DELL POWER SOLUTIONS · SEPTEMBER 2009

POWER SOLUTIONS

DELL.COM/PowerSolutions

SEPTEMBER 2009 • US$18.95 CAN$19.95

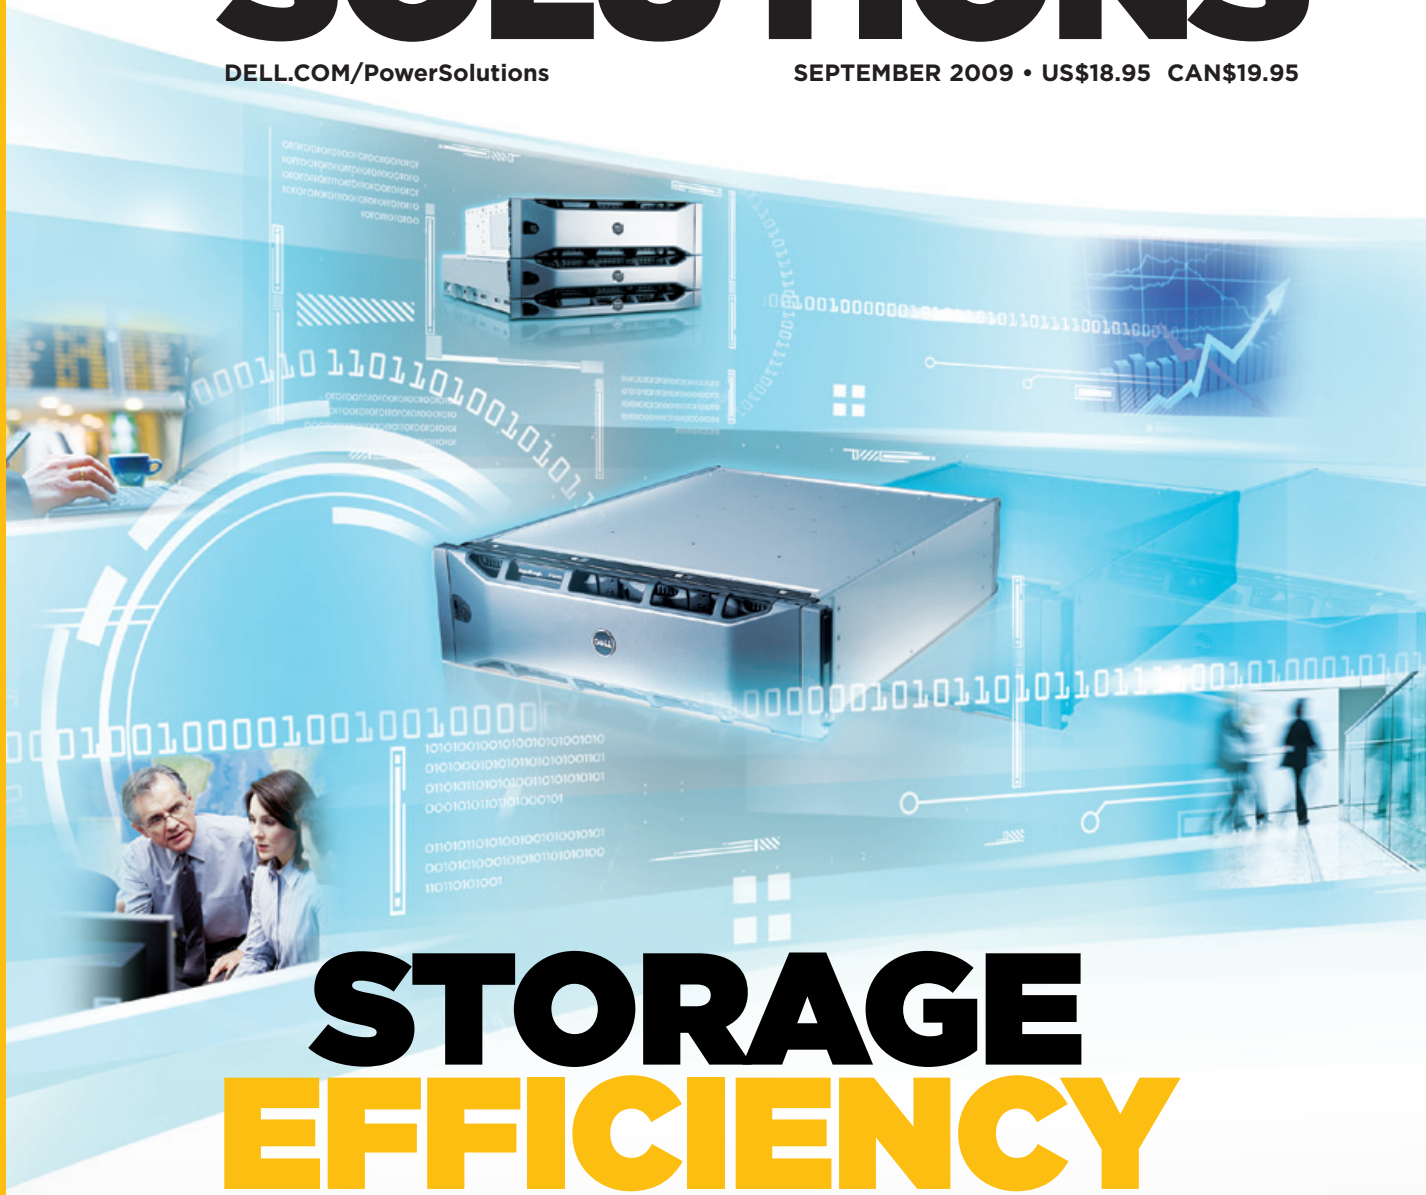

STORAGE EFFICIENCY A VIRTUAL SNAP

FEATURE SECTION

Game Changers: Darren Thomas and Paula Long on Enterprise Storage

Transitioning Dell EqualLogic iSCSI SANs to 10 Gigabit Ethernet

Allocating Storage with Virtual Provisioning on Dell/EMC CX4 Arrays

THE EFFICIENT ENTERPRISE

BUILDING INTERNAL CLOUDS WITH VMWARE vSPHERE 4

VIRTUALIZING DATA CENTERS WITH MICROSOFT WINDOWS SERVER 2008 R2 HYPER-V

STORAGE EFFICIENCY: A VIRTUAL SNAP

STK NO: DPS-200909-US

MAXIMIZE YOUR ENTERPRISE DATABASE INVESTMENT

"There are several features in SQL Server® 2008 that provide **a tremendous return on investment.** For example, database compression and backup compression.

We determined that we would save between $25K and $50K in the first year alone, and increasingly save money going forward, just because of the decrease in disk space needed."

–David P. Smith, Chief Technology Officer, ServiceU

MAXIMIZE ROI AT DELL.COM/SQL

SEPTEMBER 2009

POWER SOLUTIONS

DELL.COM/PowerSolutions

STORAGE EFFICIENCY: A VIRTUAL SNAP

An exclusive interview with Darren Thomas and Paula Long

Building on the strengths of a relationship with EMC since 2001, Dell took a significant stride forward with its US$1.4 billion acquisition of Internet SCSI (iSCSI) pioneer EqualLogic in January 2008—resulting in a combined storage portfolio spanning Fibre Channel, iSCSI, and network attached storage (NAS) technologies. The visionaries behind this game-changing strategy recently spoke with *Dell Power Solutions* about the state of enterprise storage technology today and why storage virtualization plays such a pivotal role in data center optimization.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

MIXING GIGABIT AND 10 GIGABIT ETHERNET IN A DEDICATED SAN INFRASTRUCTURE

By Tony Ansley

The recent availability of 10 Gigabit Ethernet (10GbE) and the expected enhanced Ethernet standards such as Data Center Bridging (DCB) present an opportunity for administrators to transition to a consolidated Ethernet infrastructure. The reference architecture illustrated in this article shows how enterprises may make the transition to 10GbE for storage using Dell™ EqualLogic™ PS Series arrays and Dell PowerConnect™ switches.

BREW YOUR OWN IDEASTORM

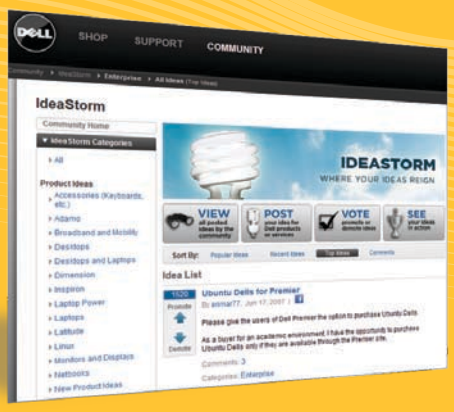

Do you have an innovative idea for improving a Dell product, service, or solution? Or want to see new offerings? Share your ideas with like-minded community members at IdeaStorm, a collaborative community where customers share, discuss, and vote on ideas. Dell monitors IdeaStorm to gauge which ideas are most important and most relevant to you and then works to implement them. Join in the discussion at **ideastorm.com**.

CONTENTS

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

The 11th-generation Dell PowerEdge server family is optimized for VMware vSphere environments. See page 54.

THE EFFICIENT ENTERPRISE

MORE ONLINE
DELL.COM/PowerSolutions

DELL POWER SOLUTIONS ONLINE EXCLUSIVES

Virtualization: Taking the Fast Track to Business Efficiency
IT belt-tightening is geared toward saving money—but you can surpass that goal by boosting business efficiency. Virtualization is the secret to this strategy.

Forecast: Cloudy with a Chance of Virtualization
Is cloud computing or in-house virtualization the answer to your scalability challenges? Maybe both.

Mission Control
Why client virtualization is the key to supporting a distributed workforce—without sacrificing security, IT control, or end-user customization.

Build or Lease?
Do the benefits of software as a service (SaaS) outweigh the risks? The short answer: maybe.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

NEXT-GENERATION IT

JOURNEY INTO THE CLOUD

Cloud computing promises to revolutionize enterprise IT and drive unprecedented levels of efficiency and flexibility. And there's nothing nebulous about the potential business benefits. To read about cloud computing alternatives in this *Dell Power Solutions* online exclusive, visit **DELL.COM/PowerSolutions**.

REDUCE YOUR CARBON FOOTPRINT WITH *DELL POWER SOLUTIONS* DIGITAL EDITION

You're just one click away from reducing your carbon footprint by subscribing to the *Dell Power Solutions* Magazine Digital Edition—featuring powerful search capabilities across archived issues. *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**.

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

DEPARTMENTS

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

CASE STUDIES

INDEX TO FEATURED DELL CUSTOMERS

INDEX TO ADVERTISERS

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

POWER SOLUTIONS
DELL.COM/PowerSolutions

EDITORIAL STAFF

EDITOR-IN-CHIEF AND PUBLISHER Tom Kolnowski
MANAGING EDITOR Debra McDonald
FEATURES EDITOR Kathryn White
ASSISTANT MANAGING EDITOR Jim Duncan
SENIOR EDITORS Dawn Davidson, James Hurd, and Terrence O'Donnell
EDITORIAL ASSISTANT Amy J. Parker
ART DIRECTOR AND COVER DESIGNER David Chan
DESIGNER AND ILLUSTRATOR Cynthia Webb
BUSINESS DEVELOPMENT MANAGER Cissy Walker

STAFF WRITERS Romy Bauer, Jeanne Feldkamp, Julie Jervis, Greg Thomas, and Chris Young

CONTRIBUTING WRITERS Anthony Ansley, Brady Black, Brian Brockway, Eric Cannell, Ron Carovano, Achmad Chadran, Gaurav Chawla, Wendy Chen, Annette Cormier, Jim Duncan, Andrew I. Fields, Franklin Flint, Andrew Gilman, Scott Hanson, Zahid Ilkal, Suresh Jasrasaria, Nick Jervis, Fred Johnson, Morten Loderup, Roger Lopez, Darren W. Miller, Mike Monthei, Ujjwal Rajbhandari, Hariharan Ramalingam, Jim Schwartz, Sanjeet Singh, Stanley L. Stevens, Wendy Torell, Bharath Vasudevan, Anand Vridhagiri, Greg White, and Rob Young

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

ADVERTISING SALES

SALES DIRECTOR Kyle Walkenhorst (323-340-8585)
WESTERN U.S. AND SOUTH/CENTRAL AMERICA SALES Shaun Mehr (949-923-1660)
EASTERN U.S. SALES Steve Branda (201-483-7871)
NORTHERN CALIFORNIA AND U.S. SALES Barbara Gough (650-595-0872)
CANADA, EMEA, AND APJ SALES Mark Makinney (805-709-4745)
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, ControlVault, Dell OpenManage, Dell Precision, EasyConnect, EqualLogic, Latitude, OptiPlex, On-Demand Desktop Streaming, 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 quarterly by Dell Inc., *Dell Power Solutions*, One Dell Way, Mail Stop RR5-05, 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 © 2009 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.

September 2009

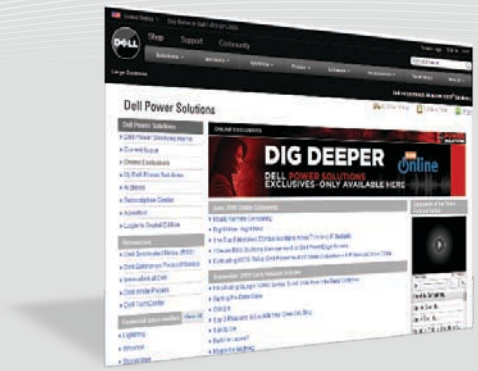

SEE IT HERE FIRST

DELL POWER SOLUTIONS ONLINE EXCLUSIVES




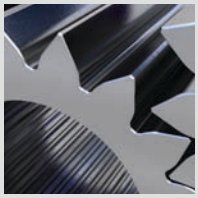


ENTERPRISE INSIGHTS

Top 3 Reasons to Launch Your Own CIO Blog

Creating a culture of innovation and inclusion starts with you.

Top 5 Reasons to Consider VoIP

Voice over IP (VoIP) outstrips standard enterprise telephone service in several key ways. Is it time to make your move?

Top 5 Enterprise 2.0 Mistakes

Avoid the missteps that can make your social media initiatives less effective.

Surfing the Data Wave

You can't fight the rising tide of data. But you can keep it at bay with an opportune strategy for information life cycle management.

CIO 2.0

Social media is changing the rules of engagement for enterprise communication—and how CIOs go about business as usual.

Talk to Me

Best practices for a unified messaging strategy help ensure security, compliance, and employee productivity without hindering the free flow of communication.

Ready for Anything

An executive survival guide for calculating potential losses—and the cost of recovery after a disaster.

People Power

How talent management software is reshaping the way organizations evaluate, train, and develop employees.

Visualize This

Live video technologies are transforming the way companies communicate with their customers and their employees. Get the picture?

Social Security

Top four issues to consider when shaping social media security policies.

Smart Talk

By pushing the envelope on IP-based infrastructure, the telecom industry is opening up a whole new world of communication-enabled business processes.

Wanted: IT Business Visionaries

It's no longer enough for IT administrators to know their way around the data center. The true measure of success is the acumen to view the company as a whole and wield technology to strategic business advantage.

New Rules

How personal technologies drive upstream change to alter the way businesses operate.

Changing the Locks

Securing business information used to be as simple as locking the door to the data center at night. In the age of mobile computing, it's not such a straightforward proposition anymore.

Building a Better Business

How the manufacturing industry is using advanced technology for strategic advantage—and why it may benefit large enterprises across a range of industries to tear a page from their IT playbook.

Customers in the Palm of Your Hand

No matter how diverse your customer base, they all likely use mobile phones—and that gives you the chance to extend the reach of your business by developing mobile applications that become part of their daily routines.

INDUSTRY SOLUTIONS

Dell Digital Cinema: The Sequel

By Franklin Flint

Advances in media technology have enabled the Dell OEM Industry Solutions Group to develop and extend efficient digital cinema architectures—helping simplify deployment, reduce power and cooling requirements, and lower both up-front and operational costs.

MORE ONLINE

DELL.COM/PowerSolutions

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

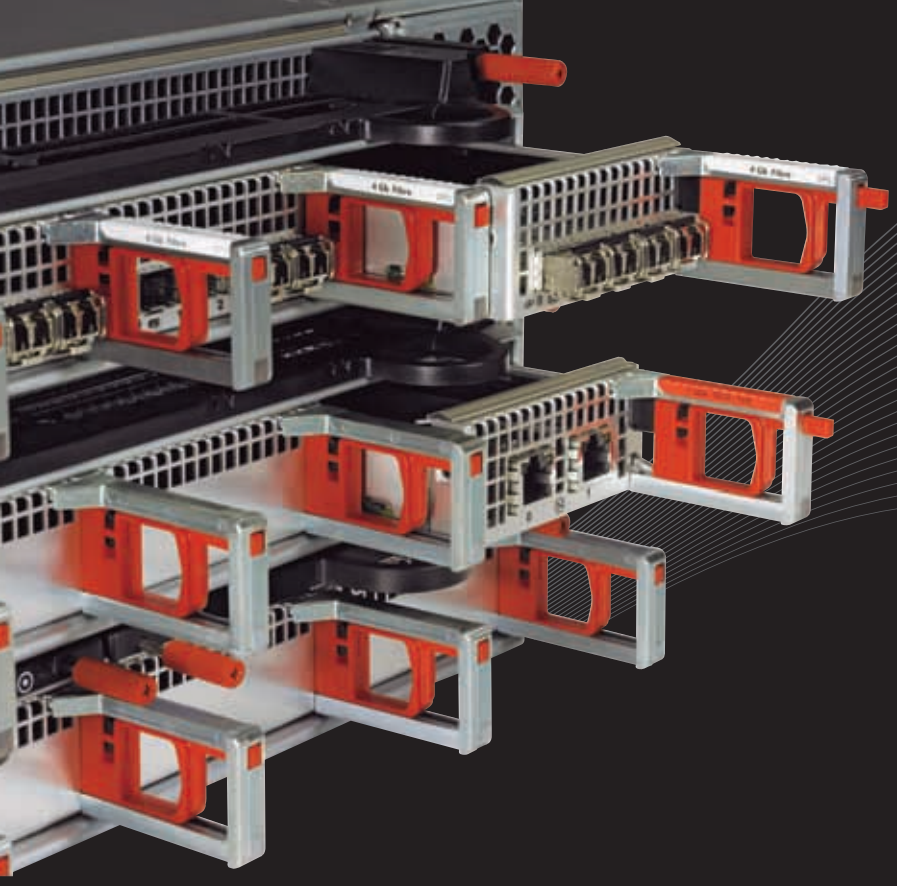

FUTURE-READY STORAGE

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

LEARN MORE ABOUT THE
DELL|EMC CX4 FAMILY
REGISTER FOR A LIVE WEBCAST AT
DELLENTERPRISE.COM/EMC

DELL EMC²

MASTERS PROGRAM, TUITION FREE

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

As an IT professional, you are justifiably proud of your credentials. Beyond the customary diploma from an educational institution for higher learning, your credentials may include one or more selections from a long list of acronymic IT certifications—the likes of Certified Information Systems Security Professional (CISSP), Cisco Certified Internetwork Expert (CCIE), IT Infrastructure Library (ITIL), Microsoft® Certified Systems Engineer (MCSE), or VMware® Certified Professional (VCP), to name a few. You know the certification drill all too well: enroll in the course, read the book, perform the exercises, pass the required exams (ideally on the first attempt), and append the credential to your résumé.

Now you may be eligible to add the Dell TechCenter Masters program to your credential list. As a vibrant online community for IT pros, the Dell TechCenter is the collaboration-savvy complement to the *Dell Power Solutions* Magazine publishing team. The Dell Masters program is all about using deep peer-to-peer networking—in tandem with your credentials, real-world experience, subject-matter-expert status, and proven track record of IT accomplishments—to help accelerate the community and its members to new levels of IT solution, innovation, and problem-solving excellence.

While anyone can join the TechCenter and interact on the site, Dell Masters create, contribute, and share content throughout the TechCenter collaboration space; participate in regular live chats with Dell experts; and work with the TechCenter on product betas, case studies, and leading-edge white papers—earning points and badges along the way. And while it will still be necessary to work hard to graduate to a level 4 badge status, the Dell Masters program is 100 percent tuition free. So don't just read the book, help write the book: submit your registration at www.delltechcenter.com/page/Dell+TechCenter+Masters. As an added membership perk, we plan to feature selected content originating from the program in future issues of *Dell Power Solutions*.

Not ready to raise the ante on your peer-to-peer networking game by joining the Dell Masters program and collaborating on content? Then settle in for reading an exceptionally strong editorial lineup in this September 2009 issue of *Dell Power Solutions*—our most expansive collection of content to date for a single issue. Within the 100 pages of the customary print edition, you'll find 22 fresh, in-depth articles from our large contingent of contributing authors and staff writers, while our More Online section features another 22 articles covering a much broader swath of IT management and technical interests—many of these quick reads of just a page or two. You can check out the list of More Online articles by turning to page 6, or click over to DELL.COM/PowerSolutions to read the very latest content additions.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SIMPLIFY YOUR STORAGE

TO IMPROVE OVERALL IT EFFICIENCY

Reduce your total cost of ownership while simplifying storage management. Virtualization is changing the way IT resources are deployed and managed, simplifying and speeding response to a changing business environment.

Learn how you can accelerate your ROI with Dell's virtual storage for the Efficient Enterprise by visiting www.dell.com/virtualization.

REGISTER FOR A FREE EQUALLOGIC PRODUCT DEMO AT
DELLENTERPRISE.COM/getmorevirtual

DELL

STORAGE EFFICIENCY: A VIRTUAL SNAP

Building on the strengths of a relationship with EMC since 2001, Dell took a significant stride forward with its US$1.4 billion acquisition of Internet SCSI (iSCSI) pioneer EqualLogic in January 2008—resulting in a combined storage portfolio spanning Fibre Channel, iSCSI, and network attached storage (NAS) technologies. The visionaries behind this game-changing strategy recently spoke with *Dell Power Solutions* about the state of enterprise storage technology today and why storage virtualization plays such a pivotal role in data center optimization.

As vice president and general manager of the Dell Enterprise Storage Group, Darren Thomas oversees storage strategy and product development across the Dell™ EqualLogic™, Dell PowerVault™, and Dell/EMC storage families. Paula Long is the vice president of engineering for the Dell Enterprise Storage Group. As a founder and vice president of products and strategy at EqualLogic, Long was a driving force in the company's growth to more than US$100 million in revenue per year. Together, Thomas and Long share their insights into the many innovations that are helping to improve enterprise efficiency, simplify management, and lower operational costs.

How has the EqualLogic acquisition advanced adoption of iSCSI storage?

"Customers have always been amazed by EqualLogic technology—but the acquisition took EqualLogic to a worldwide scale," says Thomas. "Dell brought EqualLogic products to higher-end customers who probably would have been reluctant to buy from a start-up company. But they trusted the Dell brand. So together, Dell and EqualLogic made Internet SCSI (iSCSI) a household word."

Photos of Paula Long by Bryce Vickmark

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

"EqualLogic does a quarterly independent survey of 200–300 customers, and our customer satisfaction ratings have stayed the same since the acquisition," Long adds. "We are also hearing from customers that integration with key Dell partners allows for easier deployments within their IT infrastructure. These partnerships include Citrix, Symantec, CommVault, Microsoft, and VMware."

How does the Dell/EMC relationship factor into Dell's enterprise storage strategy?

"The relationship between Dell and EMC began in 2001. Since then, Dell and EMC have shipped more than 60,000 storage area networks (SANs) to 24,000 individual customers. Technology executives love that this relationship gives them access to exceptional Fibre Channel SAN technology along with the full suite of Dell products, including servers, storage, desktops and laptops, and services," says Thomas. "Plus, the EqualLogic acquisition gives enterprises access to a world-class solution for next-generation SANs. And if you combine that with the Dell PowerVault product set—which is an extraordinary value because it can be managed alongside our server products—you begin to see the comprehensive reach of the Dell portfolio. Many of our customers have a variety of storage and application needs, so access to a full range of storage options is important to them.

"The EMC relationship keeps the cadence of innovation high," continues Thomas. "EMC was among the first to create replication and snapshot technologies, and now those technologies are standard within the industry. New EMC® product sets typically come out every 18–24 months—those are complete revamps, not just improvements. And every product is designed to be exceptionally flexible. For example, the latest Dell/EMC CX4 Series storage systems can scale from 5 to 960 drives. EMC also designs for future upgrades. They anticipated major changes on the interconnect side when designing the CX4 systems. They support 4 Gbps and 8 Gbps Fibre Channel today, with 10 Gigabit Ethernet (10GbE) following in September. And all the systems in the CX3 installed base are upgradable, which helps to provide investment protection and prevent disruption in the data center.

"Despite iSCSI's phenomenal growth, it's important to remember that Fibre Channel still represents a large share of our sales. We're finding that if customers have deployed Fibre Channel worldwide in the large enterprise, they're typically slow to transition away from that technology. Customers tell us they want access to the best Fibre Channel SAN solution in the industry, and that's what we can offer them through the relationship with EMC."

What business trends are changing the way enterprises approach storage technologies?

"Green initiatives and cost savings are huge," says Thomas. "Virtualization plays a key role in helping achieve both goals because it allows enterprises to save power and energy. Companies are dramatically lowering their power consumption by migrating to virtual servers and consolidated storage—because when you

Photos of Darren Thomas by Bryan Kuntz

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

virtualize servers, you need to aggregate the storage to go with them to get the full benefit of virtualization.

"Simplicity is another major trend. Complexity is expensive. Customers are asking how they can do what they need to do with less money, less equipment, and fewer resources managing that equipment.

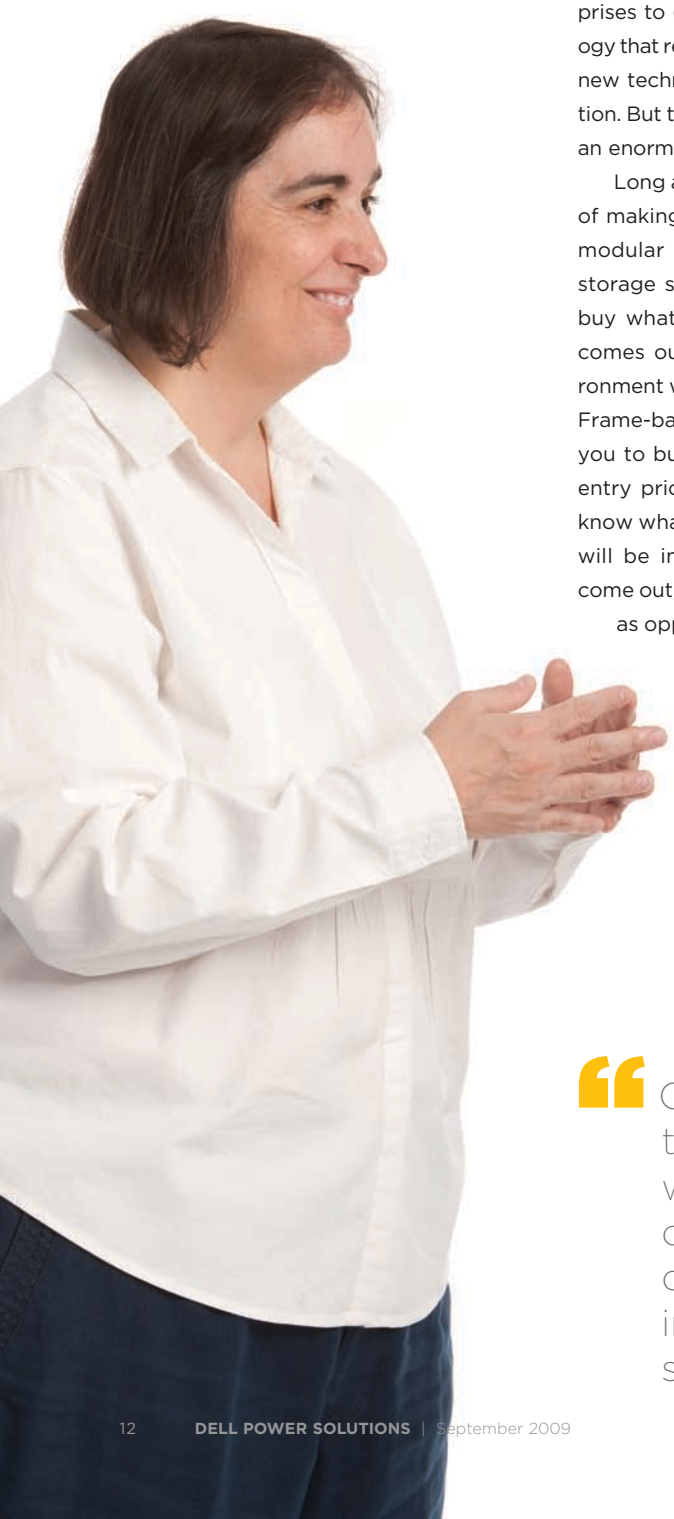

"The third trend we're seeing is automation. If you're going to simplify your IT infrastructure, you have to automate things like tiering and configuration and provisioning. The level of automation we're offering in our EqualLogic products is game-changing in terms of what organizations can do to simplify their IT.

"The current economic landscape is creating a perfect storm, driving enterprises to go out and look at new technology that reduces costs. In the past, adopting new technology could be a painful transition. But the right technology can save you an enormous amount of money."

Long adds, "Modularity is a critical part of making those savings possible. With a modular storage system, or a scale-out storage system like EqualLogic, you just buy what you need. As new technology comes out, you can augment your environment without having to retire anything. Frame-based technologies generally force you to buy more than you need, and the entry price is much higher. You need to know what your technology requirements will be in the future. So as new things come out, you may have to rip and replace as opposed to just augmenting as you would in a scale-out model."

"EqualLogic takes a far-sighted approach to help enterprises predict future requirements and costs," Thomas notes. "We include EqualLogic storage software in the price of the system. So if an organization buys storage from us today and we launch a new piece of software two years from now, the new software will be available at no additional cost to customers who have an active service agreement for the EqualLogic system. This policy helps IT managers easily calculate costs down the road because they won't have any unexpected, unintended license and support fees to pay. And those fees can be huge."

What benefits do storage innovations such as 10GbE and FCoE offer?

"I think the industry has declared that 10GbE is the wire in the future, which simplifies things for IT departments," says Long. "So now it's just about protocol. The advent of 10GbE gives enterprises a greater ability to virtualize servers and connect them to storage because they're not as bandwidth-limited as in the past. The fatter pipe makes data mobility and replication a lot more efficient. Of course, slapping a 10GbE interface on an existing storage controller may not buy you a lot unless you've got something to use with that bandwidth. Ethernet is important, but it's really just a wire."

Thomas agrees. "More than anything," he says, "Fibre Channel over Ethernet (FCoE) will help enterprises who have made huge investments in Fibre Channel to take a bridge step toward Ethernet. We're very excited about this move, since Ethernet is generally more cost-effective."

> "Our goal was to develop management tools that a common mortal could use—someone who is an IT generalist. We have taken some complex, advanced technology features like disaster tolerance and replication and tied them in with our all-inclusive suite of EqualLogic storage software."
>
> —Paula Long

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

How should data centers go about unifying the network fabric for iSCSI and Fibre Channel storage?

"The idea of a unified network fabric is all about IT simplification," says Thomas. "Putting all your networking and all your storage on the same type of network infrastructure makes technology acquisition easier because you no longer have to buy disparate networks—you buy just one kind. Your IT staff needs to develop expertise on only one type of infrastructure and tools. And you need to update only one type of product for both iSCSI and Fibre Channel systems. When IT managers start thinking in terms of simplicity and realize that iSCSI can match Fibre Channel's levels of security and performance and capability, the answer becomes very clear."

How will further technical innovations change the storage infrastructure?

"Virtualization vendors and storage vendors are working together to optimize the environment so each piece of equipment can focus on doing what it's best at," says Long. "These efforts are helping to reduce complexity and cut down on round-trips between the servers and the storage, and delegating a lot of the complex data management back down to the intelligent storage array. Basically, an optimized environment helps free up resources on the server and the network. This allows a bond between the virtualization and the storage that hasn't existed before. Optimized systems can make a copy in a fraction of the time when you delegate the work to the devices that are best suited to do it."

Thomas offers an example. "Say an engineering company with 500 people wants to give each employee a certain amount of capacity so they can run simulations. Every person's block has exactly the same characteristics: so much capacity, so much associated storage. Before, you'd have to copy that architecture 500 times so each client gets equivalent space. Some engineer would have to sit at his computer for days to do it. But with the latest generation of storage-aware virtualization tools, you just send the array controller the message that you want to perform that action 500 times. The EqualLogic storage software will go and make those copies almost instantaneously. We can create the copy virtually, in seconds, then move the data later."

Can organizations simplify management of their storage infrastructure?

"Yes. The key trends of consolidation, virtualization, advanced features, and automation dramatically simplify storage operations," says Long. "And because we virtualize everything in the EqualLogic products, we don't actually manage it like any other storage. It needs its own console."

Thomas adds, "To get the best-in-class solution that lets you do deployment and provisioning and storage upgrades with maximum efficiency, you have to use a tool that is specific to the storage infrastructure you have. And because these devices have such enormous capability, you have to use the element management tool that came with that device.

> "The current economic landscape is creating a perfect storm, driving enterprises to go out and look at new technology that reduces costs. The right technology can save you an enormous amount of money."
>
> —Darren Thomas

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

"In an ideal world, one management tool for all storage might be something to look forward to—but storage is still undergoing rapid innovation. With this rapid innovation, it is very difficult for a single tool to encompass and understand and operate the variety of devices in the market. If you had a management tool that was released even last year, it wouldn't understand what deduplication or thin provisioning means. It wouldn't understand space-efficient clones. Monitoring tools that are designed with specific storage arrays in mind are much more granular, much more specific, and as a result probably much more useful to the IT administrator."

And what makes EqualLogic arrays so easy to configure? "Our goal was to develop management tools that a common mortal could use—someone who is an IT generalist," says Long. "We have taken some complex, advanced technology features like disaster tolerance and replication and tied them in with our all-inclusive suite of EqualLogic storage software. These systems integrate with Microsoft® Exchange, Microsoft SQL Server®, VMware®, and Citrix® XenCenter™ software, and they're designed to do rapid recovery of Microsoft Hyper-V™ virtual machines in just a few seconds. The EqualLogic PS arrays are also VMware vStorage–ready so administrators can use storage, networking, and hypervisor resources for virtual machines more efficiently."

What is the outlook for virtual-aware storage solutions across all Dell storage offerings?

"EqualLogic systems were designed specifically for a virtualized environment," notes Long. "Normally, in storage, there's a tie between the data and the disk it's stored on. We've broken that connection. Everything can move anywhere within this logical grouping that makes sense. So as you move your server applications, your storage can move with them."

"Servers come and go," agrees Thomas. "But if a server goes and comes back, you want your storage to remember all of its associations, all of its snapshots. Our EqualLogic systems will re-associate everything with that server—a very important capability that didn't exist a short time ago. Being virtual-aware is a very specific set of skills."

"We're seeing a lot of technology opportunity," Long adds. "You can already combine 10,000 rpm and 15,000 rpm Serial Attached SCSI (SAS) drives, and Serial ATA (SATA) drives. And we already offer three tiers—you can pick your performance, pick your speed, pick your cost. And now you can go to solid-state drives (SSDs) for split-second storage. This performance allows new IT paradigms to emerge, such as dynamically creating virtual machines during spikes in demand to increase computing horsepower."

"Virtualization gives enterprises the ability to be flexible," Thomas concurs. "In the case of EqualLogic, we're pretty flexible with the hardware. So that's the magic: we're going to make systems that can react. If you think about the way your body works, you get hot, and you sweat a little bit, you evaporate, and you cool back off. You self-heal. Well, storage is heading in that direction too. It's getting to the point where our system

> "You want to make sure that your storage infrastructure can adapt at the same rate as your servers. If you scale by adding more servers and you're not able to add more storage bandwidth easily, your applications are just going to wait on your storage faster."
>
> —Paula Long

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

will react to what is happening in your data center.

"Our Dell/EMC offerings are becoming more virtual-aware as well. We expect those platforms to support fully automated tiering by the end of year, which offers the ability to match data with the right tier based on cost or access speed requirements. The CX4 family also has the highest attach rate to VMware environments. The EMC engineering team has spent a lot of time on addressing virtualized environments. You see evidence of this in the new virtual-aware EMC Navisphere® tools, which do storage management for the Dell/EMC line. These tools are some of the first to allow IT managers to perform server and storage management together."

Virtualized environments tend to accumulate a lot of redundant data. How can enterprises optimize the storage infrastructure?

"The question is really about deduplication technology, which is actually multiple technologies," says Thomas. "There's compression, there's single-instancing, there's deduplication within blocks and within files. The technology isn't very far along yet—and as a result, a lot of it has been focused on appliances for backup operations. It's either a virtual tape library or a disk system that is behaving like at the end of the tape—a stream of data.

"The goal should be to perform deduplication closer to the source—the application and user. That way, when enterprises move data, they move the lesser quantity—the data that's already been deduplicated. For example, say somebody creates a file. They send it to you, they send it to me, they send it to Paula, they send it to everybody else. Each copy is stored on the company's most expensive storage system. You save the most money if you can deduplicate it at that level. And whenever you finally do get ready to move it, either to replicate it or to store it in an archive, it's already deduplicated."

"EqualLogic works with our partners to provide deduplication for backup," notes Long.

What advice do you have for enterprises looking to optimize their data centers?

"Virtualization gives you the ability to consolidate a bunch of applications on a smaller footprint," Long advises. "But it also gives you the mobility to move things seamlessly as you need more computing power, or as you need more bandwidth. So you want to make sure that your storage infrastructure can adapt at the same rate as your servers. If you scale by adding more servers and you're not able to add more storage bandwidth easily, your applications are just going to wait on your storage faster. At the end of the day, balance is key for efficient IT operations, and storage needs to be flexible to maintain cost-effective performance.

"You want to have something that scales not only capacity, but also performance at the same time," says Long. "Otherwise, you're not going to get the full benefit of virtualization. You're not going to get the full benefit of being able to grow seamlessly in your server farm without growing seamlessly in your storage farm."

> "The idea of a unified network fabric is all about IT simplification. Putting all your networking and all your storage on the same type of network infrastructure makes technology acquisition easier because you no longer have to buy disparate networks."
>
> —Darren Thomas

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

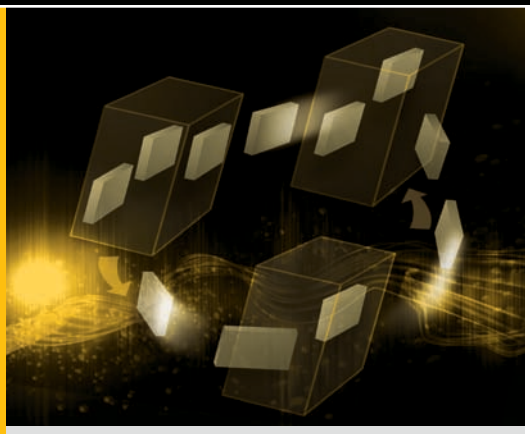

ALLOCATING STORAGE WITH VIRTUAL PROVISIONING ON DELL/EMC CX4 SERIES STORAGE ARRAYS

By Greg White
Annette Cormier
Eric Cannell

Virtual provisioning on Dell/EMC CX4 Series storage arrays is designed to dramatically increase storage utilization, simplify storage management, and minimize power and space requirements. By taking advantage of this flexible, automated approach to storage allocation, administrators can help reduce total cost of ownership and enhance the energy efficiency of the data center.

As enterprise storage requirements continue to grow, storage administrators often struggle with the challenge of meeting application capacity, performance, and availability requirements while controlling costs. Traditionally, administrators have met performance and availability requirements by allocating enough physical storage capacity to handle anticipated growth. However, this pre-allocation approach generally leads to overprovisioning storage capacity, which can in turn incur unnecessary capital costs and increased energy, space, and management costs—not to mention significant underutilization of that capacity.

Even when organizations overprovision storage capacity as a preventive measure, they still sometimes need to add capacity, leading to increased management and capital expenditures. On the other hand, allocating space on an as-needed basis can be complex, time-consuming, and difficult to manage, and can result in significant application disruption and downtime.

To help increase storage utilization, simplify management, and avoid the excess costs associated with overprovisioning storage capacity, Dell/EMC CX4 Series storage area network (SAN) arrays offer *virtual provisioning*, a thin provisioning–based storage virtualization technology that enables storage capacity to be allocated simply and automatically on an as-needed basis. One of several features available in Dell/EMC CX4 Series arrays (see the "Flexible, scalable Dell/EMC CX4 Series storage arrays" sidebar in this article), virtual provisioning provides applications access to a virtualized pool of storage capacity that is physically allocated on demand based on application usage, with no application disruption or downtime and no manual intervention required. By helping avoid the need to pre-allocate unused storage and automating key storage provisioning and management tasks, virtual provisioning can dramatically enhance utilization rates and simplify storage management—helping organizations to reduce the total cost of ownership of their storage infrastructures and enhance the energy efficiency of their data centers.

ENHANCING UTILIZATION WITH THIN LUNs

Traditional storage deployments are often overprovisioned, which in addition to cost inefficiency can lead to increased energy and space requirements in the data center. In this approach, storage is partitioned into logical units (LUNs) such that the amount of physical space allocated is the same as the capacity seen by the host servers (see Figure 1).

