching a full 12-array SAN requires as many as 48 ports per switch; including the ports required for the uplinks between the array tier and the host tier means that a 12-array SAN can be supported by two 48-port switches when using the optional dedicated 10 Gigabit Ethernet uplink ports.

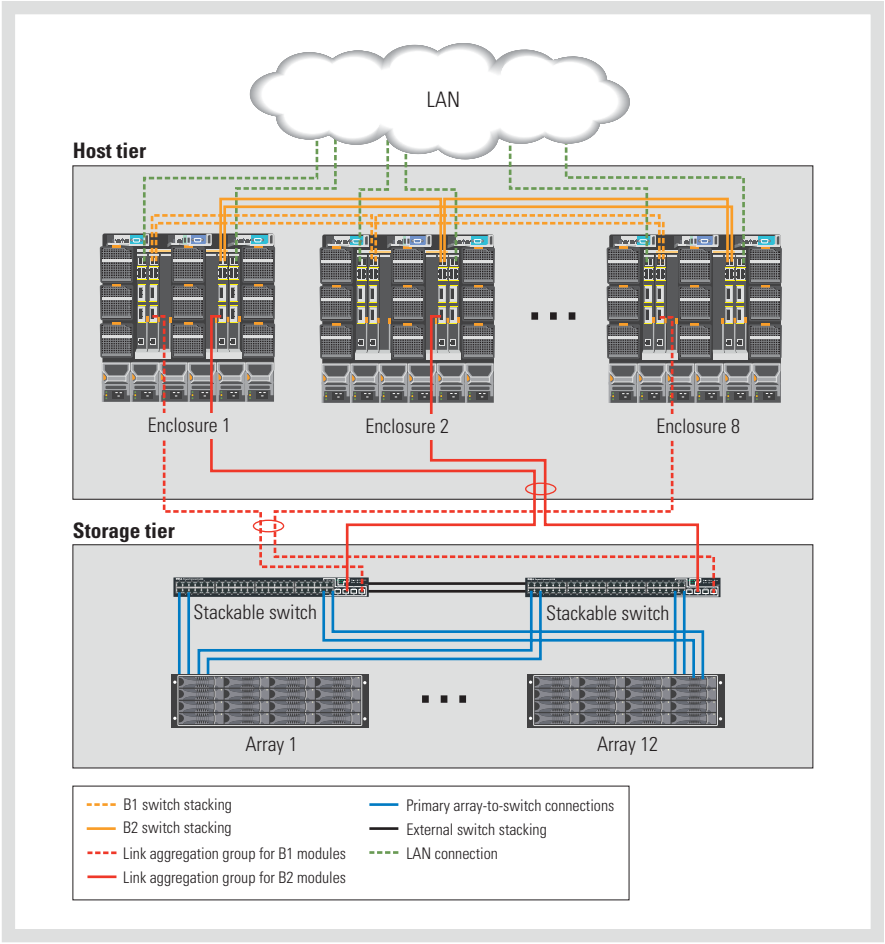

Figure 1. *Multitiered architecture*

However, using more than two switches may help increase performance for SANs with additional arrays by providing access to additional 10 Gigabit Ethernet ports. To connect the two tiers, administrators should aggregate the 10 Gigabit Ethernet ports for the external switches into a Port Aggregation Protocol (PAgP) or Link Aggregation Control Protocol (LACP) link aggregation group as shown in Figure 1, which helps maximize bandwidth and availability between tiers.

STRATEGY 2: SINGLE TIER WITH PASS-THROUGH I/O MODULES

A second option for integrating PowerEdge M1000e enclosures with EqualLogic SAN arrays is to use the PowerEdge M1000e Ethernet pass-through I/O modules with a set of Dell-recommended external stackable switches such as the Cisco Catalyst 3750-E or Dell PowerConnect 6248. Figure 2 illustrates this architecture. For clarity, this figure abstracts the connections between the enclosures and the switches; the actual connections from the enclosure pass-through ports should be evenly distributed between all external switches in the SAN to help provide full redundancy. As in the multitiered architecture, this figure shows only the connections to one array controller in each array; ports from the second array controller should be connected to the external switches in a similar fashion.

For organizations that do not primarily use Dell or Cisco switches, this option enables administrators to take advantage of existing third-party switches that support iSCSI—helping simplify management by allowing administrators to use the same management tools for SAN traffic as they

do for general-purpose network traffic. One disadvantage of this strategy, however, is the number of cables required to connect the individual blades to the external switch infrastructure, especially if the environment must be fully redundant and use two network interface cards (NICs) on each blade. As Figure 2 shows, this approach can scale to support almost any number of PowerEdge M1000e enclosures and EqualLogic arrays; however, it requires a large number of switch ports and Ethernet cables between each enclosure and the switch infrastructure. For example, a single PowerEdge M1000e enclosure would require 32 cables, and therefore require 32 ports on the SAN switch infrastructure in addition to the ports needed for the arrays.

