



POWER SOLUTIONS

DELL.COM/PowerSolutions

NOVEMBER 2008 • \$4.95



A SMART PATH TO VIRTUALIZATION

Getting Started with Microsoft
Windows Server 2008 Hyper-V

Virtualizing the Desktop with
Dell Flexible Computing Solutions

Powering Business Productivity
with New Dell OptiPlex Desktops

**SPECIAL HIGH-PERFORMANCE
COMPUTING SECTION
SEE PAGE 58**

**GLOBAL SUCCESSES IN
BLADE SERVER DEPLOYMENTS
SEE PAGE 88**

Read about the newest
Intel® 10 Gigabit Server Adapter
on page 54.



Get the most out of your virtualized servers.

Intel® 10 Gigabit Ethernet Server Adapters Featuring Virtual Machine Device Queues (VMDq)

Dell PowerEdge® Servers and Intel® 10 Gigabit Ethernet Server Adapters with VMDq technology can provide greater flexibility and increased utilization to help your virtualized server platforms perform at their best.

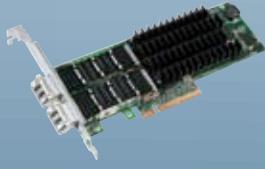
VMDq is another breakthrough technology from Intel that helps offload network I/O data processing from the hypervisor to the network silicon, improving I/O throughput for faster, more efficient networking. In a benchmark study, VMDq features and functionality more than doubled the throughput on a virtualized platform.

Learn more:

intel.com/go/10GbE and intel.com/go/vtc
www.dell.com/networking



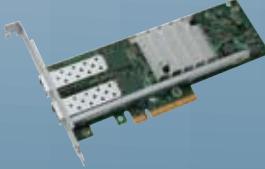
Intel® 10 Gigabit XF SR
Server Adapter



Intel® 10 Gigabit XF SR
Dual Port Server Adapter



Intel® 10 Gigabit AT
Server Adapter



Intel® 10 Gigabit AF DA
Dual Port Server Adapter



POWER SOLUTIONS

NOVEMBER 2008

DELL.COM/PowerSolutions

8

COVER STORY

A SMART PATH TO VIRTUALIZATION

By Jeanne Feldkamp, Deb McDonald, Kay Kerr,
and Tom Kolnowski

As server virtualization continues to transform the data center, enterprises have some critical decisions to make. How will they migrate to a fully virtualized environment? How should the IT infrastructure be configured for maximum business benefit? Dell is partnering with other industry leaders to build a flexible, industry-standard approach that helps simplify the transition.

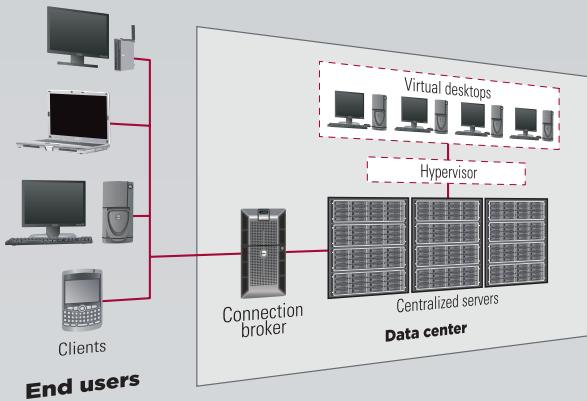
16

BEST PRACTICES

GETTING STARTED WITH MICROSOFT WINDOWS SERVER 2008 HYPER-V ON DELL SERVERS

By Ranjith Purush, Sitha Bhagvat, Ryan Weldon,
Brent Douglas, and David Schmidt

Microsoft® Hyper-V™ technology provides a simplified virtualization platform integrated directly into the Microsoft Windows Server® 2008 OS. This article provides an introduction to Hyper-V virtualization, discusses the overall architecture and underlying technologies, and offers guidance on best practices for deployment on Dell™ PowerEdge™ servers.



FLEXIBLE COMPUTING PAGE 26

ADVANCING END-USER PRODUCTIVITY WITH CENTRALIZED CONTROL

By Jeremy Ford and Roberto Ayala

Supporting an increasingly distributed workforce can be a huge drain on IT resources. By enabling a customized end-user experience with server-based control of data, settings, and applications, the Dell Flexible Computing Solutions suite of products and services is designed to simplify management, enhance security, and lower total cost of ownership.

GO GREEN WITH DELL POWER SOLUTIONS DIGITAL EDITION

Experience your information, your way—now online in eco-friendly, searchable magazine format. *Dell Power Solutions* Digital Edition offers a comprehensive IT library at your fingertips with maximum mobility: bookmark, forward, and print articles; zoom in and out on any part of the publication; and link directly to related resources. For instant accessibility, visit powersolutionsdigital.dell.com.

CONTENTS

EDITOR'S COMMENTS

6 Get Green: Go Digital

By Tom Kolnowski

FEATURE SECTION

VIRTUALIZATION AND FLEXIBLE COMPUTING

22 Highly Available Virtualization with Microsoft Hyper-V and SCVMM 2008

By Ray Weinstein and Burk Buechler

Combining Microsoft Hyper-V technology with Microsoft System Center Virtual Machine Manager (SCVMM) 2008 on Dell platforms can help organizations create highly available virtualized systems while enabling simplified management, dynamic performance, workload optimization, and reduced costs.

34 Dell Precision R5400 Rack Workstations Offer Powerful Performance and Secure Remote Access

By Richard Thwaites

The Dell Precision™ R5400 rack workstation and optional Dell FX100 remote access device enable secure access to applications and data with outstanding performance for an increasingly mobile workforce.

38 Extending Benefits of Virtual Remote Desktops Using VMware and Dell EqualLogic SANs

By Timothy Sherbak and Cris Benson

Virtual Remote Desktop solutions based on VMware® Virtual Desktop Infrastructure and virtualized Dell EqualLogic™ Internet SCSI (iSCSI) storage area networks (SANs) can help reduce the traditional cost and complexity of managing physical laptops, desktops, and workstations.

42 How Dell EqualLogic Auto-Snapshot Manager/VMware Edition Helps Protect Virtual Environments

By Andrew Gilman and William Urban

Dell EqualLogic Auto-Snapshot Manager/VMware Edition enables administrators to quickly and easily create hypervisor-aware snapshots of virtual machines—helping simplify data management, enhance scalability of data protection and recovery, and increase application performance.

SPECIAL SECTION: BLADE SERVERS

84 Introducing the Dell PowerEdge M805 and PowerEdge M905 Blade Servers

By Thomas Cloyd and Romy Bauer

Combining robust processing power, tremendous RAM capacity, and massive I/O throughput in a space- and energy-efficient, easy-to-deploy blade form factor, Dell PowerEdge M805 and PowerEdge M905 full-height blade servers help increase performance and reliability while helping simplify management.



46 Dell and Egenera Drive a New Path to Virtual Data Center Automation

By Greg Lyon

The new Dell/PAN System combines Dell PowerEdge server technology with advanced Egenera® management software in a simplified, customizable system—providing a highly available, dynamic, flexible platform for next-generation data center virtualization.

50 Best Practices for Protecting VMware Infrastructure 3 with Symantec Backup Exec

Symantec® Backup Exec™ 12.5 software, part of the new Dell PowerVault™ DL2000 – Powered by Symantec Backup Exec, is designed to provide powerful, simplified backup and recovery for VMware virtualized environments.

54 Boosting Data Center Throughput with Virtualization-Ready 10 Gigabit Ethernet

By Shefali Chinni, Omar Sultan, and Nelson Stewart

To help support the I/O requirements of multiple virtual machines in server consolidation and virtualization scenarios, Dell, Intel, and Cisco have teamed up to offer organizations a virtualization-tuned 10 Gigabit Ethernet SFP+ direct attach solution.

CASE STUDIES

88 The Business of Success

Dell PowerEdge blade servers and Dell EqualLogic iSCSI SAN arrays help the Wharton School of the University of Pennsylvania quintuple computing power while simplifying management and reducing costs.

90 Healthy IT Infrastructure

Dell PowerEdge blade and rack servers provide the foundation of a high-density virtualized infrastructure at ALERT Life Sciences Computing—supporting simplified management, high performance, and rapid deployment.

92 Blades Shift Broker into High Gear

Dell PowerEdge blade servers help PT Ciptadana Capital support future expansion of its security trading service while reducing energy use, simplifying management, and optimizing limited data center space.

94 Mortgage Lenders' Moment of Truth

A high-performance, energy-efficient solution based on Dell PowerEdge blade servers and a Dell EqualLogic iSCSI SAN array provides Optimal Blue with an anticipated 50 percent improvement in application response time.

96 Re-creating the Big Bang

Powerful Dell PowerEdge blade servers deployed in just nine days provide high levels of performance and efficiency to help put the Deutsches Elektronen-Synchrotron (DESY) research center on track for a major scientific breakthrough.

PRODUCT SHOWCASE: DELL OPTIPLEX DESKTOPS PAGE 30

POWERING BUSINESS PRODUCTIVITY WITH DELL OPTIPLEX DESKTOPS

By David Schweighofer

The new Dell OptiPlex™ family of desktop computers is designed to deliver high performance, flexibility, and reliability together with enterprise-class security and enhanced manageability in an environmentally conscious way—helping organizations to maximize employee productivity, simplify management, and lower total cost of ownership.



Dell OptiPlex 960 desktop

▶ SEE IT HERE FIRST

DELL POWER SOLUTIONS ONLINE EXCLUSIVES



Dell TechCenter: Collaborative Resources for Microsoft Hyper-V, Virtualization, and More

By Todd Muirhead

The Dell Enterprise Technology Center (TechCenter) provides a central collaborative environment offering a range of IT resources, from blogs, chats, wiki pages, and discussions to white papers, video demos, and links—now featuring virtualization, Microsoft Hyper-V technology, and other timely topics.



Virtualizing Microsoft SQL Server 2005 on Dell PowerEdge Servers

By Todd Muirhead

To demonstrate how the Microsoft SQL Server® 2005 database platform can perform in a VMware virtualized environment, Dell engineers tested its performance when scaling virtual resources and when migrating virtual machines between physical hosts.



Rethinking the Dell PERC 6 RAID Controller

By André Dumouchelle, Joe H. Trickey III, and John Seward

The Dell PowerEdge Expandable RAID Controller (PERC) 6.1 firmware release introduces multiple features designed to enhance reliability, performance, and ease of use to help organizations meet the challenges of ongoing data growth.



An Overview of Microsoft Windows Small Business Server 2008 on Dell Servers

By Perumal Raja P. and Manjunath Narayanan

Designed specifically to help meet the needs of small businesses, the Microsoft Windows® Small Business Server 2008 platform provides an all-in-one server solution offering cost efficiency, simplified data security, and enhanced productivity.

DEPARTMENTS

SPECIAL SECTION: HIGH-PERFORMANCE COMPUTING

58 **Running Microsoft Windows HPC Server 2008 on Dell PowerEdge Clusters**

By Aziz Gulbeden and Mausmi Kotecha

62 **Deploying HPC Clusters Using Platform Open Cluster Stack – Dell Edition**

By Seeram Vedantham; Arun Rajan; Shivaraj Nidoni; Li Ou, Ph.D.; Gabor Samu; and Mark Black

68 **Accelerated HPC Productivity with Intel Cluster Ready Solutions**

By Brock A. Taylor, Vishvesh Sahasrabudhe, and Onur Celebioglu



Streamlined Troubleshooting with the Dell System E-support Tool

By Shruthi Jayatheertha, Lokesh S. Balu, and Ajaykumar C. H.

By automating and consolidating hardware, software, and firmware information into a single comprehensive report, the Dell System E-support Tool can help streamline troubleshooting and accelerate problem resolution for environments using Dell PowerEdge servers.



Maintaining Dell Platforms with Dell Technical Updates

Recent firmware updates for Dell PowerEdge components and Dell PowerVault storage can help improve system functionality and minimize potential problems.



YOUR INFORMATION, YOUR WAY

To preview *Dell Power Solutions* in searchable Digital Edition format, visit powersolutionsdigital.dell.com. In addition, you can check the *Dell Power Solutions* Web site at DELL.COM/PowerSolutions for early-release articles, how-to's, implementation studies, best practices, and expert tips you won't find anywhere else. Search our **Related Categories** index and then view full-text articles from current and back issues online. Can't wait? For real-time access to content as soon as it's posted, subscribe your favorite RSS reader to the feed at DELL.COM/RSS.

72 **The Red Hat HPC Solution: Simplified High-Performance Linux Clusters**

By Wayne Slater and Gord Sissons

74 **Evaluating Third-Generation AMD Opteron Processors for HPC Workloads**

By Jacob Liberman, Mausmi Kotecha, and Vishvesh Sahasrabudhe

78 **Evaluating Mellanox ConnectX InfiniBand Performance in HPC Clusters**

By Munira Hussain, Toby Sebastian, Shivaraj Nidoni, and Gilad Shainer

**STORAGE**

98 **High-Density, Highly Scalable Storage: Dell EqualLogic PS5500E iSCSI SANs**

By Dylan Locsin and Travis Vigil

102 **Seagate SAS Drives Provide Optimized Near-Line Storage for Dell Systems**

By Barbara Craig

106 **Fast, Reliable Data Protection from Dell and CommVault**

By Sanjeet Singh and Jeff Echols

110 **Streamlining Data Management with CommVault Simpana and Dell EqualLogic PS5500E iSCSI SANs**

By Jeff Echols

112 **Simplified Data Protection with Disk-Based Backup from Dell and Symantec**

By Sanjeet Singh and Charles Butler

NETWORK AND COMMUNICATIONS

116 **Introducing Broadcom iSCSI Offload Engine Technology for Dell Servers**

By Dhiraj Sehgal

**SPECIAL SECTION:
PARALLEL APPLICATION
DEVELOPMENT SOLUTIONS**

120 **Scaling Software Architectures for the Future of Multi-core Computing**

By Matt Gillespie

INDEX TO FEATURED DELL CUSTOMERS

ALERT Life Sciences Computing	90
Deutsches Elektronen-Synchrotron (DESY)	96
HotSchedules	11, 23
Maxol Group	23
Optimal Blue	94
Peterson Sullivan PLLC	44
PT Ciptadana Capital	92
Sequestered Solutions Alaska, LLC	99
Stoneridge	12
Wharton School of the University of Pennsylvania	88

INDEX TO ADVERTISERS

American Power Conversion Corporation	29
CA, Inc.	53
Citrix Systems, Inc.	57
Dell Inc.	7, 33, 105, 109, 119, C3
Egenera, Inc.	49
F5 Networks, Inc.	15
Impulse Point	37
Intel Corporation	C2, 71
Novell, Inc.	82-83
Oracle Corporation	C4
Seagate Technology LLC	101
Symantec Corporation	64a

EDITORIAL STAFF

EDITOR-IN-CHIEF AND PUBLISHER Tom Kolnowski

MANAGING EDITOR Deb McDonald

FEATURES EDITOR Kathryn White

ASSISTANT MANAGING EDITOR Jim Duncan

SENIOR EDITORS Jennifer Malnick and Terrence O'Donnell

EDITORIAL ASSISTANT Amy J. Parker

ART DIRECTOR AND COVER DESIGNER David Chan

DESIGNER AND ILLUSTRATOR Cynthia Webb

STAFF WRITERS Romy Bauer, Jeanne Feldkamp, and Chris Young

CONTRIBUTING WRITERS Roberto Ayala; Lokesh S. Balu; Cris Banson; Sitha Bhagvat; Mark Black; Burk Buechler; Charles Butler; Ajaykumar C. H.; Onur Celebioglu; Shefali Chinai; Thomas Cloyd; Barbara Craig; Brent Douglas; André Dumouchelle; Jeff Echols; Jeremy Ford; Matt Gillespie; Andrew Gilman; Aziz Gulbeden; Munira Hussain; Shruthi Jayatheertha; Kay Kerr; Mausmi Kotecha; Jacob Liberman; Dylan Locsin; Greg Lyon; Todd Muirhead; Manjunath Narayanan; Shivaraj Nidoni; Li Ou, Ph.D.; Perumal Raja P.; Ranjith Purush; Arun Rajan; Vishvesh Sahasrabudhe; Gabor Samu; David Schmidt; David Schweighofer; Toby Sebastian; Dhiraj Sehgal; John Seward; Gilad Shainer; Timothy Sherbak; Sanjeet Singh; Gord Sissons; Wayne Slater; Nelson Stewart; Omar Sultan; Brock A. Taylor; Richard Thwaites; Joe H. Trickey III; William Urban; Seeram Vedantham; Travis Vigil; Ryan Weldon; and Ray Weinstein

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

ONLINE DESIGN Saurel Dominique and John McElhenney

ADVERTISING SALES

AD SALES EXECUTIVE Randy Byers

AD COORDINATOR Kathy Hartlove

SUBSCRIPTIONS

Subscriptions are complimentary to qualified readers who complete the online subscription form. For the Print Edition: To sign up as a new subscriber or change your existing subscription, access the Subscription Center at **DELL.COM/PowerSolutions**. For the Digital Edition: To sign up as a new subscriber or convert your Print Edition subscription to the Digital Edition, visit **powersolutionsdigital.dell.com**. For other subscription services, please e-mail **us_power_solutions@DELL.COM**.

ABOUT DELL

Dell Inc., headquartered in Round Rock, Texas, near Austin, listens to its customers and delivers innovative technology and services they trust and value. Uniquely enabled by its direct business model, Dell is a leading global systems and services company and No. 34 on the Fortune 500 list. For more information, visit our Web site at **DELL.COM**.

Dell cannot be responsible for errors in typography or photography. Dell, the Dell logo, Dell OpenManage, Dell Precision, DirectDetect, EasyConnect, EqualLogic, Latitude, On-Demand Desktop Streaming, OptiPlex, PowerConnect, PowerEdge, and PowerVault are trademarks of Dell Inc. Other trademarks and trade names may be used in this publication to refer to either the entities claiming the marks and names or their products. Dell disclaims any proprietary interest in the marks and names of others.

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

Copyright © 2008 Dell Inc. All rights reserved. Printed in the U.S.A.

 Printed on recycled paper containing 10 percent post-consumer waste. Please recycle this magazine.

GET GREEN: GO DIGITAL



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



The digital magazine you have been asking for is here. The new *Dell Power Solutions* Digital Edition—eminently green, eco-friendly, instantly accessible, and complimentary—is a natural extension of the hybrid publishing platform that we have been using to produce our conventional print and Web editions. Built on the Texterity digital publishing engine, the Digital Edition is entirely Web browser-based and requires no client-side software downloads or browser plug-ins. Just launch your browser and point to powersolutionsdigital.dell.com to begin the new digital book-reader experience.

Once you land on the *Dell Power Solutions* Digital Edition home page, you have several options on how to access the content. To get full access to the

Digital Edition—including powerful search privileges across the past year of *Dell Power Solutions* archived issues—take a moment and complete the subscription form. If you are already a print subscriber, you can convert to the Digital Edition by providing your current subscriber ID. In a few moments, you will receive an e-mail with your private URL for instant access. As a new Digital Edition subscriber, you will also receive an e-mail notification when the next issue is ready to access.

Not ready to subscribe yet? Then test-drive the Digital Edition by accessing our sample issue at powersolutionsdigital.texterity.com/powersolutions-sample/sample. You can put all of the navigation capabilities and other features through their paces, including Contents, a scrollable table of contents; Pages, a thumbnail view of the entire issue, including exclusive online-only content; Search, a powerful search feature with results returned as thumbnails and highlighted text; Links, a list of Web links on the current page or the entire issue; and Share, where you can send an e-mail to an associate with a thumbnail and direct link to the current page. At any time during your sample issue test-drive, you can click the SUBSCRIBE button in the navigation bar to join the membership ranks—and get green instantly while going digital. We hope you enjoy the digital drive.



Dell Power Solutions Digital Edition

Dell recommends Windows Vista® Business.



I WANT TO PROTECT MY DATA

**3x better data protection with Solid State Drives.
Free Fall Sensor. Strike Zone™ shock absorbers.**

THE NEW LATITUDE™

Featuring the Intel® Centrino® 2 Processor with vPro™ technology

DELL.COM/LATITUDE 866.238.DELL



Intel, the Intel logo, Centrino, Centrino Inside, Intel vPro and vPro Inside are trademarks of Intel Corporation in the U.S. and other countries.





A SMART PATH TO VIRTUALIZATION

As server virtualization continues to transform the data center, enterprises have some critical decisions to make. How will they migrate to a fully virtualized environment? How should the IT infrastructure be configured for maximum business benefit? Dell is partnering with other industry leaders to build a flexible, industry-standard approach that helps simplify the transition.

By **Jeanne Feldkamp**
Deb McDonald
Kay Kerr
Tom Kolnowski

Related Categories:

Citrix XenServer Microsoft Hyper-V
Dell PowerEdge servers Microsoft Windows Server 2008
Desktop virtualization Virtualization
Flexible computing VMware
Hypervisors

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

Recent strides in server virtualization and emerging hypervisor alternatives are changing the rules of engagement for enterprise IT departments. The way IT managers plan data centers, how and when purchase decisions are made, and—perhaps most importantly—the opportunities for IT departments to further business goals are evolving. Today, combining policy-driven automation with a virtualized environment can enable advanced functionality, including dynamic provisioning and workload balancing, high availability, and disaster recovery. And with each of these innovations, previous levels of functionality become mainstream and accessible even to small environments.

Meanwhile, the rapid progression of virtualization technology has left enterprise IT managers with more alternatives to evaluate than ever, particularly now that Microsoft is weighing in with an enterprise-class hypervisor—the Microsoft® Windows Server® 2008 Hyper-V™ platform—for the first time. Organizations that do not approach server virtualization with a clear plan risk developing a chaotic, inflexible infrastructure that wastes energy and resources. However, with appropriate planning, enterprises can leverage virtualization to enhance agility and service, simplify IT management, and empower staff to focus on strategic priorities while conserving energy and resources.

At the same time that server virtualization has gone mainstream, desktop virtualization has become a viable option for many organizations. Desktop virtualization technologies enable the Dell™ Flexible Computing Solutions suite of products and services, which combines centralized control with rich client functionality, performance, and mobility. The flexible computing model can help enterprises simplify desktop management and enhance security—while helping to improve business continuity and streamline regulatory compliance. (For more information, see the “Flexible computing on the rise” sidebar in this article.)

CHARTING A SMART PATH TO VIRTUALIZATION

Dell helps enterprises plan the transition with offerings that are grounded in industry standards. By providing standards-based, virtualization-optimized platforms and integrating tightly with key virtualization enablers such as the Microsoft Hyper-V, VMware® ESX, and Citrix® XenServer™ platforms, the Dell approach helps organizations simplify virtualization deployments without limiting their options. In addition, Dell offers a comprehensive range of management tools and services for virtualized systems.

Dell servers and storage are designed for optimal virtualization performance and simplified management. Dell PowerEdge™ servers support optimum performance with efficient power consumption, enhanced memory and I/O capability, and a broad range of form factors, hypervisors, and I/O fabrics. Dell also provides simple, capable, and cost-effective storage offerings designed to leverage existing IT infrastructures and expand automatically—helping meet a broad range of needs with the Dell PowerVault™, Dell EqualLogic™, and Dell/EMC families.

A variety of powerful Dell management tools for virtualized systems can help administrators proactively manage assets, enhance efficiency, maximize resources, and reduce total cost of ownership (TCO). To help avoid the complexity of using multiple management consoles to administer the IT environment, Dell products offer support for a range of leading management tools, including Microsoft System Center Virtual Machine Manager (SCVMM) 2008, VMware VirtualCenter, Citrix XenCenter™, and Symantec® Management Suite software. Additional Dell offerings are designed to integrate smoothly through collaborative co-development efforts with partners such as Egenera, PlateSpin, and Vizioncore.

In addition, Dell Global Services offers field-tested guidance to help enterprises accelerate virtualization deployments and

FLEXIBLE COMPUTING ON THE RISE

Desktop virtualization has gained popularity over the last several years because it allows IT departments to manage digital identities from a central location rather than managing, securing, and helping ensure compliance for hundreds or thousands of devices scattered across the globe. However, enterprise workers have diverse needs. Some users are highly mobile and require productivity tools that let them work anytime and anywhere. Some may work in an office, while others may work in highly specialized environments such as oil fields or operating rooms. In addition, some employees need only standard office applications such as e-mail, while others use demanding applications for highly targeted tasks such as data analysis or 3D rendering.

For these reasons, no single computing solution is likely to be suitable for every situation. A flexible computing model enables enterprises to deliver the right applications and the right levels of performance to the right users, on demand. To that end, Dell Flexible Computing Solutions offer a suite of hardware, software, and services designed to help enterprises increase productivity and simplify management of distributed workforces:

■ Dell On-Demand Desktop Streaming™ (ODDS) solution:

solution: The client boots from external networked storage and the OS is streamed.

■ Dell Virtual Remote Desktop (VRD) solution:

The server executes the desktop image and streams it to the client.

■ Dell Dedicated Remote Workstation (DRW) solution:

solution: The workstation is relocated to the data center and accessed through a portal device.

The flexible computing model also helps to reduce total cost of ownership, simplify image and data management, strengthen security and data protection, streamline compliance with organizational and regulatory requirements, and speed disaster recovery.*

*For more information, see "Flexible Computing: Advancing End-User Productivity with Centralized Control," by Jeremy Ford and Roberto Ayala, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090139-Flex.pdf.

train IT staff to manage ongoing operations. By working with Dell, organizations can tap into an extensive pool of best practices, innovative tools, and automated analysis to help ensure that the transition to virtualization is smooth and efficient. Dell virtualization solutions can be easily configured to address the challenges enterprises face today, with the flexibility to adapt as business conditions evolve.

DELIVERING BUSINESS-READY SERVER VIRTUALIZATION WITH MICROSOFT HYPER-V

Dell and Microsoft have teamed to offer comprehensive support for the Microsoft Hyper-V virtualization platform.¹ The two companies have jointly invested in the development of field readiness and expertise around Windows Server 2008 implementation and server virtualization on Dell hardware. In addition, Hyper-V is pretested on a wide range of Dell server and storage platforms, and is available as a factory-installed option on Dell PowerEdge servers. Reference architectures and a complimentary online Dell Virtualization Advisor tool are also available (see the "Dell Virtualization Advisor" sidebar in this article).

By providing partitioning for virtual machines (VMs), the Microsoft Windows Server 2008 OS with Hyper-V allows organizations to support both 32-bit and 64-bit VMs running side by side on a single physical server. Offering streamlined access through a unified console, Microsoft SCVMM 2008 enables Microsoft Hyper-V and other popular hypervisors to coexist in the same environment, providing organizations with a powerful solution for managing their virtual environment. When SCVMM 2008 is used with Microsoft System Center and Dell OpenManage™ software, the System Center console can manage a comprehensive environment—encompassing both virtual and physical resources—through the same intuitive interface.

Microsoft System Center Operations Manager 2007 helps automate routine, redundant tasks and provide intelligent reporting and monitoring for enhanced efficiency and control over the IT environment. Microsoft System Center Configuration Manager 2007 complements these tools by providing a way to comprehensively assess, deploy, and update servers, storage, and clients across physical, virtual, distributed, and mobile environments.

Dell plans to integrate Dell OpenManage with SCVMM 2008 through the Performance and Resource Optimization (PRO) Pack to help further simplify management of VMs on Dell PowerEdge servers. For organizations that prefer to have expert guidance during the implementation process, Dell Infrastructure Consulting Services for Hyper-V can facilitate a rapid approach to virtualization deployment and configuration. These offerings are based on proof-of-concept testing in Dell preproduction test labs. The Dell Structured Design Solution for Windows Server 2008—which includes assessment, design, and validation consulting—can also provide comprehensive services for organizations seeking a validated, repeatable migration process. (For more information on how Dell Infrastructure Consulting Services helped one company implement Hyper-V virtualization, see the "HotSchedules: Working smarter" sidebar in this article.)

ENHANCING VMWARE SUPPORT FOR SIMPLIFIED VIRTUALIZATION MANAGEMENT

As an established virtualization provider building on pervasive deployments of its software from the desktop to the data center, VMware offers the VMware Infrastructure 3, ESX, and ESXi platforms to enable continuous application availability and infrastructure security—and Dell has recently built on its strong support for VMware software.

¹For more information on Microsoft Hyper-V, see "Getting Started with Microsoft Windows Server 2008 Hyper-V on Dell Servers," by Ranjith Purush, Sitha Bhagvat, Ryan Weldon, Brent Douglas, and David Schmidt, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090140-HyperV.pdf.

HOTSCHEDULES: WORKING SMARTER

Dell Global Infrastructure Consulting Services helps HotSchedules design and deploy a Microsoft Hyper-V virtualized infrastructure, cutting power consumption by 77 percent and licensing costs by 75 percent.

HotSchedules creates online workforce-scheduling applications designed to help businesses in the restaurant, hospitality, and retail fields streamline employee scheduling, facilitate communication between managers and staff, and reduce costs. But faced with rapid growth, the company needed a data center infrastructure that would help accommodate a burgeoning customer base while also controlling hardware acquisition and power consumption costs.

"For the past few years, we have doubled the number of our customers every year, and we expect that growth rate to increase in the future," says Ray Pawlikowski, president of HotSchedules. "We quickly realized that we couldn't continue to add computer hardware for each new customer—we were running out of space and power in our data center. We needed a way to consolidate our resources while gaining scalability for future growth."

With help from Dell Global Infrastructure Consulting Services, HotSchedules deployed a virtualized server infrastructure with Microsoft Windows Server 2008 Hyper-V running on Dell PowerEdge servers. The company also uses a Dell PowerVault MD3000i storage area network, connected to the servers with Dell PowerConnect™ switches, to help make the most of the virtualized server environment.

Virtualization has helped HotSchedules dramatically consolidate its infrastructure, eliminating 36 of its 40 physical servers and leaving just 4 in the data center. The power savings alone deliver a strong return on investment. The company also expects to decrease the cost of software licensing. In addition, the new virtualized infrastructure will give the IT group the flexibility to add application servers quickly—allowing the company to pursue new accounts without having to worry about IT impediments.

"By virtualizing our server environment, we can now host up to 20 virtual servers on each Hyper-V cluster node. As a result, we have reduced the number of physical servers in our data center from 40 to just 4."

—Ray Pawlikowski
President of HotSchedules
October 2008



Dell servers have been optimized to help provide strong support for the latest ESX release, VMware ESX 3.5 Update 2, including enhanced performance for virtualized environments based on support for six-core Intel® Xeon® processors and 8 Gbps Fibre Channel host bus adapters, as well as Network File System (NFS) and Internet SCSI (iSCSI) over 10 Gigabit Ethernet. Additionally, Dell servers now enable VMware VMotion™ technology support for AMD Virtualization™ (AMD-V™) Extended Migration and Intel Virtualization Technology (Intel VT) FlexMigration technologies, and can take advantage of

multiple enhancements incorporated in VMware VirtualCenter to help increase availability and simplify virtual infrastructure management. ESX-compatible Dell server platforms have been expanded to include the PowerEdge M805, PowerEdge T100, and PowerEdge T105.

Support for VMware ESXi—which has a small 32 MB footprint and is designed to offer the same basic functionality and performance as VMware ESX—has been enhanced through the availability of ESXi embedded on the PowerEdge M805 and PowerEdge M905 blade servers. ESXi is also available after the point of sale to enable

organizations to standardize and upgrade existing infrastructures, including ESXi installable compatibility for a wide range of current and legacy PowerEdge servers.

In addition, Dell offers expanded storage support for virtualization with the Dell EqualLogic PS5500E iSCSI storage area network (SAN) and Dell/EMC CX4 Series storage arrays. And, on the Dell EqualLogic storage platform, the Dell EqualLogic Auto-Snapshot Manager/VMware Edition management tool is available and designed to provide integrated, automated, scalable, and cost-effective data protection for VMware VMs.² (For more information

²For more information, see "How Dell EqualLogic Auto-Snapshot Manager/VMware Edition Helps Protect Virtual Environments," by Andrew Gilman and William Urban, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090107-Gilman.pdf.

STONERIDGE: GREATER THAN THE SUM OF ITS PARTS

Dell PowerEdge servers, Dell EqualLogic storage, and VMware virtualization help Stoneridge reduce its server infrastructure by more than 50 percent.

As a tier-one supplier to some of the largest vehicle manufacturers in the world, Stoneridge—a global manufacturer of automotive components—has one mission: provide parts that deliver the capabilities its customers need, when they need them, in the most cost-effective way possible.

Until 2006, each business unit at Stoneridge determined its own infrastructure and IT strategy. However, senior company leadership realized that the decentralized IT strategy made it difficult to develop economies of scale and global efficiencies. Stoneridge executives saw that reducing the company's fleet of 260 servers—which were distributed among its global locations—could return significant cost and labor savings by cutting maintenance, administration, and replacement costs.

Working with Dell, the Stoneridge IT team eliminated more than 50 percent of its physical servers through server virtualization, using a combination of Dell PowerEdge servers supported by Dell EqualLogic Internet SCSI

“Having Dell as our global supplier simply eliminates many of our IT supply and integration challenges so that we can focus on our strategic architecture rather than tactical logistics.”

—Bill Johnson
CIO at Stoneridge
July 2008

on how Dell PowerEdge servers and VMware virtualization helped one company consolidate its infrastructure, see the “Stoneridge: Greater than the sum of its parts” sidebar in this article. To learn about recently released complimentary versions of Microsoft and VMware hypervisors, see the “Free hypervisors from Microsoft and VMware” sidebar in this article.)

COMPLEMENTING CITRIX XENSERVER WITH DELL HARDWARE AND TOOLS

Dell has also significantly expanded its tight integration with the Citrix XenServer hypervisor. Designed to support Citrix XenServer Dell Edition and Citrix XenServer Enterprise Edition as well as embedded versions, Dell PowerEdge servers offer

automated high-availability features that help protect against host failures. PowerEdge servers also enhance functionality for Dell EqualLogic PS Series iSCSI SAN arrays, including multipathing. Dell server platforms are designed to work seamlessly with XenServer, and the supported storage platforms have been expanded to include the Dell PowerVault MD1120 enclosure and Dell EqualLogic PS5500E iSCSI SAN.

In addition, Dell EqualLogic management tools have integrated into Citrix XenCenter to help simplify administration of virtualized servers. The Citrix XenServer Direct Storage Adapter allows out-of-the-box integration with EqualLogic PS Series iSCSI SANs. In this manner, the tools can intelligently relegate advanced

SAN capabilities such as thin provisioning, fast cloning, and automated snapshots to the EqualLogic arrays. EqualLogic management tools also facilitate integrated disaster recovery through scheduled backups of VM metadata.

DESIGNING DELL SERVER PLATFORMS FOR VIRTUALIZATION

Designed from the ground up for virtualization, Dell PowerEdge M805 and PowerEdge M905 servers offer outstanding performance for virtualization and high-end database needs in a full-height blade form factor. These servers feature three highly available, fully redundant I/O fabrics and eight high-speed ports. An internal Secure Digital (SD) card is provided



(iSCSI) storage area networks. The company chose Dell PowerEdge 2900 and PowerEdge 2950 servers with quad-core Intel Xeon processors for its pilot project, and has deployed a combination of those servers running VMware Infrastructure 3 at four sites.

As a result of its virtualization architecture on Dell PowerEdge servers, Stoneridge has reduced uninterruptible power supply draw at its main data center by 58 percent. High-performance Dell PowerEdge servers also helped Stoneridge consolidate its IT server infrastructure by more than 80 percent in the company's corporate network operations center. Moreover, the simplified infrastructure helps save approximately US\$330,000 in annual IT life-cycle replacement expenditures.

for installation of embedded hypervisors such as VMware ESXi.³

Dell PowerEdge M805 blade servers are designed to offer exceptional two-socket blade performance with the flexibility enterprise IT organizations require to support high-density, highly power-efficient data centers. Equipped with twice the memory and I/O connectivity of half-height Dell blade servers, PowerEdge M805 servers are designed to handle large workloads in a reduced physical space. Sixteen dual in-line memory modules (DIMMs) help increase cost-effective memory capacity for 32 GB and 64 GB configurations. Four embedded Ethernet controllers and four I/O dual-port mezzanine card slots provide high-availability I/O connections for each blade server.

Dell PowerEdge M905 blade servers are also designed to deliver powerful performance in virtualized environments. With enhanced performance from four quad-core processors, PowerEdge M905 servers provide a power-efficient approach for enterprises that need robust blade servers for heavy workloads. They also feature a large RAM capacity: 24 DIMMs enhance cost-efficient memory capacity for configurations greater than 48 GB.

Dell PowerEdge R900 rack servers with Intel Xeon 7400 series processors are designed for highly compute-intensive applications and architected for virtualization. Quad-core and six-core Intel Xeon 7400 processors amplify processing power, while 8 GB DIMMs can accommodate up to 256 GB of memory. These servers can deliver enhanced energy efficiency through the Dell Energy Smart high-efficiency power supply unit and come with a choice of embedded hypervisor—either VMware ESXi or Citrix XenServer.

Dell EqualLogic PS5500E iSCSI SANs offer compelling economies of scale for primary storage.⁴ Large PS Series nodes enable outstanding scalability and density

DELL VIRTUALIZATION ADVISOR

To help demystify the process of selecting a virtualization solution, Dell offers an online Virtualization Advisor tool that is designed to take into account existing server and storage infrastructure, workloads, and virtualization goals. The tool can recommend Microsoft Hyper-V, VMware ESX, or Citrix XenServer solutions.

After launching the tool, IT administrators simply select the parameters either for consolidating an existing physical server infrastructure using virtualization, or for deploying a new virtualized infrastructure, as well as the goals for their virtualized environment (see Figure A). In a matter of minutes, the Dell Virtualization Advisor can display a customized recommendation for servers, server configurations, and storage management and backup tools, as well as services—including a graphical representation of the validated configuration (see Figure B).

To use this tool, visit DELL.COM/Virtualization and select the Virtualization Advisor link in the right column. Organizations seeking an individualized approach to architecting a virtualized environment can engage the Dell Infrastructure Consulting Services team.

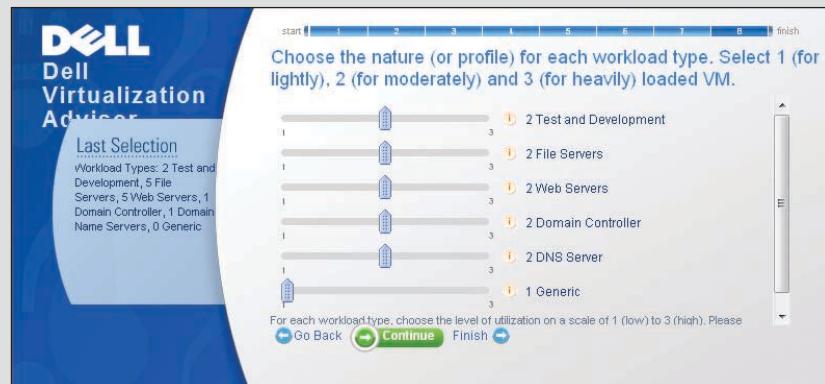


Figure A. Customizing virtual machine workload types and utilization levels for a new Microsoft Hyper-V environment



Figure B. Generating an example architecture for a Microsoft Hyper-V environment

³For more information, see "Introducing the Dell PowerEdge M805 and PowerEdge M905 Blade Servers," by Thomas Cloyd and Romy Bauer, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090110-Cloyd.pdf.

⁴For more information, see "High-Density, Highly Scalable Storage: Dell EqualLogic PS5500E iSCSI SANs," by Dylan Locsin and Travis Vigil, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20080396-Locsin.pdf.

FREE HYPERVISORS FROM MICROSOFT AND VMWARE

Both Microsoft and VMware have recently released complimentary versions of their hypervisors as downloads—the free version of Microsoft Hyper-V was christened Microsoft Hyper-V Server 2008, and was released shortly after VMware made a version of ESXi available as a complimentary download. Although these Microsoft and VMware hypervisors themselves are free, licenses are still required for any guest operating systems that are hosted on top of them as virtual machines (VMs)—and, of course, the complimentary versions offer fewer capabilities than the conventionally licensed versions.

Microsoft Hyper-V Server 2008 (not to be confused with the licensed Microsoft Windows Server 2008 Hyper-V) is a dedicated, stand-alone product that contains only the Hyper-V hypervisor with minimal virtualization components, and is designed to support basic virtualization functionality for Microsoft Windows® or Linux® OS-based VMs. Administrators can manage Hyper-V Server through a command-line interface (CLI) from the host; this tool also enables 64-bit guest and host support as well as integration with Microsoft System Center Virtual Machine Manager (SCVMM) 2008. Those who prefer a graphical user interface (GUI) can use SCVMM 2008 or the

Microsoft Management Console (MMC)-based Hyper-V Server Manager console from a separate client system, or upgrade to one of the licensed Windows Server 2008 versions running the Hyper-V role. High-availability clustering and quick migration of VMs would require Windows Server 2008 Enterprise Edition or higher. For a comparison of the various versions of Microsoft Hyper-V, visit www.microsoft.com/servers/hyper-v-server.

Administrators can manage the complimentary version of VMware ESXi using a GUI through the VMware Infrastructure Client. In addition, VMware offers an upgrade path from the complimentary version of ESXi to VMware Infrastructure 3 to provide access to enhanced enterprise management features, server consolidation, business continuity, and automated load-balancing capabilities. To help reduce the disk footprint, VMware migrated the VMware Service Console CLI in ESX to a remote CLI (RCLI) in ESXi that administrators can access from a virtual appliance; however, in the complimentary version of ESXi the RCLI is restricted to read-only access. For a comparison of the various versions of VMware ESX and ESXi, visit kb.vmware.com/selfservice/microsites/search.do?cmd=displayKC&externalId=1006543.

with a top-loading drive drawer, 48 Serial ATA (SATA) drives, and a 4U enclosure. Two models are available: the first with 48 TB (in 1 TB drives), and the second with 24 TB (in 500 GB drives).

Cost-effective EqualLogic PS Series SANs can provide support for near-line, archive, or backup-to-disk applications. Their small footprint supports large, dense installations and enhanced TCO while also providing virtualization-ready features and capabilities.

TAKING A FLEXIBLE PATH TO INDUSTRY-STANDARD VIRTUALIZATION

With a range of virtualization-optimized servers and storage, management tools, and services together with strong partnerships designed to support virtualization in the enterprise, Dell helps organizations chart a smart path to virtualization. Because every organization moves at a different pace, Dell offers flexible virtualization

solutions that can be customized for specific enterprise requirements. 

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

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

Kay Kerr is a senior marketing manager on the Dell Global Relationship Marketing team focusing on virtualization. Kay has a B.S. in Management Information Systems and an M.B.A. from the Ohio State University.

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



QUICK LINKS

Dell virtualization solutions:
DELL.COM/Virtualization

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell storage solutions:
DELL.COM/Storage

Dell and Microsoft Hyper-V:
DELL.COM/Hyper-V

Dell and VMware:
DELL.COM/VMware

Dell and Citrix XenServer:
DELL.COM/Citrix

Put total application control at your fingertips.



**FIND OUT HOW
www.f5.com**

F5 is the global leader in Application Delivery Networking. We have partnered with leading companies such as Microsoft and Oracle to ensure that your applications are always secure, fast and available.

For more info contact the F5 Dell team at dellteam@f5.com or:

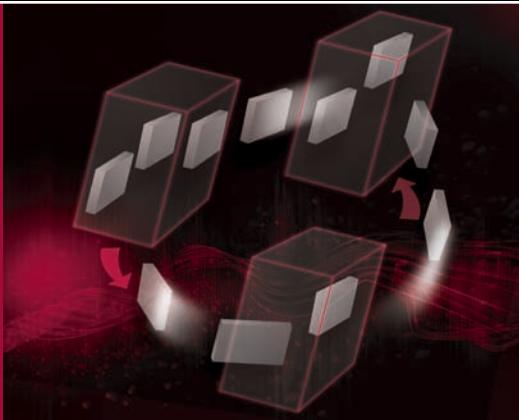
John Zygar
j.zygar@f5.com
206-272-6322

Scott Siragusa
s.siragusa@f5.com
512-818-1458

Melanie Christiansen
m.christiansen@f5.com
512-864-4813

Brian Pickering
b.pickering@f5.com
206-272-6254

Darin Spradley
d.spradley@f5.com
512-657-3172



By Ranjith Purush

Sitha Bhagvat

Ryan Weldon

Brent Douglas

David Schmidt

GETTING STARTED WITH MICROSOFT WINDOWS SERVER 2008 HYPER-V ON DELL SERVERS

Microsoft® Hyper-V™ technology provides a simplified virtualization platform integrated directly into the Microsoft Windows Server® 2008 OS. This article provides an introduction to Hyper-V virtualization, discusses the overall architecture and underlying technologies, and offers guidance on best practices for deployment on Dell™ PowerEdge™ servers.

Microsoft Hyper-V is a hypervisor-based virtualization technology integrated into all Dell-supported Microsoft Windows Server 2008 x64 Editions operating systems. As a virtualization platform, it enables organizations to run multiple virtual machines (VMs) on a single physical server to help take maximum advantage of hardware resources in their IT environments.

Windows Server 2008 introduced the concept of *roles* within the Windows Server OS. A single server OS instance is typically meant for a specific task, and a role defines the task the server performs: roles available in Windows Server 2008 include file server, Active Directory® server, Dynamic Host Configuration Protocol (DHCP) server, Domain Name System (DNS) server, and Internet Information Services (IIS) server. Hyper-V is also implemented as a role in Windows Server 2008, and offers several key advantages:

- Integration with the Windows Server 2008 OS, enabling organizations to easily take advantage of the benefits of virtualization in Windows Server environments without adopting a new technology
- A 64-bit micro-kernelized hypervisor architecture that leverages the existing device driver support in the Windows Server 2008 parent partition, extending support to a broad array of Dell servers, storage, and I/O devices

- Support for the Dell OpenManage™ suite in the Windows Server 2008 parent partition, helping simplify management of Dell servers
- Support for symmetric multiprocessing in VMs
- A high-availability feature for VMs to help minimize unplanned downtime, and a quick migration feature for VMs to help support business continuity during planned downtime
- Robust backup capabilities through native support for Microsoft Volume Shadow Copy Service (VSS)-based backups
- Simplified extensibility using standards-based Windows Management Instrumentation (WMI) and application programming interfaces

This article explores the basics of the Hyper-V architecture, discusses available management tools and their use, and outlines key considerations and best practices that can help administrators successfully plan and implement Hyper-V virtualization deployments on Dell servers and storage.

UNDERSTANDING THE HYPER-V ARCHITECTURE

The Hyper-V role is supported in both full and Server Core installations of Windows Server 2008 Standard x64 Edition, Enterprise x64 Edition, and Datacenter x64 Edition. Administrators should keep in mind that

Related Categories:

Dell PowerEdge servers

Microsoft Hyper-V

Microsoft Windows Server 2008

Operating system deployment

Systems management

Virtualization

Visit DELL.COM/PowerSolutions

for the complete category index.

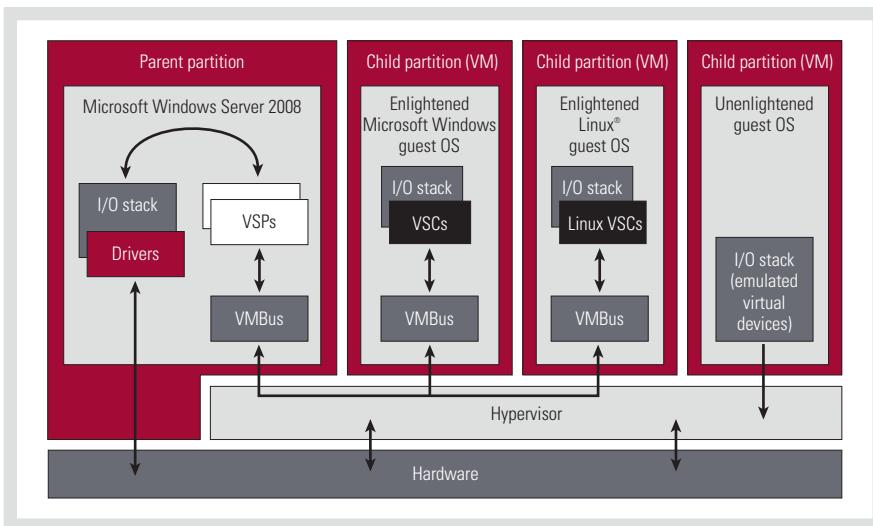


Figure 1. High-level Microsoft Hyper-V architecture

Windows Server 2008 Standard x64 Edition does not offer high-availability features; for highly available configurations with Hyper-V, they should use Windows Server 2008 Enterprise x64 Edition or Datacenter x64 Edition.¹ Dell supports Hyper-V on many servers, including the latest PowerEdge R805, PowerEdge R900, PowerEdge R905, PowerEdge M600, PowerEdge M605, PowerEdge M805, and PowerEdge M905 servers, which have been specifically designed to support virtualization workloads, as well as a broad range of previous-generation servers.²

The core component of the Hyper-V architecture is the hypervisor, which creates and manages isolated execution environments called partitions. The hypervisor sits directly on the hardware and controls access from the partitions to the physical processors. The Hyper-V server role in Windows Server 2008 also comprises several other key components, including parent and child partitions, synthetic and emulated devices, and integration services (see Figure 1).

Parent and child partitions

When the Hyper-V role is enabled, the hypervisor uses the virtualization extensions

in the processors—Intel® Virtualization Technology (Intel VT) or AMD Virtualization™ (AMD-V™) technology—to place itself under the OS. When the hypervisor loads for the first time, it creates a partition called the parent partition (or root partition), which hosts the Windows Server 2008 instance that had previously been running directly on the hardware. This partition is important for two main reasons. First, it controls hardware devices such as the network, storage, and graphics adapters and allocates physical memory to other partitions. Second, it makes requests to the hypervisor to create and delete child partitions, an

activity performed by the virtualization stack that runs in the parent partition.

Child partitions, unlike the parent partition, do not have access to the physical hardware. When a VM is created, it is assigned a newly created child partition and a set of virtual devices. I/O from the VM is routed through the parent partition to the physical hardware. This indirect I/O model enables VMs to be independent of the specific hardware devices on the physical server, and enables Hyper-V to take advantage of the broad support in Windows Server 2008 for Dell servers, storage, and I/O devices. Best practices strongly recommend downloading and installing the latest Dell-certified drivers for Windows Server 2008 x64 Editions in the parent partition. These drivers are available for download at support.dell.com.

Emulated and synthetic virtual devices

The virtual devices that a VM exposes to its guest OS fall into two broad types: emulated and synthetic (see Figure 2). Emulated virtual devices are software implementations of typical PCI devices, and to the guest OS, these devices appear to be physical PCI devices. Synthetic virtual devices, which are also implemented in software, use an architecture specific to Hyper-V based on

Emulated devices	<ul style="list-style-type: none"> Can include IDE controllers, legacy network adapters, and COM ports Take advantage of native device drivers in supported guest operating systems Typically provide lower performance than synthetic devices; however, certain emulated devices such as IDE controllers can switch to synthetic mode after Hyper-V integration services have been installed, which helps increase performance Are required for certain configurations; for example, an IDE controller is required as a boot controller for VMs, and a legacy network adapter is required for VM Preboot Execution Environment (PXE) boot
Synthetic devices	<ul style="list-style-type: none"> Can include SCSI controllers and network adapters Do not have native device drivers in guest operating systems; drivers are instead installed as part of Hyper-V integration services Use the high-performance VSP/VSC shared memory model, helping provide higher performance than emulated devices

Figure 2. Emulated and synthetic virtual devices in Microsoft Hyper-V

¹For more information on high-availability features in Hyper-V, see "Highly Available Virtualization with Microsoft Hyper-V and SCVMM 2008," by Ray Weinstein and Burk Buechler, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/p84q08-20080452-Buechler.pdf.

²For a list of supported hardware for Hyper-V, see "Dell Solutions Overview Guide for Microsoft Hyper-V," by Dell Inc., September 2008, available at support.dell.com/support/supportdocs/software/hyperv. For a list of recommended hardware for Hyper-V, see "Dell | Microsoft Windows Server 2008 Hyper-V Reference Architecture," by Dell Inc., September 2008, DELL.COM/Downloads/Global/Solutions/dell_hyper-v_reference_architecture_v1.0.pdf.

virtualization service providers (VSPs) and virtualization service clients (VSCs), and utilize a high-performance VMBus channel as the communication mechanism between the different partitions. This approach means that synthetic devices require less overhead than emulated devices.³

In general, this type of implementation that helps reduce overhead and enhance guest OS performance in virtualized environments is referred to as paravirtualization; Microsoft uses the term *enlightenment*. The synthetic device model is an example of *device enlightenment*. Hyper-V can also incorporate *kernel enlightenments*, which enhance a guest OS kernel so that it is aware of whether it is running in a virtualized environment, and can change its behavior to help reduce the overhead traditionally associated with running an OS within a VM. Microsoft Windows® operating systems that incorporate kernel enlightenments for Hyper-V environments include Windows Server 2008 and Windows Server 2003 with Service Pack 2.

Integration services

One of the primary advantages of emulated virtual devices is that all supported guest operating systems include native device driver support for these devices. This is not the case with synthetic virtual devices, however: Microsoft provides device drivers for synthetic devices in all supported guest operating systems, and administrators must install these drivers in each VM after the guest OS installation is complete. Microsoft delivers these device drivers through Hyper-V integration services.