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

FLEXIBLE, SCALABLE DELL/EMC CX4 SERIES STORAGE ARRAYS

Dell/EMC CX4 Series storage arrays incorporate a variety of features designed to combine outstanding scalability and performance with flexibility, efficiency, and ease of use to enable organizations to cost-effectively meet data center storage needs:

- **Dell/EMC UltraFlex™ architecture:** The UltraFlex architecture enables administrators to customize the network interconnects and number of ports on each array.
- **Enterprise flash drives (EFDs):** Hard drives with solid-state flash memory are designed to deliver extremely high performance for latency-sensitive applications.
- **Data-in-place migration:** Support for data-in-place migrations helps simplify upgrades to Dell/EMC CX4 Series models with additional capacity and upgrades from previous-generation Dell/EMC CX Series arrays.
- **Virtual provisioning:** Storage virtualization technology enables storage capacity to be allocated simply and automatically on an as-needed basis as opposed to all at once.
- **Replication:** Dell/EMC CX4 Series arrays support applications that enhance data protection and migration using snapshots, clones, and remote mirroring.
- **Energy efficiency:** Available low-power Serial ATA (SATA) drives consume up to 32 percent less energy than standard 7,200 rpm drives.*

Dell/EMC CX4 Series models provide a range of capacities to help meet different needs—including supporting from 5 to 960 drives for a total storage capacity of up to 954 TB as well as connections to up to 4,096 highly available storage area network (SAN)–attached host servers (see Figure A). For optimal performance, Dell/EMC CX4 Series arrays support 4 Gbps or 8 Gbps Fibre Channel as well as Internet SCSI (iSCSI) using Gigabit Ethernet or 10 Gigabit Ethernet network technologies. Designed to reduce operational and energy costs, simplify management, and enhance utilization while delivering outstanding storage performance and scalability, Dell/EMC CX4 Series arrays can help organizations create scalable, energy-efficient, and cost-effective data centers.

	Dell/EMC CX4-120	Dell/EMC CX4-240	Dell/EMC CX4-480	Dell/EMC CX4-960
Maximum cache size	6 GB	8 GB	16 GB	32 GB
Maximum number of drives per array	120	240	480	960
Maximum capacity per array	120 TB (with SATA drives) or 72 TB (with Fibre Channel drives)	234 TB (with SATA drives) or 144 TB (with Fibre Channel drives)	474 TB (with SATA drives) or 288 TB (with Fibre Channel drives)	954 TB (with SATA drives) or 576 TB (with Fibre Channel drives)
Maximum number of highly available SAN-attached host servers	256	512	1,024	4,096
Thin pools per system	20	40	40	60
Disks per thin pool	40	80	120	240
Disks in all thin pools	80	160	240	360
Thin LUNs per pool	512	1,024	2,048	2,048
Thin LUNs per system	512	1,024	2,048	2,048

Figure A. *Dell/EMC CX4 Series storage arrays offer a range of capacities to help meet different needs*

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

Although this type of provisioning can often meet application capacity, performance, and availability requirements, it can also lead to significant underutilization of storage resources.

Virtual provisioning partitions storage into *thin LUNs*, in which the physical space allocated for storage may be less than the capacity seen by the host server (see Figure 2). Thin LUNs are aggregated into shared storage pools called *thin pools*. Physical storage is assigned to applications in a capacity-on-demand fashion from shared thin pools. When additional capacity is needed, it is first added to the shared pool and then allocated to individual thin LUNs as required. Capacity can be added to or deleted from thin pools without application disruption or downtime.

Thin pools can consist of any supported drive type—Serial ATA (SATA) drive, Fibre Channel drive, or enterprise flash drive (EFD)—and can support up to 240 drives per pool and a total of 360 drives overall. The number of thin pools and the number of disks per pool vary by array (see Figure A in the sidebar in this article).

Because it avoids the need to preallocate storage in anticipation of future

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

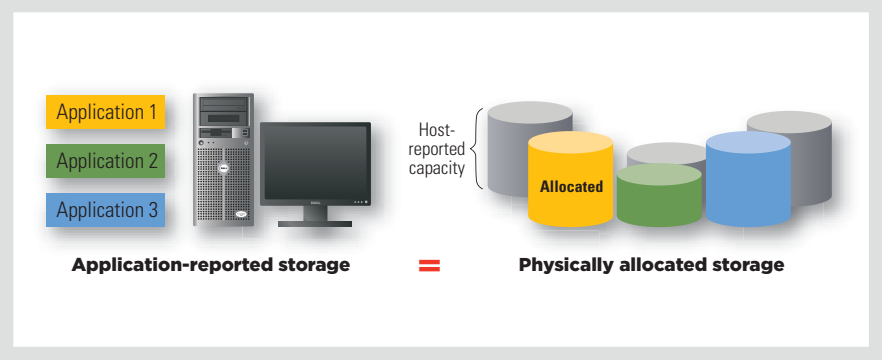

Figure 1. *Traditional provisioning allocates physical storage to match the capacity seen by host servers*

capacity demand, virtual provisioning can dramatically enhance utilization rates while still helping ensure that application capacity, performance, and availability requirements are met. Because capacity can be allocated automatically and with no application downtime or disruption, administrators do not have to overprovision storage in an attempt to avoid the management headache and application disruption often associated with expanding capacity.

The enhanced storage utilization made possible by virtual provisioning provides several key benefits compared with traditional provisioning, including enhanced energy and efficiency for data centers. For example, by avoiding the need to purchase large amounts of excess capacity, this approach can help significantly reduce capital costs, data center footprints, and power and cooling requirements.

SIMPLIFYING STORAGE ADMINISTRATION

Because the tasks associated with provisioning capacity and optimizing utilization can be complex and disruptive, traditionally provisioned storage environments can be complicated to manage—which contributes to storage overprovisioning as administrators try to avoid the need to add capacity in the future. Virtual provisioning offers a range of features designed to simplify provisioning and streamline storage management, helping increase utilization rates and reduce administrative costs.

For example, virtual provisioning adds capacity automatically on an as-needed basis, without the need for manual intervention and without application disruption or downtime. Mapping-service software can automatically configure thin LUNs and manage the placement and usage of storage within thin pools, helping avoid the need for storage administrators to configure and manage storage. And management utilities such as EMC® Navisphere® Management Suite and Navisphere Command-Line Interface (CLI) software enable real-time status monitoring and report generation for thin pools.

Dell/EMC CX4 Series storage arrays also offer organizations the flexibility of deploying different types of provisioning to meet different application and business needs. For example, Dell/EMC CX4 Series storage arrays offer support for both virtualized and non-virtualized LUNs, easy migration between virtualized and non-virtualized LUNs without application disruption or downtime, and management and replication tools such as EMC SnapView™, MirrorView™, SAN Copy™, and Navisphere Management Suite.

CHOOSING BETWEEN THIN LUNs AND TRADITIONAL LUNs

Although deploying thin LUNs in a virtually provisioned environment helps simplify management and increase energy and cost efficiency, virtual provisioning is not well suited for every application environment. Identifying application requirements and then choosing traditional deployment or virtual provisioning based on those needs can be critical when considering allocation of storage capacity.

For example, virtual provisioning is particularly well suited to environments in which data center space is at a premium; easy deployment and management are required; cost considerations, including capital costs and energy costs, are paramount; and space consumption is difficult to forecast. Traditional provisioning is well suited to environments in which performance requirements down to the level of microseconds must be met, performance must be predictable, data placement must be precise, and concern for space is minimal.

Because Dell/EMC CX4 Series storage arrays support both virtualized and non-virtualized environments and include tools to migrate easily between the two, organizations can easily deploy both virtualized and non-virtualized pools of storage within a single storage array and migrate from one to the other as application and business requirements change.

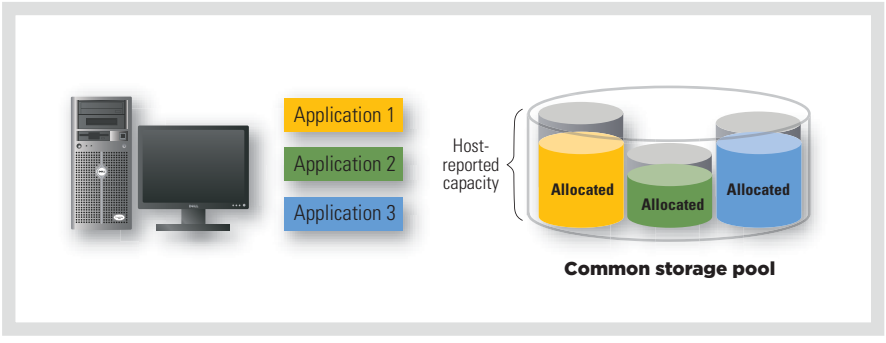

Figure 2. *Virtual provisioning automatically allocates physical storage on demand to help increase utilization while still meeting application requirements*

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

REDUCING COSTS AND ENHANCING ENERGY EFFICIENCY

Combined with the flexibility to deploy both virtualized and non-virtualized storage pools, virtual provisioning in Dell/EMC CX4 Series arrays can help organizations meet application capacity, performance, and availability requirements while controlling costs. In particular, virtual provisioning of storage can enhance utilization, simplify management, and reduce power and space requirements, allowing organizations to reduce the total cost of ownership of their storage infrastructures and enhance data center energy efficiency.

Greg White is a storage marketing manager in the Dell Global Commercial Marketing organization. He has worked for and with small and medium businesses for 14 years, and for the last several years has focused on helping businesses find solutions for their data growth, data management, and data protection problems.

Annette Cormier is a solutions marketing manager for Dell/EMC storage solutions. She has 20 years of experience in developing and bringing to market enterprise storage, network management, and security products for Dell, Hewlett-Packard, and SGI, and has previously been an IT SAS/Oracle® database programmer at the Natural Resource Ecology Lab (NREL) at Colorado State University and at Pacific Power and Light. Annette has a B.S. in Computer Science, Artificial Intelligence, from Colorado State University.

Eric Cannell is a product marketing manager for Dell/EMC storage systems. He has 17 years of experience developing enterprise technology products, before which he spent 6 years as a software engineer at the NASA Jet Propulsion Laboratory. Eric has B.S. degrees in Astronomy and Computer Science from the University of Illinois, an M.S. in Computer Science from the University of Southern California, and an M.B.A. from the Tuck School of Business at Dartmouth College.

MORE ONLINE

DELL.COM/PowerSolutions

QUICK LINKS

Dell/EMC storage:
DELL.COM/EMC

Register for a Dell/EMC event or Webcast:
www.dellenterprise.com/goto/emc/index.html

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

More agility isn't more expensive when your data center works as one.

SUSE® Linux Enterprise Server from Novell® delivers the flexibility you need to respond to a changing business environment. Its built-in virtualization gives you the freedom to run workloads on physical or virtual servers, or migrate between the two. Optimized to work with the applications you need, SUSE Linux Enterprise Server blends into your existing infrastructure—and budget—seamlessly. That way, you can add capabilities to your data center, without adding complexity. Let us make IT work as one for you.

www.novell.com/linux

Novell®
Making IT Work As One™

Copyright © 2009 Novell, Inc. All rights reserved. Novell and the Novell logo are registered trademarks and Making IT Work As One is a trademark of Novell, Inc. in the United States and other countries.

BOOSTING PERFORMANCE WITH ENTERPRISE FLASH DRIVES IN DELL/EMC CX4 SERIES STORAGE

By Bharath Vasudevan
Annette Cormier
Eric Cannell

Enterprise flash drives represent a huge leap forward over traditional mechanical drives. When deployed in Dell/EMC CX4 Series storage arrays, these drives can dramatically boost storage performance for targeted data center workloads while helping increase resource utilization, enhance energy efficiency, and reduce total cost of ownership.

Tremendous performance advances over the years have enabled nearly exponential growth in many areas of technology—processors, memory, I/O buses, and more. Until recently, however, performance improvements in storage have lagged considerably behind other technologies, often resulting in application performance bottlenecks. In particular, the rotational speeds of traditional mechanical hard disk drives (HDDs) have increased only fourfold over the past quarter century or so—from 3,600 rpm to 15,000 rpm. To compensate for the performance gap, administrators often resort to work-arounds such as short stroking and striping. While these measures help achieve the desired performance, the trade-off can be a significant underutilization of capacity. These work-arounds cannot completely overcome the rotational and seek delays inherent in traditional HDDs, which can reduce latency to an acceptable threshold.

Today, enterprise flash drives (EFDs), also known as solid-state drives (SSDs), represent a huge leap forward in storage technology. Because EFDs have no moving parts, they incur none of the rotational and seek delays inherent in traditional HDDs. As a result, EFDs offer tremendous enhancements in performance and response time relative to traditional HDDs. Also, because they require no specialized techniques such as short stroking to enhance performance, they help increase utilization and reduce the number of disks needed to meet performance and response-time requirements. Additionally, because solid-state EFDs are smaller and more energy-efficient than traditional mechanical HDDs, they help reduce data center energy and space requirements.

Because EFDs offer exceptional performance in configurations consisting of far fewer drives than traditional HDDs, EFDs in Dell/EMC CX4 Series storage area network (SAN) arrays allow organizations not only to meet application performance and response-time requirements but also to increase resource utilization and enhance energy efficiency—while helping to reduce total cost of ownership. EFDs are one of several disk drive options available in Dell/EMC CX4 Series storage arrays, which also include Fibre Channel and Serial ATA (SATA) drives.

All drives are designed to leverage the same management and replication features available in Dell/EMC CX4 Series arrays, such as virtual provisioning; EMC® MirrorView™, EMC SAN Copy™, and EMC SnapView™ software; EMC Navisphere® Management Suite software; and the EMC CLARiiON® splitter built into the EMC FLARE® code for integration into EMC RecoverPoint™ appliances. All drives are tested and validated to be used with EMC's leading-edge data protection software: Replication Manager, a host-based tool for synchronizing application and OS processes with snapshots and clones of valued data, and RecoverPoint,

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

an appliance that connects to Dell/EMC CX4 Series arrays to offload concurrent local and remote any-point-in-time recovery of Dell/EMC CX4 volumes. Dell/EMC CX4 Series arrays can also integrate with Dell OpenManage™ systems management software for seamless management of end-to-end storage and servers.

> **"EFDs in Dell/EMC CX4 Series SAN arrays allow organizations not only to meet application performance and response-time requirements but also to increase resource utilization and enhance energy efficiency—while helping to reduce total cost of ownership."**

IDENTIFYING APPROPRIATE WORKLOADS FOR EFDs

Although EFDs deliver enormous performance advantages, their cost and capacity limitations make them most effective when used as a targeted solution for particular load profiles—especially for workloads with extremely high performance and low latency requirements. Examples of this type of workload include online transaction processing (OLTP) database applications such as customer relationship management (CRM), enterprise resource planning (ERP), and order-entry applications; applications requiring extremely fast retrieval and storage of data such as currency exchange and electronic trading systems; and highly read-intensive workloads such as search engine databases.

For example, workload profiles that may benefit from EFD technology are characterized by low cache read hit rates with random I/O patterns, small I/O requests of up to 16 KB, and extremely high transaction throughput requirements. Typical usage scenarios include hot database tables that require read-response times of less than 1 ms and/or high throughput, database temp areas, and high-write-bandwidth applications such as database loads and system backups.

Many other usage scenarios are well suited to EFDs, such as small, highly active file systems that have response-time service-level agreement (SLA) requirements and metadata control areas in cluster file systems that require minimum read-response time. In general, any situation in which performance and response-time requirements would necessitate HDD work-arounds such as short stroking or striping—or which could not achieve the requisite performance and response time even by short stroking traditional HDDs—is one that might benefit from EFD technology.

Enterprise flash drives in Dell/EMC CX4 Series storage arrays offer exceptional performance and energy efficiency for targeted enterprise applications

Not all workload situations are equally suited to EFDs, however. Usage scenarios that may not be cost-efficient for EFDs include workloads with high cache read hit rates already serviced at memory access speed, workloads with large-block sequential reads such as online analytical processing (OLAP) applications, and workloads that are not constrained by rigorous performance and response-time requirements.

BENCHMARKING EFD PERFORMANCE

To quantify the performance of EFDs in Dell/EMC CX4 Series SAN arrays, and to help administrators identify optimum usage scenarios for EFDs, EMC engineers conducted a series of benchmark tests in October 2008 comparing the performance and response time of EFD-based storage arrays with traditional HDD arrays. In particular, EMC compared the performance of EFDs relative to Fibre Channel drives in three configurations: a comparison of 150 Fibre Channel drives with 6 EFDs for an extremely read-intensive workload, a comparison of 75 Fibre Channel drives with 6 EFDs using a read/write OLTP workload, and a direct comparison of 6 Fibre Channel drives with 6 EFDs using a read/write OLTP workload.

The tests were conducted using a Dell™ PowerEdge™ R900 server with four

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

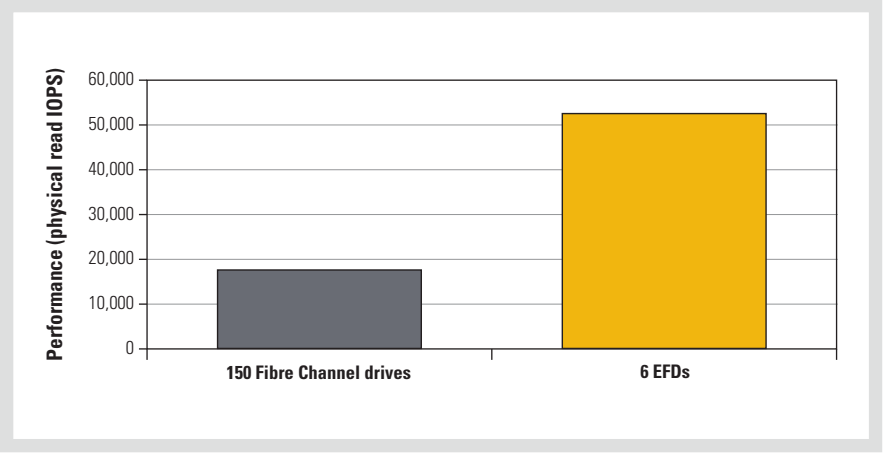

Figure 1. *Physical read I/Os per second for 150 Fibre Channel drives compared with 6 EFDs*

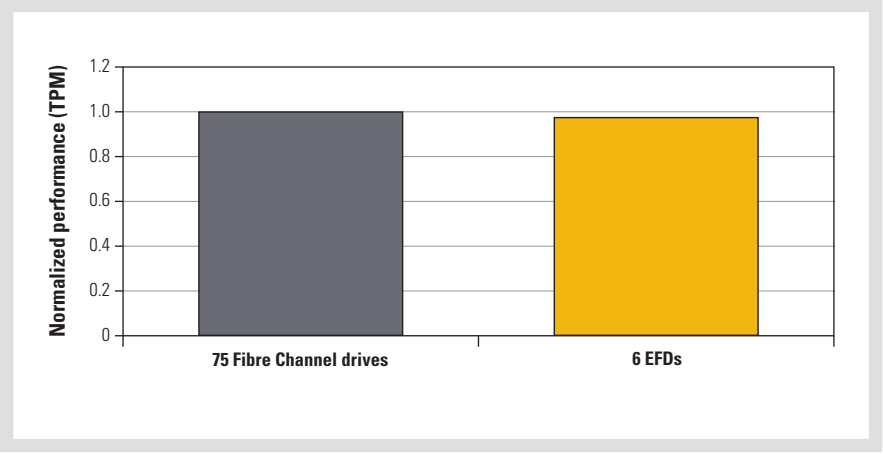

Figure 2. *Normalized transactions per minute for 75 Fibre Channel drives compared with 6 EFDs*

quad-core Intel® Xeon® processors, 32 GB of RAM, and two 146 GB, 15,000 rpm Serial Attached SCSI (SAS) drives supporting the Oracle® Enterprise Linux® 5.2 OS and Oracle database binaries. The database platform was Oracle Database 11.1.0.6 with Oracle Automatic Storage Management (ASM) for volume management and a 260 GB database with 600 warehouses and 100 concurrent users.

The storage configuration consisted of a Dell/EMC CX4-960 SAN array with the EMC FLARE 28 operating environment. The configuration utilized 300 GB, 15,000 rpm short-stroked Fibre Channel drives in a RAID-5 configuration and 73 GB EFDs in a RAID-5 configuration. Test loads were generated using a 64-bit Linux server, and the resulting system loads were measured in several ways: the storage tier load was measured by Navisphere Analyzer, the OS load was measured by vmstat and iostat, and the database usage was measured by Oracle Automatic Workload Repository (AWR) reports.

Read-intensive configuration: 150 Fibre Channel drives versus 6 EFDs

In the read-intensive configuration test, both drive configurations were subjected to a 99.9 percent read-intensive load, which would be typical of a workload such as heavy search indexing. The configuration of 150 short-stroked Fibre Channel drives subjected to this read-intensive load achieved a throughput of 17,459 physical read I/Os per second (IOPS) with an average latency of 4 ms. In contrast, the configuration of 6 EFDs subjected to the same load achieved a throughput of 53,055 physical read IOPS with an average latency of 1 ms.

In this read-intensive environment, EFDs contributed to an enormous enhancement in performance and latency—including over 3 times more read IOPS in total, and over 75 times more IOPS per drive (see Figure 1). Also, the EFD configuration delivered a 75 percent reduction in overall response time. Furthermore, these results were achieved using only 6 EFDs relative to 150 Fibre Channel drives—a 96 percent reduction in the number of drives required. These results illustrate that EFDs achieved more than 75 times more reads per drive than the Fibre Channel drives.

Read/write OLTP configuration: 75 Fibre Channel drives versus 6 EFDs

In the read/write OLTP configuration test, both drive configurations were subjected to a 60/40 read/write OLTP workload. To simulate a real-world environment, both drive configurations were also subjected to a background 40–45 percent load level with caches saturated. Redo logs were kept on the Fibre Channel drives.

In this environment, the configuration with 6 EFDs delivered almost the same transactions per minute (TPM) as the configuration with 75 Fibre Channel drives (see Figure 2). On a per-drive basis, the EFD configuration delivered 12.25 times more TPM per drive than the Fibre Channel configuration. This result was achieved using 92 percent fewer EFDs than Fibre Channel drives.

Direct OLTP comparison: 6 Fibre Channel drives versus 6 EFDs

To more directly measure the performance gains of EFDs relative to Fibre Channel drives, EMC also compared a configuration of 6 Fibre Channel drives directly against a configuration of 6 EFDs using an OLTP workload with a 60/40 read/write ratio. In this benchmark, the EFDs sustained an average of 19,000 TPM,

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

while the Fibre Channel drives sustained only 2,400 TPM (see Figure 3), delivering on average an eightfold enhancement over the Fibre Channel drives.[1]

BOOSTING PERFORMANCE COST-EFFICIENTLY

As shown by the benchmark results in this study, the EFDs in Dell/EMC CX4-960 Series SAN arrays can deliver dramatic performance gains relative to traditional HDDs in both read-intensive and high-transaction-rate OLTP environments. Because far fewer EFDs than traditional HDDs are required to achieve these results, EFD deployments in targeted applications can help increase utilization and energy efficiency while helping to reduce data center costs.

In particular, the significant reduction in drives that an EFD configuration enabled in this study indicates that EFDs can help reduce capital costs, energy requirements for power and cooling, and space requirements for data centers. And because EFDs are one of several disk drive options available in Dell/EMC CX4 Series storage arrays, EFD deployments can also benefit from the management and replication features available with Dell/EMC CX4 Series arrays, which contribute to reduced management costs as well.

> **"As shown by the benchmark results in this study, the EFDs in Dell/EMC CX4-960 Series SAN arrays can deliver dramatic performance gains relative to traditional HDDs in both read-intensive and high-transaction-rate OLTP environments."**

Bharath Vasudevan is the product manager for database programs at Dell. He previously led the solutions engineering team that delivered multiple commercial solutions including e-mail, unified communications, and flexible computing solutions. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

Annette Cormier is a solutions marketing manager for Dell/EMC storage solutions. She has 20 years of experience in developing and bringing to market enterprise storage, network management, and security products for Dell, Hewlett-Packard, and SGI, and has previously been an IT SAS/Oracle database programmer at the Natural Resource Ecology Laboratory (NREL) at Colorado State and at Pacific Power and Light. Annette has a B.S. in Computer Science, Artificial Intelligence, from Colorado State University.

Eric Cannell is a product marketing manager for Dell/EMC storage systems. He has 17 years of experience developing enterprise technology products, before which he spent 6 years as a software engineer at the NASA Jet Propulsion Laboratory. Eric has B.S. degrees in Astronomy and Computer Science from the University of Illinois, an M.S. in Computer Science from the University of Southern California, and an M.B.A. from the Tuck School of Business at Dartmouth College.

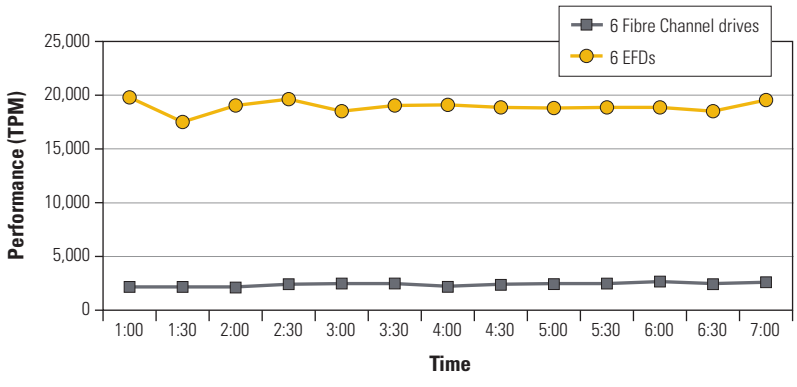

Figure 3. *Transactions per minute for 6 Fibre Channel drives compared with 6 EFDs*

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell/EMC storage:
DELL.COM/EMC

Register for a Dell/EMC CX4 event or Webcast:
www.dellenterprise.com/goto/emc/index.html

[1] For additional details and results from the tests described in this article, see "Leveraging Dell | EMC CX4 with Enterprise Flash Drives for Oracle Database Deployments," by Dell, April 2009, available at DELL.COM/EMC.

Copyright © 2009 Dell Inc. All rights reserved.

A DELL EQUALLOGIC STORAGE SOLUTION FOR ORACLE DATABASE 11*g*

By Wendy Chen
Roger Lopez
Ujjwal Rajbhandari

Dell™ EqualLogic™ PS Series Internet SCSI (iSCSI) storage arrays can provide a cost-effective, high-performance foundation for Oracle® Database 11*g*. This article outlines a recommended architecture and best practices to help enterprises create a highly flexible, scalable, and reliable clustered or single-instance database environment.

Dell EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays are designed to provide the advantages of consolidated networked storage in a cost-effective way. Designed to work together in a storage pool that evenly distributes the load across available resources, these arrays can offer high performance, reliability, scalability, intelligent automation, simplified deployment, and comprehensive data protection in enterprise environments.

Using an EqualLogic PS Series SAN as the foundation of an enterprise-class Oracle Database 11*g* solution can help simplify operations, increase performance and utilization, and enable cost-effective scalability as an organization's needs grow over time. This article provides an architectural overview of a Dell EqualLogic PS Series storage solution for Oracle Database 11*g* along with best practices that organizations can use as the basis for deployments in their own environments.

ARCHITECTURE COMPONENTS

Figure 1 shows the basic architecture of the Dell EqualLogic PS Series storage solution for Oracle Database 11*g*. This architecture comprises the following components:

- Dell OptiPlex™ desktops, which access data stored in the Oracle database
- A client/server network composed of network controllers, cables, and switches
- Dell PowerEdge™ servers running Oracle Database 11*g* with Oracle Real Application Clusters (RAC) 11*g* on the Microsoft® Windows Server® or Red Hat® Enterprise Linux® OS
- An Oracle RAC private network using redundant Dell PowerConnect™ Gigabit Ethernet switches
- Server-storage interconnects using redundant Dell PowerConnect 5400 series or PowerConnect 6200 series Gigabit Ethernet switches
- Dell EqualLogic PS Series iSCSI SAN arrays

STORAGE AND SERVER CONFIGURATION

Dell EqualLogic PS Series iSCSI SAN arrays offer highly available hardware features designed to eliminate single points of failure. Its components are fully redundant and hot swappable, and include optional dual active/standby control modules, standard dual fan trays, and standard dual power supplies. Each control module has multiple Gigabit Ethernet interfaces and a battery-backed write-back cache to help ensure cache coherency between the two control modules.

Host servers can be attached to the arrays through an industry-standard Gigabit Ethernet switch. Figure 2 shows the recommended configuration for the EqualLogic PS Series storage solution for Oracle Database 11*g*, based on a two-node cluster of Dell

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

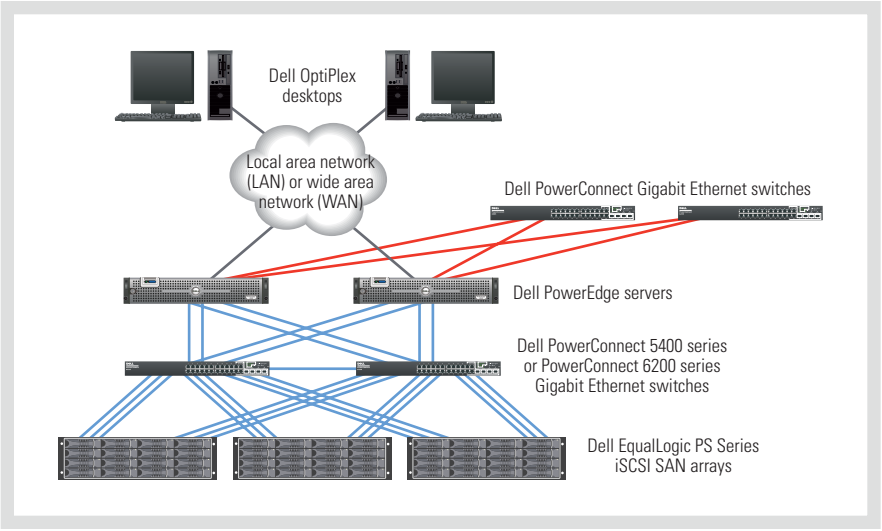

Figure 1. *Oracle Database 11g architecture using Dell EqualLogic PS Series iSCSI SAN arrays*

PowerEdge servers and three EqualLogic PS Series arrays with dual control modules. The two Dell PowerConnect 5400 series or PowerConnect 6200 series Gigabit Ethernet switches help maximize network availability and bandwidth.

Each array in a storage group is referred to as a member. Members are managed and accessed as a single storage system using the group IP address, and administrators must specify each member with a RAID level—RAID-5, RAID-6, RAID-10, or RAID-50—when the array is initialized. As shown in Figure 2, the storage group named oracle-group includes three members: oracle-member01, oracle-member02, and oracle-member03. An EqualLogic PS Series storage group can be segregated into multiple tiers or pools, helping increase administrator control over disk resource allocation. In Figure 2, pools are organized by member RAID levels, with one RAID-10 storage pool containing two RAID-10 members and one RAID-5 pool containing a RAID-5 member. From each control module in the storage arrays, the recommended configuration connects one Gigabit Ethernet interface to one Gigabit Ethernet switch and the other two Gigabit Ethernet interfaces to the other Gigabit Ethernet switch.

At the host level, the PowerEdge servers access the Oracle data through a single storage group IP address. To help provide sufficient bandwidth to support three EqualLogic PS Series arrays, each node of the Oracle RAC 11*g* cluster should have four Gigabit Ethernet network interface card (NIC) ports with independent paths to both Gigabit Ethernet switches of the iSCSI SAN. Multipath software installed on the cluster node can help balance I/O across these ports as well. As shown in Figure 2, each server has a separate connection to the redundant Dell PowerConnect Gigabit Ethernet switches and to the dual control modules on the EqualLogic PS Series arrays. This configuration helps provide pathway redundancy at the host, at the two Gigabit Ethernet switches, and at the dual control modules.

The Oracle Automatic Storage Management (ASM) feature of Oracle Database 11*g* provides vertical integration of the file system and volume manager specifically built for Oracle database files. ASM distributes the I/O load across available

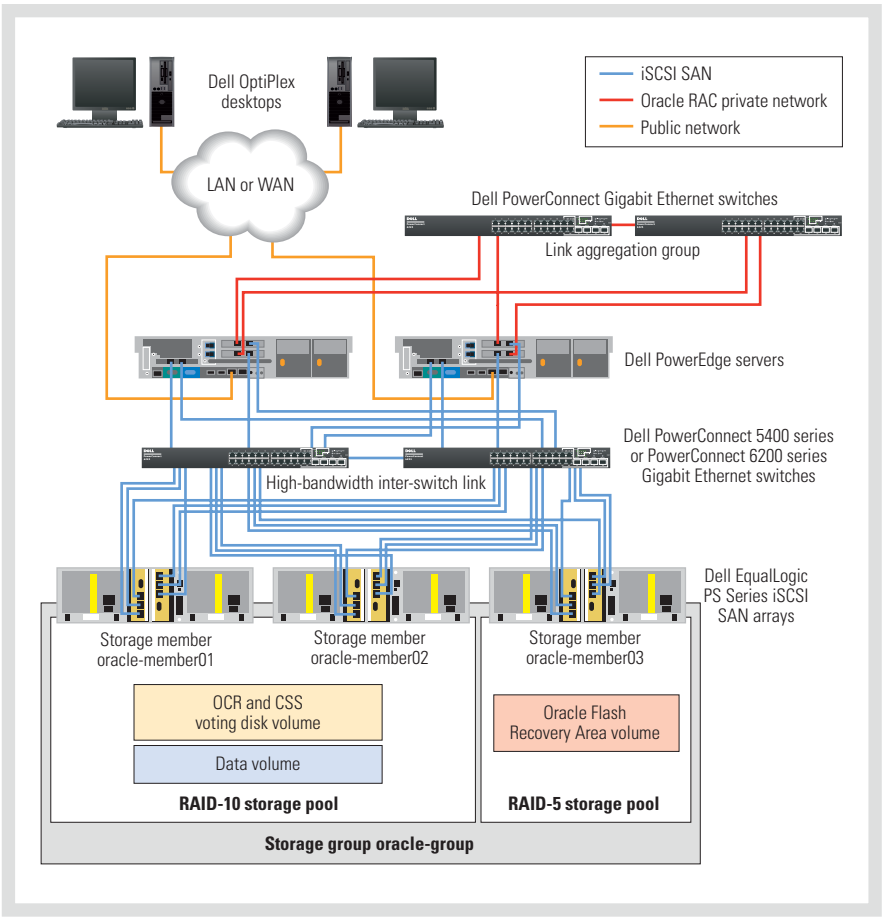

Figure 2. *Recommended configuration for a two-node cluster of Dell PowerEdge servers with three Dell EqualLogic PS Series iSCSI SAN arrays*

resources to help optimize performance while helping eliminate the need for manual I/O tuning such as spreading out the database files to avoid hot spots. ASM helps administrators manage a dynamic database environment by enabling them to increase the database size without shutting down the database.

The shared storage for the Oracle Database 11*g* database can be divided into the following areas:

- **Oracle Cluster Registry (OCR) and Cluster Synchronization Services (CSS) voting disk:** The OCR and CSS voting disk store the details of the cluster configuration and information on which nodes are currently available. These two files can be placed on storage shared by the RAC nodes. Shared storage for Oracle Clusterware is not required when using a single-instance database.
- **Oracle database:** This area holds the actual database stored in the physical files, including the data files, online redo log files, control files, SPFILE for the database instances, and temp files for the temporary tablespaces. The volumes in this area are used to create the ASM disk group and are managed by ASM instances. Although the minimal configuration is one volume per ASM disk group, multiple volumes can be created for one ASM disk group, and more than one ASM disk group can be created for a database.
- **Oracle Flash Recovery Area:** The Oracle Flash Recovery Area is an optional storage location for recovery-related files. If configured, this area stores the disk-based database backup files.

Figure 3 summarizes an example configuration using volumes for each of these three storage areas. RAID-10 is considered the optimal choice for Oracle RAC 11*g* volumes because it offers fault tolerance as well as excellent read and write performance. The EqualLogic PS Series array member on which the data is allocated should be configured with RAID-10.

NETWORK AND INTERCONNECT CONFIGURATION

To help maximize performance, administrators should isolate the iSCSI SAN from other networks. They can achieve this isolation by dedicating Gigabit Ethernet switches for iSCSI traffic or by using virtual LANs to separate networks within a switch. Administrators should also adhere to the following guidelines when configuring the Dell EqualLogic PS Series iSCSI SAN arrays:

- Use Rapid Spanning Tree Protocol (RSTP) and enable the PortFast setting on the switch ports between switches.
- Enable flow control on switch ports and NIC ports.
- Disable unicast storm control on switch ports.
- Enable jumbo frames on switches and NICs.

The servers in an Oracle RAC 11*g* cluster are bound together using the Oracle Clusterware management software, which enables the servers to work together as a single entity. Servers in the cluster communicate and monitor the cluster status using a dedicated private network known as the cluster interconnect or private interconnect. Best practices recommend that administrators implement a fully redundant interconnect network configuration, with redundant private NICs on each server and redundant private network switches.

To initiate traffic between hosts and storage arrays, administrators must install a software or hardware iSCSI initiator. The Microsoft iSCSI Software Initiator, for example, enables connections to an external iSCSI storage array on hosts running the Microsoft Windows® OS. In addition, the Dell EqualLogic Multipath I/O (MPIO) device-specific module works with Microsoft MPIO to support multiple connected paths between hosts and a SAN volume, helping to increase redundancy and performance and simplify configuration.

On hosts running Red Hat Enterprise Linux, accessing iSCSI volumes requires the Open-iSCSI initiator, with multipathing capabilities between hosts and a SAN volume supported by the Linux device mapper module. Both the Open-iSCSI initiator and the Linux device mapper module are part of the standard Red Hat Enterprise Linux distribution.

ORACLE DATABASE AND ORACLE CLUSTERWARE INSTALLATION

The recommended method to install the Oracle Database and Oracle Clusterware software is to use the Oracle Universal Installer (OUI), which provides a simple

	Minimum size	RAID level	Minimum number of partitions	OS mapping
OCR and CSS voting disk	1,024 MB	RAID-10	Two (one for the OCR and one for the CSS voting disk)	Two disk drives
Oracle database	Larger than the size of the database	RAID-10	One	ASM disk group DATABASEDG
Oracle Flash Recovery Area	Large enough to contain archived logs that have not been copied to tape	RAID-5	One	ASM disk group FLASHBACKDG

Figure 3. *Example volume configuration for a Dell EqualLogic PS Series iSCSI SAN*

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

wizard-like installation mechanism. During the installation, the installer asks for general information such as paths for the inventory directory and multi-node information. The Oracle RAC deployment feature of the installer is further enhanced with the ability to push the required binaries to multiple nodes of a RAC cluster from one master server. Best practices recommend first installing Oracle Clusterware 11*g*, then installing Oracle Database 11*g*, and then creating the listener and cluster database.

DELL EQUALLOGIC PS SERIES SNAPSHOTS

Dell EqualLogic PS Series iSCSI SAN arrays include built-in snapshot capabilities at no additional cost, enabling administrators to quickly back up and restore large Oracle databases online with minimal performance impact. Because these snapshots can create a virtually instantaneous point-in-time copy of an Oracle database, this feature can help reduce backup windows to nearly zero while maintaining content-consistent data even on databases spanning multiple volumes.

To evaluate the performance impact of these snapshots, in late 2008 and early 2009 Dell engineers performed benchmark tests using the Transaction Processing Performance Council TPC-C workload in the Quest Benchmark Factory package. The test environment consisted of two Dell PowerEdge 2950 database servers with two quad-core Intel® Xeon® X5460 processors at 3.16 GHz and 32 GB of RAM, configured in a two-node Oracle RAC cluster and running Microsoft Windows Server 2003 Enterprise x64 Edition Release 2 (R2) with Service Pack 2 (SP2) and the 64-bit version of Oracle Database 11.1.0.6 Enterprise Edition. A Dell PowerEdge R710 server with two quad-core Intel Xeon E5520 processors at 2.67 GHz and 16 GB of RAM running Microsoft Windows Server 2008 Standard x64 Edition functioned as the backup server. The servers connected to a Dell EqualLogic PS5000XV iSCSI SAN array with 15,000 rpm Serial Attached SCSI (SAS) disks in a RAID-10 configuration through two trunked 48-port Dell PowerConnect 5448 Gigabit Ethernet switches.

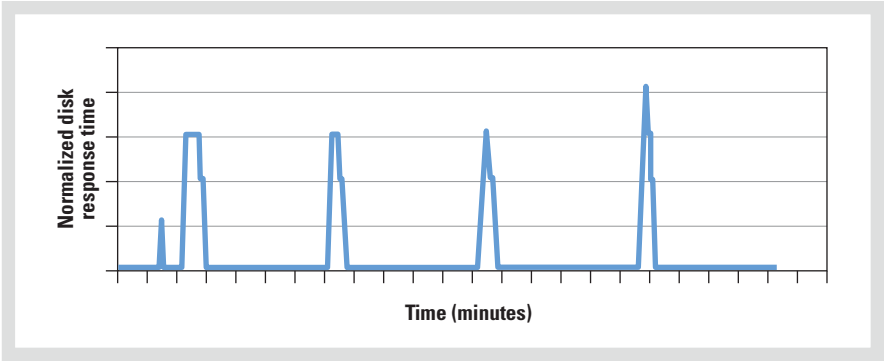

Figure 4. *Normalized disk response times when taking Dell EqualLogic PS Series snapshots of an Oracle database*

Figure 4 shows the effect of four snapshots on disk response time during a TPC-C run performed at a 2,000-user load for 30 minutes. After an initial increase when each snapshot was taken, the response times returned to the baseline within a minute, indicating that subsequent disk activity does not experience the response time latency typically associated with snapshots. The effect is similar even when multiple snapshots of the source volumes are active at the same time. Response times stayed well within an acceptable range throughout the tests.[1]

FLEXIBLE ENTERPRISE DATA ENVIRONMENT

The Dell EqualLogic PS Series storage solution for Oracle Database 11*g* is designed to provide a highly flexible enterprise database environment that administrators can rapidly deploy and easily manage. As enterprise needs grow, the EqualLogic PS Series iSCSI SAN can provide the dynamic scalability to meet those needs while helping to simplify operations and increase performance and availability.