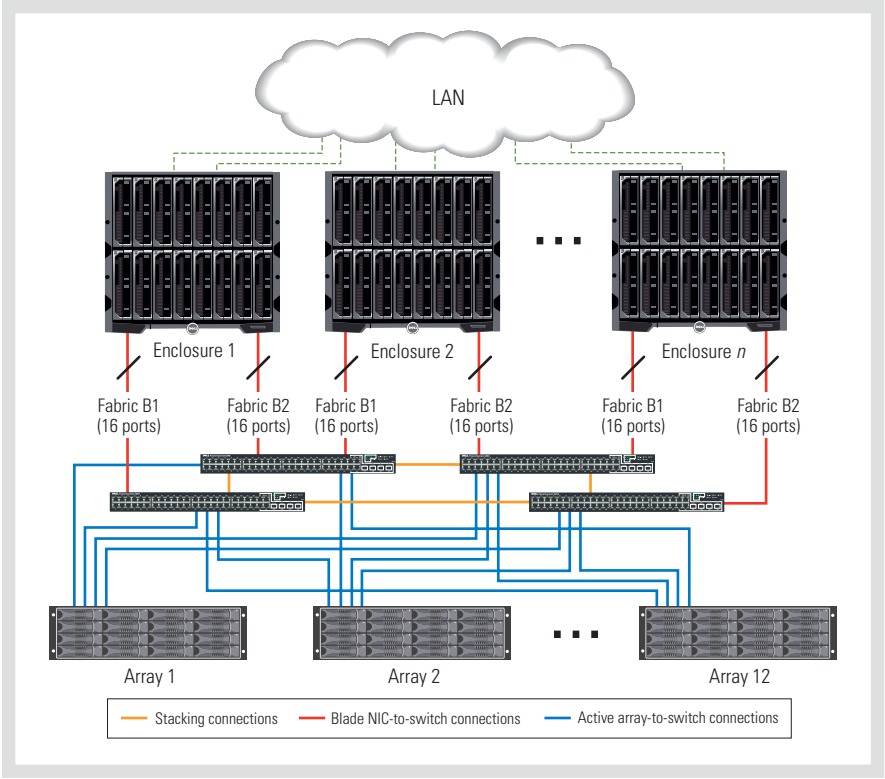

Figure 2. *Single-tiered architecture with pass-through I/O modules*

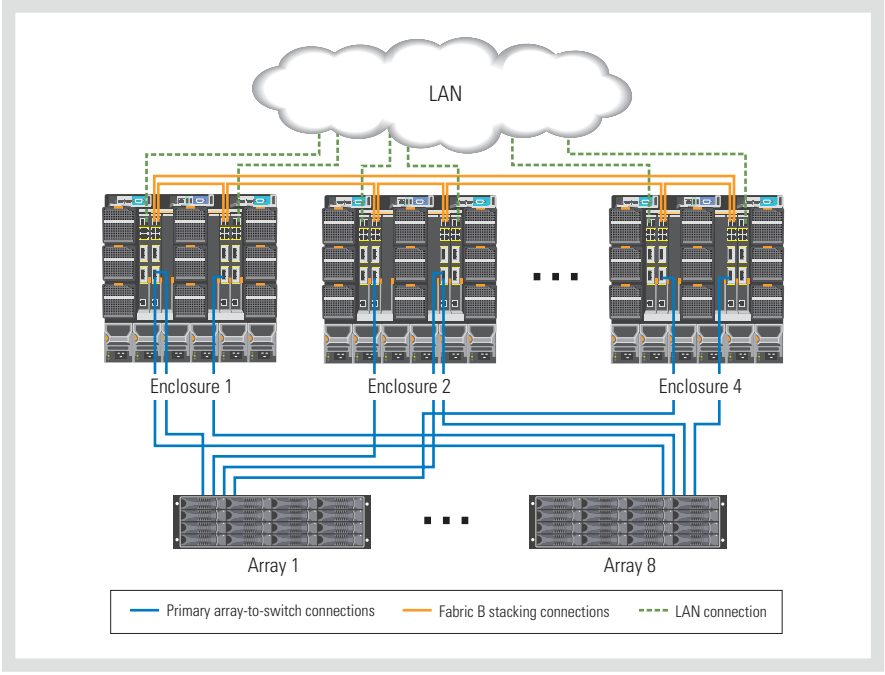

Figure 3. *Single-tiered architecture with stacking*

As with the multitiered approach, administrators can scale this architecture by expanding the external stackable switches as they add PowerEdge M1000e enclosures or EqualLogic arrays. Because of the large number of ports required for each PowerEdge M1000e, the SAN infrastructure typically requires more than two switches to support more than a few blade enclosures and a full 12-array SAN.