Hyper-V integration services are designed to enhance integration between parent and child partitions. In addition to the device drivers for synthetic devices, these services also provide enhancements such as mouse integration, time synchronization, support for VSS-based backups of VMs, and the ability to shut down guest operating systems from the management console. Available features vary depending on the guest OS. Dell strongly recommends that administrators install the integration components within VMs after the guest OS installation is complete.

MANAGING HYPER-V ENVIRONMENTS

Administrators can use two primary tools to manage Hyper-V environments: the built-in Hyper-V Manager console, which provides basic management capabilities, or Microsoft System Center Virtual Machine Manager (SCVMM) 2008, which provides advanced enterprise-class features.

Hyper-V Manager

Hyper-V Manager is a Microsoft Management Console (MMC) snap-in that enables administrators to perform basic management tasks in Hyper-V environments, including creating and managing VMs, virtual networks, and virtual hard disks; importing and exporting VMs; and creating snapshots of VMs (see Figure 3). When using Hyper-V Manager to create VM snapshots, administrators should keep in mind that doing so is not equivalent to backing up a VM. The snapshot feature is designed to capture the VM state at a particular point in time, and is intended primarily for test and development use when administrators may need to roll back a VM to a previous state.

A single Hyper-V Manager instance can connect to multiple Hyper-V hosts. However, in this case, each host and the VMs on each host are managed independently of the other hosts. For advanced enterprise-class capabilities, administrators should typically use SCVMM 2008.

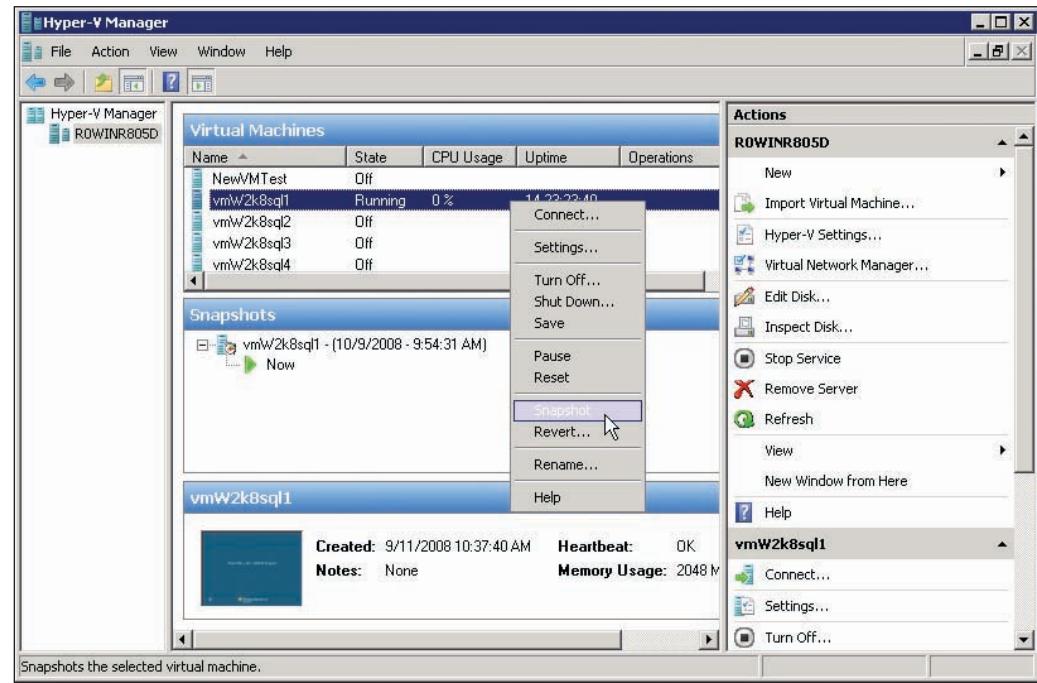


Figure 3. Microsoft Hyper-V Manager console

³For more information on emulated and synthetic devices, see "Hyper-V Architecture," by Microsoft Corporation, msdn.microsoft.com/en-us/library/cc768520.aspx.

When the Hyper-V role is enabled, administrators can launch the Hyper-V Manager console locally from Control Panel > Administrative Tools. They can also install it as a stand-alone management console on 32-bit and 64-bit editions of the Windows Server 2008 OS and Windows Vista® OS to enable remote management of Hyper-V hosts that are running either a full or Server Core installation of Windows Server 2008.

Microsoft System Center Virtual Machine Manager 2008

SCVMM 2008 is an enterprise-class management software suite that enables administrators to easily and effectively manage both physical and virtualized environments (including both Microsoft Hyper-V and VMware® ESX platforms) from a single management console, helping to significantly simplify management of the overall IT infrastructure. SCVMM 2008 integrates with the Dell OpenManage suite through Microsoft System Center Operations Manager (SCOM) 2007 to offer proactive management of Hyper-V hosts. SCVMM 2008 includes built-in support for physical-to-virtual and virtual-to-virtual migrations, works with failover clustering to support high-availability and quick migration features for VMs, and provides simplified automation through the Microsoft Windows PowerShell™ command-line shell.

The integration between SCVMM 2008, Dell OpenManage, and SCOM 2007 (or Microsoft System Center Essentials) helps further simplify the management of virtualized environments based on Dell servers and storage. This integration is based on Performance and Resource Optimization (PRO) capabilities that allow information to pass from SCOM 2007 to SCVMM 2008 and actions to be

“Hyper-V deployments require careful planning to help ensure that they can both meet current goals and scale for future growth.”

driven from SCVMM 2008 on SCOM 2007. Dell plans to provide this capability for a wide variety of scenarios in a PRO Pack in the next version of the Dell Management Pack.⁴

PLANNING HYPER-V DEPLOYMENTS

Hyper-V deployments require careful planning to help ensure that they can both meet current goals and scale for future growth. Key considerations include understanding the hardware utilization behavior of existing applications for consolidation and sizing the processors, memory, network, and storage requirements for the new virtualization solutions.

Hardware

Microsoft's "Hyper-V Planning and Deployment Guide"⁵ provides an overview of the capabilities of Hyper-V and generic deployment considerations. Before migrating workloads from a physical to a virtualized environment, best practices strongly recommend evaluating workload behavior on the existing physical servers over a period of time. Administrators should track this behavior across a typical utilization cycle to obtain the peak, average, and low utilization metrics; this assessment can then help appropriately size the VMs. Dell Infrastructure Consulting Services also offers workload assessment services to help with these efforts.⁶

Processors

In environments using multi-core processors, choosing physical processors can depend on many factors, including the type of workload, number of sockets in the server, processor utilization goals, and power consumption requirements. Hyper-V requires 64-bit processors that support Data Execution Prevention (DEP) and virtualization extensions (Intel VT or AMD-V). As of September 2008, all currently shipping two- and four-socket Dell PowerEdge servers meet these requirements. Based on the engineering evaluations, Dell recommends using two- or four-socket PowerEdge servers for all virtualization workloads.

The number of virtual processors to be assigned to a specific VM typically depends on the requirements of the workload running in the guest OS. Windows Server 2008 Hyper-V supports up to four virtual processors per VM, depending on the specific guest OS. Administrators should keep in mind that the virtual processors in VMs do not have a one-to-one mapping with the logical processors in the server. The Hyper-V hypervisor handles the scheduling of virtual processors on available logical processors. In general, when the virtual workload is highly processor intensive, using multiple virtual processors in a VM might help increase performance by enabling additional physical processors to be used; however, using multiple virtual processors can also create additional overhead.⁷

⁴For more information on PRO capabilities, visit www.microsoft.com/systemcenter/virtualmachinemanager/en/us/whats-new.aspx.

⁵Available at www.microsoft.com/downloads/details.aspx?FamilyID=5DA4058E-72CC-4B8D-BBB1-5E16A136EF42.

⁶For more information, see the Dell Services offerings at DELL.COM/Hyper-V.

⁷For more information, see the "Performance Tuning for Virtualization Servers" section in "Performance Tuning Guidelines for Windows Server 2008," by Microsoft Corporation, June 9, 2008, www.microsoft.com/whdc/system/sysperf/perf_tun_srv.mspx.

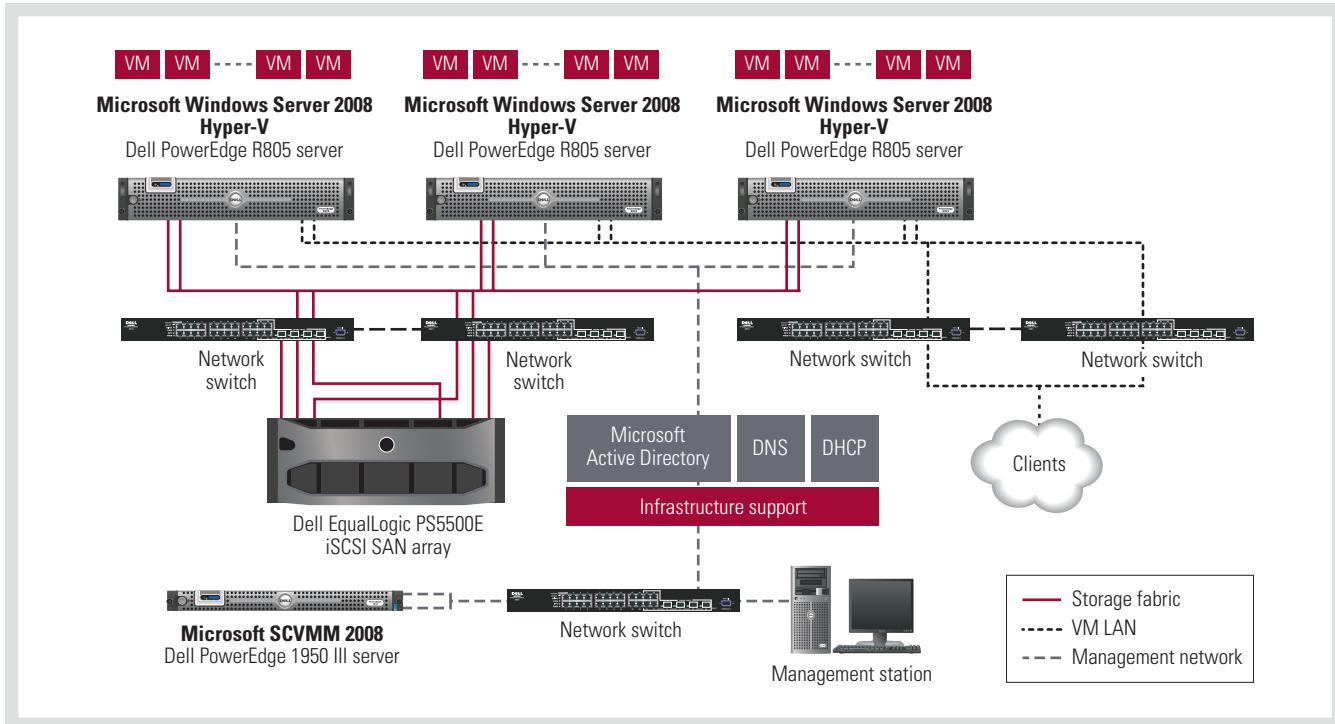


Figure 4. Example Microsoft Hyper-V deployment using Dell servers, storage, and switches

If the Hyper-V hosts will be part of a Hyper-V host cluster configuration, Dell strongly recommends that all processors on all Hyper-V hosts in the cluster be identical. If they are not identical, VM migrations between Hyper-V hosts may be unsuccessful.⁸

Memory

Virtualization workloads are memory intensive—especially when one of the goals of the virtualization deployment is consolidation. The amount of memory allocated to a specific VM depends on the workload running on the guest OS. For a physical-to-virtual migration, the information obtained from the workload assessment can provide useful information on sizing VM memory. Administrators should keep in mind that a guest OS only has access to the amount of memory that has been allocated to the VM. In Windows Server 2008 Hyper-V, changing the memory configuration of a VM requires powering it down.

After administrators have determined the amount of memory to allocate to the VMs, they can then determine the memory requirements for the physical server. Best practices strongly recommend that when administrators configure the Hyper-V server role, they should not enable any other roles in the parent partition. However, the parent partition should have sufficient memory to provide management capabilities and virtual device functionality for child partitions and to manage the physical devices on the server. In general, the amount of physical memory should be the total of the following:

- **Parent partition:** 2 GB of physical memory
- **Each child partition:** Amount of physical memory allocated to each VM, plus 32 MB of overhead for the first 1 GB of memory allocated to the VM, plus 8 MB of overhead for each additional 1 GB of memory allocated to the VM

If additional VMs may possibly migrate to the server, memory should be reserved for those VMs as well.

Networking and storage

Administrators must consider several factors when designing the network configuration and choosing the storage hardware for a Hyper-V deployment. For example, the requirements change significantly depending on whether the deployment is a highly available environment (as recommended by Dell) or a stand-alone virtualized environment. For a list of supported storage arrays and in-depth information on planning network and storage deployments, see the Dell solutions guides for Hyper-V available at support.dell.com/support/edocs/software/hyperv.

Example deployment

An actual Hyper-V deployment involves multiple servers, shared storage, and networking components; careful consideration is essential to help ensure that

⁸For more information, see "Dell High Availability Solutions Guide for Microsoft Hyper-V," by Dell Inc., September 2008, available at support.dell.com/support/edocs/software/hyperv.

the configuration provides the highest level of performance for VM workloads. Figure 4 shows an example Dell Hyper-V deployment. This type of deployment would include several key components:

- **Windows Server 2008 Hyper-V:** Tenth-generation Dell PowerEdge servers are specifically designed to support virtualization platforms like Hyper-V. Dell strongly recommends installing Dell OpenManage Server Administrator (OMSA) systems management software in the parent partition of each PowerEdge server.
- **VM storage:** To support Hyper-V high-availability and quick migration features for VMs, the VMs must be stored on external storage such as the Dell EqualLogic™ 5500E Internet SCSI (iSCSI) storage area network (SAN) array shown in Figure 4. Dell supports a wide array of storage options in Hyper-V environments, including Dell PowerVault™, Dell EqualLogic, and Dell/EMC arrays.⁹
- **SCVMM 2008:** Administrators can configure SCVMM 2008 in multiple ways depending on the implementation requirements. A basic configuration would run SCVMM 2008 on a stand-alone PowerEdge server using local disks as storage. Attaching a storage enclosure to the stand-alone server hosting SCVMM 2008 is recommended if the deployment requires a relatively large library server. The library server is a capability built into SCVMM 2008 for storing .vhdx templates, inactive VM files, ISO images, and so on.
- **Infrastructure support:** Typical deployments require a server to host key infrastructure roles such as Active Directory Domain Services, DNS, and DHCP. A Windows-based domain controller is required for SCVMM 2008 to support its key capabilities.
- **Management station:** A management station is essential to help administrators

effectively manage the virtualized infrastructure. Because the management station primarily runs client software that communicates with the server components, the management station can be a Dell client platform such as a Dell Latitude™ laptop or Dell Precision™ workstation. Typical management stations include the Hyper-V Manager MMC console, the SCVMM 2008 Administration Console, and Dell OpenManage IT Assistant, which provides one-to-many management for PowerEdge servers running OMSA.

DEPLOYING SIMPLE, POWERFUL VIRTUALIZATION

Microsoft Hyper-V technology is designed to provide a simple yet powerful virtualization platform integrated directly into the Windows Server 2008 OS. By taking advantage of architectural enhancements, using appropriate management tools, and following Dell best practices for planning and implementing Hyper-V on Dell server and storage hardware, administrators can help ensure successful Hyper-V deployments in their environments. 

Ranjith Purush is a lead systems engineer in the Dell Virtualization Solutions Engineering Group. He has led the release of numerous Dell solutions, and his current interests are in technologies that enable a dynamic data center. Ranjith has an M.S. in Computer Engineering from the University of Texas at Austin.

Sitha Bhagvat is a systems engineer in the Dell Virtualization Solutions Engineering Group. She has an M.S. in Computer and Information Sciences from the Ohio State University.

Ryan Weldon is a systems engineer in the Dell Virtualization Solutions Engineering Group focused on storage solutions. Before joining this team, Ryan spent three years as

a lead engineer in the Dell PowerVault Storage Engineering organization. He has a B.S. in Computer Science from Embry-Riddle Aeronautical University.

Brent Douglas is a systems engineer in the Dell Virtualization Solutions Engineering Group. He has a B.S. in Electrical Engineering from Purdue University and an M.B.A. from the Purdue Krannert School of Management, and holds Microsoft Certified Systems Engineer: Security (MCSE: Security), Microsoft Certified Technology Specialist: Virtualization (MCTS: Virtualization), Citrix® Certified Administrator (CCA), and VMware Certified Professional (VCP) certifications.

David Schmidt is an engineering manager for the Dell Virtualization Solutions Engineering Group. Previously, David worked as a software developer for Dell OpenManage products. David has a B.S. in Computer Engineering from Texas A&M University.

MORE

ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

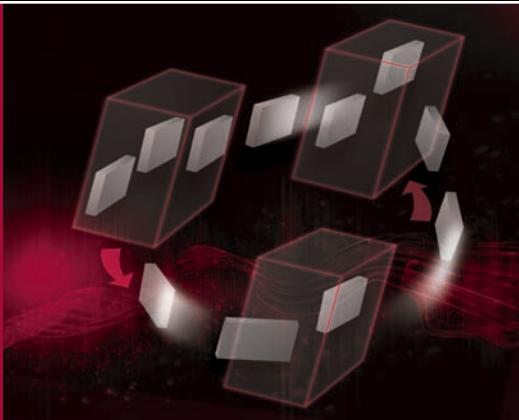
Dell and Microsoft Hyper-V:
DELL.COM/Hyper-V

Dell solutions guides for Microsoft Hyper-V:
support.dell.com/support/edocs/software/hyperv

Dell Virtualization Advisor tool:
advisors.dell.com/advisorweb/advisor.aspx?advisor=c82c3ec8-c94f-4602-9a41-c20382db1cd0

Dell reference architecture for Microsoft Hyper-V:
DELL.COM/Downloads/Global/Solutions/dell_hyper-v_reference_architecture_v1.0.pdf

⁹For more information, see "Dell Storage Solutions Guide for Microsoft Hyper-V," by Dell Inc., September 2008, available at support.dell.com/support/edocs/software/hyperv.



By Ray Weinstein
Burk Buechler

HIGHLY AVAILABLE VIRTUALIZATION WITH MICROSOFT HYPER-V AND SCVMM 2008

In virtualized environments, high availability can be critical to avoiding lost revenue, maintaining productivity, and protecting data from system failure. Combining Microsoft® Hyper-V™ technology with Microsoft System Center Virtual Machine Manager (SCVMM) 2008 on Dell™ platforms can help organizations create highly available virtualized systems while offering simplified management, dynamic performance, workload optimization, and reduced costs.

Virtualization solutions can offer significant benefits, helping boost IT efficiency and capacity and providing a robust, simplified environment while also helping significantly reduce costs. However, virtualization can also create single points of failure in critical environments: for example, if an enterprise is running 10 production workloads on 10 separate physical servers, the failure of a single server would likely be inconvenient but not disastrous. However, in a virtualized environment with 10 production workloads running in 10 virtual machines (VMs) on a single physical server, the failure of that server could be catastrophic.

In addition, computing systems have become such an integral part of enterprise environments that severe consequences can occur when the systems become unavailable. When critical functions such as the Microsoft Active Directory® directory service, the Microsoft Exchange messaging platform, or Microsoft Office SharePoint® Server services are not accessible, enterprises may face lost revenue, lost productivity, reduced customer satisfaction, data loss, or even diminished decision-making capabilities. As a result, these computing systems must be available 24/7.

Introductory virtualization projects often focus on reducing capital expenses rather than on

implementing holistic virtualized solutions. A comprehensive virtualized environment, however, goes beyond consolidation to take advantage of the encapsulation and portability of VMs as well as flexible, powerful management tools. These features enable administrators to use virtualization to help increase availability, manage workloads effectively without overprovisioning hardware, and implement disaster recovery plans.

Microsoft Hyper-V technology and Microsoft System Center Virtual Machine Manager (SCVMM) 2008 provide the core foundational components of a comprehensive solution, helping organizations simplify virtualized IT infrastructures, reduce costs, and increase availability. With integrated administration, administrators can use a single console to help centralize management of a heterogeneous VM infrastructure, increase physical server utilization, quickly provision new VMs, and provide dynamic performance and resource optimization of hardware, operating systems, and applications. Dell can help enterprises implement these features to extend their use of virtualization and help provide enhanced service levels, responsiveness, and enterprise agility (see the “Microsoft Hyper-V in the real world” sidebar in this article).

Related Categories:

High availability (HA)
Hypervisor software
Microsoft
Virtualization

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

QUICK MIGRATION

In assessing the cost-effectiveness of a high-availability environment, administrators must balance the cost of downtime against the cost of a proposed solution. The first consideration to help minimize failures should be the quality and reliability of servers and storage. To this end, Dell servers and storage incorporate a broad array of redundant features designed to maximize availability. High-availability clusters can provide the additional level of availability required for critical applications. Hyper-V takes this a step further by introducing an advanced level of integration between the Microsoft Windows Server® 2008 clustering feature and the Hyper-V quick migration feature (see Figure 1).

The quick migration feature can help increase the availability and flexibility of critical services during planned downtime for maintenance, or quickly restore services after unplanned downtime. It enables administrators to rapidly fail over a running VM from one physical host to another with minimal downtime.

HOST CLUSTERING AND GUEST CLUSTERING

Hyper-V provides two different clustering options—host clustering and guest clustering—which can be used separately or together. Host clustering is the more common of the two. It enables VM host-to-host failover without requiring the

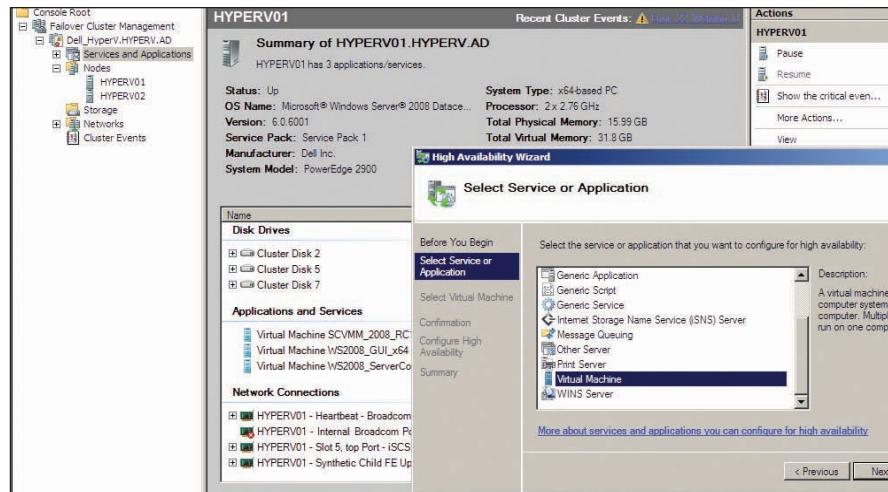


Figure 1. Microsoft Hyper-V quick migration

applications to be cluster aware; the Windows Server 2008 clustering feature manages the failover at the host level. This capability enables applications that traditionally have been incompatible with clustering to become highly available with little or no additional configuration or modification. Host clustering requires that host systems be running Windows Server 2008 Enterprise x64 Edition or Datacenter x64 Edition using a full or Server Core installation. VMs can run any cluster-enabled Windows server OS (Windows 2000 Server, Windows Server 2003, or Windows Server 2008). In addition, guest clustering requires cluster-aware applications such as the Microsoft SQL Server® 2008 Enterprise Edition database platform or the Microsoft Exchange Server 2007 Enterprise Edition messaging platform.

Guest clustering is less common than host clustering, but takes a similar approach to the one used to cluster other Microsoft Windows® OS-based applications outside

a Hyper-V environment. Cluster services within the guest OS manage the failover at the VM level. Guest clustering requires that host systems be running Windows Server 2008 Enterprise x64 Edition or Datacenter x64 Edition using a full or Server Core installation. VMs can run any cluster-enabled Windows server OS (Windows 2000 Server, Windows Server 2003, or Windows Server 2008). In addition, guest clustering requires cluster-aware applications such as the Microsoft SQL Server® 2008 Enterprise Edition database platform or the Microsoft Exchange Server 2007 Enterprise Edition messaging platform. Because this approach contains the clustered applications within a VM, it also enables ancillary benefits such as increased

MICROSOFT HYPER-V IN THE REAL WORLD

Two companies currently benefiting from highly available Microsoft Hyper-V solutions from Dell are Maxol Group, a venerable Irish-owned oil company, and HotSchedules, which provides a leading software-as-a-service (SaaS) labor scheduling and workforce management system. The Dell Infrastructure Consulting team helped deploy both solutions as part of the Microsoft Rapid Deployment Program for Hyper-V in advance of the product launch.

Maxol Group. Maxol takes advantage of Hyper-V to help increase the reliability of its data centers. "We can keep all our virtual machines available at all times," says Nicholas Merton, a member of the IT support team at Maxol Dublin. "If anything happens to one server, the cluster automatically moves the affected workloads to a new server, transparent to users. With

Hyper-V, we'll essentially eliminate downtime, which is important as we move into several e-commerce ventures that require high reliability."

HotSchedules. High availability is a critical requirement for HotSchedules. Dell helped the company's IT group design a virtualized infrastructure based on Hyper-V failover clusters. "We currently have 2,400 client installations online. We could very well triple that number in the next year, just by adding a few large chains," says Ray Pawlikowski, president of HotSchedules. "With the Dell infrastructure running Hyper-V, we could easily deploy 120 virtual servers on just 12 physical machines to accommodate that growth without maxing out our data center. The new infrastructure lets the business pursue those large clients with confidence."

MICROSOFT HYPER-V GUEST OS SUPPORT

Microsoft Hyper-V virtual machines (VMs) can run any supported guest OS that has been validated and granted *integration component* support. Integration components are sets of drivers and services designed to provide additional device support within VMs, help maintain consistent VM states, and enhance guest OS performance by enabling the guest OS to use synthetic devices. Integration components included with Hyper-V include a VMBus driver to provide transport for synthetic devices, time synchronization to keep VM clocks synchronized with the root partition, mouse integration, a video driver, a synthetic network driver, and synthetic storage drivers.

For a list of the latest guest operating systems supported by Hyper-V, visit www.microsoft.com/windowsserver2008/en/us/hyperv-supported-guest-os.aspx.

hardware utilization, quick and intelligent provisioning, and increased availability options.

SQL Server and Exchange can be made highly available in a number of ways, each with its own particular considerations. Administrators should weigh the advantages and disadvantages of each approach carefully before deciding on a specific plan. The Dell Infrastructure Consulting team can provide expert guidance to help organizations design a suitable custom solution.

PLANNED FAILOVER AND UNPLANNED FAILOVER

In addition to understanding the two types of clustering options, administrators should also be sure they understand the two types of failover events: planned failover and unplanned failover. A planned failover is an administrator-invoked migration of VMs

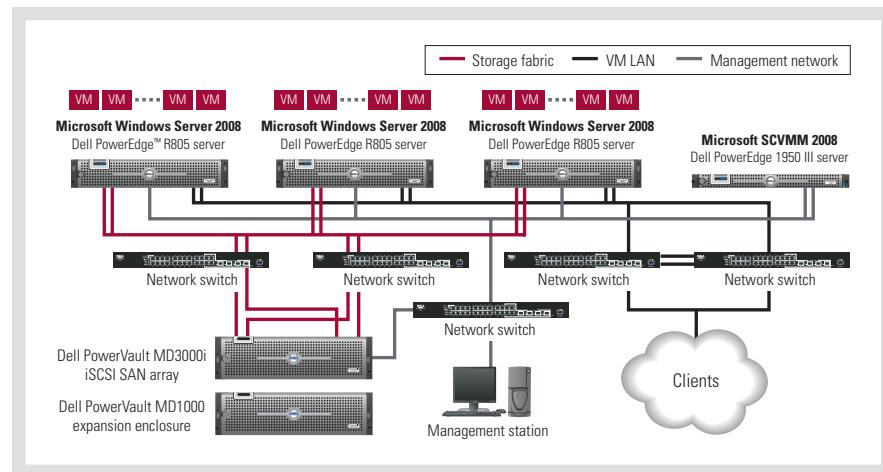


Figure 2. Example Dell high-availability reference configuration

from one or more clustered nodes to one or more other available nodes. These failovers are most commonly used during maintenance windows or to provision a workload across clustered nodes. The Hyper-V quick migration feature saves the state of a running VM by writing the VM memory to disk on shared storage, moves the storage ownership from one node to another, and then restores the VM state on the second node. The failover speed depends on several factors, including amount of memory, storage and interconnect performance, and even the guest OS version used in the VM. In a typical environment, failover would generally complete in 15-120 seconds.

Conversely, an unplanned failover occurs without administrative interaction and typically results from a hardware device failure or power outage. In the case of unplanned downtime, the workload state cannot be saved; instead, the resources written to the disk would be failed over from the shared storage automatically based on administrator-defined options. The duration of the service outage is typically longer than that of a planned failover because the VMs must complete the entire boot sequence rather than the saved state sequence used in a planned failover. (Both VMware® VMotion™ technology and the Hyper-V quick migration feature must reboot guest VMs following an unplanned failover.)

BEST PRACTICES FOR HYPER-V HIGH AVAILABILITY

Implementing Hyper-V high-availability solutions can be challenging even for experienced IT administrators. Dell helps simplify the adoption of this technology by introducing tools and services such as support matrices, reference configurations (see Figure 2), the Dell Virtualization Advisor tool, and services from the Dell Infrastructure Consulting team.

Necessary cluster components and dependencies include the following:

- **Host nodes:** Each Hyper-V cluster uses 2-16 compliant physical servers to host VM workloads. These servers should run Windows Server 2008 Enterprise x64 Edition or Datacenter x64 Edition to host the failover clustering feature as well as the Hyper-V role. The guest VMs can run any combination of client- or server-based operating systems currently supported with Hyper-V.
- **Shared storage:** Shared storage should be provisioned to support the clusters. This requirement is delivered on Internet SCSI (iSCSI)- or Fibre Channel-based storage area networks (SANs) with high-performance hard disks. Direct attach storage (DAS) can also be added on Dell PowerVault™ MD3000 enclosures. Windows Server 2008 provides connectivity natively for iSCSI through the Microsoft iSCSI Software

Initiator, while the Microsoft Multipath I/O framework helps make SAN connectivity straightforward to configure and helps increase availability.

- **Network interconnects:** Hyper-V clusters require network interconnects for a series of objectives and dependencies. These dependencies include network interface cards (NICs) for parent, guest, or iSCSI interconnects; host bus adapters (HBAs) for Fibre Channel-based storage; and Serial Attached SCSI (SAS) adapters for DAS arrays.
- **Other dependencies:** Administrators should also know the requirements of their particular environment, which may include Active Directory domain services, Domain Name System (DNS), TCP/IP addressing, and appropriate credentials.

Dell best practices for implementing highly available systems based on Hyper-V include the following:

- Use servers that carry the “Designed for Windows” logo.
- Use SCVMM 2008 to manage complex solutions and help ensure intelligent workload provisioning.
- Use tools such as PlateSpin PowerRecon to assess and establish baselines for the physical environment’s performance requirements before proceeding with a virtualization implementation.
- Install the latest Microsoft Quick Fix Engineering (QFE) updates for Hyper-V and clustering through Windows Server Update Services.
- Validate that the shared storage solution is compliant with both Windows Server 2008 and Hyper-V clustering; some solutions may be Windows Server 2008 compliant but not support Hyper-V clustering. Check with the storage vendor to learn whether updates to existing firmware and drivers are necessary.
- Use the Node and Disk Majority option for the cluster witness (called the

quorum in previous Windows Server versions) model in most single-site cluster deployments. Defer to the storage vendor for the optimal selection.

- Use Server Core deployments as often as possible for host servers and/or VMs. In addition to providing administrative benefits, VMs using Server Core installations can fail over significantly faster than VMs using comparably configured full installations with a graphical user interface. In addition, because Server Core installations typically require less memory and hard disk resources than full installations, the same host space can support additional VMs.
- For hard disks, provision one or more discrete logical units (LUNs) to each VM to enable individualized VM failover. Use tools such as Sanbolic Kayo if cluster file systems and multiple VMs per LUN are necessary. In addition, use highly available mount points for implementing more than 21 drive letters (LUNs) within a cluster, and use fixed virtual hard disks for optimal flexibility and enhanced performance.
- For scalable solutions, add a single dedicated Gigabit Ethernet NIC port for every four VMs, depending on individual VM network I/O requirements. These guest network NIC ports are in addition to the other required ports recommended for Hyper-V parent partition accessibility, cluster heartbeats, and iSCSI connectivity (if required).¹

MICROSOFT SUPPORT FOR VIRTUALIZED APPLICATIONS

Virtualization in critical enterprise environments has often been hampered by complex and confusing support policies. For example, many vendors provide “best effort” virtualization support that requires administrators to reproduce the problem on physical hardware before they will offer support.

As of August 2008, Microsoft has radically changed its policies for the most commonly used Microsoft applications—including SQL Server, Exchange, and the System Center family—to provide support in virtualized environments. These support changes apply not only to Hyper-V, but to hypervisors from other vendors as well. For Microsoft to support its applications on hypervisors from other vendors, these vendors must participate in the Microsoft Server Virtualization Validation Program.²

SIMPLIFIED HIGH AVAILABILITY FOR VIRTUALIZED ENVIRONMENTS

Microsoft Hyper-V technology combined with SCVMM 2008 can provide organizations with flexible, highly available virtualization for Dell platforms. By deploying these tools and taking advantage of expert guidance from the Dell Infrastructure Consulting team, enterprises can move beyond basic hardware consolidation toward a holistic virtualized environment designed for simplified management, dynamic performance and resource optimization, and reduced infrastructure and operational costs.³

Ray Weinstein is a global practice lead for the Dell Infrastructure Consulting Microsoft Practice.

Burk Buechler is a senior global product manager for the Dell Infrastructure Consulting Microsoft Practice.



QUICK LINK

**Dell and Microsoft Hyper-V:
DELL.COM/Hyper-V**

¹For more information on best practices, see “Hyper-V Planning and Deployment Guide,” by Microsoft Corporation, August 2008, www.microsoft.com/downloads/details.aspx?FamilyID=5DA4058E-72CC-4B8D-BB81-E16A136F42; and “Step-by-Step Guide for Testing Hyper-V and Failover Clustering,” by Microsoft Corporation, May 2008, www.microsoft.com/downloads/details.aspx?FamilyID=CD828712-8D1E-45D1-A290-7EDADF1E4E9C.

²For more information on Microsoft support policies for virtualized environments, visit support.microsoft.com/?kbid=957006. For more information on the Server Virtualization Validation Program, visit www.windowsserverscatalog.com/svvp.



By Jeremy Ford
Roberto Ayala

FLEXIBLE COMPUTING: ADVANCING END-USER PRODUCTIVITY WITH CENTRALIZED CONTROL

Supporting an increasingly distributed workforce can be a huge drain on IT resources. By enabling a customized end-user experience with server-based control of data, settings, and applications, the Dell™ Flexible Computing Solutions suite of products and services is designed to simplify management, enhance security, and lower total cost of ownership.

The traditional client computing landscape is evolving rapidly as business-critical applications and data continue to grow in size, scope, and complexity. The increasing challenges facing IT departments around manageability, data security, disaster recovery, compliance, and budget constraints are increasing interest in alternatives to traditional PCs. There is particular interest in “locked-down” server-based computing models that tend to emphasize IT control over end-user flexibility.

When it comes to enterprise computing, however, there has been a tenuous link between IT control and end-user flexibility. This already strained relationship may reach a breaking point as the expanding global workforce becomes increasingly mobile and end users are demanding anytime, anywhere access to data and productivity tools from a diverse range of client devices—including not just traditional PCs, but also PDAs, cell phones, public kiosks, and more.

Managing such environments can be a huge drain on IT time and resources. Day-to-day client management tasks such as deploying, patching, and migrating images, applications, and data can be burdensome and costly. Also, as workers employ a broadening range of client devices to access enterprise networks from remote locations, tasks such as enforcing security, ensuring organizational and regulatory compliance, and recovering systems and data following a

disaster can also increase in difficulty. It is clear that a “one size fits all,” locked-down solution is not the answer to this growing challenge.

To help simplify the task of supporting enterprise workforces, many organizations are turning to computing models that centralize control of end-user data and images while promising to offer workers the same flexibility and functionality they have come to expect from their client devices—if not better.

To help organizations realize the benefits of these emerging client computing architectures, Dell has introduced the Dell Flexible Computing Solutions suite of products and services, designed to enhance productivity, simplify management, and lower total cost of ownership. This suite includes products and technologies suited to different types of environments along with a comprehensive set of design, implementation, and support services to help organizations deploy and manage Dell Flexible Computing Solutions optimized for their particular needs.

With flexible computing, the “digital identity” of a worker—data, user preferences, applications, and even the OS—is decoupled from the individual client device and moved to the server. But unlike previous generations of server-based computing models, emerging flexible computing solutions embrace advances such as desktop virtualization to enable an expanded range of client functionality, flexibility, and

Related Categories:

Dell OptiPlex desktops
Dell Precision workstations
Dell Services
Disaster recovery
Flexible computing
Mobility
Security
Systems management
Virtualization
VMware

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

device options. This approach allows organizations to gain the benefits of centralized control—such as enhanced manageability and security—while providing end users with the functionality and mobility that can boost their productivity when using PCs and other devices.

DESIGNING FOR A DIVERSIFIED LANDSCAPE

Flexible computing has the potential to offer the benefits of centralized control combined with rich client functionality, flexibility, and mobility. Because the needs of enterprise workers can be complex and variable, however, no single flexible computing solution is likely to fit every situation.

As a result, IT organizations must consider the requirements of all stakeholders (including client and data center IT staff, facilities staff, and end users) to get a clear view of which users these technologies are best suited for. Less comprehensive approaches that are not end-to-end tend to move the problem around the enterprise without really solving the core issues that generate complexity.

To help organizations assess the appropriate solution for their needs, Dell Infrastructure Consulting Services (ICS) offers a comprehensive End User Profiling Assessment service. With this service, Dell ICS utilizes an end-user “fingerprinting” methodology to analyze the needs of both the end users and the organization as a whole. The resulting usage profile, or “End User Fingerprint,” segments end users based on their complexity of workload, level of mobility, and sensitivity of data and can be used to determine the appropriate traditional solutions or Flexible Computing Solutions for that organization.

All Dell Flexible Computing Solutions are similar in that they locate the end user's digital identity on the server. Individual solutions differ, however, in the way processing is distributed between client device and server (see Figure 1):

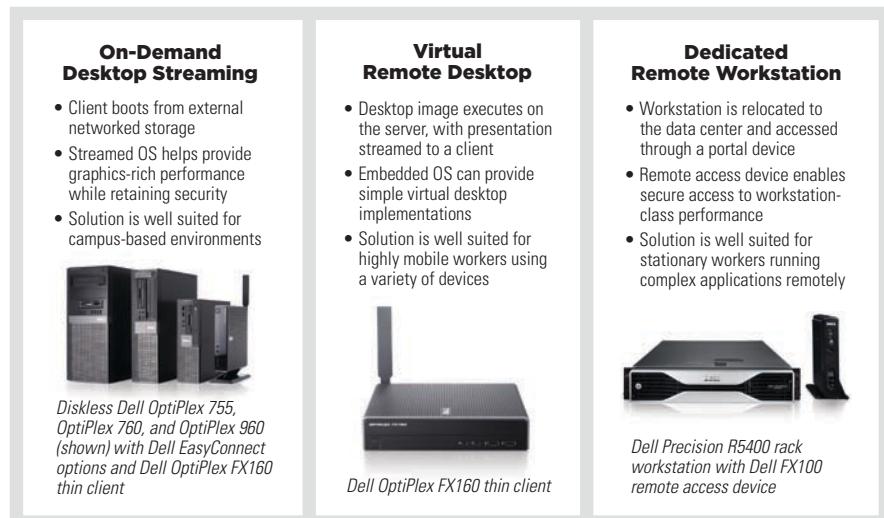


Figure 1. The Dell Flexible Computing Solutions suite helps meet a broad range of organizational requirements

- **Dell On-Demand Desktop Streaming™ (ODDS) solution:** With ODDS, storage is moved from the client to the server. Each client boots from networked storage but retains processor, graphics processing, and user interface functions. ODDS is well suited to campus-based environments in which workers require excellent graphics and processing performance.
- **Dell Virtual Remote Desktop (VRD) solution:** With VRD, storage and processor functionality—the client “desktop”—are offloaded to the server and virtualized, and the client device itself performs only graphics processing and user interface functions. VRD is well suited to highly mobile workers who may use a variety of devices from different locations.
- **Dell Dedicated Remote Workstation (DRW) solution:** With DRW, the entire physical workstation is located centrally, and the display is accessed through a portal device. DRW is well suited to stationary workers running complex applications in which the preferred workstation location may be outside the immediate physical environment—for example, specialized medical applications in an operating room.

To help with the design, implementation, and ongoing support of a flexible computing solution, Dell ICS offers several other services including a one-to-two-day flexible computing workshop—which provides an introduction to flexible computing and the range of solutions available—as well as comprehensive design and implementation services. Dell also offers a wide range of Dell ProSupport services for ongoing support of a flexible computing implementation.

OPTIMIZING FLEXIBILITY AND SESSION MOBILITY WITH VRD

For organizations that must support remote office or mobile workers who may use a variety of client devices, the VRD approach offers optimal flexibility for client devices while retaining server control of the digital identity for each end user. VRD leverages either VMware® Virtual Desktop Infrastructure or Citrix® XenDesktop™ software to enable workers to access centrally managed, server-based virtualized desktops from a range of devices, including traditional PCs, specialized Dell Flexible Computing thin clients such as the Dell OptiPlex™ FX160 desktop,¹ and many other devices such as public kiosks and PDAs (see Figure 2).

¹For more information on the new Dell OptiPlex family, including the OptiPlex FX160 thin client, see “Powering Business Productivity with Dell OptiPlex Desktops,” by David Schweighofer, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090123-OptiPlex.pdf.

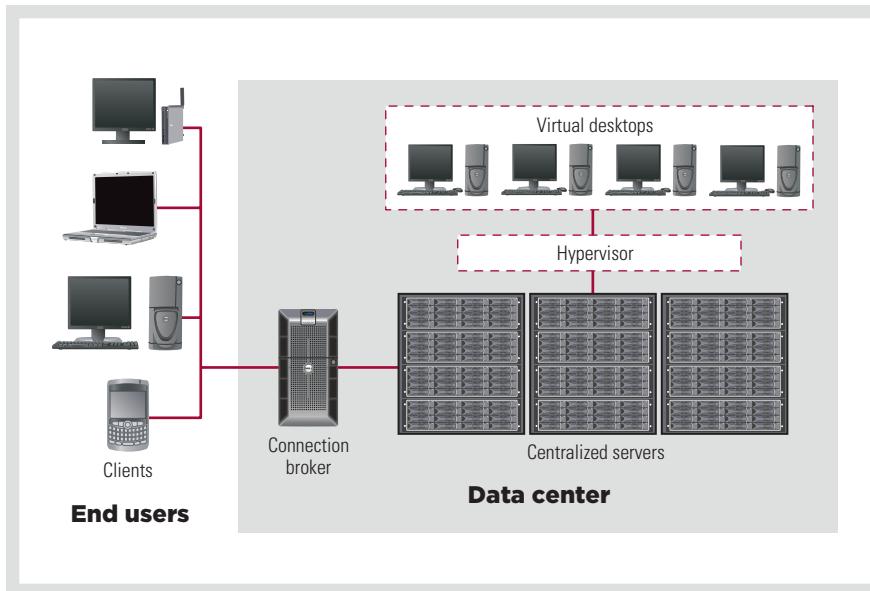


Figure 2. The Dell Virtual Remote Desktop solution offers optimal flexibility for client devices

By enabling access through a wide range of endpoint devices, VRD can dramatically enhance flexibility and mobility for workers on the go. In fact, with VRD, end users can access virtualized desktops from nearly any device or networked PC as long as it is equipped with either a connection broker or Web accessibility. And because VRD enables session mobility, users can move from one device to another without losing state.

In addition to offering enhanced end-user flexibility, the VRD approach helps simplify client management dramatically. By storing and managing virtualized desktops on the server, for example, the VRD approach helps eliminate the need to provision and manage individual client devices. As a result, administrators may provision new desktops in minutes and avoid imaging individual client devices. Also, because applications and data are stored centrally instead of on a client device, basic management tasks such as patches, upgrades, and support can be handled quickly and efficiently.

The VRD approach can also enhance security, simplify regulatory and organizational compliance, and facilitate disaster recovery. For example, because data—including the digital identity of

each user and valuable organizational data—is stored centrally, IT managers can efficiently and effectively control and manage access and help ensure regulatory and organizational compliance. Also, because data is not stored on the client device, security is not compromised in the event of device failure, loss, or theft. And because virtualized desktops are centrally located, they can be backed up and recovered in minutes, helping accelerate recovery time and enhance overall business continuity.

SUPPORTING THE END-USER EXPERIENCE WITH ODDS

For organizations that support end users who work in campus-based environments and require excellent graphics and processing performance, ODDS enables an outstanding user experience while maintaining centralized control. In particular, ODDS allows client PCs to use Dell EasyConnect™ technology to boot and run from networked storage, helping avoid the need for a local disk. Clients can be standard PCs or specialized diskless clients such as Dell OptiPlex FX160, OptiPlex 760, or OptiPlex 960 desktop computers.

PROVIDING REMOTE ACCESS WITH DRW

For organizations that require workstation-class performance in situations in which the workstation is not suited to fit in a specific physical location—for example, using specialized medical applications in an operating room—the DRW solution enables high-performance remote workstation access. In particular, DRW combines the Dell Precision™ R5400 rack workstation with the Dell FX100 remote access device to enable secure, high-performance access to workstation-class performance.

SUPPORTING DIVERSE NEEDS

Dell Flexible Computing Solutions combine leading-edge products and technologies with a comprehensive set of world-class services to help support the evolving needs of increasingly diverse enterprise workforces. By combining centralized control with enhanced client functionality, flexibility, and mobility, Dell Flexible Computing Solutions enable organizations to simplify management, enhance security and recoverability, and lower total cost of ownership. 

Jeremy Ford is a senior technology strategist in the Office of the CTO at Dell.

Roberto Ayala is a senior product marketing manager in the Dell Global Relationship Marketing Group.

MORE
 **ONLINE**
DELL.COM/PowerSolutions

QUICK LINKS

Dell Flexible Computing Solutions:
DELL.COM/FCS
DELL.COM/ICS

Virtualization is only half the battle for efficiency.

InfraStruXure
DATA CENTERS ON DEMAND



Principles of InfraStruXure® High Density-Ready Architecture...

- 1 Rack enclosures that are HD-Ready
- 2 Metered PDUs at the rack level
- 3 Temperature monitoring in the racks
- 4 Centralized monitoring software (not shown)
- 5 Operations software with predictive capacity management (not shown)
- 6 Efficient InRow® cooling technology
- 7 UPS power that is flexible and scalable

Virtualization is here to stay.

And it's no wonder – it saves space and energy while letting you maximize your IT resources. But smaller footprints can come at a cost. Virtualized servers, even at 50% capacity, require special attention to cooling, no matter their size or their location.

- 1. Heat** Server consolidation creates high densities – and high heat – per rack, risking downtime and failure.
- 2. Inefficiency** Perimeter cooling can't reach heat deep in the racks. And over-cooling is expensive and ineffective.
- 3. Power Events** Virtual loads move constantly, making it hard to predict available power and cooling, risking damage to your network.

The right-sized way to virtualize.

With the new HD-Ready InfraStruXure architecture, you can take on high density by cooling the virtualized high-density row, controlling power at the rack level, and managing the system with advanced software and simulation. Though virtualizing saves energy, true efficiency also depends on the relative efficiencies of power, cooling, and servers. Right-sizing one and not the others (See Figure 1) leaves efficiency savings on the table. To right-size, depend on the efficient, modular HD-Ready InfraStruXure and neutralize heat at the source. Equipment will be safe and efficient running closer to 100% capacity.

Don't agonize, virtualize.

What are you waiting for? With HD-Ready InfraStruXure architecture anyone can virtualize... anytime, anywhere. Just drop it in and go.

Why do leading companies prefer InfraStruXure 6 to 1 over traditional data center designs? Find out at www.xcompatible.com*



The following have been tested and work effectively with InfraStruXure Solutions. Go to www.xcompatible.com to learn more.



SQUARE D
by Schneider Electric

t.a.c.
by Schneider Electric

PELCO
by Schneider Electric

the green grid®

*According to the Green Grid, data center efficiency indicates how much data center power is consumed by servers compared to power and cooling infrastructure. If you right-size one without the other, you won't realize your full efficiency potential. See details in APC White Paper #126, "An Improved Architecture for High-Efficiency, High-Density Data Centers".

You can deploy high-density racks right now...

Deploy InfraStruXure as the foundation of your entire data center or server room, or overlay into an existing large data center.

SCHEMATIC LEGEND:

- CRAC UNITS
- STANDARD DENSITY RACKS
- CENTRALIZED UPS
- INFRASTRUXTURE HD-READY ZONES

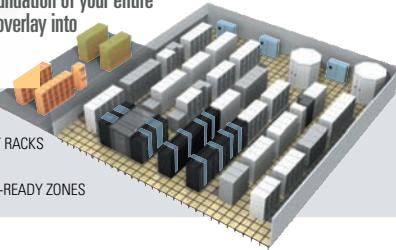


Figure 1

Efficiency and Virtualization

Your servers are efficient, but is your power and cooling?

- COOLING USAGE/CAPACITY
- SERVERS
- POWER USAGE/CAPACITY

Big gains could be made with both server and power and cooling.

- Correct Server Utilization
- Correct-sized Power
- Correct-sized Cooling



Post-Server Virtualization

- Correct Server Utilization
- Correct-sized Power
- Correct-sized Cooling



Server Virtualization with Power and Cooling

Right-sized power and cooling tip the balance back in your favor.

- Correct Server Utilization
- Correct-sized Power
- Correct-sized Cooling



Download a **FREE** copy of APC White Paper #126:

"An Improved Architecture for High-Efficiency, High-Density Data Centers"

Visit www.apc.com/promo Key Code e919w • Call 888-289-APCC x9722 • Fax 401-788-2797

For more information on this APC Solution, please contact your local Dell representative.

©2008 American Power Conversion Corporation. All trademarks are owned by Schneider Electric Industries S.A.S., American Power Conversion Corporation or their affiliated companies.

*As determined by a survey of customers who converted to InfraStruXure architecture.

e-mail: esupport@apc.com • 132 Fairgrounds Road, West Kingston, RI 02892 USA 998-0899_alt

APC
by Schneider Electric



By David Schweighofer

POWERING BUSINESS PRODUCTIVITY WITH DELL OPTIPLEX DESKTOPS

The new Dell™ OptiPlex™ family of desktop computers is designed to deliver high performance, flexibility, and reliability together with enterprise-class security and enhanced manageability in an environmentally conscious way—helping organizations to maximize employee productivity, simplify management, and lower total cost of ownership.

Today's enterprise IT departments are under constant pressure to deliver high performance and advanced functionality to the desktop computers of end users. However, managing enterprise-wide desktop infrastructures can be a challenge. Deploying, monitoring, and managing desktops to help optimize performance and safeguard business continuity can be expensive and time-consuming; protecting enterprise networks and data from malicious threats and unauthorized access can be difficult; and energy costs are skyrocketing.

The new Dell OptiPlex family of desktop computers combines performance, flexibility, and reliability enhancements that are designed to help enterprises address these challenges (see the "Built for business: Dell OptiPlex desktops at a glance" sidebar in this article). The OptiPlex family offers enterprise-class security, advanced management capabilities, and an environmentally conscious design that enable organizations to maximize employee productivity—while helping to protect enterprise data and networks, simplify manageability, and reduce total cost of ownership.

BALANCING PERFORMANCE, FLEXIBILITY, AND STABILITY

Dell OptiPlex desktops are designed to deliver excellent performance and functionality in a variety of form

factors. For example, the chassis of the OptiPlex 960—the flagship model in the family—is designed to be stronger, quieter, and faster to service than previous OptiPlex models. Based on Intel® Core™2 Duo and Core 2 Quad processors, the OptiPlex 960 features up to 8 GB of double data rate 2 (DDR2)-800 RAM and up to four simultaneous independent video displays. Optional solid-state drives help increase performance and reliability, while additional security options are designed to make the OptiPlex 960 more secure than previous generations of the same family. In addition, DisplayPort¹ and optional internal wireless capabilities are available in mini-tower, desktop, and small form-factor chassis.

OptiPlex systems can enhance flexibility for the enterprise because they can be configured easily to help meet individual organizational and employee needs. Systems are built to order at the factory, and Dell Custom Factory Integration (CFI) services are available for advanced customization and individualization. Dell also offers ImageDirect, a Web-based service that enables organizations to securely create, test, manage, and deploy images.

OptiPlex desktops are designed to support Dell Flexible Computing Solutions such as the Dell On-Demand Desktop Streaming™ solution. OptiPlex 960 and OptiPlex 760 models are available in flexible diskless configurations for On-Demand Desktop

Related Categories:

- Dell Client Manager
- Dell OptiPlex desktops
- Flexible computing
- Systems management

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

BALANCING PERFORMANCE, FLEXIBILITY, AND STABILITY

Dell OptiPlex desktops are designed to deliver excellent performance and functionality in a variety of form

¹For more information on DisplayPort technology, see "Mobility and Dynamic Graphics: New Dell Precision Workstations Shine," by Richard Thwaites, in *Dell Power Solutions*, August 2008, DELL.COM/Downloads/Global/Power/ps3q08-20080407-Thwaites.pdf.