Wendy Chen is a systems engineer on the Oracle Solutions team of Enterprise Solutions Engineering, Dell Product Group. Wendy has an M.S. in Computer Science from Brigham Young University.

Roger Lopez is a systems engineer on the Oracle Solutions team of Enterprise Solutions Engineering, Dell Product Group. Roger has a bachelor's degree in Computer Science from Florida International University.

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

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

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

Dell and Oracle:
DELL.COM/Oracle

[1] For more information and additional results, see "Backup and Recovery of Oracle Database on Dell EqualLogic PS Series iSCSI Storage," by Wendy Chen and Darren Miller, Dell Product Group, March 2009, available in the "Whitepapers" section at DELL.COM/Oracle.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SUPER SCALABILITY: SUPPORTING 90,000 MICROSOFT EXCHANGE MAILBOXES ON DELL EQUALLOGIC STORAGE

By Suresh Jasrasaria
Rob Young

Dell™ EqualLogic™ PS Series arrays feature an innovative peer storage architecture designed for manageability, data protection, and outstanding performance that scales linearly with capacity. In this way, EqualLogic arrays help organizations simply and cost-effectively meet growing storage demands of business applications with maximum efficiency.

Business applications such as Microsoft® Exchange messaging software can place a tremendous burden on enterprise storage infrastructures. For example, because applications often generate random data access patterns—which can put an increased load on not just the disk drives but also the storage controllers—scaling storage capacity efficiently while simultaneously maintaining performance levels and application uptime can be difficult. Consequently, organizations often overprovision storage to help ensure performance, which can be costly and disruptive. And when these enterprise storage environments are scaled to meet application needs, storage provisioning, ongoing management, and backup and recovery tasks can become complex and time-consuming, and can further increase costs.

Dell EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays are exceptionally well suited for business applications, such as Exchange, that keep track of large amounts of data including many small, random updates. Designed to scale the controller processing power and cache capacity with additional disks in the storage subsystems, EqualLogic arrays can virtualize disks efficiently and maintain that efficiency as the solution scales. Because organizations depend on business applications with large data sets and random access patterns on a daily basis, Exchange provides a good example of a business application that requires not only large amounts of highly available storage capacity, but also consistent response time and recovery support from storage environments.

In particular, EqualLogic arrays offer tremendous capacity and are designed to scale efficiently, without reducing performance or disrupting availability. EqualLogic arrays are also designed to support rapid installation, simplified management, and efficient backup and recovery, as well as easy deployment and management. Capacity can be added easily and with no downtime because the EqualLogic arrays employ a peer storage architecture that remains consistent from very small to very large deployments. EqualLogic arrays also come with exceptional hardware and software features such as automated load balancing, multipath I/O, and virtualization—which help simplify management, maximize utilization, and streamline backup and recovery. Overall, the scalability, performance, efficiency, and manageability of EqualLogic PS Series storage arrays can help organizations simply and cost-effectively meet the ever-growing storage demands of business applications such as Exchange.

Dell Labs has validated an Exchange messaging software configuration based on EqualLogic PS5000 series iSCSI SAN arrays for the Microsoft Exchange Solution Reviewed Program (ESRP) – Storage as capable of supporting Exchange environments with up to 90,000 mailboxes (for details, see the

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SUPPORTING 90,000 MICROSOFT EXCHANGE MAILBOXES ON DELL EQUALLOGIC STORAGE

To help Microsoft Exchange administrators assess the ability of a particular storage solution to meet the needs of their Exchange environment, the Microsoft Exchange Solution Reviewed Program (ESRP) – Storage is designed to validate the size and type of Exchange workload a particular storage configuration can handle. In November 2008, Dell Labs validated an Exchange configuration based on Dell EqualLogic PS5000 series Internet SCSI (iSCSI) storage area network (SAN) arrays as capable of supporting Exchange environments with up to 90,000 mailboxes.

The configuration tested 14 Dell EqualLogic PS5000XV arrays with two hundred and twenty-four 450 GB, 15,000 rpm Seagate Cheetah Serial Attached SCSI (SAS) hard drives, separated into one group of 12 arrays with four pools for data and one group of 2 arrays with one pool for logs (see Figure A). The configuration was supported by six Dell PowerEdge™ 1950 servers, six PowerEdge 2950 servers, and three Cisco Catalyst 3750-E series switches, and utilized Ethernet connectivity consisting of four Gigabit Ethernet (GbE) connections per server with two dual-port Intel® PRO/1000 PT GbE adapters.

The target usage profile for this configuration consisted of 90,000 heavy users, with 0.32 I/Os per second (IOPS) per user, plus 20 percent additional headroom, and a mailbox size of 300 MB. The system was configured with eight storage groups per server and five databases in each group.

The EqualLogic-based solution demonstrated outstanding performance, efficiency, and recoverability. In particular, the validated configuration was massively scalable, supporting up to 90,000 mailboxes while providing over 36,000 database IOPS. Test results indicated excellent efficiency by performing 191 database IOPS per disk and extremely fast backup—an overall backup rate of 2,696 MB/sec.

In this particular Exchange configuration, the ESRP – Storage validation also demonstrated that network bandwidth was not a limiting factor. Because the total data rate required to support 90,000 online users was an order of magnitude lower than what GbE iSCSI connections to the servers can provide, expanding network bandwidth through Fibre Channel or 10 Gigabit Ethernet would offer no operational or business benefit.*

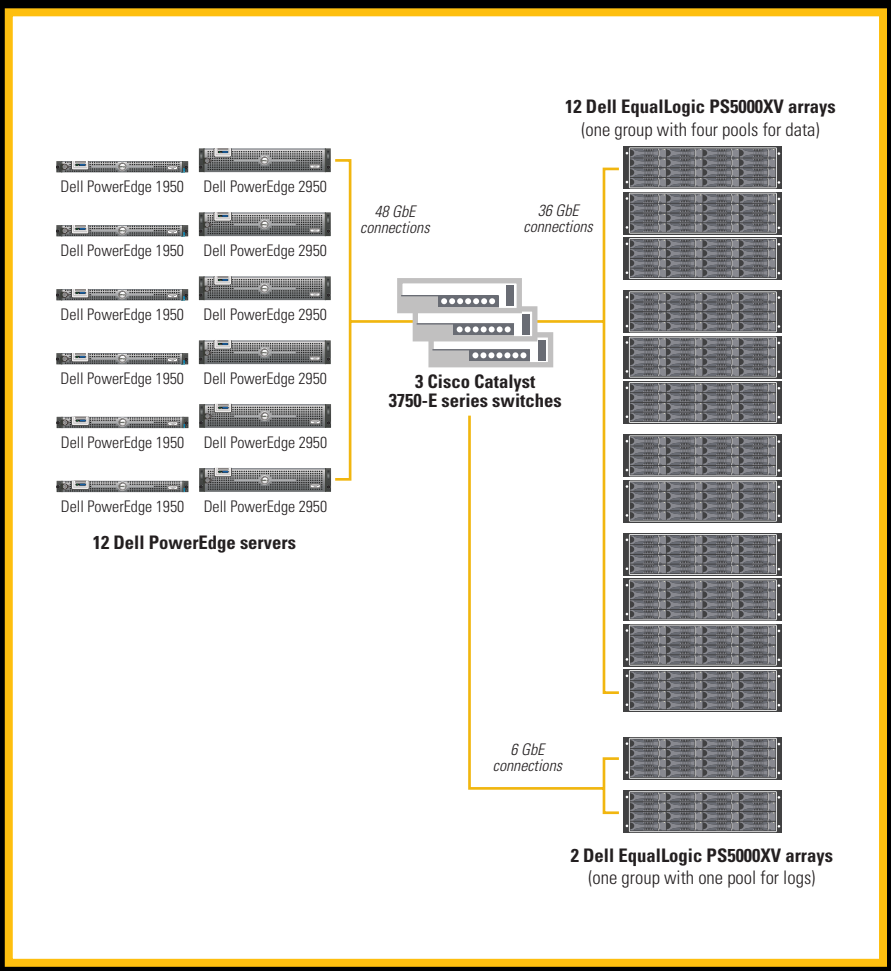

Figure A. *Validated Dell EqualLogic PS Series configuration for the Microsoft Exchange Solution Reviewed Program – Storage*

"The EqualLogic-based solution demonstrated outstanding performance, efficiency, and recoverability. In particular, the validated configuration was massively scalable, supporting up to 90,000 mailboxes while providing over 36,000 database IOPS."

* For additional details and results, see "PS Series Storage Arrays 90,000-User Storage Solution for Microsoft Exchange Server 2007," by Dell, December 2008, available at www.equallogic.com/resourcecenter/assetview.aspx?id=7195.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

"Supporting 90,000 Microsoft Exchange mailboxes on Dell EqualLogic storage" sidebar in this article). The ESRP – Storage program is designed to allow Exchange administrators to assess the ability of a particular storage configuration not only to meet a specific I/O level, but also to meet the needs of their Exchange environment, including maintaining expected response time. To that end, the ESRP – Storage program provides an opportunity for storage vendors to validate the size and type of Exchange workload that a given storage configuration can handle.[1]

The ESRP – Storage program also provides enhanced insights into the capabilities of storage configurations when compared with pure Iometer-type tests that provide maximum I/O throughput. The ESRP – Storage program can assess the ability of storage configurations to handle an I/O load while also maintaining consistently acceptable response times, which is a key metric for many business application environments.

SCALING STORAGE CAPACITY AND PERFORMANCE EFFICIENTLY

Inefficiency is often an impediment to expanding storage capacity for business applications such as Microsoft Exchange because many storage environments do not scale performance in a linear fashion with capacity—that is, doubling the number of disks does not necessarily double the amount of storage I/O throughput. As a result, Exchange administrators may need to add more disks than they expected to use for obtaining the desired I/O throughput.

Often, the bottleneck to efficient scaling is not in the capacity of the disk itself, but in the internal cache and processing capacity and algorithms of the disk controllers. For example, many refrigerator-sized traditional enclosures can hold a huge number of disks, but have only two controllers. As capacity grows, the controllers can become a bottleneck, limiting I/O throughput despite expanded disk capacity. Administrators often must deploy a data-in-place upgrade to a higher controller model and higher level of software licensing than was originally required based on data sheet specifications and numbers of disks configured.

Dell EqualLogic PS Series iSCSI SANs are designed to scale controller processing power and I/O port capacity in concert with disk capacity, enabling efficient scaling while avoiding performance bottlenecks or degradation as capacity is increased. In particular, EqualLogic SANs contain 16 drives within each array, and each array has its own controller and I/O ports. As a result, when an array is added to increase capacity, another controller with its own processor and a set of I/O ports are automatically added. The added controller also includes algorithms in its firmware that are designed to virtualize not only the disks but the controller cache and I/O ports that provide automatic load balancing across all the members in a pool. In this way, controller capacity and I/O bandwidth increase together with disk capacity to enable efficient, linear scaling.

EqualLogic arrays can also be combined to deliver a tremendous overall capacity of over 500 TB within a single group, and they can achieve additional efficiencies through included hardware and software features such as automated load balancing, storage virtualization, multipath I/O, automated tiering, integrated application backup and recovery, VMware® vSphere™ integration, dynamic storage allocation, and remote replication.

SIMPLIFYING STORAGE SCALABILITY

Complexity can be another key impediment to expanding storage capacity. For example, in many storage environments, adding storage can be hard to manage and can often significantly disrupt application availability. In some cases, storage environments need to be re-architected altogether to accommodate growth.

Dell EqualLogic PS Series iSCSI SAN arrays are designed to scale easily, and they can scale from very small to massive storage deployments without requiring any architecture redesign. When additional capacity is needed, an array can simply be added to the group, and capacity and performance can scale automatically and linearly. Capacity can also be added without incurring any system downtime. Volumes can remain available with no impact on hosts or applications, and additional storage and bandwidth are immediately available for use. Load re-balancing across multiple members in a group can be performed automatically without disruption.

A Microsoft Exchange environment based on EqualLogic arrays, for example, can be scaled from a few hundred mailboxes to tens of thousands of mailboxes without any change in storage architecture and without any application disruption. Capacity can be added efficiently and incrementally, without reducing performance or incurring downtime, and local and remote snapshot protection remains in place without disruption. EqualLogic storage arrays also include a range of hardware and software features designed to simplify management of expanded storage capacity, including automated load balancing, storage virtualization, and multipath I/O.

STREAMLINING BACKUP AND RECOVERY

Large-capacity storage environments that are typical in Microsoft Exchange deployments can also be negatively affected by sluggish backup and recovery operations. For example, slow backups in enterprise Exchange storage environments can significantly affect

[1] For more information on the Microsoft ESRP – Storage program, visit technet.microsoft.com/en-us/exchange/bb412164.aspx.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

DELL AND ISV CERTIFIED PARTNERS

HELPING YOU MANAGE IT MORE EFFICIENTLY

The Dell ISV (Independent Software Vendor) Partner Program makes management less complex and more cost-effective by certifying the integration of Dell OpenManage™ hardware management functionality in our ISV Partners' systems management solutions. Our customers get a wide range of benefits including:

- An extensive selection of management solutions that optimally manage Dell platforms
- The ability to manage hardware, software, and IT services, with fewer tools for you to manage
- Dell platform management delivered in many of the management tools you already own

DELL ISV ADVANCED CERTIFIED PARTNERS:

BMC Software
Microsoft
Symantec

DELL ISV BASIC CERTIFIED PARTNERS:

1E
Aternity
Dorado Software
KACE
LANDesk
Nagios
Novell
Platform
Quest Software
ScienceLogic
Veeam

LEARN MORE ABOUT DELL ISV CERTIFIED SOLUTIONS AT
DELL.COM/OpenManage/ISV

REGISTER TO BECOME AN ISV PARTNER AT
DELL.COM/ISV

application performance and/or require application downtime, and slow recovery times can delay application availability as well.

Using Microsoft Volume Shadow Copy Service (VSS) integrated snapshots, Dell EqualLogic PS Series iSCSI SAN arrays can perform rapid backup and recovery without affecting application availability. In fact, EqualLogic arrays demonstrated outstanding performance in Microsoft ESRP – Storage testing for backup and recovery. In this testing, the backup rate was a rapid 2,696 MB/sec.

EqualLogic arrays also come with additional hardware and software features designed to automate and streamline backup and recovery operations. For example, EqualLogic Auto-Snapshot Manager/Microsoft Edition is included with EqualLogic arrays and enables administrators to create automated, application-aware snapshots of Exchange environments on a predefined schedule (see Figure 1). Snapshots can then be used for fast local recovery of inadvertently deleted mail or a corrupted mailbox, and for off-host backup and replication as well.[2]

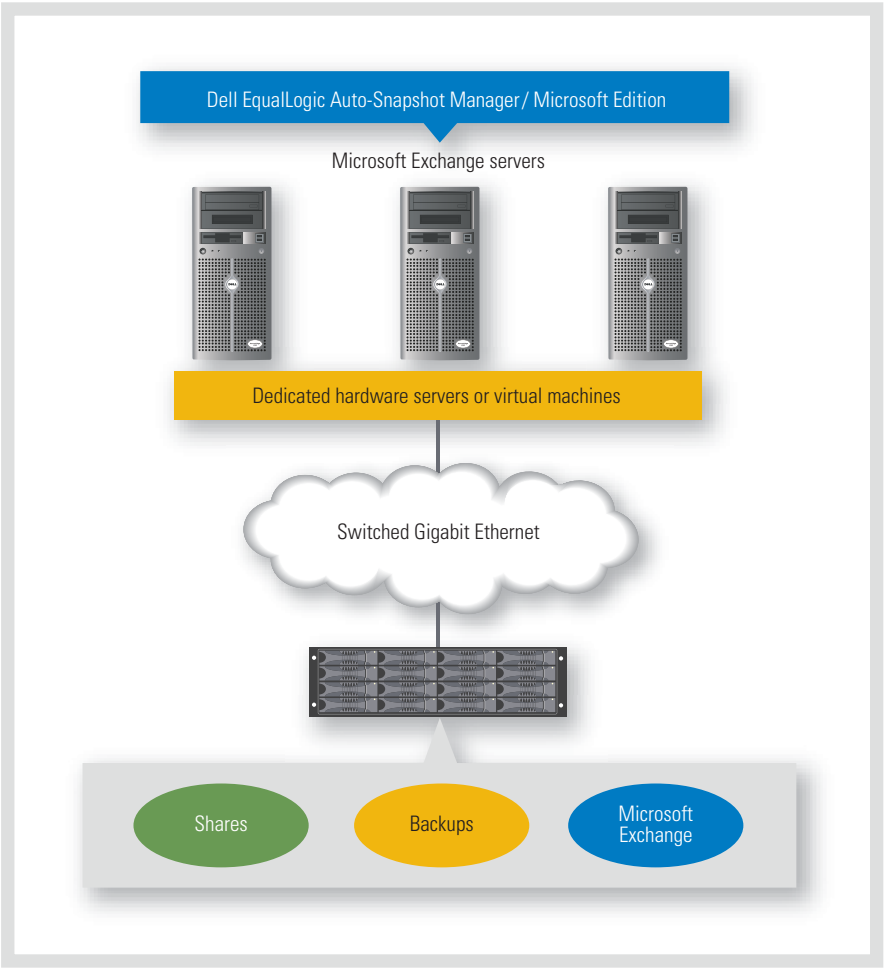

Figure 1. *Simplified protection of Exchange data with Dell EqualLogic Auto-Snapshot Manager/Microsoft Edition*

MEETING STORAGE DEMANDS IN LARGE BUSINESS APPLICATION ENVIRONMENTS

As the demand for business applications such as Microsoft Exchange continues to grow, the task of delivering the storage capacity and performance required by these applications is becoming increasingly difficult. By combining massive scalability and outstanding performance with efficiency and manageability, Dell EqualLogic PS Series iSCSI SAN arrays are designed to help organizations simply and cost-effectively meet the ever-growing storage demands of business applications such as Microsoft Exchange while maintaining application performance levels.

Suresh Jasrasaria is a product marketing senior consultant in the Dell Enterprise Storage Group. Before joining Dell, Suresh worked in the data storage industry for more than 20 years at organizations such as EMC, Veritas, and Digital Equipment Corporation. He has a bachelor's degree in Electrical Engineering from the Indian Institute of Technology Kanpur, a master's degree in Computer Science from the University of Toronto, and an M.B.A. from Boston College.

Rob Young is a technical marketing manager in the Dell Enterprise Storage Group. Rob has 14 years of experience comparing in detail different approaches to storage-related challenges. He has a B.A. from Dartmouth College and a B.E. and an M.B.A. from the Thayer and Tuck schools at Dartmouth.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

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

[2] For more information on using Auto-Snapshot Manager for backing up data, see "How Dell EqualLogic Auto-Snapshot Manager and Veritas NetBackup Enable Off-Host Backups," by Sarah Doherty and Troy Lehman, in *Dell Power Solutions*, March 2009, DELL.COM/Downloads/Global/Power/ps1q09-20090182-Doherty.pdf.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

DELL EQUALLOGIC PS4000 ARRAYS: COST-EFFECTIVE iSCSI SANs FOR SMBs AND REMOTE OFFICES

By Suresh Jasrasaria

For small and medium businesses as well as remote or branch offices in large enterprises, minimizing the cost and complexity of IT operations is essential. New Dell™ EqualLogic™ PS4000 Internet SCSI (iSCSI) storage area network (SAN) arrays are designed with these organizations in mind—offering simplified management and powerful enterprise-class reliability and data protection at a cost-effective price point.

Small and medium businesses (SMBs) and remote or branch offices (ROBOs) within large organizations have traditionally avoided the perceived cost and complexity of networked storage, relying instead on internal server storage or direct attach storage (DAS) to support their IT environments. But although internal server storage and DAS are generally easy and relatively inexpensive to deploy compared with networked storage, they also come with a variety of disadvantages. For example, in multiple-server environments, both approaches create single points of failure, increasing the risk of data loss and downtime; they cannot load balance workloads, creating islands of stranded storage resources; and they cannot scale efficiently, quickly becoming difficult to manage as storage needs increase.

The new Dell EqualLogic PS4000 series of Internet SCSI (iSCSI) storage area network (SAN) arrays is designed with SMBs and ROBOs in mind, providing a simplified way for these organizations to move beyond internal storage and DAS without requiring a large up-front investment, and without sacrificing features or the ability to scale out the SAN deployment with other EqualLogic PS Series arrays as storage needs increase. Optimized specifically for these types of small-scale environments, EqualLogic PS4000 arrays incorporate the same comprehensive enterprise-class capabilities as other EqualLogic arrays to help realize the benefits of an iSCSI SAN at a lower cost of entry.

UNDERSTANDING THE EQUALLOGIC PS SERIES ARCHITECTURE

The emergence of iSCSI has provided a key alternative to traditional Fibre Channel SANs, allowing administrators to create a SAN using existing network infrastructure and familiar Ethernet components while offering performance comparable to Fibre Channel for most workloads. For SMBs and ROBOs, it enables SAN implementation without requiring the level of hardware investment, training, and specialized knowledge typically required in a Fibre Channel SAN.

Dell EqualLogic PS Series arrays are based on a patented peer storage architecture that enables the arrays to work together seamlessly in a single iSCSI storage pool—sharing resources, evenly distributing loads, and self-optimizing performance across the storage group. Because the arrays themselves function as modular building blocks that can automatically and intelligently take on the characteristics of other arrays in the group, administrators can use them to quickly deploy and then easily expand a SAN without requiring cumbersome, complex manual configurations.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Designed for ease of use and high levels of reliability, performance, scalability, and data protection, all EqualLogic arrays include a comprehensive set of enterprise-class storage features at no additional charge to help keep both capital and operating costs down:

- **Ease of use:** EqualLogic arrays are designed for rapid deployment and include intelligent, automated management tools that help streamline SAN configuration and administration. Storage virtualization and integration with VMware®, Microsoft® Hyper-V™, and Citrix® XenServer™ server virtualization software help simplify the creation of comprehensive virtualized infrastructures for organizations of all sizes, while host integration tools help simplify integration with business applications such as the Microsoft Exchange and Microsoft SQL Server® platforms right out of the box.
- **Reliability:** EqualLogic arrays incorporate redundant, hot-swappable components and are designed to provide greater than 99.999 percent data availability. Support for enterprise-class RAID-5, RAID-6, RAID-10, and RAID-50 configurations helps maximize reliability and flexibility.
- **Performance:** EqualLogic arrays are designed to provide exceptional performance for transactional applications typical in SMB and ROBO environments, and performance can scale linearly as arrays are added. Real-time load balancing helps the arrays self-optimize performance and resource utilization, while pooling capabilities enable appropriate service levels for specific applications.
- **Scalability:** The virtualized, modular architecture enables IT administrators to add capacity and performance as they need it without service interruption, helping prevent both underutilization and overprovisioning.
- **Data protection:** Advanced snapshot, clone, and replication capabilities help provide robust, cost-effective protection for critical production data.

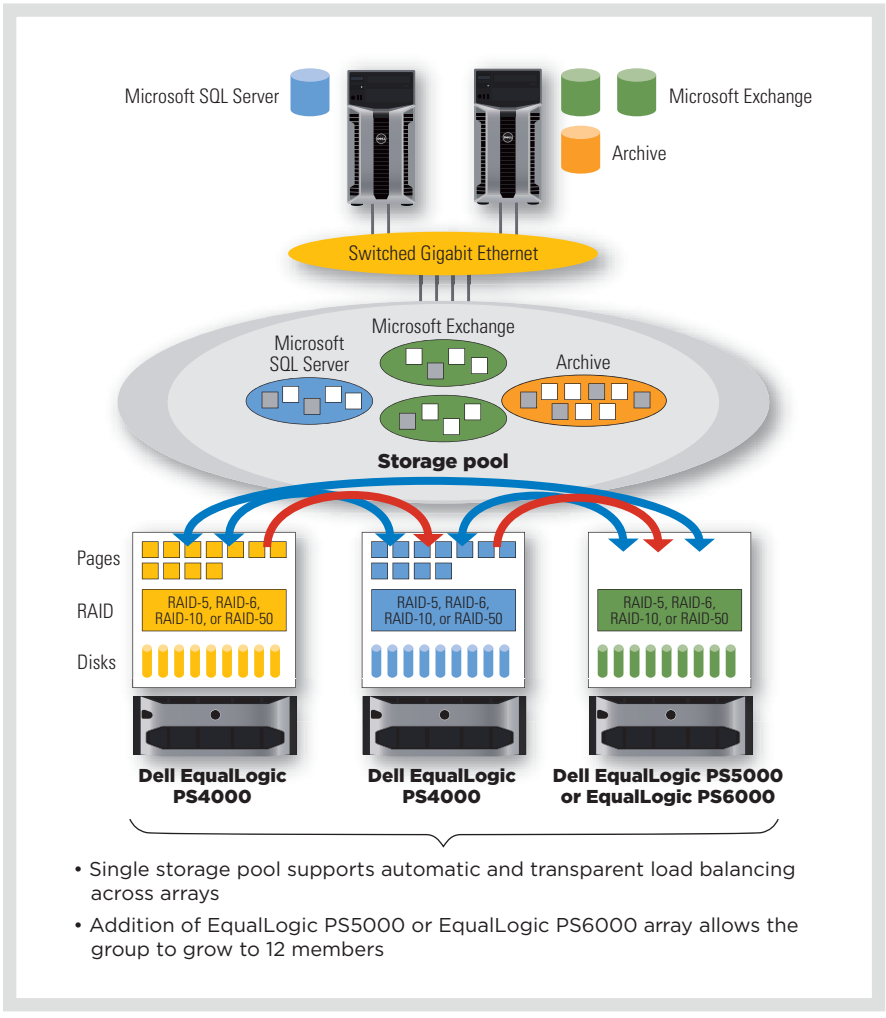

Figure 1. *Dell EqualLogic PS4000 arrays can integrate with EqualLogic PS5000 or EqualLogic PS6000 arrays to support ongoing growth*

DEPLOYING DELL EQUALLOGIC PS4000 ARRAYS

Dell EqualLogic PS4000 iSCSI SAN arrays are designed to meet the needs of SMBs and ROBOs within large organizations, as well as those of large enterprises looking to utilize an iSCSI SAN for departmental applications. The series comprises three arrays, offering a variety of cost, capacity, and performance characteristics to suit different uses. The EqualLogic PS4000E is the lowest-cost, highest-capacity array in the series, supporting up to 16 TB of raw capacity using 7,200 rpm Serial ATA (SATA) drives, and is well suited for storage consolidation projects. The versatile EqualLogic PS4000X balances cost, capacity, and performance using 10,000 rpm Serial Attached SCSI (SAS) drives, and is well suited for storage consolidation, business applications and production databases, and small server virtualization initiatives. And the EqualLogic PS4000XV, the highest-performance array in the series, uses 15,000 rpm SAS drives and is well suited for the most demanding business applications and virtualization initiatives that require strong random I/O performance.

Cost-effective, simplified SAN for SMBs

For SMBs with limited budgets and IT staff, deploying cost-effective, easy-to-manage storage can be critical. EqualLogic PS4000 arrays incorporate the same

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

enterprise-class features as EqualLogic PS6000 arrays, but optimized and scaled down for limited deployments at a lower cost of entry. They are designed for rapid deployment and intuitive management while still providing powerful data protection, self-optimizing performance, and automated load balancing.

EqualLogic PS4000 arrays themselves can scale up to 2 arrays per group, compared with 12 arrays per group for EqualLogic PS6000 arrays; however, they are also designed to integrate seamlessly with other EqualLogic arrays to support significant further SAN scaling as business needs grow. Importantly, the EqualLogic peer storage architecture is designed so that when a higher-end array such as an EqualLogic PS5000 or EqualLogic PS6000 is added to the EqualLogic PS4000 storage pool, the entire SAN adopts the software "personality" of the higher-end array, including snapshot scalability—enabling administrators, for example, to combine up to 2 EqualLogic PS4000 arrays with up to 10 EqualLogic PS5000 or EqualLogic PS6000 arrays in a single storage pool (see Figure 1). This approach enables SMBs to invest only in the level of storage they need now without losing the flexibility to upgrade and expand the SAN in the future.

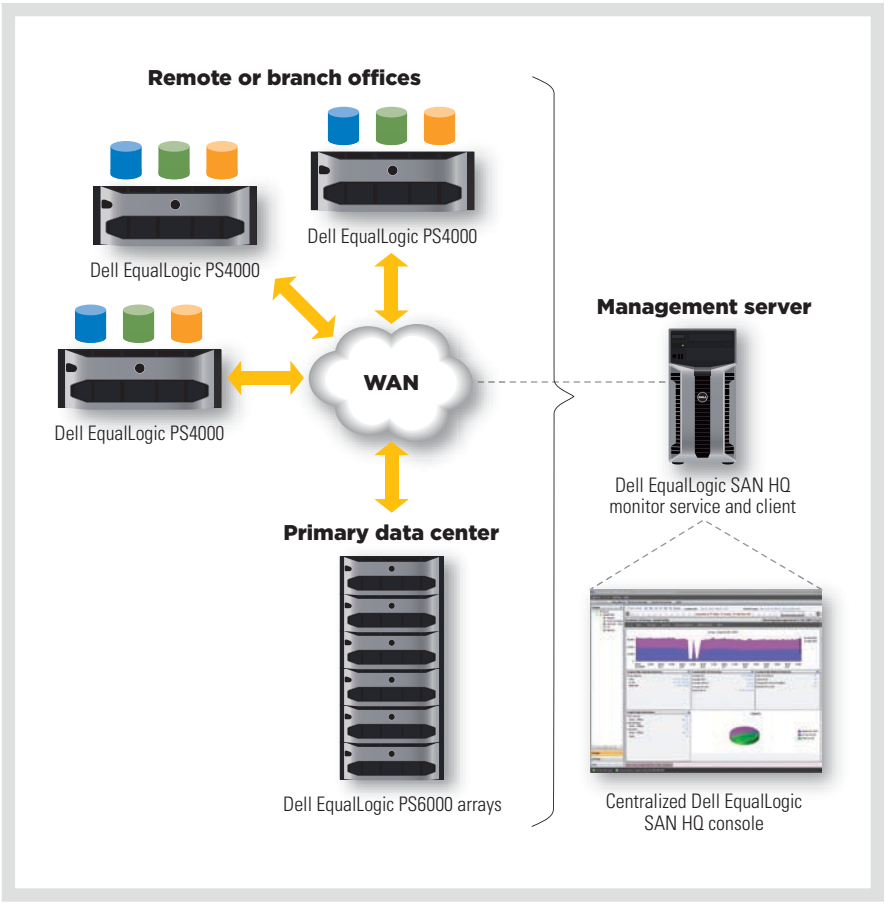

Figure 2. *Dell EqualLogic PS4000 and EqualLogic PS6000 arrays support comprehensive data protection and centralized, enterprise-wide management*

Centralized integration for ROBO SANs

ROBOs within large organizations often need the same high-end capabilities as the enterprise data center itself—including virtualization, comprehensive backup and data protection, and centralized management for IT administrators. EqualLogic PS4000 arrays are designed not only to support the industry-standard VMware, Microsoft Hyper-V, and Citrix XenServer virtualization platforms for small-site deployments, but also to work with EqualLogic arrays in the primary data center to enable simplified data protection, remote backup, disaster recovery, and administration.

For example, deploying EqualLogic PS4000 arrays in ROBOs and EqualLogic PS5000 or EqualLogic PS6000 arrays in a primary data center enables multiple sites to replicate data back to the data center (see Figure 2). This integration helps simplify and consolidate backup and disaster recovery processes across the enterprise and avoids the need to invest in a separate backup infrastructure for every ROBO, which could quickly become expensive and unwieldy to manage—a particular concern for those that lack dedicated, on-site IT staff and depend on nontechnical personnel such as office managers to handle backups. Data center administrators, meanwhile, can use centralized tools such as EqualLogic SAN HeadQuarters (SAN HQ) to help monitor their EqualLogic storage groups enterprise-wide, including arrays in the ROBOs as well as arrays in the primary data center.

Scalable SAN for departmental applications in large enterprises

In addition to helping meet the needs of SMBs and ROBOs, EqualLogic PS4000 arrays are well suited for large enterprises looking to try out iSCSI storage on a limited scale for departmental applications. These arrays not only offer an easy way to quickly get a full-featured, cost-effective iSCSI SAN up and running, but also help ensure that once administrators are ready to scale up their iSCSI SAN deployment, they can do so without requiring a forklift upgrade or complex data migration. Instead, they can begin by using the iSCSI SAN for departmental applications, then take advantage of the EqualLogic peer storage architecture to add EqualLogic PS5000 or EqualLogic PS6000 arrays and gain the scalability and performance to help handle the

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

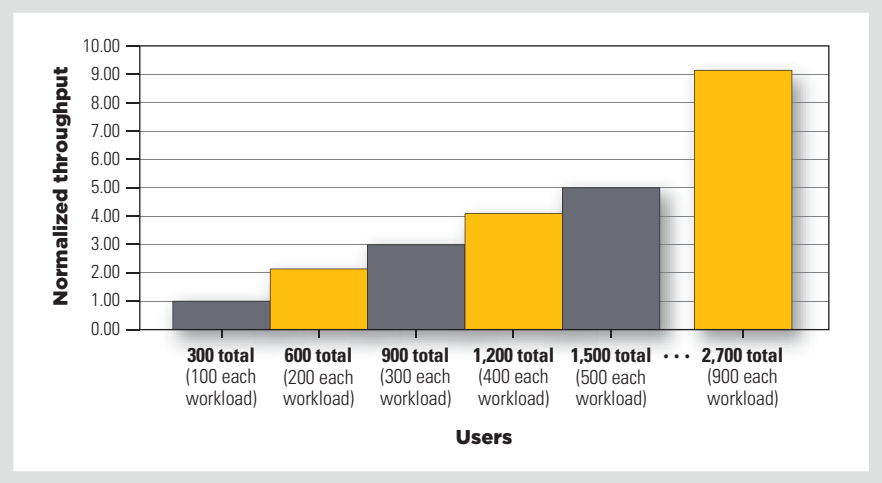

Figure 3. *Dell EqualLogic PS4000XV throughput can scale almost linearly as the number of users increases*

requirements of enterprise-grade production workloads.

EVALUATING DELL EQUALLOGIC PS4000 PERFORMANCE

In June 2009, Principled Technologies conducted benchmark tests commissioned by Dell to determine how many users one Dell EqualLogic PS4000XV iSCSI SAN array can support when running multiple workloads typical of SMB or ROBO environments. The test environment consisted of three Dell PowerEdge™ 2950 servers, each configured with two quad-core Intel® Xeon® E5405 processors at 2.0 GHz; 16 GB of RAM; one 80 GB, 7,200 rpm SAS hard drive; and the Microsoft Windows Server® 2008 Enterprise x64 Edition OS. These servers connected to the Dell EqualLogic PS4000XV array through a Dell PowerConnect™ 6248 Gigabit Ethernet switch using the Microsoft iSCSI Software Initiator; the EqualLogic array contained sixteen 450 GB, 15,000 rpm SAS hard drives in a RAID-50 configuration. Three sets of high-end, custom-built, desktop-class clients as well as a controller system provided the client network for the workloads.

The tests were based on three benchmarks: Dell DVD Store Version 2 (DS2) using a Microsoft SQL Server 2008 database workload, Microsoft Exchange Load Generator (LoadGen) using an Exchange Server 2007 e-mail messaging workload, and Microsoft File Server Capacity Tool (FSCT) as a file server workload. In these tests, the SQL Server database used with the DS2 workload and the Exchange server used with the LoadGen workload each ran in a single Microsoft Hyper-V virtual machine on its own PowerEdge 2950 server.

The test team ran all three benchmarks simultaneously, adding users to each workload in increments of 100 while tracking the throughput provided by the EqualLogic PS4000XV array. Figure 3 illustrates the results, normalized against the baseline of 300 total users (100 users for each of the three workloads). As this figure shows, the throughput of the EqualLogic PS4000XV array scaled almost linearly with the number of users, without overhead, up to 2,700 total users (900 users for each of the three workloads). At 3,000 total users, although the DS2 and LoadGen benchmarks passed the test criteria, the FSCT benchmark showed an overload condition—indicating that even at 2,700 users, the array still had sufficient I/O headroom to handle additional tasks.

For each user count tested, the workloads put greater stress on the storage than the same number of real users typically would under real-world conditions. Therefore, the throughput provided by an EqualLogic PS4000XV array in a RAID-50 configuration would generally be able to support even more users than the maximum number supported in the test environment—which would already be more than sufficient for the number of users typical of SMB and ROBO environments.[1]

MINIMIZING COST AND COMPLEXITY

For SMBs and ROBOs within large enterprises, minimizing costs and complexity is critical in today's economic environment. Dell EqualLogic PS4000 iSCSI SAN arrays are designed with these organizations in mind—offering powerful enterprise-class features and high levels of reliability, performance, scalability, and data protection in a simple, capable, cost-effective storage array.

Suresh Jasrasaria is a product marketing senior consultant in the Dell Enterprise Storage Group. Before joining Dell, Suresh worked in the data storage industry for more than 20 years at companies such as EMC, Veritas, and Digital Equipment Corporation. He has a bachelor's degree in Electrical Engineering from the Indian Institute of Technology Kanpur, a master's degree in Computer Science from the University of Toronto, and an M.B.A. from Boston College.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

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

[1] For the complete report, including detailed information on the test environment, benchmark workloads, methodology, and results, see "Dell EqualLogic PS4000XV Performance with Varying User Counts Running a Combined Database, Email, and File Server Workload," by Principled Technologies, June 2009, available at www.equallogic.com/resourcecenter/assetview.aspx?id=8449.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Windfall

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

Shred the notion of complex desktops.

Centralize desktop management and let ‘er rip

CITRIX

Deliver Desktops. Not confusion.

Visit www.citrix.com today.

CITRIX®

© 2009 Citrix Systems, Inc. All rights reserved. Citrix® and Citrix XenDesktop™ are trademarks of Citrix Systems, Inc. and/or one or more of its subsidiaries, and may be registered in the United States Patent and Trademark Office and in other countries. All other trademarks and registered trademarks are property of their respective owners.

MIXING GIGABIT ETHERNET AND 10 GIGABIT ETHERNET IN A DEDICATED SAN INFRASTRUCTURE

By Tony Ansley

The recent availability of 10 Gigabit Ethernet (10GbE) and the expected enhanced Ethernet standards such as Data Center Bridging (DCB) present an opportunity for administrators to transition to a consolidated Ethernet infrastructure. The reference architecture illustrated in this article shows how enterprises may make the transition to 10GbE for storage using Dell™ EqualLogic™ PS Series arrays and Dell PowerConnect™ switches.

As organizations transition to 10 Gigabit Ethernet (10GbE) technology, the proposed Data Center Bridging (DCB) specification is expected to provide a unified network fabric.[1] These advances can enhance the network storage infrastructure by helping to simplify inefficient deployments of Ethernet adapters, Fibre Channel host bus adapters (HBAs), and cabling that populate many data centers today. Moving to an infrastructure that mixes Gigabit Ethernet (GbE) and 10GbE switches with Dell EqualLogic PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays enables administrators deploying Dell iSCSI storage to incrementally integrate servers and storage with 10GbE network access as they become available while maintaining legacy GbE hosts and storage.

SIMPLIFYING NETWORKED STORAGE ENVIRONMENTS

Providing network access for application servers can be very complex. Typical servers often connect to multiple networks to support client/server, storage, and possibly dedicated management functions. Each of these connections can require multiple ports and cables for each server and potentially require different networking technologies. The introduction of 10GbE is driving a movement toward consolidation of these various usage and technology requirements onto a single Ethernet-based medium. Moving from GbE to 10GbE allows organizations to consolidate multiple GbE LAN connections into a reduced number of 10GbE network connections. LAN and SAN connections can be further consolidated using a unified DCB fabric in conjunction with Fibre Channel over Ethernet (FCoE) and iSCSI, allowing administrators to manage a single Ethernet-based, multipurpose network infrastructure.

Although several standards in the DCB specification have not yet been ratified, early adoption of 10GbE offers several advantages compared with implementing new deployments of Fibre Channel or GbE solutions. Organizations can position themselves for future DCB technology while taking advantage of the increased bandwidth and performance enabled by 10GbE connections today. By eventually consolidating LAN and SAN networks as the standards are adopted, IT managers can minimize adapter, cable, and power usage requirements—thereby helping to reduce total cost of ownership.