STRATEGY 3: SINGLE TIER WITH STACKING

The third option for integrating PowerEdge M1000e enclosures with EqualLogic SAN arrays relies only on the PowerEdge M1000e Ethernet switch I/O modules for the SAN. This architecture allows a single interface for network management similar to the architecture described in the preceding section, but is not as scalable as the other two architectures. Figure 3 illustrates this architecture. As with the other two architectures, this figure shows only the connections to one array controller in each array; ports from the second array controller should be connected to the array switches in a similar fashion.

Using Cisco CBS 3130X or CBS 3130G switches, this approach enables the SAN to scale to support a maximum of four PowerEdge M1000e enclosures and eight EqualLogic arrays, depending on which EqualLogic models are deployed. Because each array requires three or four ports per redundant switch and each CBS 3130 I/O module can have up to eight external Gigabit Ethernet ports available for array connection using Cisco TwinGig converter modules, each

enclosure can support up to two EqualLogic arrays. Administrators can use the switch stacking ports to integrate additional PowerEdge M1000e enclosures into the SAN; each additional PowerEdge M1000e enclosure would then enable them to add two EqualLogic arrays.

As Figure 3 shows, when adding PowerEdge M1000e enclosures, administrators should ensure that the connections from each EqualLogic array are redistributed so that each of the array's active Ethernet ports connect to at least two different blade switches. Figure 3 shows only the active port connections for each array; administrators should use an identical connection plan to connect each array's passive ports to the blade switches, helping ensure that a switch or enclosure failure does not break a connection to the EqualLogic arrays.

CONSOLIDATED, HIGHLY SCALABLE INFRASTRUCTURE

Each of the three strategies presented in this article has its own trade-offs in cost, performance, complexity, and scalability. Ultimately, successful integration of a Dell PowerEdge M1000e modular blade enclosure with a Dell EqualLogic iSCSI SAN hinges on a few criteria: the solution should provide full redundancy; enough inter-switch bandwidth to support hosting SAN traffic, inter-array management, and load balancing; and enough I/O and minimal latency between the blades and arrays to meet the requirements of the attached host applications. The PowerEdge M1000e enclosure with Cisco CBS 3130 or Dell PowerConnect M6220 I/O modules and Ethernet pass-through modules along with external Ethernet switches such as the Cisco Catalyst 3750-E and Dell PowerConnect 6248 can provide the flexibility and scalability necessary to build a high-performance virtualized storage architecture that can meet almost any storage need.

Tony Ansley is a senior storage consultant on the Dell iSCSI Solutions Marketing team with 24 years of experience in the computer industry. He has a bachelor's degree in Information and Computer Sciences from the Georgia Institute of Technology.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell PowerEdge blade servers:
DELL.COM/Blades

Dell EqualLogic PS Series:
DELL.COM/EqualLogic
DELL.COM/PSSeries

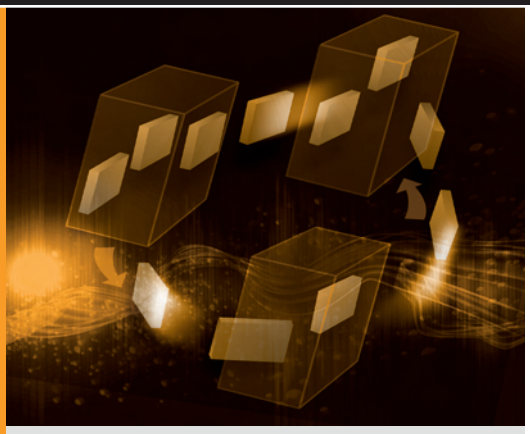

ADVANCING APPLICATION AVAILABILITY WITH PAN MANAGER SOFTWARE BY EGENERA FOR DELL

By Christine Crandell

As application services become increasingly interdependent, how can IT departments help ensure continuous availability without creating a costly and complex infrastructure? PAN Manager® software by Egenera® for Dell—part of the Dell™ PAN System for blades—enables high availability and disaster recovery for applications running on Dell PowerEdge™ blade servers.

Related Categories:

Blade servers

Business continuity

Dell PAN System

Dell PowerEdge blade servers

Disaster recovery

Egenera

High availability (HA)

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

Once reserved for only the most mission-critical services, high availability and disaster recovery have become essential requirements across the spectrum of enterprise applications. Historically, these capabilities were available as integrated services built into RISC and mainframe systems. Then, as x86-based servers came into prominence, vendors began offering supplemental high-availability and disaster recovery products in the form of layered software and multiple standby servers for each application.