Streaming. In addition, the OptiPlex FX160 enables thin client usage models such as the Dell Virtual Remote Desktop solution. The OptiPlex FX160 thin client is designed specifically for flexible computing and can be used with either On-Demand Desktop Streaming or Virtual Remote Desktop.²

OptiPlex desktops are also designed for stability, including support for long life cycles of up to 36 months. In addition, Dell global standard platforms help simplify worldwide standardization of desktop configurations.

SAFEGUARDING INFORMATION WITH MULTILAYERED SECURITY

To help protect sensitive enterprise data and networks, Dell OptiPlex desktops offer comprehensive security features. Starting at the system protection level, anti-chassis-intrusion features help protect internal components from theft. A Chassis Intrusion Switch can give the IT department automatic alerts about theft attempts, and I/O interface security features can flag potential breaches for IT administrators. Intel Trusted Execution Technology provides hardware-based mechanisms that help protect against software-based attacks. In addition, CFI asset tags and security labels can help IT departments track equipment easily.

At the data protection level, optional encrypted hard drives are available for the

OptiPlex 960 model. Dell ProSupport offers hard drive data protection and recovery as well as data destruction services.

Some OptiPlex systems incorporate smart card and biometric readers to help prevent unauthorized users from accessing the system. The OptiPlex 960 model supports secure wireless networking and preboot authentication for enhanced security. As an additional safeguard layer, Dell ControlPoint software³ with Trusted Platform Module (TPM) 1.21 technology supports centralized management of security preferences and accessibility settings—helping enterprise IT managers maintain consistent security policies even in physically distributed environments.

PROVIDING CUSTOM MANAGEABILITY AND SUPPORT

Dell OptiPlex desktops feature a wide range of serviceability, manageability, and support options to help organizations simplify problem diagnosis and resolution, reduce the need for desk-side visits, and streamline deployment and management. Serviceability features include a chassis that does not require tools to access, Dell DirectDetect™ LEDs designed to facilitate efficient hardware diagnostics, an intuitive Dell client configuration utility, and Dell Client Manager™ software from Altiris.

Remote management capabilities also help simplify and automate maintenance

tasks involving OptiPlex desktops. For example, the OptiPlex 760 features basic Intel Active Management Technology (AMT) 5.0 for remote management tasks. Going one step further, the OptiPlex 960—the most advanced member of the OptiPlex family—offers Intel vPro™ technology that helps provide robust hardware-based security and enhanced maintenance and management capabilities. Because these capabilities are built into the hardware, vPro technology enables OS-absent management and increased security even when the desktop is off, the OS is unresponsive, or software agents are disabled.

Microsoft® System Center Operations Manager (SCOM) software provides a simplified monitoring environment that helps IT departments track thousands of servers, applications, and clients to provide a comprehensive view of the health of an organization's IT environment, including OptiPlex desktops. The Dell Management Pack for Microsoft System Center complements SCOM 2007 by extending its management capabilities to OptiPlex operating systems, applications, and other technology components.

Dell ProSupport helps further simplify management of OptiPlex infrastructures with customizable support and service offerings. Asset and data protection services can help guard against losses, while support options such as same-day response



The new Dell OptiPlex family combines high performance, flexibility, and reliability to help maximize productivity, simplify management, and lower total cost of ownership

²For more information on Dell Flexible Computing Solutions, see "Flexible Computing: Advancing End-User Productivity with Centralized Control," by Jeremy Ford and Roberto Ayala, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090139-Flex.pdf.

³For more information on Dell ControlPoint software, see "Mobility Redefined," by Jeanne Feldkamp, Daniel Bounds, Terry Myers, and Tom Kolnowski, in *Dell Power Solutions*, August 2008, DELL.COM/Downloads/Global/Power/ps3q08-20080388-CoverStory.pdf.

and on-site service allow enterprise IT departments to address critical issues with minimal disruption to the business.

DELIVERING ECO-FRIENDLY ENERGY EFFICIENCY

In addition to delivering high performance, outstanding security, and simplified manageability, Dell OptiPlex desktops are engineered to be environmentally friendly. With an energy-efficient processor design, power supplies designed to be up to 88 percent efficient, thermally efficient Balanced Technology Extended (BTX)

chassis, and a range of Dell Energy Smart options, the OptiPlex family can play a role in reducing power consumption in day-to-day operations, helping to conserve energy and reduce total cost of ownership. Remote power-settings management enables administrators to implement energy-saving measures across the enterprise. In addition, Dell is planning Energy Star- and Electronic Product Environmental Assessment Tool (EPEAT)-certified configurations.

OptiPlex desktops also feature environmentally friendly chassis, packaging, and services. The OptiPlex 960 chassis are

designed to incorporate 10 percent post-consumer recycled content and to be brominated flame retardant (BFR) and polyvinyl chloride (PVC) free. Dell packaging is designed to be up to 89 percent recyclable. Dell also enforces strict compliance with worldwide Restriction of Hazardous Substances (RoHS) guidelines, which restrict the use of certain hazardous substances (such as lead and mercury) in electrical and electronic equipment. In addition, optional eco-delivery and asset recovery services are available.

POWERING BUSINESS PRODUCTIVITY

The new Dell OptiPlex family of desktop computers is designed to meet the myriad IT management challenges of enterprise desktop infrastructures. Combining high performance, flexibility, and reliability with an energy-efficient and environmentally friendly design, these desktops can help maximize employee productivity, simplify management, and lower total cost of ownership. 

David Schweighofer is the worldwide out-bound marketing manager for Dell client solutions on the Dell Global Relationship Marketing team. He has a degree in Marketing and Organization as well as an M.B.A. from the HEC Lausanne business school.

BUILT FOR BUSINESS: DELL OPTIPLEX DESKTOPS AT A GLANCE

Designed to optimize performance, flexibility, and reliability while providing outstanding security and enhanced manageability—all in an eco-friendly way—the Dell OptiPlex family of desktop computers offers a range of models that can help enterprises enhance employee productivity, streamline administration tasks, and reduce total cost of ownership.

	Features	Suitability to task
OptiPlex FX160	Intel Atom™ processor; up to 4 GB of RAM; optional 80 GB hard drive; up to 2 GB of flash storage; native VGA and Digital Visual Interface (DVI) support; Altiris® Deployment Solution™ management software	Flexible thin client in tiny desktop form factor offering robust memory and storage with low energy consumption
OptiPlex 360	Intel Celeron® up to Intel Core 2 Duo processor; up to 4 GB of RAM; support for two drives in mini-tower chassis; native VGA, optional DVI, and PCI Express (PCIe) x16 add-in card; Dell Client Manager software	Solid performer offering essential business value with core productivity and reliability features
OptiPlex 760	Intel Pentium® D or Intel Core 2 Duo processors; up to 8 GB of RAM; support for two drives in mini-tower chassis; native VGA and DisplayPort support and optional DVI and PCIe x16 add-in card; Dell Client Manager software; basic Intel AMT 5.0 technology	Mainstream performer offering security, manageability, reliability, energy efficiency, and productivity features
OptiPlex 960	Intel Core 2 Duo or Core 2 Quad processors; up to 8 GB of RAM; support for two drives in mini-tower, desktop, and small form-factor chassis; native VGA and DisplayPort support and optional DVI and PCIe x16 add-in card; Dell Client Manager software; Intel vPro technology	Technologically advanced performer offering security, management flexibility, and reliability features with energy efficiency and long product life cycle

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell OptiPlex:
DELL.COM/OptiPlex

Dell Energy Smart:
DELL.COM/EnergySmart

Dell ImageDirect:
DELL.COM/ImageDirect

Dell OptiPlex Serious Business:
DELL.COM/SeriousBusiness



FROM THE WORLD'S #1 BUSINESS DESKTOP PROVIDER THE NEW OPTIPLEX

NOW MORE
SECURE
THAN EVER



OPTIPLEX SERIOUS BUSINESS

FEATURING THE INTEL® CORE™2 DUO PROCESSOR WITH VPRO™ TECHNOLOGY

SIMPLIFY YOUR DESKTOP COMPUTING AT DELL.COM/OptiPlex

Source: IDC WW PC Tracker Q2, 2008

Intel, the Intel logo, Intel Core, Intel vPro, Core Inside and vPro Inside are trademarks of Intel Corporation in the U.S. and other countries.





DELL PRECISION R5400 RACK WORKSTATIONS OFFER POWERFUL PERFORMANCE AND SECURE REMOTE ACCESS

By Richard Thwaites

Organizations in manufacturing, energy, digital content creation, and other industries require workstation-class performance for compute- and graphics-intensive applications. The Dell Precision™ R5400 rack workstation and optional Dell™ FX100 remote access device enable secure access to applications and data with outstanding performance for an increasingly mobile workforce.

Although desktop systems excel at a wide range of tasks, some compute- and graphics-intensive applications demand more processing power than these systems are equipped to handle. Whether an automaker needs to accelerate the analysis of crash test simulations, an energy firm needs to increase the detail of reservoir models, or a digital production studio wants to enhance the realism of animated characters—organizations are increasingly turning to workstations to achieve performance that surpasses what desktops can offer without incurring the cost or complexity of a large-scale infrastructure.

Such organizations need workstations with high-performance processors, chipsets, and memory to process large data sets quickly. And they require graphics performance that can provide high-resolution models and deliver creative output quickly. Failing to meet those requirements can result in lost productivity, missed revenue opportunities, and failed designs.

At the same time, organizations are challenged to extend workstation-class performance to an increasingly remote workforce. By enabling remote designers, geographically dispersed oil exploration groups, or remote teams of scientists to access

high-performance systems in a centralized location, organizations can optimize resources without having to deploy and manage systems in distant locations. Remote access can also help organizations protect workers and foster productivity by moving employees out of dangerous or difficult work environments such as oil well heads or factory floors. Whatever the goals, achieving secure remote access can be difficult using traditional remote access solutions. Organizations need to provide easy access to shared data and deliver performance that is on par with local workstation access (one-to-one access) while keeping data secure.

Together, the Dell Precision R5400 rack workstation and the optional Dell FX100 remote access device are designed to provide the performance required for compute- and graphics-intensive applications along with secure access for remote employees. The Dell Precision R5400 workstation and Dell FX100 remote access device are also key components of the Dell Flexible Computing Solutions offering, a suite of solutions designed to help organizations meet the demands of distributed workforces while maintaining security and manageability through centralized control of user data and images.¹

Related Categories:

Dell Precision workstations

Flexible computing

Mobility

Performance

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

¹For more information on Dell Flexible Computing Solutions, see "Flexible Computing: Advancing End-User Productivity with Centralized Control" by Jeremy Ford and Roberto Ayala, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090139-Flex.pdf.

“Together, the Dell Precision R5400 rack workstation and the optional Dell FX100 remote access device are designed to provide the performance required for compute- and graphics-intensive applications along with secure access for remote employees.”

DELIVERING POWERFUL PERFORMANCE IN HIGH-DENSITY RACK ENVIRONMENTS

The Dell Precision R5400 rack workstations combines powerful processors and large-memory configurations to help deliver the performance required for compute-intensive tasks. Equipped with a two-socket motherboard, the Dell Precision R5400 can be configured with either dual- or quad-core Intel® Xeon® processors for up to eight total cores. Compared with single-core systems, the multi-core architecture of the Dell Precision R5400 can dramatically increase performance for multi-threaded applications and multitasking environments. The quad-channel, fully buffered dual in-line memory module (DIMM) architecture supports large-memory configurations—with up to 32 GB of total memory—that enable users to work with large data sets.

The Dell Precision R5400 can also provide enterprises in computer-aided design (CAD), animation, energy, and other fields with outstanding graphics performance to view high-resolution models and see creative output quickly. The Dell Precision R5400 is designed to deliver uncompromised OpenGL 3D performance and reliable 2D performance with industry-standard workstation graphics cards. The Dell Precision R5400 offers dual full-height, full-length x16 PCI Express (PCIe) graphics slots to provide a broad range of graphics card choices while also enabling organizations to scale up to high-end 150 W graphics cards for rigorous visualization tasks.

To help optimize workstation resources, organizations can take advantage

of general-purpose graphics processing unit (GPU) capabilities to use the processing power of high-end graphics cards for general-purpose floating-point-intensive computing in a clustered environment. For example, by using the Dell Precision R5400 in conjunction with an established rendering solution such as the multi-processor NVIDIA Tesla GPU computing processor, designed to provide up to a teraflop of parallel computation, organizations can create an energy- and space-efficient high-performance computing (HPC) system.

With a standards-based design, Dell Precision R5400 rack workstations offer standards-based I/O for simple expandability and flexibility for integrating the workstations into an existing IT environment. A 2U rack form factor allows the Dell Precision R5400 to be placed at the desk side or in a dense rack infrastructure. The Dell Precision R5400 workstation can deliver a highly scalable, cost-effective architecture that is easily housed in a rack, which offers a choice of PCI, PCIe, and PCI Extended (PCI-X) slot combinations that help IT administrators and



The Dell FX100 remote access device (center) enables secure remote access to the Dell Precision R5400 rack workstation (left).

support technicians add cards easily. Because the Dell Precision R5400 shares a common design with standards-based Dell PowerEdge™ servers, IT organizations can use Dell server racks and accessories for installation. The Dell Precision R5400 can be configured as part of a high-performance render farm or HPC cluster, enabling enhanced performance without the complexity of deploying and configuring an HPC system.

RUNNING APPLICATIONS THAT REQUIRE OPTIMIZED WORKSTATION COMPATIBILITY

Organizations that run crash test, reservoir modeling, airflow analysis, 3D animation, scientific, or other high-performance applications know that they

Dell also provides software vendors with hardware platforms to facilitate multi-threaded and 64-bit application development. Dell and application providers offer qualified technical support should problems arise.

SECURING REMOTE ACCESS WITH THE DELL FX100 REMOTE ACCESS DEVICE

The Dell FX100 remote access device is designed to help organizations offer remote access to applications and data while enabling remote employees to capitalize on the powerful performance of Dell Precision R5400 rack workstations. Used in conjunction with the Dell Precision R5400, the Dell FX100 enables a cost-effective, flexible, and secure

The Dell FX100 remote access device is designed to overcome the limitations of traditional remote access solutions, providing workstation-class performance with plug-and-play simplicity. It consists of a PCIe card housed in the host workstation and a compact remote portal device. By using Teradici PC-over-IP technology, it helps overcome distance restrictions—allowing organizations to capitalize on standard IP networks to reach anywhere with Internet access. The dedicated hardware performs the compression and encryption algorithms, freeing the host workstation to focus on running applications. The result is enhanced performance with minimized latency, typically limited only by network performance and bandwidth.

The Dell FX100 also provides the flexibility to support numerous types of applications and a wide variety of work environments. Designed to support any host OS and all image content (including 3D graphics, video, animation, Microsoft® ClearType® technology, the Microsoft Aero™ interface for the Windows Vista® OS, and Microsoft DirectX® technology), the Dell FX100 can accommodate a comprehensive range of application requirements. In addition, it supports two Digital Visual Interface (DVI) ports, with one DVI-Integrated (DVI-I)-to-VGA adapter supplied so organizations can use up to two screens per device. The remote unit also offers numerous I/O ports, including four USB ports and a 10/100/1,000 Mbps Ethernet jack as well as audio microphone in, audio headphone out, and audio line out jacks. An on-board audio controller enables use of standard audio codecs.

The Dell FX100 also incorporates security features designed to support remote access while helping ensure protection of intellectual property, assets, and data. The device includes a range of software security features, such as the Trusted Platform Module (TPM) specification and setup and BIOS passwords. Organizations can use USB port disabling,

“By enabling organizations to extend secure access to applications, data, and workstation-class performance to an increasingly remote workforce, the Dell Precision R5400 rack workstation supports a comprehensive range of usage models that can help maximize computing and personnel resources.”

cannot risk application or hardware compatibility problems with workstations. Even relatively minor compatibility issues could result in lost productivity and expensive, time-consuming hardware or software modifications. Critical applications must run successfully and consistently, and software must be backed with qualified technical support.

Dell works closely with application providers and other software vendors to help ensure reliable application performance on Dell Precision R5400 rack workstations. Both Dell and application providers conduct thorough testing to help ensure applications run successfully.

solution for sharing applications, data, and workstation-class performance with remote workers while helping avoid the limitations of traditional remote access solutions.

Many traditional remote access solutions have inherent restrictions that may limit their utility. For example, analog and digital KVM (keyboard, video, mouse) solutions and PCI/PCIe extender solutions can support solid performance, but only at short distances. Software-based KVM over IP works over long distances, but typically consumes valuable host processor cycles and generates latency in the process.

or filtering by device type or by user profile, to offer flexible user authorization methods. An internal front panel enclosure lock helps prevent tampering. With the Dell FX100 and Dell Precision R5400, organizations can keep the workstation in a secure data center or central location while facilitating secure access for a remote workforce.

Additionally, the Dell FX100 portal is a small, fanless device with minimal power consumption requirements. The simple architecture of the portal and host card helps organizations avoid the need for a second client—which would also require a second OS, specialized drivers, a graphics subsystem, and local IT support. By using PC-over-IP technology, the Dell FX100 provides a cost-effective alternative to wired KVM, digital KVM, or Category 5 (Cat 5) approaches.

OPTIMIZING WORKSTATION-CLASS PERFORMANCE AND REMOTE ACCESS

Whether working locally or remotely, professionals in manufacturing, energy, digital content creation, finance, science, software development, and other fields require uncompromising performance for compute- and graphics-intensive applications. The scalable Dell Precision R5400 rack workstation offers a flexible 2U rack form factor and optional Dell FX100 remote access device. By enabling organizations to extend secure access to applications, data, and workstation-class performance to an increasingly remote workforce, the Dell Precision R5400 rack workstation supports a comprehensive range of usage models that can help maximize computing and personnel resources. 

Richard Thwaites is the worldwide out-bound marketing manager for Dell Precision workstations on the Dell Global Relationship Marketing team. He has a degree in Engineering from Coventry University and a postgraduate diploma in Marketing from the Chartered Institute of Marketing.



QUICK LINK

Dell Precision workstations:
DELL.COM/Precision

Reprinted from *Dell Power Solutions*, November 2008. Copyright © 2008 Dell Inc. All rights reserved.

VIRTUALLY SIMPLE

Safe•Connect

Network Access Control

SIMPLIFY NAC AT IMPULSE.COM/DELL

Impulse Point is a Dell Preferred Partner



The Impulse Point and Safe-Connect logos are trademarks of Impulse Point. All other marks and names are trademarks of their respective companies.
Copyright ©2008 Impulse Point. All rights reserved. Unpublished rights reserved under U.S. copyright laws.





By Timothy Sherbak

Cris Benson

EXTENDING BENEFITS OF VIRTUAL REMOTE DESKTOPS USING VMWARE AND DELL EQUALLOGIC SANs

Virtualization can offer increased scalability, reliability, and availability while helping simplify management and reduce operating costs. Now, Virtual Remote Desktop solutions based on VMware® Virtual Desktop Infrastructure and virtualized Dell™ EqualLogic™ Internet SCSI (iSCSI) storage area networks (SANs) can offer similar benefits for the desktop environment—helping reduce the traditional cost and complexity of managing physical laptops, desktops, and workstations.

Prompted by the success of virtualization in the data center and the growing challenges presented by enterprise desktop management, many IT organizations are preparing for the transition to desktop virtualization. Like server virtualization, desktop virtualization requires a virtual storage environment to help deliver the full range of benefits. For example, desktops can become highly available only if they are backed by highly available storage in the form of a storage area network (SAN). Determining the appropriate SAN to optimize the performance, manageability, and scalability of the storage system is particularly important when aggregating hundreds or thousands of desktops.

The combination of VMware Infrastructure 3 virtualization software and virtualized Dell EqualLogic PS Series Internet SCSI (iSCSI) SANs can provide an excellent foundation for virtualizing desktop environments. Virtualized EqualLogic PS Series SANs help extend the capabilities of VMware Infrastructure 3 and simplify virtual desktop deployments, mitigating project risk. iSCSI connectivity helps organizations reduce the cost of infrastructure and take advantage of existing networking knowledge to help lower operational costs. In addition, the PS Series virtualized

storage architecture is designed to provide simple deployment, management, and growth to help reduce ongoing costs.

VMWARE VIRTUAL DESKTOP INFRASTRUCTURE

Organizations evaluating desktop virtualization environments must address several challenges, including escalating support and maintenance costs for physical PCs, frequent security patches and software upgrades, unregulated downloads of personal software, a dramatically increased number of mobile users, and a lack of comprehensive planning for data and application backups. Administrators must support growing numbers of local and remote users while performing the many day-to-day tasks that desktop systems require, such as upgrading applications, refreshing hardware, troubleshooting user problems, and securing data.

VMware Virtual Desktop Infrastructure helps address these challenges in a way that can benefit both the organization and end users. It is designed to provide users with convenient access to virtual desktop systems that are hosted in a central data center but behave just like physical PCs. It also helps

Related Categories:

- Data consolidation and management
- Dell EqualLogic storage
- Flexible computing
- Internet SCSI (iSCSI)
- Storage
- Storage area network (SAN)
- Virtualization
- VMware

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

provide a secure, flexible, and easy-to-manage IT infrastructure designed to be cost-effective and helps reduce the complexity of managing physical laptops, desktops, and workstations (see the “Step-by-step: Sizing the solution” sidebar in this article).

VMware Virtual Desktop Infrastructure leverages VMware Infrastructure 3 to consolidate desktop environments onto data center servers (see Figure 1). It works closely with VMware VirtualCenter to provide advanced virtual desktop management capabilities, such as automatic suspend and resume, that help reduce the memory and processing power required to host virtual desktops. By taking advantage of the capabilities of VMware Infrastructure 3, desktops can run even when server hardware fails and recover quickly from unplanned outages without duplicate hardware.

VMware Virtual Desktop Manager, a key component of Virtual Desktop Infrastructure, can securely connect authorized users to centralized virtual desktops. It is designed to work with VMware Infrastructure 3 to enhance control and manageability while offering a familiar desktop experience. Users can log in to centralized virtual desktop images from a Web browser or through the VMware Virtual Desktop Manager client.

Emerging technology advances are expected to further enhance the power of VMware Virtual Desktop Infrastructure. For example, VMware View Composer desktop image management software, based on VMware vStorage Linked Clones technology, allows administrators to significantly reduce storage requirements while helping to reduce desktop provisioning time from hours to minutes. View Composer is designed to make updating and patching seamless to desktop end users and efficient for IT administrators, while retaining individual users' data and settings. As a result, IT staff can streamline management of thousands of desktops while helping ensure that all user systems are up-to-date.

ESSENTIAL ELEMENTS FOR VIRTUAL DESKTOP ENVIRONMENTS

The combination of VMware Virtual Desktop Infrastructure software and virtualized Dell EqualLogic PS Series iSCSI SANs can help meet evolving enterprise requirements and lay the foundation for an end-to-end virtualized IT infrastructure. By deploying EqualLogic storage arrays, IT departments can circumvent the cost and management complexity of traditional SAN infrastructures without compromising the enterprise-class storage requirements of a virtual desktop infrastructure. Designed for enterprise deployment and bundled with a comprehensive suite of enterprise software features, EqualLogic PS Series SANs offer the elements required for a successful desktop virtualization environment: scalability, high availability, data protection, simplified management, and high performance.

Scalability: Enabling performance and capacity expansion without disruption

VMware virtualization enables organizations to dynamically scale computing

resources by seamlessly adding physical servers into a VMware ESX server cluster. The VMware Distributed Resource Scheduler (DRS) feature enables existing virtual machine (VM) workloads to be automatically and nondisruptively distributed across the new physical resources. EqualLogic PS Series SANs help extend that same dynamic scalability to back-end storage, enabling organizations to easily expand both the server and storage infrastructure as the number of hosted users on the system increases.

Adding storage can be fast and simple. Administrators can expand EqualLogic PS Series SAN resources online without disrupting users (see Figure 2). EqualLogic PS Series SANs added to an EqualLogic group can “learn” the configuration from currently deployed arrays, with the recently added arrays then automatically integrated into the existing virtualized pool of storage. The SAN distributes data volumes across available storage resources, rebalancing workloads automatically to help provide high storage resource utilization and efficiency. As the

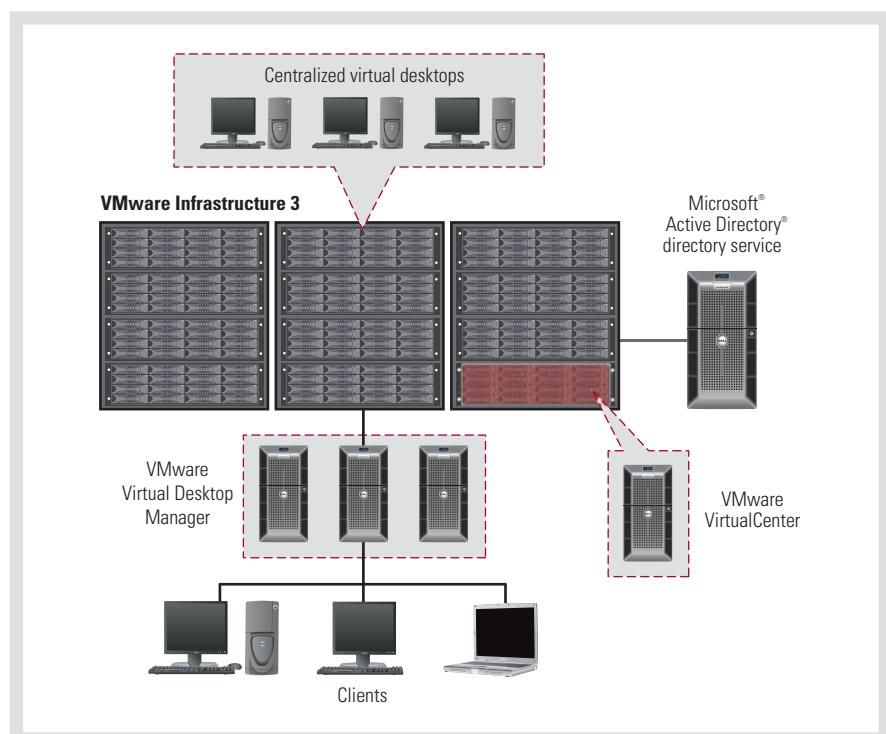


Figure 1. Example environment based on VMware Infrastructure 3 and Virtual Desktop Manager

STEP-BY-STEP: SIZING THE SOLUTION

Sizing a desktop virtualization environment with VMware software and Dell server and storage platforms depends on several key factors. On the client side, the desktop configuration and user workload profile largely determine how many desktops can be hosted on a single VMware ESX server. Administrators must appropriately size the Dell PowerEdge™ servers hosting the desktops and the Dell EqualLogic storage area network (SAN) storing the virtual disks to help ensure that an optimal number of desktops can run with adequate user response times. The following steps can help collect the necessary information:

- 1. Measure physical desktop usage.** Understanding how different applications affect physical processor, memory, network, and storage resources is key. In large desktop environments, VMware Capacity Planner can collect the physical metrics that help estimate virtual machine (VM) requirements.
- 2. Estimate processor requirements.** Administrators can use Microsoft Performance Monitor to measure average processor utilization for the physical desktops, then multiply the measurement by the targeted number of virtual desktops for each VMware ESX server to estimate their processor requirements. Other factors to consider include storage and virtualization, connection protocol, and additional capacity for usage spikes.
- 3. Estimate memory requirements.** Because memory usage can vary based on VMware ESX page sharing, calculating memory requirements requires some estimates. The high watermark can be estimated by ignoring the effect of page sharing: for example, to create 64 VMs with 512 MB of RAM each, a hypervisor without page sharing would require at least 32 GB of RAM. With page sharing on the ESX server, actual memory requirements in production would be less.
- 4. Calculate network requirements.** Dividing the estimated traffic by the speed of the network interface card (NIC) yields the number of NICs typically required. Administrators should also consider the type of NICs used, multimedia applications, shared folders, and the connection broker protocol.
- 5. Estimate storage capacity requirements.** Administrators can estimate the storage capacity requirements for VM system drives using the following formula:

$$(\text{.vmdk file size}) + (\text{VM RAM}) + (\text{size of suspend/resume space for VM}) + (100 \text{ MB per VM for logs})$$

Beyond VM system disk requirements, administrators should also take into account the virtual disks hosting files and data for each user. For performance requirements, they should consider both the I/Os per second (IOPS) and storage throughput data gathered from the physical workload. Given the significant capacity of disk drives today and modest throughput requirements, in practice, administrators should give careful attention to creating a storage configuration with enough disk drives to meet the aggregate IOPS requirements.

Using the information acquired in these steps, IT organizations can create an optimal configuration for virtual desktop environments. As a final step, administrators should deploy the configuration in a proof-of-concept pilot test environment to monitor and validate performance.

SAN is expanded and the added resources are applied to the existing virtual desktop workloads, virtual desktop users can continue their work, operationally unaffected by the infrastructure changes.

High availability: Avoiding downtime

Large-scale enterprise deployments can yield significant cost savings but require a robust feature set and configuration. VMware Virtual Desktop Infrastructure enables organizations to increase the baseline level of availability provided for all applications and provide increased levels of availability easily and cost-effectively. EqualLogic PS Series SANs are designed with fully redundant, hot-swappable components and include a media error detection and correction feature to help preempt disk failure. If a failure occurs, controller failover, RAID protection, hot spare components, and I/O multipathing capabilities can help keep operations running.

Data protection:

Providing built-in capabilities

When desktop resources are consolidated, ensuring that they are fully protected at all times is critical. To help meet this requirement, snapshot, clone, and replication capabilities are included in the software of EqualLogic PS Series SANs to enable advanced data protection for VMware Virtual Desktop Infrastructure environments.

EqualLogic Auto-Snapshot Manager/VMware Edition helps provide space-efficient, hypervisor-aware, SAN-based snapshots and fast recovery of VMs and Virtual Machine File System (VMFS) data stores.¹ An easy-to-navigate graphical user interface displays SAN- and VMware-centric views and provides simple, centrally managed data protection for the entire virtual desktop infrastructure. Through this interface, consolidated

¹For more information, see "How Dell EqualLogic Auto-Snapshot Manager/VMware Edition Helps Protect Virtual Environments," by Andrew Gilman and William Urban, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090107-Gilman.pdf.

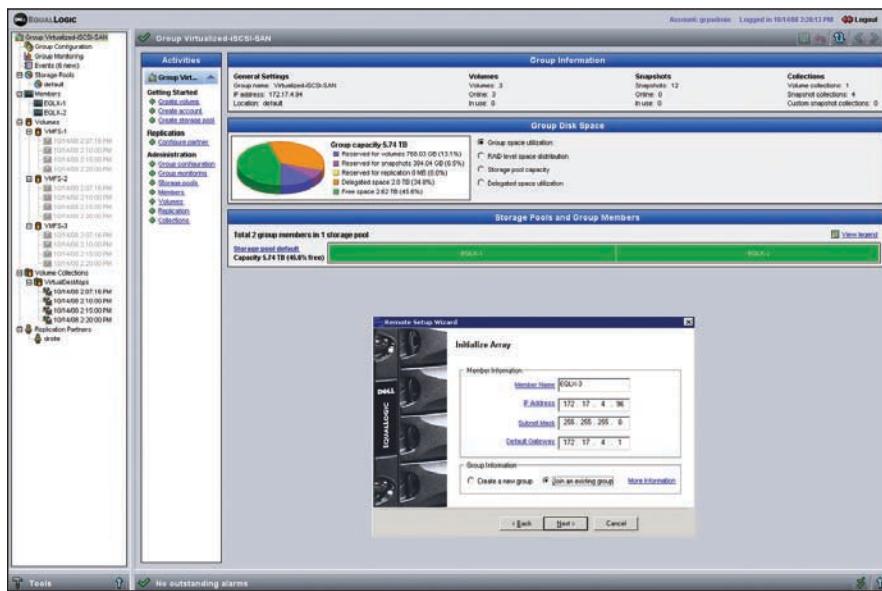


Figure 2. Dell EqualLogic Group Manager and Remote Setup Wizard for nondisruptively adding SAN resources

scheduling of periodic point-in-time copies of virtual desktops in the environment can be achieved using space-efficient SAN-based snapshots for quick online retrieval of the virtual desktop in case of data loss or corruption.

Simplified management: Automating IT tasks

Virtual desktop infrastructures help streamline IT tasks including configuration, deployment, and administration, and EqualLogic PS Series SANs are designed to extend those benefits by simplifying consolidated storage. Storage configuration, management, pooling, and tiering are automated to help save administrative time and expense.

Within the SAN environment, administrators can quickly and easily provision storage with a few button clicks. The key functions needed to configure, manage, and scale storage are automated, helping reduce administration tasks for volume and capacity growth. RAID type and data layout are created automatically at the time of provisioning, then optimized as workload patterns for the recently provisioned storage evolve. By avoiding complex tasks and enabling fast and flexible storage provisioning,

EqualLogic PS Series SANs can help dramatically reduce deployment and ongoing operational costs.

A primary advantage of these features is that storage administration remains stable as the group expands. Administrators perform management tasks at the group level through the intuitive, single-pane EqualLogic Group Manager console. Storage and network I/O load balancing can occur automatically across multiple storage tiers as resources and performance metrics change. Administrators are not required to manually map application data to specific physical devices and controllers. Because the group can automatically tune itself and is managed as a single logical system, the ongoing operational costs of storage management can remain fixed even as storage grows.

High performance: Delivering efficient I/O for aggregated environments

I/O performance and scalability are especially important in virtual desktop environments because increasing numbers of desktop images are sharing storage resources. Individually, virtual desktops have relatively modest I/O requirements—but in aggregate, hundreds or thousands of concurrent desktops can require a

significant storage infrastructure to help maintain responsiveness. Administrators can stripe volumes hosting virtual desktops and their associated data sets across multiple controllers, cache memories, network connections, and disk drives to distribute workloads across available resources for optimal utilization and concurrent processing.

FLEXIBLE, EASY-TO-MANAGE DESKTOP ENVIRONMENT

Moving toward an end-to-end virtualized IT infrastructure that includes servers, desktops, storage, and disaster recovery helps organizations realize the efficiencies and cost advantages of virtualization. Dell EqualLogic PS Series virtualized iSCSI SANs can help organizations take full advantage of the simplicity, security, performance, and efficiency of a VMware Virtual Desktop Infrastructure environment, while also gaining the flexibility to easily grow and change an enterprise-wide virtualized infrastructure based on evolving needs. 

Timothy Sherbak is senior manager of virtualization solutions marketing for the Dell EqualLogic product family.

Cris Benson is a senior technical storage specialist at VMware.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

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

VMware Virtual Desktop Infrastructure:
www.vmware.com/vdi



By Andrew Gilman

William Urban

HOW DELL EQUALLOGIC AUTO-SNAPSHOT MANAGER/VMWARE EDITION HELPS PROTECT VIRTUAL ENVIRONMENTS

Protecting and restoring virtual machines (VMs) can be slow and inefficient, and can take precious server resources away from critical applications. Dell™ EqualLogic™ Auto-Snapshot Manager/VMware Edition enables administrators to quickly and easily create hypervisor-aware snapshots of VMs—helping simplify data management, enhance scalability of data protection and recovery, and increase application performance.

As organizations continue to consolidate their IT infrastructures through virtualization, reliability and recoverability are becoming increasingly important. In particular, because physical servers running multiple virtual machines (VMs) can support more applications than non-virtualized single-application servers, uptime requirements for virtualized servers are typically much higher than those for non-virtualized servers. In fact, consolidated environments are often expected to support operations requiring 24/7 uptime. As a result, rapid recovery and restore times can be critical to maintaining business continuity.

However, data protection and recovery can be a complex process in consolidated virtualized environments, and high uptime requirements typically reduce the time available for backup operations. Although traditional tape- and disk-based methods are essential to comprehensive backup strategies, they can also be slow and disrupt business continuity. For that reason, many organizations are adding a step in their data protection methods by utilizing the built-in capabilities of VMware® virtualization software such as VM snapshots, which provide an additional layer of data protection. Taking a VM snapshot involves placing

the VM in a quiesced state, taking a point-in-time copy, and storing it within the VMware Virtual Machine File System (VMFS). Although this method can be well suited for small deployments, as the virtualized infrastructure grows, advanced tools often become necessary to support increased scalability, automation, and recoverability. Software has been introduced that integrates the advantages of VM snapshots with storage area network (SAN)-based solutions; however, many of these products are complex and require costly add-on licenses.

Dell EqualLogic PS Series Internet SCSI (iSCSI) SAN arrays include EqualLogic Auto-Snapshot Manager/VMware Edition, a data management tool that enables administrators to quickly and easily create and recover hypervisor-aware SAN-based snapshots of VMs. This approach helps simplify data management, enhance scalability and recoverability, and increase application performance. Auto-Snapshot Manager/VMware Edition is part of a comprehensive suite of data management and protection tools available on EqualLogic PS Series iSCSI SANs at no additional charge (see the “Providing multi-layer data protection for virtualized environments” sidebar in this article).

Related Categories:

Data consolidation and management

Dell EqualLogic storage

Storage

Virtualization

VMware

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

DEPLOYING TRADITIONAL SERVER-BASED VM SNAPSHOTS

To enhance recoverability and help maximize server uptime, organizations often choose to augment traditional backup methods with VM snapshots. However, creating and recovering VM snapshots at the server or hypervisor level can present challenges. For example, the creation of server-based snapshots can be slow and may take server resources away from critical applications. Server performance can be negatively affected, especially as the environment grows. And recovery can be slow and negatively affect performance.

Snapshots initiated by a virtualized server can also take up substantial amounts of data storage space. For example, when a VM snapshot is taken through VMware software, the software stores the journal log files in addition to the VM itself. If the server operates continually in this mode, the VMware software eventually fills up the volume, leading to excessive resource constraints on the storage, host processor, and memory.

SUPPORTING SAN-BASED VM SNAPSHOTS

To help maintain the performance of virtualized servers and enhance the scalability of creating and recovering VM snapshots, Dell EqualLogic PS Series iSCSI SANs include EqualLogic Auto-Snapshot Manager / VMware Edition, a data management tool that enables administrators to create and recover hypervisor-aware VM snapshots at the SAN level.

Auto-Snapshot Manager / VMware Edition is a stand-alone Web-based application that integrates directly with the VMware application programming interface to create hypervisor-aware VM snapshots. It takes advantage of the built-in snapshot facilities of EqualLogic PS Series iSCSI SANs and utilizes an intuitive, easy-to-use graphical user interface (GUI) that enables administrators to create VM snapshots on demand or to

PROVIDING MULTILAYER DATA PROTECTION FOR VIRTUALIZED ENVIRONMENTS

In response to the challenge of protecting today's complex virtualized environments, many organizations are exploring a multilayered approach to disaster recovery and data protection—including data protection at the application, virtual machine (VM), physical infrastructure, and data center levels (see Figure A).

For data protection at the application level, Dell EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays include EqualLogic Auto-Snapshot Manager / Microsoft Edition, a data management tool that enables administrators to quickly and flexibly create and recover application-aware snapshots of Microsoft® Windows®, Exchange, and NT File System (NTFS) data as well as standard PS Series snapshots for non-Microsoft operating systems. For data protection at the VM level, EqualLogic PS Series iSCSI SANs also include Auto-Snapshot Manager / VMware Edition, a data management tool that enables administrators to create and recover hypervisor-aware SAN-based snapshots of VMs.

To help maximize the uptime of physical storage components, EqualLogic PS Series iSCSI SANs help achieve 99.999 percent availability, and when used with VMware High Availability (VMware HA) failover protection software help ensure the availability of virtualized IT environments. Finally, to help protect the data center as a whole, EqualLogic PS Series iSCSI SANs include PS Series Auto-Replication and VMware Site Recovery Manager integration to support automated disaster recovery of virtualized data centers.

The comprehensive data protection and disaster recovery tools available in EqualLogic PS Series iSCSI SAN arrays also follow the design tenets—easy to implement, easy to install, and easy to procure—of the arrays themselves. Administrators can implement comprehensive data protection and disaster recovery easily and cost-effectively, and because EqualLogic PS Series arrays are designed for ease of use, organizations can allow additional time and resources for developing their data protection and disaster recovery plans. Together, the data management tools and hardware reliability features provided by EqualLogic PS Series iSCSI SANs can help simplify data management, increase uptime, and facilitate fast, scalable data protection and recovery throughout the enterprise.

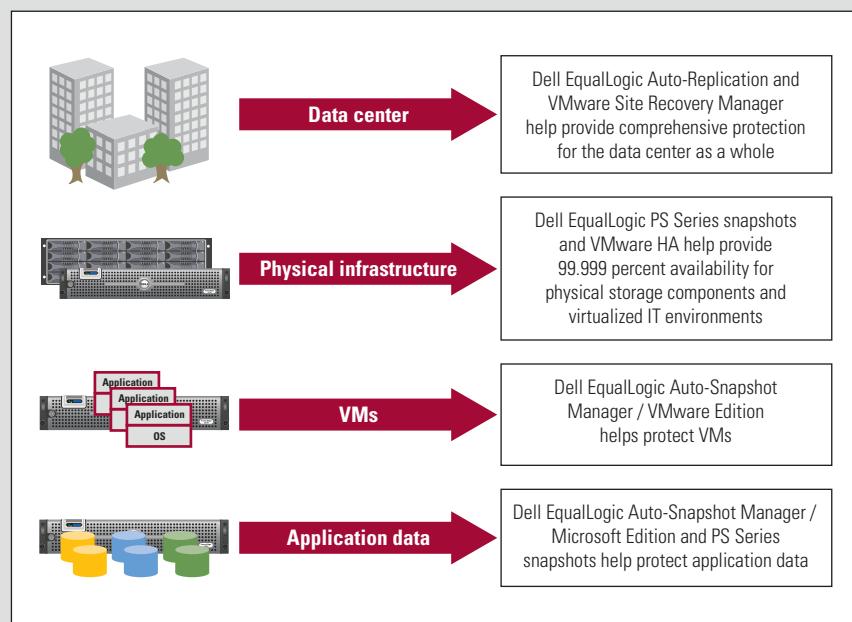


Figure A. Taking a multilayer approach to data protection and disaster recovery in virtualized environments with Dell EqualLogic PS Series iSCSI SANs

HOW PETERSON SULLIVAN PROTECTS THE PAPERLESS OFFICE

Dell EqualLogic Auto-Snapshot Manager/ VMware Edition enables a Seattle accounting firm to protect its data in seconds—rather than hours.

Peterson Sullivan PLLC may have been founded during the Eisenhower era, but today this large Seattle accounting firm uses the latest technology to help its professionals complete work on time and on budget. For example, the firm recently completed one of its largest IT initiatives ever—a multiyear project to convert to a completely digital, paperless office.

To accomplish this goal, the firm consolidated and virtualized its server and storage infrastructures, deploying Dell PowerEdge™ 1950 servers virtualized with VMware Infrastructure 3 software to run a wide array of applications, and a Dell EqualLogic Internet SCSI (iSCSI) storage area network (SAN) to store the thousands of PDF documents scanned in by members of the firm each day. For backup and recovery, the IT group utilized local disk backups and deployed an additional SAN to enable remote data backups and archiving.



“With the previous process, it could take minutes to several hours to back up our virtual machines. With the Auto-Snapshot Manager/VMware Edition, it takes just seconds to protect and recover from hypervisor-aware SAN-based snapshots.”

—Cody Page
IT manager at Peterson Sullivan PLLC
September 2008

Although the IT group already had a solid backup strategy in place, the team decided to explore ways to capture virtual machine (VM) data frequently throughout the day to achieve more granular and faster recovery of the environment. “We started backing up full virtual machine images to the SAN and sending backups off-site for disaster recovery,” says Cody Page, the firm’s IT manager. “The application we were using to create backups worked fine, but the process was lengthy and complicated.”

The Dell Support team recommended evaluating Dell EqualLogic Auto-Snapshot Manager/VMware Edition software. Page’s group was able to start using the new tool quickly. “We had to invest hundreds of hours into our current backup application to get it to work,” says Page. “With the Auto-Snapshot Manager/VMware Edition, we were up and running in about an hour.”

REDUCING TIME TO PROTECT DATA FROM HOURS TO SECONDS

By creating space-efficient point-in-time snapshots of data instead of entire VM images, the IT group expects to accelerate the process of protecting and restoring data while also preserving storage capacity. “The Auto-Snapshot Manager/VMware Edition will help us speed up virtual machine protection dramatically,” says Page. “With the previous process, it could take minutes to several hours to back up our virtual machines. With the Auto-Snapshot Manager/VMware Edition, it takes just seconds to protect and recover from hypervisor-aware SAN-based snapshots.”

Auto-Snapshot Manager/VMware Edition also enables the IT group to automate the creation of snapshots over the course of the day. “If you back up virtual machines just once or twice a day, you stand to lose lots of work in the event of a problem,” says Page. “We can use Auto-Snapshot Manager to create snapshots every hour or two, throughout the day, all without interrupting application availability.”

Auto-Snapshot Manager/VMware Edition also makes it simple to create distinct snapshot schedules for different VMs. “We can use the folder structure to logically group certain virtual machines together and then schedule snapshots for each group at different frequencies,” says Page. “We might want to capture one group every two hours while another once a day. Being able to fine-tune the scheduling helps us optimize our resources.”

Easily integrated into the existing backup process, the Auto-Snapshot Manager/VMware Edition tool is expected to facilitate an added layer of data protection. “We will continue to conduct image-level backups. But now we have a way to capture and restore document changes at a finer level,” says Page.

READYING FOR ONGOING DATA GROWTH

As the document scanning continues, the firm sees no end in sight for its data growth. With a new IT infrastructure and easy-to-use management tools, the IT group is now confident that it can safeguard that data. “Protecting our employees’ work and our clients’ data is essential for our business,” says Page. “We now have the hardware and the management tools to protect data without excessive costs or administrative burdens.”

automate snapshot creation through a built-in scheduler (see Figure 1).

Auto-Snapshot Manager/VMware Edition interprets and preserves VMware folder structures, and can create snapshots at several levels—including snapshots of VMs, VM folders, data stores, clusters, and even entire VMware deployments. To create a VM snapshot, Auto-Snapshot Manager/VMware Edition sets the VM in snapshot mode, takes a SAN-level snapshot of the volume the VM resides on, and then returns the VM to its native production mode.

Once taken, snapshots can be rapidly restored from multiple recovery points. In particular, Auto-Snapshot Manager/VMware Edition has a quick rollback feature for fast recovery and supports automated rollbacks and selective restores using native VMware tools. Snapshots can also be offloaded and used for test and development purposes, including performance analysis, data mining, and testing of software builds, upgrades, and patches.

ENHANCING SCALABILITY AND RECOVERABILITY

By moving snapshot processing from virtualized servers to SANs, Dell EqualLogic Auto-Snapshot Manager/VMware Edition can significantly reduce the load on virtualized servers—which in turn helps free processing power for critical applications and enhance overall server scalability. In addition, because Auto-Snapshot Manager/VMware Edition snapshots are designed to be space efficient, they enable organizations to preserve storage capacity and enhance storage scalability.

Auto-Snapshot Manager/VMware Edition can also help organizations enhance recoverability (see the “How Peterson Sullivan protects the paperless office” sidebar in this article). For example, because Auto-Snapshot Manager/VMware Edition is designed to create snapshots without affecting server performance, organizations can reduce the interval between snapshots to help minimize the risk of data loss and help reduce

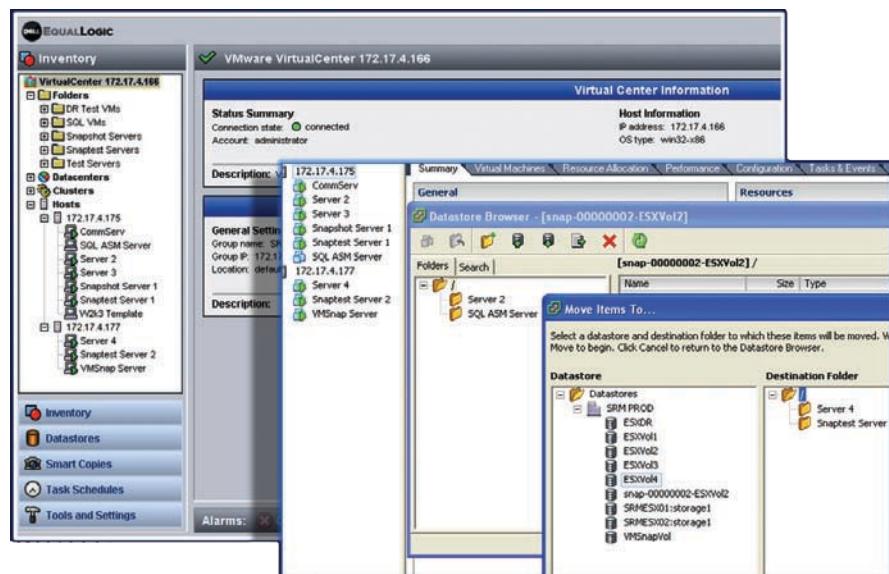


Figure 1. Streamlining administration with the Dell EqualLogic Auto-Snapshot Manager/VMware Edition GUI

recovery times. And because Auto-Snapshot Manager/VMware Edition supports rapid rollback to the most recent snapshot, organizations can recover quickly following a failure, helping avoid the long recovery time associated with recovering from traditional sequential tape media.

Additionally, Auto-Snapshot Manager/VMware Edition can help simplify data management tasks. For example, the simple, easy-to-use GUI enables administrators to create and recover snapshots on demand or schedule automated snapshots. Furthermore, the ability to create snapshots at each level of the VM hierarchy enables administrators to target snapshots for the particular needs of their organizations.

PROTECTING VIRTUALIZED ENVIRONMENTS

For many organizations, VM snapshots have become a vital part of a comprehensive data protection and recovery strategy. By integrating space-efficient SAN-based snapshots with VM snapshots to create hypervisor-aware data protection, Dell EqualLogic Auto-Snapshot Manager/VMware Edition can help organizations not only protect their

virtualized environments, but also simplify data management, enhance scalability and recoverability, and increase application performance. 

Andrew Gilman is a solutions marketing manager at Dell responsible for virtualization marketing activities for the Dell EqualLogic product family. Andrew has a degree in Business Administration from the Boston University School of Management.

William Urban is a product marketing consultant at Dell working on VMware solutions integration for the Dell EqualLogic product family. William has a degree in Information Technology and is a VMware Certified Professional.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell EqualLogic and VMware:
www.equallogic.com/vmware



By Greg Lyon

DELL AND EGENERA DRIVE A NEW PATH TO VIRTUAL DATA CENTER AUTOMATION

The new Dell/PAN System combines Dell™ PowerEdge™ server technology with advanced Egenera® management software in a simplified system—providing a highly available, dynamic, flexible platform for next-generation data center virtualization.

Enterprise data centers are often filled with underutilized hardware systems rigidly tied to specific applications—a result of static, often proprietary legacy architectures that can make IT environments difficult and costly to manage. Hardware and software overprovisioning, inefficient resource utilization, excessive expenses for energy and floor space, and labor-intensive infrastructure management can all contribute to this problem.

Server virtualization has partly addressed some of these challenges through hardware consolidation. The new Dell/PAN System is designed to extend these benefits into next-generation virtualization technology. By creating a processing area network (PAN) that connects pools of compute, storage, and network resources, this system offers a highly available, dynamic, flexible platform to help eliminate the need to dedicate servers to specific applications. PAN Manager® Software by Egenera for Dell, part of the Dell/PAN System, adds sophisticated management capabilities that enable administrators to allocate, monitor, and control Dell/PAN System resources, bringing a critical level of availability to applications.

The flexibility of the Dell/PAN System and PAN Manager enable IT departments to offer automatic high availability and verifiable disaster recovery

without the complexity often associated with static configurations. Together, they help both simplify data center virtualization and greatly reduce capital and administration costs.

UNDERSTANDING THE PAN ARCHITECTURE

A PAN removes physical connections that processors and memory have with each server configuration and application—in other words, it abstracts server and network resources just as a storage area network (SAN) abstracts storage resources (see Figure 1). This design separates the processor and memory from other noncore functions and removes state and identity from the server. Resources typically associated with a server, such as the OS, drivers, network addresses, and SAN logical unit (LUN) IDs, are defined using an XML specification dynamically assigned to a PAN server at boot time.

A PAN virtualizes processing, network, and storage resources, enabling the pooling and dynamic assignment of processing resources to applications. The foundation of the PAN architecture was designed using three components: I/O virtualization, an intelligent communication fabric, and stateless servers.

I/O virtualization. The Dell/PAN System includes redundant PAN controllers that run PAN Manager, store abstract server definitions, and provide physical

Related Categories:

Data center technology

Dell PowerEdge servers

Dell/PAN System

Egenera

Systems management

Virtualization

Visit DELL.COM/PowerSolutions

for the complete category index.

connectivity to networks and SANs. The PAN controllers serve as a physical I/O bridge that consolidates the I/O infrastructure in a data center. I/O between the Dell/PAN System and external storage and data networks travels through ports on the PAN controllers. Servers running in a Dell / PAN System environment have no direct connections to networks and SANs; the I/O is virtualized through PAN Manager. External host addresses, including Media Access Control (MAC) and World Wide Name (WWN) addresses, are abstracted by PAN Manager.