Separate Fibre Channel and Ethernet networks deployed by many organizations today require duplicate infrastructure and resources, and moving to 10GbE helps set the stage for reducing this duplication as well. Organizations can use FCoE to connect legacy Fibre Channel infrastructures to the 10GbE

[1] For more information about 10 Gigabit Ethernet and a unified network fabric, see "10 Gigabit Ethernet: Unifying iSCSI and Fibre Channel in a Single Network Fabric," by Achmad Chadran, Gaurav Chawla, and Ujjwal Rajbhandari, in *Dell Power Solutions*, September 2009, DELL.COM/Downloads/Global/Power/ps3q09-20090392-Chadran.pdf.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

network. FCoE enhances flexibility by mapping Fibre Channel frames over Ethernet while preserving the Fibre Channel protocol (see Figure 1). Storage traffic, whether iSCSI or FCoE, can then be carried in Ethernet frames.

Putting 10GbE switches in place today allows organizations to move to converged FCoE and iSCSI using Ethernet on the same adapters. It also allows organizations to help enhance performance over current Fibre Channel connectivity, in many cases bypassing the move to 8 Gbps Fibre Channel and moving directly to 10GbE. FCoE is particularly suitable for organizations with an existing Fibre Channel infrastructure and Fibre Channel skill sets, allowing them to continue to leverage their investment in skills and experience.

The iSCSI protocol also benefits organizations moving to 10GbE. iSCSI allows organizations to take advantage of 10GbE bandwidth today and simultaneously prepare for the expected DCB standard. In addition, iSCSI can work over both legacy and DCB Ethernet networks, offering enhanced performance and delivery characteristics over DCB.

MAKING THE CASE FOR 10 GIGABIT ETHERNET

Several factors make 10GbE implementations a compelling option, including interoperability, cost efficiency, low power consumption, communication quality, and hardware availability. Each of these factors merits careful consideration.

Interoperability leveraging existing technology

During infrastructure upgrades, 10GbE and the TCP/IP protocol are designed to interoperate seamlessly with GbE links, enabling a relatively easy and nondisruptive transition to 10GbE. Two different types of 10GbE connectors are expected to facilitate these links, including 10GBase-T copper and the 10GbE small form-factor pluggable + (SFP+) interconnect. SFP+ supports different physical port types such as twinaxial copper and various types of fiber connections.

By helping ensure that the 10GbE components can cooperatively communicate with GbE devices, switch vendors can deliver interoperability between GbE and 10GbE. Data transitioning from 10GbE to GbE links potentially requires additional buffering on the 10GbE switch to temporarily store the data while it is being transmitted to a low-speed device. In addition, support can be provided for the expected Ethernet standard pause frames (IEEE 802.3x) and priority flow control standards that are part of the enhanced Ethernet standards.

Cost efficiency resulting from fewer connections

Over time, as 10GbE becomes commonplace, one 10GbE port is expected to be more cost-efficient than multiple GbE ports and Fibre Channel ports. Current GbE storage normally requires multiple ports to provide acceptable storage bandwidth between hosts and arrays. Based on industry best practices for redundancy, a minimum of two connections are used to provide a failover path between host and storage. Additional bandwidth may be required by the application—for example, the performance of sequential data applications such as data warehouses is typically gated by bandwidth. Another best practice is to isolate storage traffic on the SAN from client/server traffic on the LAN, which requires a separate LAN port. Dedicated management ports are often required as well. Just two 10GbE connections (for minimal redundancy) in conjunction with enhanced Ethernet standards such as DCB can handle these requirements while still upholding the best practices just described.

Low power consumption with SFP+

Since the 10GBase-T standard was adopted in 2007 for twisted-pair copper cabling, efforts have been underway to help reduce 10GBase-T power consumption—with a goal of reaching power levels per port that are equivalent to the current 1GBase-T standard. First-generation 10GBase-T adapters have higher wattage demands than their short-reach optical counterparts. Currently, prototype second-generation 10GBase-T implementations are designed to bring wattage demand per port down to reasonable levels.

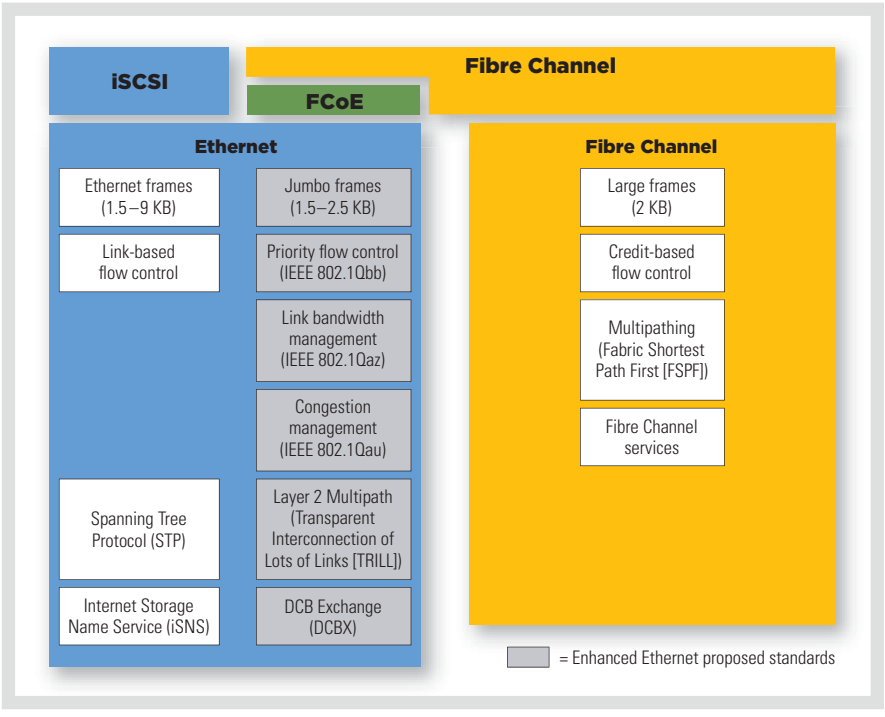

Figure 1. *Transitioning to 10GbE with iSCSI, Fibre Channel, and FCoE connections*

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

10GbE SFP+, which today is an early implementation choice for network and storage vendors such as Dell, has very low wattage requirements per port, and SFP+ direct attach copper connections can provide a power-efficient, cost-effective 10 m cabling reach between rack-mounted servers and a top-of-rack switch. In SFP+ direct attach connections, the module is built into the cables. This effort, along with the reduction in the number of separate connections required to manage multiple networks, should help significantly reduce the power requirements of the network.

Communication quality with compliant standards

The DCB standard is expected to encompass several IEEE 802.1 standards to help ensure communication quality for 10GbE and iSCSI deployments. Priority flow control (802.1Qbb), a link-level flow-control mechanism, is designed to ensure zero loss under congestion in DCB networks. Another standard, 802.1Qau, is intended to provide end-to-end congestion management.

Hardware availability to mix GbE and 10GbE

Hardware such as the Dell PowerConnect series of switches is available today for mixing GbE and 10GbE. For example, the Dell PowerConnect M8024 blade I/O switch modules can configure ports to run at GbE or 10GbE speeds and provide several options for physical connection types, including SFP+. When used in conjunction with an external 10GbE switch, such as the planned PowerConnect 8024F SFP+ switch, and legacy GbE switches, such as the PowerConnect 6200 series, this hardware is expected to offer several options for configuring iSCSI storage solutions that utilize mixed Ethernet speeds.

MIXING GIGABIT ETHERNET TECHNOLOGIES IN A DEDICATED SAN

Administrators can now begin implementing 10GbE technologies in a dedicated SAN to transition to a consolidated Ethernet environment. The reference architecture shown in Figure 2 illustrates one example of a mixed infrastructure that can support GbE and 10GbE hosts today and is capable of integrating expected 10GbE iSCSI arrays into the existing GbE SAN group. This reference architecture can leverage currently available and planned switch and storage components, but similar architectures can be configured using switches from other vendors as well (see the "Best practices for mixing GbE and 10GbE protocols with Dell EqualLogic arrays" sidebar in this article).

The storage in this demonstration consisted of 7 Dell EqualLogic PS6000 series iSCSI arrays. Up to 12 arrays can be included in a single EqualLogic PS Series SAN group, depending on the needs of the organization, and expected 10GbE arrays can be efficiently added to the SAN. The virtualized iSCSI SANs in the reference architecture include both Serial Attached SCSI (SAS) and Serial ATA (SATA) devices for a mix of high-speed and high-capacity drives.

The single PowerEdge M1000e modular blade enclosure was configured with 16 blades. Each blade contained either GbE or 10GbE mezzanine cards that were used to attach to the external iSCSI storage. Internally, these mezzanine cards are connected to two Dell PowerConnect M8024 blade I/O modules that are positioned in the

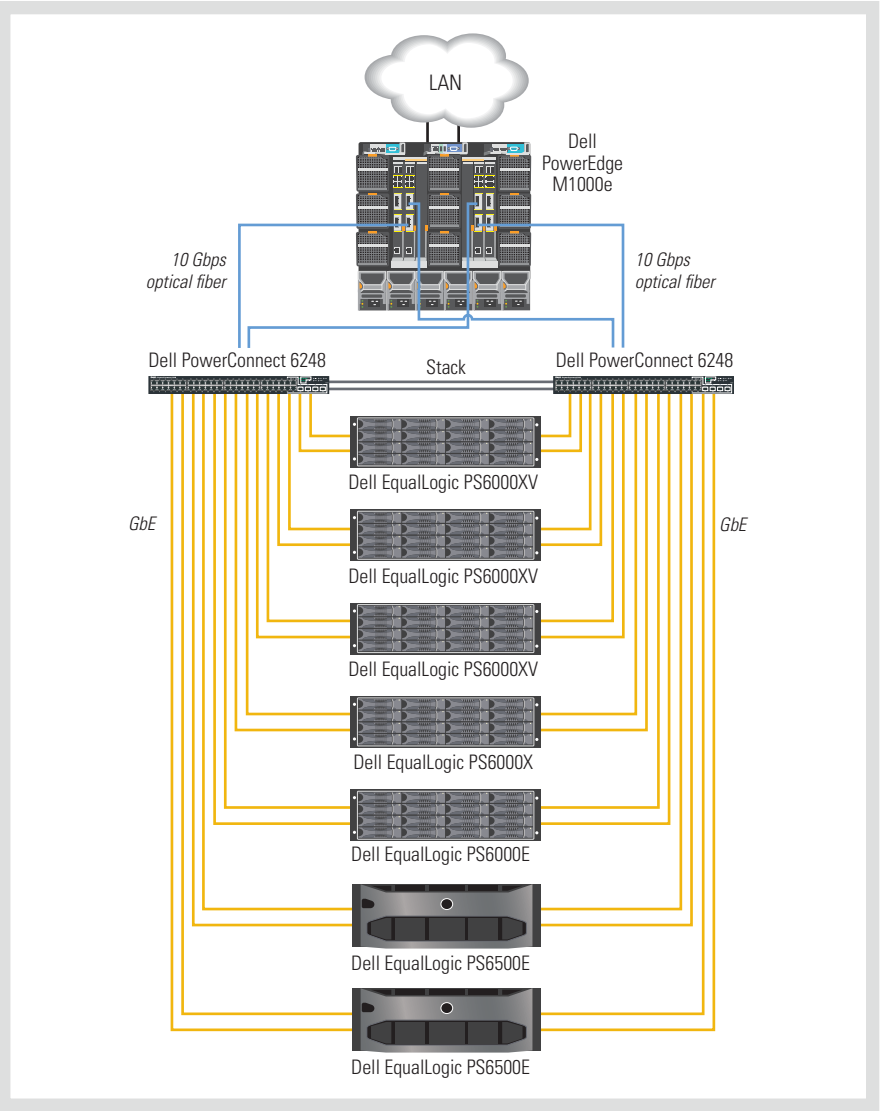

Figure 2. *Mixing GbE and 10GbE in a dedicated SAN architecture*

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

BEST PRACTICES FOR MIXING GbE AND 10GbE PROTOCOLS WITH DELL EQUALLOGIC ARRAYS

Findings gleaned from the reference architecture illustrated in this article include four best practices for mixing Gigabit Ethernet (GbE) and 10 Gigabit Ethernet (10GbE) protocols:

- Enable jumbo frames on the Internet SCSI (iSCSI) storage area network (SAN) switch and Ethernet ports running the Microsoft® iSCSI Software Initiator implementation to help increase throughput and reduce processor overhead.
- Determine the number of inter-switch links to be trunked between switches by determining the number of storage systems attached. Up to four active GbE ports are available on each storage array depending on the actual array model deployed.
- Activate flow control on the hosts, targets, and switches in the SAN to help ensure consistent performance and help ensure that GbE and 10GbE paths cooperate on managing the data transmission through the SAN.
- Ensure that the inter-switch link connection between each switch is connected to at least two other switches for redundancy.

This example scenario shows how organizations may explore similar ways to configure networks, servers, and storage in their own IT environments.

fabric B slots of the PowerEdge M1000e blade chassis. Depending on whether the mezzanine card in a blade server is GbE or 10GbE, the PowerConnect M8024 switches can automatically match the appropriate speed. This approach allows GbE and 10GbE hosts to connect to the 10GbE infrastructure and communicate with the SAN.

Externally, the EqualLogic PS Series arrays in the reference architecture are connected to a pair of PowerConnect 6248 switches that provide GbE connectivity to the arrays. These switches are stacked together to provide a high-performance interconnect between the switches that are directly connected to the EqualLogic PS Series arrays. These external switches are then linked to the 10GbE PowerConnect M8024 switches using 10GbE SFP+ copper or optical links. Additional non-blade servers can be attached to additional external 10GbE switches such as the planned PowerConnect 8024F. These 10GbE switches can be integrated into the infrastructure for solutions in which there are large numbers of stand-alone hosts that need access to the SAN or are connected to the external 10GbE ports on the PowerConnect M8024 blade I/O module—if the number of stand-alone servers is small.

BUILDING ON THE INITIAL DEPLOYMENT

The reference architecture for the SAN configuration discussed in this article illustrates how organizations may leverage GbE arrays now while building out the storage infrastructure and preparing to plug 10GbE arrays into the same infrastructure over time. Deploying 10GbE ports on the host server can enhance access to GbE arrays today—and pave the way for 10GbE iSCSI SANs when they become available. These SANs include Dell EqualLogic PS Series arrays, Dell/EMC CX4 arrays, Dell NX4 network attached storage (NAS) devices, and Dell PowerVault™ MD3000i iSCSI arrays.

By following a similar approach to the dedicated SAN described in this article, organizations may gradually implement tiered storage using GbE and 10GbE technologies. In addition, this approach enables IT managers to boost performance by moving important applications and data onto 10GbE storage targets.

Organizations that currently have Fibre Channel in their data centers are expected to be able to migrate to 10GbE as well by incrementally adding new 10GbE switches that support DCB and provide legacy Fibre Channel ports to merge the existing technology into a consolidated network fabric. This incremental migration capability allows organizations to begin migrating to 10GbE today and ultimately achieve the full benefits of a consolidated network.

CAPITALIZING ON 10 GIGABIT ETHERNET

The reference architecture illustrated in this article shows how it can be viable to mix GbE and 10GbE technologies in a dedicated SAN that helps create a foundation for future growth. As a result, organizations considering iSCSI do not need to wait for 10GbE iSCSI arrays—administrators can roll out a 10GbE switch infrastructure and leverage existing GbE iSCSI storage and hosts today. As 10GbE host network adapters and storage targets—both FCoE and iSCSI—become available, these technologies can be incrementally introduced into the SAN infrastructure. In this way, familiar Ethernet implementation helps simplify IT and reduce total cost of ownership, while helping IT staffs to prepare for expected storage and networking technology advancements.

Tony Ansley is a senior storage consultant on the iSCSI Solutions Marketing team at Dell with over 25 years of experience in the computer industry. He has a bachelor's degree in Information and Computer Sciences from the Georgia Institute of Technology.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell storage solutions:
DELL.COM/PSSeries

Dell Global Services:
DELL.COM/Services

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

10 GIGABIT ETHERNET: UNIFYING iSCSI AND FIBRE CHANNEL IN A SINGLE NETWORK FABRIC

By Achmad Chadran
Gaurav Chawla
Ujjwal Rajbhandari

The advent of 10 Gigabit Ethernet (10GbE), Data Center Bridging (DCB), and Fibre Channel over Ethernet (FCoE) offers enhanced performance and throughput for connecting networked storage to servers. Whether using the Internet SCSI (iSCSI) or Fibre Channel protocols, organizations now have a clear path for unifying a network fabric in 10GbE environments.

Networked storage can provide several key advantages for organizations, including cost reduction and increased efficiency—but it also presents challenges. Storage area networks (SANs) can add their own complexity, and organizations often require increasing levels of throughput for connecting networked storage to servers as the enterprise grows.

The arrival of 10 Gigabit Ethernet (10GbE) along with the Data Center Bridging (DCB) and Fibre Channel over Ethernet (FCoE) specifications holds the promise of a truly converged network fabric. These technologies offer IT administrators a clear path for unifying Internet SCSI (iSCSI) and Fibre Channel SANs while providing enhanced levels of storage efficiency, increased throughput, and cost-effective network storage deployment in their data centers.

SAN GROWTH IN NETWORKED STORAGE ENVIRONMENTS

SANs are essential elements of the move to data center virtualization. In virtualized environments, images and data are stored on a shared SAN to facilitate live migration of virtual machines. SANs are also growing because they deliver value in key areas including storage consolidation, enhanced disk utilization, disaster recovery, and centralized data protection.

Deploying SANs introduces a number of challenges for IT administrators. As the virtualized environment scales, for example, SANs require multiple networks; each network calls for the addition of ports and cables from each server, which can increase costs and power consumption. Servers and storage require advanced integration and management to realize the full benefits of virtualization, further increasing costs. And a virtualized, consolidated infrastructure also creates increased I/O requirements: running multiple virtual machines means supporting multiple I/O streams, and the aggregate of the streams increases the I/O bandwidth and throughput needs for physical servers and storage arrays.

Still another source of complexity is the fact that many organizations deploy two types of networks: Fibre Channel for storage and Ethernet for data. Organizations typically maintain both types because each protocol has its own advantages and disadvantages. The latest Fibre Channel storage devices provide relatively high throughput—hardware is currently available for 8 Gbps Fibre Channel, and is expected to become available for 16 Gbps Fibre Channel—but Fibre Channel can also have high acquisition and administration costs. Ethernet is typically more cost-efficient than Fibre Channel and connects with IP networks to help overcome long distances, but the Gigabit Ethernet (GbE) networking prevalent in

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

today's data centers typically has less throughput and higher latency compared with Fibre Channel.

With dual networks, managing growth and optimizing utilization can become increasingly difficult, costly, and complex. The two network types require separate IT resources, including different hardware and technical expertise, increasing the costs of infrastructure and management. The emergence of iSCSI has allowed cost-effective Ethernet infrastructure to be used as a SAN fabric, and has fueled increasing adoption of iSCSI SANs such as Dell™ EqualLogic™ PS Series arrays.

EMERGENCE OF 10GbE

10GbE is expected to emerge as the future of data center networking because it retains the advantages of Ethernet while opening up new possibilities. For example, new 10GbE components becoming available are expected to preserve the existing Ethernet cost advantage over Fibre Channel, and cost-efficient 10GbE interfaces can help reduce management complexity.

10GbE offers an effective way to expand bandwidth for virtualized environments—providing highly scalable and simplified connectivity by enabling multiple virtual networks to be streamed onto the same physical connection. Using 10GbE connectivity is also generally more power efficient and more cost-effective than using multiple GbE network interface cards.

10GbE also offers a clear path for unifying iSCSI and Fibre Channel storage on a single network fabric. It enables the increased throughput required to unify communications and allow network consistency while building on the familiar, cost-effective Ethernet and IP technology generally already in place in the enterprise. Original equipment manufacturers (OEMs) are readying products supporting the emergence of 10GbE; Dell has introduced a 10GbE iSCSI I/O module for Dell/EMC CX4 Series storage, and plans to add 10GbE capability to its comprehensive range of storage arrays.

UNIFIED NETWORK FABRIC FOR A CONVERGENCE PARADIGM

Organizations are looking for ways to combine their storage and data networks into a single converged fabric to help reduce the total cost of ownership of the data center infrastructure, connect multiple storage islands, and enhance storage scalability. 10GbE offers the necessary throughput to help accomplish this goal, and the DCB specification is the last piece of the convergence paradigm that is falling into place (see the "Enhancing Ethernet bridging" sidebar in this article). The DCB specification provides a set of standards-based extensions to traditional Ethernet, offering a lossless data center transport layer that allows the convergence of LANs and SANs onto a single unified network fabric.

DCB provides advantages for both iSCSI and Fibre Channel storage. Organizations can use the FCoE specification—which depends on DCB capabilities, and is supported by a large number of network and storage vendors—to connect legacy infrastructures to the 10GbE and DCB Ethernet network. FCoE maps Fibre Channel frames over Ethernet while preserving the Fibre Channel protocol, and the DCB specification is designed to maintain the assured-delivery characteristics of the Fibre Channel physical and link layers. The DCB specification also helps ensure delivery for iSCSI storage by providing enhanced congestion management and end-to-end high bandwidth for Ethernet traffic, including iSCSI traffic.

NETWORK FABRIC OPTIONS WITH 10GbE AND DCB

Together, 10GbE and DCB can help provide a single efficient network fabric that offers organizations strategic choices in planning data center networking.[1] Organizations can use iSCSI, FCoE, or both: utilizing FCoE would require new infrastructure investments, while iSCSI simply awaits the forthcoming availability of cost-effective 10GbE interfaces.

Organizations that opt to use FCoE as a bridge to legacy Fibre Channel SANs

ENHANCING ETHERNET BRIDGING

Cisco Systems originally created the term Data Center Ethernet for a set of enhancements to Ethernet bridge standards designed to boost Ethernet Layer 2 congestion management and enable convergence of different traffic types—including not only storage area network (SAN) traffic, but also LAN, management, and Interprocess Communication (IPC) traffic—on the same network. These enhancements led to the IEEE 802.1 Data Center Bridging (DCB) working group, and Cisco now refers directly to the DCB specification, which is expected to soon be formally adopted and compliant with the following standards:

- Priority-based flow control (IEEE 802.1Qbb) at the link level helps ensure no packets are lost under congestion in DCB networks.
- Enhanced transmission selection (IEEE 802.1Qaz) enables administrators to reserve a specific amount of bandwidth for each traffic type to help ensure high quality of service.
- Congestion notification (IEEE 802.1Qau) helps enhance end-to-end congestion management and avoid recurring congestion and frame loss.
- The DCB Exchange (DCBX) protocol helps ensure consistent configuration across the network.

Together with other IEEE 802.1 standards, the DCB specification is expected to help IT organizations take advantage of enhanced communication quality for converged networking.

[1] For information on using DCB in mixed GbE and 10GbE environments, see "Mixing Gigabit Ethernet and 10 Gigabit Ethernet in a Dedicated SAN Infrastructure," by Tony Ansley, in *Dell Power Solutions,* September 2009, DELL.COM/Downloads/Global/Power/ps3q09-20090416-Ansley.pdf.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

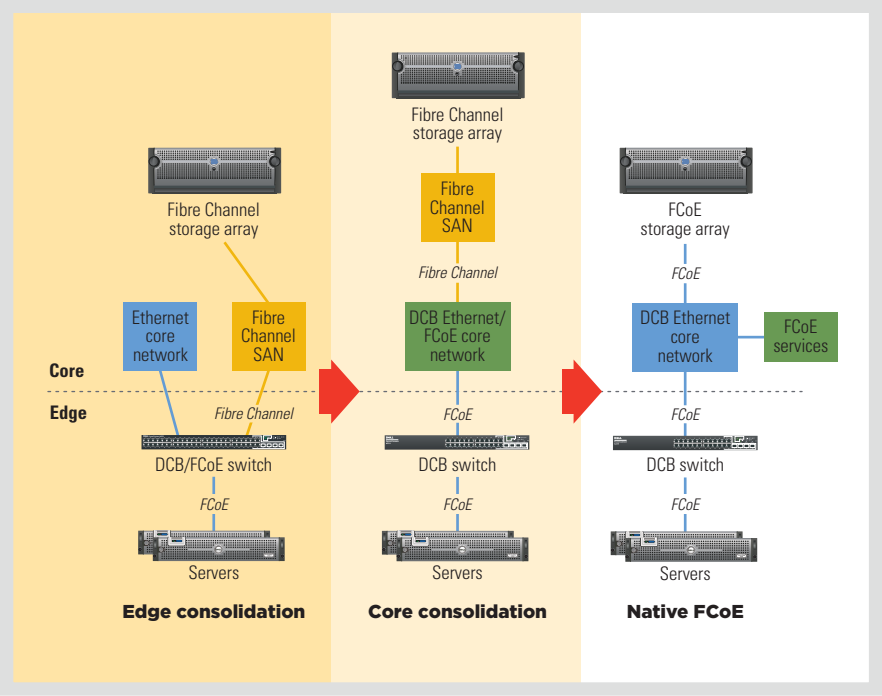

Figure 1. *FCoE offers a bridge to legacy Fibre Channel SANs in support of evolving network consolidation*

are expected to be able to take advantage of the ongoing evolution of FCoE (see Figure 1). The first generation of FCoE-enabled devices is expected to focus on I/O convergence on the server using an Ethernet switch. In the second phase of this evolution, large FCoE networks supported with DCB-enabled switches are expected to provide assured-delivery characteristics over Ethernet that are equivalent to Fibre Channel switches. Finally, the third phase is expected to provide availability of native FCoE storage for connectivity to the FCoE network, which requires FCoE services to run on the DCB network.

CONVERGENCE TO CONSOLIDATE NETWORK STORAGE

Network convergence allows organizations to consolidate storage, providing enhanced levels of efficiency and cost reduction. Multiple networks can share one 10GbE host connection, helping to minimize server adapters, cabling, and power consumption (see Figure 2). Furthermore, combining SAN and LAN traffic on the same network helps significantly reduce the number of adapters, cables, and switches. SAN traffic on the converged network can use either iSCSI or FCoE.

The additional bandwidth availability provided by the unified 10GbE fabric helps to address I/O challenges presented by virtualization of servers and storage. Networks that deploy 10GbE and DCB are expected to support bandwidth up to 20 Gbps using two 10GbE adapters for redundancy. Other benefits include the following:

- **Low support costs:** Convergence can reduce management complexity, and resources no longer need to be divided between Ethernet and Fibre Channel.
- **Expanded high-performance computing (HPC) bandwidth:** 10GbE is designed to expand bandwidth for connecting HPC clusters to the network.
- **Energy savings:** Using fewer adapters, cables, and switches in a unified network fabric than in a legacy data center network helps reduce physical infrastructure, which enhances power and cooling efficiency.
- **Security:** The Ethernet features of virtual LANs and Ethernet bridge access control lists can be used to provide traffic isolation and security for various traffic flows. Security remains robust because storage traffic (iSCSI or FCoE) is simply carried in Ethernet frames.

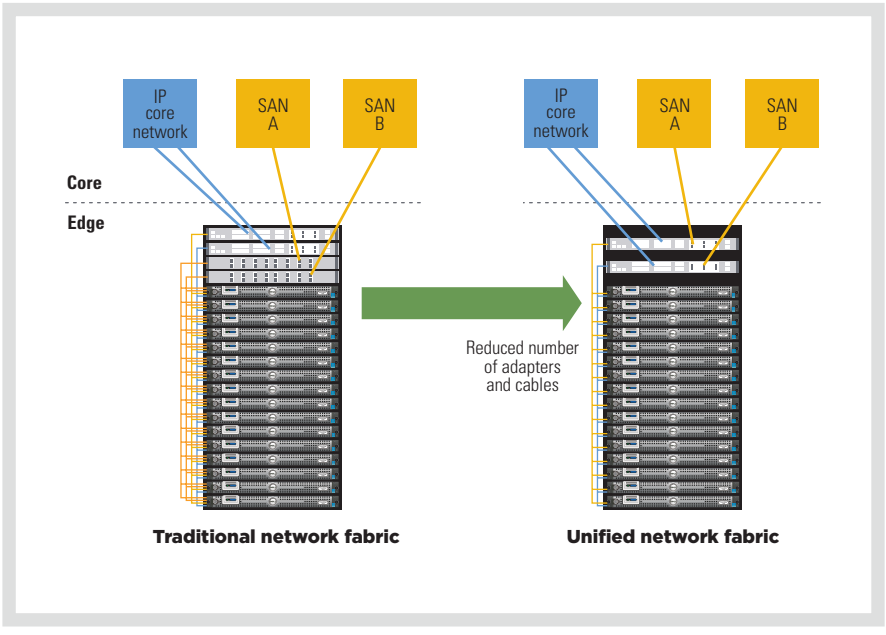

Figure 2. *A unified network fabric helps reduce the number of adapters, cables, and switches*

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

FLEXIBILITY FOR MIGRATING TO UNIFIED NETWORK STORAGE

Organizations have considerable flexibility as they prepare to unify network storage. FCoE can be used to connect FCoE servers to legacy Fibre Channel SANs through Ethernet, preserving the Fibre Channel user experience as the organization migrates to 10GbE, while iSCSI offers the ability to run storage in native Ethernet environments and to route traffic across both LANs and wide area networks (WANs). As the organization migrates from GbE to 10GbE, iSCSI can work in a mixed GbE and 10GbE network environment, and the iSCSI traffic can take advantage of enhanced network features when the infrastructure is upgraded to the DCB specification.

Many enterprises can ultimately benefit by migrating to either iSCSI or FCoE connectivity. Now is a good time for organizations to conduct an extensive review of storage strategy. Enterprises can continue to consolidate operations without anxiety over stranded investments, because the DCB specification for Ethernet can equally benefit both iSCSI and FCoE networked storage.

Achmad Chadran is a storage solution marketing manager in the Dell Large Enterprise Business Unit. Before joining Dell, Achmad held positions including industry analyst, market consultant, and product marketing manager in IT and telecommunications. He has a bachelor's degree from the University of Virginia and a master's degree from Ohio University.

Gaurav Chawla is a technology strategist in the Enterprise Storage Architecture and Technology Group in the Dell Office of the CTO. In this role, he leads the technology initiatives for networked storage and also participates in associated industry standards organizations. He has a B.S. in Computer Engineering from Manipal Institute of Technology, Manipal, and an M.S. in Computer Engineering from Santa Clara University.

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

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell storage solutions:
DELL.COM/Storage

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

NOW AVAILABLE!
Three new Intel® Ethernet Gigabit adapters for Dell PowerEdge™ servers and blades.

GET THE MOST BANDWIDTH OUT OF YOUR DATA CENTER.

The new Intel® Ethernet server and blade adapters feature advanced I/O virtualization, native iSCSI acceleration, and iSCSI Boot technologies for customers who need high levels of performance for their LAN and storage connections.

Intel® Virtualization Technology for Connectivity reduces I/O bottlenecks in a virtualized server and improves performance, flexibility and virtual machine scalability. Optimized for virtualization, Intel's native iSCSI acceleration solution delivers iSCSI performance using customer-preferred, native iSCSI initiators and intelligent stateless offloads, making it the easiest and most cost-effective way to connect servers to iSCSI SANs.

Intel® 10 GbE and Multi-port GbE Server Adapters on Dell PowerEdge™ servers—visit **www.IntelEthernet-Dell.com** to learn more.

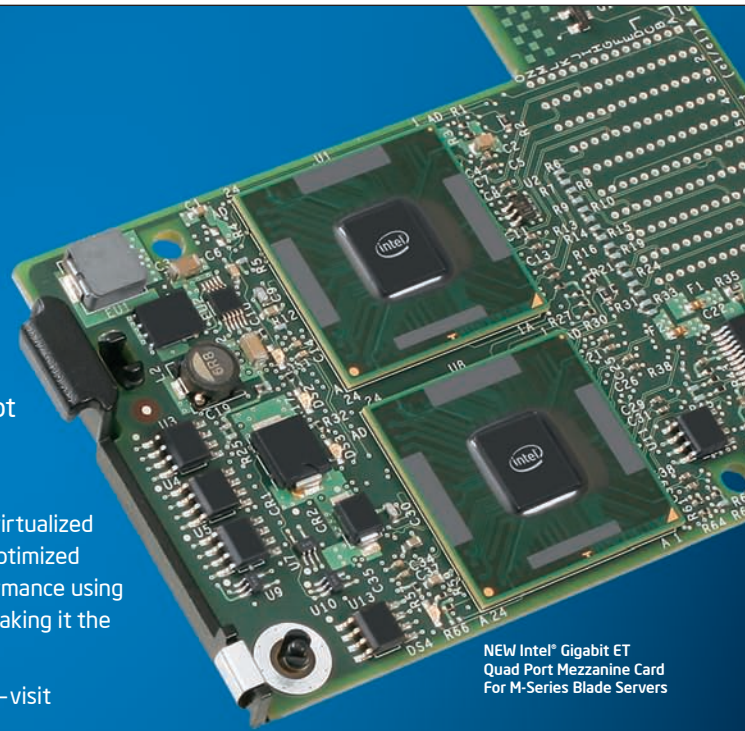

NEW Intel® Gigabit ET
Quad Port Mezzanine Card
For M-Series Blade Servers

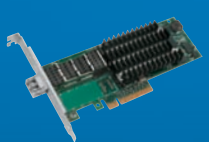

Intel® 10 Gigabit XF SR
Server Adapter

Intel® 10 Gigabit AT
Server Adapter

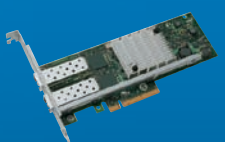

Intel® 10 Gigabit AF DA
Dual Port Server Adapter

NEW Intel® Gigabit ET
Dual Port Server Adapter

NEW Intel® Gigabit ET
Quad Port Server Adapter

intel®

©2009, Intel Corporation. All rights reserved. Intel and the Intel logo are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

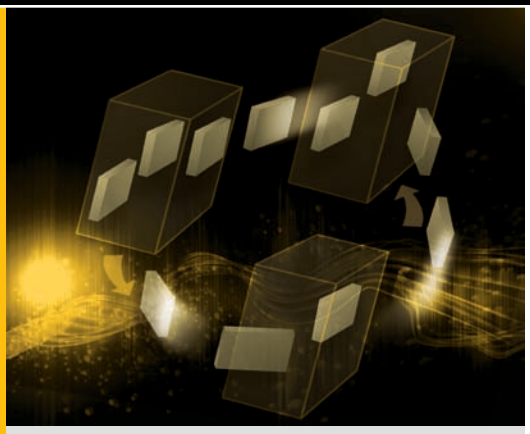

STREAMLINING DATA PROTECTION IN VMWARE ENVIRONMENTS USING COMMVAULT SIMPANA

By Brian Brockway
Zahid Ilkal

Using traditional backup methods in virtualized environments can quickly stretch processing, storage, and networking resources to their limits. CommVault® Simpana® software offers next-generation data protection designed specifically for virtualization, helping streamline backup processes, reduce storage requirements, and support rapid data recovery.

As virtualization continues to sweep through the enterprise—helping to consolidate server workloads and increase resource utilization throughout data centers, remote offices, and recovery sites—administrators are struggling to adjust legacy backup methods to this flexible, ever-changing infrastructure. Backup tools designed to protect traditional physical servers can introduce a variety of problems and limitations when deployed in virtualized environments, including taxing server processing resources, stretching networks and storage capacity to their limits, and slowing recovery times.

CommVault Simpana software, included in the Dell™ PowerVault™ DL2000 – Powered by CommVault, offers next-generation data protection designed specifically for VMware® and Microsoft® virtualization platforms to help overcome these challenges. This article outlines how following best practices for protecting VMware virtual machines (VMs) and taking advantage of key Simpana features such as incremental VM backup capabilities can help administrators streamline their backup processes, reduce storage requirements, and support rapid recovery of both individual files and entire VMs.

DESIGNING FOR VIRTUALIZED ENVIRONMENTS

Implementing data protection strategies in VMware environments presents a number of common challenges:

- **VM sprawl:** As VMs proliferate—often without the knowledge of data protection administrators—it becomes increasingly difficult to manage and apply consistent data protection services based on organizational needs.
- **Infrastructure overload from concurrent backup processing:** Not only can traditional backup agents be cumbersome to implement and maintain inside VMs, but the resulting concurrent backup operations can overwhelm host servers.
- **Balance between disaster recovery and granular recovery needs:** VMware Consolidated Backup (VCB) supports two distinct methods for backing up VMs: image level and file level. In large environments, however, having multiple policies for every VM can complicate management and increase backup media consumption.
- **Retention on and recovery from disk target:** Typical approaches to VM backup utilize the VCB image-level method to create full backups every day, resulting in an excessive volume of data transfer. The lack of incremental backup capability—a feature that is generally taken for granted for physical servers—can cause massive scalability problems. In backup-to-disk (B2D) configurations, the disk target can be exhausted rapidly, forcing backups to move to secondary and slower storage tiers. This in turn reduces the ability to perform fast

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

restores from disk and can negatively affect service-level agreements (SLAs).

Legacy backup software places a clear emphasis on the backup process, generally ignoring factors that contribute to scalability, resiliency, and operational efficiencies. Next-generation tools such as the CommVault Simpana Universal Virtual Server Agent (UVSA) provide a slew of advanced capabilities including automatic discovery, multi-streaming for concurrency, resiliency controls, incremental backup and deduplication for disk efficiency and rapid recovery, extended data retention and recovery options such as granular restore without restaging, and cross-platform restore. The UVSA is designed from the ground up to ease operational overhead for administrators; automate mundane, day-to-day tasks; and allow crucial resources to focus on business problems while providing extended disk-based retention and rapid restore capabilities to help meet aggressive SLAs.

CONFIGURING ROBUST BACKUP PROCESSES

In large environments, it is common for VMs to "float" across physical servers using advanced capabilities such as VMware vMotion™ technology, VMware High Availability (HA), and VMware Distributed Resource Scheduler (DRS). Integration with VMware vCenter™ Server (formerly VMware VirtualCenter) enables administrators to enumerate VMs as well as the VMware ESX servers and data stores hosting them.

Many backup solutions can integrate with vCenter Server during the initial backup policy configuration to locate and list existing VMs. In addition, however, the CommVault Simpana UVSA incorporates features such as automatic VM discovery to automatically assign new VMs to backup policies, backup process offloading, multiple backup options, and multi-streaming to help optimize backups in virtualized environments.

Automatic VM discovery

In large environments, new VMs can be added to the virtualized infrastructure almost daily, often outside the control of data protection administrators. Without automated tools, these administrators are left to manually hunt down new VMs and add them to the appropriate backup policy—a process that is both time-consuming and error-prone, potentially leaving VMs backed up multiple times or (worse) not at all.

The UVSA supports automatic VM discovery to automatically assign new VMs to backup policies based on predefined rules. Administrators create an initial set of backup policies based on desired protection levels, with discovery rules then applied that associate VMs with an appropriate backup policy. The discovery rules need to be defined only once, during the initial setup.

Backup process offloading

VMs are typically configured to balance and maximize physical server resource utilization for regular application workloads. A traditional backup model using an agent in every VM can quickly destabilize this balance. Scheduled simultaneous backups on VMs can consume all available ESX server resources, overwhelming the physical infrastructure and slowing production systems to a crawl.

While all VCB-based data protection tools benefit from some built-in VCB capabilities, VCB does have scalability constraints—and many ungoverned simultaneous VCB snapshot requests can cause failures that then lead to delayed or failed backups. The UVSA includes built-in controls designed to streamline simultaneous VCB snapshot requests and dramatically increase scalability and resiliency, even in very large environments.

Multiple backup options

VCB-based data protection tools typically require two backup policies for each VM: image-level backups for disaster recovery and file-level backups for granular recovery. However, this approach typically doubles the number of backups, increases storage requirements, and complicates recovery by requiring administrators to track two backup sets for each VM.

The UVSA offers three backup options to help balance disaster recovery with granular recovery and to help meet disk retention and recovery needs:

- **Disk-level backup:** This option backs up the full VM image, moving only the occupied portions of a VM image. (For example, for a VM image that is provisioned at 20 GB but consumes only 10 GB of space, the UVSA backs up only 10 GB.) Files and folders are indexed during backup, allowing granular recovery from an image backup. In addition, true block-level incremental (BLI) capability backs up only changed extents of a VM image, helping limit the amount of data that must be transferred and stored.
- **File-level backup:** This option enables backup of individual files within a VM. A separate file-level backup policy is useful when administrators need to separately retain both a disaster recovery copy and a subset of files in the VM. In addition, the ability to index files inside the VM also makes it available for search and e-discovery—a key capability as an increasing amount of critical enterprise data shifts to VMs, given that other e-discovery solutions typically require agents inside VMs that then complicate deployment, management, search, and discovery processes.
- **Volume-level backup:** This hybrid option takes advantage of the VCB file-level backup mode to mount the file system instead of copying the full image over, and then backs up blocks of the file systems for an image-like backup. Recovery options include the ability to restore individual partitions to either a physical system or a VM running on VMware or Microsoft Hyper-V™ platforms. This option also supports BLI backup.