Meanwhile, blade servers have brought RISC- and mainframe-class reliability and serviceability features to x86 server architectures. However, many blade systems still require additional software and dedicated backup servers to help maintain application availability. This approach, while workable, often contributes to decreased performance and to delays in bringing applications back online. In addition, time-consuming and labor-intensive operations are necessary to keep system configurations synchronized, thus adding further cost, complexity, and risk to IT operations. Consequently, IT managers remain hesitant to migrate high-end applications to blade servers, opting instead to retain their RISC and mainframe systems for those selected applications considered essential to business continuity.

PAN Manager software by Egenera for Dell—part of the Dell PAN System for blades—helps deliver mainframe-class high availability and verifiable disaster recovery to applications running on Dell PowerEdge blade servers.[1] PAN Manager is a factory-integrated, noninvasive solution that helps avoid the need for IT administrators to install separate software products, accept degraded system performance, and sustain cumbersome management practices to help ensure business continuity, and is designed to allow organizations to standardize their computing environment while complying with strict enterprise availability requirements.

INTRODUCING THE DELL PAN SYSTEM FOR BLADES

The Dell PAN System for blades is based on a Dell PowerEdge M1000e modular blade enclosure containing up to 16 PowerEdge M600 or PowerEdge M610 blade servers (see Figure 1). The PAN Manager software runs on two PowerEdge 2950 III servers

[1] For information on PAN Manager features and its support for Dell PowerEdge rack servers, see "Dell and Egenera Drive a New Path to Virtual Data Center Automation," by Greg Lyon, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090120-Egenera.pdf.

that act as redundant, load-balanced processing area network (PAN) controllers. The blade servers and controllers are designed to communicate over a secure, three-way redundant fabric using Dell PowerConnect™ M6220 Ethernet switches. Together, the servers, controllers, switches, and PAN Manager software constitute the Dell PAN System.

PAN Manager stores each server's configuration information as a server definition file in a storage area network (SAN)–resident database. Server definitions include application and OS images, processor and memory requirements, and other information typically associated with server hardware, such as the IP address, Media Access Control (MAC) addresses, Fibre Channel World Wide Names (WWNs), and storage logical units (LUNs). This approach enables PAN Manager to quickly and automatically move applications from one server to another, helping eliminate the need for physical reconfiguration or other manual intervention.

MAINTAINING RESPONSIVENESS THROUGH HIGH AVAILABILITY

PAN Server Portability™ software by Egenera for Dell, a component of PAN Manager, is designed to provide *n* + 1 high availability for applications without incurring the cost, overhead, and IT complexity of many alternative approaches. With PAN Server Portability, PAN Manager can automatically invoke application recovery operations in response to a variety of conditions such as resource overutilization, communication failure, hardware failure, and unrecoverable software errors. Monitoring, decision-making, and control operations take place on the redundant PAN controllers, thus enabling successful recovery while helping eliminate performance impact on the blade servers.

Using PAN Manager, IT administrators can create policies that specify the conditions and severity thresholds that should trigger the appropriate recovery operations, such as restarting an application, adding a new application instance to a cluster, or initiating application failover to a backup server. For example, in an environment with a cluster configured to support several instances of a Web application spanning multiple servers, administrators might create a policy that is triggered when user response time increases to a specified level. If this threshold is reached, PAN Manager would automatically start a new application instance on a designated blade server. The cluster software detects the new instance and begins directing incoming transactions to it, thereby enhancing service availability.

Administrators might also specify an escalating set of actions that PAN Manager should take once a server's processor utilization reaches a specific threshold, or when a hardware error is detected. In these cases, the policy might indicate that PAN Manager should first attempt to restart the application. If the restart action does not resolve the problem, the policy could then specify a failover operation, causing PAN Manager to transition the application to an alternate server. If necessary, PAN Manager can fail over an entire virtualized server running multiple virtual machines.

Although PAN Manager is designed to move any application to any server in the Dell PAN System, administrators typically identify one server to function as backup for all of the applications running on the system. This backup server can remain in a passive state, awaiting a failover, or can actively run nonessential workloads. In the latter case, PAN Manager first shuts down the active workloads and then automatically boots the appropriate OS image on the failover server before restarting the application. This *n* + 1 approach to high availability helps ensure efficient, flexible, and cost-effective utilization of computing resources.