Intelligent communication fabric. The Dell / PAN System connects Dell PowerEdge servers and controller hardware using a high-performance, low-latency fabric. The fabric transports disk and network traffic using a physical switching layer that enables point-to-point connectivity among the control and processing nodes in a PAN. Two fabric switch interconnects help provide redundancy and support data transfer rates of up to 2 Gbps between servers and up to 8 Gbps between the fabric switch and PAN controllers. The resulting network supports higher throughput between servers than a typical data center would have, which can benefit multilayered applications.

Stateless servers. The Dell/PAN System uses only processors and memory on Dell PowerEdge servers. Information

that might typically be associated with physical disk drives, network interface cards, and host bus adapters is placed within the controllers and abstracted through software drivers. This approach is designed to eliminate specific physical identities or states and allow a server to automatically assume the identity of any application at any time.

CREATING A DYNAMIC DATA CENTER

Because the Dell/PAN System comprises compute, network, and storage resources that can be allocated flexibly, it can deliver a dynamic data center that supports server failover, verifiable application recovery for multiple data centers, automated load balancing, rapid server provisioning and reconfiguration, a single interface for managing both physical and virtual servers, and consolidation of servers and I/O connectivity. These advantages are made possible by a variety of capabilities inherent in the PAN architecture.

Built-in networking and I/O consolidation. The Dell/PAN System is designed to support thousands of internal virtual switches, controlled by PAN Manager, that enable large-scale I/O consolidation and help reduce complexity. To connect system resources, a low-latency interconnect is built into the enclosure that is designed to be significantly faster than a

typical enterprise backbone, supporting point-to-point data transfer speeds of up to 8 Gbps. This feature is designed to consolidate I/O resources and provide enhanced throughput and latency performance compared with standard enterprise networks.

Built-in multipath support. PAN Manager integrates multipath support that consolidates and load balances I/O. Multipath support is not typically integrated in legacy architectures, requiring the installation and maintenance of third-party software.

Critical-point redundancy. The physical and virtual components of the Dell/PAN System are designed to be fully redundant, with no single point of failure. The PAN active/active architecture can provide high availability and increased throughput, resulting in exceptional speed and efficiency.

Advanced power and cooling. Because each server in the Dell/PAN System is self-contained, the onboard power supply can flexibly handle increases in processor power requirements. Only cooling active servers helps further reduce power requirements. The rack can have open slots without affecting cooling, and the removal of extraneous components can dramatically enhance airflow across the active servers.

Modular configurations. The Dell/PAN System is modular, providing a consistent form factor that comes in a variety of configurations. Because the servers are hot swappable, the Dell/PAN System supports rapid server replacement. Administrators can increase or decrease processing capacity without having to shut down the Dell/PAN System environment. This modular design facilitates the rapid, cost-effective replacement of failed components and enables processor and OS migrations without a forklift upgrade.

Simplified systems management. PAN Manager offers centralized management of multi-server environments from a graphical user interface (GUI) or command-line interface (CLI) accessible

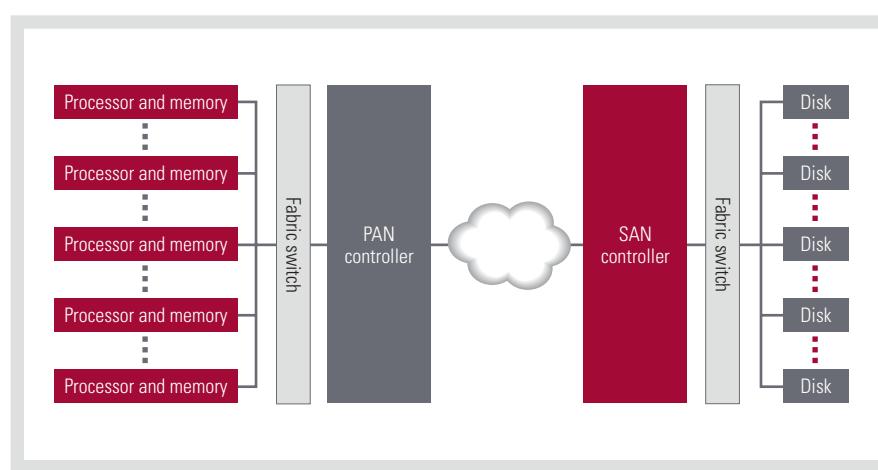


Figure 1. The Dell/PAN System abstracts server and network resources just as a SAN abstracts storage resources

through a Web browser or Telnet client, enabling remote, lights-out control. SANs and LANs are configured once, enabling administrators to reprovision servers without additional SAN or LAN configuration. System event and monitoring capabilities enable the continuous monitoring of system health at both the hardware and application level as well as failover when needed. The robust security model can allocate resources among different organizations while isolating them from one another.

DEPLOYING NEXT-GENERATION VIRTUALIZATION

The Dell/PAN System is designed to maximize the inherent capabilities and performance advantages of the PAN architecture. Its hardware components are housed in an enclosure that helps provide a secure environment for the resident servers as well as maximum visibility, accessibility, and airflow. The package includes up to 24 stateless and anonymous Dell PowerEdge 1950 servers. Two redundant Dell PowerEdge 2950 controllers are included in the enclosure to handle I/O, health monitoring, and dynamic allocation of system resources. The two controllers perform I/O for the system. Fully redundant Dell PowerConnect™ 6248 Gigabit Ethernet switches provide a converged low-latency fabric for disk and network traffic.

Four separate software modules provide the key PAN Manager capabilities: PAN Builder™ Software by Egenera for Dell, PAN Server Portability™ Software by Egenera for Dell, PAN Portability™ Software by Egenera for Dell, and PAN vmBuilder™ Software by Egenera for Dell.

PAN Builder. PAN Builder is the foundation software package for PAN Manager, providing an intuitive GUI that helps administrators configure, control, and monitor PAN resource pools. It includes a chargeback feature for both physical and virtual server environments as well as customizable security management.

PAN Server Portability. PAN Server Portability enables applications and server

images to shift between servers. Because failover resources are shared across production environments, this feature has a number of uses. First, in the event of a hardware or software interruption, PAN Server Portability is designed to automatically restart the failed application on a different server within the time it takes to reboot an OS. Second, applications can be rightsized to organizations' needs and automatically scaled during peak demand by moving them to servers that can provide increased performance. Third, physical servers can be repurposed to run different applications and operating systems at varying times.

PAN Portability. PAN Portability enables whole PAN configurations to move between Dell/PAN System platforms. An XML archive file is stored on the SAN and is automatically replicated between sites using existing storage data replication solutions. As a result, PAN Manager is designed to reduce site recovery times from days or hours to minutes.

PAN Portability enables one secondary site to support multiple primary production sites. And because an entire Dell/PAN System environment can be repurposed quickly and easily, backup resources can be shared with different user groups such as development and quality assurance groups. For example, recovery resources can be repurposed for training and then assume the role of supporting production environments if a site outage occurs.

PAN vmBuilder. PAN vmBuilder integrates Citrix® XenServer™ software to manage physical and virtual servers within a single management domain. Using PAN vmBuilder software helps eliminate the need to learn Xen commands running in the Dell/PAN System, providing administrators with a single environment for configuring, allocating, repurposing, and managing both types of resources. Virtual servers within the Dell/PAN System can take advantage of the same levels of application availability, $n + 1$ failover, dynamic repurposing, and

other critical services available through PAN Manager.

BUILDING A COMPREHENSIVE VIRTUALIZED INFRASTRUCTURE

Creating a data center infrastructure that is responsive to changing organizational needs requires an approach to IT that reduces the time and complexity of deploying new enterprise services while efficiently utilizing available resources and maintaining application availability. The Dell/PAN System is designed to meet these requirements by integrating compute, storage, and networking resources under a single simplified management system. The resulting dynamic data center helps eliminate the time-consuming and error-prone work of statically configuring servers, enabling administrators to flexibly manage resources to help meet the needs of their environment. 

Greg Lyon is the national technology director at Egenera for the Dell/Egenera original equipment manufacturer (OEM) partnership, and has more than 25 years of experience in software development, project management, sales engineering, and business development. He has worked previously at Egenera as a senior sales engineer and sales engineer manager, before which he worked in a technical sales capacity for several startups. He has a B.S. in Computer Science and Mathematics from the State University of New York Maritime College.



MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell and Egenera:
www.egenera.com/dell

A Proven Path to the Dynamic Data Center



Dell/PAN System

Everything you need for a reliable dynamic data center

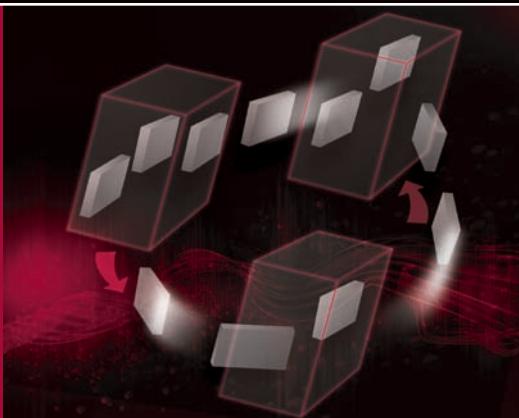
Available today from Dell and Egenera®

Pursue a clear path towards meeting the changing needs of your business. PAN Manager® software by Egenera for Dell is purpose built with reliability, flexibility and simplicity in mind. Ready to deploy on Dell PowerEdge™ servers, it virtualizes and pools data center resources, including virtual and physical servers, into scalable, agile assets. Now you can dynamically and securely allocate and optimize your computing resources as requirements change, all while reducing total cost of ownership. See how PAN Manager can get you on the fast track.

Looking for a solution to quickly get you to your destination?

Go to www.egenera.com/dell and see how easily you can achieve a reliable dynamic data center.





BEST PRACTICES FOR PROTECTING VMWARE INFRASTRUCTURE 3 WITH SYMANTEC BACKUP EXEC

Effectively protecting virtualized environments using traditional backup tools can be both difficult and time-consuming. Symantec® Backup Exec™ 12.5 software, part of the new Dell® PowerVault™ DL2000 – Powered by Symantec Backup Exec, is designed to provide powerful, simplified backup and recovery for VMware® virtualized environments.

Server virtualization is quickly becoming a key element of enterprise data centers, offering advantages ranging from efficient resource utilization to reduced hardware and operational costs. But although data created and utilized in virtual machines (VMs) is typically no less important than data in a single physical server, effectively protecting and recovering data in VMs can be a significant challenge—especially when trying to do so using traditional backup tools not designed for virtualized environments.

Symantec Backup Exec 12.5, part of the new Dell PowerVault DL2000 – Powered by Symantec Backup Exec,¹ introduces several features designed specifically to help protect virtualized environments based on the VMware Infrastructure 3 suite, including the new Agent for VMware Virtual Infrastructure (AVVI). Using these features and following best practices can help organizations overcome some of the challenges of protecting virtualized environments to help implement powerful, simplified backup and recovery.

VIRTUALIZATION BACKUP AND RECOVERY CHALLENGES

As hardware consolidation continues to accelerate, planning backup and recovery strategies for virtualized environments becomes an essential part of data protection. Enterprises are becoming increasingly dependent

on efficient backup and quick recovery of their virtual systems and the host systems they run on to help maintain the productivity and cost savings server virtualization can deliver. These virtual systems include not only the VMs themselves, but also the applications that have been installed on those VMs, which might include Microsoft® Exchange, Microsoft SQL Server®, and Microsoft Office SharePoint® Server software. A failed virtualized server could affect productivity for up to several hours, or even days, in multiple departments while administrators struggle to recover the virtualized environment and the individual VMs.

Protecting a virtualized environment with backup technologies not built specifically for virtualization can be both frustrating and time-consuming, typically requiring the following:

- Installing a backup agent inside each VM or directly on a virtualized server
- Taking time to restore an entire VM just to recover a single file
- Maintaining separate backups for system-level and individual file-level recovery
- Taking VMs offline during backup operations to help provide comprehensive protection
- Ensuring that applications running inside VMs can be recovered

Related Categories:

Backup

Data consolidation and management

Dell PowerVault storage

Symantec

Virtualization

VMware

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

¹For 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.

- Requiring separate backup products for physical and virtual systems

The VMware Infrastructure suite has quickly become an industry-standard virtualization platform. VMware Infrastructure 3 introduced VMware Consolidated Backup (VCB) to help overcome some of these challenges. But VCB can also introduce challenges of its own that administrators should consider before implementing it—including potentially requiring them to manage cumbersome and complicated scripts to integrate with existing backup systems, install proprietary integration modules that require additional testing and setup, and continue maintaining separate backups for system-level and file-level recovery.

SYMANTEC AGENT FOR VMWARE VIRTUAL INFRASTRUCTURE

The Symantec Backup Exec 12.5 AVVI is designed to provide the advantages of VCB, such as off-host backup, while removing some of the challenges of implementing scripted VCB-based backup systems. Beyond basic VCB script-level integration, it provides integration with key VMware application programming interfaces to help avoid the need for VCB scripting or integration modules, helps eliminate separate VCB backups for system-level and file-level recovery from VMs, and helps protect Microsoft Volume Shadow Copy Service (VSS)-aware applications such as Exchange, SQL Server, and SharePoint as part of a VM.

The AVVI does not require an agent to be installed on a VMware ESX server, and no additional configuration is necessary for VCB backups to occur. All support necessary to perform backups of the virtualized environment is included (see Figures 1 and 2).

After administrators have installed the AVVI license on the Backup Exec server, the simplified Backup Exec interface can communicate with VMware VirtualCenter or with individual ESX servers and guide administrators through the process of identifying the necessary ESX hosts, groups, and

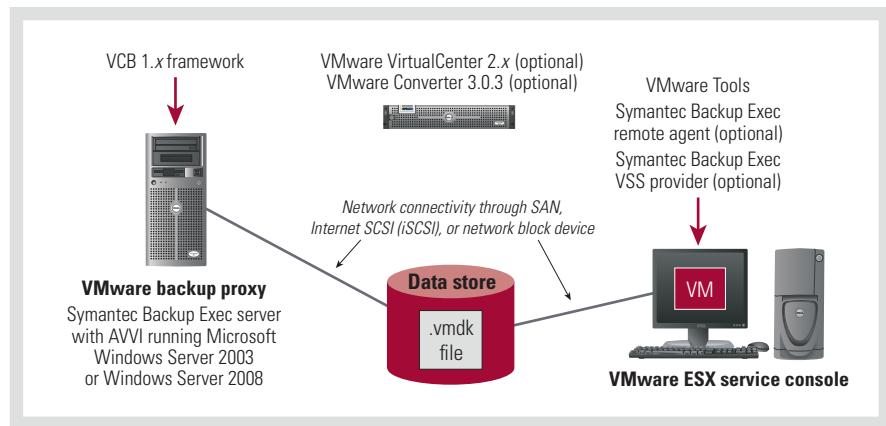


Figure 1. Example environment using the Symantec Backup Exec 12.5 Agent for VMware Virtual Infrastructure

Integration with VMware Infrastructure 3	The AVVI supports and integrates with key VMware tools and features, including VCB, VirtualCenter, VMotion™, VMware Converter, ESX and ESXi, and VMware Tools technologies.
Integration with Symantec Backup Exec	The AVVI can automatically discover both VMware virtualized environments and physical environments to help provide seamless protection for both.
Scriptless VCB integration with Symantec Backup Exec	The AVVI integrates directly into the Backup Exec 12.5 console, and requires neither VCB scripts nor integration modules to help protect a VMware virtualized environment.
Agentless VM backup	Backups can be performed without installing a Backup Exec agent inside VMs or on an ESX server.
Simplified licensing and pricing	A single AVVI license supports protection for an unlimited number of Microsoft Windows and Linux® OS-based VMs on an ESX server.
Embedded GRT	GRT technology enables recovery of individual files and folders inside a Windows-based VM without needing to restore the entire VM.
Application protection with Microsoft VSS	When protecting an entire Windows-based VM, the AVVI can also protect applications through VSS—enabling the entire server and application to be recovered simultaneously.
Flexible recovery	Flexible recovery features enable administrators to restore VMs to their original or alternate data store locations, including specifying a different VM name and virtual network use after the restore.

Figure 2. Key advantages of the Symantec Backup Exec 12.5 Agent for VMware Virtual Infrastructure

VMs for fast, simplified backup and recovery (see Figure 3). The VM and its necessary components—including the .vmdk files, .vmx files, .log files, and .nvram files—are automatically selected for backup.

When administrators must recover an entire VM, they can browse to protected VM systems in the Backup Exec console to restore the entire VM or individual .vmdk files. As an alternative, they can use the built-in Granular Recovery Technology (GRT) in Backup Exec to recover individual files and folders from within a .vmdk file without needing to run a separate backup.

Administrators can restore backups to their original locations or to different locations—including alternate data stores, host ESX servers, and virtual networks—and can use the original VM name or a different name.

BEST PRACTICES FOR USING SYMANTEC BACKUP EXEC 12.5

Best practices such as the following can help ensure successful backup and recovery in VMware environments:

- To perform backups of exported data from ESX servers, install Backup Exec

12.5 for Windows Servers on a Microsoft Windows Server® 2003 OS-based VCB proxy server. If Backup Exec is installed on a separate server from the VCB proxy server, Backup Exec 12.5 for Windows Servers or the Backup Exec 12.5 for Windows Servers Agent for Windows Systems (AWS) must be installed on a Windows Server 2003-based VCB proxy server.

- When performing VCB image-level backups, ensure that sufficient disk space exists on the VCB proxy server for all .vmdk files that will be copied directly to it for off-host backup.
- Install either the Backup Exec VSS provider or the VCB 1.5 VSS requester on a given VM; do not use both.
- Keep in mind that VSS-enabled VCB backups of VMs that contain applications such as Exchange, SQL Server, and SharePoint should not replace traditional application- or database-level backups through Backup Exec application or database agents. VSS-enabled VCB backups do not support application- or database-level full, incremental, or differential backup methods; these backups are copy backups that do not truncate application log files or provide granular application recovery.
- For storage area network (SAN) backups, ensure that the off-host VCB proxy server is zoned properly to enable it to see the Virtual Machine File System (VMFS) logical units (LUNs) used by the ESX servers. VCB mounts a .vmdk file to a directory on the centralized Windows Server-based VCB server and enables backup of the .vmdk file contents.
- To help avoid snapshot-related problems, schedule backups during times of relatively low I/O activity on the VM. Reducing the number of simultaneous backups (and, in turn, VCB snapshots) can help avoid problems as well.
- Upgrade to the latest version of the VMware Infrastructure suite, including the latest versions of ESX, VirtualCenter, and VCB. Updated versions of

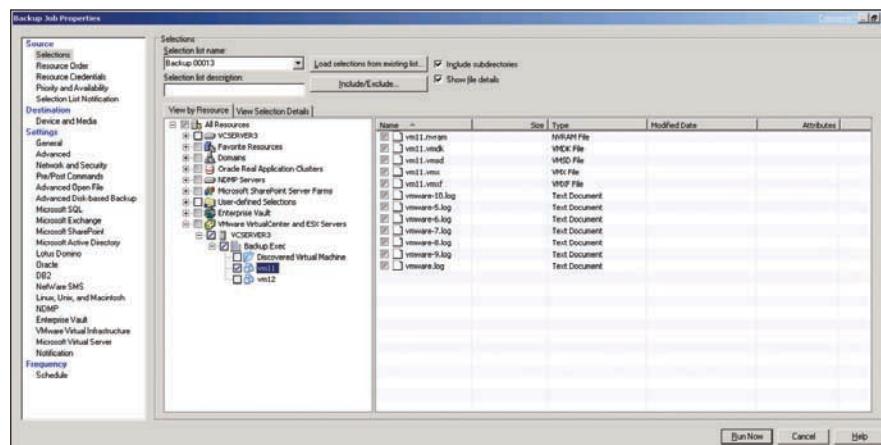


Figure 3. Virtual machine discovery and selection in Symantec Backup Exec 12.5

VirtualCenter components can often enhance VCB snapshot reliability.

- Configure the snapshot mount point over as many dedicated spindles as possible. After a VCB snapshot is created, data is transferred from the VM data store to the backup proxy mount point, and ensuring that the data path from the data store to the snapshot mount point is as fast as possible can help significantly accelerate the completion of the snapshot process.
- Keep in mind that raw device mapping (RDM) disks are not currently supported through VCB backups, and are skipped automatically.

Administrators should also keep the following in mind regarding data recovery:

- VCB does not enable direct recovery of individual files and folders to individual VMs. Administrators must install the Backup Exec AWS on the target VM to perform GRT-enabled recovery of individual files and folders. Alternatively, they can perform a client recovery to a Microsoft Windows® OS-based share, then access the restored files and transfer them to the VM through this share.
- Granular recovery of individual files and folders from within a .vmdk file typically works best when restoring from a disk-based backup. Although granular recovery from a tape-based backup is supported, it requires admin-

istrators to temporarily stage the entire .vmdk file to a disk location during the restore process and then remove it. To recover the .vmdk file, sufficient disk space must exist on the temporary staging location specified in the Restore Job Properties window.

SIMPLIFIED DATA PROTECTION FOR VIRTUALIZED ENVIRONMENTS

The Symantec Backup Exec 12.5 software in the new Dell PowerVault DL2000 - Powered by Symantec Backup Exec introduces a variety of flexible capabilities to help protect VMware virtualized environments as part of an overall backup strategy. Taking advantage of features such as the Symantec Backup Exec AVVI and following best practices can help organizations implement powerful, simplified backup and recovery for their VMware virtualized environments. 



QUICK LINKS

Symantec Backup Exec:
www.backupexec.com

Dell PowerVault DL2000:
DELL.COM/DL2000



In about three seconds, this man will lose what he's working on. But with Continuous Data Protection, he can employ

ATM TASHTLSFSA

(A Time Machine To A Simpler, Happier Time, Like, Say, Four Seconds Ago)

Introducing CA XOsoft™ Replication r12. It's how your data is saved just in the nick of time — all the time. Because it's the first recovery management solution to integrate Continuous Data Protection (CDP) capabilities. So now, historically painful data recovery validation can be done automatically, in mere seconds. And with the optional Interactive CDP Repository, your exchange users can even retrieve their own data back to the last second before it corrupted. And when your users are free to retrieve their own data, you are free to do more important things like run an IT department. View the webcast replay or download the flash at ca.com/partners/xosoft.

For immediate assistance,
contact a Dell Software Specialist
at 1-800-388-2785
or download your FREE trial at:
ca.com/xosoft/dell

GOVERN • MANAGE • SECURE

ca Transforming
IT Management



By Shefali Chinni

Omar Sultan

Nelson Stewart

BOOSTING DATA CENTER THROUGHPUT WITH VIRTUALIZATION-READY 10 GIGABIT ETHERNET

As virtualization becomes pervasive, increasing the number of high-bandwidth network ports is not always enough to meet the bandwidth needs of enterprise data centers. To help support the I/O requirements of multiple virtual machines in server consolidation and virtualization scenarios, Dell, Intel, and Cisco have teamed up to offer organizations a virtualization-tuned 10 Gigabit Ethernet SFP+ direct attach solution.

The 10 Gigabit Ethernet (10GbE) networking standard established in 2002 was, at that time, primarily applicable as part of specific solutions that demanded extremely high bandwidth. Today, however, the broad adoption of server virtualization is helping reshape networking in the data center. Because multiple virtualized servers support heavy workloads on a single physical server, they compete for I/O capacity. This demand often exceeds the multiple Gigabit Ethernet (GbE) connections commonly used for servers today.

The shift toward rack-optimized servers in the data center has changed network architecture. Today, top-of-rack switches are deployed to aggregate network connections from rack-mounted servers. Using rack-mounted equipment has shortened the distance between switch and servers to an ideal length of less than 10 m—helping provide a corresponding decrease in overall costs.

Dell, Intel, and Cisco have responded to the demand for increased I/O bandwidth and cost efficiency with a joint approach that combines virtualization, reliability, and performance features with a cost-effective 10 m cabling reach for connections between rack-mounted servers and the top-of-rack switch. Based on Intel® 10GbE small form-factor plugable + (SFP+) direct attach network adapters and

Virtual Machine Device Queues (VMDq), Cisco Nexus™ 5000 Series Switches, and Intel architecture-based Dell™ PowerEdge™ servers, this ecosystem is designed to provide an outstanding platform for virtualization.

CABLING INTEL ADAPTERS, CISCO SWITCHES, AND DELL SERVERS

Intel adapters and cabling can provide a simplified alternative to multiple GbE server adapters in slot-constrained environments. The Intel 10 Gigabit AF DA dual-port server adapter is a cost-effective, energy-efficient adapter that provides direct attach copper twinaxial cable connection (compliant with the SFP+ MSA SFF-8431 specification) between a server and a top-of-rack switch. The adapter offers two ports and a low-profile PCI Express (PCIe) form factor to provide cost efficiency and energy efficiency for slot-constrained data centers. Designed in a low-profile PCIe form factor, these adapters offer 10GbE network performance with inexpensive SFP+ direct attach copper connectors and cabling up to 10 m.

Intel 10 Gigabit XF SR server adapters, based on the Intel 82598EB 10GbE controller, are also designed to meet the throughput and latency requirements of bandwidth-intensive applications. These adapters are designed to provide a significantly reduced power envelope for additional energy efficiency.

Related Categories:

Cisco Systems

Dell PowerEdge servers

Ethernet

Fibre Channel

Intel

Simplify IT

Virtualization

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

Cisco Nexus switches are another key component of the 10GbE SFP+ networking solution. Designed to use standard SFP+ connectors, Cisco Nexus 5000 Series Switches support SFP+ direct attach Intel 10GbE copper server adapters for in-rack or adjacent-rack use. They integrate transceivers with twinaxial cables into an energy-efficient, cost-effective, low-latency solution, and are designed to use only 0.1 W of power per transceiver—which typically introduces only 0.25 microseconds of latency per link.

Dell PowerEdge servers with dual- and quad-core Intel Xeon® processors form the foundation for the solution. Intel supports the 10GbE SFP+ server adapter on the PowerEdge R900, a four-socket, quad-core 4U rack server; the PowerEdge 2900 III, a two-socket, quad-core 5U rack-mountable tower server; the PowerEdge 2950 III, a two-socket, quad-core 2U rack server; and the PowerEdge 1950 III, a two-socket, quad-core 1U rack server.

OPTIMIZING IN-RACK CONNECTIONS WITH 10GbE ADAPTERS

Intel 10GbE adapters can provide a range of benefits for enterprise data centers. Network interface cards (NICs) come with support for both Microsoft® Windows® and Linux® operating systems. In addition, the SFP+ form factor is typically used for fiber connections, where the optical connector is plugged into the cage to which the cable is connected. Direct attach twinaxial cables feature a factory-terminated connector that slides directly into the cage and offers up to a 10 m length for rack-mounted servers connecting to the top-of-rack switch.

Intel 10GbE adapters can also extend significant cost benefits to enterprise data centers. Because the direct attach cables do not exceed 10 m in length (compared with 100 or 300 m for stand-alone cables) and are already terminated, they are typically not as costly as traditional fiber options—often approaching the same price point as GbE cabling. In addition,

dual-port SFP+ direct attach server adapters help reduce per-port energy consumption and costs.

Intel Virtualization Technology for Connectivity (Intel VT-c) supports Intel 10GbE adapters by providing a set of virtualization-specific enhancements designed to accelerate network I/O performance and reduce processor utilization. Intel VT-c incorporates two key technologies that help reduce latency, accelerate I/O throughput, and offload network overhead tasks for virtualized servers: Intel I/O Acceleration Technology (Intel I/OAT) and Intel VMDq.

Intel I/OAT is available on Dell PowerEdge servers, and it offers a suite of features that help enhance data acceleration across a platform that is independent of the OS on the server. Intel I/OAT includes features such as checksum and segmentation offloads, Message-Signaled Interrupts Extended (MSI-X), receive-side scaling, receive-side coalescing, Intel QuickData Technology, and low-latency interrupts.

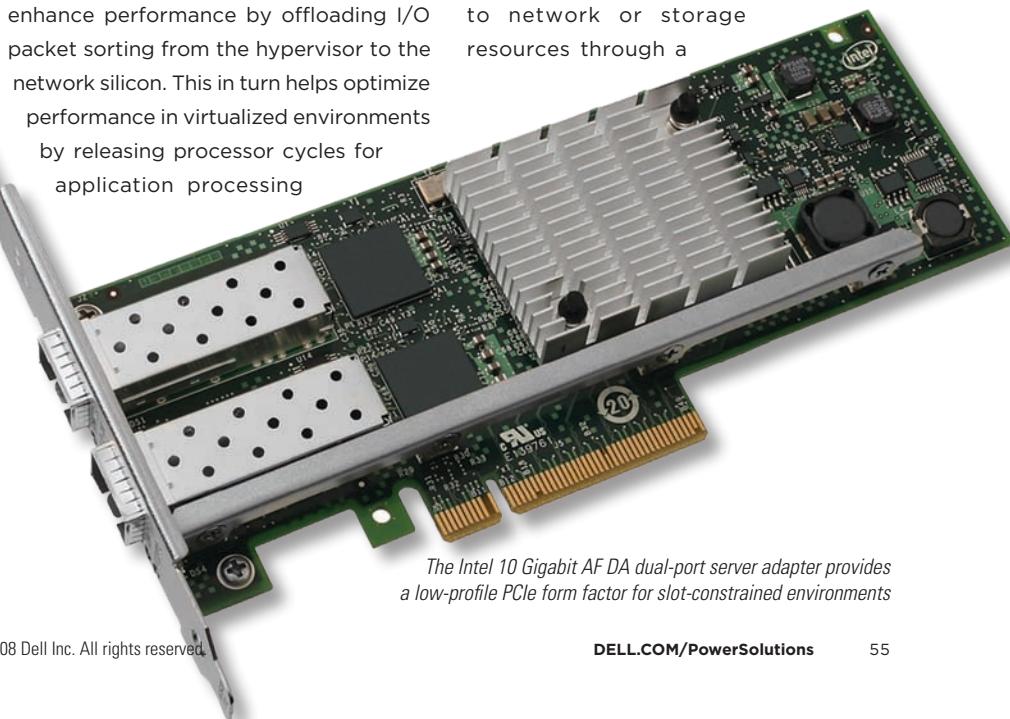
The Intel VMDq networking hardware feature on Intel server adapters provides further acceleration by assigning packets to different virtual machines (VMs) in a virtualized server. The received packets are sorted and collected into queues for the appropriate VM and are then transferred up to the hypervisor, helping enhance performance by offloading I/O packet sorting from the hypervisor to the network silicon. This in turn helps optimize performance in virtualized environments by releasing processor cycles for application processing.

rather than network I/O processing to make efficient use of available resources. Intel VMDq can also handle transmission of packets from the VMs on the host server to enable timely and fair delivery to the network, which helps further enhance speed and throughput.

CONSOLIDATING NETWORK ENVIRONMENTS WITH CISCO NEXUS 5000 SERIES SWITCHES

The Cisco Nexus 5000 Series, part of the Cisco Nexus family of data center-class switches, is part of an innovative architecture designed to enable a high-performance, low-latency, and cost-effective standards-based Ethernet unified fabric. Cisco Nexus 5000 Series Switches facilitate consolidation and enhanced utilization of previously separate infrastructure and cabling, and are designed to reduce the number of adapters and cables by up to 50 percent while helping avoid redundant switches. In this manner, these switches can help simplify cable management and contribute significantly to power usage efficiency and cost-effective cooling—especially for rack-optimized servers.

Cisco Nexus 5000 Series Switches also help create a consistent set of network and storage services across the data center, allowing hosts to connect to network or storage resources through a



The Intel 10 Gigabit AF DA dual-port server adapter provides a low-profile PCIe form factor for slot-constrained environments



Cisco Nexus 5000 Series Switches help consolidate network environments into a single unified fabric (shown: Nexus 5020 switch)

unified Ethernet interface and helping accelerate rollout of new applications and services. They are designed to provide VM-aware services for increased asset utilization, simplified server connections, rapid server provisioning, security, and quality of service. In addition, Cisco Nexus 5000 Series Switches can map VMs to network profiles, allowing network services to be managed at the VM level. This mapping helps enhance manageability, simplify troubleshooting, and maintain compliance. Consistent network and security policies can also be managed centrally over the unified fabric.

The distributed virtual software Cisco Nexus 1000V Series Switches, which are expected to be available in the first half of 2009, help address one contradiction of the evolving data center: the need to exert a fine level of control over bandwidth versus the need to deliver high levels of aggregate bandwidth to the server. In an environment in which the VM is essentially the new building block in the data center, organizations must be able to provision, manage, and troubleshoot network connectivity with VM-level granularity.

The Cisco Nexus 1000V Series is designed to enable IT managers to set and enforce security and connection policies, conduct automated provisioning, and leverage diagnostics features for thousands of live VMs throughout the data center. When available, the Cisco Nexus 1000V Series is expected to help IT administrators easily manage VMs as they migrate them across physical servers during routine hardware maintenance or balance server workloads for optimized application performance and availability.

DEPLOYING VIRTUALIZED DATA CENTERS WITH DELL SERVERS

Dell offers valuable data center benefits in a powerful portfolio of Intel processor-based rack servers, including the Dell PowerEdge R900, PowerEdge 2900 III, PowerEdge 2950 III, and PowerEdge 1950 III servers. Designed to deliver enhanced performance, these highly efficient rack servers optimize space and productivity by utilizing virtualization technologies.

These PowerEdge servers incorporate Intel Xeon processor-based platforms integrating Intel I/OAT. When used together with Intel 10GbE server adapters, I/OAT helps enhance server I/O performance and optimize bandwidth. These capabilities combine to offer excellent platforms for virtualization that help streamline deployment across virtualized infrastructures.

PowerEdge servers with Intel 10GbE server adapters are supported by VMware ESX 3.5 Update 1, which enables organizations to enhance performance across platforms, optimize network traffic management in virtualized servers, and efficiently balance adapter workloads.

BUILDING A PLATFORM FOR HIGH-BANDWIDTH NETWORKING

Escalating deployments of servers with multi-core processors and running applications requiring ever-increasing bandwidth demands—such as virtualization, high-performance computing, database clusters, and video on demand—are contributing to the demand for 10GbE connections in data center server racks. Intel 10GbE adapters enable a clear upgrade and migration path to high-performance

networks. And with direct attach copper twinaxial cabling, platforms comprising Intel 10GbE AF DA dual-port server adapters, Cisco Nexus 5000 Series Switches, and Dell PowerEdge rack servers allow organizations to take advantage of an ecosystem

tuned for high-bandwidth networking—and at an attractive price point. 

Shefali Chinni is the marketing manager for future I/O virtualization technologies and future Gigabit Ethernet silicon at Intel. Before joining Intel, she held various product marketing management and strategic planning positions at Hewlett-Packard. She has an M.B.A. in Marketing from Osmania University.

Omar Sultan is the solution manager for data center switching for the Data Center Solutions team at Cisco Systems. Omar has over 20 years of experience in the IT industry and has been an active Cisco Certified Internetwork Expert (CCIE) since 1994.

Nelson Stewart is a storage and networking product manager for enterprise NICs and blade server I/O, and has over 29 years of experience in the IT industry. Before joining Dell, he was a member of Hewlett-Packard's BladeSystem Business team.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Intel Virtualization Technology:
www.intel.com/go/vtc

Cisco Nexus 5000 Series Switches:
www.cisco.com/go/nexus5000

Intel adapters for Dell servers:
DELL.COM/SNP

How long will. Users.

Struggle with slow. Access.

To Web apps. Before they. Complain?

Stop the. Madness.

Citrix NetScaler accelerates Web apps up to 5x.



Now you can accelerate delivery of Web apps up to 5x to all users—anytime, anywhere—with Citrix® NetScaler®. With integrated technologies like layer 4-7 load balancing, advanced compression

and caching, and TCP optimization, Citrix NetScaler significantly speeds your Web applications, while it improves availability and security, and lowers data center costs.

So don't let another day of slow Web applications get you down. Learn how to dramatically accelerate the Web apps your users depend on.



Learn how to accelerate your Web apps by up to 5x. Get the guide at www.citrix.com/NetScaler

CITRIX®



By Aziz Gulbeden
Mausmi Kotecha

RUNNING MICROSOFT WINDOWS HPC SERVER 2008 ON DELL POWEREDGE CLUSTERS

The Microsoft® Windows® HPC Server 2008 high-performance computing (HPC) platform is designed to simplify the deployment, configuration, and management of HPC clusters while also integrating multiple performance enhancements. Running this platform on clusters of Dell™ PowerEdge™ servers can provide a highly available, high-performance foundation for HPC applications.

Microsoft Windows HPC Server 2008 is designed to provide a simplified, highly available platform for high-performance computing (HPC) cluster environments. Built on Microsoft Windows Server® 2008 64-bit technology and integrating multiple enhancements to help increase network performance, it includes both Windows Server 2008 HPC Edition as the cluster OS and Microsoft HPC Pack 2008 to provide the necessary cluster utilities for running and managing HPC applications. By deploying this platform on clusters of Dell PowerEdge servers, cluster administrators can minimize the time required for cluster management and focus on running HPC applications productively.

Related Categories:

Dell PowerEdge servers
High-performance computing (HPC)
Microsoft
Microsoft Windows HPC Server 2008

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

SIMPLIFIED CLUSTER CONFIGURATION AND MANAGEMENT

Microsoft Windows HPC Server 2008 includes HPC Cluster Manager, an integrated cluster management utility that helps simplify cluster deployment, configuration, and management, including key HPC tasks such as setting up head nodes and managing compute nodes. It incorporates wizard-based configuration tools, compute node templates, node monitoring and management tools, job monitoring and management tools, integrated diagnostic and

reporting utilities, and the Windows PowerShell™ command-line shell.

Wizard-based configuration tools. Administrators can use wizards to perform many initial configuration tasks in HPC Cluster Manager. When the application launches for the first time, it displays a To-Do List screen showing the wizards available for required configuration steps (see Figure 1). After the cluster is configured, administrators can perform cluster management tasks through corresponding Configuration, Node Management, Job Management, Diagnostics, and Charts and Reports panes of HPC Cluster Manager.

Compute node templates. Compute node templates specify the operations to be performed during compute node provisioning. To help simplify compute node management when compute nodes have different hardware configurations or different roles, Windows HPC Server 2008 allows administrators to create multiple compute node templates. These templates do not need to include an OS; administrators can include the OS as part of the template or install it separately, depending on their needs. During OS deployment, the templates specify whether the copy operation is to be performed through multicast, which helps reduce overall network traffic and accelerate

the simultaneous installation of multiple compute nodes.

Administrators can customize templates using the template editor, including installing additional patches or applications or running additional commands during deployment. Deployments are based on Windows Imaging (WIM) images, which administrators can create from the installation media. HPC Cluster Manager also allows administrators to inject drivers into these images. For advanced setups, administrators can also create images using the Microsoft Windows Automated Installation Kit. As updates for compute nodes become available, administrators can apply them using a node template.

Node monitoring and management tools.

The Node Management pane in HPC Cluster Manager includes different views of the compute nodes for monitoring and lists the available actions administrators can perform on cluster nodes. A heat map allows real-time monitoring of selected metrics on the cluster nodes, which can be grouped and filtered for management or job scheduling.

Job monitoring and management tools.

Administrators can schedule or monitor jobs using the Job Management pane in HPC Cluster Manager. A separate job scheduler application is available for nodes that are not part of the cluster but will be submitting jobs. Only domain users identified as cluster users are allowed to submit jobs, and only users identified as cluster administrators can perform administrative tasks such as canceling a job submitted by another user.

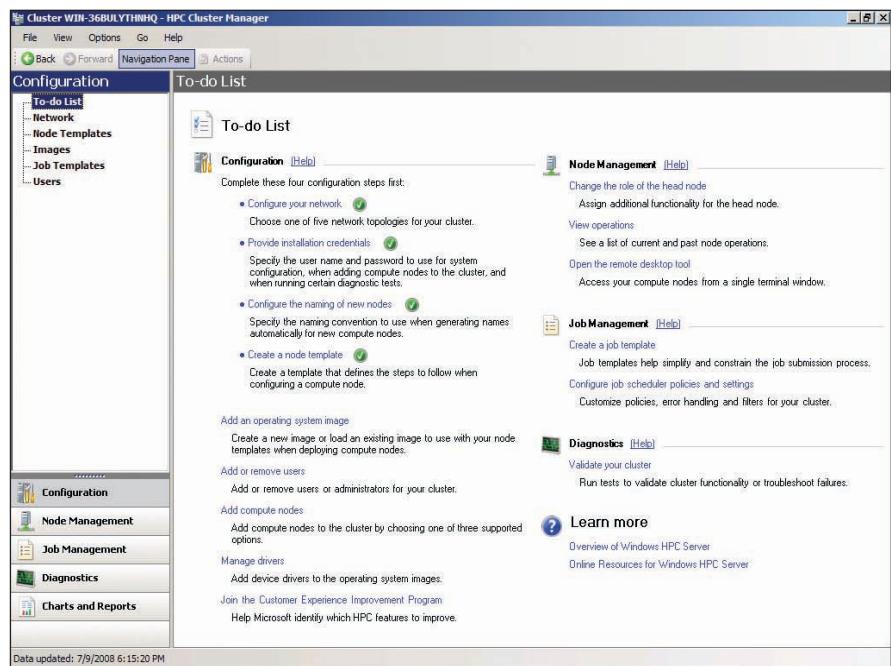


Figure 1. Wizards in Microsoft HPC Cluster Manager help simplify initial configuration tasks

The job scheduler also includes policies for service-oriented architecture applications and adaptive multilevel resource allocation, and allows administrators to launch jobs on a large number of nodes without any additional penalties compared with small jobs.

Integrated diagnostic and reporting utilities.

HPC Cluster Manager provides a variety of tools for monitoring cluster performance and node health and to help troubleshoot the cluster if problems arise. The tests are categorized as scheduler tests, services tests, connectivity tests, system configuration tests, and service-oriented application tests. The reports monitor cluster metrics such as job throughput, cluster processor usage,

cluster network usage, node availability, and job turnaround.

Windows PowerShell. Windows PowerShell offers comprehensive management utilities through a command-line shell, enabling administrators to use commands or scripts to help effectively perform cluster operations.

PERFORMANCE AND AVAILABILITY ENHANCEMENTS

Microsoft Windows HPC Server 2008 includes multiple enhancements to performance and availability, including Microsoft Message Passing Interface (MS MPI), event tracing, the Network Direct interface, a variety of networking changes and enhancements, and head node high availability.

Microsoft Message Passing Interface.

Windows HPC Server 2008 comes with MS MPI, which provides the networking interface to the cluster applications and is compatible with the MPICH2 reference implementation. MPI-1 and MPI-2 are standard messaging interfaces defined for cluster applications that implement the required communication operations used by parallel applications. Multi-core

“By deploying Microsoft Windows HPC Server 2008 on clusters of Dell PowerEdge servers, cluster administrators can minimize the time required for cluster management and focus on running HPC applications productively.”

processor-based systems can benefit from the enhancements to shared memory communications in MS MPI.

Event tracing. To help increase cluster application performance, MS MPI includes integrated event tracing that compilers can use to optimize the code for a specific cluster. Event tracing helps simplify developer tasks by enabling them to tune performance and debug the application if they encounter any errors.

Network Direct. Network Direct is a Remote Direct Memory Access (RDMA) networking interface that uses a more direct path than Windows Sockets (Winsock) Direct to support networking hardware. This approach helps increase the performance and efficiency of MPI applications running over high-speed network fabrics.

Networking enhancements. Windows Server 2008 includes many networking changes and enhancements. The updated implementation of the TCP/IP stack included in Windows Server 2008 is called the Next Generation TCP/IP stack and is based on current networking technologies.¹ Key features of this stack include the following:

- **Receive Window Auto-Tuning:** The receive window size sets the maximum number of packets that have been received but not acknowledged. Receive Window Auto-Tuning allows the OS to dynamically adjust the receive

window to an optimal size based on the network condition, enabling increased network throughput between cluster nodes and increased network utilization during data transfer.

- **Compound TCP:** Compound TCP tunes the amount of data sent at a time in a packet by monitoring the bandwidth-delay product, network variations, and packet losses, with an increase in packet size that helps accelerate data transfer. Together, Receive Window Auto-Tuning and Compound TCP can help increase network utilization and optimize network performance for HPC applications.

- **Explicit Congestion Notification (ECN):** Network congestion can potentially cause dropped packets. ECN support for TCP enables routers experiencing congestion to mark their packets, with TCP peers then reducing their transmission rate on receiving marked packets to help avoid packet loss.

- **Server Message Block (SMB) 2.0:** Windows uses SMB for file sharing. Many HPC applications use file shares to communicate input, output, or intermediate result files. SMB 2.0 supports a larger buffer size than SMB 1.0, allows multiple SMB commands to be sent at once, and has increased limits on the number of concurrent open file handles.
- **TCP Chimney Offload:** TCP Chimney Offload delegates TCP traffic processing to TCP/IP Offload Engine (TOE)-capable

network adapters, helping free processor cycles for other application tasks.

- **Receive-side scaling:** Receive-side scaling distributes the processing of incoming traffic among multiple processors, enabling multi-core processor-based servers to handle incoming traffic more quickly than they could otherwise.

Head node high availability. Windows HPC Server 2008 enables administrators to set up the head node in a failover cluster configuration for high availability. If the cluster head node fails, the standby server in the failover cluster becomes active and then serves as the head node, enabling submitted jobs to continue running. Other head node tasks also migrate automatically to the standby server after a short pause. This high-availability configuration requires that the head node be running the Microsoft Windows Server 2008 Enterprise Edition or Datacenter Edition OS as well as the Microsoft SQL Server[®] 2005 database platform.²

NETWORKING PERFORMANCE ON DELL POWEREDGE SERVERS

Dell supports Microsoft Windows HPC Server 2008 on Intel[®] Xeon[®] processor-based Dell PowerEdge 1950 III and PowerEdge 2950 III servers as head nodes and on Intel Xeon processor-based PowerEdge 1950 III, PowerEdge 2950 III, and PowerEdge M600 servers as compute nodes. Dell also supports Windows HPC Server 2008 on AMD Opteron[™] processor-based PowerEdge 2970 servers as head nodes and on AMD Opteron processor-based PowerEdge SC1435, PowerEdge M605, PowerEdge M805, and PowerEdge M905 servers as compute nodes. For cluster storage, Dell PowerVault[™] MD1000, PowerVault MD3000, and PowerVault MD1120 RAID arrays in direct attach storage mode can be attached to the head node.³

“Using Microsoft Windows HPC Server 2008 on clusters of Dell PowerEdge servers can enable organizations to deploy highly manageable, highly available, high-performance clusters in their HPC environments.”

¹For more information about TCP/IP networking enhancements, visit technet.microsoft.com/en-us/library/bb726965.aspx.

²For more information, see the Windows HPC Server 2008 documentation.

³For more information on supported configurations, visit DELL.COM/HPCC.

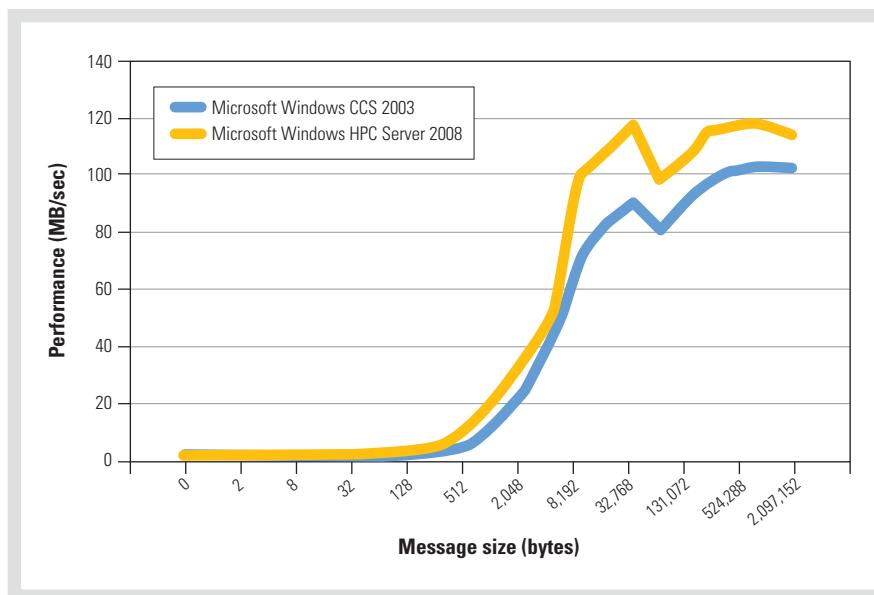


Figure 2. Microsoft Windows HPC Server 2008 provides increased network performance over Windows Compute Cluster Server 2003

To evaluate the networking performance enhancements in Windows HPC Server 2008, in July 2008 Dell engineers tested an example cluster consisting of 16 Dell PowerEdge 1950 servers, each configured with two quad-core Intel Xeon E5450 processors at 3.00 GHz, 4 GB of RAM, and Gigabit Ethernet interconnects. The tests were designed to compare networking performance on these nodes when running the Windows Compute Cluster Server (CCS) 2003 platform and when running Windows Server 2008 Standard Edition as the OS with Microsoft HPC Pack Release Candidate 1 (RC1)⁴ as the cluster package.

The tests were based on the SendRecv parallel transfer benchmark in the Intel MPI Benchmarks 3.0 suite,⁵ which is based on MPI_SendRecv and is well-suited for measuring bidirectional bandwidth. The processes in this benchmark form a virtual ring, with each process sending messages to the process on the right and receiving messages from the process on the left in the communication chain. For each process, the turnover

count is two messages per sample (one send and one receive). The test runs with varying message lengths starting at 0 bytes, then increasing to 1 byte, then doubling up to 4,194,304 bytes, with timings averaged over multiple samples.

Figure 2 shows the results, which demonstrate that Windows HPC Server 2008 performed better than Windows CCS 2003 for all message sizes. Performance did dip at around the 128 KB message size, which was attributable to the MPI protocol changing from the Eager protocol to the Rendezvous protocol. The Eager protocol transfers the MPI header and the message without waiting for the receiver to be ready, and is typically suitable for small message sizes. As message size increases, however, MPI switches to the Rendezvous protocol, in which an initial handshake occurs and the sender waits for the receiver to be ready with a buffer. Administrators can modify the point at which this switch occurs using the MPICH_SOCKET_EAGER_LIMIT environment variable; the default value for this variable is 128 KB.

SIMPLIFIED, HIGHLY AVAILABLE COMPUTE CLUSTER PLATFORM

Microsoft Windows HPC Server 2008 is designed to greatly simplify cluster management, helping reduce the time administrators spend managing a cluster and enabling them to focus on productively running HPC applications. Using this platform on clusters of Dell PowerEdge servers can enable organizations to deploy highly manageable, highly available, high-performance clusters in their HPC environments. 

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

Mausmi Kotecha is an engineering analyst in the Enterprise Solutions Group at the Dell Bangalore Development Center. Her areas of interest include HPC clustering packages, performance analysis of parallel applications, and cluster file systems. Mausmi has a bachelor's degree in Computer Science from Atmiya Institute of Technology and Science.

ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell HPC cluster solutions:
DELL.COM/HPCC

Microsoft Windows HPC Server 2008:
www.microsoft.com/hpc

⁴ Available at connect.microsoft.com.

⁵ Available at www3.intel.com/cd/software/products/asmo-na/eng/219848.htm.



DEPLOYING HPC CLUSTERS USING PLATFORM OPEN CLUSTER STACK – DELL EDITION

By Sreeram Vedantham

Arun Rajan

Shivaraj Nidoni

Li Ou, Ph.D.

Gabor Samu

Mark Black

As high-performance computing (HPC) environments continue to grow in size and complexity, cluster deployment and management have become increasingly challenging. Platform™ Open Cluster Stack – Dell Edition helps address these challenges with an open source, modular, hybrid architecture designed to easily integrate third-party software and help simplify deployment and management for cluster administrators.

High-performance computing (HPC) environments often comprise a huge data center housing thousands of compute nodes designed to incorporate powerful multi-core processors. To help them deploy and manage these complex environments, administrators often use multiple open source tools for different aspects of administration. Integrating these tools seamlessly, however, can be time-consuming and difficult.

Platform Open Cluster Stack (OCS) – Dell Edition is designed to simplify this process. By providing a unified framework that supports a variety of tools and the seamless integration of third-party products—including both open source software and Dell-specific tools such as the Dell OpenManage™ suite—this comprehensive software stack can help administrators efficiently deploy and manage Linux® OS-based HPC clusters even in highly complex environments.

Related Categories:

High-performance computing
(HPC)

Linux

Platform Computing

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

the consistent deployment and management of scale-out Linux clusters. Integrated components include the Platform Lava job scheduler, the Platform Load Sharing Facility (LSF®) application (optional), Mellanox OpenFabrics Enterprise Distribution (OFED) (for Dell™ PowerEdge™ servers using Mellanox InfiniBand cards), QLogic OFED (for Dell PowerEdge servers using QLogic InfiniBand cards), Platform OFED, Message Passing Interface (MPI) (MVAPICH and Open MPI), Intel® runtime libraries, and so on. The latest version of the Platform OCS – Dell Edition software stack is based on Red Hat® Enterprise Linux 5 Update 1 and Platform OCS 5.1.