Multi-streaming

VM proliferation can stretch the limits of already-shrinking backup operating windows—and sequential VM backup processing combined with full image processing of

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

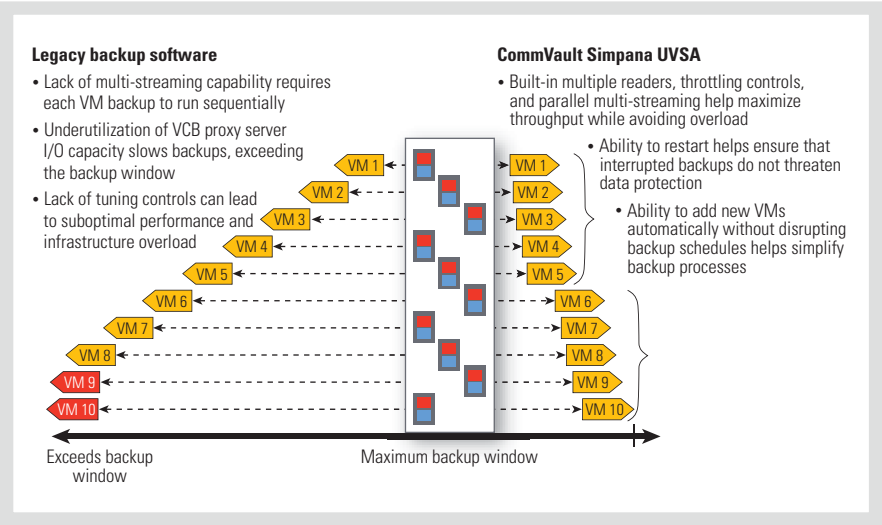

Figure 1. *Multi-streaming helps accelerate backups and avoid infrastructure overload*

every backup is almost certain to exceed operating windows. In these environments, multi-streaming is virtually a necessity for concurrent backups.

Uncontrolled concurrency, however, presents a different set of challenges. Simultaneous backup activity kicks off a period of intense processing, coordination, and movement within a confined operating window. In a virtualized environment, the shared operations and resources yield a large number of variables and conditions that can contribute to failure, rendering the diagnostic process long and tedious and making concurrency a critical unknown factor.

Figure 1 illustrates the results of running a set of backup jobs over 10 individual VMs using a single backup policy, highlighting the operational challenges of consolidated backup policies. The profile on the left, a traditional file system backup policy extended to support full VCB backups, performs sequential processing of each VM and therefore exhibits a much longer overall backup window than the profile on the right. Although the method on the left can help reduce the amount of concentrated load on the system, even underutilizing available resources on the physical server and backup targets, it also risks missing backup windows and exposing the environment to data loss.

The profile on the right uses multi-streaming to support the same set of full VCB backups in significantly less time. Because ungoverned parallel jobs can overload the infrastructure and cause system errors, the UVSA offers two controls to mitigate this problem and help ensure optimal performance. First, administrators can set the number of VMs that can be processed concurrently within a given policy. Second, the UVSA includes a built-in time delay between successive VCB snapshot calls, which helps limit the impact on the system. As each VM backup completes, resources are reapplied to the next VM in the policy. This approach imposes an operational governor across the overall process, helping minimize the impact of unknown conditions through the standard Simpana job management resiliency and job restart features.

OPTIMIZING DATA TRANSFER AND RETENTION

Environments using VCB typically utilize the image-level mode with granular recovery options. As more servers are converted to VMs, therefore, the total volume of nightly data transfer grows in full VM size increments, which can have a major impact on scalability. For example, under a typical schedule of weekly full backups and daily incremental backups, 50 traditional servers averaging 20 GB of data each with a 10 percent daily incremental change rate would require 1,600 GB of data transfer each week. If these servers were converted to VMs in an environment that supports only full VM image backups, the need to back up all 50 VM images on a daily basis would result in 7,000 GB of data transfer each week—an increase of 430 percent. In many environments, this massive inflation can quickly break the networks and storage budgets.

The BLI capabilities provided by the CommVault Simpana UVSA allow administrators to back up only the changed sections of a VM, enabling them to continue following weekly full backup cycles as physical servers are converted to VMs and helping to limit nightly data volumes to comparable pre-virtualization levels. Supporting continued use of existing infrastructure helps limit the need to invest in additional tier 2 disks or tape drives as part of the virtualization project.

Figure 2 illustrates the advantages of BLI backups and the trade-offs imposed when such features are lacking. As with the previous example, this environment has 50 VMs averaging 20 GB each, for a total of 1 TB of data. A typical B2D profile would include two to three times this production disk space for backups and a 30-day disk-based recovery point objective (RPO), or 30-day data retention; reflecting this typical profile, administrators have provisioned 2–3 TB of B2D disk space.

The workflow on the left reflects legacy backup software that does not support incremental VM backups, requiring daily full backups of each VM. In addition to stretching the limits of data transfer capacity each day, this example would fit only two days' worth of backups on the B2D disks—2 TB for two days of full backups, and 0.5 TB of room to recover granular file-level data from previous backups (a typical amount depending on VM size). Meeting the 30-day RPO would typically force administrators to store data older than two days on tape, add disk capacity, or invest in a deduplication appliance to help reduce storage requirements.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

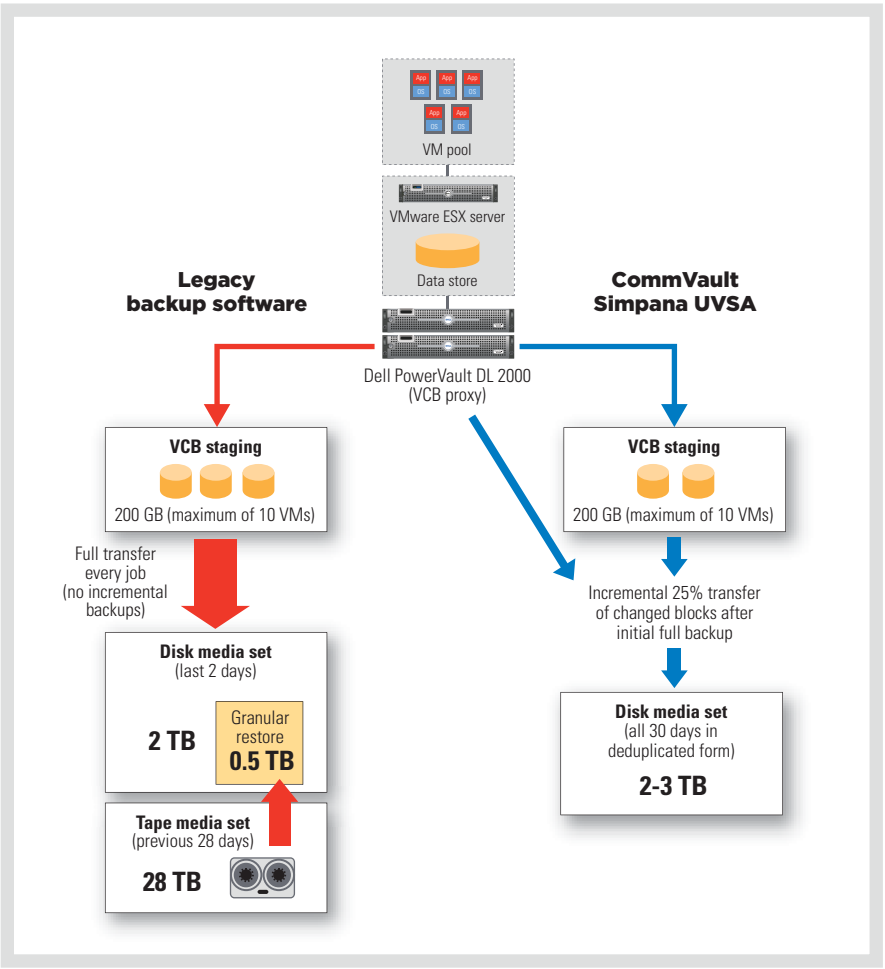

Figure 2. *Block-level incremental backup helps reduce storage requirements and accelerate recovery*

The workflow on the right illustrates the advantages of the BLI approach. In this case, the first full backup transfers the equivalent of 1 TB of data. Then, even assuming a high daily change rate of 25 percent, subsequent backup jobs incrementally transfer changed segments equivalent to 0.25 TB each day. Assuming a typical software compression of 50 percent, the 2–3 TB target can easily accommodate 10–14 days' worth of recovery copies. Adding built-in, in-stream, block-level deduplication would help reduce the volume of data transfer and increase disk utilization even further. With common workloads, and assuming a low 3:1 deduplication ratio, this solution can easily support 30 days' worth of backups on disk, helping meet the 30-day RPO for rapid recovery.

ENABLING RAPID DATA RECOVERY

Administrators often struggle to find the right balance between disaster recovery and granular recovery policies at acceptable performance levels. For example, as shown in Figure 2, a lack of incremental backup capability can quickly force data to tape. In this case, recovering files older than two or three days would require staging an entire VM image on disk—a process that could potentially take hours, defeating the purpose of granular recovery. Administrators face a difficult trade-off: investing in a large disk target to help maintain acceptable granular recovery performance, or using cost-effective tape media that can compromise recovery speed and SLAs.

The CommVault Simpana UVSA helps eliminate this decision. By enabling administrators to maintain a B2D capacity comparable to that required in a pre-virtualization environment while still meeting RPOs, it helps avoid the need to push data off to tape media and supports rapid granular recovery.

DEPLOYING FLEXIBLE DATA PROTECTION FOR VMs

Although virtualization can provide a range of benefits in enterprise data centers, continuing to use legacy backup tools not designed for these types of environments can lead to increased storage costs for data retention and slow, inflexible recovery. CommVault Simpana offers next-generation data protection designed for virtualization—helping organizations keep pace with the emerging challenges of disparate systems, rapid data growth, shrinking operational windows, and tightened recovery objectives.

Brian Brockway is vice president of product management at CommVault, and has over 10 years of experience in the software and storage industry. He has a B.S. in Aerospace Engineering from the University at Buffalo and an M.B.A. from New York University.

Zahid Ilkal is a senior product manager at CommVault, and has over 10 years of experience in the storage industry. He has a master's degree in Computer Science from the University of Pune and an M.B.A. from New York University.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and CommVault:
DELL.COM/CommVault

Dell PowerVault DL2000:
DELL.COM/DL2000

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SIMPLIFYING IT: THE DELL INTELLIGENT DATA MANAGEMENT STRATEGY

By Sanjeet Singh

Rapid data growth and compliance requirements have made data management increasingly challenging, especially for small and medium businesses. The Dell Intelligent Data Management approach incorporates a comprehensive strategy and a set of products and services to help organizations reduce costs and develop a simplified framework for managing enterprise information throughout its life cycle.

Data has become increasingly valuable as both a business asset and a competitive differentiator, and its intrinsic value is being acknowledged in the form of strict government, industry, and enterprise standards for protecting and preserving it. As a result, organizations are generating and keeping increasing amounts of data, with data growth expected to only accelerate in the future.

In challenging economic times, this confluence of rapid data growth and the need to protect it for long periods has created significant challenges for IT leaders, who must find resources to store and manage this data while keeping costs under control. To effectively overcome these challenges, organizations must reconsider how they manage and use their data throughout its life cycle. Current data management solutions that can meet these needs, however, are typically costly and complex. This problem is especially acute for small and medium businesses (SMBs), which are generally restricted to a limited and poorly integrated set of options.

Dell is working to simplify IT with products, services, and partnerships that enable SMBs to easily and cost-effectively manage their data from creation to end of life using a strategy called Intelligent Data Management (IDM). IDM is designed to encompass growth planning and analysis, data protection, archiving, discovery and compliance, and optimization of server and storage resources in the data center. This strategy includes a line of integrated TierDisk systems—including Dell™ PowerVault™ DL2000 disk-based backup solutions based on CommVault® Simpana® and Symantec™ Backup Exec™ software—along with tiered storage based on Dell and Dell/EMC arrays and Dell consulting services.

DATA MANAGEMENT CHALLENGES

IT leaders in organizations of all sizes are feeling the pressure of data growth and increased regulation regarding how long, how securely, and how accessibly that data must be stored. They need reliable backup and restore operations and a way to classify data according to its use, value, and compliance requirements. They must be able to search this data and find what they need quickly. And they want to be able to plan for data growth, rather than respond to it after the fact.

Current efforts to manage data are often insufficient. Too much time and money are spent on incomplete, poorly integrated, and unreliable systems. All too often, IT managers are forced to simply keep all data, and to store it indefinitely. Although deploying additional storage resources can provide a short-term solution, in the long run, this model entails increasing

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

both management complexity and operational costs.

In many organizations, responding to a myriad of department and user needs has created an "island" infrastructure made up of point solutions with multiple management interfaces. An organization may have separate vendors and solutions for tape backup, disk-to-disk backup, archiving, data classification, and search and discovery—and may depend heavily on expensive consulting services to integrate these myriad systems within their data centers or expend significant resources to manage them.

Today, for the most part, well-integrated data management solutions for the full life of an organization's data do not exist. In fact, current solutions focus primarily on large enterprises: high-performance, highly complex solutions that must be managed by trained IT specialists. The cost and complexity of these solutions mean that they are often not feasible options for SMBs. Lacking the budget or expertise to implement complex solutions, many SMBs simply do not deploy data recovery solutions—leaving them exposed to data loss or legal action when data cannot be easily discovered or recovered.

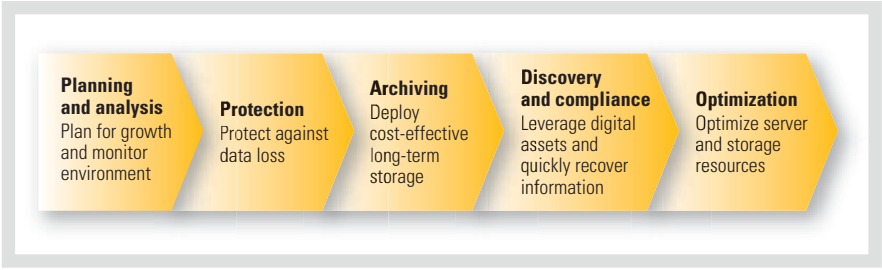

Figure 1. *The Dell Intelligent Data Management strategy provides a comprehensive approach to managing and protecting information*

DELL INTELLIGENT DATA MANAGEMENT STRATEGY

The Dell IDM strategy is designed to help SMBs overcome these challenges by focusing on simplifying data management and disaster recovery—with the goal of keeping data as safe and available as it needs to be at every stage of its life cycle, and helping organizations retain, discover, and recover information based on their specific requirements. Dell is working to develop fully integrated and automated solutions designed to be simple to install, configure, and use, and able to help meet common needs such as data protection and recovery, legal and regulatory compliance, disaster recovery and site recovery, leveraging of digital content or assets, long-term data retention, and IT optimization.

Goals of Intelligent Data Management

To manage and leverage information from creation to end of life, organizations must plan adequately for growth. Administrators must monitor their environments to understand where current systems and architectures are inadequate or wasteful, and to detect potential issues and resolve them before they become problems. Doing so can help IT departments implement the necessary resources to scale gracefully and help reduce or eliminate the possibility of data loss at each phase of the data's life cycle.

Information must be protected both in primary storage and in cost-effective long-term storage, and must be available through search and discovery mechanisms that enable organizations to utilize their digital assets and recover information quickly. Cost-effectiveness is vital—and encompasses, beyond the initial purchase price, the cost of long-term management and efficient use of resources. Complex solutions, for example, require more administrator hours than simple ones, which can add dramatically to total cost of ownership, as can inefficient use of server and storage resources. Optimization is a hallmark of IDM.

The most important hallmark of IDM, however, is IT simplification. The resources to implement effective data management exist today, but as disparate products. Organizations end up struggling to manage fragmented systems with multiple manual consoles and minimal integration—backup and archiving, for example, are often separate operations. Using consulting services to integrate these products may still result in multiple management interfaces and ongoing manual work to manage the data.

> **"Dell is laying out a road map for a set of IDM solutions under its TierDisk family of products, which can help SMBs solve their data management problems by integrating multiple functions."**

Key components of Intelligent Data Management

IDM enables systems that perform individual functions to work together seamlessly, with no data silos that would make data migration a complicated manual task. In addition, organizations should not have to purchase an overconfigured solution—instead, they should be able to scale their solutions as the amount of data grows. Existing technologies can help control the growth of resources required to maintain exponentially growing data, but should be well integrated with solutions performing other functions.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Dell is laying out a road map for a set of IDM solutions under its TierDisk family of products, which can help SMBs solve their data management problems by integrating multiple functions. The IDM strategy encompasses several key components (see Figure 1):

- **Planning and analysis:** Planning and analysis can be as simple as understanding current workloads or as complicated as trying to predict what an organization's needs will be at some time in the future, and can help organizations identify the problems that they need to address. Enterprises typically accomplish such tasks through manual tracking and predictive analysis. Others may utilize storage resource management software. Solutions may employ storage assessment consulting or managed services that can take advantage of outside expertise for planning and analyzing IT environments.
- **Protection:** Data protection typically involves storing multiple copies of data to traditional disk- or tape-based systems using backup software, or using virtual tape libraries—disk-based devices that emulate tape systems.
- **Archiving:** Archiving is often equated with long-term retention of backup data using cost-effective tape-based systems. However, traditional tape-based backups often lack the value of a truly effective archiving solution—one that can remove inactive or infrequently accessed data from expensive primary storage to lower-cost media while maintaining appropriate levels of access to and governance of that data throughout its life cycle.
- **Discovery and compliance:** Effective archiving solutions offer robust data classification and policy-based rules engines for efficient, compliant content retention, search, and discoverability. This feature is especially important for complying with industry, government, or enterprise records-retention policies or to support litigation. Traditional approaches to discovery and compliance are often reactive, manually intensive, and ultimately ineffective efforts to locate data across disparate systems and media types. Even large enterprises often outsource discovery requests to consulting companies that specialize in performing such tasks.
- **Optimization:** Many technologies can help optimize IT resources. Deduplication, for example, can help reduce the data footprint by maintaining single copies of redundant data. Organizations can also tier their storage, moving infrequently accessed data from high-performance storage arrays to more cost-effective, lower-performance arrays. Energy-efficient systems can help reduce power requirements. Server virtualization can enable the consolidation of multiple workloads onto a single physical server.

Dell PowerVault DL2000 solutions offer integrated, automated disk-to-disk-to-tape backup

DELL INTELLIGENT DATA MANAGEMENT SOLUTIONS

Dell is partnering with industry leaders to deliver integrated products and services designed to simplify information management through IDM, including the Dell PowerVault DL2000, tiered Dell and Dell/EMC storage, data deduplication, and consulting services.

Dell PowerVault DL2000

The Dell PowerVault DL2000 – Powered by CommVault and the Dell PowerVault DL2000 – Powered by Symantec Backup Exec are integrated, automated disk-to-disk-to-tape solutions designed to manage and protect information at multiple points in its life cycle. They enable retention and recovery according to an organization's specific requirements while helping reduce the cost and complexity of data management.

The systems come factory installed with CommVault Simpana or Symantec Backup Exec software, along with wizards designed to help administrators set them up quickly and easily. They include automated dynamic disk provisioning that configures the disks without administrator intervention, and provide integrated support for tape libraries to help simplify offsite storage for disaster recovery. The systems are also available with Dell Storage Assessment Services, which can help plan the overall data management strategy and PowerVault DL2000 deployment.

The PowerVault DL2000 – Powered by CommVault additionally comes with content-aware, block-based deduplication capabilities integrated into the software, combining two components of IDM—protection and optimization—in a single system. Archiving functionality can also be added, creating a single solution that can protect and archive data while optimizing disk usage.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

> **"IDM enables systems that perform individual functions to work together seamlessly, with no data silos that would make data migration a complicated manual task."**

Tiered Dell and Dell/EMC storage

Implementing tiered storage arrays can help provide important cost savings over a one-size-fits-all storage model. By using high-performance disks or arrays to host primary data and cost-effective, lower-performance disks or arrays to host secondary or infrequently accessed data, organizations can appropriately optimize each storage investment for its intended use.

Dell storage includes dual-protocol Fibre Channel and Internet SCSI (iSCSI) arrays such as the Dell/EMC CX or AX Series, iSCSI arrays such as the Dell EqualLogic™ PS Series and PowerVault MD Series, and direct attach storage and tape—a broad range to enable tiered storage. Dell consulting services can help organizations classify their data and optimize how they store, protect, and manage it in a tiered architecture.

Dell/EMC CX4 arrays and EqualLogic PS6000 series arrays with Serial Attached SCSI (SAS), Fibre Channel, or solid-state drives are well suited for hosting primary data for high-performance applications such as e-mail systems and databases. For secondary storage, Dell recommends Dell/EMC CX4 Series arrays with Serial ATA (SATA) drives or EqualLogic PS5500 arrays with mirroring to disaster recovery sites or software for disk-based backups.

Data deduplication

Deduplication is a storage optimization technology that helps reduce the data footprint by deleting multiple copies of redundant data and storing only unique data, replacing data copies with references to the original data. Deduplication at the file, sub-file, or block level can help provide the benefits of disk-based backup at costs approaching those of tape.

Consistent with the IDM strategy, Dell believes that deduplication is a storage feature rather than a point solution; as deduplication technology becomes widespread, this feature becomes increasingly ubiquitous. Dell is currently delivering content-aware, block-level deduplication embedded within the software of the PowerVault DL2000 – Powered by CommVault,[1] and plans to continue integrating deduplication into other products within this portfolio.

Consulting services

Dell has introduced services to help analyze enterprise IT environments and develop plans designed to optimize storage and improve data recovery. Dell Storage Optimization Services can help identify opportunities to control costs using technologies such as tiering, deduplication, and archiving. Dell Data Protection Services can help identify flaws in existing data protection processes, design backup and restore systems, and set up disaster recovery infrastructures to help meet data protection requirements. Dell uses a tools-based approach to assess backup environments so that the design is based on existing implementations as well as Dell's real-world expertise. This approach and standard methodologies enable Dell to complete the tasks in a cost-effective way, thereby helping to free IT staff to focus on planning and related strategic projects.

COST-EFFECTIVE, SIMPLIFIED DATA MANAGEMENT

Effectively managing rapid data growth and protecting data for extended periods requires organizations to integrate how they manage and use their data from creation to end of life. The Dell IDM strategy encompasses products and services designed to fit seamlessly into existing IT environments while providing the advanced tools necessary to help easily and cost-effectively safeguard data. In the future, Dell anticipates developing a range of additional IDM-based solutions—focused on storage resource management, data replication, continuous data protection, server virtualization, fixed-content retention and governance, and more—designed to provide further cost-effective, simplified solutions that can help meet IDM goals.

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.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerVault DL2000:
DELL.COM/DL2000

[1] For more information, see "Streamlining Storage with CommVault Deduplication in the Dell PowerVault DL2000," by Jeff Echols, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090305-CommVault.pdf.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

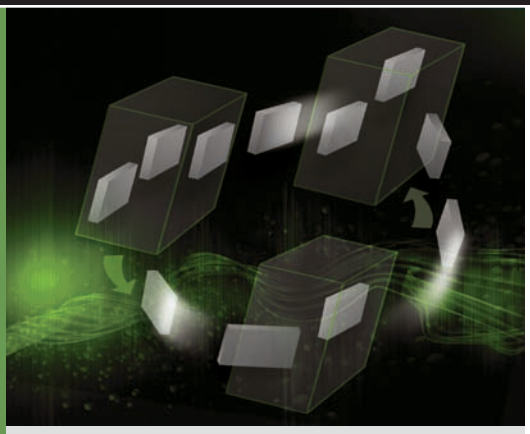

BUILDING THE EFFICIENT ENTERPRISE WITH DELL, INTEL, AND VMWARE vSPHERE 4

By Andrew Gilman
Mike Monthei
Andrew I. Fields

The next generation of virtualization solutions goes beyond just servers to include virtualized storage, networking, and I/O. When IT managers combine VMware® vSphere™ software, 11th-generation Dell™ PowerEdge™ servers with the Intel® Xeon® processor 5500 series, and advanced Dell EqualLogic™ storage arrays, they can create the ultimate in powerful, flexible computing models: an internal cloud.

In today's highly automated, technology-enabled economy, business success has become inextricably linked to IT capability. But IT infrastructures have grown increasingly complex, inflexible, and costly. Enterprises can spend a major part of their IT budgets simply to maintain the status quo—often leaving the business starved for new and updated capabilities.

At the same time, power and cooling costs are catching up to expenditures on hardware and applications to support new users, and many IT managers are facing power and cooling capacity constraints that limit data center growth. The cost of managing IT infrastructures generally far outstrips the cost of initial acquisition. This waste is magnified by poor server utilization, excessive power consumption, and ever-increasing consulting costs to manage the infrastructure. In addition, as enterprises expand to meet evolving business demands and usage models, many data centers are running out of space for server infrastructure.

How can organizations tackle these cost, resource, and time pressures while also ensuring continued innovation and meeting application service-level agreements (SLAs)? With the emergence of standards-based architectures, virtualization, and ubiquitous high-speed connectivity, cloud computing technology is primed to form the basis for a new, highly efficient enterprise IT infrastructure. The cloud computing model drives efficiency through virtualization and economies of scale. Just as virtualization helps to maximize utilization of IT resources, control operational costs, and automate management in enterprise data centers, it also enables massive levels of scalability—helping to reduce administrative burden and optimize performance in business-critical environments.

Why are business and technology strategists so excited about this computing model? Beyond breakthrough scalability, cloud services enable significant cost advantages together with business opportunities—helping enterprises to reduce capital and operational expenses while enabling rapid scale-up and scale-down capacity to enhance business agility and minimize up-front technology investments. Finally, heightened enterprise efficiency enables organizations to shift IT resources from simply keeping the data center running to advancing strategic business goals.

The key to reaching these goals is to take a holistic approach to building an efficient enterprise—which is why Dell, Intel, and VMware have collaborated to establish a vision of the fully virtualized data center that includes not only servers but also storage, networking, and I/O. By virtualizing resources across the data center,

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

enterprises can optimize efficiency and flexibility by creating an internal cloud.

In an economy in which expenditures can be difficult to justify, the question is not whether enterprises can afford to create an internal cloud—it is whether they can afford not to. The operational efficiencies to be gained from comprehensive virtualization are simply too great to ignore. By consolidating the physical computing infrastructure and simplifying management, virtualization helps to reduce capital expenditures, minimize operational expenditures through process and management automation, and limit lost revenue caused by application or infrastructure downtime.

An efficient enterprise model helps achieve these benefits by approaching virtualization holistically. Built on virtualization-optimized hardware and software platforms such as VMware vSphere 4, 11th-generation Dell PowerEdge servers with the Intel Xeon processor 5500 series, and advanced Dell EqualLogic Internet SCSI (iSCSI) storage area network (SAN) arrays, the efficient enterprise model can help IT departments increase the value of both new and existing virtualization investments.

COLLABORATION PROVIDES POWERFUL, INTEGRATED INFRASTRUCTURE

Virtualization is a key way to increase overall enterprise efficiency. Dell offers a comprehensive set of services and prebuilt, validated configurations designed to accelerate virtualization adoption throughout the organization—helping enterprises dramatically lower total cost of ownership while advancing business responsiveness and productivity. Based on extensive field experience and a long-standing partnership with VMware, Dell can also help IT strategists develop an end-to-end approach to virtualization.

Together, Dell and VMware provide a powerful, integrated cloud infrastructure that offers exceptional control and simplicity. By helping to automate the management of resources and SLAs, VMware vSphere 4 helps shift the data center

The 11th-generation Dell PowerEdge server family is optimized for VMware vSphere environments

DELL OFFERS A COMPREHENSIVE RANGE OF VMWARE vSPHERE–OPTIMIZED PLATFORMS

A wide variety of Dell server, storage, and management platforms are optimized to work with VMware vSphere 4 to create powerful, easy-to-administer internal clouds. The 11th-generation Dell PowerEdge server family, for example, is designed for virtualization. These servers feature more memory capacity and integrated I/O than previous-generation PowerEdge servers, in addition to the latest Intel Xeon processors and an embedded VMware ESXi hypervisor.

Dell OpenManage™ systems management software utilizes the vSphere Installation Bundle (VIB) feature to integrate Dell agents into VMware ESXi. This capability allows IT departments to manage, monitor, and update Dell servers running VMware ESXi with a choice of management tools—including the Dell Management Console Powered by Altiris™ from Symantec™.

The Dell EqualLogic PS Series storage family is designed to provide a broad choice of cost-efficient, highly scalable Internet SCSI (iSCSI) storage area network (SAN) arrays that integrate seamlessly with VMware software and help simplify deployment, management, and data protection in virtualized environments. In the SAN, volumes are automatically provisioned as a single scalable pool of storage that leverages all storage resources within the SAN. As workload demands shift or migrate between virtual machines and their physical host servers, the virtualized SAN can automatically apply storage resources where and when they are needed without intervention or manual tuning.

EqualLogic SANs offer an optimized approach to data management within cloud computing environments. EqualLogic PS4000 and EqualLogic PS6000 series SANs incorporate VMware vStorage application programming interface (API) compatibility out of the box—helping to deliver high performance, simplified management, and robust data protection in virtualized environments by offloading performance-hungry operations from the hypervisor to the SAN itself.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

FREDERICK MEMORIAL HOSPITAL USES VMWARE vSPHERE 4 TO TRANSFORM ITSELF INTO AN EFFICIENT ENTERPRISE

Frederick Memorial Hospital (FMH) began providing health services to Maryland residents over 100 years ago. But when the number of servers in the data center grew to the point where maintaining 24/7 uptime became a challenge, the IT team decided to virtualize and consolidate its computing resources to eliminate older servers and boost availability.

Dell ProConsult Services initially helped FMH transform its data center using VMware ESX 3.5 and VMware VirtualCenter 2.5, deploying 106 virtual servers on six Dell PowerEdge R900 servers with the quad-core Intel Xeon processor 7300 series.

Based on the success of this transition, FMH asked Dell to help upgrade the data center to VMware vSphere 4. Today, the IT team at FMH is pleased with the enhanced performance, increased fault tolerance, and dynamic workload allocation capabilities enabled by the vSphere implementation.

"We were able to retire 90 physical servers, and dropped the power consumption in our data center by about 20 percent as a result," says Roy Turner, server systems engineer at FMH. "Beyond consolidation, the failover and load-balancing services in VMware made it easier to maintain uptime around the clock. VMware was the right choice for us to make, and with the new VMware products and Dell's help, we continue to reap the benefits of that decision."

"We were able to retire 90 physical servers, and dropped the power consumption in our data center by about 20 percent as a result."

—Roy Turner
Server systems engineer at
Frederick Memorial Hospital
July 2009

paradigm from infrastructure maintenance to service delivery. Dell server and storage platforms are designed for seamless integration with vSphere to help ensure a high degree of control and administrative simplicity—encompassing management of SLAs as well as the underlying hardware platforms.

Joint Dell and VMware engineering, architectural design, and solution validation efforts can smooth the way for enterprise IT departments to quickly and efficiently deploy a vSphere architecture in production environments. Collaboration between the two companies helps to boost system performance, increase consolidation ratios, and ensure that servers across three generations of Intel processors can be easily combined in the same virtualized server pool. As a result, the platform can help enterprises build a solid foundation for virtualization and the deployment of both internal and external clouds. It can also help avoid the costly and time-consuming trial-and-error process that sometimes occurs when designing and implementing a complex infrastructure.

VMware vSphere 4 and virtualization-optimized Dell platforms are designed to help companies achieve several key objectives. An optimized virtual infrastructure can help reduce space requirements, provide exceptional performance, and boost energy efficiency—all of which help to minimize operating expense and increase return on investment (see the "Dell offers a comprehensive range of VMware vSphere–optimized platforms" sidebar in this article). In addition, because the environment is tightly integrated, the solution is easy to deploy and manage. Furthermore, extensive best practices guidance can help empower IT staff to simplify infrastructure management and refocus on innovation that drives the business.

DELL PLATFORMS DELIVER NEAR-NATIVE VIRTUALIZATION PERFORMANCE

Cloud infrastructures built on Intel Xeon processor 5500 series–based Dell PowerEdge servers can help IT departments implement a high-performance, power-efficient, flexible virtualization platform. When running VMware vSphere 4, these servers can deliver up to 168 percent higher performance in a virtualized environment and achieve more than twice the virtual machine (VM) density compared with previous-generation servers.[1] In addition, they can help organizations achieve near-native performance for software such as enterprise resource planning (ERP) and business intelligence applications running in a virtualized environment.

IT managers should consider upgrading targeted systems to help increase performance, enhance energy efficiency, or simplify administration. vSphere also offers additional benefits to organizations already running VMware software when they refresh hardware platforms—by performing a server refresh and upgrading vSphere at the same time, IT organizations can create a synergistic multiplier effect. Upgrading from vSphere 3.5 to

[1] Based on VMmark benchmark testing performed by Dell Labs in March 2009 on a Dell PowerEdge R710 server running VMware ESX 4.0 (results published April 21, 2009), compared with September 2008 tests on a Dell PowerEdge 2950 III server running VMware ESX 3.5.0 Update 2 (results published September 29, 2008). For details and complete results, visit www.vmware.com/products/vmmark/results.html.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

vSphere 4.0 on existing hardware helps significantly increase consolidation ratios, and combining this upgrade with a hardware refresh can increase those ratios even further.

VMWARE CLOUD OS HELPS SIMPLIFY MANAGEMENT

Designed to provide a highly efficient, simplified computing model, VMware cloud operating systems represent a new category of software that has been specifically developed to manage large collections of virtualized infrastructure components—including processors, storage, and networking—as a seamless, flexible, and dynamic operating environment.

Strategic alignment between Dell, Intel, and VMware makes this seamlessness possible. Joint development processes have resulted in automated, intelligent, real-time management tools such as the Dell Management Console Powered by Altiris from Symantec, Dell EqualLogic PS Group Manager, and VMware vCenter™ software that can dramatically simplify management of virtualized environments and internal clouds. In addition, the close engineering relationship has yielded advanced integration across multiple layers of the infrastructure, including VMware vStorage integration.

A cloud OS manages the complexity of a data center in the same way that a standard OS manages the complexity of an individual server. The cloud OS allows enterprise IT departments running Intel Xeon processor 5500 series–based Dell PowerEdge servers to automatically manage applications to predefined SLAs by enabling high levels of availability, security, and performance for applications. This capability helps companies meet SLA specifications cost-effectively, in a way that requires minimal maintenance. The cloud OS also enables enterprises to run applications on a highly unified, reliable, efficient infrastructure made up of industry-standard components designed for easy replacement. And by moving applications with the same service-level expectations across on- or off-premise computing clouds, IT departments can help reduce total cost of ownership and increase operational efficiency.[2]

VMware vSphere 4 is optimized to run clouds built on PowerEdge servers with Intel Xeon processors (see Figure 1). By combining these platforms in the data center to create a seamless virtualized infrastructure, companies can reap the benefits of a tested virtualization platform as the foundation for internal and external clouds. Federation and standards enable IT professionals to create a secure private cloud and deliver high levels of availability, reliability, scalability, and security. (For an example of how one organization used vSphere 4 to help simplify backup and recovery, see the "Frederick Memorial Hospital uses VMware vSphere 4 to transform itself into an efficient enterprise" sidebar in this article.) The cloud model also helps to facilitate dramatic reductions in capital and operating costs by enabling highly efficient delivery of business services. And the combination of vSphere 4 and virtualization-optimized Dell hardware allows IT departments to retain the flexibility to choose the hardware, operating systems, application stack, and service providers that are best suited to each workload.

Outstanding density and cost efficiency

Joint Intel and VMware engineering and development initiatives have led to multiple enhancements and optimizations for Intel Xeon processors and the vSphere platform. Together, they are designed to deliver outstanding consolidation ratios and exceptional VM density.

When used in conjunction with Intel Xeon processor 5500 series–based PowerEdge servers with Intel Virtualization Technology (Intel VT), vSphere 4 helps to enhance responsiveness for applications in virtualized environments and minimize power and cooling requirements. Joint optimizations enable the VMware hypervisor to fully utilize Intel Hyper-Threading Technology, which is designed to double processing performance for each core and ultimately increase the number of VMs that can run in each socket. Consequently, increased VM density enables IT managers to pack additional computing power into a reduced space—helping to reduce data

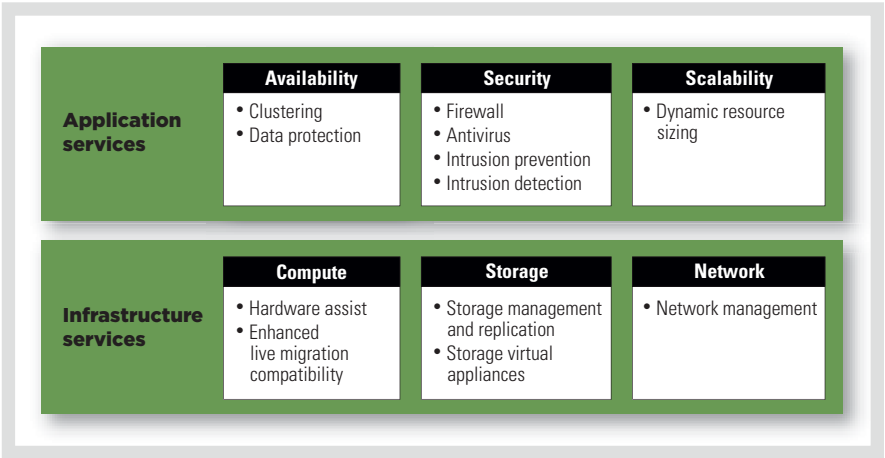

Figure 1. *Combining the VMware vSphere 4 cloud OS with Intel processor–based Dell servers helps create a seamless virtualized infrastructure*

[2] For more information on building foundations for intelligent data centers, see "Computing Pods: Large-Scale Building Blocks for Intelligent, Automated Data Center Deployments," by Timothy Sherbak and Chris Auger, in *Dell Power Solutions*, June 2009, DELL.COM/Downloads/Global/Power/ps2q09-20090238-Sherbak.pdf.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

center space requirements as well as power and cooling expenditures.

High performance

When organizations began attempting to extend the benefits of virtualization to transactional workloads like the Microsoft® Exchange and SQL Server® platforms, they often encountered performance, availability, and security challenges. Virtualization of processors, memory management units (MMUs), and I/O devices often caused an increase in processing overhead and a decrease in the overall performance and scalability of the virtualized environment.

Today, advanced virtualization technologies are making these challenges a thing of the past. With the introduction of the Intel Xeon processor 5500 series, this performance gap is bridged by three key technologies: Extended Page Tables (EPT), which provides hardware support that incorporates MMU virtualization; Intel Virtualization Technology for Connectivity (Intel VT-c); and Intel Virtualization Technology for Directed I/O (Intel VT-d). Together, these technologies can deliver workload transaction latency periods for virtualized environments that are comparable to latencies seen in native environments. These enhancements also help to accelerate network performance and simplify VM migration. Used in conjunction with vSphere 4, virtualization-optimized PowerEdge servers and EqualLogic iSCSI SANs can provide an outstanding platform for virtualizing mission-critical applications in a scalable, secure, high-performance environment. And although providing adequate, cost-effective data protection and quick recovery for tier 1 applications can be challenging within the underlying infrastructure, EqualLogic arrays provide integrated application and VM protection with Auto-Snapshot Manager/VMware Edition and Auto-Snapshot Manager/Microsoft Edition with no additional licensing costs.[3]

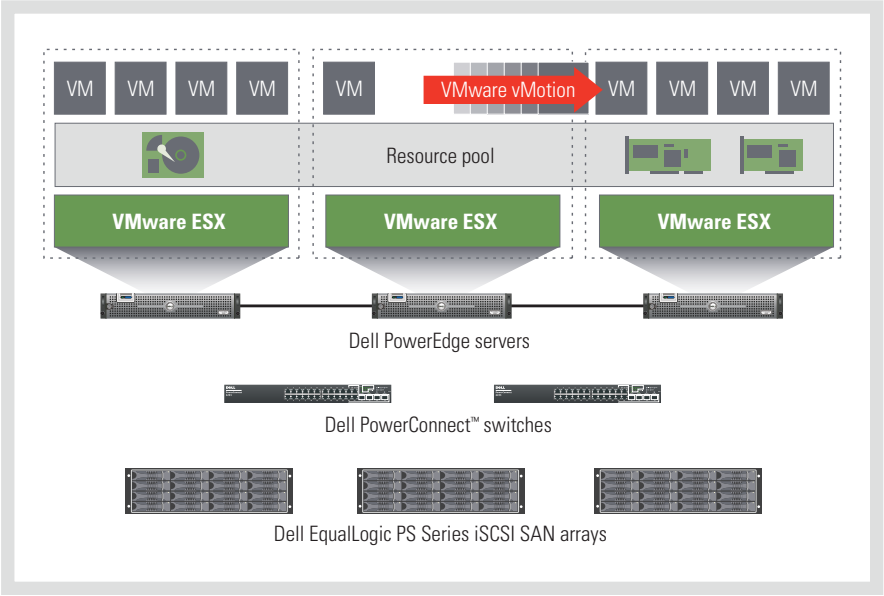

Figure 2. *Virtualized infrastructures using VMware vSphere 4 and Intel processor–based Dell platforms can intelligently allocate resources based on business priorities*

Extensive flexibility

Organizations require flexibility and agility in the data center—but with available space and power capacity dwindling, many enterprises have a limited ability to deploy new applications and respond to changing business conditions. A vSphere 4 virtualized infrastructure built on Intel Xeon processor 5500 series–based PowerEdge servers with Intel VT helps to address these issues in several ways.