PAN Manager is also designed to simplify high-availability management in several additional ways. First, using just a few mouse clicks, administrators can specify policies and attach those policies to a server definition. Second, PAN Manager can automatically monitor system health and initiate recovery actions without requiring administrator intervention, thus helping to minimize application downtime. Third, PAN Manager provides customization features that allow administrators to tailor monitoring criteria. Finally, PAN Manager is designed to ensure that failover operations work every time, because of its redundant architecture and because

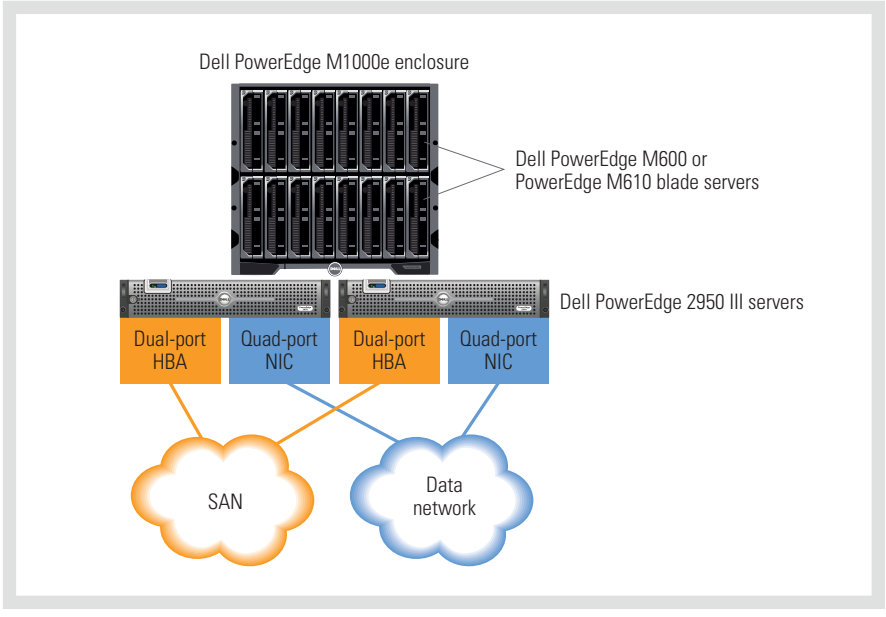

Figure 1. *The Dell PAN System for blades combines PAN Manager software with efficient, high-performance Dell PowerEdge blade servers*

the configuration and software information necessary to restart an OS and application is contained within the server definition files—thereby rendering unnecessary the painstaking and error-prone measures usually required to both keep backup systems synchronized and perform a failover operation.

SUPPORTING RELIABLE DISASTER RECOVERY

PAN Portability™ software by Egenera for Dell, another component of PAN Manager, enables verifiable disaster recovery protection for production data centers. PAN Manager is designed to ensure that disaster recovery operations replicate the production environment identically, every time, even in data centers where there are frequent modifications to system configurations. In addition, PAN Manager can transition a complete PAN System automatically, without requiring complex manual operations found in typical runbook-based disaster recovery plans.

PAN Manager maintains the hardware and software configuration for each PAN System in a database called the PAN Archive. PAN Manager automatically replicates this archive, thus helping ensure that the disaster recovery site is always synchronized with the production site. Each PAN Archive database contains individual server configuration files as well as information global to the PAN System, such as hardware addresses, log files, and administrative security settings.

Once an administrator initiates a disaster recovery operation, PAN Manager automatically merges the hardware, network, and SAN mappings from the disaster recovery system into the PAN Archive. PAN Manager then installs the new archive file on the PAN System and begins the process of booting individual servers and starting applications. No manual intervention is necessary beyond initiating the operation. The entire recovery process is designed to complete, error free, in a matter of minutes—typically well within the time frames of stringent recovery time objectives.

This design enables administrators to create disaster recovery plans based on an *n* + 1 model, allowing one disaster recovery site to support multiple production sites. This capability can contribute not only to low hardware and software capital costs, but also to low operational costs associated with updating and testing multiple disaster recovery sites. In addition, a PAN System designated for disaster recovery may be used for non-production functions such as development or quality-assurance testing. PAN Manager automatically shuts down active workloads running on the system and saves the local PAN Archive before starting the recovery process.

ENHANCING BLADE SERVER RESILIENCY

The advanced PAN Manager application availability features are based on a resilient and comprehensive system foundation. The PAN controllers perform the internal and external I/O operations on behalf of the blade servers, and are designed to provide high-performance, secure, and reliable communication switching among software applications running within the blade system. To support communication outside the blade system, each PAN controller connects to the SAN through a dual-port Fibre Channel host bus adapter (HBA) and to the data network through a quad-port network interface card (NIC).

The PAN controllers are designed to share all of the system's I/O operations using advanced load-balancing techniques. In addition, both of the controllers have access to the configuration and management database for the entire blade system. If one controller fails, the other controller can automatically detect the failure and assume the management and I/O workload. After the failed controller has been replaced, the load-balancing procedures can automatically resume.

Finally, PAN Manager allows administrators to manage both physical and virtual servers using the same tools and procedures. In addition, PAN Manager supports industry-standard server virtualization products, including the VMware® Infrastructure 3 and Citrix® XenServer™ platforms. These features help administrators ensure the same levels of resiliency for data centers spanning physical and virtual environments.