Architecture

Platform OCS – Dell Edition is designed as a comprehensive HPC software stack (see Figure 1). Kusu, the middle layer of the stack, comprises the tools that build the cluster and acts as a cluster toolkit foundation for the stack. The bottom layer, which Kusu makes use of, contains different protocols for a variety of functionalities. The upper layer includes the cluster middleware technologies such as job schedulers, MPI applications, external cluster health-monitoring tools, and end-user applications. These layers are closely tied to and integrated within the Kusu framework.

INTRODUCING THE KUSU-BASED PLATFORM OCS – DELL EDITION STACK

Platform OCS – Dell Edition is an open source, modular, hybrid software stack that integrates other open source software as well as commercial software to form a unified cluster environment. This seamless integration of different components helps support

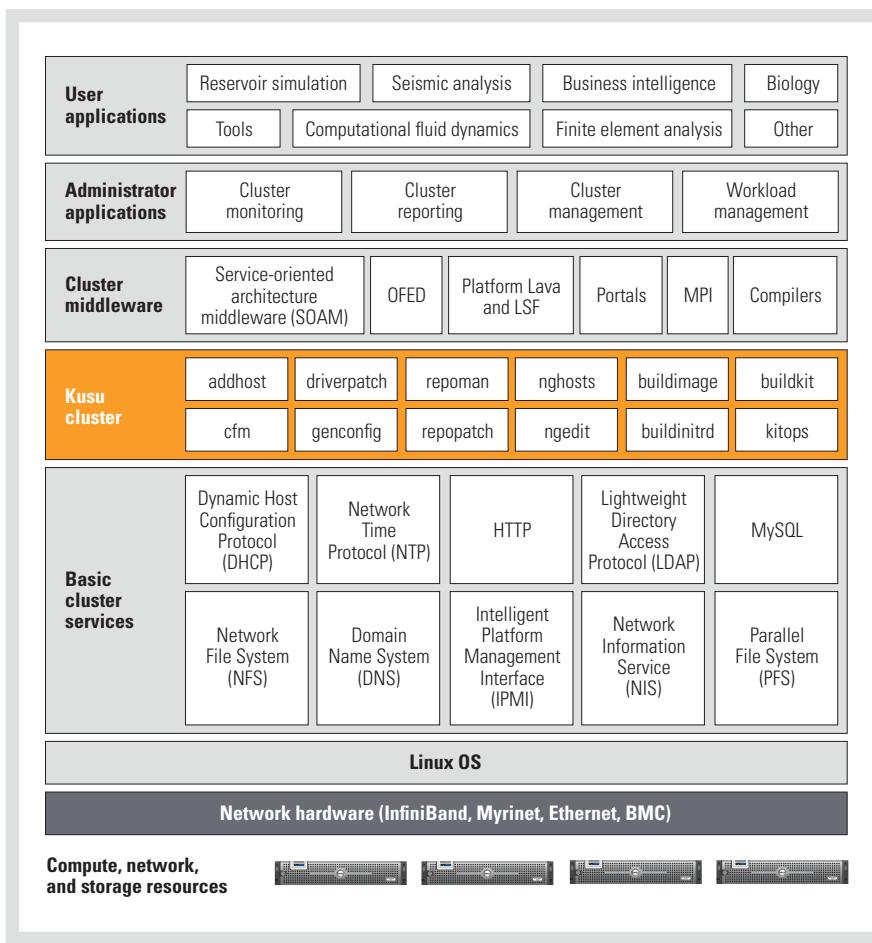


Figure 1. Platform OCS – Dell Edition framework

The Kusu framework is built on several basic elements:

- **Installer node:** HPC cluster setup primarily requires a master node known as an installer node. Installer nodes maintain the kusudb MySQL database for all functionalities.
- **Node group:** The node group is the template that defines the various properties of the compute nodes so that one set of compute nodes can be distinguished from the others based on the node grouping. Kusu provides different node groups such as installer, compute-rhel-5-x86_64, compute-imaged-rhel-5-x86_64, compute-diskless-rhel-5-x86_64, and unmanaged node groups. From these node groups, administrators can make copies and create customized node groups—for

example, one group for InfiniBand and another for Gigabit Ethernet. This approach enables administrators to distinguish sets of nodes to take

specific actions on, including selecting components to install; setting network configurations; changing the disk partitioning, OS, and kernel parameters; or adding custom scripts.

- **Kits:** To help ease the installation of different components, the Kusu framework packages different software components into kits. A kit may contain Red Hat Package Manager (RPM™) packages, scripts, applications, or services that must be installed on compute nodes, and includes meta-RPM components with configuration scripts and dependency information. Administrators can install kits from an ISO image or a mount point. Figure 2 shows an example of a kit framework.

Two advantages of using this mechanism are that the MySQL database stores information related to the different kits, and that it eliminates the need to configure applications and services manually. Installing the kit components automatically generates and updates the configurations the service or application needs, which can help save time.

A variety of kits are available by default with Platform OCS – Dell Edition; administrators can add these kits during installation of the installer node or after installation. Available kits include the following:

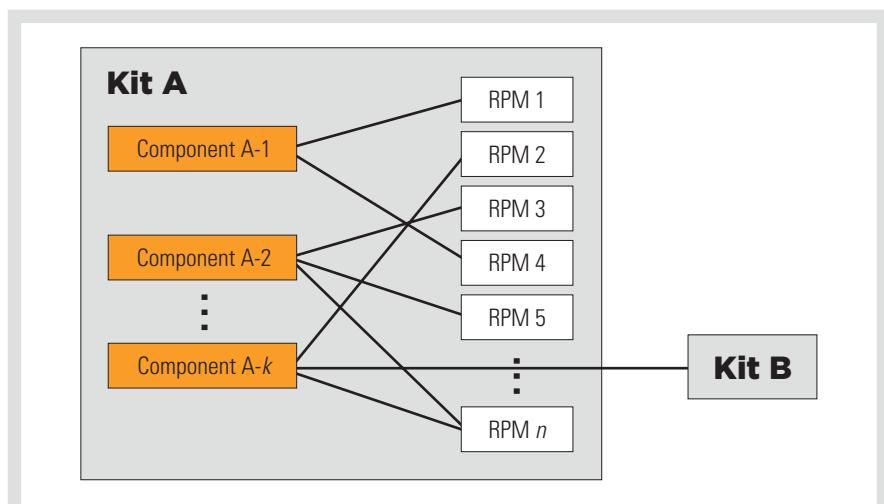


Figure 2. Example Kusu kit framework

- **Base Kit:** Contains the basic components needed for the Kusu framework, as well as the tools for managing the kit and node groups and installing the nodes
- **OS Kit:** Contains the Red Hat Enterprise Linux 5 Update 1 ISO image on DVD
- **OFED Kit:** Includes the OFED stack and drivers for InfiniBand configuration, as well as the opensmbd subnet manager
- **HPC Kit:** Contains math libraries such as Automatically Tuned Linear Algebra Software (ATLAS), Basic Linear Algebra Communication Subprograms (BLACS), and Scalable Linear Algebra Package (ScalAPACK), as well as MPI distributions such as MVAPICH 1 and 2, Open MPI, and MPICH1 and 2 and precompiled benchmarks such as High-Performance Linpack (HPL) and IOzone
- **Platform Lava Kit:** Provides scheduling for workload management on the compute nodes (based on an open source version of Platform LSF)

The Base Kit and OS Kit are mandatory in Platform OCS - Dell Edition clusters. For more information on management and monitoring kits, see the “Managing and monitoring clusters” section in this article.

During installation, Platform OCS - Dell Edition copies the OS media into a kit it uses to make a repository—a container that can store individual software components, RPM packages, custom scripts, and even a complete OS. To help resolve conflicts between kits and help provide fail-safe upgrades, administrators can create custom repositories to distinguish one repository from another. The default repository in Kusu is called “Repo for rhel-5-x86_64”; the default repository path is /depot/repos/1000. Administrators can take snapshots of a repository without affecting the original repository.

During installation of the installer node, associations are created between the kits and node groups—meaning that many applications (provided by kits) require no manual configuration. If

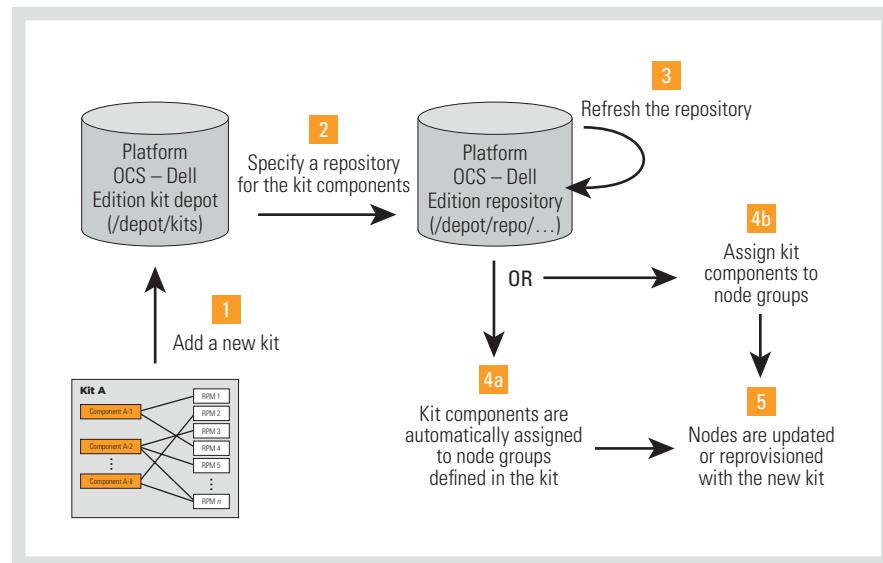


Figure 3. Sequence for creating custom Kusu kit, node group, and repository associations

administrators create custom repositories, kits, or node groups, they should do so sequentially to help ensure that the components are updated correctly in the MySQL database. Figure 3 shows the correct sequence:

1. Add a new kit A using the kitops --add A command.
2. Specify a repository for the kit components using the repoman command. Documentation and kit component configuration are stored on the installer node.
3. Refresh the repository after adding the new kit.
4. If the kit components are not automatically assigned to node groups defined in the kit, use the ngedit tool to assign the components to node groups.
5. Nodes in the node group are then updated or reprovisioned with the new kit. Adding a new kit typically does not require reprovisioning.

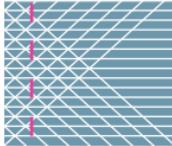
This framework architecture enables administrators to build heterogeneous software bundles for various node groups, and to take advantage of its flexibility in defining the behavior and characteristics of individual components in the cluster.

Installation and deployment

Installing cluster nodes using the Kusu framework is a network-based process; only the primary installer node requires administrators to select from various options and start the installation. Administrators can install the other compute nodes using a Preboot Execution Environment (PXE) mechanism, without needing to perform manual intervention.

Platform OCS - Dell Edition uses the Red Hat Anaconda installer for the compute nodes. A kickstart file tells Anaconda the desired RPM packages and the root password needed on compute nodes. During installation, the installer node shows different options to define the partitioning schema, which is completed outside Anaconda. For compute nodes, the ngedit node group editor enables administrators to modify the partitions needed for compute nodes.

The automatic installation mechanism in Platform OCS - Dell Edition enables administrators to easily deploy the components, node groups, and other behavior predefined in the installer node to the requisite compute nodes. By integrating the ability to define functionality and the installation mechanism within one framework, Platform OCS - Dell Edition helps significantly simplify cluster deployment and management.



Symantec Enterprise Messaging Management for Microsoft® Exchange

**A comprehensive approach to effectively
managing messaging environments**

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



symantec™

Available now...

The New Symantec Yellow Book™

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

Highlights of the new Yellow Book:

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

Retail value: \$65.00

Obtain a copy by electronic download.

E-mail: korbie_tomlinson@symantec.com

Subject of e-mail: Dell YB – Protecting Exchange 2007



symantec™

MANAGING AND MONITORING CLUSTERS

Managing a cluster effectively can be challenging—requirements vary depending on the complexity of cluster setup and operations performed throughout the cluster life cycle. The Dell OpenManage suite offers several options ranging from command-line interfaces to intuitive graphical user interfaces for both in-band and out-of-band systems management. In addition, the Kusu-based Platform OCS – Dell Edition stack makes a variety of utilities available to support complementary cluster software deployment, resource management, and monitoring of cluster-level performance, cluster job state, and resource utilization.¹

Cluster management tools

The Platform OCS – Dell Edition Base Kit contains a number of useful command-line tools and utilities to help simplify cluster management, including the following:

- **addhost:** A tool that allows administrators to add and remove nodes from a specified node group
- **boothost:** A tool to create the PXE configuration files used when starting up a node
- **cfm:** A configuration file manager that synchronizes node group packages and files
- **netedit:** A tool that defines available network interfaces in a cluster, which are then associated with node groups
- **ngedit:** The node group editor
- **nghosts:** A tool that moves hosts between node groups
- **repoman:** A tool that manages the repositories associated with different node groups in a cluster
- **repopatch:** A tool that updates repositories using the yum (Yellowdog Updater, Modified) tool and/or the Red Hat Network (RHN)

- **kitops:** A tool to manipulate kits in a cluster, including adding, removing, and listing kits
- **buildkit:** A tool that allows administrators and developers to build new kits by creating a template workspace for the kit
- **genconfig:** The database-reporting and file-generating tool
- **pdsh:** A high-performance, parallel remote-shell utility used to run commands across a cluster²

Platform OCS – Dell Edition is also bundled with Platform Lava, an open source job scheduler based on Platform LSF that provides a distributed batch system for submitting jobs and managing workloads. It helps manage the day-to-day workload of a whole cluster, providing simplified job execution, management, and accounting. This package is part of the optional Platform Lava Kit and is made up of two components: component-lava-master (which installs Platform Lava on the Platform OCS – Dell Edition primary installer node and sets up the Platform Lava master node with a default configuration) and component-lava-compute (which installs Platform Lava on Platform OCS – Dell Edition cluster compute nodes). Administrators can use Platform Lava to submit jobs within a cluster using the bsub command, view job status using the bjobs command, and view job output using the bpeek command.

Cluster monitoring tools

The Kusu-based Platform OCS – Dell Edition stack bundles several monitoring applications, including open source utilities:

- **Cacti:** A comprehensive network graphing solution that uses the data storage and graphing capabilities of RRDTool, Cacti provides a fast poller, advanced graphing templates, multiple data

acquisition methods, and various user management features. Cacti can display almost any kind of data graphically and is part of the optional Cacti Kit, which includes two main components: component-cacti (installed on a node that becomes the Cacti master monitoring node that runs the Cacti Web service and allows administrators to view collected Cacti data) and component-cacti-monitored-node (which instructs Cacti to include the package-based compute node group in the collected data, and to monitor only specified node groups). The Cacti monitoring node performs all Cacti configurations.

- **Nagios:** A network management server that runs on the primary installer node and monitors cluster hosts, services, and networks, Nagios provides a Web interface to display the collected node information and to generate alerts if problems occur within configured thresholds. Nagios monitors the various services (such as disk usage, users, and ping) on both the primary installation and compute nodes in a Platform OCS – Dell Edition cluster. This package is part of the optional Nagios Kit, which includes two main components: component-nagios-installer (which monitors hosts and services and reports exceeded thresholds and problems) and component-nagios-compute (which installs a monitoring agent on compute nodes that collects information and sends it back to the network management service). Nagios provides a Web-based console accessible from compute nodes through component-nagios-installer.
- **Ganglia:** A scalable statistics collector that monitors node availability and displays system load, network usage, and other resource information for HPC clusters, Ganglia displays detailed information about Platform OCS – Dell Edition clusters and their day-to-day functions,

¹For more information, see "Customizing Management of HPC Clusters," by Arun Rajan, Tong Liu, Yung-Chin Fang, Garima Kochhar, and Ron Pepper, in *Dell Power Solutions*, February 2007, DELL.COM/Downloads/Global/Power/ps1q07-20060427-Fang.pdf.

²For a complete list of commands and detailed information on usage and functionality, see the media kit available on the master node and the online documentation available at www.platform.com.

including uptime, available memory, number of processors, and processor usage. Administrators can drill down within the data to view information about each node in a cluster. This package is part of the optional Ganglia Kit, which includes two main components: component-ganglia-server (which acts as the server to which collected cluster information is returned) and component-ganglia-agent (which installs a small daemon that runs on compute nodes and returns collected data to the server). All Ganglia configurations are carried out on the server (by default, the primary installer node) through the Ganglia console.

- **Ntop:** A tool that monitors network bandwidth and analyzes traffic, ntop allows administrators to examine the network patterns of a Platform OCS - Dell Edition cluster and positions the network interfaces into a passive listening mode so administrators can watch traffic to and from each interface. The tool plots the data into a database, then displays this information in the Web console. This package is part of the optional Ntop Kit, which has only one component—component-ntop. Many ntop configurations, such as changing the default listening interface, adding users, restricting user access to pages, and stopping ntop, are made on the ntop Web console.

USING THE DELL KIT AND DELL OPENMANAGE TOOLS

The Platform OCS - Dell Edition media includes the Dell Kit (called dell-vendor), which provides drivers and custom scripts for Dell PowerEdge servers. This kit automates the installation and deployment of Dell drivers and utilities on installer nodes and compute nodes. It also contains scripts and RPM packages that administrators can use to install Dell drivers and Dynamic Kernel Module Support (DKMS) packages, deploy Dell

OpenManage software, and configure the baseboard management controllers (BMCs) and BIOSs of cluster nodes to support out-of-band systems management and console redirection.³

Understanding kit components

The Dell Kit includes four major components: component-Dell-compute, component-Dell-installer, component-Dell-pxeboot, and component-Dell-bmc-setup. Component-Dell-compute and component-Dell-installer are used for installing and configuring Dell drivers and DKMS packages for Dell PowerEdge servers. These components include e1000, igb, bnx2, tg3, mptlinux, and megaraid_sas. During installation, the components check the PowerEdge server model and install appropriate drivers and DKMS packages that are supported and required for that model. Component-Dell-installer supports an automatic installation process for installer nodes. Component-Dell-compute is used to deploy Dell drivers on compute nodes.

Platform OCS - Dell Edition requires that compute nodes have PXE ahead of the hard drive in the BIOS boot sequence, which can be a problem for servers that have PXE after the hard drive in this sequence. To help avoid this problem, component-Dell-pxeboot contains scripts to configure the PXE boot order (for the next boot only) of the compute node using ipmitool. To reinstall compute nodes without manual intervention (pressing F12 for PXE), administrators must pre-enable these components.

Component-Dell-bmcsetup installs Dell OpenManage and configures the BMCs and BIOSs of the compute nodes during deployment. For systems with limited Dell OpenManage support, such as PowerEdge SC servers, the ipmitool utility can only configure the BMCs. For systems that include comprehensive Dell OpenManage support and are supported as part of the Dell HPC program, the script

allows configuration of BMCs and BIOSs using the Dell OpenManage omconfig utility. After the system has been configured, it can employ features such as remote power cycling of servers and console redirection. The script must be enabled to execute on the compute nodes during deployment.

Enabling kit components

Administrators can add the Dell Kit by using the Kit window in the Kusu front-end installation wizard or by using the kitops command. After kit installation, the component-Dell-compute and component-Dell-installer components are automatically enabled and associated to the default node group (installer-rhel-5-x86_64) and default compute node group (installer-rhel-5-x86_64), respectively. If administrators deploy compute nodes directly without any further configuration, the Dell drivers and DKMS packages included in those components are automatically installed on the servers.

By default, component-Dell-pxeboot is disabled after front-end installation; administrators must perform additional steps to enable the component using the ngedit node group editor for a selected node group. After running ngedit and choosing a node group, administrators would then navigate to the Components window and select "dell-vendor," then select "component-Dell-bmcsetup," and then update the nodes in the group. Component-Dell-pxeboot is then activated on the compute nodes without reinstallation.

Administrators must enable component-Dell-bmcsetup manually, following the same steps as for component-Dell-pxeboot; component-Dell-bmcsetup contains the scripts administrators need to automate the installation and the configuration, but not the Dell OpenManage software itself, which allows Platform OCS - Dell Edition to be independent of Dell OpenManage version changes.

³For more information, see "Configuring the BMC and BIOS on Dell Platforms in HPC Cluster Environments," by Garima Kochhar, Rizwan Ali, and Arun Rajan, in *Dell Power Solutions*, November 2005, DELL.COM/Downloads/Global/Power/ps4q05-20050222-Kochhar.pdf.

Component-Dell-bmcsetup provides a tool that allows administrators to customize the behavior of BMCs and BIOSs before deployment, including defining the BMC IP addresses, enabling and disabling console redirection, and so on. If administrators customize the BMCs and BIOSs, they must then reinstall the existing compute nodes.

Integrating the predefined vendor- and cluster-specific parameters and bundles with the compute node installation process can be challenging. Administrators can use the components of the Dell Kit to help them meet their systems management requirements on Dell clusters.

MIGRATING FROM PLATFORM OCS 4 TO PLATFORM OCS – DELL EDITION

Platform OCS 4 was based on the Rocks® cluster toolkit developed by the Rocks Cluster Group.⁴ Platform OCS – Dell Edition is based on the open source Kusu cluster toolkit created by Platform Computing. A document is available from Platform Computing Support to help simplify the migration process from Platform OCS 4.x to Platform OCS – Dell Edition.

Administrators can perform this migration by backing up and then restoring important information in three categories: system information, custom RPM packages, and XML configurations for compute nodes. They can use migration tools to help them back up and restore information during this process. First, they must copy custom RPM packages, scripts, and XML files manually from the Platform OCS 4.x front-end node to the new installer nodes. The optional packages and custom script windows of the ngedit tool in Platform OCS – Dell Edition allow administrators to specify custom RPM packages and scripts deployed on compute nodes. Administrators must also back up and restore all other important data on the existing front-end node.

After all information is ready on the new Platform OCS – Dell Edition installer

node, administrators can start the compute node deployment process. To start the deployment, they can simply reboot (to PXE boot) all Platform OCS 4.x compute nodes.

SIMPLIFYING HPC CLUSTER DEPLOYMENT AND MANAGEMENT

The Kusu-based Platform OCS – Dell Edition software stack provides a comprehensive set of cluster features and tools that help simplify the deployment and maintenance of HPC clusters. Built on a hierarchical open source architecture, Platform OCS – Dell Edition provides a unified framework that supports the seamless integration of third-party products, including both open source utilities and proprietary software, as well as pre-bundled tools to help administrators monitor, manage, and administer clusters. By taking advantage of tools such as the Dell Kit and Dell OpenManage software, administrators can help simplify the deployment and management of Dell PowerEdge server-based HPC clusters. 

Sreeram Vedantham is a development engineer adviser in the Dell High-Performance Computing Group. His current interests include high-speed interconnects, HPC cluster stacks, and schedulers.

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

Shivaraj Nidoni is a senior engineering analyst in the Dell High-Performance Computing Group. His current interests include high-end interconnects, parallel

file systems, and HPC clustering packages. Shivaraj has an M.Tech. in Computer Network Engineering from Visvesvaraya Technological University.

Li Ou, Ph.D., is a systems engineer in the Scalable Systems Group at Dell. His current interests include parallel rendering, visualization, parallel file systems, high-performance storage systems, and HPC clustering. He has a B.S. in Electrical Engineering and an M.S. in Computer Science, both from the University of Electronics Science and Technology of China, and a Ph.D. in Computer Engineering from Tennessee Technological University.

Gabor Samu is a product manager at Platform Computing. His current interests include cluster management and multi-core programming. He has a B.S. in Computer Science from the University of Toronto.

Mark Black is development manager at Platform Computing. His current interests include cluster management, cluster provisioning, and virtualization. He has a B.E. from Ryerson University and an M.E.S. in Signal Processing from the University of Western Ontario.



QUICK LINKS

Platform Open Cluster Stack:
www.platform.com/products/platform-open-cluster-stack5

Project Kusu:
www.hpccommunity.org

Dell HPC cluster solutions:
DELL.COM/HPCC

⁴This 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.



By Brock A. Taylor
Vishvesh Sahasrabudhe
Onur Celebioglu

ACCELERATED HPC PRODUCTIVITY WITH INTEL CLUSTER READY SOLUTIONS

The Intel® Cluster Ready program provides a standardized, replicable way to build and run high-performance computing (HPC) clusters, helping simplify cluster deployment and management. By using Intel Cluster Ready-certified Dell™ HPC clusters, organizations can quickly install and configure clusters to begin running registered HPC applications.

High-performance computing (HPC) clusters can be extremely effective for solving large, complex problems in areas such as computational fluid dynamics, weather modeling, bioinformatics, and oil and gas exploration. However, they can also be extremely complex and time-consuming to deploy, manage, and maintain. To help reduce this complexity and the time required to manage cluster resources, Intel—with the help of Dell and other partners—launched the Intel Cluster Ready program. This program is designed to help organizations create compliant cluster software environments to run registered applications, and provide the tools that can help maintain systems at their expected levels of operations. It also aims to speed adoption of clusters by organizations otherwise limited to single workstations or symmetric multiprocessing systems because of the complexities of deploying and operating HPC clusters.¹

To help demonstrate how this program can help get verified clusters quickly up and running registered applications, a team of Dell and Intel engineers undertook a project to simulate an organization deploying and using an Intel Cluster Ready-certified Dell HPC cluster. The project team deployed a 16-node cluster based on Dell PowerEdge™ M600 blade servers, then ran eight registered applications to showcase the efficiency and effectiveness of Intel Cluster Ready solutions.

Using similar clusters in their own environments enables organizations to focus on determining which registered applications to run rather than on handling the complexities of cluster deployment—ultimately helping increase productivity and accelerate the return on their investment in computing resources.

UNDERSTANDING THE INTEL CLUSTER READY PROGRAM

The Intel Cluster Ready program is based on a specification that defines a common interface for HPC applications and incorporates industry-wide practices for building and managing HPC clusters. The defined application-layer interface enables application portability between certified clusters. Compliant clusters are designed to run any registered application, providing a one-to-many cluster-to-application environment. Requiring industry standards and practices helps certified clusters provide the tools and functionality necessary to deploy clusters efficiently and provide the administrative functionality that fosters simplified management. Each certified cluster must be built and verified to be compliant before being released, which shifts the intricacies of designing and building compliant clusters to the vendor, allowing organizations to focus on using the cluster for its intended purpose—running applications and solving problems.

Related Categories:

Dell PowerEdge blade servers
High-performance computing (HPC)
Intel
Linux
Visit DELL.COM/PowerSolutions
for the complete category index.

¹For more information on the Intel Cluster Ready program, visit DELL.COM/Content/Topics/Global.aspx/Sitelets/Solutions/Cluster_Grid/En/US/Clustering_HPCC?c=us&cs=555&l=en&s=biz§ion=008.

As of July 2008, the Intel Cluster Ready program defines a specification for Linux® OS-based clusters; Figure 1 shows the Intel Cluster Ready architecture. Although other operating systems may be considered in the future, the July 2008 version is a specific Linux binding. Example software interface requirements for each cluster node include the following:

- A Linux kernel that conforms to the Linux Standard Base
- A Portable Operating System Interface for UNIX® (POSIX) command system
- Intel 64 binaries for the base libraries, GNU C library (glibc), and X11 libraries
- An implementation of Message Passing Interface (MPI) 1.2 or later
- A set of common runtime components, including Intel compilers and libraries
- The OpenFabrics Enterprise Distribution (OFED) software stack as the network interface for network fabrics beyond Ethernet

Example hardware characteristics and capabilities include the following:

- At least 512 MB of memory per core
- At least 10 GB of storage space available to each compute node

- Identical compute node hardware (only nominal differences allowed)
- Identical file trees across compute nodes
- Node access to a single shared file space for user home directories
- Cluster management network fabric

The Intel Cluster Ready program uses a certification process for HPC clusters and a registration process for applications to enable its one-to-many ecosystem. A cluster recipe is a process or mechanism to build a solution that is compliant with the specification. A reference system built from a recipe is certified to meet the requirements and can provide the software interface to the application layer. Separately, HPC application vendors register that their codes operate on top of the defined minimum interface provided by a compliant cluster. Applications with components not required by the specification must be identified on the Intel Cluster Ready Web site to be a registered code, or those additional components must be bundled with the application itself. Combining these two processes helps provide a one-to-many ecosystem and avoid potential problems when a cluster is built specifically around the requirements of a single application. Compliant clusters are

designed to run a single registered application or multiple registered applications without modifications or rebuilding the cluster software stack.

Intel Cluster Checker—a preconfigured tool included in compliant clusters and designed to check system health and performance—is a key component of the program. This tool systematically runs through test modules, checking basic system functionality such as Secure Shell (SSH) capabilities between nodes as well as high-level performance across the cluster—including exhaustive link checking for expected bandwidth and latency performance between each node pair in the cluster. It is designed to support efficient cluster management, reduce the need for downtime for maintenance, and provide quick isolation of problems. Its output provides definitive pass or fail reporting of cluster health, helping confirm that the cluster is operating as expected and helping reduce maintenance time.

DEMONSTRATING CLUSTER SETUP AND USE

In April 2008, a team of Dell and Intel engineers launched a project to showcase the rapid deployment of an Intel Cluster Ready-certified Dell HPC cluster and the

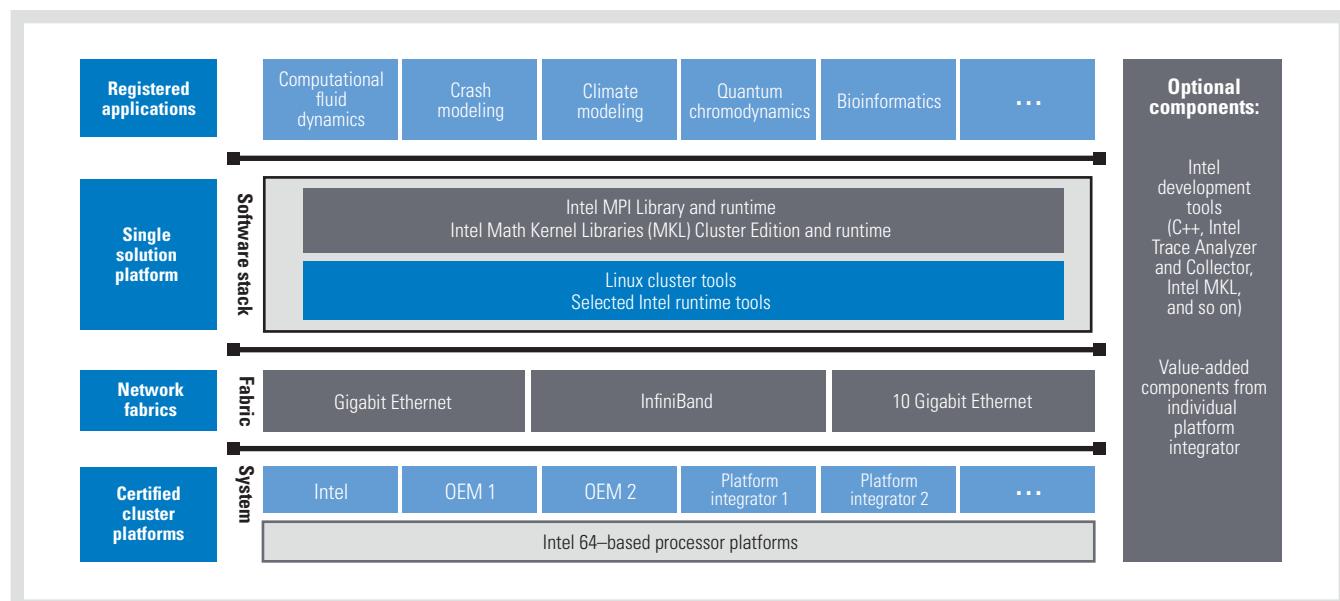


Figure 1. Intel Cluster Ready architecture

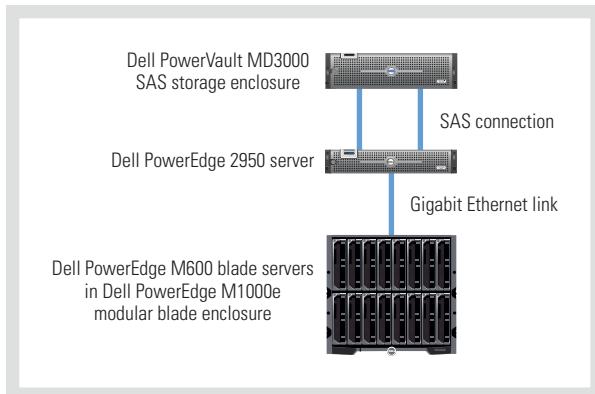


Figure 2. Cluster configuration in the example environment

quick, useful results this resource could provide. Dell has invested significant effort and expertise in creating and certifying cluster configurations designed to provide consistently compliant systems. Installation and configuration steps are documented, and automation helps avoid manual steps that may introduce problems in the cluster build. The components are integrated and then verified to help ensure the cluster presents the necessary interface to the applications.

The team deployed a certified cluster of 16 Dell PowerEdge M600 Intel 64-based blade servers in a PowerEdge M1000e modular blade enclosure as compute nodes, with a PowerEdge 2950 server as the master node (see Figure 2). Each PowerEdge M600 contained two quad-core Intel Xeon® E5450 processors at 3.00 GHz and 16 GB of RAM; the PowerEdge 2950 was equipped with two quad-core Intel Xeon E5450 processors at 3.00 GHz and 8 GB of RAM, and was attached to a Dell PowerVault™ MD3000 Serial Attached SCSI (SAS) storage enclosure. The cluster included two network fabrics: Gigabit Ethernet between all nodes and InfiniBand between compute nodes. The InfiniBand fabric included a Cisco M SFS7000E InfiniBand switch module for the enclosure with 16 internal and 8 external 4X, 20 Gbps, double data rate (DDR) InfiniBand ports. Because the InfiniBand fabric was contained in the PowerEdge M1000e enclosure for this cluster, no external InfiniBand cable was needed.

The cluster used Platform Open Cluster Stack 4.5.1 to provision the software stack across the cluster.² Each node ran the Red Hat® Enterprise Linux 4 Update 5 OS and required components such as Intel runtime libraries and the Java Runtime Environment (JRE). The software stack also included OFED 1.3 for the InfiniBand fabric. The

Intel Cluster Checker tool was installed on the master node along with preconfigured input files specifically designed to check the health of the cluster.

The deployment simulation was carried out in two stages. For the first stage, the project team mimicked the creation of an Intel Cluster Ready-certified Dell HPC cluster by performing actions carried out at the engineering level at Dell. By following the cluster recipe, the team moved the cluster hardware from bare metal to a fully functional system in approximately two hours. The Intel Cluster Checker tool verified that the software stack was a copy of the original certified reference cluster and that the cluster was operating as expected. All nodes were functional, the Gigabit Ethernet and InfiniBand fabrics were operational and performing at the expected levels, and the cluster was ready for use.

After the Intel Cluster Checker tool had verified the cluster, the team began the second stage—simulating an organization using the cluster for real work using registered applications. Eight HPC applications from six independent software vendors (ISVs) were installed on the cluster. One by one, the team launched real workload analysis on ANSYS, FLUENT, STAR-CD, LS-DYNA, Abaqus/Standard, Abaqus/Explicit, NASA Structural Analysis System (NASTRAN), and PAM-CRASH. Each application ran without modifications to the cluster, and each ran to

completion. Apart from the effort to bring the cluster to an Intel Cluster Ready-certified stage, no additional time was necessary to customize the cluster configuration to fit the needs of individual applications. This study demonstrates that the Intel Cluster Ready program combined with certified Dell cluster configurations can provide a turnkey HPC platform that end users can take advantage of rapidly and efficiently.

SIMPLIFYING CLUSTER DEPLOYMENT AND MANAGEMENT

The Intel Cluster Ready program is designed to take advantage of the expertise of independent hardware and software vendors to provide a turnkey solution for organizations deploying HPC clusters. Dell and Intel are working closely with ISVs to provide cluster resources that meet organizations' needs without requiring in-depth knowledge of how to make those systems work—helping simplify the installation, configuration, and management of clusters and enabling organizations to focus on their productivity. 

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

Vishvesh Sahasrabudhe is a member of the HPC Engineering Group at Dell.

Onur Celebioglu is an engineering manager in the HPC Engineering Group at Dell.

MORE


ONLINE

DELL.COM/PowerSolutions

QUICK LINKS

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

Dell HPC cluster solutions: DELL.COM/HPCC

²This 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.



IT Utopia.

Hardware and Software in Harmony.
That's IT as it should be.

Simplify clustering.

The Intel® Cluster Ready program simplifies buying, deploying, and managing an HPC cluster. It helps you ensure the interoperability of HPC components and applications, so you can deploy and operate a cluster faster, more easily, and with greater confidence.

Experience the productivity and problem-solving power of high-performance computing (HPC) based on Intel® processors. Make your HPC solutions Intel Cluster Ready today.

Get certified solutions and registered applications at www.intel.com/go/cluster.





By Wayne Slater
Gord Sissons

THE RED HAT HPC SOLUTION: SIMPLIFIED HIGH-PERFORMANCE LINUX CLUSTERS

High-performance computing (HPC) cluster deployments can be complicated, time-consuming, and costly. The Red Hat® HPC Solution enables organizations to deploy enhanced systems quickly and run them efficiently, helping reduce the cost and complexity of their Linux® cluster deployments.

High-performance computing (HPC) clusters on x86 servers are a cost-effective alternative to supercomputers and other architectures running proprietary operating systems. In fact, Linux-based x86 servers such as Dell™ PowerEdge™ 1950 III, PowerEdge 2950 III, PowerEdge 2970, PowerEdge SC1435, PowerEdge M600, and PowerEdge M605 servers have come to dominate the TOP500 Supercomputing Sites list (www.top500.org). To make the power of HPC clusters available to organizations of all sizes, Red Hat and Platform Computing have collaborated closely to create the Red Hat HPC Solution.

The root of the Red Hat HPC Solution is Project Kusu, an open source development effort sponsored by Platform Computing. Project Kusu is designed as a source kit for simplified cluster management and deployment, and supports a range of different Linux distributions, including Red Hat Enterprise Linux.

Project Kusu calls for a standards-based solution combining open source and commercial software into a single operating environment. The Red Hat HPC Solution achieves that objective by integrating Red Hat Enterprise Linux with Platform™ Open Cluster Stack (OCS) 5. The Red Hat HPC Solution combines these components into a single cluster operating environment designed to provide IT administrators with what they need to deploy, run, and manage an HPC cluster (see Figure 1).

Related Categories:

High-performance computing (HPC)
Intel
Linux
Platform Computing
Red Hat Enterprise Linux
Visit DELL.COM/PowerSolutions
for the complete category index.

PROVIDING A COMPREHENSIVE HPC CLUSTER SOLUTION

The Red Hat HPC Solution features an Intel® Cluster Ready-certified software suite. The Intel Cluster Ready specification includes tools and functionality to help simplify cluster deployment and management. Intel Cluster Ready-certified clusters are also designed to run all registered applications, providing a one-to-many cluster-to-application environment. The Red Hat HPC Solution provides application portals tailored to a variety of widely used HPC applications, many of them Intel Cluster Ready registered.

Along with cutting-edge cluster management tools from the open source community, the Red Hat HPC Solution includes Platform Lava—an open source version of the Platform Load Sharing Facility (LSF®) application—for integrated, open source workload management and accounting; the Platform Management Console (PMC), a Web browser-based graphical user interface (GUI); automated software maintenance; and a variety of HPC tools, libraries, and developer tools.

SIMPLIFYING CLUSTER MANAGEMENT

The PMC was introduced in Platform OCS 5. An intuitive GUI available as a complimentary download option and designed to administer all components of the cluster environment, it supports cluster monitoring and reporting by node, service, workload, and network. Availability and inventory reports are standard;

interfaces are provided for commercial solutions such as the Platform LSF job scheduling solution. Platform LSF is designed to intelligently schedule parallel and serial workloads to help make optimum use of available computing resources.

Job submission templates for common HPC applications are also included in the PMC. Users can customize them with optional Platform components such as the EnginFrame grid computing portal. By logging in to the portal, users can access and control their computing resources from almost anywhere using the Internet or an intranet.

QUICKLY PROVISIONING AND DEPLOYING CLUSTERS

Using the Red Hat HPC Solution enables administrators to provision nodes quickly and deploy them immediately. The solution supports diskless nodes, image-based nodes, customizable package-based provisioning, IP address assignment, and node-naming conventions of hosts or groups. The Red Hat HPC Solution provides comprehensive node configuration templates—unlike typical provisioning solutions that offer partial templates. The Red Hat HPC Solution is also distinguished by its support for multiple operating systems and versions.

The Red Hat HPC Solution offers a variety of kits to help simplify software installation and maintenance. Administrators can automate upgrades and patches using standard Linux tools such as yum (Yellowdog Updater, Modified) and the Red Hat Network (RHN) service. RHN makes updates and patches for packages included within Red Hat Enterprise Linux available to subscribers. Administrators can then use the yum program to download and install updates from RHN.

AUTOMATICALLY INSTALLING RESOURCE MONITORING TOOLS

The Red Hat HPC Solution enables administrators to automatically install and configure management tools such as the Cacti node and cluster monitor and

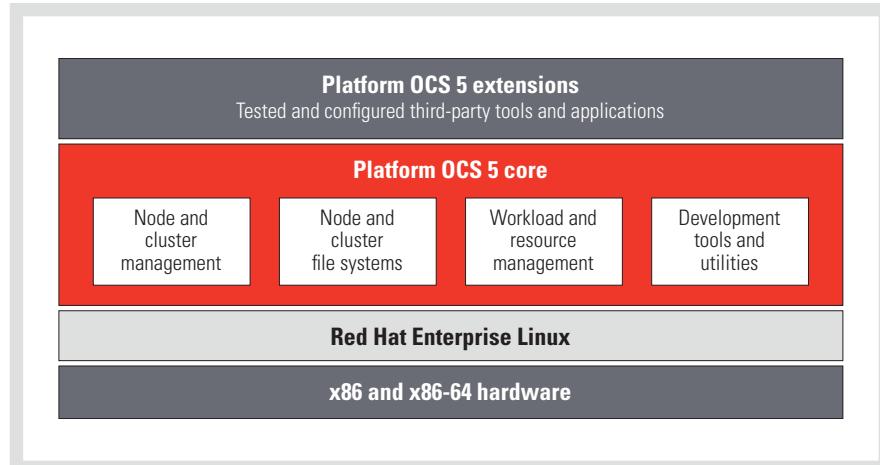


Figure 1. Components in the Red Hat HPC Solution combine open source and commercial software into a single operating environment

the Nagios host and service monitor. The tools can be configured to track resources and alert administrators if user-defined thresholds are exceeded. Nagios is fully integrated. These monitoring and reporting tools, combined with the Web browser-based workload management capabilities in Platform Lava, help simplify resource and performance optimization.

Platform Analytics and Platform Real Time Monitoring (RTM) also integrate seamlessly with the Red Hat HPC Solution. Platform Analytics offers tools for collecting, analyzing, and visualizing information for decision making based on usage patterns and loads. Platform RTM provides a dashboard to monitor physical devices, application software functions, and individual job statistics for workload management.

SIMPLIFYING LIFE CYCLE MANAGEMENT

As clusters grow, they become increasingly complex. Change management and cluster capacity can become significant challenges. The Red Hat HPC Solution helps meet these challenges while helping reduce costs. For example, changes to cluster nodes can be propagated without reprovisioning systems—the necessary tools are already included. The Red Hat HPC Solution also helps simplify capacity planning and expansion. Administrators

can easily visualize trends with PMC graphing tools, document usage with Platform Lava or Platform LSF, and add nodes or groups to the cluster without disrupting operations.

The Red Hat HPC Solution enables analysts, engineers, and scientists to employ the power of HPC with an out-of-the-box, pre-integrated, vendor-certified software solution. Users and administrators alike can enhance productivity with a simplified HPC cluster that can be up and running quickly, and can be optimized for enhanced performance throughout the cluster life cycle. 

Wayne Slater is manager of partner marketing for Platform Computing.

Gord Sissons runs a small software company and technology consulting practice located near Toronto.



QUICK LINK

Red Hat HPC Solution:
www.redhat.com/hpc



By Jacob Liberman
Mausmi Kotecha
Vishvesh Sahasrabudhe

EVALUATING THIRD-GENERATION AMD OPTERON PROCESSORS FOR HPC WORKLOADS

Third-generation AMD Opteron™ processors are designed to optimize multi-threaded application performance, and include multiple architectural enhancements over second-generation AMD Opteron processors. To help organizations understand the performance increases possible when upgrading to third-generation processors, Dell engineers performed benchmark tests against a variety of high-performance computing (HPC) workloads.

AMD Opteron processors have enjoyed considerable success in high-performance computing (HPC) during the last several years. According to the TOP500 list—one indicator of HPC trends—AMD Opteron processors at one point powered 23 percent of the world's fastest supercomputers.¹ An innovative dual-core design enhanced the rapid adoption of these processors: by integrating the memory controller directly on the chip, second-generation AMD Opteron processors bypassed the memory contention inherent to systems using multi-core processors based on legacy technology. Additionally, high-speed bidirectional HyperTransport™ links provided a scalable bandwidth interconnect between the computing cores and the I/O subsystem. These enhancements were collectively known as the Direct Connect Architecture. The balanced processing capability and memory bandwidth of second-generation AMD Opteron processors translated into increased efficiency and scalability for HPC applications.

Third-generation AMD Opteron processors include key architectural enhancements designed to boost

performance and energy efficiency over the second-generation processors without sacrificing memory bandwidth. These enhancements include the following:

- An increased number of processor cores per socket (four instead of two) designed to operate without an increase in power envelope
- An innovative cache structure that includes a 512 KB dedicated level 2 (L2) cache per core and a 2 MB shared L3 cache
- 128-bit Streaming SIMD (single instruction, multiple data) Extensions (SSE) execution width that allows four floating-point operations per clock cycle
- 128-bit integrated memory controller divided into two independent 64-bit channels
- Reduced power consumption through multiple-core voltage and frequency scaling²

Figure 1 illustrates the architectural differences between second- and third-generation AMD Opteron processors. Ninth- and tenth-generation Dell™ PowerEdge™ servers support both second- and

¹As listed in the November 2006 TOP500 list. For more information, visit www.top500.org.

²For a detailed description of third-generation AMD Opteron processors, visit www.amd.com/us-en/Processors/ProductInformation/0_30_118_8796_15223,00.html.

third-generation AMD Opteron processors, which are both socket-F compatible, helping simplify upgrades for HPC systems.

To demonstrate the performance increases made possible by upgrading from second-generation to third-generation AMD Opteron processors, in July 2008 Dell engineers compared the performance of each type of processor in an HPC test environment. The results indicate that third-generation AMD Opteron processors can provide significant performance increases across a broad spectrum of HPC workloads.

TEST ENVIRONMENT

The test team carried out two rounds of testing. In the first round, they compared the performance of second- and third-generation AMD Opteron processors using a single Dell PowerEdge M605 blade server, an energy-efficient two-socket server that supports up to 64 GB of double data rate 2 (DDR2) RAM. In the second round, the test team compared processor performance across a 16-node cluster of PowerEdge M605 blade servers. Figure 2 shows the hardware and software used in the test environment for each round of testing.

SINGLE-SERVER TEST RESULTS

The single-server tests compared the performance of second- and third-generation AMD Opteron processors in a single Dell PowerEdge M605 blade server using the following synthetic and application benchmarks:

- **Double Precision General Matrix Multiply (DGEMM):** A Basic Linear Algebra Subprograms (BLAS) subroutine from the AMD Core Math Library (ACML) that performs matrix multiplication in double precision
- **Basic Local Alignment Search Tool (BLAST):** A tool that provides algorithms for rapidly searching nucleotide and protein databases

- **STREAM:** A synthetic benchmark that measures sustainable memory bandwidth across four threaded operations: COPY, SCALE, ADD, and TRIAD
- **ANSYS Benchmarks:** A popular structural dynamics application that evaluates performance across codes with varying data access patterns and problem sizes using benchmark data sets³

The test team ran each benchmark three times using the maximum number of threads or processes for each processor—four for the second-generation dual-core

processors and eight for the third-generation quad-core processors. The test team then took the mean score for each processor type and normalized the results to show performance relative to the second-generation processors.

Figure 3 shows the results. As this figure indicates, in most cases the third-generation AMD Opteron processors provided significantly increased performance over the second-generation processors. The DGEMM benchmark resulted in the greatest performance increase, exceeding the baseline performance by

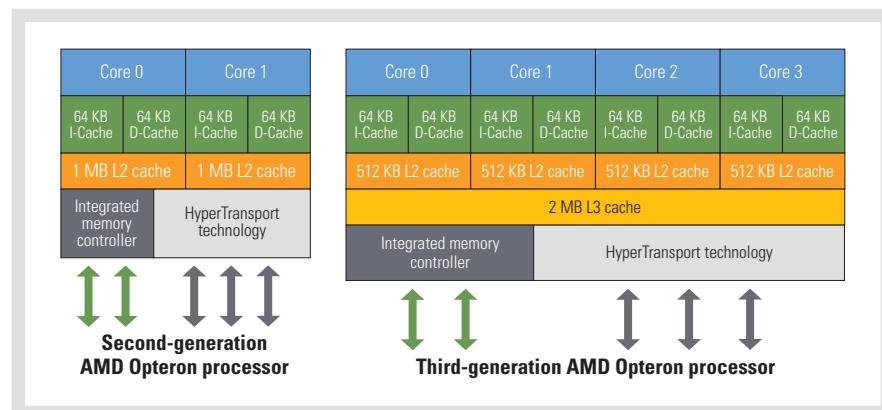


Figure 1. Second- and third-generation AMD Opteron processor architectures

	Single-server tests	Cluster tests
Servers	1 Dell PowerEdge M605 blade server	16 Dell PowerEdge M605 blade servers
Processors	▪ Two second-generation dual-core AMD Opteron 2218 processors at 2.6 GHz ▪ Two third-generation quad-core AMD Opteron 2356 processors at 2.3 GHz	
Memory	Eight 2 GB DDR2 dual in-line memory modules (DIMMs) at 667 MHz	
Disk	One 36 GB, 15,000 rpm Serial Attached SCSI (SAS) drive	
OS	Red Hat® Enterprise Linux® 5 Update 1 OS	
Compilers	GCC 4.2.1 and PGI 7.2.1	
Message Passing Interface (MPI)	N/A	Open MPI 1.2.5
Interconnect	N/A	Mellanox MT25408 ConnectX DDR InfiniBand with the OpenFabrics Enterprise Distribution (OFED) 1.3 software stack

Figure 2. Single-server and cluster benchmark configurations

³For more information on these benchmarks, visit developer.amd.com/cpu/libraries/acml/pages/default.aspx, www.ncbi.nlm.nih.gov/Education/BLASTinfo/information3.html, www.cs.virginia.edu/stream, and www.ansys.com/special/news-images/high_per_computing.htm.

nearly 250 percent. This increase over the second-generation processors was a result of the additional cores and wider SSE register, which allowed each core to complete four floating-point operations per clock cycle. Memory bandwidth, as measured by the STREAM benchmark, doubled in the case of the SCALE operation and increased by narrower margins across the COPY, ADD, and TRIAD operations. Taken together, the results indicate that the third-generation processors' expanded floating-point capability did not come at the expense of the memory subsystem.

The third-generation processors delivered substantial performance increases while running the synthetic DGEMM and STREAM benchmarks; however, because these benchmarks represent idealized workloads, they often fail to accurately predict application performance. The Dell team therefore performed additional tests using authentic applications. In these tests, changes in application performance across processor generations varied according to the application characteristics. Four of the seven ANSYS benchmarks showed performance increases in the range of 30-40 percent. These results were

favorable, because the ANSYS benchmark suites do not typically scale well. The ANSYS benchmarks with larger problem sizes—including bmd-4 and bmd-5—did not fare as well. The bmd-5 benchmark showed some performance increase; bmd-4 showed slight performance degradation. And although bmd-5 can make effective use of the processor cache, which generally results in good parallel performance, a larger problem size does not decompose into cacheable elements, which in bmd-4 increases the amount of I/O relative to computation. This increased I/O had the overall effect of minimizing performance gains related to the processing subsystem. BLAST, which also includes a substantial I/O component, showed similar results.