First, an efficient enterprise built with these components is optimized to intelligently allocate computing resources based on business priorities (see Figure 2). VMware Distributed Resource Scheduler (DRS) running on Intel Xeon processor 5500 series–based PowerEdge servers can continuously monitor utilization across data center resource pools and intelligently allocate VMs among available compute resources according to business needs. In this manner, data centers built on the efficient enterprise model also help to deliver outstanding virtualization performance and utilization of legacy 32-bit guest applications running in VMs.

The virtualized infrastructure also allows IT organizations to easily respond to new opportunities and changing business priorities, and to preserve flexibility and choice in managing and allocating virtualized workloads across new and existing servers. Intel VT FlexMigration allows IT administrators to expand the pool of resources in the virtualized environment, further enhancing agility and cost efficiency.

The Dell and VMware model for the fully virtualized data center also provides a tested, enterprise-ready live migration solution to enable data center agility and evolving usage models. The VMware Enhanced vMotion™ Compatibility feature works with Intel VT FlexMigration to enable IT organizations to move VM workloads from one physical server to another on the fly—avoiding planned downtime and allowing organizations to readjust loads quickly based on changing requirements. These tools are the key to expanding compatibility pools, which helps to provide investment protection and simplify upgrade paths.

Dell and VMware platforms can further help to ensure cost-effective high availability

[3] For 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.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

by automatically restarting VMs on servers that have spare capacity. VMware High Availability (HA) tools and Intel VT FlexMigration help organizations to minimize downtime and IT service disruption while avoiding the need for dedicated standby hardware. In addition, VMware Fault Tolerance features allow organizations to take advantage of Intel VT FlexMigration to provide near-instantaneous failover and continuous availability for applications.

Groundbreaking power efficiency

Data centers in the United States account for a significant portion of the country's total consumption of electricity, and in many enterprises power and cooling costs actually outstrip hardware expenditures. Using Intel Xeon processor 5500 series–based PowerEdge servers running vSphere 4, organizations can enhance performance while consuming less energy and space than traditional data centers configured with all-physical servers—both of which contribute to reduced operating costs.

The vSphere infrastructure enhances both individual server efficiency and overall data center efficiency through innovative Intel Xeon processor 5500 series design in PowerEdge servers. Intel Intelligent Power Technology is designed to automatically step down individual cores to near-zero power independent of other operating cores, helping reduce energy costs. Integrated power gates power down unused cores during low-use periods, and VMware Distributed Power Management (DPM) features can turn off entire servers when there is unneeded capacity. Servers are powered back up when their capacity is needed—allowing data centers to power up and down quickly and efficiently.

DELL BUSINESS READY CONFIGURATIONS STREAMLINE DEPLOYMENT

Dell service offerings help organizations plan for the gamut of virtualized infrastructure implementations. To help enterprises simplify virtualization deployments, Dell offers engineer-tested, pre-built virtualization configurations based on VMware vSphere 4. These Dell Business Ready Configurations can be deployed as designed or modified to meet individual requirements. They also help to simplify design, ordering, and deployment of virtualized infrastructures while helping ensure adherence to best practices recommendations and guidelines.

For enterprises that require broad guidance in designing an overall virtualization strategy, the Dell ProConsult Services team can help IT departments develop a comprehensive approach to virtualization that encompasses people, processes, and technology. The result: action-oriented plans that deliver high-impact, short-duration projects with predictable and measurable outcomes. Dell ProManage™ Managed Services complements these capabilities by using combinations of automated technology, cloud delivery, and Dell services infrastructure to enhance configuration, deployment, and management of end-user and data center environments.

END-TO-END CLOUD INFRASTRUCTURE HELPS BUILD THE EFFICIENT ENTERPRISE

As companies look for innovative ways to cut costs and exploit opportunities, building an end-to-end cloud infrastructure as the basis for an efficient enterprise is an outstanding place to start. IT departments should take advantage of the current buyer's market to strategically position themselves for agility and growth—and Dell, Intel, and VMware can provide the virtualized infrastructure and cloud OS that make this effort possible.

Dell, Intel, and VMware plan to continue joint technology development and optimization. This collaboration is expected to continue providing enterprises with clear upgrade paths including expanded VMware vStorage integration, expanded Intel integration with Dell and VMware platforms, and expanded Dell Business Ready Configurations. Organizations that have installed the Dell EqualLogic PS Series iSCSI SAN architecture can also obtain complimentary upgrades and enhancements because Dell does not charge licensing fees.

By accelerating planned investments in virtualization now, organizations stand to immediately reduce operational expenditures while also laying a strong foundation for growth. VMware vSphere 4 is designed to deliver unparalleled efficiency, control, and choice, and vSphere cloud computing infrastructures built with virtualization-optimized, Intel Xeon processor 5500 series–based Dell PowerEdge servers and Dell EqualLogic storage provide the flexibility and agility required for enterprise data centers—both today and in the future.

Andrew Gilman is a solutions marketing manager at Dell responsible for virtualization marketing activities.

Mike Monthei is a global marketing manager at VMware responsible for marketing activities for the Dell alliance.

Andrew I. Fields is an alliance marketing manager at Intel responsible for collaborative marketing programs.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and VMware:
DELL.COM/VMware
www.vmware.com/dell

Dell virtualization solutions:
DELL.COM/Virtualization

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

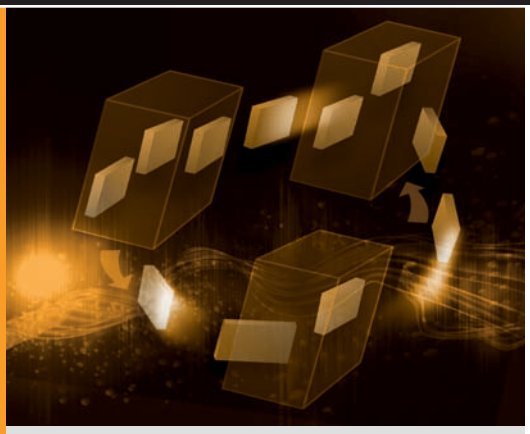

HIGH-AVAILABILITY VIRTUALIZATION WITH DELL EQUALLOGIC ARRAYS AND MICROSOFT WINDOWS SERVER 2008 R2 HYPER-V

By Stanley L. Stevens
Darren W. Miller
Jim Schwartz

The Microsoft® Windows Server® 2008 Release 2 (R2) Hyper-V™ platform introduces capabilities designed to deliver high availability for virtualized data centers. Integrating a Dell™ EqualLogic™ PS Series storage area network into Hyper-V virtualized environments helps organizations capitalize on new Hyper-V capabilities, including live migration, hot-add storage, and storage platform features that enhance usability and performance.

For enterprises in a broad range of fields, virtualization has become a popular IT strategy for increasing hardware utilization, enhancing flexibility, and decreasing power, cooling, and data center footprint costs. But while virtualization can offer these and other benefits, running multiple applications on each physical server can place application availability at risk: if a physical server fails, virtualized applications running on that server could become unavailable. In many situations, even temporary downtime for applications can impair worker productivity and jeopardize revenue. At a university, for example, a network outage could prevent administrators from viewing student records or researchers from working with data sets. At a hospital, the inability to access patient records or medical imaging could delay treatment.

Microsoft Windows Server 2008 Release 2 (R2) with Hyper-V virtualization technology introduces capabilities designed to increase application availability and avoid downtime in virtualized environments. Windows Server 2008 natively supports Internet SCSI (iSCSI) and Fibre Channel storage area networks (SANs). Of the two, iSCSI has a natural affinity for virtualization, and the Microsoft iSCSI Software Initiator available in Windows Server 2008 and Windows Server 2008 R2 is designed to provide ubiquitous SAN connectivity for organizations using existing Ethernet infrastructure and leveraging existing networking skills and tools.

Organizations can design an infrastructure supporting high availability and a robust, reliable shared-storage solution with built-in redundancy and tight integration of virtualization management capabilities. The Microsoft Multipath I/O (MPIO) framework helps provide high availability and dynamic load balancing for SAN devices through redundant network or fabric connections. MPIO is designed to dynamically route I/O to the best path and helps protect against connection-point failures between a Hyper-V host and shared storage, including network interface cards and adapters, switches, and array ports.

The Dell EqualLogic Multipath I/O device-specific module works with MPIO to provide automated management and enhanced end-to-end load balancing for EqualLogic SAN volumes. EqualLogic PS Series iSCSI SAN arrays offer a highly automated shared-storage foundation for virtualized data centers that enables organizations to take full advantage of the high-availability features incorporated into Windows Server 2008 R2 Hyper-V.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

BUILDING A VIRTUALIZED ENVIRONMENT

Microsoft helped bring virtualization to a wide audience when it introduced Windows Server 2008 Hyper-V. This platform offers virtualization capabilities within a cost-effective and familiar Windows® OS–based environment that can extend from the desktop to the data center. In addition, the Microsoft System Center management suite enables organizations to access a comprehensive array of management tools to help simplify administration of virtualized environments. Windows Server 2008 R2 extends the capabilities provided in Windows Server 2008, and Dell servers and storage arrays are designed to take advantage of these new features.

For environments virtualized with Hyper-V technology, 11th-generation Dell PowerEdge™ servers can provide a particularly robust and reliable hardware foundation. Featuring the Intel® Xeon® processor 5500 series architecture and designed to provide more integrated I/O than previous-generation PowerEdge servers, 11th-generation PowerEdge servers enable the processing performance, memory bandwidth, and I/O capabilities required to host numerous virtualized applications on each physical server. Because PowerEdge servers incorporate multiple redundant components, they help to deliver high availability for virtualized environments by avoiding unplanned outages.

PowerEdge servers also help simplify the deployment of virtualized server environments. This latest generation of servers includes a Hyper-V hypervisor so administrators do not have to download and install additional virtualization software. In addition, the Intel Virtualization Technology (Intel VT) FlexMigration capabilities built into the processors allow organizations to integrate 11th-generation PowerEdge servers featuring the latest processing architecture into an environment with previous generations of Intel Xeon processor–based PowerEdge servers. In this way, FlexMigration helps simplify new deployments and protect existing investments.

Windows Server 2008 Hyper-V technology and PowerEdge servers help ensure high availability while also providing high performance, outstanding scalability, and simplified management. Additionally, failover clustering can help provide redundancy in the event of a server outage. Hyper-V technology allows clustering to operate from the parent partition of the OS and from within individual virtual machines (VMs) (with iSCSI direct). Connecting to iSCSI disks from the Hyper-V parent partition is designed to work just as it does from a physical host, including leveraging all existing tools and features. On the parent partition, MPIO provides high availability to shared storage. In this configuration, the Hyper-V host can create Microsoft virtual hard disk (VHD) files for each VM on iSCSI-connected disks or create pass-through disks that are individually owned by the VM. Administrators can use the quick migration and live migration features to move VMs from one physical server to another in the event of a planned or unplanned outage.

In addition, iSCSI storage disks can be connected from within the VM itself. Support for guest clustering requires running the Microsoft iSCSI Software Initiator within the VM to allow services and features to run in the application context within the virtual environment as they do in a physical server environment with full visibility by the VM—helping provide for the transparent operation of all functions because applications can instantiate common workflows and data movement from within the VM. These services and features include the cluster service, MPIO, Microsoft Volume Shadow Copy Service (VSS), hardware providers, and Microsoft Virtual Disk Service (VDS) hardware providers.

To take advantage of clustering and the new high-availability capabilities available with Windows Server 2008 R2 Hyper-V, organizations must use shared common storage with high-performance disks. In many cases, an iSCSI SAN can be a suitable and cost-effective approach, because iSCSI connectivity helps simplify management compared with Fibre Channel systems.

Dell EqualLogic PS Series SAN arrays offer robust storage in Microsoft Windows Server 2008 R2 Hyper-V environments

INCORPORATING DELL EQUALLOGIC ARRAYS INTO THE VIRTUAL DATA CENTER

Designed as high-availability virtualized storage solutions, Dell EqualLogic PS Series iSCSI SAN arrays can be an excellent fit for data centers virtualized with Hyper-V technology. Storage is virtualized in EqualLogic PS Series arrays, and data volumes are provisioned automatically from a single scalable storage pool. The EqualLogic peer storage architecture enables the arrays to share resources, evenly distribute workloads, and provide data protection for the VMs in virtualized environments. Resources are applied automatically even as VMs and their workloads change.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

EqualLogic PS Series SANs also enhance application availability and help protect data. They feature redundant, hot-swappable components, including power supplies, controllers, enclosures, and disk drives, plus disk failover capabilities to help ensure high availability. Cloning and replication capabilities are also included in the firmware of EqualLogic PS Series SANs, offering advanced data protection for virtualized environments.

EqualLogic and Microsoft began working together in 2002 during the early stages of iSCSI implementation, and the 2003 releases of the Microsoft iSCSI Software Initiator and the first EqualLogic PS Series array helped a broad range of organizations to begin deploying SANs in their data centers. Additionally, Windows storage frameworks such as MPIO, VSS, and VDS offered tight integration and scalability, automated backups, and ease of management for organizations deploying Microsoft Windows Server and EqualLogic SANs.

The EqualLogic PS Series is particularly well suited for supporting Microsoft Windows Server 2008 environments. In fact, Microsoft has certified the EqualLogic PS Series for Windows Server 2008 x64 Editions, indicating that it meets Microsoft standards for compatibility and best practices. The EqualLogic PS Series has also earned the Microsoft Simple SAN designation, which means that it meets a range of ease-of-use, ease-of-installation, and ease-of-management criteria.

Joint Dell and Microsoft efforts to increase IT efficiency and ensure high availability have streamlined the process of administering the entire environment using Microsoft System Center together with the Dell OpenManage™ systems management suite. This collaboration has yielded several capabilities that help simplify IT when using EqualLogic PS Series SANs with Windows Server 2008 Hyper-V. Dell offers the EqualLogic Auto-Snapshot Manager/Microsoft Edition tool to help simplify management and deliver high availability by providing fast online backups and rapid restores for Hyper-V VMs as well as NT File System (NTFS), Microsoft SQL Server®, and Microsoft Exchange data. For Hyper-V, Auto-Snapshot Manager/Microsoft Edition provides SmartCopy snapshots of VMs and VHDs. Those snapshots can be restored to an EqualLogic SAN on a different remote or local server for new application usage or test and development when the application supports it.

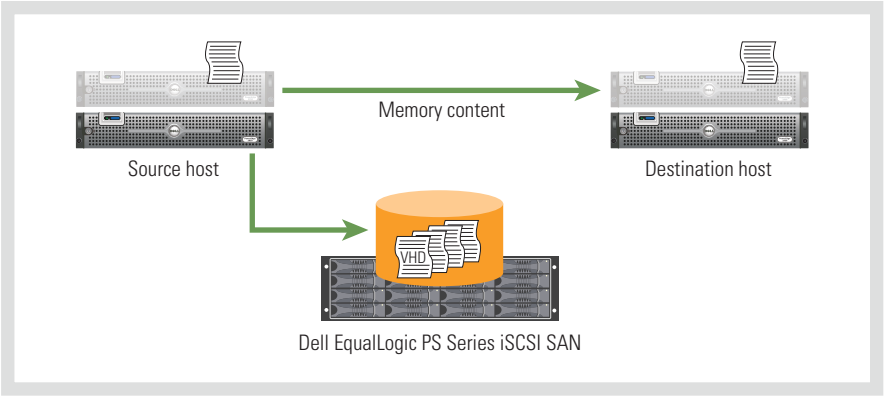

Figure 1. *Live migration helps accommodate planned outages without downtime*

INNOVATING SHARED STORAGE PLATFORMS

Several new features in Microsoft Windows Server 2008 R2 are designed to enhance user experience and performance with shared storage to benefit virtualized environments. The Quick Connect feature allows a single-click connection from a system using the Microsoft iSCSI Software Initiator to iSCSI storage. The iSCSI user interface is also now included in Windows Server 2008 Server Core installations, which provides the option of configuring through a lightweight graphical user interface or a command line. Microsoft iSCSI Software Initiator performance has been optimized for 10 Gigabit Ethernet networks to help deliver wire-speed performance with reduced processor utilization. Additionally, Hyper-V includes support for jumbo frames and VM queues (VMQs), which helps increase iSCSI direct throughput performance in Hyper-V for demanding VM workloads. In addition, as diskless servers continue to emerge, a growing number of organizations are booting virtualized servers remotely—and Windows Server 2008 R2 includes enhanced support for iSCSI boot that provides up to 32 redundant paths and integrated support with Windows setup.

Storage resource monitoring features in Windows Server 2008 R2 enable IT administrators to efficiently diagnose and gather information on storage I/O performance and storage path health, and to take action to move workloads or VMs to other Hyper-V hosts to help optimize performance and avoid downtime. Storage configuration reporting gathers and exports SAN component information, including the number of disks connected and multipath parameters. MPIO data center automation allows administrators to preconfigure multipath settings before connecting storage to help automate storage deployment.

The Cluster Shared Volumes (CSV) capability—part of the failover cluster feature in Windows Server 2008 R2—enables access to the same logical unit (LUN) on a SAN for Windows Server 2008 R2 Hyper-V–based systems. CSV enables administrators to map multiple VMs to a single volume or LUN, supporting live migration of VMs from one Hyper-V host to another.

MANAGING VMs WITH LIVE MIGRATION

The live migration capability of Microsoft Windows Server 2008 R2 Hyper-V enables IT administrators to move running VMs from one physical host to another without disrupting service or causing downtime. Live migration helps administrators ensure

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

high availability while performing maintenance on physical servers. Administrators can also create an increasingly dynamic data center by using live migration to help with load balancing.

Live migration complements quick migration—an existing capability of Windows Server 2008 Hyper-V. Quick migration was designed to save, move, and restore VMs to help minimize the impact of downtime during unplanned outages; live migration provides an easy way to avoid disruptions during planned maintenance.

To optimize the agility gained from live migration, organizations need to use shared storage such as a Dell EqualLogic PS Series SAN. Live migration allows VM memory pages to be transferred from the source Hyper-V physical host to the destination physical host while the source VM draws from the shared storage of the SAN (see Figure 1). Modifications to the VM memory pages during this process are tracked and transferred to the destination server. Hyper-V then transfers control of the storage for the VHD files of the VM to the destination physical server and brings the destination VM online on the destination physical server. Windows Server 2008 R2 Hyper-V allows administrators to queue multiple live migrations without having to keep track of other migrations within the cluster.

SCALING STORAGE DYNAMICALLY WITH HOT-ADD STORAGE

Working together, Microsoft Windows Server 2008 R2 Hyper-V and Dell EqualLogic PS Series arrays can also provide seamless scalability for virtualized environments. Traditionally, administrators needed to take host servers offline to upgrade storage. Windows Server 2008 R2 Hyper-V is designed to allow administrators to add and remove VHDs and pass-through disks in the existing SCSI controller for a running VM without requiring downtime. As a result, administrators can respond quickly to the changing storage requirements of VMs without causing interruptions. They can also deploy mission-critical applications that might call for fast storage growth and require high availability.

EqualLogic PS Series SANs offer an exceptional fit for this scalable environment. The modular design of the EqualLogic PS Series helps make adding or removing disk capacity simple. Arrays can be added without disrupting operations. By combining the Hyper-V hot-add storage capabilities with the EqualLogic PS Series SAN, administrators can create a dynamic environment in which they can quickly scale their storage infrastructure or reconfigure VMs to meet changing requirements.

BRINGING HIGH AVAILABILITY TO THE VIRTUAL DATA CENTER

To help avoid the risk of running multiple applications on a reduced number of physical servers, IT administrators must implement multiple high-availability strategies into their virtual data center environments. Microsoft Windows Server 2008 R2 Hyper-V introduces several new features that help ensure high availability in virtualized environments. Designed to provide virtualized storage and to protect data with a wealth of redundant components, Dell EqualLogic PS Series iSCSI SAN arrays can help organizations capitalize on the new capabilities of Hyper-V to help maximize availability.

Stanley L. Stevens is a virtualization solutions marketing manager at Dell and a 16-year veteran of the IT industry.

Darren W. Miller is a product consultant at Dell and has worked in the storage industry for over 12 years.

Jim Schwartz is director of virtualization solutions for Microsoft.

MORE ONLINE

DELL.COM/PowerSolutions

QUICK LINKS

Dell EqualLogic PS Series:
DELL.COM/EqualLogic
DELL.COM/PSSeries
DELL.COM/PSSeries/Hyper-V

Microsoft Windows Server 2008:
www.microsoft.com/windowsserver2008

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Momentus® 7200-RPM Drives—
Smart, Green and Built for the Future

Seagate®

The Freedom of Mobility With Exceptional Performance and Staggering Capacity

- Up to 500 GB of pure, digital creativity in your hands
- 7200-RPM spin speed enables the productivity you demand
- Energy-efficient power gives you the freedom to go mobile anytime
- Innovative features: Self-Encrypting Drives, free-fall sensor

www.seagate.com/www/en-us/products/laptops/momentus/

Copyright © 2009 Seagate Technology LLC. All rights reserved. Seagate, Seagate Technology, the Wave logo and Momentus 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 gigabyte, or GB, equals one billion bytes and 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.

HIGH VALUE, HIGH PERFORMANCE: 11TH-GENERATION DELL POWEREDGE SERVERS FOR SMBs

By Nick Jervis
Hariharan Ramalingam

The high-value, high-performance Dell™ PowerEdge™ R410 and PowerEdge T410 servers are designed to meet the electrical, mechanical, thermal, and acoustic standards of the 11th-generation PowerEdge server family while incorporating cost-effective components and features to help meet the needs of small and medium businesses (SMBs).

When defining the next generation of servers, Dell starts by creating a set of consistent behaviors and specifications across all servers in that generation, and then differentiates specific models to help meet the needs of different usage models and applications based on feedback from Dell customers. The features that define the 11th-generation Dell PowerEdge server family—from high-end, mainstream-class two- and four-socket servers to cost-optimized one- and two-socket servers—include both obvious elements such as the Intel® Xeon® processor 5500 series, Double Data Rate 3 (DDR3) memory, disk drives, power supply units (PSUs), and industrial design as well as less obvious technical implementations. Among the features common to 11th-generation PowerEdge servers are the following:

- **Energy efficiency:** PSUs that meet 80 PLUS and Climate Savers Computing standards,[1] efficient voltage regulators, efficient DC-to-DC components, zoned cooling, and so on
- **BIOS:** PowerEdge BIOS core, diagnostics, and factory customization options
- **Management:** Baseboard management controller (BMC) components and integration with the Dell OpenManage™ suite of tools for deployment and life cycle management, including the Unified Server Configurator enabled by the Lifecycle Controller standard for 11th-generation PowerEdge servers
- **Electrical specifications:** Circuit board design and layout guidelines, including network components and backplanes designed to support speeds of up to 6 Gbps, as well as component derating to help fulfill reliability and high quality demands
- **Validation:** Regulatory, reliability, thermal, acoustic, and software compatibility objectives to help ensure design and quality goals are met
- **Mechanical specifications:** Chassis and chassis sub-assemblies conforming to portfolio behavioral guidelines and standards as well as industrial design and usability requirements

The high-value, high-performance PowerEdge R410 and PowerEdge T410 server models are designed to meet the stringent electrical, mechanical, functional, thermal, and acoustic standards of the 11th-generation PowerEdge server family while incorporating cost-effective configuration options to help

[1] For more information, visit www.80plus.org and www.climatesaverscomputing.org.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Dell PowerEdge R410 rack server in hot-plug chassis configuration with LCD

meet the specific requirements of small and medium businesses (SMBs) as well as organizations deploying high-performance computing (HPC) clusters. By focusing on features, components, and other criteria designed for optimal value, while still offering optional upgrades for organizations that need them, these servers can help meet a variety of needs in SMB and HPC cluster environments.

DESIGNING FOR VALUE

In designing the high-value, high-performance Dell PowerEdge R410 and PowerEdge T410 servers, Dell took advantage of SMB feedback, historical purchasing data, and analysis of technology and industry trends to define appropriate specifications for the majority of SMBs. For application or organizational usage models that require features not typically incorporated in this value-oriented class of server, the design commonalities and consistency across the PowerEdge family enable rapid integration and qualification.

The key to creating cost-effective servers—while still incorporating some of the same key design tenets as high-end models—lies in understanding how specific components can affect cost and complexity in the rest of the system. For example, each memory slot, hard disk drive, PCI slot, rear or front I/O feature, and so on requires incorporating the connectors as well as the associated electrical components, motherboard space, reserve power capacity, cooling capacity, and associated mechanical structures. Similarly, each hard drive bay requires many of those same components as well as backplane board space for hot-pluggable configurations and Serial ATA (SATA) cables for cabled configurations. And each integrated LAN on Motherboard (LOM) port or pair of ports requires supporting electrical components and motherboard space as well.

PowerEdge R410 and PowerEdge T410 servers are designed to help meet the needs of SMBs in a cost-effective way through base configurations that incorporate the minimum feature set, including the following:

- **Cabled PSUs:** Because not all organizations or applications require redundant, highly available PSUs, the base configurations include an efficient cabled PSU. Redundant, hot-pluggable PSUs are available as an optional upgrade (two redundant, hot-pluggable PSUs and the power distribution board).
- **Cabled hard drives:** Using cabled hard drives in the base configurations provides a value-oriented solution that helps avoid the need for a hard drive backplane and additional cabling. Hot-pluggable hard drives are available as an optional upgrade.
- **BMC management:** All 11th-generation Dell PowerEdge servers use a common BMC architecture integrated on the motherboard for entry-level BMC features. The additional management features offered in the Integrated Dell Remote Access Controller 6 (iDRAC 6) Express and iDRAC 6 Enterprise are available as optional upgrades after purchase.
- **LED indicators:** For cost-optimized configurations with cabled hard drives, a four-LED interface is standard, with the combination of lit LEDs indicating the server status. Systems with hot-pluggable hard drives include an LCD for advanced status and fault indication and diagnosis.

The Intel Xeon processor 5500 series in 11th-generation PowerEdge servers is well suited for SMBs seeking high performance and energy efficiency in a cost-effective system that can scale and virtualize in the future. Compared with existing servers using single-core processors, these intelligent processors can enable up to nine times higher performance per server, up to 90 percent lower operating costs, and an estimated eight-month payback on investment.[2] Key features include Intel Turbo Boost Technology, Intel Hyper-Threading Technology, 8 MB shared level 3 (L3) cache with Enhanced Smart Cache, Intel QuickPath Technology, Intel Intelligent Power Technology, and Intel Virtualization Technology.

In addition, like other 11th-generation PowerEdge servers, the PowerEdge R410

Dell PowerEdge R410 rack server in cabled chassis configuration with LED

[2] Based on March 2009 Intel comparison replacing nine four-year-old servers configured with single-core Intel Xeon processors at 3.8 GHz with 2M cache with one server configured with Intel Xeon X5570 processors. Results have been estimated based on internal Intel analysis and are provided for information purposes only.

 Copyright © 2009 Dell Inc. All rights reserved.

and PowerEdge T410 incorporate the Lifecycle Controller—an embedded flash chip containing systems management components such as the system BIOS, firmware, drivers, and Dell OpenManage tools that can function independently of both media and platform OS. The Dell Unified Server Configurator offers a single simplified interface to this controller to help administrators perform tasks such as firmware updates, OS deployment, and diagnostics.[3]

VALIDATING TO MEET HIGH STANDARDS

Dell validates its servers using a common test methodology that incorporates a test case library, specification, and process. Dell PowerEdge R410 and PowerEdge T410 servers are comprehensively tested in minimum, typical, and maximum configurations under a variety of environmental conditions to help ensure compliance with stringent 11th-generation PowerEdge server standards across electrical, mechanical, thermal, and acoustic design criteria. Design quality is verified to high Dell standards before the servers begin shipping.

Dell PowerEdge T410 tower server in cabled chassis configuration with LED

Electrical design

PowerEdge R410 and PowerEdge T410 electrical designs focus on meeting typical SMB requirements for two-socket rack and tower servers, and Dell selects the electrical components to help maintain quality, performance, robustness, reliability, and compatibility. The servers are also designed to maximize cost-effectiveness, including optimizing the size of the motherboard. In printed circuit board manufacturing, virtually any size or shape can be created to meet the needs of a given system, but a few commonly used standard sizes help maximize the number of boards cut from the large panel used in the manufacturing process. In addition, designing the system to have separate boards for the electrical circuitry—thereby keeping the motherboard size small—helps support optional features such as the hot-pluggable hard drive backplane and redundant PSU power distribution board. This approach helps reduce costs for configurations without these features.

Mechanical design

PowerEdge R410 and PowerEdge T410 mechanical designs focus on simplification and high value. Organizations can choose from two basic configurations for each model: one with cabled hard drives and an LED, and one with hot-pluggable hard drives and an LCD. Both configurations can incorporate cabled or redundant PSUs.

Dell PowerEdge T410 tower server in hot-plug chassis configuration with LCD

Like mainstream-class PowerEdge servers, PowerEdge R410 and PowerEdge T410 servers are designed for ease of use and simplified servicing. Serviceable parts—including the optical disk drive, system fan, extension cards, planar, backplane, power distribution board,

> **"Dell PowerEdge R410 and PowerEdge T410 servers are comprehensively tested to help ensure compliance with stringent 11th-generation PowerEdge server standards."**

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

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

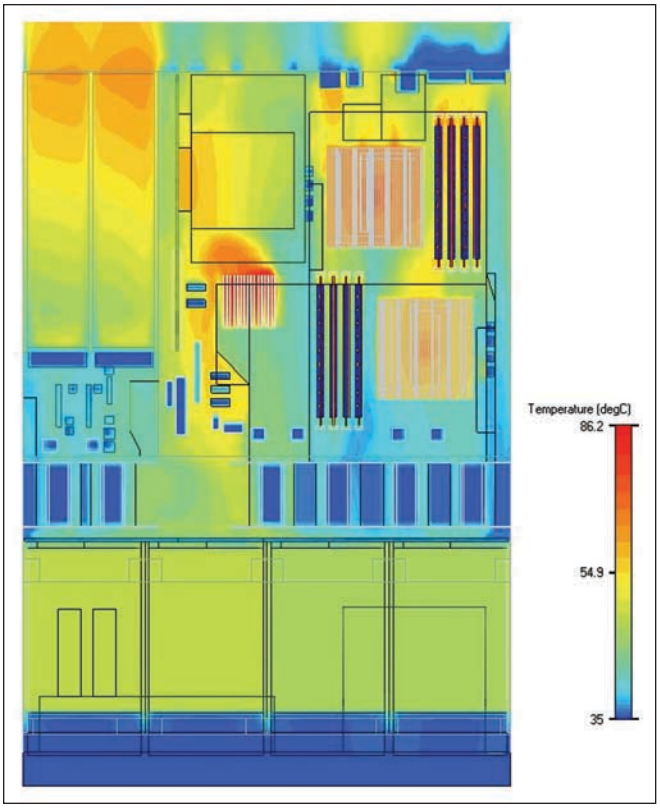

Figure 1. *Dell PowerEdge R410 thermal profile*

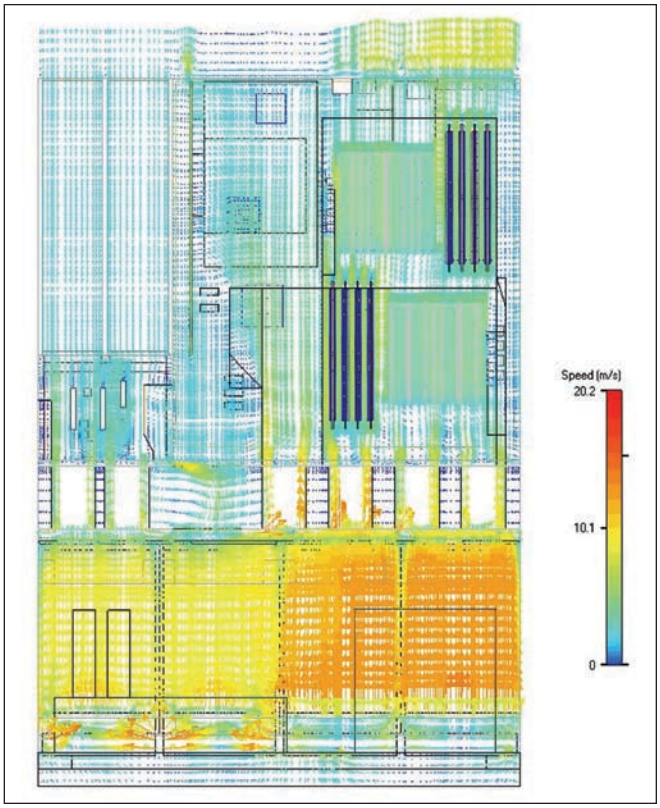

Figure 2. *Dell PowerEdge R410 airflow vector profiles*

redundant PSUs, and so on—use a tool-less implementation when possible. Other key differences from mainstream-class models that help optimize value include not supporting rack mounting for the PowerEdge T410 tower server, a non-rotational Dell logo on the bezel, a non-removable cosmetic top cover, a reduced number of standoffs on the chassis for sliding rails, a one-piece plastic fan shroud with a molded Dell logo engraved into the plastic instead of a separate logo, a heat sink mounted with screws, a simplified design for the tower cover latch on the PowerEdge T410, a simplified but still robust rack ear on the PowerEdge R410, and simplified removable card-mounting features. Mechanical validation involves performing a range of testing to help ensure high levels of reliability; tolerance for shock and vibration (including rotational vibration); ease of insertion and removal from the system for mechanical components, boards, and commodities; and so on.

Thermal and acoustic design

PowerEdge servers share a set of thermal and acoustic requirements, and the thermal capabilities and reliability of the PowerEdge R410 and PowerEdge T410 servers are designed to be comparable to those of the mainstream-class PowerEdge R610 and PowerEdge T610 servers even though the PowerEdge R410 and PowerEdge T410 require less total thermal capacity than those servers.

Figures 1 and 3 show the PowerEdge R410 and PowerEdge T410 thermal profiles in the maximum configuration with minimum airflow; even under these worst-case conditions, the systems are cooled sufficiently. Figures 2 and 4 show the PowerEdge R410 and PowerEdge T410 airflow vector profiles under typical ambient conditions; fan speed is optimized automatically by the BMC based on temperature readings to help ensure sufficient airflow to all components in the system.

Comprehensive testing and evaluation during the design and validation phases help the PowerEdge R410 and PowerEdge T410 meet 11th-generation PowerEdge thermal and acoustic design criteria, including the following:

- **Worst-case testing:** Both the theoretical and actual system validation are performed using "worst-case devices"—those that potentially have the greatest risk. These devices include systems with the maximum supported configurations, including the fastest processors, memory with the maximum speed and capacity, the maximum number of hard drives with the highest speed and capacity, and fully loaded PCI slots, including RAID, network, and other supported devices.
- **Thermal simulation:** Thermal simulation analyzes the system's hot spots at typical and maximum configurations and with integrated and anticipated add-on devices to help ensure compliance with thermal specifications.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

- **Fan-speed analysis:** The FloTHERM computational fluid dynamics (CFD) tool from Mentor Graphics helps perform detailed analysis of fan speeds and temperatures. Iterative adjustments are made to the fan duty cycle, with fan speed adjusted based on different inlet temperatures. Thermal sensors monitor these temperatures to help ensure that the fan curve is optimized across the range of supported configurations at different temperatures.
- **Acoustic specifications:** Dell research on enterprise office and data center environments as well as industry research helped develop an internal acoustic specification. Measurements of a number of acoustic criteria—including fan speed optimized to environmental conditions to help lower noise levels, reduced vibration noise from mechanical parts, and optimized PSU fan noise levels—during the system validation help ensure conformance with this specification.

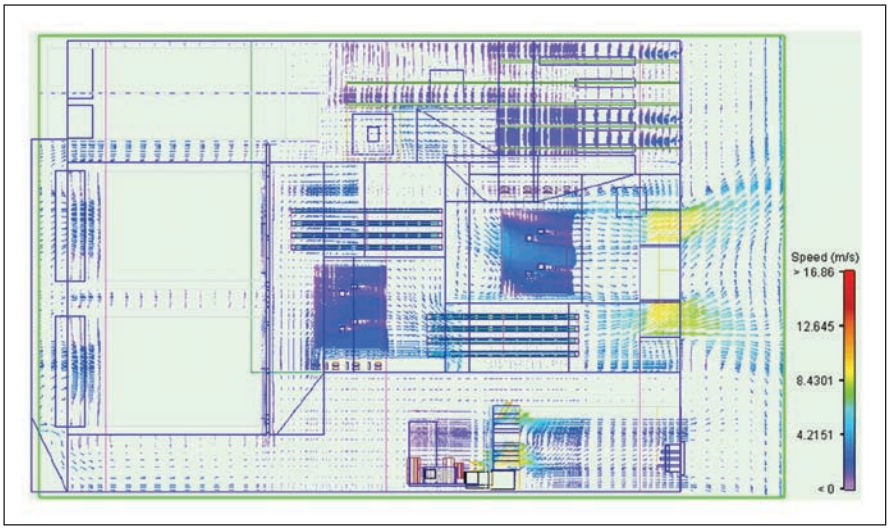

Figure 4. *Dell PowerEdge T410 airflow vector profiles*

Both the PowerEdge R410 and PowerEdge T410 servers use passive aluminum-base heat sinks with embedded heat pipes for the processor cooling; the servers draw cool air in through the front and vent air through the rear. During system validation, custom algorithms are created to help ensure both adequate cooling across the supported temperature range as well as conformance with Dell acoustic specifications.

MEETING THE NEEDS OF SMBs

As part of the 11th-generation Dell PowerEdge server family, the high-value, high-performance PowerEdge R410 and PowerEdge T410 include a variety of features common across that generation of servers—including the powerful Intel Xeon processor 5500 series; simplified management; high levels of efficiency, reliability, and performance; and conformance with Dell design standards. By focusing on value-optimized components while still making advanced features available as upgrades, these servers can help meet the specific needs of SMBs in a cost-effective way.

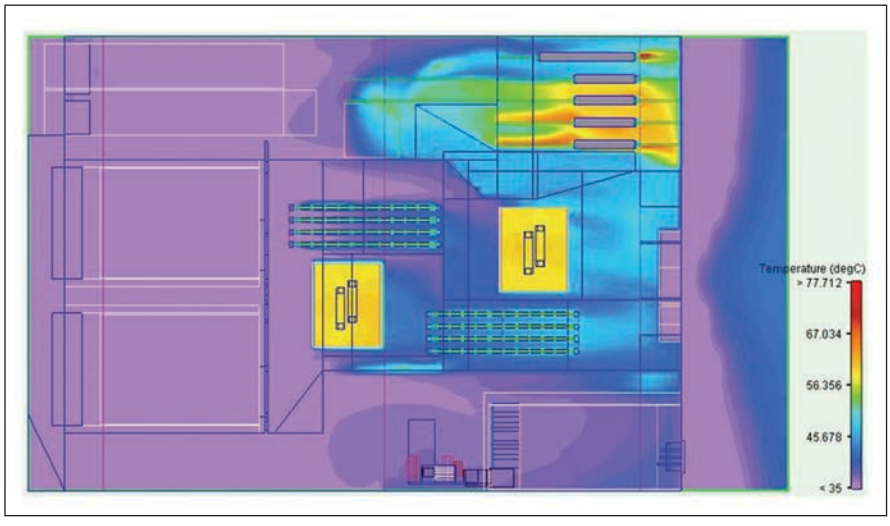

Figure 3. *Dell PowerEdge T410 thermal profile*

Nick Jervis is a senior development manager working for the Dell Taiwan Design Centre Product Development team, and previously held electrical engineering, system engineering, and advanced engineering leadership roles. He has an M.B.A. in Engineering Management and an M.S. in Electronics Systems Design.

Hariharan Ramalingam is a development manager working for the Dell Taiwan Design Centre Product Development team. He has a postgraduate degree in Electronics and Communication Technology and a postgraduate diploma in Management Studies.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell PowerEdge servers:
DELL.COM/PowerEdge

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

There is nothing more stimulating than great value.

Dell™ PowerEdge™ Servers

Extreme Networks® Ethernet Switches

Dell EqualLogic™ Storage Array

Dell and Extreme Networks

have changed the economics of iSCSI to create an affordable and accessible solution for consolidated storage.

Get a FREE copy of ESG's Ethernet Storage whitepaper and see how you can set the standards for the next generation of storage solutions. To register for your copy and see our limited time discounts go to: www.extremenetworks.com/howstimulating

---

extreme networks™

© 2009 Extreme Networks, Inc. All rights reserved. Extreme Networks and the Extreme Networks logo are either registered trademarks or trademarks of Extreme Networks, Inc. or its subsidiaries in the United States and other countries. All other names and marks are the property of their respective owners.

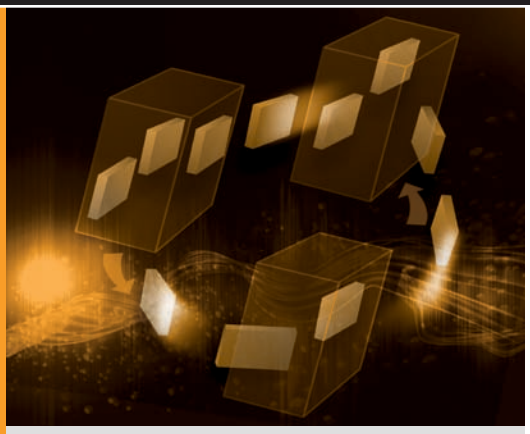

STRATEGIES TO OPTIMIZE VIRTUAL MACHINE CONNECTIVITY WITH XSIGO VIRTUAL I/O

By Scott Hanson

Server virtualization places demands on I/O infrastructures that can limit overall performance and efficiency. By using virtual connectivity and a high-speed converged interconnect to remove traditional barriers, Xsigo® virtual I/O can help increase virtual machine performance and flexibility, reduce power and cooling requirements, and simplify management.

