DELL™

MAY 2006 · $12.95

POWER SOLUTIONS

THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS

Enhancing Microsoft Exchange to Keep Business Flowing

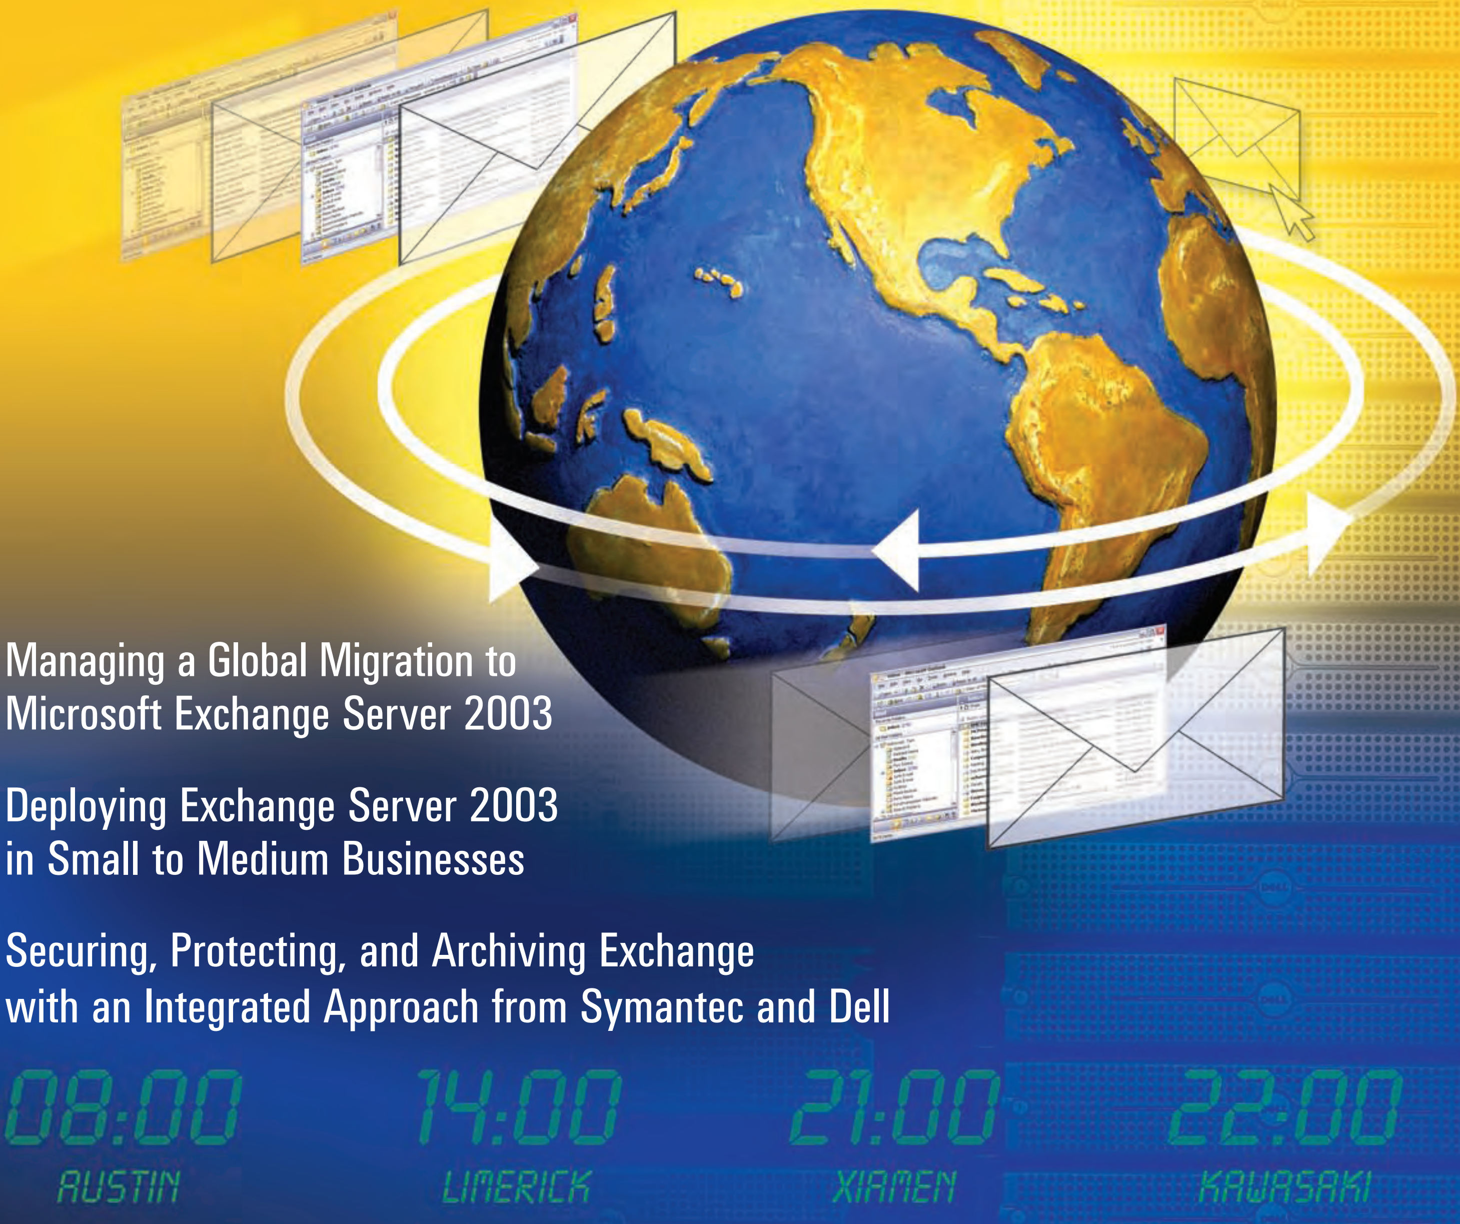

Managing a Global Migration to Microsoft Exchange Server 2003

Deploying Exchange Server 2003 in Small to Medium Businesses

Securing, Protecting, and Archiving Exchange with an Integrated Approach from Symantec and Dell

DELL POWER SOLUTIONS · MAY 2006

Enhancing Microsoft Exchange to Keep Business Flowing

Uptime expert.

It could be you.

Want to achieve a new level of reliability while increasing server throughput? Team multi-port Intel® PRO Server Adapters with onboard connections.

Improved network uptime? Yes.
Increased bandwidth and balanced traffic? Yes.
Bottlenecks? No way.

Intel® PRO
Network Connections

Whatever your infrastructure needs, Intel® PRO Server Adapters can help make network design easier. Way easier.

Learn more: intel.com/go/adapters

intel®

DELL™
POWER SOLUTIONS
THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS

MAY 2006

FEATURE SECTION: MICROSOFT EXCHANGE

COVER STORY | PAGE 10

Best Practices for Managing a Global Migration to Microsoft Exchange Server 2003

By Jesse Freund, Tyrone Freitas, and Kathryn White

The Dell IT operations team is the first to hear about lost revenue when the company's business-critical messaging infrastructure goes down. So over the years it has developed well-honed strategies for managing an enterprise-wide Microsoft Exchange Server migration. This article shares Dell's insider perspectives on how to mitigate the risks associated with such a huge, worldwide undertaking—including best practices that show how enterprises can minimize the time and expense of their own migration to Microsoft Exchange Server 2003.