CLUSTER TEST RESULTS

Although the performance of the third-generation AMD Opteron processors exceeded that of the second-generation processors across most of the single-node benchmarks, single-server performance enhancements do not always translate to the cluster level. Certain applications may not scale well beyond a single node because of their lack of inherent parallelism or overhead

introduced by communication. To evaluate performance at the cluster level, the Dell team used the following synthetic and application benchmarks:

- **FLUENT:** A popular commercial computational fluid dynamics application for simulating fluid flow and heat transfer; the Dell team used three FLUENT benchmark data sets of increasing size
- **High-Performance Linpack (HPL):** A highly configurable and floating-point-intensive benchmark used to rank the world's fastest computers; HPL solves a dense linear system in double precision and represents an extreme HPC workload because it involves very little communication and I/O
- **NASA Advanced Supercomputing (NAS) Parallel Benchmarks (NPB):** An application-centric set of benchmarks that tests the performance of multiprocessor computers; the Dell team ran the Class C Embarrassingly Parallel (EP), Fast Fourier Transform (FT), Integer Sort (IS), Lower-Upper Symmetric Gauss-Seidel (LU), and Multigrid (MG) applications from the NPB suite
- **Weather Research and Forecasting (WRF):** A numerical weather prediction system⁴

Figure 4 compares the measured performance of the second- and third-generation AMD Opteron processors running benchmarks across the cluster, with the average of the three benchmark results again normalized relative to the performance of the second-generation processors. The FLUENT results in Figure 4 show a positive correlation between relative performance gains for third-generation AMD Opteron processors and increasing problem size. The larger benchmark data sets—truck and truck_poly, at 13 million cells

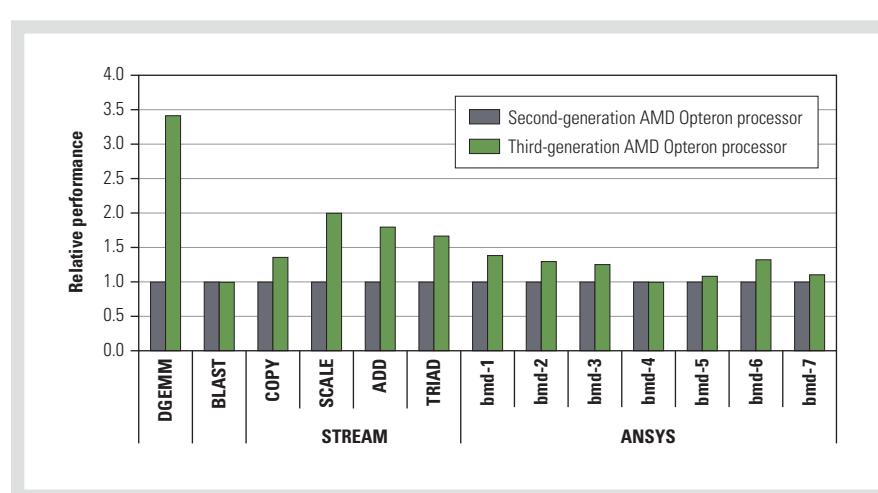


Figure 3. Relative single-server performance results for each AMD Opteron processor

⁴For more information on these benchmarks, visit www.fluent.com/software/fluent/fl6bench, www.netlib.org/benchmark/hpl, www.nas.nasa.gov/Resources/Software/npb.html, and www.mmm.ucar.edu/wrf/users/download/get_source.html.

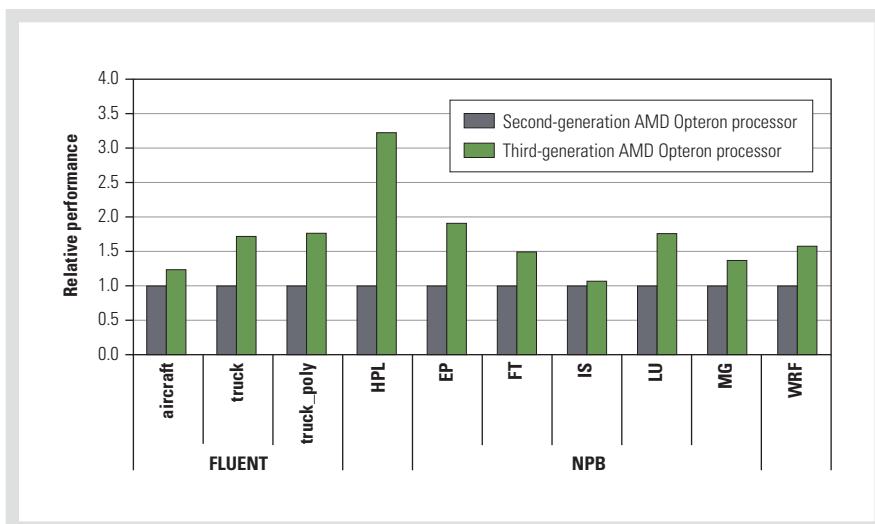


Figure 4. Relative cluster performance results for each AMD Opteron processor

each—showed performance increases of 71 percent and 74 percent, respectively. The smaller data set—aircraft, at 1.8 million cells—showed a performance increase of only 22 percent. These results indicate that the cluster was able to effectively utilize the additional cores of the third-generation processors, providing increased performance as the size of the data set grew.

HPL showed the largest performance increase of the benchmarks used. With a large problem size and a data distribution scheme that can make effective use of the processor cache, HPL is optimized to exploit both the additional cores and the increased floating-point operations per clock cycle, which is reflected in a more than threefold increase in HPL performance. The results also suggest that the DDR InfiniBand network was able to handle the increased traffic to provide near-optimal scaling.

The five NPB results show performance increases across a series of tests with wide variation in inter-node communication, latency, and floating-point requirements. The EP benchmark showed the largest increase because of a lack of communication between nodes combined with the increased number of cores. The IS benchmark, which involves intensive global communication with mixed

message sizes, showed the smallest increase. This benchmark is communication bound, providing a counterexample to the EP results, and does not take advantage of the additional SSE width because integer sorting does not involve floating-point calculations. The remaining NPB benchmarks, which feature balanced communication and floating-point calculation, showed increases in the range of 40-70 percent.

WRF showed a 63 percent performance increase with a constant problem size. Although both HPL and WRF are floating-point intensive, WRF fell short of the increase when using HPL. This disparity results from the larger memory footprint and I/O overhead of WRF compared with HPL.

SCALABLE PERFORMANCE FOR HPC WORKLOADS

Dell clusters equipped with third-generation AMD Opteron processors can offer substantial performance increases over those equipped with previous-generation processors. In the Dell tests, the third-generation processors demonstrated increased performance across a broad spectrum of HPC workloads with varying I/O and communication patterns, with memory-bound and floating-point-intensive application performance improving by more than 80 percent in

most cases. Because 9th- and 10th-generation Dell PowerEdge servers support both processor generations within the same power envelope, organizations have a clear upgrade path that can provide tangible performance gains in their environments. 

Jacob Liberman is a development engineer in the Dell Scalable Systems Group. He has a master's degree in Instructional Technology from the University of Texas at Austin and holds multiple industry certifications.

Mausmi Kotecha is an engineering analyst in the Enterprise Solutions Group at the Dell Bangalore Development Center. Her areas of interest include HPC clustering packages, performance analysis of parallel applications, and cluster file systems. She has a bachelor's degree in Computer Science from Atmiya Institute of Technology and Science.

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



QUICK LINKS

Dell HPC cluster solutions:
DELL.COM/HPC

Dell PowerEdge blade servers:
DELL.COM/Blades

AMD Opteron processors:
www.amd.com/opteron



EVALUATING MELLANOX CONNECTX INFINIBAND PERFORMANCE IN HPC CLUSTERS

By Munira Hussain

Toby Sebastian

Shivaraj Nidoni

Gilad Shainer

Next-generation Mellanox ConnectX double data rate (DDR) InfiniBand technology provides a low-latency, high-bandwidth interconnect that can help increase scalability and performance in high-performance computing (HPC) cluster environments. This article outlines the basics of the ConnectX architecture and presents test results comparing its latency and bandwidth with those of InfiniHost III Lx and Gigabit Ethernet.

Increases in the number of processors and cores within system architectures have created a constant challenge to balance computational and communications data in environments using powerful multi-core processors. An efficient communications subsystem is critical to help sustain high application performance. InfiniBand—a low-latency, high-bandwidth interconnect widely used in high-performance computing (HPC) cluster environments for Interprocess Communication (IPC)—can provide a key component in these environments.

ConnectX is the fourth-generation InfiniBand architecture from Mellanox Technologies, following the third-generation InfiniHost III architecture, and includes multiple enhancements designed to increase performance in HPC clusters. To help demonstrate these enhancements, Dell engineers ran a variety of benchmarks on an HPC cluster of Dell™ PowerEdge™ servers using Mellanox ConnectX double data rate (DDR) InfiniBand, InfiniHost III Lx DDR InfiniBand, and Gigabit Ethernet interconnects. The results illustrate the advantages ConnectX can provide in reduced latency and increased bandwidth in HPC cluster environments.

UNDERSTANDING THE INFINIBAND ARCHITECTURE

The designs and specifications for the InfiniBand switched-fabric I/O architecture are created by the InfiniBand Trade Association. The nodes connect to the switch through a host channel adapter (HCA) that plugs into a PCI Express (PCIe) slot. HCAs have programmable direct-memory-access engines that process packets transferred to and from the host.¹ The packets are transferred through Remote Direct Memory Access (RDMA), which can help reduce processor overhead. Data transfers directly from host memory to sender memory without using host processor cycles.

Many software stacks and libraries can support InfiniBand hardware; one commonly used unified software stack is the OpenFabrics Enterprise Distribution (OFED), available from the OpenFabrics Alliance. Major InfiniBand vendors and developers, including research and national labs, have collaborated and contributed to the OFED stack. The most recent release, OFED 1.3, includes a comprehensive InfiniBand stack, including drivers, libraries, diagnostic tools and utilities, and message-passing libraries. OFED 1.3 also includes

Related Categories:

High-performance computing (HPC)

InfiniBand

Performance

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

¹For more information, see "MPI over InfiniBand: Early Experiences," by Jiuxing Liu et al., August 2003, www.mellanox.com/pdf/whitepapers/mpi_over_ib_early_experiences.pdf.

other advanced features supported by InfiniBand in the kernel and user space, such as Internet SCSI (iSCSI) Extensions for RDMA (iSER), IP over InfiniBand (IPoIB), SCSI RDMA Protocol (SRP), and Sockets Direct Protocol (SDP).²

ConnectX HCAs are 4X dual-port adapters designed to support PCIe 1.1 and 2.0. ConnectX includes enhanced packet processing and offload capabilities compared with the InfiniHost III chipset, which can help utilize additional bandwidth out of the PCIe slot. The advanced packet scheduler for ConnectX is implemented at the hardware level, and does not involve software or firmware components for its RDMA and send and receive operations.

The enhanced ConnectX architecture is designed to support a messaging rate of up to 25 million messages per second, which helps reduce latency and increase bandwidth compared with InfiniHost III Lx DDR InfiniBand. The theoretical peak bidirectional throughput for a DDR InfiniBand HCA is 4 GB/sec. The ConnectX architecture is designed for enhanced performance and bandwidth, providing a greater messaging rate compared with the previous-generation InfiniBand chipset.

COMPARING HPC INTERCONNECTS

In July 2008, engineers from the Dell HPC team tested the performance of ConnectX DDR InfiniBand, InfiniHost III Lx DDR InfiniBand, and Gigabit Ethernet interconnects on an HPC cluster. Figure 1 shows the cluster configuration used in the test environment.

Server	Four Dell PowerEdge SC1435 servers
Processors	Two dual-core AMD Opteron™ 2216 processors at 2.4 GHz per server
Memory	Eight 1 GB DDR2 dual in-line memory modules (DIMMs) at 667 MHz per server
OS	Red Hat® Enterprise Linux® 5.1 OS running kernel version 2.6.18-53.el5
InfiniBand HCA	<ul style="list-style-type: none"> ■ Mellanox InfiniHost III Lx DDR HCA (MHGS18-XTC) in an x8 PCIe slot ■ Mellanox ConnectX DDR HCA (MHGH28-XTC) in an x8 PCIe slot
InfiniBand switch	24-port Cisco SFS 7000D switch
Gigabit Ethernet network interface card (NIC)	On-board Broadcom NetXtreme II BCM5708 NIC
Gigabit Ethernet switch	Dell PowerConnect™ 6248 switch

Figure 1. HPC cluster configuration in the test environment

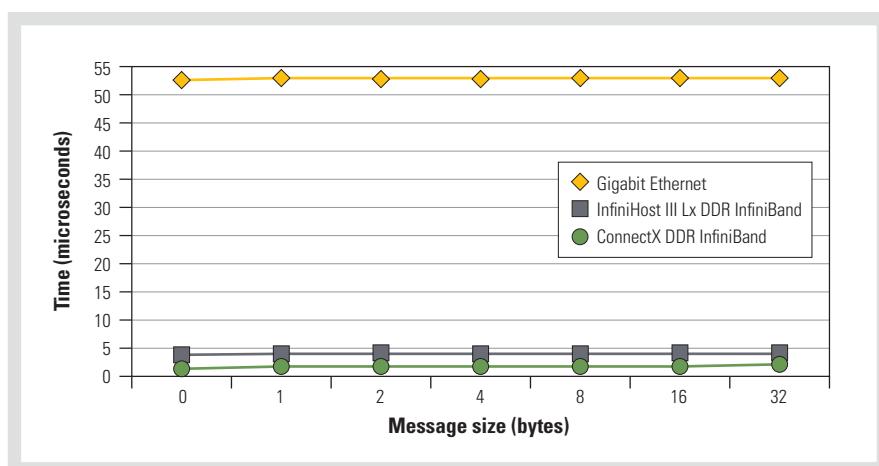


Figure 2. Latency for different interconnects with varying message sizes measured using the Ohio State University MPI-level latency test

The test team used three Ohio State University Message Passing Interface (MPI)-level benchmarks—latency, unidirectional bandwidth, and bidirectional bandwidth—to compare the point-to-point performance of ConnectX DDR InfiniBand

with that of InfiniHost III Lx DDR InfiniBand and Gigabit Ethernet.³ The team performed these tests with two nodes connected through a single switch. In the latency test, the sender node sends a message of a certain data size to the receiver node. The receiver node receives the message and sends back a reply of the same data size, with latency then calculated as half the round-trip time. As shown in Figure 2, ConnectX DDR InfiniBand had a lower latency than InfiniHost III Lx DDR InfiniBand: the small-message latency for ConnectX was approximately 1.6 microseconds, while the small-message latency for

“ConnectX is a fourth-generation InfiniBand architecture and includes multiple enhancements designed to increase performance in HPC clusters.”

²For more information, see “Configuring InfiniBand on HPC Clusters Using the OpenFabrics Enterprise Distribution,” by Munira Hussain, Arun Rajan, Sreeram Vedantham, and Jacob Liberman, in *Dell Power Solutions*, November 2007, DELL.COM/Downloads/Global/Power/ps4q07-20070551-Hussain.pdf.

³For more information on these benchmarks, visit mvapich.cse.ohio-state.edu/benchmarks.

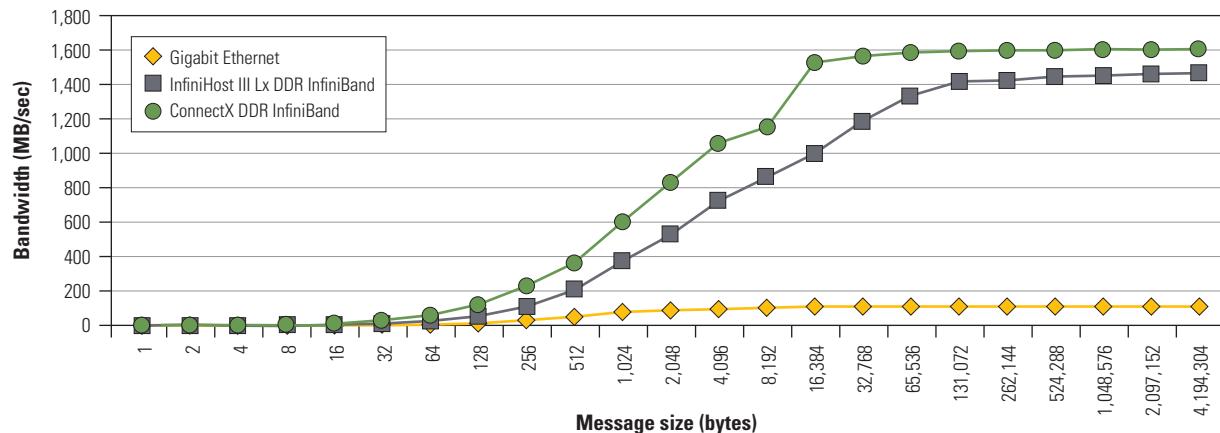


Figure 3. Unidirectional bandwidth for different interconnects and message sizes measured using the Ohio State University MPI-level unidirectional bandwidth test

InfiniHost III Lx was approximately 3.75 microseconds.

The test team next compared the unidirectional and bidirectional bandwidth for ConnectX DDR InfiniBand with those of InfiniHost III Lx DDR InfiniBand and Gigabit Ethernet. In the unidirectional bandwidth test, the sender node sends messages of varying sizes ranging from 1 MB to 4 MB. These messages saturate the interconnect to the receiver node by sending back-to-back messages. As shown in Figure 3, the unidirectional bandwidth for ConnectX was considerably

higher than that of InfiniHost III Lx: the maximum unidirectional bandwidth for ConnectX was approximately 1,600 MB/sec, while the maximum unidirectional bandwidth for InfiniHost III Lx was approximately 1,470 MB/sec.

In the bidirectional bandwidth test, both nodes send back-to-back messages to each other simultaneously. As shown in Figure 4, the bidirectional bandwidth was also much higher for ConnectX than for InfiniHost III Lx: the maximum bidirectional bandwidth for ConnectX was approximately 3,050 MB/sec, while the maximum

bidirectional bandwidth for InfiniHost III Lx was approximately 2,400 MB/sec. Bidirectional bandwidth results can be tuned and optimized by, for example, changing the message size when moving from the Eager protocol to the Rendezvous protocol from 8 KB to 16 KB, which helps increase the throughput for ConnectX.

The test team also used four other benchmarks—Conjugate Gradient (CG), Fast Fourier Transform (FT), Integer Sort (IS), and Lower-Upper Symmetric Gauss-Seidel (LU)—from the NASA Advanced Supercomputing (NAS) Parallel

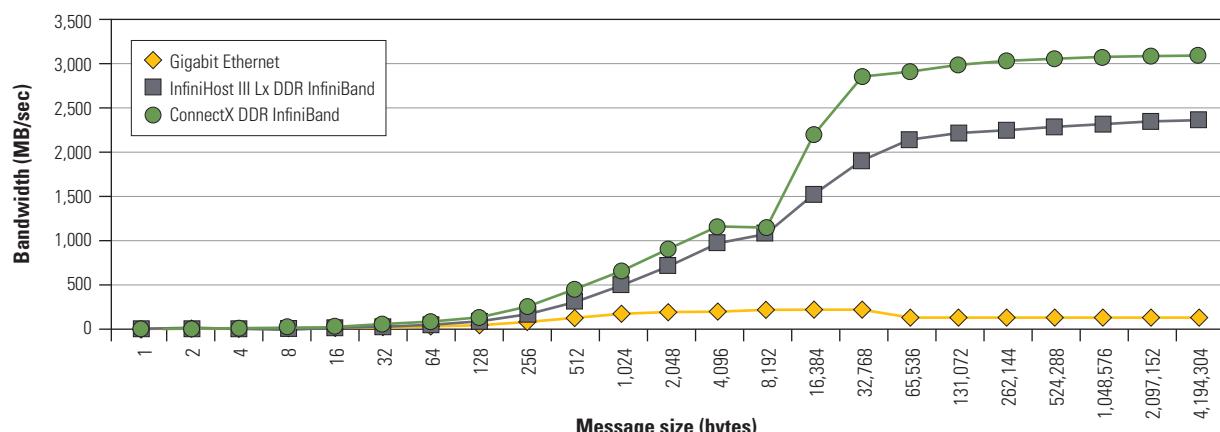


Figure 4. Bidirectional bandwidth for different interconnects and message sizes measured using the Ohio State University MPI-level bidirectional bandwidth test

Benchmarks (NPB) suite to simulate applications with different message-passing characteristics and to analyze the effect of ConnectX DDR InfiniBand on these applications.⁴ The CG, FT, and IS benchmarks use a mix of large and small messages for IPC; the LU benchmark primarily uses small message sizes. The test team compiled the benchmarks for 16 processors and used the Class C problem size, then executed them on the four nodes of the cluster using ConnectX DDR InfiniBand, InfiniHost III Lx DDR InfiniBand, and Gigabit Ethernet. Figure 5 shows the average performance speedup results for InfiniHost III Lx and ConnectX compared against a Gigabit Ethernet baseline. ConnectX provided consistently higher performance than InfiniHost III Lx at various scales, depending on the message-passing characteristics of the specific benchmark.

This performance increase was achieved with four nodes, and will vary with different numbers of nodes, processors, and cores. Actual speedup for a particular application depends on the application's exact message-passing characteristics. Because ConnectX DDR InfiniBand has extremely low latency and enhanced bandwidth, latency- and bandwidth-intensive applications in particular can benefit from use of this interconnect.

OPTIMIZING INTERCONNECT PERFORMANCE IN HPC CLUSTERS

Because different HPC applications have unique characteristics, patterns, and requirements, generalizing optimal HPC configurations is difficult. For communication-intensive applications, however, low latency and high bandwidth have become two key factors that contribute to overall cluster performance. As the number of processors and cores in HPC systems continues to increase, the next-generation Mellanox ConnectX DDR InfiniBand architecture can help meet the

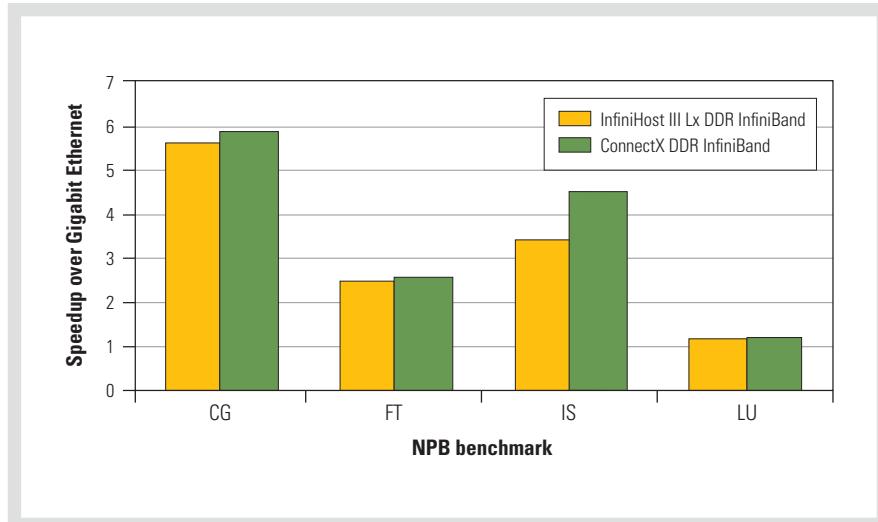


Figure 5. Speedup for InfiniHost III Lx and ConnectX DDR InfiniBand over Gigabit Ethernet measured using four NAS Parallel Benchmarks tests

I/O requirements for these increasingly powerful clusters.

Munira Hussain is a systems engineer and adviser in the Dell High-Performance Computing Group specializing in HPC systems and architecture design. She has a bachelor's degree in Electrical Engineering from the University of Illinois at Urbana-Champaign.

Toby Sebastian is a senior engineering analyst in the Enterprise Solutions Group at the Dell Bangalore Development Center. His current interests include high-end interconnects, HPC clustering packages, and performance analysis of parallel applications. Toby has a B.Tech. in Computer Science and Engineering from the University of Calicut.

Shivaraj Nidoni is a senior engineering analyst in the Dell High-Performance Computing Group. His current interests include high-end interconnects, parallel file systems, and HPC clustering packages. Shivaraj has an M.Tech. in Computer Network Engineering from Visvesvaraya Technological University.

Gilad Shainer is a director of HPC and technical marketing at Mellanox Technologies focusing on HPC, high-speed interconnects, and performance characterization. He has B.S. and M.S. degrees in Electrical Engineering from the Technion – Israel Institute of Technology, and is a member of the PCI-SIG PCI Extended and PCI Express Working Groups, cochair of the OpenFabrics Legal Working Group, and chairman of the HPC Advisory Council.

ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell HPC cluster solutions:
DELL.COM/HPCC

Mellanox Technologies:
www.mellanox.com

OpenFabrics Alliance:
www.openfabrics.org

⁴For more information on these benchmarks, visit www.nas.nasa.gov/Resources/Software/npb.html and see "Performance Comparison of MPI Implementations over InfiniBand, Myrinet and Quadrics," by Jiuxing Liu et al., November 2003, nowlab.cse.ohio-state.edu/publications/conf-papers/2003/liu-sc03.pdf.



SIDE-BY-SIDE SUPPORT FOR WINDOWS SERVER 2008 AND SUSE LINUX ENTERPRISE.

FITS NICELY WHERE ALL
THOSE SERVERS USED TO BE.



The virtualization solution that brings Windows® Server 2008 and SUSE® Linux Enterprise Server together is here. And so is joint customer support from Microsoft® and Novell®. So you can run two, three or even four applications all on the same server with your choice of operating system — and get more reliability, flexibility, efficiency and utilization than ever before. All with clearly defined intellectual property rights and no support headaches.

Need the best Linux? Request your SUSE Linux Enterprise deployment kit and workshop now at moreinterop.com

Novell **Microsoft**®



By Thomas Cloyd
Romy Bauer

INTRODUCING THE DELL POWEREDGE M805 AND POWEREDGE M905 BLADE SERVERS

Dell™ PowerEdge™ M805 and PowerEdge M905 full-height blade servers are designed to deliver outstanding performance in virtualized and database-intensive environments. Combining robust processing power, tremendous RAM capacity, and massive I/O throughput in a space- and energy-efficient, easy-to-deploy blade form factor, these blade servers help increase performance and reliability while helping simplify management.

In enterprise data center environments, IT organizations often rely on virtualization to enhance flexibility, increase resource utilization, and simplify management. At the same time, database-intensive applications such as enterprise resource planning (ERP), customer relationship management (CRM), and messaging applications are placing heavy demands on IT infrastructures. Although many enterprise-class servers can offer the processing power required to meet these needs, I/O throughput and RAM capacity are often limiting factors to delivering outstanding performance and availability in virtualized and database-intensive environments.

The new Dell PowerEdge M805 and PowerEdge M905 full-height blade servers are designed specifically for these virtualized and database-intensive environments. In particular, these blade servers augment robust processing power with expanded dual in-line memory module (DIMM) capacity and three highly available, fully redundant I/O fabrics, which can deliver tremendous RAM capacity and massive I/O throughput in a space- and energy-efficient, easy-to-deploy form factor.

The PowerEdge M805 is a two-socket blade server that provides quad-core AMD Opteron™ processors; 16 DIMM slots for up to 128 GB of RAM; an internal

Secure Digital (SD) card for embedded hypervisors; three highly available, fully redundant I/O fabrics; four integrated network interface cards (NICs); and up to eight high-speed 10 Gigabit Ethernet ports. The PowerEdge M905 is a four-socket blade server that provides quad-core AMD Opteron processors, 24 DIMM slots for up to 192 GB of RAM, and the same SD card capability and I/O benefits as the PowerEdge M805.

Both blade servers are engineered to fit seamlessly within the Dell PowerEdge M1000e modular blade enclosure. This combination of enterprise-class capabilities comparable to those of rack-mounted servers in a blade form factor provides a compelling platform that can help organizations increase performance and reliability, simplify management, and lower total cost of ownership.

DESIGNING FOR HIGH I/O THROUGHPUT

The proliferation of external storage such as storage area networks has increased the importance of I/O reliability, flexibility, and throughput for many organizations. In virtualized environments, I/O reliability and availability are especially critical. To enable true enterprise-class data access, Dell PowerEdge M805 and PowerEdge M905 blade servers offer exceptional I/O throughput, utilizing three highly available, fully

Related Categories:

Blade servers

Dell PowerEdge blade servers

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

redundant I/O fabrics.¹ Figure 1 shows the network fabric architecture for the PowerEdge M805 blade server; the architecture for the PowerEdge M905 blade server is identical, but includes four processors instead of two.

Fabric A consists of four LAN on Motherboard (LOM) Ethernet ports separated into two Broadcom dual-port NICs capable of supporting TCP/IP Offload Engine (TOE) and Internet SCSI (iSCSI) Offload Engine (iSOE) technology.² Fabrics B and C comprise four flexible mezzanine cards that can be customized for Ethernet, 4 Gbps Fibre Channel, or InfiniBand. In particular, two separate Fabric B cards and two separate Fabric C cards are supported with two 1-4 lane, 10 Gbps-capable ports per card that can support both 10 Gigabit Ethernet and 8 Gbps Fibre Channel.

PowerEdge M805 and PowerEdge M905 blade servers contain several features specifically designed to deliver high availability and full redundancy on all three I/O fabrics. For example, LOM and mezzanine cards are physically separated and have physically separate routes and connections to the motherboard. In addition, the top three cards and bottom three cards plug into physically separate enclosure sockets, and each port on each card follows separate physical traces to different I/O switch modules.

PowerEdge M805 and PowerEdge M905 blade servers are also designed to deliver massive I/O throughput on all ports. With four 10 Gbps-capable lanes per port, each mezzanine card port is already capable of supporting quad data rate (QDR) InfiniBand communication. And because the PowerEdge M805 and PowerEdge M905 blade servers offer eight ports in two redundant pairs of four ports each, they can also support four fully redundant QDR InfiniBand ports. Furthermore, all mezzanine card

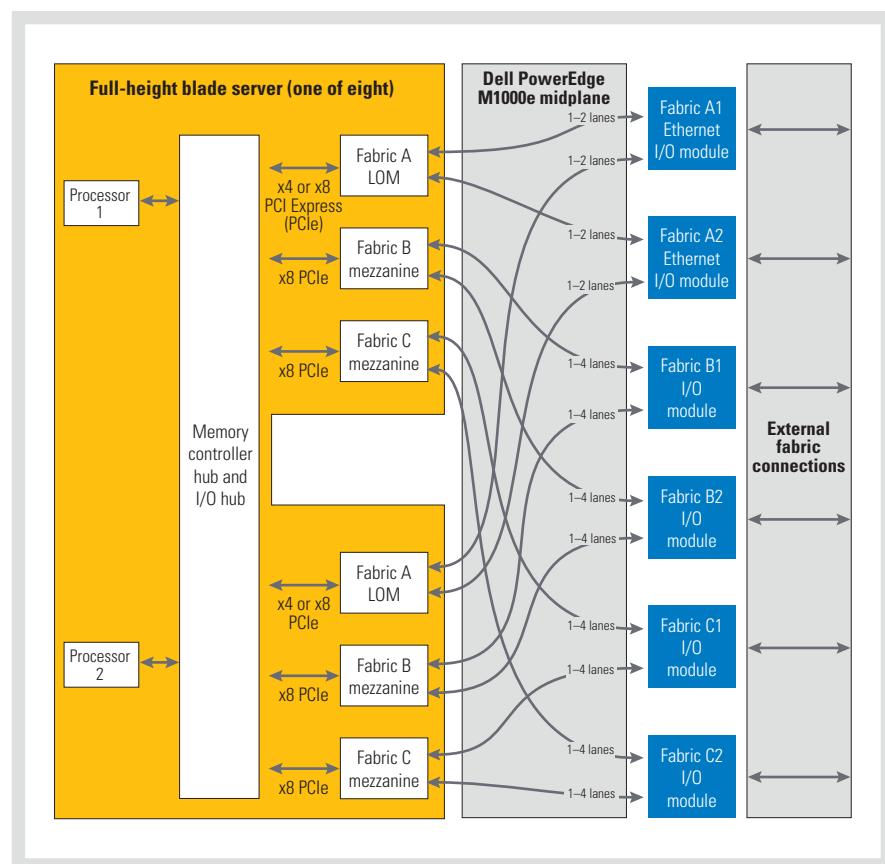


Figure 1. Network fabric architecture for the Dell PowerEdge M805 full-height blade server

ports support 8 Gbps Fibre Channel and 10 Gigabit Ethernet.

In total, the I/O throughput, availability, and redundancy offered by PowerEdge M805 and PowerEdge M905 blade servers are designed to be unmatched by any other blade server currently available. In particular, no other competitive blade server offers three highly available redundant I/O fabrics, and no other competitive blade server offers as much total I/O throughput.

DELIVERING ENHANCED MEMORY CAPACITY

Another key limiting factor in virtualized and database-intensive environments is memory. For example, applications such as Oracle® database software, the Microsoft® Exchange messaging platform,

and the Microsoft SQL Server® database platform can make tremendous demands on memory resources. To help meet increasing memory requirements, Dell PowerEdge M805 and PowerEdge M905 blade servers offer enhanced DIMM capacity that enables organizations to expand RAM configurations cost-effectively and to scale memory without having to increase the number of processors.

In particular, the two-socket PowerEdge M805 has 16 DIMM slots for a total of up to 128 GB of RAM, offering the same DIMM count and capacity as competitive four-socket blade servers. The four-socket PowerEdge M905 has 24 DIMM slots for a total of up to 192 GB of RAM, offering 50 percent more DIMM slots than HP and IBM® blade servers.

¹For more information on the Dell PowerEdge M1000e network fabric architecture, see "Exploring the Dell PowerEdge M1000e Network Fabric Architecture," by John Loffink, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20070500-Loffink.pdf.

²iSOE technology requires purchase of an additional hardware key to implement. For more information on iSOE, see "Introducing Broadcom iSCSI Offload Engine Technology for Dell Servers," by Dhiraj Sehgal, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20080208-Broadcom.pdf.

OPTIMIZING FOR VIRTUALIZATION

In addition to offering significant I/O throughput and DIMM capacity, Dell PowerEdge M805 and PowerEdge M905 blade servers are also optimized for virtualization. For example, the AMD Opteron processors in these servers include Rapid Virtualization Indexing, a feature that utilizes nested paging technology to enable virtual machines to directly manage memory, helping to reduce hypervisor cycles and increase virtualization performance. In addition, optional embedded flash SD hypervisor cards—available for VMware®, Citrix® XenServer™, and Microsoft Hyper-V™ virtualization platforms—can help increase virtualization flexibility and performance.

INCREASING EFFICIENCY AND MANAGEABILITY

Dell PowerEdge M805 and PowerEdge M905 blade servers are designed specifically to fit within the Dell PowerEdge M1000e modular blade enclosure³—enabling enterprise-class performance with the efficiency and ease of a blade form factor. The PowerEdge M1000e enclosure combines

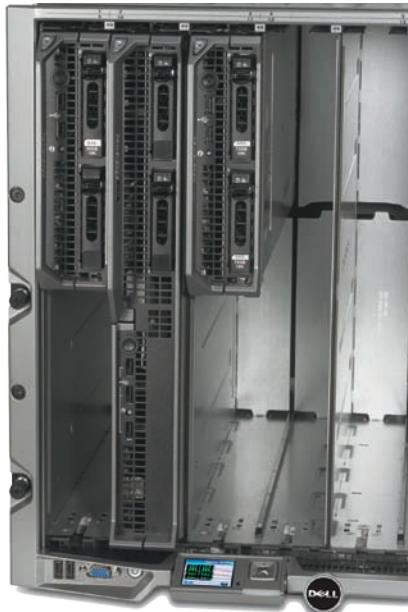


Figure 2. Full-height and half-height blade servers side by side in the Dell PowerEdge M1000e

space- and energy-efficient modularity with scalability, reliability, and ease-of-use features to help simplify server and network management and help reduce downtime.

To help optimize energy efficiency, the PowerEdge M1000e enclosure utilizes

Dell Energy Smart technologies. For example, the enclosure employs a power supply designed for up to 91 percent efficiency that includes Dynamic Power Supply Engagement functionality to help optimize power utilization based on system demands.⁴ In particular, Dynamic Power Supply Engagement puts lightly loaded power supplies into standby mode, thus enhancing the utilization and efficiency of active power supplies. The PowerEdge M1000e also takes advantage of high-efficiency fans designed by Dell and an optimized airflow design to help minimize the amount of power required to cool the enclosure.

The PowerEdge M1000e is also designed for optimal space efficiency. For example, for every five rack-dense 2U servers, administrators can deploy eight PowerEdge M805 servers, and for every five 4U servers, they can deploy 16 PowerEdge M905 servers. Full-height and half-height blade servers can be mixed freely in a single PowerEdge M1000e enclosure (see Figure 2). In addition, reduced cabling and front-side,

DELL FLEXADDRESS STREAMLINES DEPLOYMENT

Traditional blade and rack-dense server environments can be a challenge to deploy and manage. A key contributing factor to this challenge is the complexity of managing storage and network addresses, which can make deploying, upgrading, and replacing blade servers and blade server components time-consuming and disruptive. To help simplify the task of managing addresses, Dell now offers Dell FlexAddress, a patent-pending tool that enables administrators to manage network and storage addresses at the enclosure level instead of at the server level—helping dramatically simplify server management and reduce downtime. Dell FlexAddress is currently available on the Dell PowerEdge M1000e modular blade enclosure both at time of sale and as an upgrade to existing enclosures.

FlexAddress works by locking in a server's Fibre Channel World Wide Name (WWN) and Ethernet and Internet SCSI (iSCSI) Media

Access Control (MAC) addresses to a specific blade slot rather than the blade server itself. Because FlexAddress transfers network and storage identity from individual servers to the enclosure, administrators can deploy, upgrade, and replace individual components or servers without the time-consuming task of changing identity on the network. The technology's approach is to deliver this chassis-centric technology through the chassis itself, making it switch or pass-through agnostic rather than switch delivered and dependent. FlexAddress is enabled through the Chassis Management Controller (CMC) on the PowerEdge M1000e enclosure, and works automatically with I/O modules including Cisco, Brocade, Dell PowerConnect™, and pass-through switches and with I/O mezzanine cards including QLogic, Emulex, and Broadcom cards without requiring a new switch.*

*For more information on FlexAddress, see "Dell FlexAddress for PowerEdge M-Series Blades," by Rick Ward, Mike J. Roberts, and Samit Ashdhir, Dell Inc., June 16, 2008, DELL.COM/Downloads/Global/Products/Pedge/En/FlexAddress%20for%20PowerEdge%20M1000E%20Whitepaper.pdf.

³For more information on the Dell PowerEdge M1000e enclosure, see "The Next-Generation Dell PowerEdge M1000e Modular Blade Enclosure," by Chad Fenner, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080206-Fenner.pdf.

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

Benchmark	Measurement	Dell PowerEdge R805	Dell PowerEdge M805	Dell PowerEdge R905	Dell PowerEdge M905
SPECjbb2005	Java transaction performance	218,882 business operations per second (bops)	219,269 bops	380,461 bops	383,456 bops
VMmark	VMware ESX virtualization performance	7.96 at 6 tiles (36 VMs)	7.88 at 6 tiles (36 VMs)	15.35 at 11 tiles (66 VMs)	15.09 at 11 tiles (66 VMs)

Figure 3. Processor, business transaction, and virtualization performance results for Dell PowerEdge rack-mounted and blade servers

enclosure-level crash cart access help reduce complexity and simplify access.

The PowerEdge M1000e enclosure also offers a range of features designed to simplify deployment and streamline management. For example, it includes Dell FlexAddress technology that can help dramatically simplify server address management (see the “Dell FlexAddress streamlines deployment” sidebar in this article). The PowerEdge M1000e also utilizes FlexI/O technology, which supports modular I/O switches that can be upgraded in place with Gigabit Ethernet stacking and 10 Gigabit Ethernet ports, as well as integrated KVM (keyboard, video, mouse) functionality to support concurrent, enclosure-level access to all 16 slots. Also available are redundant Chassis Management Controllers (CMCs) with automatic failover.

To further enhance manageability, the PowerEdge M1000e is integrated with a range of systems management software products such as Dell OpenManage™, Altiris® Deployment Solution™, Microsoft System Center Operations Manager, and LANDesk Management Suite software. Dell blade servers can be customized, built, and tested in the Dell factory to facilitate rapid, easy deployment.

EVALUATING BLADE SERVER PERFORMANCE

In July 2008, the Standard Performance Evaluation Corporation (SPEC) conducted benchmark tests on Dell PowerEdge M805 and PowerEdge M905 blade servers to evaluate their processing and business transaction performance. The PowerEdge

M805 was equipped with two quad-core AMD Opteron 2360 SE processors at 2.5 GHz and 16 GB of RAM, and the PowerEdge M905 was equipped with four quad-core AMD Opteron 8356 processors at 2.3 GHz and 32 GB of RAM. When compared with April 2008 SPEC benchmark tests of similarly equipped PowerEdge R805 and PowerEdge R905 rack-mounted servers, these benchmarks demonstrate that the blade servers had little or no performance penalty (see Figure 3).⁵ In addition, in SPECjbb2005 benchmark tests for business transaction performance, the PowerEdge M905 delivered better overall performance and performance per watt than comparably configured HP ProLiant BL680c G5, HP ProLiant BL685c G5, and IBM BladeCenter LS41 blade servers.⁶

To evaluate performance in virtualized environments, in July 2008 Dell engineers conducted VMmark 1.1 benchmark tests on a PowerEdge M805 blade server equipped with two quad-core AMD Opteron 2360 SE processors at 2.5 GHz and 32 GB of RAM, and in September 2008 they conducted VMmark 1.1 benchmark tests on a PowerEdge M905 blade server equipped with four quad-core AMD Opteron 8360 SE processors at 2.5 GHz and 64 GB of RAM. Comparing these results with May 2008 VMmark 1.0 tests on a similarly equipped PowerEdge R805 and with September 2008 VMmark 1.1 tests on a similarly equipped PowerEdge R905 again indicates that the blade servers had little or no performance penalty. The PowerEdge M905 was the first four-socket blade or rack server to reach 11 tiles (66 VMs) in

VMmark testing, and as of October 9, 2008, was the highest-performing 16-core blade server tested.⁷

SIMPLIFYING IT WITH FLEXIBLE BLADE SERVERS

Dell PowerEdge M805 and PowerEdge M905 blade servers are designed to deliver robust processing power, tremendous RAM capacity, and massive I/O throughput in a space- and energy-efficient, easy-to-deploy blade form factor. Designed to fit flexibly and seamlessly with other Dell M-Series blade servers in the scalable, highly available, and easy-to-use PowerEdge M1000e modular blade enclosure, these blade servers can help organizations increase performance and reliability, simplify management, and reduce total cost of ownership in enterprise data centers. 

Thomas Cloyd is a member of the Dell Global Relationship Marketing team responsible for launching new blade servers.

Romy Bauer is a business and technology writer based in San Francisco.



MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades

⁵For complete SPECjbb2005 results, visit www.spec.org/jbb2005/results/jbb2005.html.

⁶“SPECjbb Performance and Power Consumption on Multi-processor Intel- and AMD-Based Blade Servers,” by Principled Technologies, August 2008, www.principledtechnologies.com/clients/reports/dell/DellM905SPECjbb0808.pdf.

⁷For complete VMmark results, visit www.vmware.com/products/vmark/results.html.



CHALLENGE

The Wharton School of the University of Pennsylvania needed to refresh its IT infrastructure with ultra-dense servers and storage to minimize its footprint and maximize its compute and storage capabilities in a new data center.

SOLUTION

Wharton consolidated its existing IBM® BladeCenter server and storage infrastructure onto Dell PowerEdge M1000e modular blade enclosures and Dell EqualLogic PS5000X Internet SCSI (iSCSI) storage area network (SAN) arrays.

BENEFITS

- Dell PowerEdge blade servers reduced the number of managed servers by half, with projected performance gains beyond expectations.
- Dell EqualLogic iSCSI SAN delivers high performance at one-quarter the cost of a Fibre Channel SAN.
- Infrastructure minimizes power and cooling needs while providing its user base with productivity-boosting computing opportunities.

Related Categories:

Blade servers, case study, Dell EqualLogic storage, Dell PowerEdge blade servers, Dell ProSupport Services, Internet SCSI (iSCSI), storage, storage area network (SAN), virtualization, Wharton School

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

THE BUSINESS OF SUCCESS

Dell™ PowerEdge™ blade servers and Dell EqualLogic™ Internet SCSI (iSCSI) storage area network arrays help the Wharton School of the University of Pennsylvania quintuple computing power while simplifying management and reducing costs.

With its unmatched faculty and academic programs, the Wharton School of the University of Pennsylvania is one of the most comprehensive sources of business knowledge in the world. Founded in 1881 as the first collegiate school of business, Wharton continues its heritage of innovation today with leading-edge technology in a new data center based on Dell PowerEdge blade servers and Dell EqualLogic Internet SCSI (iSCSI) storage area network (SAN) arrays. The end result will be fast applications for students, faculty, and researchers and additional opportunities to increase learning productivity.

POWERFUL, EFFICIENT SERVERS

"We outgrew the power and cooling capacities of our old data center, and had to prepare for growing data volumes in the coming years," says Dan Alig, senior IT director at Wharton. "So we are refreshing our server and storage infrastructure so that our new facilities will take advantage of all the tremendous compression of processing and storage technologies that Dell now offers with its blade server systems and storage. We're creating a highly efficient data center model that is low in its power consumption and cooling needs and has plenty of room for growth."

One PowerEdge M1000e modular blade enclosure containing 16 PowerEdge M600 blade servers with quad-core Intel® Xeon® processors will house Wharton's EMC® Documentum® eRoom® collaboration software for students as well as VMware® ESX servers. A second PowerEdge M1000e enclosure with 16 similarly equipped blade servers will run a Linux® OS-based grid for performing complex financial algorithms.

Wharton uses PowerEdge 2950 servers as Microsoft® Exchange Server 2007 Hub Transport and Client Access servers and Veritas™ NetBackup™ servers, and PowerEdge 1950 servers as domain controllers. In addition, Wharton is consolidating eight IBM servers down to four PowerEdge R900 servers to run its Exchange environment. The Dell OpenManage™ suite provides a set of standards-based tools for proactive management. "We've essentially quintupled our compute power per

node in the Exchange environment," says Joe Cruz, senior IT leader at Wharton. "In addition to the consolidation benefit, we're now only managing four servers instead of eight and reducing power and cabling on those four. And we're expecting performance to far exceed anything that we could possibly imagine."

"It would have been prohibitively expensive to expand using our existing IBM technology, and ultimately it did come down to cost and service," says Alig. "Dell wins on both cost and service. We want to stretch our dollar as far as it will go, and Dell's letting us do more for less than we could do with other server providers."

COST-EFFECTIVE, HIGH-PERFORMANCE STORAGE

For storage, Wharton upgraded from IBM DS4000 series Fibre Channel and Serial ATA (SATA) array enclosures to Dell EqualLogic PS5000X iSCSI SAN arrays. The main reason the school chose iSCSI SAN storage was the ease of deployment. "We're a big fan of virtualization," says Alig. "And it made sense to marry the server virtualization platform with storage virtualization. It would have been extremely expensive to do that with other vendors' storage technology, but with Dell EqualLogic technology, it's really easy and not that expensive."

Scaling an iSCSI SAN is also easy. "In raw throughput, we're getting—across the nine arrays that we have—easily 250 MB per second sequential and 17,000 to 20,000 I/Os per second," says Cruz. "And that's just with the out-of-the-box performance. It really far exceeded all my expectations in terms of performance."

In addition, the EqualLogic arrays include features such as auto-replication, thin provisioning, writable snapshots, automatic load balancing, volume management, and many others at no additional cost. "Dell EqualLogic delivers Fibre Channel performance in an iSCSI

"We want to stretch our dollar as far as it will go, and Dell's letting us do more for less than we could do with other server providers."

—Dan Alig

Senior IT director at the Wharton School of the University of Pennsylvania
August 2008

SAN at a quarter of the cost," says Cruz. "If you factor in the software features, the cost advantage is awesome." The replication feature utilizes Wharton's existing Ethernet backbone, so the new data center will replicate to one of the existing data centers at no additional cost, with the savings adding up over time as the school takes advantage of all the EqualLogic software features to help reduce administration time and avoid data losses. "Over time we're talking about hundreds of thousands of dollars worth of savings," says Alig.

SIMPLIFIED DEPLOYMENT AND PROACTIVE SUPPORT

Deliveries of new servers and storage from other vendors came in about 60 boxes, requiring approximately 24 hours simply to unbox the equipment. The Dell shipment, in contrast, came fully loaded in 1 box, which took just 4 hours to unbox and assemble. In addition, Dell negotiated with its partner Delaware Valley Liebert for room-wide uninterruptible power supply units as well as racks and power distribution equipment, helping provide a total solution in a cost-effective way.

Wharton also plans to migrate to the VMware ESXi 3.5 embedded firmware hypervisor through Dell. "Right now, provisioning an ESX host is a two-to-three-hour process, partly because we have to provision Fibre Channel storage, and partly because we don't have any imaging

capabilities for ESX," says Cruz. "So moving to an embedded hypervisor is going to save us a lot of time."

Wharton chose Dell ProSupport to fast-track dispatch parts and labor, bypassing basic troubleshooting and helping ensure 24/7 direct access to Dell Expert Centers. "We had a problem with another vendor's SAN and Dell responded before they did," says Alig. "It meant a lot to us that Dell offered us the use of loaner storage. It showed how responsive Dell could be as a partner."

The total solution that Dell provides is finding a warm reception at Wharton. "We're able to do a lot more for less using the Dell blade systems and Dell EqualLogic storage," says Alig. "We've just been amazingly impressed with how responsive Dell has been, whether it's sales support or helping us find technological resources and sharing their vision of the technologies that Dell is moving toward in the future." 

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades



CHALLENGE

With the IT infrastructure at ALERT Life Sciences Computing growing by more than 100 percent annually, the IT department needed to take action to avoid running out of rack space or spending increasing amounts of time on basic server management.

SOLUTION

ALERT deployed a Dell PowerEdge M1000e modular blade enclosure with PowerEdge M600 blade servers to support its internal IT infrastructure. The company connected the enclosure to a Dell/EMC CX3-10c storage area network and implemented PowerEdge R805 servers to run VMware® virtualization software.

BENEFITS

- Dell engineers installed the new PowerEdge M1000e enclosure in just over one hour.
- Virtual machines can be deployed up to 75 percent faster than non-virtualized physical servers.
- Simplified management tools help reduce reboot time from 15 minutes to less than a minute.
- Two virtualized servers do the work of 20 non-virtualized physical servers.
- Dell ProSupport for IT helps protect investments with rapid support.

Related Categories:

ALERT Life Sciences Computing, blade servers, case study, Dell PowerEdge blade servers, virtualization

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

HEALTHY IT INFRASTRUCTURE

Dell™ PowerEdge™ blade and rack servers provide the foundation of a high-density virtualized infrastructure at ALERT Life Sciences Computing—supporting simplified management, high performance, and rapid deployment.

ALERT Life Sciences Computing, a medical services provider based in Portugal, has developed rapidly since its launch in 1999, with annual growth rates averaging 100 percent. "In just three years, we went from 4 to 37 servers," says Filipe Pinto, system administrator at ALERT. When the company moved to new offices with a new data center, Pinto saw an opportunity to consolidate its servers and help streamline its server management process. "We said that this was the time to reorganize our server infrastructure so we could continue to meet the growing demands of the business without running out of rack space," he says.

HIGH-DENSITY INFRASTRUCTURE

Pinto wanted servers that not only delivered high-density computing, the latest Intel® Xeon® processors, and high energy efficiency, but also enhanced remote manageability. During the week, manually rebooting servers in the new offices took at least 15 minutes, requiring him to descend nine floors to the data center. And on weekends, when Pinto traveled 200 miles away to his hometown, the task was thoroughly impracticable.

ALERT decided to implement virtualization technology to help overcome these challenges—providing an effective way to expand its infrastructure with limited rack space while helping simplify management and enabling rapid deployment of new applications. Pinto compared several leading IT vendors, looking specifically at their blade server technology and virtualization offerings. "Compared with the blade servers of other IT solution providers, Dell blade servers offered the best mix of high-density computing, processing power, energy efficiency, and manageability," says Pinto.

ALERT deployed a Dell PowerEdge M1000e modular blade enclosure and 14 Intel Xeon processor-based PowerEdge M600 blade servers. "The servers arrived prebuilt along with the chassis," says Pinto. "A couple of days later, Dell Infrastructure Consulting Services arrived and an engineer installed everything in just over an hour. We'd not experienced anything like it before." He adds, "My time is expensive, so for Dell to install all the blades and chassis and do it so quickly is of great benefit to me and the company."

EFFICIENT, HIGH-PERFORMANCE BLADE SERVERS

Dell PowerEdge M600 blade servers are designed to provide high levels of performance and reliability as well as low power consumption. In addition to providing up to 60 percent

greater computing density than traditional 1U servers, tests have shown that PowerEdge M600 blade servers consume up to 19 percent less power than other leading blade servers.¹ "We are very pleased with the performance of the Dell PowerEdge M600 blade servers," he says. "They are incredibly stable and high performing."

Pinto has also deployed several other Dell systems, including three Dell/EMC CX3-10c storage area networks and numerous PowerEdge servers. The blade servers run Microsoft® Navision, Microsoft Office SharePoint® Server, Microsoft SQL Server®, and VMware ESX software, as well as Oracle® Database 10g databases. All the blade servers are connected to a Dell/EMC CX3-10c storage array, with the entire infrastructure supporting approximately 4 TB of data.

SIMPLIFIED SERVER MANAGEMENT

According to Pinto, server management has become incredibly simple with Dell technology. For example, the Avocent iKVM (integrated keyboard, video, mouse) switch in the Dell PowerEdge M1000e modular blade enclosure enables Pinto to view, monitor, and manage the blade servers. "I have control over all my blades from a single console," Pinto says. Other key features of the switch include local access to each blade server, integration with the Avocent On-Screen Configuration and Activity Reporting (OSCAR) graphical user interface for configuration and to work with existing Avocent KVM infrastructures, and integration with the Microsoft Active Directory® directory service.

The iKVM switch also allows access to the integrated Dell Remote Access Controller (iDRAC) on each blade server, which helps simplify remote management. "I can control and monitor each blade using a simple Web browser-based graphical user interface," Pinto says. "I can inspect

"We are very pleased with the performance of the Dell PowerEdge M600 blade servers. They are incredibly stable and high performing."

—Filipe Pinto

System administrator at ALERT Life Sciences Computing
October 2008

event log records, manage local users, and shut down or reboot the machine."