Increased resource utilization is the primary objective of enterprise consolidation initiatives—and typically the greatest single benefit of server virtualization. But IT managers are now finding that deploying virtualization with a traditional server I/O architecture introduces its own problems that can prevent consolidation projects from reaching their full potential.

There are several reasons why server virtualization places greater demands on connectivity than traditional non-virtualized environments. The increased hardware utilization, storage traffic, and network traffic increases the risk of I/O bottlenecks. Application mobility requires that each server be able to access storage and network resources across the environment, increasing the number of connections per server. And connectivity for virtual machine (VM) migration and management networks further increases connectivity demands.

Faced with these challenges, IT staff commonly try to enhance connectivity by adding network and storage ports to the servers. But this approach has its limitations. The increased connectivity can increase management complexity, costs, and server space requirements—and even then, performance and efficiency may be suboptimal. In some cases, connectivity may be limited by the space constraints of the servers themselves. Failure to address these issues, meanwhile, can lead to unpredictable performance due to traffic congestion and an inflexible infrastructure that may not be able to accommodate the needs of particular applications.

Xsigo virtual I/O is designed to overcome these challenges to server connectivity, complementing Dell™ PowerEdge™ rack-optimized or blade servers by enabling fast, easily managed connectivity to Gigabit Ethernet and 10 Gigabit Ethernet networks and to storage systems such as Dell EqualLogic™ PS Series Internet SCSI (iSCSI) storage area network (SAN) arrays and Dell/EMC arrays. By transitioning to a next-generation virtual I/O approach in virtualized environments, organizations can greatly simplify the infrastructure while providing the necessary I/O capacity to maximize efficiency and VM performance.

APPLYING VIRTUALIZATION TECHNOLOGY TO SERVER I/O

As fundamental a change as server virtualization itself, virtual I/O can be thought of as two technologies: I/O resource virtualization and converged connectivity. Together, they can deliver dramatic management and performance advantages in virtualized environments.

I/O resource virtualization is somewhat analogous to server virtualization. Just as server virtualization allows one processor to function as multiple virtual processors, virtual I/O allows one physical adapter card to appear as multiple virtual network interface cards (vNICs) and virtual host bus adapters (vHBAs).

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

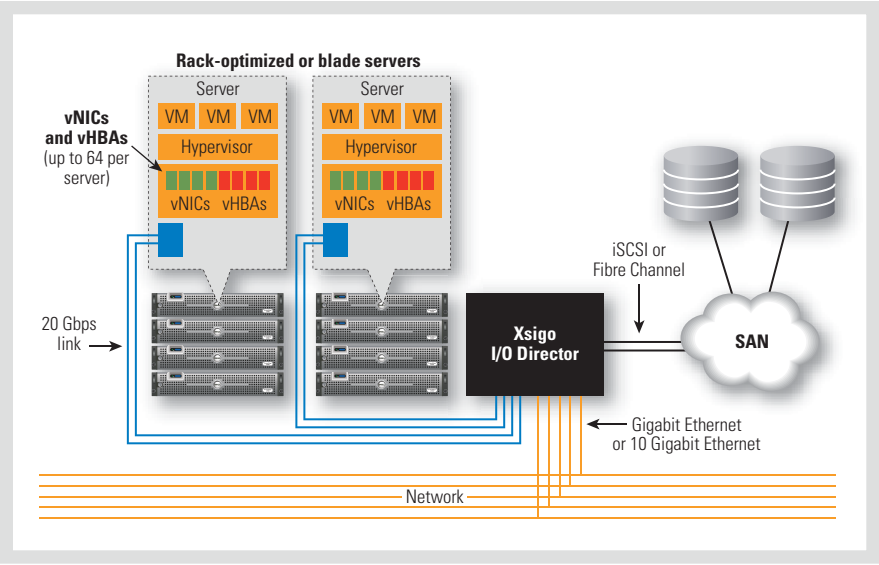

Figure 1. *The Xsigo I/O Director connects servers to conventional LAN and SAN resources through up to 64 vNICs and vHBAs without downtime*

These vNICs and vHBAs function just as conventional NICs and HBAs would: they are designed to be compatible with existing operating systems, hypervisors, and applications, and appear to LANs and SANs as conventional cards.

Converged connectivity replaces the numerous I/O cables of a typical server with a single cable that provides shared transport for all network and storage connections. That cable (or two cables for redundancy) connects to the external Xsigo I/O Director™ device, which then provides connections to the conventional data center Ethernet and Fibre Channel switches (see Figure 1). Up to 64 virtual connections can be deployed to each server, which means that each server can be attached to 64 distinct Ethernet and Fibre Channel networks.

INCREASING VM FLEXIBILITY

Virtual I/O helps address the demands of server virtualization in several ways, but one key attribute is that virtual I/O is *dynamic*: new connections can be deployed on demand without server downtime. VM mobility, for example, often presents a management challenge because different applications require access to different network and storage resources: one application may require iSCSI over 10 Gigabit Ethernet for storage connectivity, while another needs Fibre Channel. But configuring every server with both these resources can quickly become expensive.

Virtual I/O offers a simplified, powerful, cost-effective solution for VM mobility. Because vNICs and vHBAs are software elements, they can be deployed to different servers as needed—any server can potentially connect to any network or storage asset, in real time, without regard for the physical configuration of that device. This approach helps eliminate operational boundaries that could otherwise limit flexibility and VM density.

Blade systems in particular can benefit from this capability. Because virtual I/O helps reduce the dependency on mezzanine cards, blade servers can be dynamically configured with the specific I/O required for a given workload. This in turn helps increase the utility of the blades by enabling them to support a wider range of applications and more VMs per blade than they could otherwise.

ENHANCING VM PERFORMANCE

Traditional I/O can introduce performance bottlenecks when heavily trafficked connections become overloaded. IT staff could add I/O resources to help alleviate this problem, but this approach comes at the cost of additional equipment and increased complexity.

Virtual I/O actually addresses this type of bottleneck in the opposite way: it *reduces* the number of physical resources. Xsigo virtual I/O uses a single 20 Gbps connection that is dynamically allocated in real time across multiple virtual connections to both storage and network resources. Because this 20 Gbps capacity typically exceeds the I/O capacity of the server itself, the link is not a limiting factor. In I/O-intensive applications, this approach can help increase both VM performance and the potential number of VMs per server.

Server-to-server data transfer can also benefit from virtual I/O. For example, when an application server accesses data from a database server, the data transfer occurs entirely within the virtual I/O environment over the 20 Gbps link, without utilizing external networks—helping both increase application performance and decrease network congestion.

ACCELERATING AND SIMPLIFYING VM MIGRATION

The high-speed Xsigo server-to-server communication link helps accelerate VM migration as well. Moving a VM from one server to another requires transferring the application's state information—which could be gigabytes of data. Best practices recommend using a dedicated interconnect for this type of transfer, typically by implementing VM migration over a separate Gigabit Ethernet network. Virtual I/O enables administrators to configure a dedicated network that has access to the full 20 Gbps bandwidth when required, which can help significantly accelerate the migration process.

Virtual I/O can also help simplify VM migration by minimizing hardware constraints. Migration usually requires identical I/O on the source and destination servers, a restriction that can limit flexibility. Virtual I/O enables connectivity to be configured on the fly, allowing VM migration to proceed without regard for the physical I/O resources.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

MAINTAINING QUALITY OF SERVICE

Within a virtualized server, the virtual switch provides a useful tool for sharing I/O among multiple VMs—but can also result in unpredictable performance because of resource contention at the physical I/O level. Virtual I/O can help eliminate resource contention without the need to add I/O cards and cables.

Because vNICs and vHBAs can be deployed on demand, administrators can easily provision critical VMs with dedicated storage and network connectivity. Quality-of-service (QoS) controls help ensure that critical applications receive the bandwidth required from each connection. Administrators can define QoS controls per virtual resource (on both storage and network connections), and can control both the committed information rate (the minimum allowed bandwidth) and peak information rate (the maximum allowed bandwidth). These controls help ensure that critical applications can coexist on a shared hardware platform without resource contention.

INCREASING ENERGY EFFICIENCY

Virtualized servers tend to be more I/O intensive than traditional non-virtualized servers, and those I/O resources can require a surprising amount of power. In a typical environment with 120 servers, Xsigo virtual I/O can require 70 percent fewer I/O cards, cables, and switch ports than a traditional I/O infrastructure—helping to reduce power consumption by up to 30 percent and significantly lowering related power and cooling costs in the data center.[1]

CENTRALIZING I/O MANAGEMENT

Virtual sprawl typically refers to an overabundance of VMs, but it could just as easily apply to connectivity. Managing the physical connections, virtual switches, and data center switches is becoming increasingly challenging. But because using virtual I/O means that those virtualized resources exist in software, they can be monitored and managed from a software utility.

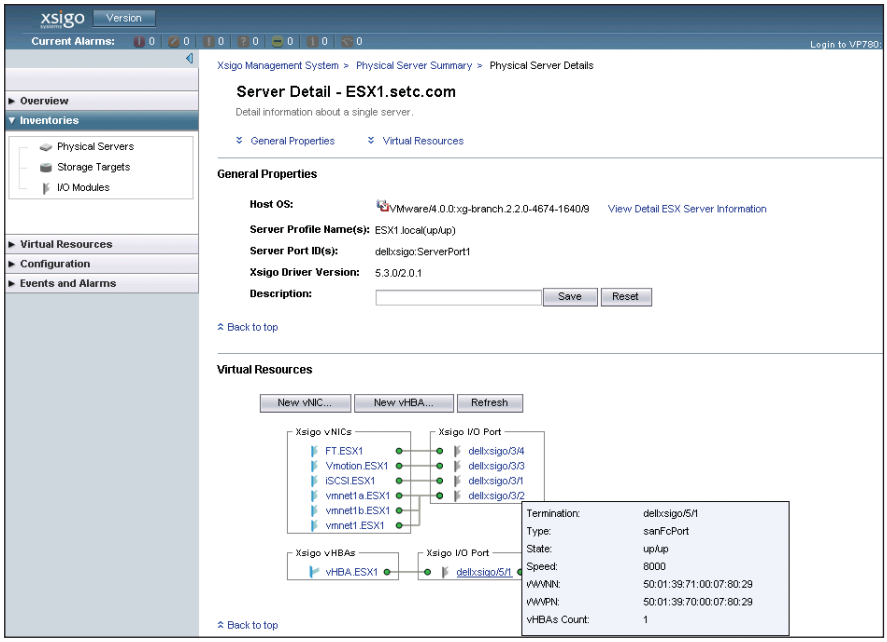

Figure 2. *The Xsigo Management System provides reporting and monitoring tools to help manage virtual connectivity*

The Xsigo Management System (XMS) provides reporting views that enable administrators to see the currently configured servers, vNICs, and vHBAs from a single management console (see Figure 2). Traffic monitoring tools in the XMS help administrators understand how their resources are being used. And in addition to running as a stand-alone system, the XMS has also been integrated with the VMware® Infrastructure Client to enable comprehensive management of VMs and virtual I/O from a single application.

OPTIMIZING CONNECTIVITY IN VIRTUALIZED ENVIRONMENTS

The dynamic data center enabled by virtualization technology can offer major advantages, but also introduces multiple performance and management challenges. Virtual I/O is designed to meet those challenges—helping increase VM flexibility and performance, reduce power and cooling requirements, and simplify management through an approach that matches the technological elegance of server virtualization itself.

Scott Hanson is a senior engineering consultant on the Dell Enterprise Technology Center team. Scott has a B.S. in Computer Science from the University of North Texas and holds Microsoft® Certified Systems Engineer + Internet (MCSE+I) and Red Hat Certified Engineer® (RHCE®) certifications.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Xsigo and Dell:
www.xsigo.com/dell

Dell TechCenter:
www.delltechcenter.com

[1] Comparison based on 120 servers each configured with four Ethernet ports and two Fibre Channel ports. Typical server connectivity for this configuration includes 120 quad-port Ethernet cards, 120 dual-port HBAs, 10 Ethernet switches, and 6 Fibre Channel switches, and draws approximately 4,200 W; Xsigo connectivity includes 120 dual-port host channel adapters (HCAs), 2 Xsigo I/O Director devices, and 6 Xsigo expansion switches, and draws approximately 2,912 W. Savings will vary based on configuration, usage, and manufacturing variability.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Adapt quickly to any environment.

Get any-to-any connectivity with Xsigo's I/O Director.

The Xsigo I/O Director lets you connect any server to any network or storage device in seconds, without the maze of cards, cables, and switch ports that add cost and complexity. With Xsigo, you can dynamically configure up to 64 network and storage connections to each server, without downtime. Xsigo's fast 20Gbs link and fine-grained QoS controls let you deliver I/O performance where it is needed.

- **50% less cost than conventional I/O**[1]
- **70% fewer I/O cards, cables, and switch ports**
- **100X faster I/O configuration management**

For a FREE analyst report on the cost and performance benefits of virtual I/O, call 408-329-5600, or visit **www.xsigo.com/dell.**

The Xsigo I/O Director is available at Dell.com, part #A2210294

xsigo systems

[1] *For complete details, go to www.xsigo.com/savings. Savings vary depending on configuration.*

BEST PRACTICES FOR SCALING DATA CENTER POWER CAPACITY AND AVOIDING HIDDEN COSTS

By Wendy Torell

As energy costs rise, excess capacity in enterprise data centers has become a financial and ecological burden. Using a stepped, phase-in approach and deploying scalable uninterruptible power supply (UPS) systems can help organizations avoid the hidden costs of conventional UPS designs—creating an infrastructure that can expand efficiently and cost-effectively as IT loads grow.

When designing a data center power infrastructure, organizations must consider not only their current IT loads, but also the loads they expect to support in the future. Historically, a full up-front build-out designed to support the maximum projected load has provided a way to ensure that data centers have room to grow without exceeding available capacity. When energy costs were relatively low, incorporating enough spare power capacity to handle even extreme workloads was often a matter of pride and preparedness—but today, with a stressed energy supply and skyrocketing costs, excess unused capacity has become a financial and ecological burden.

Although conventional approaches to uninterruptible power supply (UPS) systems offer a certain amount of scalability, these infrastructures also come with hidden costs that can increase maintenance requirements and downtime for critical systems. By instead using a stepped, phase-in approach to UPS deployment designed to match actual loads and expand only when needed, organizations can create scalable data centers that can expand efficiently and cost-effectively as IT requirements grow.

PLANNING FOR GROWTH

Enterprise data centers face multiple challenges that a stepped, phase-in approach to UPS deployment can help overcome:

- **Energy costs:** Energy has become a major expense, and the electric bill is a powerful incentive to avoid excess capacity wherever possible. Green initiatives and demand-side management programs that reward efficient operation further increase the incentive to run a lean data center. A right-sized data center—in which power and cooling capacity tracks the growth of the IT load—is typically much more efficient than one with excessive unused capacity.
- **Unused capacity:** Data centers rarely build out to support the maximum load levels projected during planning. A stepped, phase-in approach can mitigate the risk of installing capacity that may never be used and help avoid unnecessary capital expense—making this the largest benefit of incremental deployment for many data centers.
- **Unnecessary maintenance costs:** Unused capacity generates unnecessary maintenance costs, because installed equipment must be maintained and repaired even if it is idle. Installing only components needed to support the current load can help avoid significant service expenses.

Figure 1 illustrates the traditionally accepted excess capacity that occurs over a data center's lifetime, in which the actual IT load might be as little as half the capacity built into the data center. Why do data centers end up this way? Part of the

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

reason is that new data center plans are often vague, and are unable to account for the constantly evolving nature of technical developments in IT equipment. In addition, the further into the future an IT load projection extends, the lower the confidence in the projection. For many organizations, oversized capacity is perceived as a better option than disruptive future upgrades that could lead to downtime.

Figure 2 illustrates the stepped, phase-in approach: when the future size of the IT load is uncertain, phase-in steps can provide organizations with stopping points for reassessment before committing to further deployment. The higher the degree of uncertainty, the smaller and more frequent the steps should typically be, to provide additional opportunities for adjustment based on developing conditions. For extreme uncertainty, this bail-out feature can become the primary consideration in designing the size and frequency of steps.

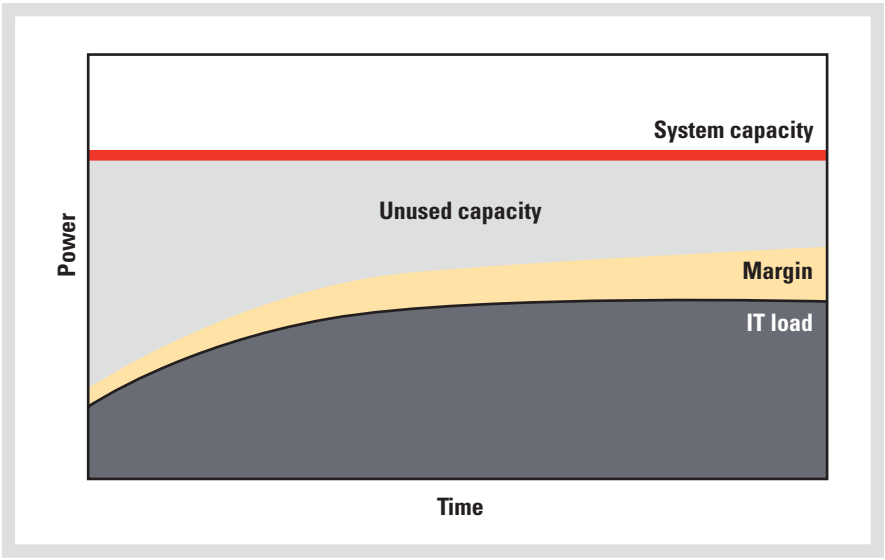

Figure 1. *Wasted capacity over a data center's lifetime when IT load does not reach the projected maximum*

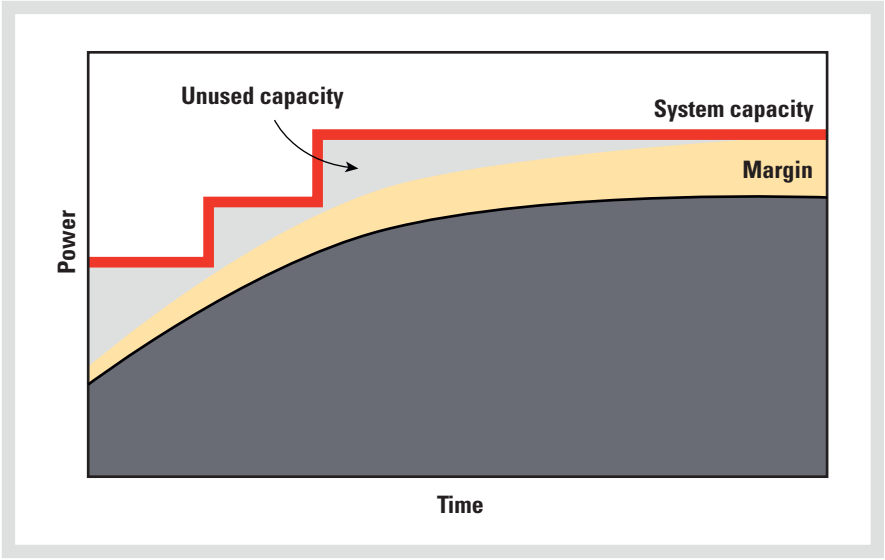

Figure 2. *Efficient resource utilization using a stepped, phase-in plan that matches capacity to actual IT load*

COMPARING UPS APPROACHES TO SCALABILITY

Some might argue that because it is technically *possible* to expand the capacity of an existing UPS architecture, virtually any UPS is scalable. Although this may be true, it is important to understand the constraints associated with scaling certain UPS designs—and the resulting costs.

Understanding conventional and scalable designs

Two typical approaches to UPS scalability are conventional UPSs and scalable UPSs. Conventional UPSs can be designed to provide a degree of scalability by adding UPS power modules to a parallel bus, provided that the paralleling switchgear deployed during the initial installation has sufficient power-handling capacity to support the final, full-size power-handling capacity of the final UPS system configuration. Scalable UPSs, such as the APC® Symmetra® PX models, can afford the same advantage in that the supply can match the load growth in the racks, but can also help avoid the need to purchase a large-scale paralleling system cabinet as well as the accompanying cost.

As an example, consider a hypothetical environment in which an IT manager has identified an initial capacity requirement of 80 kW of UPS power and a planned ultimate capacity of 240 kW. The environment includes sufficient redundancy to take one of the power modules offline for service or repairs while staying on conditioned power.

In the conventional approach, the IT manager plans to make the system expandable by buying it in stages, which helps minimize up-front project costs. The manager assumes that the load growth will occur in 80 kW increments. On the first day, then, the plan is to use two 80 kW UPS modules in parallel so that the anticipated 80 kW initial load can be supported by either module in an $n + 1$ configuration. When the load begins to exceed the 80 kW capacity of the redundant system, another 80 kW module would be installed

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

to maintain UPS redundancy while meeting the capacity requirements of the load, which would then grow to 160 kW. When the need for the final tier of growth approaches, the last 80 kW UPS module would be installed, bringing the total capacity to 240 kW with one redundant 80 kW UPS module. In its final configuration, this UPS system would be a four-module paralleled system. This approach offers a degree of scalability, with the understanding that the costs of installing the full-size paralleling equipment necessary for the final configuration are included in the initial purchase.

In the scalable approach, the IT manager purchases an 80 kW UPS to support the initial 80 kW load. The scalable system has redundant 10 kW power modules within the unit, providing an *n* + 1 configuration without the need for a second 80 kW unit; the scalable system also does not require large parallel electrical distribution equipment, helping avoid the cost of purchasing and installing that equipment. As the load grows past the 80 kW level, a second UPS—also with built-in *n* + 1 redundancy—would be deployed, with the new loads then dedicated to its output. A similar third UPS would follow as the data center grows to full capacity.

Evaluating ongoing costs

In this example scenario, the two approaches appear very similar on the surface. However, the conventional approach comes with a hidden cost. Conventional UPS systems installed with equipment to parallel the necessary modules and provide scalability and redundancy have common tie points where the output of each module is attached to the load distribution equipment. These tie points are the *critical bus* where the power supplied by a given UPS module is coupled with the power from the other modules. When initially commissioning the system, organizations typically verify the ability of the modules to parallel, share load, and demonstrate redundancy. As they add modules, IT staff must power down the critical bus to help safely tie in the new module, and then re-verify the function of the system as a whole through another commissioning exercise.

Implementing the planned expansion of the conventional system as described in the example scenario involves a minimum of two shutdowns of the critical bus. In addition to the time that the UPS bus must be down or bypassed, organizations must consider the impact that a major construction project has in a critical facility. Tasks such as rigging a large UPS into the room while the system is supporting critical loads, hanging conduits from the ceiling, and pulling cables, as well as having a large number of people working in the vicinity of critical power equipment, can pose an even greater downtime risk than the amount of time necessary to make the electrical cutovers.

Expanding a scalable UPS system, by contrast, typically has no impact on operations under normal conditions. These systems are designed to be load specific within the data center, with each UPS supplying a dedicated number of racks. As long as provisions are made to supply power to the expansion systems during the initial data center construction, installing and testing the new UPS systems should have minimal impact on concurrent operations. IT staff can load test the systems that provide the additional capacity using smaller load banks than they could in the conventional system. This is because they do not need to test the paralleling function and the capacity of each test is limited to 80 kW. No shutdown of existing processing is required, because no common critical bus exists that must be tied in to provide power—the critical bus of each UPS system is dedicated to the 80 kW load it is designed to support. The overall costs associated with the expansion of a scalable design are typically much smaller than those of the conventional design because the system can be expanded without taking loads offline.

AVOIDING THE HIDDEN COSTS OF EXPANSION

Several best practices can help organizations when designing and deploying a scalable power infrastructure:

- Apply a stepped, phase-in approach to help manage uncertainty and avoid the wasteful underutilization of capacity caused by overbuilding to an inflated target capacity.
- Use total cost of ownership to evaluate different architectures—up-front pricing is only part of the system costs over time.
- Use a scalable UPS system to help avoid the downtime typically required when upgrading power capacity in a conventional system.

By following best practices and deploying scalable UPS systems, organizations can significantly simplify planned expansion while avoiding downtime for critical IT resources—helping eliminate the difficulties and hidden costs associated with conventional approaches and enabling cost-effective data center growth.

Wendy Torell is a senior research analyst with American Power Conversion (APC) by Schneider Electric's Data Center Science Center. She has a bachelor's degree in Mechanical Engineering from Union College in Schenectady, New York, and an M.B.A. from the University of Rhode Island, and is an American Society for Quality (ASQ) Certified Reliability Engineer.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

APC scalable UPS systems:
www.apc.com/products

Dell™ Energy Smart:
DELL.COM/EnergySmart

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Now, upgrade your server room without overhauling it

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

Is your server room a barrier to adopting new technologies?

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

Introducing the APC® by Schneider Electric total server room solution

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

Future-proof your server room easily, cost-effectively

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

The integrated, cooled, managed server room

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

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

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

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

APC® by Schneider Electric

© 2009 American Power Conversion Corporation. All trademarks are owned by Schneider Electric Industries S.A.S., American Power Conversion Corporation, or their affiliated companies.
e-mail: esupport@apc.com • 132 Fairgrounds Road, West Kingston, RI 02892 USA • 998-2028 *Full details are available online.

INTRODUCING QLOGIC 12000 SERIES QUAD DATA RATE INFINIBAND SWITCHES

By Anand Vridhagiri
Brady Black

QLogic® 12000 series Quad Data Rate (QDR) InfiniBand switches offer an energy-efficient, highly scalable way to build low-latency InfiniBand fabrics for high-performance computing clusters.

Most high-performance computing (HPC) applications are highly sensitive to latency, and in HPC environments, reduced latency can lead to significantly accelerated performance. The efficient InfiniBand I/O interconnect is designed to provide high-speed data transfers and extremely low latencies over a single scalable fabric, offering a way to meet the needs of even highly demanding HPC software.

QLogic 12000 series Quad Data Rate (QDR) InfiniBand edge switches and director switches are designed to help maximize flexibility, scalability, and bandwidth in InfiniBand-based HPC environments. These advanced switches, based on the 36-port QLogic TrueScale™ QDR application-specific integrated circuit (ASIC) platform, support HPC, database clustering, and grid utility computing applications and are designed to maximize HPC interconnect bandwidth while helping to simplify fabric management, minimize switch power consumption, and reduce data center costs.

ADVANCED SWITCH DESIGN

QLogic 12000 series switches adhere to version 1.2 of the InfiniBand Trade Association specification, helping ensure interoperability with other compliant devices. Combining QLogic 12000 series switches with the QLogic InfiniBand Fabric Suite can provide powerful, simplified fabric management, including automated fabric installation, configuration, and monitoring. Advanced quality-of-service features, adaptive routing, and QLogic vFabric™ virtual InfiniBand fabrics are designed to enhance utilization of InfiniBand resources.

For example, the advanced QLogic 12000 series design allows a fabric to be shared by mission-critical applications to help maximize bandwidth utilization, while the segmentation feature allows multiple subnet managers to run within a single fabric to provide additional segmentation of virtual clusters. The switches also allow administrators to implement virtual fabrics on a port-by-port basis, enabling them to create multiple virtual InfiniBand fabrics at the port level. And adaptive routing allows the switches to analyze traffic patterns, identify congestion, and make routing table changes to intelligently utilize available bandwidth.

The major components of QLogic 12000 series switches are designed to be field replaceable and hot pluggable, helping simplify and accelerate repairs and maintenance. Other advanced features include nondisruptive firmware upgrades, port-to-port and module-to-module failover, component-level diagnostics and alarms, and both in-band and out-of-band remote management to help simplify administrative tasks. QLogic 12800 series director switches incorporate QLogic StarPower™ technology to help support green data center solutions, including high port densities and low per-port power consumption that can help reduce the data center footprint as well as power and cooling requirements.

FLEXIBLE EDGE SWITCHES

The QLogic 12000 series includes two 36-port edge switch models: the QLogic 12300 and the QLogic 12200. The QLogic 12300 is designed to be used as an edge switch in large clusters or to help build small workgroup clusters. It includes hot-pluggable and redundant fan

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

modules as well as optional redundant power supplies, and supports the advanced features available across the QLogic 12000 series—including the QLogic InfiniBand Fabric Suite, QLogic vFabric virtual fabrics, the QLogic InfiniBand Fabric OS, QLogic TrueScale ASICs, and wizard-based installation and configuration. The QLogic 12300 is designed to support aggregate bidirectional bandwidth of up to 2.88 Tbps at the full specified QDR InfiniBand bandwidth of 40 Gbps per port, as well as scalable latency designed to be constant at traffic levels beyond 90 percent of the total available bandwidth. The QLogic 12200 switch is designed to be used as an edge switch in large clusters, offering many of the same standard features as the configurable QLogic 12300 at a lower cost.

Both the QLogic 12300 and QLogic 12200 models also help minimize switch energy use. Designed for extremely low per-port power consumption, these switches can help administrators control power and cooling requirements and related operational costs while supporting green data center initiatives.

SCALABLE DIRECTOR SWITCHES

QLogic 12800 series director switches take advantage of a common design that can accept the same spines, I/O leaf modules, management cards, power modules, and fan modules. In configurations using these modular components, this approach helps simplify shelf sparing by enabling administrators to order from a single stock-keeping unit (SKU) regardless of chassis type.

For large clusters, QLogic 12800 series director switches feature a modular design based on port, management, power, and cooling building blocks that can be used to scale from entry-level 18-port director switches up to massive 864-port core switches. These switches are also designed for high availability, including redundant components, support for automatic failover, and hot-swappable, field-replaceable port modules, management modules, power supplies, and fans to help simplify and accelerate repairs.

	12800-040	12800-060	12800-120	12800-180	12800-360
Maximum number of ports with UHD leaf	96	144	288	432	864
Maximum number of ports with UHP leaf	72	108	216	324	648
Maximum bandwidth with UHP leaf	5.76 Tbps	8.64 Tbps	17.28 Tbps	25.92 Tbps	51.84 Tbps

Figure 1. *Ports and maximum bandwidth supported by QLogic 12800 series director switches*

QLogic 12800 series switches can support either 18-port Ultra High Performance (UHP) leaf modules designed to maximize bandwidth or 24-port Ultra High Density (UHD) leaf modules designed to maximize connectivity (see Figure 1). For organizations with high performance and bandwidth requirements, the UHP leafs can support up to 648 ports at QDR InfiniBand speeds for an aggregate bidirectional bandwidth of up to 51.84 Tbps in the QLogic 12800-360 model, making this the fastest rack-mount QDR InfiniBand switch available. The QLogic 12800-360 model is also the only switch that supports building a single switch cluster of 325 ports up to 648 ports, avoiding the need for a multitiered solution using additional switches to achieve this port count, which can increase both latency and costs.

UHD leaf modules can take advantage of QDR InfiniBand speeds within the switch chassis or support additional ports at a 2:1 oversubscribed QDR speed. When using the UHD leafs, 12 ports connect from the backplane and 24 ports are external, helping increase port density and enabling support for up to 864 ports in the QLogic 12800-360 model. The UHD leafs enable cost-effective support for Double Data Rate (DDR) InfiniBand speeds while still providing QDR InfiniBand connectivity and helping simplify future expansion.

Importantly, QLogic 12800 series director switches are designed not only for high port counts and bandwidth, but also high levels of energy efficiency. As with the QLogic edge switches, these QLogic director switches are designed for low per-port power consumption—helping administrators control power and cooling requirements and related operational costs while supporting green data center initiatives.

EFFICIENT, HIGH-PERFORMANCE QDR INFINIBAND SWITCHES

Low latency, energy efficiency, scalability, and simplified management can be critical in HPC environments. QLogic 12000 series QDR InfiniBand switches along with the QLogic InfiniBand Fabric Suite are designed to help organizations—even those with little or no previous HPC experience—to quickly, efficiently, and effectively install and configure InfiniBand-based HPC clusters while helping simplify administration, reduce power and cooling requirements, and reduce operational costs.

Anand Vridhagiri is a senior OEM marketing manager at QLogic.

Brady Black is an HPC solutions architect at QLogic.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

QLogic HPC networking:
www.qlogic.com/products/hpc_products_landingpage.aspx

QLogic and Dell:
www.qlogic.com/go/dell

Dell™ HPC solutions:
DELL.COM/HPC

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

NEXT-GENERATION APPLIANCE SOLUTIONS POWERED BY DELL POWEREDGE SERVERS

By Franklin Flint

The 11th-generation Dell™ PowerEdge™ server family introduces myriad features designed to meet the needs of original equipment manufacturers (OEMs) developing and selling custom appliances. Combined with comprehensive services from the Dell OEM Industry Solutions Group, these servers can provide robust, reliable, cost-effective platforms for a broad range of OEM appliances.

In the 10 years since the Dell OEM Industry Solutions Group began providing hardware, services, and supply chain management to original equipment manufacturers (OEMs) developing and selling custom appliances, the group has helped hundreds of companies get what they need for their hardware solutions—ranging from network appliances to medical solutions, digital media devices, manufacturing and process control solutions, and more. (For one example, see the "Small business, big success" sidebar in this article.) The Dell OEM division offers a wide range of branding solutions to help meet a variety of OEM needs depending on specific applications, marketing goals, and acceptable cost models. These solutions range from standard Dell-branded systems to generic-looking hardware to customized products with unique bezels, logos, colors, and packaging. Dell also has the capability to create full-fledged custom supply chains to handle hardware design, manufacturing, assembly, integration, inventory control, delivery, and post-sales support, enabling the OEMs to focus on software development and marketing.

The direct model has allowed Dell to learn firsthand what these OEMs (and their end users) need by actively requesting, compiling, and studying information from OEMs to find opportunities for improvement, and then applying this knowledge when designing and manufacturing Dell hardware platforms. New 11th-generation Dell PowerEdge servers—including currently available models and those expected to be released soon—incorporate many new features requested by appliance OEMs, and are designed to provide robust, reliable, cost-effective platforms for their appliances.

FLEXIBLE LIFE CYCLE OPTIONS

For many OEMs, consistent hardware built with a long life cycle along with advance information on platform updates is critical to appliance development and marketing. For example, when a new platform is released, many OEMs must be able to continue selling the previous hardware until they can develop on the new hardware and perform the necessary qualification tests to help ensure that it is stable with their applications. Previously, the Dell change management program allowed three months of overlap with legacy platforms, but with the introduction of the 11th-generation PowerEdge server family, this period has been extended to a full six months—doubling the time available for OEMs to qualify their applications on new hardware.[1]

[1] Product transition periods will vary depending on customer forecasting, component availability, and demand. To achieve a full six-month transition, accurate and complete forecasts with commitments must be provided.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Consistent hardware compatibility throughout the platform life cycle is also important in helping to reduce the need for new hardware development and qualification. To help meet this need for platforms that support full hardware RAID, Dell plans to offer the current PowerEdge Expandable RAID Controller (PERC) 6 for the entire life of the 11th-generation PowerEdge platforms. For OEMs that do not require the features and capabilities introduced in next-generation controllers released during the lifetimes of these servers, the continued availability of the PERC 6 can help them save both development time and engineering resources by avoiding the need to recompile custom kernels to support next-generation RAID controllers. Dell also offers versions of 11th-generation PowerEdge appliance platforms with long-life Intel® Xeon® processors from Intel's embedded processor road map, helping minimize transitions for OEMs sensitive to processor changes.

In addition to these enhancements to hardware consistency, the Industry Solutions Group also offers change management services through online tools, regular road map presentations, direct account management teams, and even project managers when appropriate. Ultimately, the goal of the Industry Solutions Group team is to be sure that OEMs know what is changing in the latest hardware, when those changes are expected, and how Dell can help smooth transitions for OEMs.

VERSATILE HARDWARE ENHANCEMENTS

Although rack and tower servers are typically intended for installation in data center racks in traditional locations, the reality is that many appliance OEMs must place these servers in locations that are less than ideal—inside closets, in frame relay telco racks with minimal rear clearance, or inside a larger system. To help accommodate these types of uses, many entry-level 11th-generation Dell PowerEdge servers (currently available and upcoming models) are shorter than their predecessor models by several inches, including one upcoming 1U rack-mount model that is expected to be just 15.5 inches deep.

11th-generation Dell PowerEdge servers are designed to provide robust, reliable, cost-effective platforms for OEM appliances

The 11th-generation PowerEdge appliance platforms also offer myriad other enhancements, including the following:

- **Internal persistent flash storage:** Bootable, managed, and available in multiple options, internal USB and Secure Digital (SD) storage connections enable diskless configurations or permanent utility storage options for when disks become corrupted or fail.
- **Four network interface cards (NICs) on the motherboard:** High-performance platforms include four NICs embedded on the motherboard, helping to free valuable PCI Express (PCIe) slots for other I/O options, enhance compatibility, and reduce costs.
- **Chipset RAID option:** Certain platforms are expected to offer RAID support from the system chipset rather than the RAID controller, helping to cut costs and maintain reliability.
- **Expanded RAM options:** By offering more dual in-line memory module (DIMM) slots and higher-capacity DIMMs than were available in comparable previous-generation PowerEdge models, 11th-generation PowerEdge platforms help increase flexibility for developers by enabling them to optimize their applications for performance, cost, or a balance of the two.
- **PCIe 2.0 slots:** PCIe 2.0 slots are designed to support twice the I/O throughput when the next-generation PCIe add-in cards are available.
- **x16 PCIe slots:** Many 11th-generation PowerEdge platforms have options for x16 PCIe slots, helping support additional I/O options and higher-end video cards than they could otherwise.
- **Enhanced LCDs:** The LCDs support higher resolution, more characters, and more options for alerting and customization than previous-generation LCDs.
- **IP version 6 (IPv6) support:** Support for IPv6 in management hardware options helps increase compatibility with anticipated future technologies.
- **ENERGY STAR certification:** Dell PowerEdge R610 and PowerEdge R710 rack servers meet the U.S. government's ENERGY STAR program requirements, and the family as a whole is designed for energy efficiency to help meet the standards of environmentally conscious OEMs.
- **Headless operation:** Headless operation with "no halt on errors" support in the BIOS enables appliances to continue booting when a power-on self-test (POST) error occurs that is not critical to system operation.
- **Open source drivers for Linux® kernels:** On platforms running the Red Hat® Enterprise Linux or Novell® SUSE® Linux Enterprise operating systems, open source drivers for embedded server components are available to help simplify development of nonstandard or custom operating systems.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SMALL BUSINESS, BIG SUCCESS

ScienceLogic EM7 Meta-Appliances are designed to take the complexity out of IT management by incorporating the applications, tools, and reporting features that IT managers need into one flexible, easy-to-deploy appliance. Despite earning rave reviews from customers, in many cases the benefits were offset by issues with the hardware platform, which was supplied by a white-box server manufacturer—nearly 20 percent of customers experienced hardware-related issues that required them to contact the ScienceLogic help desk.

After less than a year, it became clear that ScienceLogic needed to find a new hardware source that could meet several key criteria: excellent hardware reliability, timely support capabilities, the ability to consistently fulfill orders that fluctuated in size depending on customer demand, and the ability to offer a broad range of platform sizes to meet varying customer needs. After carefully evaluating several options, the ScienceLogic team chose the Dell OEM Industry Solutions Group. "Once we talked to Dell, it was clear they were the best fit for us," says Louis DiMeglio, a ScienceLogic partner and senior vice president of customer engineering. "Most of the other manufacturers we met wanted us to go through resellers because they didn't consider us large enough for direct business at that time. The Dell OEM Industry Solutions Group was ready to set up a direct relationship and help us quickly grow our appliance business."

> **"Our growth of over 800 percent in the past four years has certainly been enabled by Dell's global support. We made a great decision when we went with Dell, and we haven't looked back since."**
>
> —Louis DiMeglio
> Senior vice president of customer engineering at ScienceLogic
> February 2009

The Dell OEM team helped ScienceLogic select a range of platform sizes from the Dell server line: the PowerEdge 1950 server for its entry-level product, the PowerEdge 2950 server for its mid-range product, and the PowerEdge R900 and PowerEdge R905 servers for its high-end products, all with quad-core Intel Xeon processors. "The Dell server models gave us just the right additional capacity at each of our product levels," says DiMeglio.

The switch to Dell platforms brought a host of benefits—including reducing hardware failures from 20 percent to less than 1 percent, accelerating hardware repairs with on-site service from a Dell technician within four hours of diagnosis, cutting order fulfillment time by an average of 10 days, reducing inventory costs by 100 percent, and providing a clear road map for technological advances. DiMeglio also credits Dell with helping to enable ScienceLogic's global expansion and rapid growth. "Using the white-box manufacturer, I do not believe we would have considered going beyond the U.S. and the UK, where we have a local presence," he says. "But working with Dell and knowing that we can get worldwide support for the hardware has allowed us to expand into other countries such as Australia and South Korea. Our growth of over 800 percent in the past four years has certainly been enabled by Dell's global support. We made a great decision when we went with Dell, and we haven't looked back since."