TABLE OF CONTENTS

EDITOR'S COMMENTS

DELL SCALABLE ENTERPRISE TECHNOLOGY CENTER SERIES

SCALABLE ENTERPRISE

SYSTEMS MANAGEMENT

DELL POWER SOLUTIONS | THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS
WWW.DELL.COM/POWERSOLUTIONS

EDITORIAL

EDITOR-IN-CHIEF | Tom Kolnowski
MANAGING EDITOR | Debra McDonald
FEATURES EDITOR | Kathryn White
ASSOCIATE MANAGING EDITOR | Liza Graffeo
EDITORIAL ASSISTANT | Amy Hargraves
CONTRIBUTING AUTHORS | Tim Abels; Mahmoud Ahmadian; Avital Arora; Alexander Arthur; Rudramuni B.; Bala Beddhannan; Max A. Benham; Robert Bernstein; Patrick Bourke; Scott Chambers; Babu Chandrasekhar; Abhijit Chattopadhyay; Narayan Devireddy; Puneet Dhawan; Mohammad Dhedhi; Jake Diner; Steve Fagan; Anthony Fernandez; Tyrone Freitas; Jesse Freund; Jordan Gardner; Bhushan Gavankar; Kelly Harriman-Polanski; Christopher Hearn; Wolf Hengevoss; John Hull; Balasubramaniam J.; Harish Jayakumar; Dave Jaffe, Ph.D.; Garrett P. Jones; Anoop K.; Carl Kagy; Zain Kazim; Michael Kimble; Siobhan Kennedy; Chethan Kumar; Rich Lacey; Nadine Latief; Jim Lee; Min-John Lee; J. Craig Lowery, Ph.D.; Kit Lou; Cynthia Lovin; Neil MacLean; David Mar; Jon McGary; Todd Mitchell; Todd Muirhead; QuocDat Nguyen; Jason D. Norman; J. Marcos Palacios, Ph.D.; Jimmy Pike; Ujjwal Rajbhandari; David Schmidt; Rohit Sharma; Uday Shet; Amresh Singh; Sanjeet Singh; Sanjeev S. Singh; Suman Kumar Singh; Scott Stanford; Jake Swim; Teresa Taylor; Bharath Vasudevan; Mahesh Vellore; Kevin R. Webb; Phil Webster; Mike Weimann; Kathryn White; Tony Yaptangco; Brian Zhang; and Werner Zurcher