Pinto says it takes less than a minute to reboot a system—a drastic reduction from the 15 minutes this task previously took. "I really notice the difference toward the end of each month when I'm writing up reports," Pinto says. "I can get on with my work instead of being interrupted by server management tasks." Through the company's virtual private network, Pinto can also seamlessly complete management tasks such as rebooting servers from home without having to make a long trip back to the office.

RAPID DEPLOYMENT AND RESPONSIVE SUPPORT

The new Dell infrastructure highlights the value of VMware virtualization technology. "I wanted to introduce VMware for a long time," says Pinto. "The alliance between Dell and VMware gave us confidence to deploy virtualization technology in our infrastructure."

ALERT deployed two Dell PowerEdge R805 servers, which are designed for optimum virtualization performance. In addition to offering twice the memory and I/O scalability of previous 2U servers, the PowerEdge R805 combines high performance with low energy consumption. By running the VMware ESX virtualization platform, these 2 servers can do the same work as 20 non-virtualized servers while helping significantly reduce deployment

times, Pinto says. "I can set up a virtual machine in 30 minutes, when it took me up to two hours to install a physical machine in the past," he says. "Imagine how much time that frees up to work on more strategic, higher-value tasks."

Pinto estimates that a system failure would cost ALERT more than US\$1,000 (€676) an hour. For this reason, the company's IT systems must be reliable and backed by responsive support services. At the heart of the Dell ProSupport for IT service used by ALERT is a 24/7 commitment from Dell Expert Centers in which IT professionals with case management skills can help Pinto troubleshoot. In addition, Next Business Day On-Site Service After Remote Diagnosis helps ensure delivery and installation of new or replacement parts. With the Dell infrastructure and Dell ProSupport for IT, Pinto believes he has the right combination of products and services to help avoid periods of expensive downtime. 

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell virtualization solutions:
DELL.COM/Virtualization

¹ "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.



CHALLENGE

PT Ciptadana Capital needed to scale its existing blade server infrastructure to support the extension of its online trading services. With around 5,000 customers connected to the Ciptadana trading application, the company needed a solution that would support further expansions and allow new users to connect directly to the trading application.

SOLUTION

Ciptadana deployed Dell PowerEdge M600 blade servers to power its online trading and brokerage application.

BENEFITS

- Ease of deployment helped provide simple and efficient consolidation from rack servers to blade servers.
- Centralized management console helps simplify server management.
- Dell consultants helped provide rapid server deployment and migration and assisted with migration of Oracle® database server configuration.
- Dell PowerEdge M600 blade servers help reduce energy use while helping save on infrastructure costs and space.

Related Categories:

Blade servers, case study, Dell PowerEdge blade servers, energy efficiency, green IT, PT Ciptadana Capital

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

BLADES SHIFT BROKER INTO HIGH GEAR

Dell™ PowerEdge™ blade servers help PT Ciptadana Capital support future expansion of its security trading service while reducing energy use, simplifying management, and optimizing limited data center space.

PT Ciptadana Capital, one of the top 10 brokerage firms in Indonesia, relies on its world-class data center to provide timely and responsive information to its customers and enable the execution of real-time trading. Provisioning of real-time trade data, reinforced by seamless back-office operations that handle settlements, corporate actions, dividends, and cash management aspects of trade executions, has made Ciptadana a leading broker in the high-net-worth retail equity brokerage market.

A key challenge faced by the company was finding a way to meet future expansion requirements stemming from the anticipated growth of its customer base. "Our ability to execute trades and provide critical financial information is tied to the efficiency of our IT infrastructure," says Benny Haryanto, president director of Ciptadana. "In particular, trading online with the stock exchange demands a robust and scalable data center that our customers can trust to deliver real-time trades whenever, wherever they need it."

SUPPORTING CUSTOMER GROWTH

With more than 200 employees and around 5,000 customers connected to its online trading platform across two branch offices, Ciptadana began looking for the high-performance processing capacity it needed to supplement its existing data center and support future expansion plans. The company's existing data center comprised blade servers in two separate server enclosures running a remote trading platform, domain controller services, and an e-mail server.

The selection of blade servers over rack servers was a critical decision made by the IT department. "With two existing server enclosures, we had limited space available in the data center," Haryanto says. "At the time, we were looking for several additional servers to support our expansion plans, and given the space limitations, blades were a definite consideration."

After conducting an internal review, Ciptadana determined that it would require servers with two quad-core processors and at least 8 GB of RAM. "Our older servers were a quad-core system with 4 GB of RAM," Haryanto says. "Since we required fast I/O access with quick database searches and network response, we decided we would need more processors to scale performance."

Ciptadana worked with Dell consultants and chose Dell PowerEdge M600 blade servers, each with two quad-core processors and 8 GB of RAM. The blade servers

were deployed at Ciptadana's Jakarta-based data center within a day, and were in testing and development phases for five months before going live. "Dell consultants were quite helpful throughout the acquisition process," says Haryanto. "They worked with us to refine the configuration and also assisted us with the migration of our Oracle database server configuration."

REDUCING ENERGY CONSUMPTION

Because of space limitations in its data center, Ciptadana had significant restrictions on use of power sources for reliable server availability. "Our existing power redundancies for the enclosures we purchased previously were operating well," Haryanto says. "However, with the impending expansion, energy consumption was a serious consideration and a challenge we needed to address."

Energy consumption was a key criterion for Ciptadana when deciding to select an all-blade configuration. An ultra-efficient power supply; high-flow, low-power fans; and optimized airflow are some of the integrated energy-saving features of the PowerEdge blade servers that helped reduce energy consumption. Haryanto says, "The Dell PowerEdge M600 blade servers offer enhanced design features that brought savings in energy consumption. We noticed a difference during the testing and development phase."

"Our existing power redundancies for the enclosures we purchased previously were operating well. However, with the impending expansion, energy consumption was a serious consideration and a challenge we needed to address."

—Benny Haryanto
President director of PT Ciptadana Capital
September 2008

SIMPLIFYING SERVER MANAGEMENT

With the two existing server enclosures and additional blade servers, Ciptadana also wanted simplified management. "Our IT team is relatively small, given the scale of the data center and the customers we were expecting to serve," Haryanto says. "It's a very time-consuming and labor-intensive job because you cannot afford any component in the system to break down during trades that can take place in split seconds and involve millions of dollars."

The Dell PowerEdge blade server system has helped simplify server management for Ciptadana. In particular, the company can now easily manage the entire system from the built-in centralized management console. From a single console, the IT staff has redundant, secure access paths to manage multiple enclosures

and blade servers. Dell consultants also provided Ciptadana with expert advice in configuring the Altiris® Deployment Solution™ for Dell Servers server management application. "The system is very user-friendly, giving us a single point of control to manage everything in one place," Haryanto says.

ENABLING EFFICIENT USE OF SPACE

Shifting to an all-blade configuration helped Ciptadana optimize limited data center space. "We were already experiencing a tight squeeze with our two server enclosures in the original configuration," Haryanto says. "Originally, we used four racks of non-blade servers, but with this solution we've replaced them with just two enclosures of blade servers. That's almost a 50 percent savings in space, and no new enclosures were needed." 

"The Dell PowerEdge M600 blade servers offer enhanced design features that brought savings in energy consumption. We noticed a difference during the testing and development phase."

—Benny Haryanto
President director of PT Ciptadana Capital
September 2008

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades



CHALLENGE

Optimal Blue sought to optimize its technology platform by adopting energy-efficient servers and flexible storage.

SOLUTION

The company upgraded from stand-alone servers with direct attach storage to a Dell PowerEdge M1000e modular blade enclosure with PowerEdge M600 blade servers and a Dell EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) array.

BENEFITS

- Optimal Blue assembled the blade system and SAN storage in 15 minutes with no previous SAN experience.
- Dell blade servers and storage provide an anticipated 50 percent reduction in application response time.
- Elimination of single points of failure helps improve reliability and availability.
- A 50 percent reduction in hardware space helps save US\$1,000 per month.
- Energy-efficient systems help save US\$500 per month on power.

Related Categories:

Blade servers, case study, Dell EqualLogic storage, Dell PowerEdge blade servers, Internet SCSI (iSCSI), Optimal Blue, storage area network (SAN)

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

MORTGAGE LENDERS' MOMENT OF TRUTH

A high-performance, energy-efficient solution based on Dell™ PowerEdge™ blade servers and a Dell EqualLogic™ PS Series Internet SCSI (iSCSI) storage area network array provides Optimal Blue with an anticipated 50 percent improvement in application response time.

The mortgage industry is in a state of crisis. Secondary markets have been disrupted and are often liquid—and volatility is reflected in the cost of funding. Optimal Blue is confronting this turmoil by providing accurate information on the pricing and timing of mortgage products.

"If the pricing on a loan is out of date by a couple of minutes, it's worthless," says Dr. Ivan Darius, founder and co-CEO of Optimal Blue. "In challenging market conditions such as these, when investors' Web sites are nonresponsive and investor guidelines and rate sheets are changing rapidly, our technology and service differentiate us from the rest."

ENERGY-EFFICIENT SERVERS AND SCALABLE STORAGE

Darius and cofounder Larry Huff have created a comprehensive service that provides brokers and lenders with the information they need to sell into the secondary market. The company's 140 customers purchase the information through a software-as-a-service (SaaS) model.

Optimal Blue needs its IT platform to perform to the highest standards to help ensure that the company can provide fast and accurate mortgage market information to customers. "We found that we wanted to improve reliability at peak times without having to double our hardware," says Darius. "Also, in our previous storage solution the database servers each had their own Fast SCSI drives. This led to a lot of difficulties with growing databases and the inability to add disk capacity because of chassis and other constraints. Also, the failure rate of the disks we were using was higher than we expected. We wanted a more power-efficient solution which would cost less to cool and occupy less space. It was obvious that a blade system would be the answer."

To support a new release of the Optimal Blue search engine, Darius switched to a Dell PowerEdge M1000e modular blade enclosure containing PowerEdge M600 blade servers with quad-core Intel® Xeon® processors. He also upgraded from his direct attach storage to a Dell EqualLogic PS5000XV Internet SCSI (iSCSI) storage area network (SAN) array.

"We saw from published numbers that Dell was the leader in the performance-per-watt equation with the M1000e offering," says Darius. "We also like the fact that Dell provides us with tools that enable us to manage the system remotely. We are using the Dell Chassis Management Controller and the integrated Dell Remote Access Controller and the Dell EqualLogic Group Manager to manage the SAN."

Darius plans to virtualize the servers to achieve maximum use of the blades and help keep down power and space consumption. "We will turn all these blades into a virtual computer," he says. "With management tools such as VMware and a good SAN, we can be a lot more flexible."

The EqualLogic iSCSI SAN array was key to the plan. As the company's customer base grew, the process of moving data on the existing systems to new and larger servers was painful without scalable, flexible storage. Darius did his due diligence and considered the alternatives. "We were looking for a more reliable and expandable storage mechanism," he says. "We found that EqualLogic gives us more redundancy, not only at the disk level, but also at the server level. We can have multiple servers in a backup mode to run certain databases if we need them. With RAID-50 storage, we also increase our reliability by a significant factor. And the EqualLogic SAN is inherently stackable; if you need another one, you just add it into your switch network."

FAST, FLEXIBLE SETUP

With the flexibility and power of the Dell PowerEdge M1000e system and Dell EqualLogic iSCSI SAN array, Darius feared the setup process would be laborious, but he was in for a pleasant surprise. "We put the whole system together—the blades and the storage—in a minimal amount of time," he says. "It took less than 15 minutes to get the Dell

"With our current system, our response time is 3 to 5 seconds. With the Dell blades and storage, we expect to cut response time in half."

—Ivan Darius
Founder and co-CEO of Optimal Blue
September 2008

blades and the storage set up. We are specialists, but we are not SAN specialists, and the fact that we could do this without help in such little time was remarkable." With the built-in storage virtualization capabilities and advanced automation of the EqualLogic PS Series, Optimal Blue can also seamlessly expand its SAN without disruption of service.

Optimal Blue's previous server solution filled four cabinets, and Darius was close to renting more space. Once again, he was pleasantly surprised. "With the Dell blades and storage, we'll need one and a half cabinets—about 50 percent less space, which saves \$1,000 per month," he says. "In addition, we're also saving \$500 per month on power, so we're benefiting from a greener and more cost-friendly solution.

"We're going green as a company, and the Dell solution is helping us take that step," says Darius.

a day, and at any given time we have between 1,500 and 3,000 users logged in. However, we have 35,000 licensed users, so there's a potential for much higher usage," Darius says.

By eliminating stand-alone servers, each with its own access to storage disks, Optimal Blue will eliminate single points of failure, a key cause of potential downtime. Although Optimal Blue is still in the process of implementing the Dell solution, it is already planning to enhance stability and reliability. Darius intends to purchase another Dell EqualLogic iSCSI SAN array and put it in a backup site for a hot disaster recovery solution. "Having a failover solution will give us a lot of assurance as we grow and add customers," says Darius. "It's an important step for the future, and we trust the Dell EqualLogic iSCSI solution and our production Dell blades to do the job." 

ENHANCED RESPONSE TIME AND RELIABILITY

Optimal Blue is expecting a dramatic reduction in the response time of its search engine. "With our current system, our response time is 3 to 5 seconds," Darius says. "With the Dell blades and storage, we expect to cut response time in half." Reducing the response time is a benefit to customers because it makes their workflows faster and more productive. "We support about 40,000 searches



QUICK LINKS

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



CHALLENGE

Scientists at the Deutsches Elektronen-Synchrotron (DESY) research center rely on IT to support major research into the structure of matter. With the arrival of new data from the world's largest particle accelerator—the Large Hadron Collider—DESY needed to expand its server infrastructure and ensure high performance in limited space.

SOLUTION

Dell PowerEdge blade servers provide DESY scientists with a powerful environment that helps maximize uptime for critical calculations. Dedicated Dell services including Premier Pages and Custom Factory Integration help ensure a smooth implementation and ongoing simplified management.

BENEFITS

- Blade server solution was up and running in just nine days.
- Thousands of physicists worldwide can analyze petabytes of data.
- Efficient blade servers help minimize power, cooling, and space requirements.
- Dell Premier Pages and Custom Factory Integration services help significantly speed installation and simplify asset management.

Related Categories:

Blade servers, case study, data consolidation and management, Dell PowerEdge blade servers, Deutsches Elektronen-Synchrotron (DESY), green IT

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

RE-CREATING THE BIG BANG

Powerful Dell™ PowerEdge™ blade servers deployed in just nine days provide high levels of performance and efficiency to help put a German research center on track for a major scientific breakthrough.

The Deutsches Elektronen-Synchrotron (DESY) research center, located in Hamburg and Zeuthen in Germany, hosts more than 3,000 scientists working to identify the structure of matter. For almost 50 years, DESY has developed and created particle accelerators used in photon science and particle physics. The center plays an important role in the latest experiments conducted at the European Organization for Nuclear Research (CERN) Large Hadron Collider (LHC)—the world's highest-energy particle accelerator. From the information gathered at the LHC, physicists hope to examine how matter behaved a fraction of a second after the Big Bang, at the moment the universe began, and find explanations for dark matter and the origins of mass.

At approximately 10 PB every year, the volume of data produced by the LHC is vast. And, given the size of the project, the LHC is unable to process all the data in Geneva. Instead, it sends raw data to regional distributors. Grid computing centers worldwide—including DESY—then receive and process this data, performing complex calculations and analysis.

To support its scientists' work, DESY has two data centers that house more than 2,000 servers and currently store more than 5 PB of data produced by the center's own particle accelerators. But in preparation for rapidly increasing data volumes—from the LHC project in particular—DESY was keen to enhance its infrastructure. "We constantly review our IT infrastructure to ensure that it supports our critical work," says Dr. Knut Woller, deputy head of IT at DESY. "We undertake fundamental research projects that can go on for more than a decade, with results that have a worldwide impact in the scientific community and beyond."

The location and design of the data centers presented the DESY IT group with a further challenge. Following expansion, DESY created a second data center in the computer center in 2005. It provides room for 1,500 rack units and an air-conditioning capacity of 500 kW. It was essential that the new infrastructure make optimum use of this space. "The less space and energy each new server consumes, the more computing power we can host without expanding into a second site," Woller explains.

OPTIMIZED CONNECTIVITY AND ENERGY EFFICIENCY IN REDUCED SPACE

DESY conducted a European tender in search of a solution to meet its specific requirements. For the past three years the center has worked closely with Dell. "We've found Dell servers more reliable than those of our previous providers. What's more, we have a great partnership," says Woller. The measurements, tests, and vendor discussions during the tender convinced the DESY team that Dell PowerEdge M600 blade servers, together with PowerEdge M1000e modular blade enclosures, would offer the best price/performance and result in a favorable total cost of ownership compared with competing products.

Crucially, the blade servers are designed for energy efficiency. "Based on our own measurements, we found Dell blades 10 to 15 percent more energy efficient than competitor products," Woller says. And by providing 60 percent greater density than traditional 1U servers, the blade servers offer the required computing power in the data center's restricted space.

In total, DESY has more than 300 blade servers in modular enclosures, and an additional 80 have been ordered. Approximately 90 percent of the capacity is used for scientific calculations on particle physics. The remaining 10 percent runs daily IT tasks such as databases and e-mail.

"Connectivity is also very important," adds Woller. "The multiple I/O interfaces in the Dell blades mean we can access our data through different technologies, including iSCSI and Fibre Channel. The blades also support InfiniBand, which offers a higher bandwidth and lower latency so that we can rapidly exchange data within a set of servers." Woller adds, "Dell provides the most cost-effective solution as our project needs continue to grow. The blade servers meet our needs with regard to space as well as being highly energy efficient."

"Dell provides the most cost-effective solution as our project needs continue to grow. The blade servers meet our needs with regard to space as well as being highly energy efficient."

—Dr. Knut Woller
Deputy head of IT at DESY
October 2008

RAPID IMPLEMENTATION THROUGH DEDICATED DELL SERVICES

The Dell Premier Pages and Custom Factory Integration (CFI) services helped DESY dramatically cut the implementation time with its latest blade server deployment. The Premier Pages service enables the DESY team to order and configure systems online, while CFI services play a vital role in speedy deployment. "Dell support, and in particular Dell CFI, has halved the time taken to deploy and install servers compared with any previous provider's solution," says Woller. "The first delivery was of 32 servers. These were fully operational in less than seven days, which set a new record at the center for server setup."

A week before a server delivery, CFI sends an e-mail message informing DESY exactly which IT assets are scheduled to arrive. "With Dell CFI, we can configure the network and management ports before the products arrive," says Woller. "The servers arrive with individually assigned addresses, and we can immediately place them in the racks. The whole process is simple and the system is up and running fast, saving us valuable time."

POWERFUL, RELIABLE SYSTEMS

The staff at the computer center continues to be impressed by the support and products provided by Dell. "With previous solutions, we've encountered a number of failures over a three-month period, but

with Dell blades we haven't had a single issue," says Woller.

For additional reassurance, DESY has Dell ProSupport for IT with the Mission Critical option, which includes a four-hour on-the-spot response time for enclosure issues and next-business-day support for individual blade servers. "We have excellent relationships with Dell technical experts and are very satisfied with the comprehensive levels of support," says Woller.

Crucially, the powerful system supports the center's pivotal role in the LHC investigations into the origins of the universe. "With our current Dell infrastructure, thousands of LHC scientists worldwide can run analysis on petabytes of data stored here at DESY," says Woller. As the LHC project continues, and while preparing future accelerators for photon science research, DESY knows it can rely on Dell systems to support scientists working toward the next breakthrough. 

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge blade servers:
DELL.COM/Blades



By Dylan Locsin

Travis Vigil

HIGH-DENSITY, HIGHLY SCALABLE STORAGE: DELL EQUALLOGIC PS5500E iSCSI SANs

The Dell™ EqualLogic™ PS5500E Internet SCSI (iSCSI) storage area network (SAN) array allows organizations to simply and cost-effectively deploy high-density, highly scalable, and consolidated SANs for tiered storage, primary and secondary applications, and disaster recovery operations.

In today's business environment, organizations must store, retain, and protect ever-increasing amounts of data. The skyrocketing amount of e-mail, file, and digital media data and the constant growth of databases necessary to perform critical functions have dramatically increased the storage requirements for production and near-line environments. Stringent regulatory legislation such as the Sarbanes-Oxley Act has intensified data retention and archiving requirements for many organizations as well. Furthermore, the need to preserve business continuity in such a regulated and challenging environment underscores the importance of data protection and disaster recovery.

At the same time, deploying ever-growing storage infrastructures can be costly and complex and can take up valuable data center space. To help address these challenges, the highly scalable Dell EqualLogic PS5500E Internet SCSI (iSCSI) array enables organizations to consolidate storage and deploy high-density storage area networks (SANs) simply and cost-effectively. This approach of consolidation and simplification helps organizations to save space in their data centers, enhance efficiency and manageability, and ultimately lower the total cost of providing storage services for tiered primary applications, secondary applications, and disaster recovery operations.

Related Categories:

Dell EqualLogic storage

Internet SCSI (iSCSI)

Storage

Storage area network (SAN)

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

OVERCOMING OBSTACLES TO LARGE-SCALE DEPLOYMENTS

As storage infrastructures within data centers continue to grow, the physical space required to house them has become an increasingly significant contributor to overall cost. In many data centers, especially those in urban areas, available floor and rack space is limited, and the cost of adding space can be prohibitively expensive. When organizations turn to hosted data center facilities for key operations such as disaster recovery, the amount of space utilized can greatly affect overall cost as well. In addition, the power and cooling needs of sprawling server and storage infrastructures can significantly drive up energy costs.

Large-scale storage infrastructures based on traditional SAN architectures can also be more difficult to operate efficiently and manage effectively than small storage infrastructures. With traditional physically oriented storage systems, management operations such as load balancing, performance optimization, capacity planning, and expansion can be cumbersome and time-consuming; at best they are performed periodically with careful planning, and at worst they are performed infrequently with well-intentioned guesswork. Also, day-to-day administrative tasks such as provisioning logical units (LUNs) and supporting data protection and recovery can become excessively burdensome and require large IT staffs as environments grow.

SECURE WITH DELL

The scalable Dell EqualLogic PS5500E array helps an Alaskan data security company manage IT services for more than 160,000 potential new users.

Sequestered Solutions Alaska, LLC (SSA), provides data security for a wide range of commercial clients as well as state and federal government agencies. In selecting storage systems to support its current managed IT services and storage-as-a-service offerings, the company requires reliable hardware that can ensure 24/7 availability. "Whether our clients are running a patient management system, providing software as a service, or backing up critical government data, we need to ensure that they can access information whenever they need it, without fail," says Lara Baker, CTO at SSA.

SSA needed a new storage solution quickly. The company was planning to offer a new service called DataCenter Vault, a comprehensive, online, endpoint data-protection and backup application for laptops and desktops. SSA anticipated supporting up to 3,000 clients on that service within 60 days and more than 160,000 individual clients within six months.

After exploring a range of storage options, SSA selected the Dell EqualLogic PS5500E Internet SCSI (iSCSI) storage area network (SAN) for its scalability, reliability, and cost-effectiveness. "We considered purchasing several small storage arrays for the DataCenter Vault service, but we learned that the Dell EqualLogic PS5500E iSCSI SAN could offer a simpler, more cost-effective solution," says Baker. "We have used Dell EqualLogic SANs for several years, and we were confident in the reliability of the hardware."

SSA decided to test a beta unit of the EqualLogic PS5500E, and Baker's team was able to get the SAN up and running rapidly. "We were able to deploy the new system in just 15 minutes, after installing the physical array in our

"By housing up to 48 TB of raw capacity in a single 4U space, the Dell EqualLogic PS5500E SAN leaves us lots of rack space to support other clients and other revenue-generating services."

—Lara Baker
CTO at Sequestered Solutions Alaska, LLC
October 2008

ENABLING MASSIVE SCALABILITY WITH VIRTUALIZED STORAGE

The Dell EqualLogic PS5500E is a high-density, highly scalable iSCSI storage array in the Dell line of virtualized EqualLogic SANs. It enables organizations to simply and cost-effectively deploy high-capacity storage within an extremely small form

factor—24 TB to 48 TB of raw data within a single 4U, 48-drive device. EqualLogic PS5500E arrays can be configured with forty-eight 500 GB or 1 TB Serial ATA (SATA) disk drives and can be combined to create a SAN with up to 576 TB of raw storage under a single management interface. The EqualLogic peer storage architecture

provides enterprise-class reliability and redundancy with dual controllers in each modular array. Each controller in the SAN contains three Gigabit Ethernet ports and 2 GB of memory. Because of the 48 drives on the system, dynamic load balancing, virtual architecture, and controller design, the EqualLogic PS5500E can deliver good



advanced rack technology," says Baker. "It took longer to unpack the box than it did to deploy the SAN."

The density of the Dell EqualLogic SAN enables SSA to optimize data center real estate. "The more space that a system occupies, the less space we have for other services," says Baker. "By housing up to 48 TB of raw capacity in a single 4U space, the Dell EqualLogic PS5500E SAN leaves us lots of rack space to support other clients and other revenue-generating services."

REDUCING OPERATING COSTS AND SIMPLIFYING PROVISIONING

With simplified configuration and management, the Dell EqualLogic PS5500E SAN can help SSA deploy new storage capacity rapidly. "With other SANs, it might take hours or days to deploy new storage," says Baker. "Once the Dell EqualLogic PS5500E was installed, we deployed new storage in 15 minutes after powering up the PS5500E."

With the potential to scale out to 576 TB by integrating additional units under a single management interface, the EqualLogic PS5500E also provides an easy way to accommodate growth without interrupting service to clients. "If we run out of space on a single system, we can just buy an additional unit, connect it to the existing system, and it is recognized immediately as additional available storage," says Baker. "The newly expanded system operates as a single pool of storage, so it's very easy to manage. We don't have to spend time worrying about where data is located on the SAN. We can just allocate capacity and keep moving."

SSA estimates that the EqualLogic PS5500E will help the company save significant acquisition and management costs compared with competing SANs. "We estimate this new high-density SAN technology will save us in excess of one million dollars over the next three years versus midrange fibre-based competitive products," says Baker. "And because the Dell EqualLogic SAN is so much simpler to use than other products, we estimate that it will be seven to eight times less expensive in total operating costs. Those are tremendous savings that will help us stay competitive in the field."

performance for random I/O applications (see the “Secure with Dell” sidebar in this article).

In addition to being operated as a stand-alone SAN with one pool of storage, EqualLogic PS5500E arrays can be used in a tiered storage model. In this model, multiple pools of storage within a single SAN contain EqualLogic Serial Attached SCSI (SAS) and/or SATA arrays for different levels of performance or different service levels, which are provided according to application or user group needs. For example, Oracle® volumes for online transaction processing and Microsoft® Exchange mailboxes for executives might reside in one storage pool of 15,000 rpm SAS drives in a RAID-10 configuration, with moderate-I/O Microsoft SQL Server® databases, general workgroup e-mail services, file storage, and development volumes residing in another pool containing an EqualLogic PS5500E iSCSI SAN in a RAID-50 configuration (see Figure 1). Administrators can manage multiple arrays as a single SAN, configure them with the same or different RAID level, and move LUNs between pools and arrays without service interruption.

Like other storage arrays in the EqualLogic PS Series, the EqualLogic PS5500E utilizes a modular, virtualized peer storage architecture designed to enhance disk utilization, deliver high-I/O performance, and simplify management compared with non-virtualized storage architectures. The EqualLogic architecture provides these advantages in part by automating load balancing across disks, RAID sets, connections, caches, and controllers. EqualLogic arrays automate several key functions to help further simplify storage administration, including array configuration, volume setup and expansion, and storage pooling.¹

The EqualLogic PS5500E also comes standard with a comprehensive suite of data management and protection

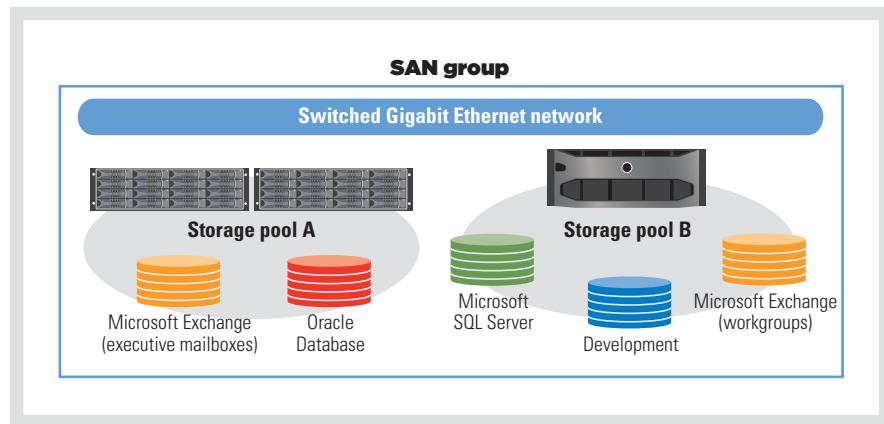


Figure 1. Mixing drive capacity and performance profiles helps increase storage flexibility

software features designed to simplify management, help ensure availability, and facilitate recovery from local data corruption or site disasters. Notably, an auto-replication feature provides data replication for off-site disaster recovery, and Auto-Snapshot Manager provides application-consistent data protection and recovery for Microsoft Exchange, SQL Server, and Windows® file systems, as well as hypervisor-aware protection of VMware® Infrastructure 3 environments.² These software features are included at no additional cost.

MEETING BOTH PRIMARY AND SECONDARY STORAGE NEEDS

Because of its outstanding capacity and density, solid random I/O performance, and enterprise software features, the Dell EqualLogic PS5500E iSCSI SAN array can be especially well suited to large-scale, primary production data centers where space and management are key considerations. In fact, given the storage consolidation and enhanced utilization enabled by the EqualLogic PS5500E, an organization could potentially deploy an entire enterprise-class data center in a single rack.

The EqualLogic PS5500E is also well suited to consolidation of storage for

secondary production applications, backup, or development. Its cost-effectiveness for large capacities and density makes it appropriate for online storage for files, backup copies, data archives, and development environments, which can be a safe and easy-to-manage alternative to offline or direct attach storage systems. With its extremely high density and capacity, the EqualLogic PS5500E can scale to meet evolving data retention and archiving requirements. Furthermore, the EqualLogic PS5500E is certified with a variety of backup vendors including CommVault, Symantec, and others.

The consolidation made possible by the EqualLogic PS5500E is well suited to hosted, co-located, minimally staffed (or *lights-out*) disaster recovery sites as well. In a hosted environment, where space is often rented by the rack or square foot, space utilization can significantly affect overall cost. The EqualLogic PS5500E enables massive capacity in a small footprint, enabling organizations to dramatically reduce space requirements for disaster recovery operations. Furthermore, automation helps simplify management when needed most—during disaster recovery. Because EqualLogic arrays are designed to replicate to any other EqualLogic array, a production SAN can be on

¹For more information, see “Inside the EqualLogic PS Series iSCSI Storage Arrays,” by John Joseph, Eric Schott, and Kevin Wittmer, in *Dell Power Solutions*, February 2008, DELL.COM/Downloads/Global/Power/ps1q08-20080249-EqualLogic.pdf.

²For more information, see “How Dell EqualLogic Auto-Snapshot Manager/VMware Edition Helps Protect Virtual Environments,” by Andrew Gilman and William Urban, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20090107-Gilman.pdf.

performance-oriented arrays such as the EqualLogic PS5000XV, but still be protected cost-effectively with minimal rack space with the EqualLogic PS5500E.

By helping reduce space requirements and simplify management, the EqualLogic PS5500E enables organizations not only to consolidate storage, but also to lower the cost of deploying, managing, and growing a large-scale storage infrastructure. With the EqualLogic architecture's ability to drive high utilization rates and high performance with reduced disk space, high-density consolidation on the EqualLogic PS5500E can contribute to reduced power and cooling costs as well. Like all EqualLogic storage arrays, the EqualLogic PS5500E provides a comprehensive suite of data monitoring, management, and protection features that help reduce ownership costs further by eliminating add-on license fees.

Reprinted from *Dell Power Solutions*, November 2008. Copyright © 2008 Dell Inc. All rights reserved.

SIMPLIFYING STORAGE MANAGEMENT

The ever-increasing demand for storage can strain IT space and budgets. The density, performance, and manageability of the Dell EqualLogic PS5500E iSCSI SAN array can help organizations simply and cost-effectively deploy highly consolidated, high-capacity storage for a range of data center operations including production and near-line applications, data backup, and disaster recovery. 

Dylan Locsin is a product marketing consultant for the Dell EqualLogic PS Series. He has 8 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.

Travis Vigil is a product marketing strategist for Dell storage solutions. He has nearly 10 years of experience with technology companies including Intel and Dell. He has a B.S. from Stanford University and an M.B.A. from Northwestern University's Kellogg School of Management.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell EqualLogic PS5500E iSCSI SAN arrays:
DELL.COM/EqualLogic
DELL.COM/PSSeries

Seagate 

The highest-capacity hard drive designed specifically for Tier 2/Nearline enterprise data center requirements



Get maximum value at an affordable price. The new **Seagate® Barracuda® ES.2 SAS** hard drive lets you reduce power consumption while you boost performance.

- NEW! Combines SAS (Serial Attached SCSI) enterprise capabilities with capacities up to 1 TB
- SUPERIOR! Data integrity over SATA
- GREENER! Lower power consumption than SATA with interposer card
- FASTER! System performance with SAS over SATA
- TRUE! Enterprise host compatibility with no single point of failure

www.seagate.com/barracudaes2



Copyright © 2008 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, the Wave logo and Barracuda are registered trademarks of Seagate Technology LLC in the United States and/or other countries. All other trademarks are the property of their respective owners. When referring to hard drive capacity, one terabyte, or TB, equals one trillion bytes. Your computer's operating system may use a different standard of measurement and report a lower capacity. In addition, some of the listed capacity is used for formatting and other functions, and thus will not be available for data storage. Seagate reserves the right to change, without notice, product offerings or specifications.



By Barbara Craig

SEAGATE SAS DRIVES PROVIDE OPTIMIZED NEAR-LINE STORAGE FOR DELL SYSTEMS

New Seagate® Barracuda® ES.2 Serial Attached SCSI (SAS) hard drives, available in Dell™ PowerEdge™ servers and Dell PowerVault™ storage, are designed to provide the exceptional performance, efficient operation, robust data integrity, and seamless integration required by near-line ecosystems in enterprise data centers.

Growing data access requirements for users and applications, internal policies on system uptime and reliability, and other challenges have resulted in increasing demand for cost-effective, high-capacity near-line storage (also known as *tier 2* storage). Low-cost storage that must be frequently replaced, however, may force other systems to operate in a degraded mode, failing to deliver the required level of service and even driving up costs by demanding repeated maintenance and additional redundancy at the application or system layer. Although Serial ATA (SATA) hard drives can provide a capable, cost-effective foundation platform for near-line storage, IT organizations are constantly seeking higher levels of performance, efficiency, and reliability than these drives can provide.

The new Seagate Barracuda ES.2 Serial Attached SCSI (SAS) hard drive can help meet these needs, providing significant and comprehensive advantages for virtually every aspect of capacity-intensive enterprise applications (see Figure 1). Available in Dell PowerEdge servers and Dell PowerVault storage (see the “High-performance Seagate SAS drives in Dell systems” sidebar in this article), these drives are designed to maintain high performance even in the presence of rotational and linear vibration commonly encountered in high-density enterprise IT environments, as well as efficient operation, robust data integrity, and seamless integration into existing infrastructures.

Related Categories:

Seagate
Serial ATA (SATA)
Serial Attached SCSI (SAS)
Storage

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

EXCEPTIONAL PERFORMANCE AT THE DRIVE AND SYSTEM LEVEL

In April 2008, Seagate engineers tested the performance of the Seagate Barracuda ES.2 SAS hard drive using the Storage Performance Council (SPC) component-level SPC-1C and SPC-2C tests with an on-site SPC auditor present. These SPC benchmarks are designed to measure hard drive performance by simulating real-world workloads: SPC-1C simulates random workloads, while SPC-2C simulates large-scale sequential data movement of near-line storage applications characterized by large I/Os organized into one or more concurrent patterns, such as large file processing (large computer-aided design files or file copy operations), large database queries (scans or joins of large relational tables, or data mining or business intelligence), and video on demand (individualized video entertainment or drawing from a digital film library).

The test system was a server with two Intel® Xeon® processors at 2.8 GHz, 2 GB of RAM, a Super Micro X6DH8-XG2 motherboard, and an LSI SAS3041X-R SAS/SATA host bus adapter running the Microsoft® Windows Server® 2003 Enterprise Edition OS. The test team ran the benchmark tests against five different 1 TB hard drives in this system: a Barracuda ES.2 SAS drive, a Barracuda ES.2 SATA drive, and three third-party SATA drives. In the SPC-1C tests, the Barracuda ES.2 SAS drive provided 8-54 percent higher throughput than the third-party SATA drives.

In the SPC-2C tests, the Barracuda ES.2 SAS drive provided 38–93 percent higher composite throughput than the third-party SATA drives (see Figure 2).¹

The performance advantages of near-line Barracuda ES.2 SAS drives over SATA drives extend to the system level, where the full-duplex, bidirectional SAS architecture allows systems to transmit data on two data channels simultaneously. This transmission might entail a write command on one channel and a read command on the other, or, for maximum throughput, the two channels can be combined into a single wide port that can concurrently transmit write or read commands.

OPTIMIZED PERFORMANCE PER WATT

A growing emphasis on reducing energy consumption and lowering data center operating costs has made hard drive power consumption and performance per watt key metrics for enterprise IT departments. Near-line SAS drives typically use similar amounts of power as comparable SATA drives—for example, 1 TB and 750 GB Seagate Barracuda ES.2 SATA drives are specified at 8.0 W when idle and 11.6 W when operational, with the addition of an interposer card adding 1.2 W for a total of 9.2 W when idle and 12.8 W when operational. Comparable 1 TB and 750 GB Barracuda ES.2 SAS drives are specified at 9.0 W when idle and 12.5 W when operational, and do not require an interposer card.

To demonstrate the power optimization of near-line SAS drives, in July 2008 Seagate

Exceptional performance	<ul style="list-style-type: none"> 33 percent average I/O performance boost over comparable SATA drive High levels of throughput in sequential read and write commands Full-duplex, bidirectional I/O and 7,200 rpm platter speeds Enhanced rotational vibration tolerance in high-density enclosures
Optimized performance per watt	<ul style="list-style-type: none"> Similar power consumption to a comparable SATA drive, and lower operational power consumption than a comparable SATA drive with an interposer card Seagate PowerTrim™ technology, which helps reduce power consumption, heat, and energy costs without sacrificing performance
Robust data integrity	<ul style="list-style-type: none"> End-to-end initiator-target nexus checking that helps ensure data moving to or from the drive is not misdirected, helping eliminate a key cause of data corruption Dual ports that enable failover and data recovery in multi-host storage environments
Simplicity and continuity in deployment	<ul style="list-style-type: none"> Easy enterprise-level management with support for up to 16 hosts per drive queue rather than one host for SATA drives Seamless SAS connectivity, avoiding the need for SATA interposer cards and helping reduce costs while minimizing servicing complexity Compatibility with existing SCSI software and middleware
Optimization for near-line enterprise storage	<ul style="list-style-type: none"> Support for up to 1 TB of storage for capacity-intensive applications Designed for 24/7 operation and a mean time between failures of 1.2 million hours Hot-swappable drives and four-wire cabling for quick, simple serviceability Increased reliability by eliminating single points of failure imposed by single-ported SATA drives

Figure 1. Seagate Barracuda ES.2 SAS drives can provide significant benefits in enterprise IT environments

engineers used the Iometer benchmark to compare the throughput of a Barracuda ES.2 SAS drive with that of a Barracuda ES.2 SATA drive in a Dell PowerVault MD1000 disk expansion enclosure configured with a Dell SAS 5/E controller. The write cache was disabled for these tests. The results demonstrated that the SAS drive provided a 7 percent improvement in sequential reads, 14 percent improvement in sequential writes, 100 percent improvement in random reads, and 9 percent improvement in random writes over the SATA drive, with an average 33 percent performance boost over the SATA drive.

As these results indicate, the SAS drive provided significantly higher performance

than the SATA drive with only fractionally greater power consumption, resulting in higher performance per watt. Adding an interposer card (frequently used to provide dual-port connectivity to SATA drives) can make the power consumption of a SATA drive equal to or even greater than that of a SAS drive, helping increase the performance-per-watt advantage of the SAS drive even further.

ROBUST DATA INTEGRITY AT THE SYSTEM LEVEL

Physical interconnect malfunctions can cause up to 68 percent of storage subsystem failures.² Lacking inherent dual-port failover capability, SATA systems often report such interconnect failures as a drive-not-found error, which can result in administrators misdiagnosing the problem as a drive failure rather than silent corruption while the data was being transmitted.

Given this limitation, it follows that near-line SATA drives and adapters typically develop silent data corruptions significantly more often than Fibre Channel or SAS drives. RAID storage arrays offer

HIGH-PERFORMANCE SEAGATE SAS DRIVES IN DELL SYSTEMS

High-performance, cost-effective Seagate Barracuda ES.2 Serial Attached SCSI (SAS) drives are currently supported in the following Dell systems:

- Dell PowerEdge server models 840, 860, 6850, 1900, 1950, 2900, 2950, 2970, 6950, 6970, SC440, SC1430, SC1435, R200, R300, R900, T105, T200, and T605
- Dell PowerVault storage models MD1000, MD3000, MD3000i, NF500, and NF600

¹For complete details, visit www.storageperformance.org/results.

²“Are Disks the Dominant Contributor for Storage Failures? A Comprehensive Study of Storage Subsystem Failure Characteristics,” by Weihang Jiang, Chongfeng Hu, Yuanyuan Zhou, and Arkady Kanevsky, FAST ‘08: 6th USENIX Conference on File and Storage Technologies, February 26–28, 2008, San Jose, CA, www.usenix.org/event/fast08/tech/full_papers/jiang/jiang.pdf.

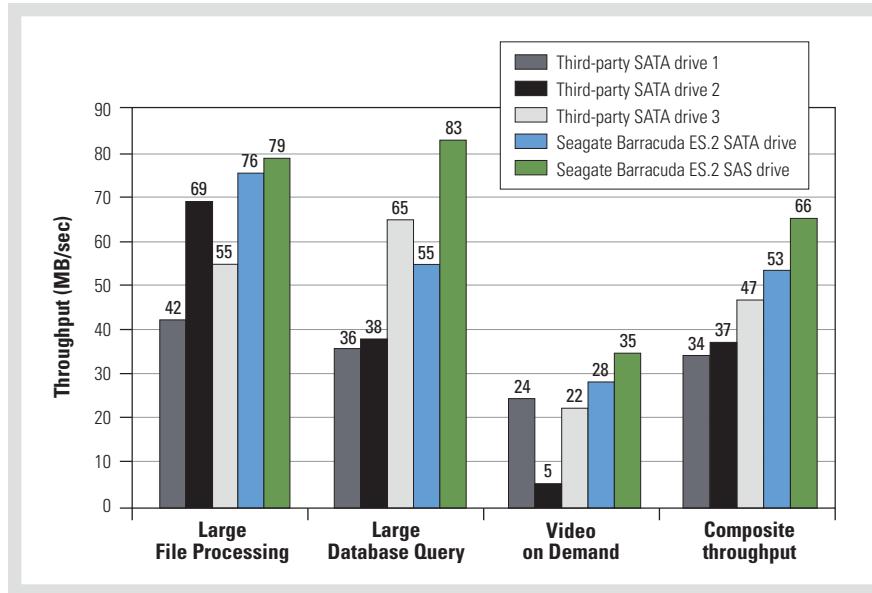


Figure 2. Seagate Barracuda ES.2 SAS drives can provide significantly higher performance than SATA drives

no protection from silent data corruption, because current file systems use checksums to help detect corrupt data when the data is read back, which may not happen until months after the corrupted data was written. Such delays can increase the chances that the uncorrupted data has been lost forever.

The goal of a RAID architecture is to provide data redundancy in the event of a failure, but by definition data corruption is an undetected failure—one that RAID offers no protection against. SAS drives help reduce data corruption with end-to-end initiator-target nexus checking, a key component of data integrity that helps ensure that data traveling to or from the drive is not misdirected. SATA drives cannot achieve this level of enterprise-class data integrity because they lack native addressability (although purchasing and deploying interposer cards can provide that capability, at additional cost).

Furthermore, SAS drives help decrease storage system failure rates by reducing the number of physical interconnects and adding dual-port redundancy (see the “Simplicity and continuity in deployment” section in this article). Administrators also gain enhanced access to advanced error reporting, which can

provide valuable information for effective troubleshooting.

SIMPLICITY AND CONTINUITY IN DEPLOYMENT

IT professionals can be reluctant to deploy storage systems that may disrupt operations. Near-line SAS drives, however, are designed to seamlessly integrate into the same SAS infrastructures that currently support critical tier 1 storage.

In addition, near-line SAS drives not only help eliminate the additional expense and complexity of SATA interposer cards, but also offer enhanced connectivity and failover capabilities. SAS command queuing enables one drive to accommodate up to 16 hosts, compared with the single host that a standard SATA hard drive can accept through native command queuing. Interposer cards can enable command queuing for a second host connection, but this approach is typically still inadequate for enterprise storage environments.

Support for a large number of host connections helps SAS drives avoid the single-point-of-failure risk that characterizes SATA drives. By avoiding the need for a SATA interposer card, SAS drives can also reduce total system parts count—a key consideration when designing for

high reliability. In addition, SAS hard drives can help ensure operational continuity and investment protection because of their compatibility with existing SCSI software and middleware.

OPTIMIZED, COST-EFFECTIVE SAS DRIVES

Enterprise-ready Seagate Barracuda ES.2 SAS hard drives are designed, tested, and optimized to handle the rigors of 24/7 operation in enterprise data centers. Providing exceptional performance, efficient operation, robust data integrity, and seamless integration with existing infrastructures, these drives can help organizations deploy optimized, cost-effective storage in their near-line/tier 2 storage ecosystems. 

Barbara Craig is a senior product marketing manager at Seagate, where she is responsible for product marketing for Seagate’s Enterprise Compute products, which comprise a broad range of high-capacity, high-performance, low-power hard drives for 24/7, business-critical IT environments. She has held a number of marketing positions over a 30-year career at various high-tech companies, and was most recently outbound product marketing manager at QLogic, where she was responsible for global marketing for original equipment manufacturer accounts. Barbara attended the State University of New York, Pepperdine University, and Vanderbilt University.

MORE
 **ONLINE**
DELL.COM/PowerSolutions

QUICK LINKS

Seagate Technology:
www.seagate.com

Dell PowerEdge servers:
DELL.COM/PowerEdge

Dell PowerVault storage:
DELL.COM/PowerVault

UP TO **48TB** IN ONE LITTLE ARRAY



**INTRODUCING THE DELL EQUALLOGIC™ PS5500E,
AN ECONOMICAL, ENTERPRISE-CLASS STORAGE SOLUTION.**

Highly-scalable iSCSI SAN solution

3x the drives with only 1U more space than previous PS5000 Series arrays

All-inclusive feature set delivers the functionality you need right away

SIMPLIFY YOUR STORAGE AT
DELL.COM/PSseries/PSNovPS5500E





By Sanjeet Singh
Jeff Echols

FAST, RELIABLE DATA PROTECTION FROM DELL AND COMMVAULT

Dell and CommVault have teamed up to create a next-generation disk-based backup system that combines leading-edge Dell™ hardware with innovative CommVault® data protection software. The new state-of-the-art Dell PowerVault™ DL2000 – Powered by CommVault integrates disk-based backup and recovery with de-duplication technology to help deliver fast, reliable data protection.

Faced with exponential data growth, many IT organizations are struggling with the amount of critical information they must store, manage, administer, and protect on a daily basis. At the same time, they are also seeking solutions to help increase backup reliability, decrease data protection costs and complexities, and meet stringent compliance and e-discovery service-level agreements (SLAs).

For example, a 30 percent annual data growth rate would mean that the amount of data that backup and recovery systems must support would more than double in the next three years. Because organizations may experience annual growth up to or exceeding 60 percent, backing up and restoring vital data within acceptable time frames can become increasingly difficult. To help increase reliability and streamline the process, many organizations have looked beyond traditional tape-based backup to disk technologies, which are typically both faster and more reliable than tape. According to a Gartner user survey analysis on storage management software usage driven by replication, de-duplication, and virtualization, the overall trend toward backup to disk continues, with 48 percent of the survey's respondents stating that their

first-line backup is to disk compared with only 35 percent backing up directly to tape.¹

To help meet the backup needs of organizations of all sizes, Dell has worked with two leading backup software companies, CommVault and Symantec, to introduce two new backup-to-disk appliances: the Dell PowerVault DL2000 – Powered by CommVault and the Dell PowerVault DL2000 – Powered by Symantec Backup Exec.² The new Dell PowerVault DL2000 – Powered by CommVault is a next-generation disk-based backup solution that combines high-performance Dell hardware with advanced CommVault data protection software. It provides a state-of-the-art solution that integrates disk-based backup and recovery with de-duplication technology to help provide fast, reliable data protection in a cost-effective way.

INCREASING STORAGE EFFICIENCY WITH DE-DUPLICATION

Disk-based backup has generally been available in two primary forms: traditional backup software written to disk, and virtual tape libraries (VTLs). VTLs, which use disks to emulate tape drives, gained popularity for their ability to create multiple libraries and

¹"User Survey Analysis: Storage Management Software Usage Driven by Replication, Deduplication and Virtualization," by Alan Dayley, Gartner, Inc., March 4, 2008.

²For 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.

write multiple streams of data simultaneously. Many VTLs are compatible with leading backup software and function as plug-and-play appliances, which can be deployed easily in data centers and remote locations.

Although VTLs are adequate short-term resolutions to backup and recovery bottlenecks, they are separate systems that can result in disjointed silos of isolated data that end up being managed as disparate point solutions. For administrators at midsize organizations, VTLs can also lack sufficient capacity, neither scaling well nor taking advantage of the true random I/O nature of native disk technology.

The introduction of innovative data de-duplication technologies has accelerated the migration to disk-based backup and recovery. Data redundancy is a major contributor to ever-increasing data sprawl. For example, if a single file is created and shared with 10 people, the environment now contains 10 extra copies of that file. These 10 copies may then be backed up to disk, resulting in a total of 20 extra copies, and then replicated, increasing the total to 30 extra copies. Over time and multiplied by all the data generated by an organization, these extra copies can begin consuming a large amount of storage resources—requiring organizations not only to purchase, deploy, and maintain the necessary storage hardware, but also to pay for the power, cooling, and other infrastructure to support that hardware. De-duplication is designed to eliminate these redundancies to help accelerate backups, reduce hardware costs, and reduce recovery times while alleviating the administrative burden of managing duplicate data.

As with VTLs, however, not all de-duplication solutions are equally effective. Common deployment challenges are related to performance, increased management complexity, and islands of de-duplication. Block-based de-duplication can affect recovery times because each file must be reassembled, generating high processing overhead during a recovery

window that is often short. It can also be difficult for administrators to extract data from a de-duplication system for long-term retention on tape or other storage media—information can become trapped inside the system, forcing a continuous cycle of adding proprietary hardware or software that must then be managed, increasing administrative complexity.

Finding and eliminating redundant data with high-end, server-based de-duplication can be expensive. Fortunately, the far-reaching benefits of de-duplication are now within the reach of midsize organizations seeking high-performance, simplified, cost-effective solutions.

DEPLOYING NEXT-GENERATION DISK-BASED BACKUP

Unlike traditional backup systems and VTLs that focus on short-term goals such as accelerated backups, the Dell PowerVault DL2000 – Powered by CommVault is a comprehensive backup and recovery platform that uses storage policies to move data copies to the optimal storage device—including tape—automatically. This turnkey solution also provides automated storage policy setup and centralized management, which can ease many aspects of backup,

recovery, and de-duplication. Organizations can also seamlessly integrate advanced add-on features such as archiving and replication to help scale and increase functionality as demands dictate.

The PowerVault DL2000 – Powered by CommVault is designed for simple, powerful, integrated disk-based backup, recovery, and de-duplication. Setup and operation do not require separate software or devices; instead, administrators can manage their data protection through a centralized console to help unify administration across an entire enterprise. The backup software also comes integrated with dynamic disk provisioning, which is designed to set up un-configured disks and put them into immediate use (see Figure 1). The addition of application-aware, file-based de-duplication helps ensure rapid data recovery, improved storage efficiency with streamlined manageability, and high performance.