- **Custom boot screen image utility:** The Industry Solutions Group has developed a simple utility for applying a custom boot screen image to a server's BIOS, providing a cost-effective customization option.
- **Expanded server lineup:** To help meet the specific performance, cost, and feature needs of appliance OEMs, the 11th-generation PowerEdge server lineup has been expanded with more platform options than were available in previous generations.

ROBUST PLATFORM FOR OEM APPLIANCES

The 11th-generation Dell PowerEdge server family introduces a variety of enhancements to help meet the needs of appliance OEMs—incorporating not only myriad hardware improvements in the servers themselves, but also an extended overlap period with legacy platforms, long-life PERC and Intel processor options, and enhanced support services and capabilities. Combined with the supply chain, manufacturing, integration, and support services offered by the Dell OEM Industry Solutions Group, 11th-generation PowerEdge servers are designed to offer robust, reliable, feature-rich, and cost-effective platforms for OEM appliances.

Franklin Flint is the technology evangelist for the OEM Sales Group and a systems consultant in the Advanced Systems Group supporting the Dell OEM Industry Solutions Group. He has been at Dell for 14 years, and currently focuses on digital media, OEM appliances, audio/video technology, digital cinema, video surveillance, and audio recording.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell OEM appliances:
DELL.COM/Appliances

Dell OEM Industry Solutions Group:
DELL.COM/OEM

Direct2Dell blog:
www.direct2dell.com

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

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

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

www.novell.com/workload

Novell®
Making IT Work As One™

Copyright © 2009 Novell, Inc. All rights reserved. Novell, the Novell logo and PlateSpin are registered trademarks and Making IT Work As One is a trademark of Novell, Inc. in the United States and other countries.

ENHANCING SAP APPLICATION PERFORMANCE WITH DELL AND F5 SOLUTIONS

By Morten Loderup
Fred Johnson
Ron Carovano

Combining Dell™ servers and storage with F5® BIG-IP® Local Traffic Manager™ application delivery systems and the BIG-IP WebAccelerator™ module can help significantly increase the performance of SAP® applications—helping organizations enhance end-user experience, increase productivity, and reduce infrastructure costs.

Dell, F5, and SAP have developed long-term, in-depth partnerships designed to simplify integration and interoperability, optimize enterprise infrastructures, and maximize performance in environments using Dell servers and storage, F5 application delivery solutions, and SAP business software. In addition to the long-standing partnership between Dell and F5, Dell is an SAP global technology partner. As such, Dell operates various Dell SAP Competence Centers located around the world that serve as central points of contact for SAP application benchmarking, custom sizing for SAP implementations, performance engineering, and support services. F5 and SAP, meanwhile, have worked together to develop SAP-certified solutions that help ensure interoperability between F5 solutions and SAP software.[1]

In February 2009, Dell and F5 carried out proof-of-concept testing to evaluate the performance advantages of combining Dell systems with F5 BIG-IP Local Traffic Manager (LTM) systems and the BIG-IP WebAccelerator module in an SAP application environment. As the results demonstrate, using Dell and F5 solutions for SAP applications can help significantly increase performance over a wide area network (WAN)—helping to maximize end-user productivity, lower network bandwidth requirements, reduce operating costs, and increase return on investment.

TESTING DELL AND F5 SOLUTIONS FOR SAP

F5 BIG-IP LTM is an application delivery networking system designed to secure, optimize, and deliver applications, and incorporates load balancing, TCP connection management, intelligent compression, fast caching, and routing features. F5 BIG-IP WebAccelerator is an add-on software module for BIG-IP LTM designed to increase the performance of Web browsers, Web application platforms, and WANs. The combination is certified by SAP for integration with the SAP NetWeaver® platform.

The Dell and F5 test environment compared two configurations that simulated a remote-office user accessing the SAP Enterprise Portal component hosted in a central office over a WAN. In the first configuration, the remote-office client used a Web browser to directly access a single SAP Enterprise Portal Web and application server; in the second configuration, the remote-office client used a Web

[1] For more information on SAP-certified Dell servers and custom sizings from the Dell SAP Competence Centers, visit DELL.COM/SAP. For more information on F5 solutions for SAP applications, visit www.f5.com/sap.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

browser to access two SAP Enterprise Portal Web and application servers through a highly available pair of BIG-IP LTM systems with WebAccelerator modules in an asymmetric deployment (see Figures 1 and 2). In this second configuration, the BIG-IP LTM device replaced the SAP Web dispatcher, providing application load balancing and fault-tolerant high availability. Figure 3 details the hardware and software used in the test environment.

The initial set of tests, performed in the Dell Enterprise Reference Architectures lab in Austin, Texas, used a script manually executed eight times on the remote-office laptop to log in to SAP Enterprise Portal, navigate to the Knowledge Management page, download a document, and log out. The first configuration, in which the client was directly accessing SAP Enterprise Portal, took an average of 67.4 seconds to complete these tasks. The second configuration, using the F5 BIG-IP LTM systems with WebAccelerator, took an average of 54.1 seconds—a 20 percent faster elapsed time compared with the direct-access configuration.

Based on these promising initial results, more exhaustive, automated testing was conducted at the F5 Technology Center in Seattle, Washington. In addition to completing multiple runs using each of the two configurations, the test team evaluated performance under four different WAN conditions to provide a more comprehensive picture of real-world performance (see Figure 4). Because performance can depend on whether a user has recently visited the portal site and has content cached locally, the team also tested both first-visit and repeat-visit performance under each network condition.

Figures 5 and 6 illustrate the results. In first-visit tests, the BIG-IP LTM configuration completed the test script 23–36 percent faster than the direct-access configuration. Even in the repeat-visit tests, in which content is already cached in the client browser,

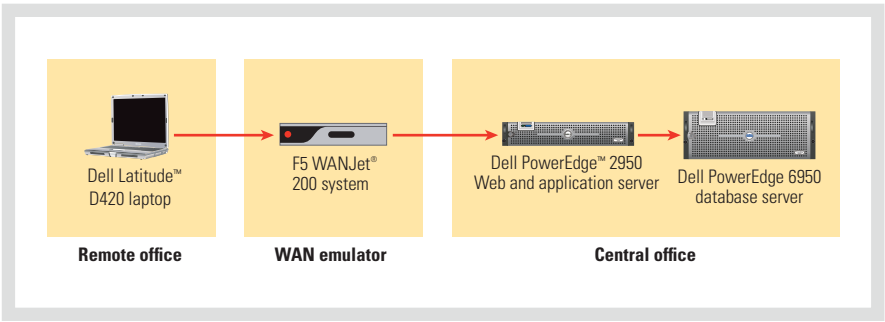

Figure 1. *Test configuration using direct access to SAP Enterprise Portal*

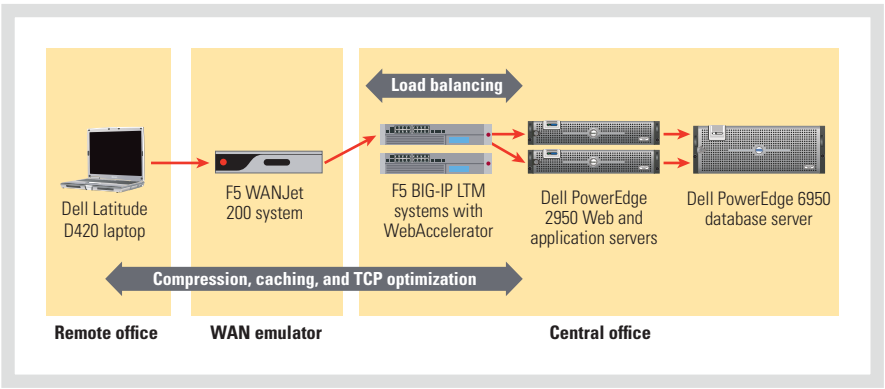

Figure 2. *Test configuration using F5 BIG-IP LTM systems with WebAccelerator to access SAP Enterprise Portal*

SAP Enterprise Portal Web and application servers	Two Dell PowerEdge 2950 servers with two quad-core Intel® Xeon® X5355 processors at 2.66 GHz and 3.24 GB of RAM running the Microsoft® Windows Server® 2003 OS with Service Pack 2 (SP2) and SAP Enterprise Portal 7.0
Database server	Dell PowerEdge 6950 server with four dual-core AMD Opteron™ 8220 SE processors at 2.8 GHz and 16 GB of RAM running the Microsoft Windows Server 2003 Enterprise Edition OS with SP2
Application delivery and central office	F5 BIG-IP LTM 6400 systems with WebAccelerator modules running the F5 TMOS™ 9.4.5 OS
Network switch	Dell PowerConnect™ 5224 switch
WAN emulator	F5 WANJet 200 running LANforge 4.3.3
Network configuration	T1 line with 1.544 Mbps of bandwidth, 300 ms of latency, and 0% packet loss
Remote-office client	Dell Latitude D420 laptop with two Intel Core™ Duo U2500 processors at 1.2 GHz and 1.49 GB of RAM running the Microsoft Windows® XP Professional Edition OS with SP3 and the Microsoft Internet Explorer® 7 Web browser
Laptop logging and decoding tool	HttpWatch 6.0 Basic Edition

Figure 3. *Hardware and software used in the test environment*

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

the BIG-IP LTM configuration completed the script 11–17 percent faster than the direct-access configuration.

Importantly, these results represent only a single user accessing the SAP Portal environment. Unlike the direct-access configuration, the BIG-IP LTM systems with WebAccelerator could continue to offload and load balance the SAP Enterprise Portal servers even under a load of hundreds or thousands of users—helping to significantly increase scalability and performance while enhancing end-user experience. In other words, the test conditions highlight the minimum performance advantage that could be provided in a real-world deployment.

	Bandwidth	Latency		Packet loss
		Range	Average	
Satellite	6–10 Mbps	500–900 ms	700 ms	2.5%
E1 (international)	1.5–2 Mbps	250–400 ms	300 ms	1.0%
Cable/DSL	4 Mbps	85–95 ms	90 ms	0.5%
OC-3	155 Mbps	10–40 ms	30 ms	0.005%

Figure 4. *WAN conditions in the test environment*

OPTIMIZING SAP APPLICATION PERFORMANCE

As the proof-of-concept tests demonstrate, the combination of Dell and F5 solutions can help significantly increase SAP application performance—helping to enhance end-user experience and increase productivity while reducing network bandwidth requirements. Taking advantage of these solutions can help organizations implement a secure, high-performance, highly available infrastructure that can help reduce operating costs and increase return on investment in SAP environments.

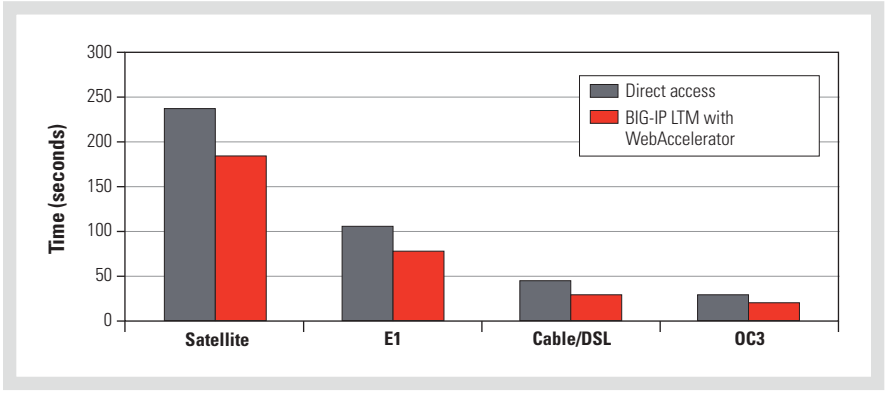

Figure 5. *First-visit performance under each WAN condition*

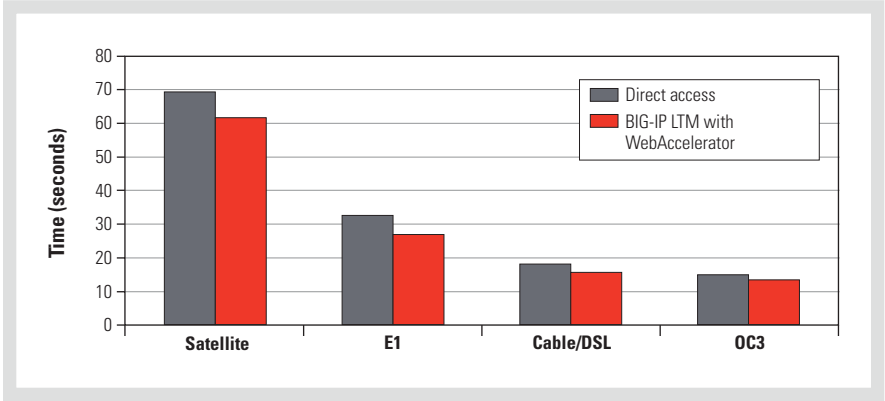

Figure 6. *Repeat-visit performance under each WAN condition*

Morten Loderup is a systems engineer at the Dell SAP Competence Center. He has a bachelor's degree from Brigham Young University, and an M.S. in Computer Information Systems and an M.B.A. from the University of Phoenix.

Fred Johnson is a strategic partner engineer with F5 Networks dedicated to Dell Labs. He has 20 years of experience in IT and has a B.A. in Psychology from the University of Texas at Austin.

Ron Carovano is a business development manager with F5 Networks responsible for the global relationship with SAP. He has 20 years of experience in high-tech industries, has a B.S. in Electrical Engineering and an M.B.A. from the University of Florida, and is listed as an inventor on 14 U.S. patents.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell and SAP:
DELL.COM/SAP

F5 and SAP:
www.f5.com/sap

F5 and Dell:
www.f5.com/dell

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

Slow apps mean slow business.

Lose application performance and you lose revenue, reputation, and opportunity. It's time to take care of business with a more agile infrastructure that puts you in control of your application delivery.

www.f5.com/dell

f5 IT agility. Your way.

OPTIMIZING EFFICIENCY AND VALUE WITH VOLUME SOFTWARE LICENSING AND PERIPHERALS FROM DELL

Dell™ volume software licensing, laser printers, and energy-efficient flat-panel monitors advance enterprise efficiency and user productivity. At the same time, on-demand IT asset tracking provides usage and license reconciliation reports that help simplify management and significantly reduce operating expenses.

Many organizations respond to changes in the business environment by provisioning equipment on an ad hoc basis, often introducing multiple suppliers to meet various software and hardware requirements as they come along. However, a diverse collection of IT assets introduces layers of complexity that require valuable resources to deploy, monitor, and manage. Hidden costs—for power consumption, toner supplies, and more—often have an unexpectedly negative impact on budgets, and the time-consuming process of negotiating software licenses can be hampered by uncertainty that may result in significant overspending.

Dell offers a complement of volume software options, laser printers, flat-panel monitors, and other peripheral devices that help address these challenges. Taking advantage of key interoperability features and tools designed to simplify asset management can help organizations optimize the performance and efficiency of their IT investments.

SIMPLIFYING VOLUME SOFTWARE LICENSING

To help organizations streamline the software acquisition process, Dell offers an extensive software publisher selection—more than 150,000 titles—together with volume software licensing (VSL) expertise for assessing requirements and recommending the appropriate software licensing arrangements. Because Dell negotiates directly with more than 2,000 software publishers, VSL enables organizations to benefit from reduced costs and the convenience of a single point of contact. Through a vast network of international locations and partners, Dell is able to provide service and support around the world, with the convenience of local costs, currency, and language.

Dell also offers a Custom Factory Integration (CFI) service, which installs software purchased through Dell so each newly purchased Dell system arrives ready to plug in and power up. CFI installation speeds deployment and helps organizations avoid configuration issues, errors, and help-desk tickets.

Helping enterprises understand exactly what IT assets exist within their organization and how much those assets are used, and helping ensure that assets are properly licensed, are key aspects to simplifying IT asset management. Dell offers an on-demand software-as-a-service (SaaS) asset management solution, eSMART, that features patented tracking technology and an easy-to-use, wizard-driven inventory. Usage and compliance reports in eSMART help manage software and hardware assets efficiently. For example, the eSMART License Reconciliation report displays the variance between purchased and installed software to help ensure compliance with software licensing and manage spending.

REDUCING ENVIRONMENTAL IMPACT AND ENERGY CONSUMPTION

The energy-efficient, high-performance Dell G-Series flat-panel and wide-screen flat-panel monitors combine recycled materials and other eco-friendly components with LED technology. G-Series monitors are compliant

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

with U.S. ENERGY STAR program requirements, and use a high-efficiency power supply unit and white LED (WLED) backlight technology. In addition, G-Series monitors help reduce energy consumption by up to 60 percent[1] through innovative power management features, including the following:

- **PowerNap screen saver:** Dims the monitor based on preferences and avoids drawing power from the host computer in sleep mode
- **Ambient light sensor:** Adjusts on-screen brightness automatically
- **Dynamic dimming:** Dims on-screen brightness based on image brightness
- **User-defined energy modes:** Sets preferences for the correct balance of energy-efficient display and comfort

As a result, G-Series flat-panel monitors can extend IT budgets by helping to reduce energy use and enhance environmental stewardship.

PROVIDING SEAMLESS ACCESS TO COST-EFFICIENT PRINTING

Controlling printing costs is not just about acquiring quality printers at a cost-effective price—it also requires effectively operating and managing those printers. Over time, for example, the costs for printer consumables and services can balloon dramatically if not managed appropriately, representing a major proportion of overall printing costs.

Dell business printers and services can help organizations proactively reduce the total cost of printing, enabling them to reinvest those savings in other areas. For example, Dell's expert printer team can help organizations assess and streamline their printer fleet to strategically optimize costs—first by breaking down individual printing costs and identifying ways to manage assets, consumables, and warranties over the life of the printers, then by helping to consolidate assets and deploy energy-saving printers and multifunction devices. And by helping select appropriate models; helping predict hardware warranty coverage, service, and response time to repair; and offering options such as cost-effective, high-yield toner cartridges, Dell can help optimize printer acquisition, maintenance, and operating costs throughout the enterprise.

Dell business laser printers can help organizations proactively reduce the total cost of printing

In addition, each Dell laser printer carries a next-business-day on-site or advanced-exchange warranty—with options to extend the warranty or customize support—to help maximize printer uptime and minimize maintenance costs. Rapid access to a specialized service group focused on business printers and multifunction devices helps ensure that problems can be identified quickly. And the quality, reliability, and performance of its printer line enables Dell to cover ongoing maintenance and parts throughout the warranty period to help avoid unexpected costs—even including replacement of expensive components such as the fuser kit.

HELPING ENSURE QUALITY AT KEY TOUCH POINTS

Standardizing on Dell for VSL, monitors, and printers enables organizations to take advantage of simplified management and customer service options that help keep cost of ownership low. This extensive collection is based on best-of-breed technology, drawing on key partnerships with leading specialists in software, imaging, and printing. And as business requirements change over time, Dell VSL and peripherals are designed to allow organizations to thrive, regardless of the economic environment.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell software licensing:
DELL.COM/Licensing

Dell flat-panel monitors:
DELL.COM/Monitors

Dell printers:
DELL.COM/Printers

Dell eSMART:
DELL.COM/eSMART

[1] Based on Dell G2410 typical power consumption of 20 W versus Dell E248WFP typical power consumption of 50 W.

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

THE SCIENCE OF SCALABILITY

The Texas Advanced Computing Center, a leading scientific computation facility at The University of Texas at Austin, tackles the research data explosion using high-performance Dell™ PowerEdge™ servers and massively scalable storage from DataDirect Networks.

CHALLENGE

The Texas Advanced Computing Center (TACC) required a high-speed, scalable research data repository that was friendly to both their computer center and their budget. The solution needed to not only meet current needs, but also be able to scale for the next generation of scientific research.

SOLUTION

Together with DataDirect Networks, Dell provided a highly scalable turnkey system using Dell PowerEdge servers with Intel® Xeon® processors to support the sophisticated TACC compute and visualization systems.

BENEFITS

- Massively scalable storage provides 1.2 PB of capacity, with a file system that can potentially support up to 10 PB.
- High performance, including up to 6 GB/sec of aggregate file system performance and over 1 GB/sec of per-client performance, helps support supercomputing sites worldwide across a wide range of scientific disciplines.
- Rack-dense, energy-efficient solution helps minimize physical footprint and power consumption.

Outfitted with some of the most powerful computation systems in the open science world, the Texas Advanced Computing Center (TACC) at The University of Texas at Austin is helping lead a class of research institutions that champion computer simulation as a cornerstone of the scientific discovery process—an element on par with theory and experimentation. But large-scale simulations and data analysis tools have challenged TACC with controlling the scientific data explosion. To help prepare for the next generation of scientific data resolution, TACC turned to Dell and DataDirect Networks to provide a scalable, cost-effective research storage solution designed for the petabyte era.

BUILDING MASSIVELY SCALABLE STORAGE

Because TACC runs several of the world's largest supercomputers, including the Ranger system ranked as the world's eighth-fastest clustered computer on the June 2009 TOP500 Supercomputing Sites list,[1] the need for a scalable data repository that could deliver petascale performance was critical when selecting storage for the project. Dubbed Corral, the TACC platform incorporates eight Dell PowerEdge 1950 servers, each configured with one quad-core Intel Xeon processor at 2.3 GHz, 16 GB of RAM, and Community Enterprise Operating System (CentOS) 5.3; eight Dell PowerEdge 2950 servers, each configured with two quad-core Intel Xeon processors at speeds ranging from 2.3 GHz to 3.2 GHz, 16–32 GB of RAM, and CentOS 5.3; and 1.2 PB of storage using the S2A9900 storage system from DataDirect Networks. The massive platform—four times larger than any other current data-collection resource on the U.S. National Science Foundation's distributed research network, the TeraGrid—is designed to effortlessly handle the challenges and opportunities of data-driven science.

Chris T. Jordan, a senior operating systems specialist in the Advanced Systems Group at TACC, notes that Corral complements the TACC system portfolio by enabling users to gain additional insights from the systems already in place. For example, users can access all of the Corral storage capabilities from the Ranger or

[1] For more information, visit www.top500.org.

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

> **"Dell and DataDirect Networks have always been a good choice for us to deploy and manage very scalable data sets. The solution delivered by Dell is designed to grow as we grow—and the price is right to get the most out of our budget."**
>
> —Chris T. Jordan
> Senior operating systems specialist
> in the Advanced Systems Group at TACC
> June 2009

Lonestar high-performance computing (HPC) systems and from the Spur or Stallion visualization systems.

Able to support over 1,200 hard drives behind a single storage system, the S2A9900 storage used in Corral is well suited for scale-out file systems. The open source Lustre file system uses the Dell PowerEdge servers as file system gateways and aggregates the disks in the one-to-many S2A9900 systems to deliver high-speed parallel file access to TACC researchers and supercomputing systems. This solution can scale to over 10 PB of capacity in a single file system to help satisfy the demands of the world's largest supercomputing systems. "Dell and DataDirect Networks have always been a good choice for us to deploy and manage very scalable data sets," says Jordan. "The solution delivered by Dell is designed to grow as we grow—and the price is right to get the most out of our budget."

OPTIMIZING PERFORMANCE FOR DATA-DRIVEN SCIENCE

The Corral solution is highly optimized to help serve the diverse performance requirements of many research disciplines, including climate research, high-energy physics, space and cosmological research, disease studies, and more. It is designed to deliver up to 6 GB/sec of aggregate file system performance and over 1 GB/sec of per-client performance for demanding single-stream workloads. Additional storage and gateway building blocks can be added to the Lustre namespace to increase capacity, performance, or both.

More than just a storage repository, the Corral platform enables TACC to compute and visualize data in the system. It is used worldwide by supercomputing facilities requiring greater performance than what is available from traditional network attached storage systems. DataDirect Networks also packages high-performance Dell PowerEdge servers with its award-winning storage platforms and the Lustre file system in its ExaScaler™ parallel file storage system.

MAXIMIZING SPACE AND ENERGY EFFICIENCY

Because large data sets often lead to large data storage complexities, DataDirect Networks has developed a series of advanced capabilities to help minimize the impact of petabyte-scale storage. With its innovative 4U, 60-drive storage enclosure, the S2A9900 is designed to store up to 600 hard drives in a single data center rack. And by using 2 TB Serial ATA (SATA) drives, a single S2A9900 system can store up to 2.4 PB total, and as much as 1.2 PB in a single rack.

The S2A9900 also features high levels of energy efficiency. The high-density, 60-drive storage enclosure is designed to significantly reduce the number of fans, power supplies, and I/O modules compared with systems that use low-density storage enclosures. This reduction in components also helps reduce overall power consumption—which in turn helps both reduce costs and support environmentally conscious IT practices.

DESIGNING FOR THE FUTURE

The high-performance solution from Dell and DataDirect Networks was delivered and running in less than two weeks, and researchers are now enjoying the benefits of a massive storage repository that helps provide high-performance, cost-effective access to scientific data. "Corral provides online storage at the petabyte scale—it's all online, accessible, and high-speed so that researchers can store and use much more data as part of their computation or visualization," says Jordan. Dell and DataDirect Networks have implemented a system at TACC that can grow in nearly any direction to help meet its research needs both now and in the future—enabling TACC to corral the next generation of scientific discovery data.

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINKS

Dell PowerEdge servers:
DELL.COM/PowerEdge

DataDirect Networks:
www.ddn.com

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

A MODEL OF EFFICIENCY

Dell helps Brazilian IT consulting firm CPMBraxis build an SAP® environment running on Dell™ PowerEdge™ servers and supported by a Dell/EMC storage area network and Dell PowerVault™ tape libraries—saving the company up to US$1 million while providing a model solution for customers.

CHALLENGE

CPMBraxis needed to create a powerful, scalable enterprise resource planning (ERP) infrastructure to enhance business efficiency, simplify IT, and provide a model solution for customers.

SOLUTION

The company worked with Dell to design and deploy an IT infrastructure that includes SAP applications and VMware® virtualization software running on Dell PowerEdge 6950 servers with AMD Opteron™ processors and PowerEdge 2950 servers with Intel® Xeon® processors, supported by a Dell/EMC storage area network and Dell PowerVault tape libraries.

BENEFITS

- Streamlined Dell hardware environment helps save up to US$1 million in employee time, avoid an anticipated US$150,000 per year in support costs, and reduce energy costs by 5 percent.
- Simplified infrastructure helps decrease backup time by 42 percent and cut IT administration time by 30 percent, saving approximately US$100,000 per year.
- Dell servers provide 15–20 percent better processing performance for SAP compared with other servers evaluated, providing the capacity to accommodate 50 percent growth.

It is no surprise that Brazil-based CPMBraxis—Latin America's largest IT consulting business—relies heavily on technology for its continued success. In particular, its enterprise resource planning (ERP) systems have played a vital role in the company's operations. Its existing ERP environment, however, could no longer provide the necessary performance, reliability, or scalability as the company expanded. "Our company went through three significant mergers and acquisitions in just over a year—we grew from 2,700 employees to 5,000," says Fabiano Marques, general infrastructure manager at CPMBraxis. "Our existing ERP system could not handle the increase in users or transactions. We decided to move to an SAP solution with a more powerful, robust, and scalable hardware infrastructure."

The company also realized that moving to SAP could provide an important business opportunity. "To sell SAP solutions, we need to prove to customers that we have in-house experience running SAP," says Marques. "We wanted to build a new infrastructure that could serve as a model for promoting our SAP solutions."

DESIGNING A MULTILAYERED HARDWARE INFRASTRUCTURE

The CPMBraxis IT group had previously had positive experiences with Dell, but it was the in-depth assessment from the Dell Global Infrastructure Consulting Services team that convinced them to work with Dell on the project. "We have deep expertise in IT consulting, but we had not designed an SAP infrastructure on such a large scale before. The Dell team spent significant time with us asking about our particular requirements," says Marques. "No other vendor was as thorough. We realized that Dell could offer us more than just hardware for this project—the company could become a valuable business partner as we design and sell SAP solutions for our own customers in the future."

The Dell team helped CPMBraxis design a multilayered hardware infrastructure. Dell PowerEdge 6950 servers with AMD Opteron processors provide the backbone for the SAP ERP Central Component (SAP ECC) application, while PowerEdge 2950 servers with the Intel Xeon processor 5300 series support the individual SAP applications. "The AMD Opteron processors in the 6950 servers deliver the processing power to run SAP ECC, while the Intel Xeon processors in the 2950 servers

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

> **"By choosing more reliable hardware and using Dell support, we anticipate that we will save US$150,000 per year in support costs."**
>
> —Fabiano Marques
> General infrastructure manager at CPMBraxis
> June 2009

provide outstanding, energy-efficient performance to run each of the individual applications," says Marques. "The processing performance and memory capacity of these servers give us plenty of room to add more users over the next year." The IT group also built a virtualized SAP test and development environment by running VMware virtualization software on additional PowerEdge servers.

A Dell/EMC CX3-80 storage area network (SAN) provides cost-effective, scalable storage to support growing production, testing, and development environments, while Dell PowerVault TL4000 tape libraries help reliably archive and restore a growing collection of data. "We needed something that could grow to more than 4 TB without having to make any significant changes," says Marques. "This SAN costs less per gigabyte than other solutions, and it gives us a way to support our growth for several years." The Dell/EMC SAN also offers easy integration and simplified management. "We didn't have any trouble connecting this SAN to our Dell-based SAP environment," says Marques. "Compared with some of the other SANs we have used in the past, the Dell/EMC SAN is much simpler to use. Unlike the other systems, we can manage this SAN without any specialized training."

The Dell team helped the IT group design and deploy the new environment quickly. "The Dell team was instrumental in helping us collaborate with a variety of companies and enabling us to meet our deployment goals," says Marques. "In the end, we were able to design and implement the solution in four months. The final step—setting up the hardware and deploying all the software—took just four weeks."

ENHANCING PERFORMANCE, SCALABILITY, AND EFFICIENCY

The CPMBraxis IT team has seen significant increases in performance and scalability in the new environment. "Compared with the other servers we evaluated, the Dell PowerEdge servers deliver 15 to 20 percent better processing performance for SAP," says Marques. "We also have the capacity to support more users and transactions. We could easily accommodate 50 percent growth before having to buy more hardware." The Dell tape libraries, meanwhile, have helped ease backups and accelerate data recovery. "It used to take us 12 hours to back up ERP data," says Marques. "The Dell tape libraries automate the backup process so we can save time. Now we can back up the same amount of data in just 7 hours. That's a 42 percent reduction."

The company anticipates increased efficiency with the SAP solution. "The SAP environment will enable us to streamline many business processes, and that is essential as we incorporate the three new companies into CPMBraxis," says Marques. "We estimate that we will save approximately US$1 million in employee time by moving to the SAP solution." Plus, the new infrastructure will help reduce IT burdens. "We used to have five or six administrators to manage the network and the servers. Now we can manage the infrastructure with just three people. We're saving about 30 percent of the time and money with the new environment—that's about US$100,000 per year," says Marques.

The IT group also anticipates significant savings for hardware support. "The Dell hardware is much more reliable than our previous equipment, so we expect fewer hardware problems," says Marques. "By choosing more reliable hardware and using Dell support, we anticipate that we will save US$150,000 per year in support costs."

Creating a dense, energy-efficient environment has helped save on power and cooling costs as well. "So far, we have decreased our energy consumption 5 percent by moving to Dell hardware and by implementing a virtualized server environment," says Marques. "We could save more than 10 to 15 percent on energy alone in the future."

DEMONSTRATING SAP EXCELLENCE

The new Dell environment will help CPMBraxis increase efficiency while accommodating growth, and help expand the business by offering tangible evidence of the company's IT expertise. "When we sell an SAP solution, our customers ask us whether we run our own SAP environment in-house," says Marques. "Now we can bring them to our facility and show them a powerful, effective SAP solution. Dell helped us create that infrastructure, and we look forward to working with Dell to deliver similar solutions to our customers."

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell and SAP:
DELL.COM/SAP

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SETTING THE PACE

Southern Company, the premier energy company serving the southeastern United States, saves an estimated US$1.3 million and sets an example for energy efficiency using virtualized Dell™ PowerEdge™ servers to consolidate its data centers.

CHALLENGE

Southern Company's rapid business growth had created server sprawl, threatening to outstrip the available data center space and driving up costs by consuming more energy each year.

SOLUTION

Virtualizing and consolidating on Dell PowerEdge R900 servers with the Intel® Xeon® processor 7300 series enables the Southern Company IT team to save data center space, reduce costs, and increase energy efficiency.

BENEFITS

- Virtualized Dell PowerEdge servers enable up to 26:1 server consolidation to help save data center space and improve resource utilization.
- Consolidated Dell infrastructure helps to avoid over 2 million kilowatt-hours of energy use, eliminate over 1,340 tons of CO_2 emissions, and save an estimated US$1.3 million in capital expenditures for equipment.
- Virtualization helps automate the server build process and reduce deployment time from more than a week for a physical server to one or two days for a virtual machine.

Atlanta-based Southern Company is the premier energy provider in one of America's fastest-growing regions, serving nearly 4.4 million customers across the Southeast. Even as its customer base has expanded, Southern Company has been able to keep its retail electric prices significantly below the national average while maintaining exceptional customer service. In fact, Southern Company has been listed as the top-ranking U.S. electric service provider for nine consecutive years by the American Customer Satisfaction Index.

One of the secrets of Southern Company's success is continually exploring new technologies that can increase efficiency and reduce costs. "As a power company, we know that one of the best ways to hold down costs is through energy efficiency," says Dan Traynor, IT infrastructure director at Southern Company. "We're committed to keeping our own house in order from a green perspective and providing technology that can help our customers save energy as well."

By establishing high standards for itself, Southern Company has helped set the pace for energy conservation among the nation's utilities. For example, Southern Company was one of the first to adopt automated metering infrastructure (AMI) on a large scale. The company is putting more than 4 million smart meters into the field at homes and businesses throughout its service area, with over a million installed to date. The meters send power usage information through radio towers to Southern Company's data centers for processing. The company anticipates the program will deliver huge benefits, as there is no longer a need to dispatch meter readers to businesses and residences where the AMI units have been installed. "When the system is fully implemented, not only will we see significant labor savings, we also estimate we're going to get about 500 vehicles off the road and save about 12.5 million miles of driving annually," says Traynor. "That will contribute to cleaner air, reduce energy use for travel, and lower our costs."

RAPID BUSINESS GROWTH LEADS TO SERVER SPRAWL

Southern Company found one of the biggest opportunities to save energy and reduce costs in its own data centers. Rapid annual business growth was causing

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

the Southern Company IT team to add approximately 30 percent more servers to its two main data centers each year, resulting in a mounting problem of server sprawl. Consolidation was difficult because many of the IT team's internal customers required their own dedicated servers, and concern about application conflicts kept most business units from sharing space on the same physical systems. "We needed to continually add more servers, and at the same time we weren't able to fully utilize the resources we had," says Traynor.

Virtualized Dell PowerEdge R900 servers have helped Southern Company save dramatically on energy and hardware costs

The Southern Company IT team realized that the pace of growth in the data center wasn't sustainable. Although the team had not yet run out of space, they could see that problem coming if they kept growing at the same rate. What also caught the team's attention was the rising power usage. "With more servers, the monthly bill we were paying our sister company, Georgia Power, was steadily increasing," says Traynor. "So was the environmental impact. When the company instituted a green data center program, we saw that we were consuming as much as 15 percent more kilowatt-hours and producing more of the corresponding CO_2 emissions every year."

The IT team decided to ask Dell for ideas on optimizing the servers and reducing server proliferation in the data centers. The company had standardized on Dell servers to power its Microsoft® Windows® OS–based systems several years ago and had a very positive experience. "Any time we've asked for advice in the past, the Dell team has been very knowledgeable and provided useful suggestions that helped us achieve our goals," says Traynor.

DELL SERVERS POWER VIRTUALIZATION PROJECT

With help from Dell, the IT team decided to virtualize and consolidate its data center servers using VMware® ESX technology. The powerful partitioning provided by the VMware platform would help eliminate application conflicts and enable servers to be shared. But the team also determined that it would need more powerful servers to achieve a sufficiently high density of virtual machines. The Dell team recommended that Southern Company replace its aging dual-processor servers with a quad-processor architecture. "We're on a lease agreement for our servers and have a regular refresh cycle, so we were able to begin replacing them right away," says Traynor.

The team decided to use Dell PowerEdge R900 servers with the Intel Xeon processor 7300 series to build a virtual server farm. The team had already used PowerEdge R900 servers for applications with high transaction volumes, and they were impressed by the servers' performance. "With four processors and 16 cores each, we knew the Dell PowerEdge R900 was powerful enough to host a large number of virtual servers while ensuring fast server response for our users," says Traynor. "The R900 is also more energy efficient than the servers we were replacing."

The Southern Company IT team has been able to place as many as 26 virtual machines on each of the PowerEdge R900

"The new Dell servers have enabled us to create approximately 390 virtual servers over a year and a half, avoiding an estimated US$1.3 million in capital expenditures for equipment."

—Dan Traynor
IT infrastructure director at Southern Company
May 2009

Reprinted from *Dell Power Solutions,* September 2009. Copyright © 2009 Dell Inc. All rights reserved.

servers. As a result, the team has eliminated hundreds of legacy physical servers, alleviating overcrowding and creating space in the data centers for future growth. "The increased performance provided by the Dell servers enables high consolidation ratios and gives us much better utilization of our data center resources," says Traynor.

> **"The efficiencies we've gained by consolidating on Dell contribute to the company's bottom line and ultimately help keep electric rates low for Southern Company customers."**
>
> —Dan Traynor
> IT infrastructure director at Southern Company
> May 2009

CONSOLIDATION HELPS SAVE ENERGY AND COSTS

The reduced number of physical systems and the energy efficiency of the Dell servers has helped the Southern Company IT team achieve big energy savings. As part of the company's green data center program, the team has created a software tool to aggregate all of the kilowatt-hour savings and the corresponding CO_2 emission reductions. "As a result of the virtualization, we avoided more than 2 million kilowatt-hours, worth approximately US$200,000 over a period of a year and a half," says Traynor. "We've eliminated over 1,340 tons of CO_2 emissions, which have a significant impact on the environment. That's a pretty big deal."

Southern Company also saves money because consolidating on Dell servers helps reduce the need to purchase new hardware. "The new Dell servers have enabled us to create approximately 390 virtual servers over a year and a half, avoiding an estimated US$1.3 million in capital expenditures for equipment," says Traynor. "The efficiencies we've gained by consolidating on Dell contribute to the company's bottom line and ultimately help keep electric rates low for Southern Company customers."

The IT team can now deploy a new virtual server in one or two days instead of spending more than a week to deploy a new physical server. "Virtualizing the data center enabled us to automate the server build process, and we're also more efficient because we don't have to order a physical machine," says Traynor. "By configuring new virtual machines on the Dell servers, we can more easily keep up with internal demand as the company grows—which helps keep the business unit managers happy with IT."

DELL EQUIPS SOUTHERN COMPANY FOR FUTURE GROWTH

Traynor feels that the company is now well equipped for the future. "With help from Dell, we've been able to head off potential data center space and power supply problems," says Traynor. "That gives us room to grow while saving energy and reducing costs."

Southern Company is still seeing rapid growth in business demand, but only modest growth in data center power consumption year to year. "We still have growth in our network, so we're not seeing an actual drop in the monthly power bill," says Traynor. "But we certainly have slowed the rate of increase. Today the IT department is down to single-digit growth in kilowatt-hours. Before, we were well into double digits."

As an energy utility, Southern Company is always looking for ways to use energy more efficiently and to set examples for its partners and business customers for achieving greater energy sustainability. Virtualizing on Dell servers has been a success on both counts.

> **"We've eliminated over 1,340 tons of CO_2 emissions, which have a significant impact on the environment. That's a pretty big deal."**
>
> —Dan Traynor
> IT infrastructure director at Southern Company
> May 2009

MORE ONLINE
DELL.COM/PowerSolutions

QUICK LINK

Dell virtualization solutions:
DELL.COM/Virtualization

Reprinted from *Dell Power Solutions*, September 2009. Copyright © 2009 Dell Inc. All rights reserved.

SYMANTEC PROTECTS MORE

THAN ANYONE.

SYMANTEC IS THE WORLD LEADER IN SECURITY.

Read the latest Tolly Group Report on security and learn how Symantec outperforms the competition at **securityleader.com**

Confidence in a connected world.

symantec™

© 2009 Symantec Corporation. All rights reserved. Symantec and the Symantec Logo are registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

Enterprise Linux Support

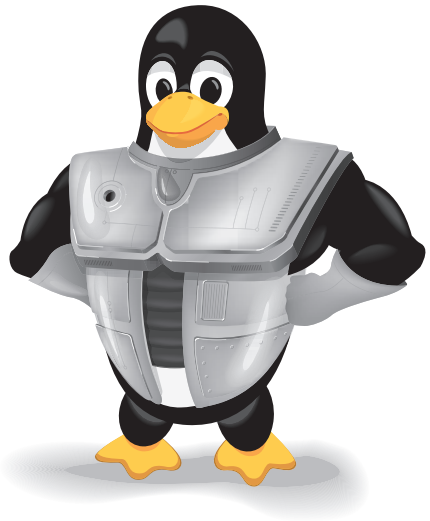

From $119
Per server/per year*

Unbreakable Linux

ORACLE®

oracle.com/goto/store
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.

Copyright © 2008, Oracle. All rights reserved. Oracle is a registered trademark of Oracle Corporation and/or its affiliates.
Other names may be trademarks of their respective owners. The Tux design is copyrighted work of Larry Ewing.