RECONSIDERING BUSINESS CONTINUITY SOLUTIONS

The application services that support today's dynamic business processes have become progressively interdependent, blurring the lines that previously separated traditional continuity categories of mission critical, business critical, and edge. Consequently, today's data centers require robust continuity plans designed to ensure availability for all applications. PAN Manager enhances the hardware reliability and serviceability of blade systems by enabling mainframe-class high availability and verifiable disaster recovery for all applications—helping to automatically recover lost production environments accurately and reliably and to restore productivity rapidly even in the event of a total site outage.

Christine Crandell is the executive vice president of global marketing and chief marketing officer at Egenera. Christine has a B.S. in Business Administration and an M.B.A. from Florida Atlantic University, and has completed doctoral studies at Golden Gate University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

PAN Manager Software by Egenera for Dell:
www.egenera.com/dell

Dell PowerEdge blade servers:
DELL.COM/Blades

SOLARIS™ ON DELL™ POWEREDGE™ M-SERIES SERVERS

OFFERS UP TO 47% MORE PERFORMANCE PER WATT.

Dell is committed to helping our customers run IT infrastructures more efficiently, while lowering operating costs. By running Solaris on Dell M-Series Blades, customers can achieve up to 38% better performance/watt than HP ProLiant® BL685c G5 and 47% better performance/watt than IBM BladeCenter® LS42.* Get more out of Your IT infrastructure with Solaris on Dell.

SIMPLIFY OPERATING SYSTEMS AT DELL.COM/Solaris

solaris™

* Based on Principled Technologies report commissioned by Dell, "SPECjbb2005 performance and power consumption on Sun Solaris 10 multi-processor AMD-based blade servers from Dell, HP, and IBM" January 2009.

DELL DELIVERS PERFORMANCE AND ENERGY EFFICIENCY OVER HP AND IBM BLADE SERVERS

Dell™ PowerEdge™ blade servers are designed for high performance and efficient energy use. In Principled Technologies® blade server tests, the Dell PowerEdge M905 provided both higher peak performance and lower power consumption than the HP ProLiant BL685c G5 or IBM® BladeCenter LS42—delivering significantly higher performance per watt to help optimize data center efficiency.

Related Categories:

AMD Opteron processors

Benchmarks

Blade servers

Dell PowerEdge blade servers

Energy efficiency

Green IT

Performance

Performance characterization

Power and cooling

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

When evaluating a server, performance and power consumption are two key factors that organizations should take into account. Perhaps most important, however, is the efficiency with which the server can translate energy into useful work performed in the data center—a factor captured by the key metric of *performance per watt*. Maximizing performance per watt can lead to significant improvements in IT productivity without increasing power consumption.[1]

In January 2009, Principled Technologies performed tests commissioned by Dell to compare the performance and power consumption of three AMD Opteron™ processor–based blade servers: the Dell PowerEdge M905, HP ProLiant BL685c G5, and IBM BladeCenter LS42.[2] As the results show, the Dell server not only achieved the highest peak performance of the three tested servers, but also did so while using the least amount of power—providing significantly higher performance per watt than the HP and IBM servers to help optimize data center efficiency.

TEST ENVIRONMENT AND METHODOLOGY

The Dell PowerEdge M905 was housed in a Dell PowerEdge M1000e enclosure, the HP ProLiant BL685c G5 was housed in an HP BladeSystem c-Class enclosure, and the IBM BladeCenter LS42 was housed in an IBM BladeCenter H (8852) enclosure. Each blade server was configured with four quad-core AMD Opteron 8356 processors at 2.3 GHz and 32 GB of RAM, and ran the Sun Solaris 10 OS.

The tests were based on version 1.07 of the Standard Performance Evaluation Corporation (SPEC) SPECjbb2005 benchmark. SPECjbb2005 utilizes multiple special data groups and multiple threads as it runs, with each data unit referred to as a *warehouse*—a roughly 25 MB collection of data objects. Each thread represents an active user posting transaction requests within a warehouse. The benchmark run begins with one warehouse and then increases the number of warehouses until it has saturated the server's processor capacity. As the number of warehouses increases, so does the number of threads. The benchmark's results portray

[1] For more information on how best practices for energy efficiency can help increase IT productivity, see "Compute More, Consume Less: Smart Policies Unleash Data Center Productivity," in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090176-Esser.pdf.

[2] See "SPECjbb2005 Performance and Power Consumption on Sun Solaris 10 Multi-processor AMD-Based Blade Servers from Dell, HP, and IBM," by Principled Technologies, January 2009, DELL.COM/Downloads/Global/Solutions/DellM905SPECjbb0109.pdf.