The system is available in two cost-effective configurations designed to help meet the needs of midsize Microsoft® Windows®, Novell® NetWare®, Linux®, or UNIX® OS-based physical or virtualized server environments. Both configurations include a Dell PowerEdge™ 2950 server,

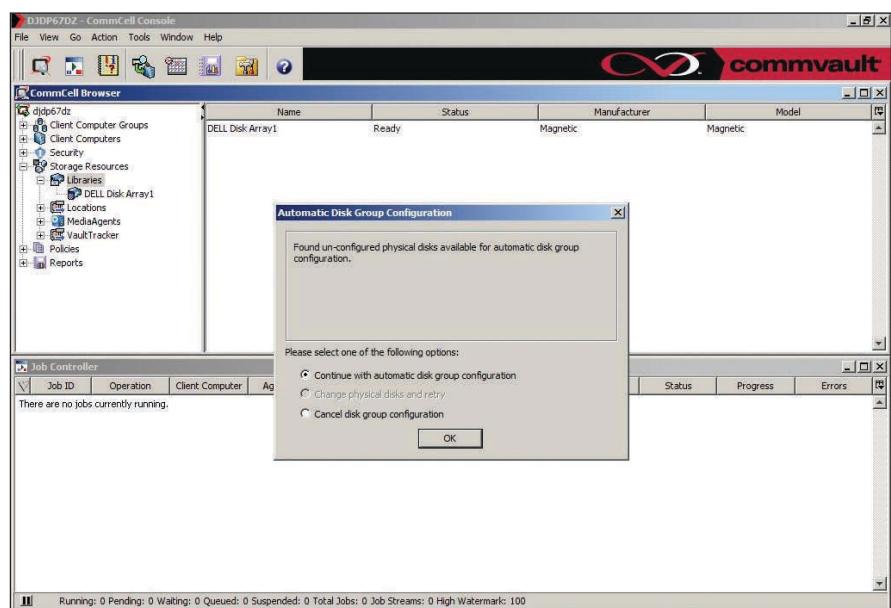


Figure 1. The PowerVault DL2000 – Powered by CommVault can automatically recognize and set up the RAID configuration for new disks and put them into immediate use

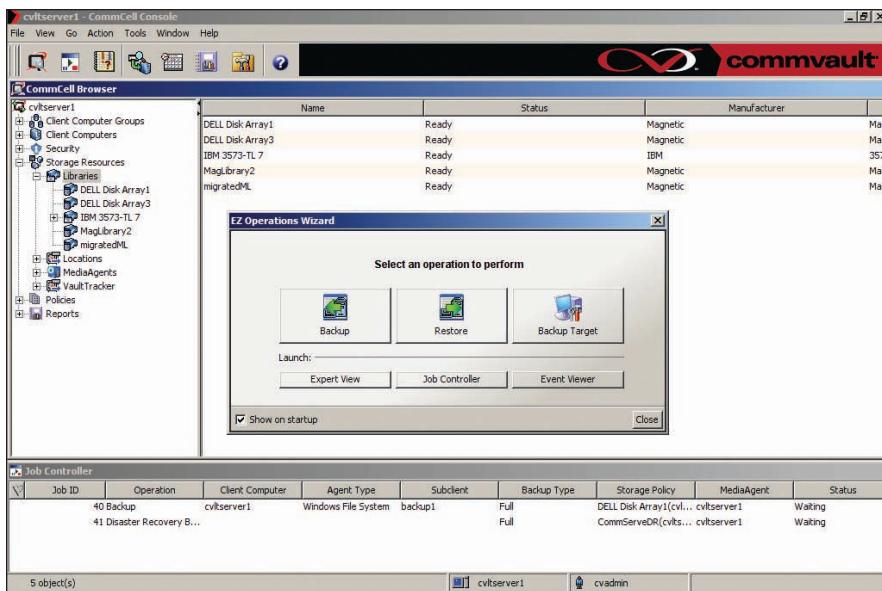


Figure 2. An easy-to-use wizard helps administrators manage backup and restore operations

a Dell PowerVault MD1000 disk expansion enclosure, and the Microsoft Windows Server® 2008 OS with CommVault software consisting of CommServe® and MediaAgent software, a disk-to-disk license for 3 TB of capacity, and five client agents with the CommVault Data Classification Enabler. The advanced configuration adds file de-duplication to help reduce redundant data copies during backup and archive jobs. Additional application agents are available for systems running Microsoft Active Directory®, Microsoft Exchange, Microsoft Office SharePoint® Server, Microsoft Windows SharePoint Services, Microsoft SQL Server®, Oracle® Database, IBM® Lotus® Notes, and Novell GroupWise® software to help deliver granular attribute, e-mail, or document recovery. File and e-mail archiving agents are also available for removing stale data from a primary disk while still leaving stubs for recalls.

The PowerVault DL2000 - Powered by CommVault can also provide a variety of other benefits, including the following:

- Simplified backup operations, including built-in reporting, single-console management, and an auxiliary copy feature, which can move data copies from disk to disk or disk to tape

automatically without affecting the host (see Figure 2)

- Host software compression to help reduce network traffic and increase storage efficiency
- Minimized downtime because of the ability to recover a single file, e-mail, or document
- Flexibility to de-duplicate data across backup, archived, and replicated data, helping reduce storage requirements by up to 15 times compared with tape
- Reduced network loads when replicating backup data to a centralized environment, helping reduce the need for IT resources in branch offices
- Scalability that can grow from a few to thousands of servers without requiring expensive hardware upgrades
- Seamless integration with advanced functionality, including archiving to manage primary and e-mail storage growth as well as replication to help safeguard remote office data
- Ability to assist and accelerate upgrades of Windows, Exchange, and SharePoint by restoring items from previous versions to updated versions (for example, Exchange Server 2003 items can be restored directly to Exchange Server 2007)

- 64-bit-optimized CommVault solution designed to take full advantage of the PowerVault DL2000 and Windows Server 64-bit architectures

PROVIDING FAST, RELIABLE DATA PROTECTION

The combination of Dell hardware and CommVault software in the new Dell PowerVault DL2000 - Powered by CommVault can help organizations overcome the challenges of increasing data growth in their environments. By providing a simplified platform for backup and recovery and taking advantage of deduplication technology to help eliminate redundant data, this system is designed to increase storage efficiency, reduce backup and recovery windows, and help IT administrators meet SLAs while enhancing overall manageability. 

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

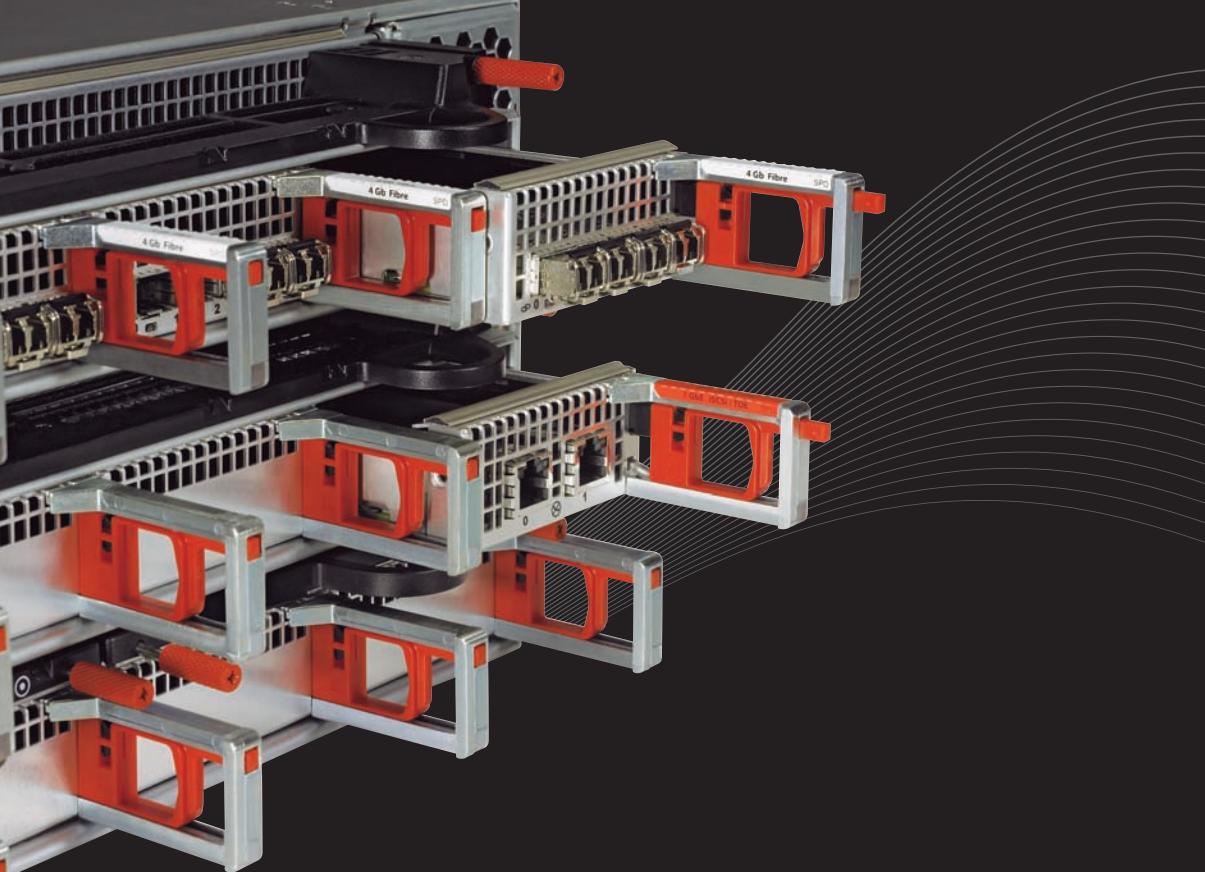
Jeff Echols is a director of business development at CommVault. He has over eight years of storage marketing and business development experience. Jeff has a B.S. in Mechanical Engineering and an M.B.A. from the University of Texas at Austin.

MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and CommVault:
DELL.COM/CommVault

Dell PowerVault DL2000:
DELL.COM/DL2000



FUTURE-READY STORAGE

**THE DELL/EMC CX4 ARRAYS
UNPRECEDENTED FLEXIBILITY IN A SAN**

SCALES UP TO 960 DRIVES

iSCSI AND FIBRE CHANNEL IN ONE ARRAY

TIER DATA ON UP TO FIVE DIFFERENT DRIVE TYPES

**SIMPLIFY STORAGE AT
DELL.COM/EMC**



EMC²



By Jeff Echols

STREAMLINING DATA MANAGEMENT WITH COMMVAULT SIMPANA AND DELL EQUALLOGIC PS5500E iSCSI SANs

Ongoing data growth can present major challenges for IT organizations. The combination of CommVault® Simpana® backup-to-disk software and Dell™ EqualLogic™ PS5500E Internet SCSI (iSCSI) storage area network (SAN) arrays can provide a simplified, scalable, cost-effective way to gain control over enterprise data and create a platform to support future growth.

Backup to tape used to be a simple, straightforward process—full backups on the weekend, followed by incremental backups throughout the week. Increasing data requirements combined with regulations such as the Sarbanes-Oxley Act and Federal Rules of Civil Procedure mean that many organizations must now retain huge amounts of information. However, many IT departments have found that managing this data growth with tape backup can be too cumbersome, too expensive, and too unreliable. And although disk-based appliances can help overcome some of these challenges, without an effective overall strategy, adding appliances can simply result in installing additional disks—disks that require increasing amounts of power and cooling to support while driving up management complexity and overall operational costs.

The combination of CommVault Simpana software and Dell EqualLogic PS5500E Internet SCSI (iSCSI) storage area network (SAN) arrays can help organizations meet their data management needs in a simplified, scalable, cost-effective way. Simpana backup-to-disk, archive, and de-duplication technologies provide a centralized way to help manage data growth, while EqualLogic PS5500E arrays are designed for high levels of efficiency and exceptional return on investment (ROI).

Related Categories:

CommVault

Data consolidation and management

Dell EqualLogic storage

Storage

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

SIMPLIFYING DATA MANAGEMENT WITH COMMVAULT SIMPANA SOFTWARE

CommVault Simpana software uses plug-in modules to manage data—an approach that enables administrators to perform backup and recovery, archive, and replication operations through a single integrated interface. Because the modules share the same backend engine, both the modules and storage devices can share one storage policy. For example, creating a storage policy to de-duplicate file data from virtualized servers to a Dell EqualLogic PS5500E array works for both the backup and archive modules—a useful feature when consolidating physical servers to virtual machines running on an EqualLogic SAN.

In many organizations, a majority of data stored on primary storage is infrequently used and no longer changes on a regular basis. Storing this data can slow application performance; clog backup, recovery, and archive jobs; and waste valuable primary storage space. To help overcome these challenges, organizations can deploy Simpana Archive software to clean out old or unwanted data, which can help significantly reduce primary storage space while helping increase application, backup, and archive performance.

CommVault next recommends transitioning from frequent full backup jobs to frequent incremental backup jobs. In addition to performing incremental

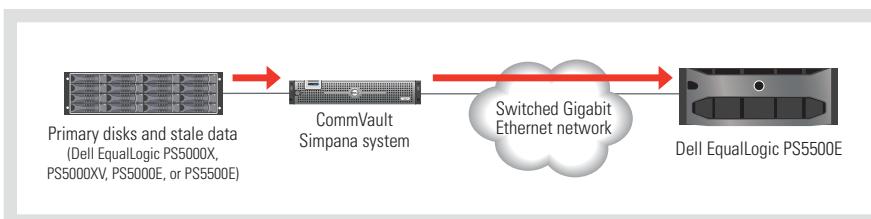


Figure 1. Simplified backups with CommVault Simpana software and Dell EqualLogic iSCSI SAN arrays

backups, Simpana software can use Smart Client agents that classify data to help eliminate file scans that may take hours to complete before an incremental backup job begins.

The final component is creating the full backup copy. Simpana software can create a *synthetic full* backup by using the previous full backup and adding the incremental backups. The full backup copy is created using a CommVault MediaAgent server, helping remove the load from the client servers.

REDUCING DATA FOOTPRINT WITH DE-DUPLICATION

Although archiving and incremental backups are not new, data de-duplication technology combined with CommVault backup and archive modules can greatly increase their value in enterprise environments while significantly increasing ROI.

Before the introduction of data de-duplication, archive operations could remove stale files and e-mail from a primary disk to a secondary disk. Although the secondary disk was typically less expensive and presumably slower in performance than the primary disk, the archived data required the same disk footprint as the primary data. Data de-duplication, however, can enable data archived to a secondary disk to achieve de-duplication ratios of as much as 15:1 or more, depending on the type of data—helping dramatically increase the effective capacity of existing storage resources.

The same holds true for backups. Prior to de-duplication, backing up to disk was primarily a caching mechanism to

expedite backups before sending them to tape. However, CommVault Simpana de-duplication software helps administrators keep many backup copies on the same amount of disk space, which can greatly accelerate data search and recovery.

DEPLOYING SCALABLE DELL EQUALLOGIC iSCSI SAN ARRAYS

Dell EqualLogic PS5500E iSCSI SAN arrays are designed to provide massively scalable storage that can allow organizations to simply and cost-effectively deploy high-capacity, highly consolidated SANs supporting primary data center, data protection and retention, archiving, and disaster recovery operations. Their scalability helps organizations begin with small- and medium-size storage requirements, and then grow to match their retention requirements over time—without performance degradation. Scalable performance enables EqualLogic arrays to support backup data copies as well as archive copies, which typically demand higher performance than backup copies because of usage patterns in which end users retrieve files and e-mail directly from secondary systems. Organizations can deploy EqualLogic PS5500E arrays in combination with CommVault Simpana software to help create a simplified, scalable, cost-effective backup-to-disk system (see Figure 1).

EqualLogic PS5500E arrays are also designed to address related issues in data center environments—including demands for improved energy efficiency, shrinking floor space availability, and lowered burden on administration teams. SANs built with EqualLogic PS5500E arrays can scale from

24 TB to 576 TB raw. All EqualLogic arrays can also maintain performance and scale capacity because of their virtualized architecture, and come with a comprehensive suite of advanced data management and protection software features at no additional cost. The form factor is designed for maximum space efficiency, provided as a 4U device that supports 48 drives in 500 GB or 1 TB sizes, while fully redundant hot-swappable controllers help ensure storage availability.¹

CONTROLLING DATA SPRAWL

The combination of CommVault Simpana software and Dell EqualLogic PS5500E iSCSI SAN arrays can help administrators back up and store data in a simplified, scalable, cost-effective way. By effectively managing backup and recovery, archive, and storage operations through a single interface and taking advantage of the flexibility and high performance of EqualLogic PS5500E arrays, organizations can gain control over data sprawl while helping create an efficient platform for future growth. 

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



MORE

ONLINE

DELL.COM/PowerSolutions

QUICK LINKS

Dell and CommVault:
DELL.COM/CommVault

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

¹For more information, see "High-Density, Highly Scalable Storage: Dell EqualLogic PS5500E iSCSI SANs," by Dylan Locsin and Travis Vigil, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20080396-Locsin.pdf.



By Sanjeet Singh
Charles Butler

SIMPLIFIED DATA PROTECTION WITH DISK-BASED BACKUP FROM DELL AND SYMANTEC

Combining high-performance Dell™ hardware with market-leading Symantec® Backup Exec™ software, the new Dell PowerVault™ DL2000 – Powered by Symantec Backup Exec can help organizations of all sizes deploy simplified, cost-effective data protection—helping accelerate backup and recovery, enhance media reliability, reduce total cost of ownership, and minimize the need for IT staff intervention and management.

As backup windows shrink and system availability requirements increase, IT organizations must find ways to protect and manage their ever-growing data with limited staff and hardware resources. Tape has been the traditional media of choice for backup and recovery, but it comes with a variety of disadvantages, including a lack of flexibility and time-consuming data recovery. Although cost-effective disk technology has been available for some time, only recently has it become practical to implement disk-based backup as part of an overall data protection strategy in enterprise IT environments.

To help meet the backup needs of organizations of all sizes, Dell has worked with two leading backup software companies, Symantec and CommVault, to introduce two new backup-to-disk appliances: the Dell PowerVault DL2000 – Powered by Symantec Backup Exec and the Dell PowerVault DL2000 – Powered by CommVault.¹ The new Dell PowerVault DL2000 – Powered by Symantec Backup Exec provides flexible, cost-effective disk technology to support simple, manageable data protection. Combining high-performance Dell hardware and market-leading Symantec backup and recovery software, this appliance helps accelerate backup and recovery, enhance media reliability, lower

total cost of ownership, and minimize the need for IT staff intervention and management.

TAPE- AND DISK-BASED BACKUP

Although tape has made significant improvements in throughput and capacity, its sequential-access nature typically makes it inflexible compared with disks. Because the raw throughput of disk is typically faster than tape, in most cases, backing up to and restoring from disk is faster than using tape—disk drives can begin transferring files instantly, whereas tape drives require that the tape be loaded, accessed, and sequentially written. Disk volumes, especially RAID volumes, can have very fast read performance, rivaling the throughput of even the newest tape drives. Disk snapshot technology, meanwhile, enables disk backups and restores to be virtually instantaneous. Importantly, these high levels of performance and efficiency enable administrators to schedule frequent backups, helping reduce the risk of data loss, while high levels of flexibility mean that, unlike tape, disks can support simultaneous backup, restore, and duplication operations.

Backing up multiple sources to a single tape drive requires a technology called multiplexing, or interleaving, which helps increase tape device efficiency

Related Categories:
Backup
Backup to disk
Continuous data protection (CDP)
Data consolidation and management
Dell PowerVault storage
Storage
Symantec

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

¹For more information on the Dell PowerVault DL2000 – Powered by CommVault, see “Fast, Reliable Data Protection from Dell and CommVault,” by Sanjeet Singh and Jeff Echols, in *Dell Power Solutions*, November 2008, DELL.COM/Downloads/Global/Power/ps4q08-20080393-CommVault.pdf.

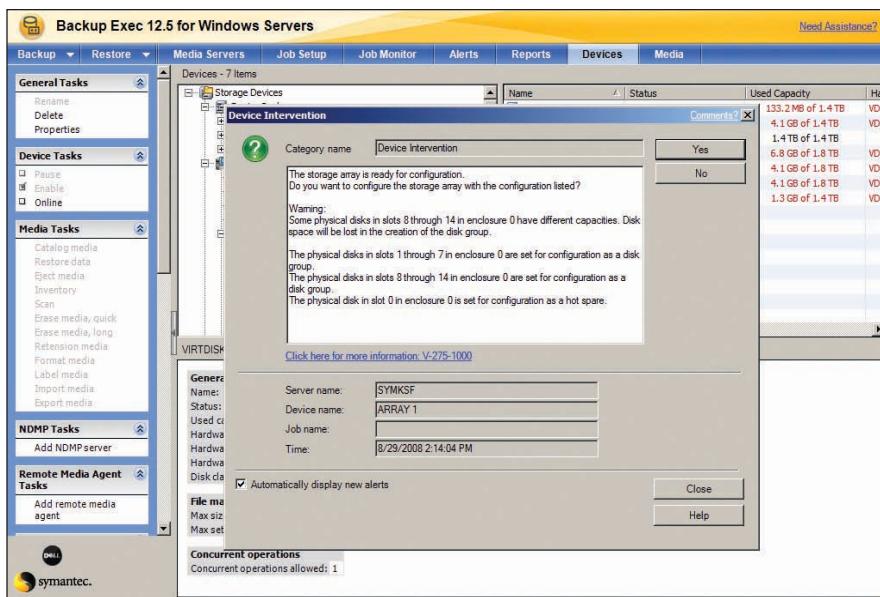


Figure 1. Symantec Backup Exec can automatically recognize new storage devices and help simplify setup

but also increases the time for restore operations. Using disks helps eliminate the need for multiplexing because disks are inherently random-access devices, designed to write multiple backup jobs simultaneously to individual backup files (one file per job)—helping provide both exceptional performance and storage granularity that even tape solutions with sophisticated multiplexing and multi-threading cannot provide.

For some types of data protection, disks are also typically more reliable than tape drives and tape libraries, which can cause delays or failures in backup and restore operations. Using disks as a high-frequency, short-term data protection medium—and tape as an archival medium—can help minimize these problems.

The speed and random-access capabilities of disk media can also provide the foundation for continuous data protection, helping protect data in real time as changes to data occur. Continuous data protection not only enhances the overall level of protection, but also helps reduce the administration and complexity associated with traditional data protection practices. For example, it can help eliminate the need for full, incremental, or differential backups currently in place, helping

protect data immediately and continuously by backing it up to disk.

Despite these advantages, however, successful deployment of disks as part of a data protection solution does present several challenges for IT administrators, including the following:

- **Lack of familiarity:** Backup administrators may be unfamiliar with configuring and managing disks, and must rely on server administrators to help select the appropriate disk technology and configure, set up, and diagnose problems.
- **Complex management:** Administrators typically manage tape devices through a backup application. Adding disk hardware to this environment may require introducing a separate application, increasing management complexity.
- **Disparate solutions:** Organizations typically purchase server hardware, backup hardware, and backup software from separate vendors. When a problem arises, it may be unclear which product is causing the problem, or which vendor backup administrators should contact to help resolve it.
- **Proprietary solutions:** In the past, virtual tape libraries have been deployed to take advantage of disk capabilities by

emulating tape devices. However, these devices may present a shortcoming—they must act like tape, which prevents backup applications from taking advantage of disk capabilities.

INTEGRATED DATA PROTECTION

To help organizations deploy simplified, cost-effective disk-based backup, Dell and Symantec have partnered to develop the Dell PowerVault DL2000 – Powered by Symantec Backup Exec. This appliance, based on the Dell PowerEdge™ 2950 server, the Dell PowerVault MD1000 disk expansion enclosure, the Microsoft® Windows Server® 2008 OS, and Symantec Backup Exec 12.5 for Windows Servers, can provide a variety of key benefits:

- **Integrated solution:** The appliance provides an integrated disk backup solution that also supports backing up to tape. Dell and Symantec have created and tested the components as part of a solution designed to be ready for deployment right out of the box.
- **Simplified setup:** The appliance comes factory installed with Symantec Backup Exec as well as automated wizards that guide administrators through the process of configuring the system name, login credentials, appliance IP address, and other items—helping them quickly get the system up and running.
- **Simplified disk management:** Integrated hardware and software management capabilities provide a simplified way to manage disk resources alongside backup and restore operations. Backup Exec can automatically discover, provision, manage, and enable disk resources as targets for backup and restore operations (see Figure 1).
- **Central management console:** A single management console serves as the launch point for appliance operations. The console is designed for simplicity, centralizing key tasks such as backup and restore functionality, configuration utilities, and management in a single location rather than burying them in

sub-folders—which can be especially important for managing data protection in remote offices with limited IT resources.

- **Disk usage analysis:** The appliance can use statistical information based on previous backup performance to predict disk space needs and help ensure the necessary space is available to complete backup operations. If additional space is needed, the system can notify administrators in advance and provide suggestions to help remedy the situation without affecting future backups.
- **Simplified alerting:** Administrators no longer need separate interfaces to evaluate component health—instead, a single console provides integrated alerting for backup destinations (see Figure 2). The system can send alerts when available disk space reaches each of three thresholds administrators can specify for a virtual disk. These alerts help prevent jobs from failing because of low disk space on a virtual disk.

COMPREHENSIVE PROTECTION WITH SYMANTEC BACKUP EXEC

The Dell PowerVault DL2000 – Powered by Symantec Backup Exec includes Symantec Backup Exec 12.5 for Windows Servers, the latest release of Symantec backup technology. This software provides comprehensive disk-to-disk-to-tape backup and recovery for Microsoft Windows® OS-based physical and virtual systems, including those running Windows Server 2008. Patent-pending Granular Recovery Technology (GRT) and the continuous data protection provided by Continuous Protection Server (CPS) can deliver reliable point-in-time recovery for critical Microsoft applications and enable the rapid recovery of individual Microsoft Exchange e-mails, Microsoft Office SharePoint® Server 2007 documents, and Microsoft Active Directory® user profiles.

Agent for VMware Virtual Infrastructure

Data protection software has typically treated virtualized environments as physical

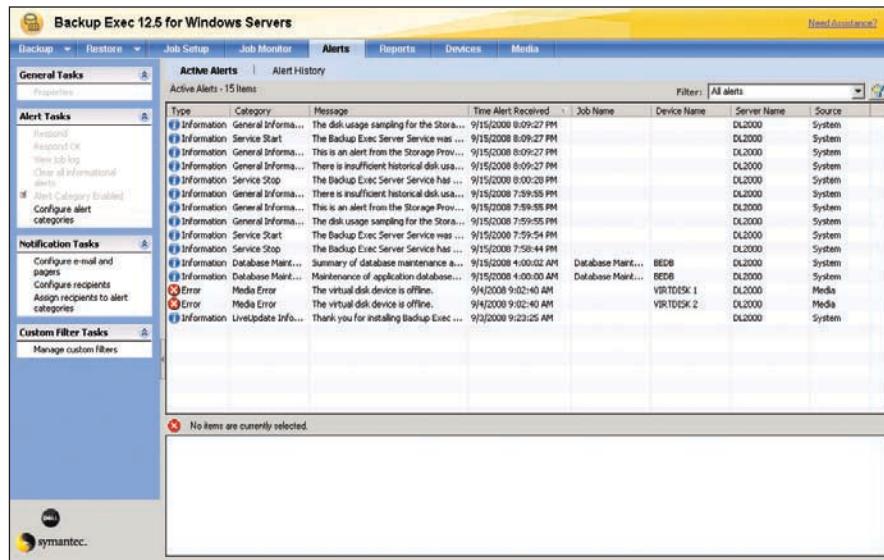


Figure 2. Integrated alerting helps administrators easily monitor backup and restore operations

environments, with agents often loaded onto each virtual machine (VM) and backups performed over the LAN. Although this paradigm was sufficient for a physical environment, it typically could not meet backup and recovery objectives for virtualized systems, often resulting in performance degradation and missed backup windows.

For environments based on VMware® virtualization software, Backup Exec can now utilize the Agent for VMware Virtual Infrastructure, which supports data protection for an unlimited number of VMs on an individual VMware ESX host. In addition, Backup Exec is designed to use numerous VMware technologies to enhance backup and recovery in virtualized environments. For example, it uses VMware VirtualCenter to automate discovery and presentation of VMware ESX servers and VMs as part of the backup environment. Administrators can also use VMware Consolidated Backup (VCB) to help minimize impact on VMs by taking advantage of the VCB framework to perform off-host VM backups—VCB scripting is no longer required, and an agent need not be loaded into the virtualized environment to perform backups. Finally, GRT enables the recovery of individual files and folders from VCB image backups.

For environments based on Microsoft virtualization software, Backup Exec can now also utilize the Agent for Microsoft Virtual Server, which supports data protection for an unlimited number of VMs on an individual host server. In addition, the agent supports both Microsoft Virtual Server and Microsoft Hyper-V™ technology. GRT enables the recovery of individual files and folders from full VM backups.

Granular Recovery Technology

Critical applications can present their own specific backup and recovery requirements. Recovery objectives, for example, often dictate the ability to restore both a full application and individual objects within each application. As a result, organizations might choose from a variety of methods for protecting applications—methods that often consist of backing up each application with a full database-level backup and then performing a separate backup of the individual items in these databases. Because two or more backups may be required to meet recovery objectives, this approach can double the amount of data and time required for backup operations.

Backup Exec provides GRT for key Microsoft applications, including Exchange,

Active Directory, SharePoint Server, and Windows SharePoint Services—providing administrators with a simplified approach to quickly recovering granular data from a single-pass backup:

- **Microsoft Exchange:** Exchange is not only a critical application for many organizations, but can also generate a huge amount of data that must be immediately available and protected. Even a single lost message can reduce productivity or disrupt operations. GRT enables the recovery of both databases and individual items such as public folders and mailboxes. Administrators can also recover granular objects from individual mailboxes such as specific e-mail messages, calendar items, folders, attachments, and notes.
- **Microsoft Active Directory:** As the standard foundation for organization and management in Windows-based environments of all sizes, Active Directory also requires comprehensive data protection and quick recovery. GRT enables the recovery of both databases and individual objects and attributes without an authoritative or non-authoritative full restore, and without requiring a reboot.
- **Microsoft Office SharePoint Server:** SharePoint Server is becoming a vital link for internal communications, and a lost or corrupted component could cause a major disruption. GRT enables the recovery of both databases and content databases, sites, subsites, lists, and individual documents.

Continuous Protection Server

CPS is designed to eliminate the need for full, incremental, or differential backups by protecting data immediately and then continuously backing it up to disk. This continuous data protection helps eliminate backup windows, accelerate backups, and enhance backup reliability, and features an innovative Web-based end-user file retrieval interface. Designed

specifically for disk, it integrates with Backup Exec for Windows Servers to deliver a comprehensive disk-to-disk-to-tape solution, helping increase data protection, reduce the administrative complexity associated with traditional data protection practices, and reduce the cost of the media used through efficient block-level data protection. End users can restore their own files without contacting IT departments, which can help improve service levels without increasing the number of IT staff or administrative costs.

GRT-enabled backups help protect Exchange at the level of both storage groups and mailbox stores while providing granular recovery of individual mailboxes, messages, and private and public folders from a single-pass backup. Many organizations run these traditional full or incremental backups of Exchange databases nightly using Backup Exec. However, as Exchange has become increasingly important to many organizations, the need for frequent Exchange data recovery beyond daily backups has also increased. The Backup Exec Continuous Protection of Exchange feature uses the same GRT-enabled technology for full database or granular recovery, but extends it by supporting frequent backups to help ensure quick recovery of recent data. Administrators can use the Backup Exec console to create GRT-enabled recovery points for Exchange at specified intervals to help meet their requirements.

Using this approach to continuous protection, administrators would typically perform a full backup each week or each month. Backup Exec can then provide continuous data protection for Exchange transaction logs, automatically consolidating them into easily managed recovery points to help ensure that the Exchange databases are protected up to the latest complete transaction log. When administrators enable recovery points to run at intervals between the weekly or monthly full backups, they can restore individual mailboxes, messages, and folders of

Exchange components, including embedded objects and attributes, to a time when the recovery point was created. This approach helps ensure that Exchange database and transaction logs are protected and can be recovered quickly in a disaster recovery situation, helping provide comprehensive protection for the Exchange environment.

SIMPLIFIED DATA PROTECTION FROM DELL AND SYMANTEC

The Dell PowerVault DL2000 – Powered by Symantec Backup Exec is designed to meet the backup and recovery needs of organizations of all sizes, including those running Microsoft applications and operating systems, Oracle® databases, Linux® and UNIX® operating systems, and VMware virtualization software. Integrating high-performance disk-based technology from Dell and market-leading backup and recovery software from Symantec, this appliance can help organizations deploy simplified, cost-effective data protection to help accelerate backup and recovery, enhance media reliability, lower total cost of ownership, and minimize the need for IT staff intervention and management. 

Sanjeet Singh is a senior global product marketing manager in the Dell Enterprise Storage Group.

Charles Butler is a technical director in the Data Protection Group at Symantec.

ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Symantec Backup Exec:
www.backupexec.com

Dell PowerVault DL2000:
DELL.COM/DL2000



By Dhiraj Sehgal

INTRODUCING BROADCOM iSCSI OFFLOAD ENGINE TECHNOLOGY FOR DELL SERVERS

Internet SCSI (iSCSI) is rapidly becoming a convergent data center technology for disparate types of networking. The Broadcom® iSCSI Offload Engine technology available in 9th- and 10th-generation Dell™ PowerEdge™ servers can free server resources to help increase performance and throughput, enabling organizations to maximize the benefits of iSCSI in their environments.

Increasing storage and networking demands have made Internet SCSI (iSCSI) technology a key component of enterprise IT infrastructures. Built on the familiar TCP/IP protocol, iSCSI helps increase storage flexibility by allowing access to content on a server through an Ethernet fabric. As a result, servers can either converge data applications and storage on the same network to help lower total cost of ownership, or dedicate one network for data applications and another for storage. In addition, by enabling an OS to be initialized across a storage area network and accessed over an IP network, multiple client systems can access the same available storage space over the network while allowing individual clients to access a different storage space over the same network.¹

By enabling organizations to use the same standard Ethernet equipment for multiple purposes—including high-speed networking, storage, clustering, and remote management—iSCSI can help lower power consumption, enhance performance, localize patch management, and conserve data center space. However, a standard server equipped with Layer 2 Ethernet controllers cannot efficiently run network, storage, and cluster traffic simultaneously over

Ethernet at the full line rate without consuming a significant amount of processing power. This heavy network traffic can then consume vital system resources required to process critical applications.

Broadcom converged network interface controllers (C-NICs) with iSCSI Offload Engine (iSOE) technology, available in 9th- and 10th-generation Dell PowerEdge servers, are designed to overcome this drawback. By handling traffic from disparate network functions in a unified Ethernet fabric and offloading iSCSI processing from host processors to LAN on Motherboards (LOMs), network interface cards (NICs), or host bus adapters (HBAs), these controllers can help increase both performance and throughput while helping optimize server processor utilization.

INTRODUCING BROADCOM ISOE TECHNOLOGY

Broadcom BCM5708C, BCM5708S, BCM5709C, and BCM5709S C-NICs with iSOE technology are featured in Dell PowerEdge server models 840, 860, 900, 1950, 2900, 2950, 2970, 6950, R200, R300, R805, R900, R905, SC1435, T105, T300, and T605. By offloading iSCSI protocol processing overhead from host processors to a BCM5708 or BCM5709 C-NIC,

¹For more information on the iSCSI protocol, visit www.rfc-archive.org/getrfc.php?rfc=3720. For more information on the advantages of iSCSI, see the three-part series “iSCSI: Changing the Economics of Storage”: “Part 1—Understanding iSCSI in Enterprise Environments,” by Travis Vigil, in *Dell Power Solutions*, May 2007, DELL.COM/Downloads/Global/Power/ps2q07-20070335-Vigil.pdf; “Part 2—Deploying iSCSI in Virtualized Data Centers,” by Matt Baker and Travis Vigil, in *Dell Power Solutions*, August 2007, DELL.COM/Downloads/Global/Power/ps3q07-20070401-Baker.pdf; and “Part 3—Using iSCSI in Small and Medium Businesses,” by Travis Vigil, in *Dell Power Solutions*, November 2007, DELL.COM/Downloads/Global/Power/ps4q07-20070402-Vigil.pdf.

LOMs and NICs with iSOE technology can help free processor cores and memory resources to help increase I/Os per second (IOPS) and reduce processor utilization at line rate for various I/O sizes. This offloading helps increase performance for file-oriented storage, block-oriented storage, backups, database transactions, and tightly coupled distributed applications such as high-performance computing.

The minimum system requirements for iSOE functionality are as follows:

- Broadcom BCM5708 or BCM5709 add-in iSCSI HBA or integrated LOM
- Microsoft® Windows Server® 2003, Microsoft Windows Server 2008, Red Hat® Enterprise Linux® 5, or Novell® SUSE® Linux Enterprise Server 10 OS, which can be factory-installed on supported Dell PowerEdge servers for an iSOE-enabled LOM or installed as part of an add-in iSCSI HBA kit
- iSOE hardware license key installed on LOMs before server boot (this license key is built into select add-in NICs and is an optional feature on Dell PowerEdge servers equipped with Broadcom BCM5708C-, BCM5708S-, BCM5709C-, and BCM5709S-based LOMs)
- Activation of embedded iSOE feature set (a Broadcom LOM is integrated into the base Dell PowerEdge server configuration at no additional cost, but organizations must elect to activate iSOE as a purchase option at the point of sale; licenses are not optional on all Broadcom controller-based NICs, but certain Broadcom controller-based NICs contain the license)
- Broadcom NetXtreme® II drivers for specific operating systems (pre-installed or provided for installation)

NICs operate in non-iSOE mode if any of the preceding components are missing or if the system is running an unsupported OS. After organizations have activated iSOE as purchase option, no additional

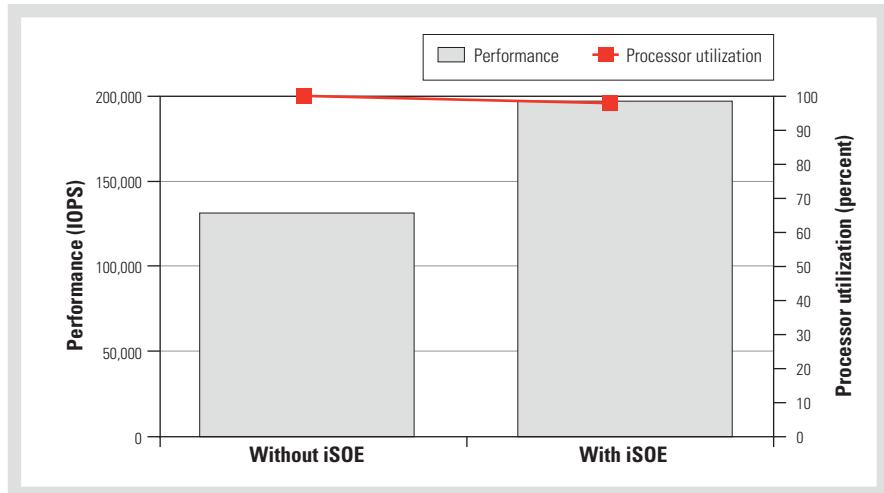


Figure 1. Performance and processor utilization with and without Broadcom iSOE technology

work is typically required to enable the hardware feature and the iSCSI software components; for iSOE-capable NICs, no additional hardware key is required.

Organizations that want to use iSOE functionality in their LOMs can purchase iSOE keys as an upgrade to the initial basic configuration. To enable iSOE, they must install the appropriate iSOE key in the TCP/IP Offload Engine (TOE) key connector on the server system board and update the LOM firmware and drivers to the latest versions.²

EVALUATING iSOE PERFORMANCE AND THROUGHPUT

Increased performance for key metrics—such as IOPS and processor utilization—was a key focus during iSOE development. The result is that for applications with small I/O sizes, such as Web servers and file database servers, Broadcom iSOE technology can provide higher bandwidth and lower processor utilization than a software solution such as the Microsoft iSCSI Software Initiator. For applications with I/O sizes greater than or equal to 4 KB, such as remote storage data backup, iSOE can provide tremendously increased bandwidth and correspondingly reduced processor utilization. In addition, iSOE technology is designed to address both I/O reads and writes, enhancing

performance for both types of operations as well as freeing processor cycles for other critical applications such as databases, media streaming, and file sharing.

To demonstrate the performance advantages of Broadcom C-NICs with iSOE technology, in October 2007 Broadcom engineers configured two Dell PowerEdge 2950 servers with dual-core Intel® Xeon® 5160 processors at 2.66 GHz, 8 GB of double data rate 2 (DDR2) RAM at 2.66 GHz, a frontside bus at 1,066 MHz, and Microsoft Windows Server 2003 with Service Pack 1. Each server had two onboard Broadcom BCM5708C C-NICs, with two software-based iSCSI targets connected through switches to each server; one server had iSOE enabled on the C-NICs, while the other had iSOE disabled and used the Microsoft iSCSI Software Initiator. The benchmark used in the tests included the Softpedia lometer 2006.07.27 network performance analyzer (1 worker per target and 32 outstanding I/Os).

Figure 1 shows the performance results for iSCSI read operations using a 512-byte I/O size along with processor utilization. In these tests, the server with iSOE technology handled 50.2 percent more IOPS than the server without iSOE technology, at comparable levels of processor utilization.

²Dell uses four part numbers for iSOE keys: YR232 for single-port TOE and iSOE technology, WY733 for dual-port TOE and iSOE technology, CR774 for quad-port TOE and iSOE technology in Dell PowerEdge R805 and PowerEdge R905 servers, and C402D for quad-port TOE and iSOE technology in other supported PowerEdge servers. For more information on installing iSOE software components, refer to the user's guide included with the network driver package.

Figures 2 and 3 show the throughput results for sequential iSCSI read and write operations, respectively, along with processor utilization. For sequential read operations, the iSOE-enabled server provided up to 28 percent higher throughput and up to 85 percent lower processor

utilization than the iSOE-disabled server using the Microsoft iSCSI Software Initiator; at large I/O sizes, in fact, the iSOE-enabled server had a processor utilization of less than 10 percent. For sequential write operations, the iSOE-enabled server provided up to 38 percent

higher throughput with up to 85 percent lower processor utilization than the iSOE-disabled server using the Microsoft iSCSI Software Initiator, with the iSOE-enabled server again having processor utilization of less than 10 percent at large I/O sizes.

These results demonstrate the performance and throughput increases possible when using Broadcom C-NICs with iSOE technology compared with a software-based iSCSI solution such as the Microsoft iSCSI Software Initiator. Organizations should keep in mind that real-world performance will vary based on the specific configuration running in a given environment.

OPTIMIZING iSCSI ENVIRONMENTS

iSCSI technology can offer significant benefits in enterprise IT environments, including reduced power consumption, increased performance, and reduced total cost of ownership. Broadcom C-NICs with iSOE technology are designed to enhance these benefits even further, enabling organizations to unify disparate network functions over Ethernet while increasing performance compared with software-based iSCSI initiators, minimizing the burden on host processors, and enabling the efficient use of enterprise resources. 

Figure 2. Throughput and processor utilization for sequential iSCSI read operations with and without Broadcom iSOE technology

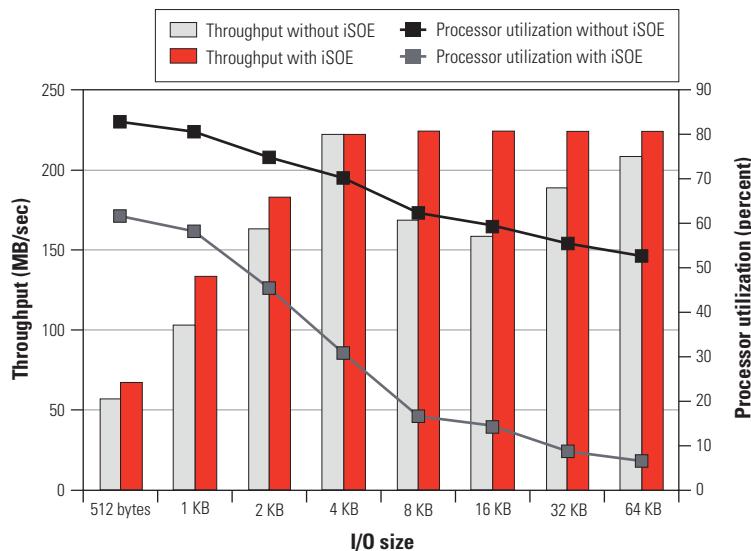


Figure 3. Throughput and processor utilization for sequential iSCSI write operations with and without Broadcom iSOE technology

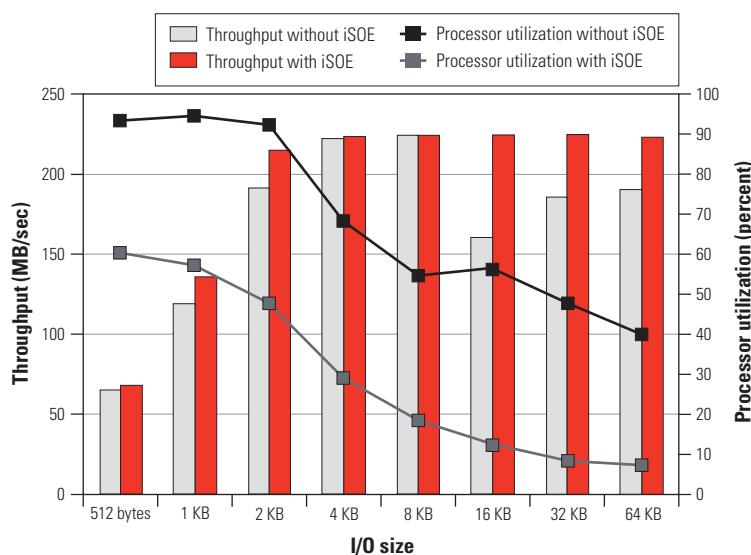


Figure 3. Throughput and processor utilization for sequential iSCSI write operations with and without Broadcom iSOE technology

Dhiraj Sehgal is a senior product line manager for Ethernet controllers at Broadcom. Dhiraj has an M.S.E.E. from North Carolina State University, Raleigh.



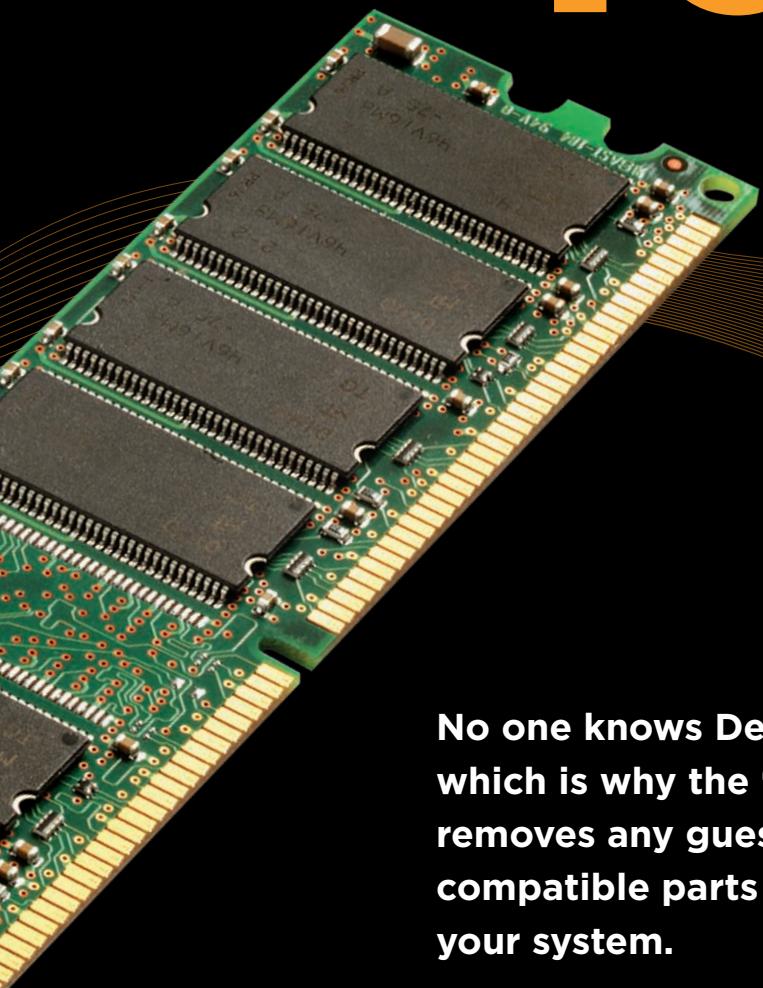
MORE
ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Broadcom controllers: www.broadcom.com/products/Enterprise-Networking/Gigabit-Ethernet-Controllers

Dell PowerEdge servers: DELL.COM/PowerEdge

LEAVE IT TO US



**No one knows Dell™ parts better than Dell,
which is why the “Parts for Dell” Website
removes any guesswork from finding
compatible parts and upgrades for
your system.**

**SIMPLIFY PARTS AND UPGRADES AT
DELL.COM/Upgrades OR CONTACT YOUR
DELL REPRESENTATIVE AT 1-800-357-3355**





Get the best results.



Scaling Software Architectures for the Future of Multi-Core Computing By Matt Gillespie

As the number of execution cores in multi-core processors continues to increase, software that isn't properly threaded will suffer performance penalties and become less competitive in the marketplace. Software publishers need to act now to position their products for success.

In February 2007, Intel demonstrated an 80-core processor. While a proof of concept, this chip demonstrated that massive parallelism opens the door to application possibilities that are only hinted at today. As more multi-core products are brought to market, software publishers will need to thread their applications to leverage the hardware's capabilities.

On a dual-core system, for example, unthreaded software can lead to significantly reduced utilization—with utilization decreasing even further for quad-core or eight-core systems. Similarly, as the number of processor cores increases, performance penalties associated with improper threading also tend to increase geometrically.

The key to accommodating multi-threading is to achieve a balance between near-term time-to-market and cost requirements on one hand, and long-term strategy on the other. For that reason, formally identifying and prioritizing highly threaded application designs both as a near-term and a long-term goal is an important step.

Characterize the Threading Process in Sophisticated but Simple Terms

In simple terms, software multi-threading is breaking down larger tasks addressed by a software application into subtasks that can be processed separately. Each of these tasks can then be assigned to a separate core. Since multiple subtasks are performed simultaneously, the overall result can be achieved faster, which is the main benefit of multi-threading.

Applying the correct threading methodology is critical, both in terms of dividing the task into the right subtasks and coordinating the threads that perform work on those subtasks. Considering the first of these

issues—dividing the task into the right subtasks—it helps to be familiar with two general ways of breaking tasks into smaller pieces: data decomposition and functional decomposition.

Data decomposition is based on the idea that performing the same work over and over on different pieces of data can be subdivided by giving pieces of the overall data set to separate threads for parallel execution.

Functional decomposition means identifying each of the discrete things that a piece of software needs to do at any given time, and assigning a thread to each of them. Workload balancing can be somewhat more complex in functional-decomposition problems than in data-decomposition ones, since it can be difficult to gauge how much work is required for each subtask.

Characterize the Difficulty of Parallelizing Specific Workloads

Another way of characterizing the threadability of workloads is by the amount of developer effort involved in creating the threaded version. Workloads fit into one of three broadly defined categories:

- Easily threaded workloads: This category is sometimes referred to as "embarrassingly parallel."
- Moderately difficult-to-thread workloads: This category includes some database applications, data mining, synthesizing, and text and voice processing.
- Very difficult-to-thread workloads: This category includes workloads that are very difficult to parallelize, due to linear arrangements where the input data of one subtask is generally dependent upon the output data from another.

Correlating the difficulty of threading tasks with the commercial value of doing so makes it possible to create a framework for identifying parallelization priorities.

Architect Software with Future Hardware Innovation in Mind

Application logic should be built to take advantage of the number of cores available to it. Such flexibility will be increasingly important as processor architectures become larger and more complex. The 80-core chip mentioned above uses a "tile" design that allows replication of many sets of identical structures on the silicon, potentially leading to the development of chips with an open-ended number of cores. Another promising innovation includes technologies that allow software to turn cores on and off as needed to save power. Robust and flexible software design can accommodate these trends, enabling today's software to run efficiently on tomorrow's hardware.

Additional Resources

- Intel® Multi-Core Developer Community includes a wide range of developer resources for creating software that takes advantage of multi-core processing.
- Intel® Multi-Core Technology and Research Portal provides access to a variety of resources about current multi-core technology at Intel.
- Intel® Software Development Products help to simplify the development of high-quality parallel software with tools that integrate with popular development environments.

Buy Intel Software Products at:
www.dell.com/intelsoftware

Learn more at:
www.intel.com/software/products



SIMPLIFY YOUR PATH TO VIRTUALIZATION

Dell can help make your transition to a Dell™ and Microsoft® virtual infrastructure smooth and efficient, helping to accelerate your deployment and empower your staff.

**RAPID CONSULTING ENGAGEMENTS BASED ON EXPERTISE GAINED
FROM THOUSANDS OF HOURS INVESTED IN IMPLEMENTATIONS**

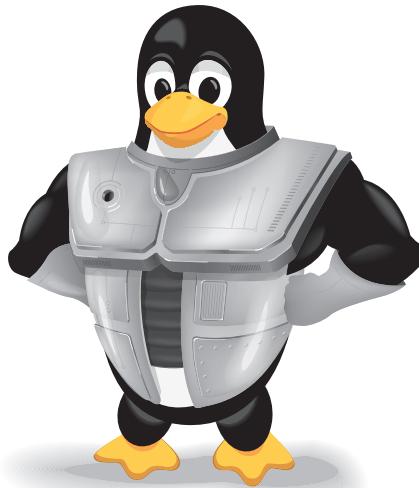
**VALIDATED, REPEATABLE MIGRATION PROCESS TO
DEPLOY POWERFUL SOLUTIONS**

**SAMPLE REFERENCE CONFIGURATIONS TO ACCELERATE
SOLUTION IMPLEMENTATION**

**SIMPLIFY YOUR OPERATIONS AT
DELL.COM/Hyper-V**



Enterprise Linux Support



From \$119
Per server/per year*

Unbreakable Linux

ORACLE®

oracle.com/goto/linux
or call 1.800.ORACLE.1

*Terms, conditions and restrictions apply. \$119 price solely includes access to Linux binaries and updates.
24x7 global support starts at \$499 per year/per server.