ART

ART DIRECTOR | Iva Frank
DESIGNER AND ILLUSTRATOR | Cynthia Webb
COVER DESIGN | Iva Frank

MARKETING

SPECIAL INSERTS MANAGER | Erin Stolle

ONLINE

WEB PRODUCTION | Brad Klenzendorf

Subscriptions are free to qualified readers who complete the online subscription form. To sign up as a new subscriber, renew an existing subscription, change your address, or cancel your subscription, submit the online subscription form at **www.dell.com/powersolutions_subscribe**, return the subscription reply form by surface mail, or fax the subscription reply form to +1 512.283.0363. For subscription services, please e-mail **us_power_solutions@dell.com**.

ABOUT DELL

Dell Inc., headquartered in Round Rock, Texas, near Austin, is the world's leading direct computer systems company. Dell is one of the fastest growing among all major computer systems companies worldwide, with approximately 47,800 employees around the globe. Dell uses the direct business model to sell its high-performance computer systems, workstations, and storage products to all types of enterprises. For more information, please visit our Web site at **www.dell.com**.

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

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

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

May 2006

Talk Back

We welcome your questions, comments, and suggestions. Please send your feedback to the *Dell Power Solutions* editorial team at **us_power_solutions@dell.com**.

Now you can help make your e-mail environment **even more secure.**

Upgrade to Microsoft® Exchange Server 2003 on award-winning* Dell™ PowerEdge™ servers and you can strengthen your messaging security. Enhance the mobility of your workforce. Maximize the manageability of your systems. And help your organization comply with regulations. Easily and affordably.

Click www.dell.com/exchange now to use Dell's Exchange Advisor Tool. Or get advice you can trust from your Dell Account Executive at 1-866-213-5042. And take your first step toward a customized solution.

Dell Hardware Awards: Dell PowerEdge 2850 server wins *IT Week Magazine*'s Editor's Choice Award, September 26, 2005; Dell PowerEdge 830 server wins Recommended Award by Trusted Reviews UK, October 17, 2005.

Dell/EMC AX150i networked storage array

BLADE SERVERS

APPLICATION DEVELOPMENT: SAP

MICROSOFT WINDOWS

STORAGE

STORAGE: PRODUCT SHOWCASE

DISASTER RECOVERY

DATABASES: ORACLE

SAP BUSINESS ONE

ADVERTISER INDEX

intel® Leap ahead™
Remember when
the sky was the limit?
With Intel® Software Development Products, the Swinburne Center for Astrophysics is showing today's kids a universe filled with unlimited possibilities.

See It Here First!

Check the *Dell Power Solutions* Web site for our late-breaking exclusives, how-to's, case studies, and tips you won't find anywhere else. Plus: These *Dell Power Solutions* articles are available only online at **www.dell.com/powersolutions**.

Installing Dell OpenManage Server Administrator Using DSA and DRAC 4

By Kit Lou and Zain Kazim

In Dell OpenManage 4.3, Dell OpenManage Server Administrator (OMSA) introduces native Microsoft Windows and Linux OS installation capabilities that enable quick and flexible OMSA deployment using Dell OpenManage Server Assistant and the Dell Remote Access Controller 4.

Expediting the Change-Management Process Using Dell OpenManage Server Update Utility 1.3

By Steve Fagan and J. Marcos Palacios, Ph.D.

Dell OpenManage Server Update Utility 1.3 supports the SUSE Linux Enterprise Server 9 and Microsoft Windows Server 2003 R2 operating systems—allowing administrators to customize component updates, including rollback, to help simplify operations and streamline change management.