	IBM BladeCenter LS42	HP ProLiant BL685c G5	Dell PowerEdge M905
JVM 1	90,957 bops	94,186 bops	101,453 bops
JVM 2	92,262 bops	94,081 bops	102,010 bops
JVM 3	91,671 bops	92,834 bops	92,936 bops
JVM 4	93,223 bops	91,250 bops	91,974 bops

Figure 1. *Median SPECjbb2005 results for each Java Virtual Machine on the tested blade servers*

the server's throughput in business operations per second (bops).[3]

SPECjbb2005 requires a Java Virtual Machine (JVM) on the test system. The test team used the Sun Java Platform, Standard Edition (SE) 6 Runtime Environment (build 1.6.0_06-p-b04) and Java HotSpot Server VM (build 13.0-b04, mixed mode) with the default installation settings. After installing these components, editing the SPECjbb_config.props file to include disclosure and license information, and editing the SPECjbb.props file to allow the servers to each run four JVM instances, the test team created a shell script to issue the Java run command to launch the benchmark.

The team used the Solaris `psrinfo -pv` command to retrieve system processor information before beginning the testing, using the processor assignments to set the processor affinity appropriately on each system. The team then set the `psrset` command to assign groupings of cores to JVMs. The core numbers differed among the systems, so the `psrset -c` commands differed slightly among the three blade servers; the test team configured these settings to provide maximum performance on all servers.

The shell script also included several Java options that controlled JVM performance:

- `-Xms3350m`: Set the minimum heap size to 3,350 MB
- `-Xmx3350m`: Set the maximum heap size to 3,350 MB (because the minimum and maximum heap sizes were the same in these tests, the heap size stayed at a constant 3,350 MB)
- `-Xmn2600m`: Set the JVM nursery size to 2,600 MB
- `-XX:+UseParallelOldGC`: Set Java to use parallel garbage collection
- `-XX:+AggressiveOpts`: Enabled performance compiler optimizations
- `-XX:AllocatePrefetchStyle=2`: Set code-style prefetch instructions
- `-XX:ParallelGCThreads=2`: Set the number of GC threads to two

Server power consumption during each test was recorded using an Extech Instruments 380803 power analyzer and data logger. Before starting the benchmark, the test team logged into the servers and allowed them to sit idle for 8 minutes. To gauge idle power consumption, the test team then recorded server power consumption for 2 minutes while the server was running only the Solaris 10 OS. This process meant that each system was idle for 10 minutes before the tests began. After the tests began, the team recorded server power consumption at one-second intervals. Average power consumption was calculated by averaging power consumption during the time the server was producing its peak performance results.

TEST RESULTS: PERFORMANCE AND EFFICIENCY

During the tests, each server ran four JVMs simultaneously. To compute the overall score for each server, SPECjbb2005 sums the scores of its JVMs; the scores of each JVM are calculated by taking the average of the results when the server is running at peak performance. Figure 1 shows the median SPECjbb2005 results for each JVM across three runs of the SPECjbb2005 benchmark, while Figure 2 illustrates the total median SPECjbb2005 results for each blade server. A higher SPECjbb2005 score indicates that the server handled more Java requests and thus delivered greater throughput.

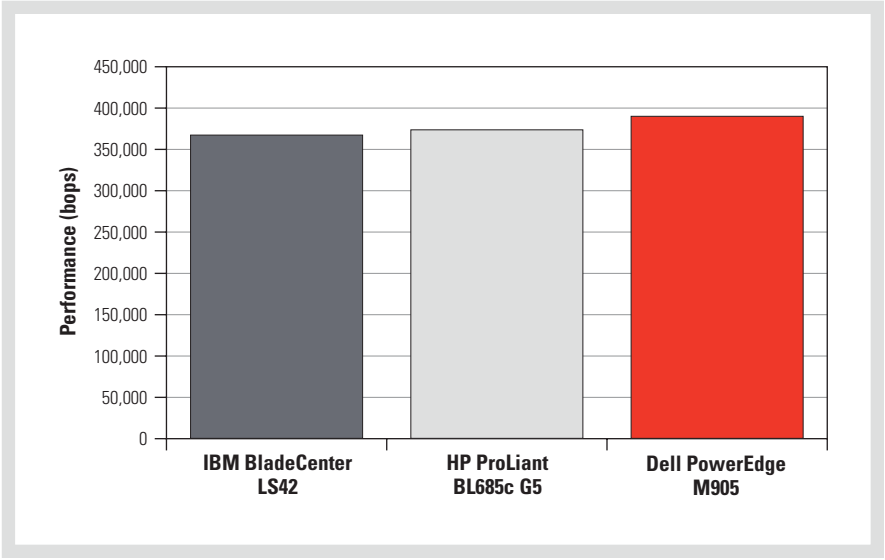

Figure 2. *Total median SPECjbb2005 results for the tested blade servers*

[3] For more information on SPECjbb2005, visit www.spec.org/jbb2005.

As these figures show, the Dell PowerEdge M905 blade server produced the highest peak performance of the three systems tested, handling a total median of 388,373 bops. This represented 4.3 percent higher performance than the HP ProLiant BL685c G5 (which achieved a total median performance of 372,351 bops) and 5.5 percent higher performance than the IBM BladeCenter LS42 (which achieved a total median performance of 368,113 bops).

Figure 3 shows the average power consumption for each server when idle and during peak performance. Figure 4 shows the performance-per-watt results, which have been normalized to show relative increase over the server with the lowest result, the IBM BladeCenter LS42. Performance per watt was calculated by dividing the total median peak SPECjbb2005 performance by the average power consumption in watts when the server was delivering peak performance during the median run. As these figures show, the Dell PowerEdge M905 blade server provided not only the highest overall performance of the three tested servers, but also the lowest power consumption—delivering 37.6 percent higher performance per watt than the HP ProLiant BL685c G5 and 46.7 percent higher performance per watt than the IBM BladeCenter LS42.

> **"As the results show, the Dell server not only achieved the highest peak performance of the three tested servers, but also did so while using the least amount of power—providing significantly higher performance per watt than the HP and IBM servers to help optimize data center efficiency."**

	IBM BladeCenter LS42	HP ProLiant BL685c G5	Dell PowerEdge M905
Idle	856.2 W	707.1 W	484.1 W
During median benchmark run	995.9 W	945.0 W	716.2 W

Figure 3. *Average power consumption for the tested blade servers*

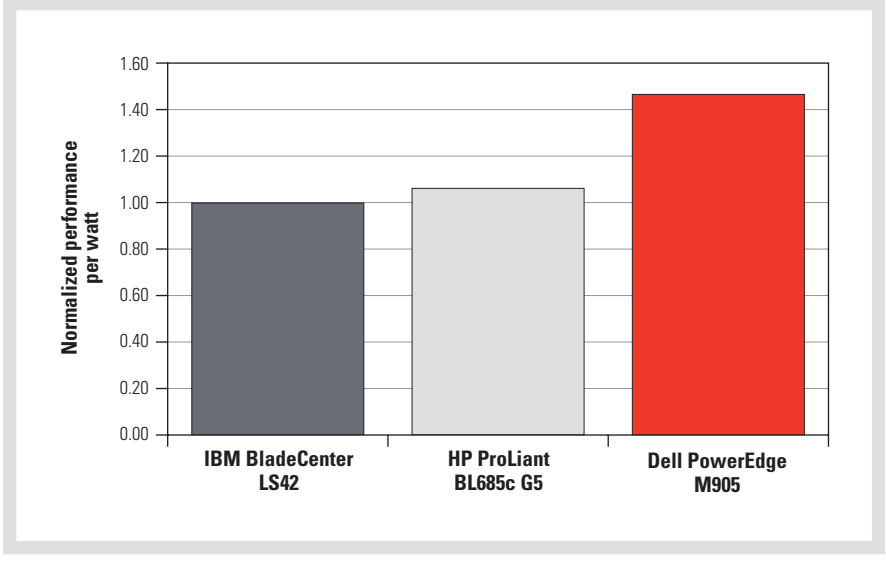

Figure 4. *Normalized performance per watt for the tested blade servers*

POWERFUL, ENERGY-EFFICIENT BLADE SERVERS

Optimizing energy use can be critical to controlling operational costs in enterprise data centers. As these Principled Technologies tests demonstrate, Dell PowerEdge blade servers are designed to offer high performance while minimizing power consumption—providing the foundation of an effective and comprehensive strategy for data center efficiency.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Full Principled Technologies test report:
DELL.COM/Downloads/Global/Solutions/DellM905SPECjbb0109.pdf

Dell PowerEdge blade servers:
DELL.COM/Blades

Dell Solaris solutions:
DELL.COM/Solaris

SIMPLIFY YOUR PATH TO VIRTUALIZATION

Proactively monitor and manage your entire IT infrastructure through a common interface.

Use Microsoft® System Center Virtual Machine Manager 2008 and System Center Operations Manager plus Dell™ Management Packs and the NEW Dell PRO-Pack to manage:

- Servers
- Chassis/Blades
- Storage
- Business clients
- Virtual machines
- Multi-vendor hypervisors including Microsoft Hyper-V™ technology and VMware® ESX

Let Dell and Microsoft help you streamline the virtualization process.

Microsoft® System Center

SIMPLIFY YOUR OPERATIONS AT
DELL.COM/SystemCenter

DELL