Security Best Practices for Dell OpenManage Applications

By Jason D. Norman and Rohit Sharma

Dell OpenManage applications enable management, monitoring, and updating of Dell PowerEdge servers and attached storage. This article explores best practices for configuring and managing systems running Dell OpenManage software.

Enhancing Remote Management with the Dell OpenManage Remote Access Controller Virtual Media Command-Line Interface

By Kevin R. Webb, Avital Arora, Siobhan Kennedy, and Harish Jayakumar

The Dell OpenManage Remote Access Controller Virtual Media Command-Line Interface helps administrators use scripts to automate repetitive systems management tasks, enhancing flexibility for managing remote systems.

Dynamic DNS Updates Using the Dell Remote Access Controller 4

By Phil Webster and Brian Zhang

The Dynamic Domain Name System (DDNS) update feature of the Dell Remote Access Controller 4 (DRAC 4) combines Dynamic Host Configuration Protocol, Domain Name System, and remote access controller client registration. This article discusses the DRAC 4 DDNS feature and different interfaces used to configure it.

Enhancing Remote Management with Localization Support in the Dell Modular Server Enclosure

By Babu Chandrasekhar and Bala Beddhannan

The Dell Modular Server Enclosure contains the Dell Remote Access Controller/Modular Chassis and the Avocent Digital Access KVM (keyboard, video, mouse) switch. This article describes the localization support available for these remote-management modules.

Management by Contract in the Virtual Data Center

By Jimmy Pike and Tim Abels

Increased competition and globalization are forcing traditional data centers to provide more services while reducing resource costs. This article explores approaches that can enable contract management for disparate services in massive resource pools.

LAMP Performance on Dell PowerEdge 1855 Blade Servers Running VMware ESX Server

By Amresh Singh; J. Craig Lowery, Ph.D.; Rudramuni B.; and Scott Stanford

The Dell Scalable Enterprise Technology Center Labs team ran performance characterization and sizing tests on a Dell PowerEdge 1855 blade server running VMware ESX Server 2.5.1 software to simulate a virtualized LAMP (Linux, Apache, MySQL, and PHP) environment. These test results can help determine the appropriate number of virtual machines to host on ESX Server for a typical transactional LAMP-based workload.

Live Migration with Xen Virtualization Software

By David Schmidt and Puneet Dhawan

For enterprises planning to evaluate Xen in their own environment, this article outlines an evaluation deployment of a LAMP application stack for a live migration setup on a virtual machine using Xen open source software.

Using VMware ESX Server Virtual CPU Shares in a Microsoft Outlook Web Access Environment

By Amresh Singh and Scott Stanford

Dell engineers explored virtual CPU resource-management options and the benefits of using a dynamic resource-allocation mechanism to help meet quality of service requirements for various Internet mail protocol scenarios. In this study, the test team ran the Microsoft Exchange Server 2003 Outlook Web Access front-end service on VMware ESX Server virtual machines that were configured on a Dell PowerEdge 2850 server.

Using Dell OpenManage for Easy Integration and Management of SAS and SATA Storage Hardware

By Nadine Latief and Teresa Taylor

Leveraging advances in Serial Attached SCSI (SAS) technology, Dell OpenManage software can help simplify the configuration and management of SAS and Serial ATA (SATA) storage hardware. This article discusses how the enterprise storage infrastructure can benefit from enhanced Dell OpenManage features.

Migrating from Microsoft Windows 2000 Server to Windows Server 2003 with SP1 on Dell PowerEdge Servers

By Bhushan Gavankar

Microsoft Windows Server 2003 with Service Pack 1 (SP1) offers enhanced file, print, application, Web, and communication services. This article explains best practices for determining system requirements and upgrading to Windows Server 2003 with SP1.

Your Information, Your Way

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

Message As Medium

The Internet, e-mail, instant messaging, and blogs were nonexistent in 1964 when Canadian-born Marshall McLuhan published the book *Understanding Media: The Extensions of Man*—a work that is widely regarded as foretelling the cultural impacts of the evolving electronic world. In that book, and in the more widely known follow-on published in 1967, *The Medium is the Massage,* McLuhan coined the phrase "the medium is the message" and propelled those words into the lexicon of popular culture. Some 40 years later in 2006—with e-mail firmly entrenched as a mainstream communications vehicle and global e-mail traffic expected to exceed 160 billion messages per day[1]—it may be more appropriate to say the e-mail message is the medium.

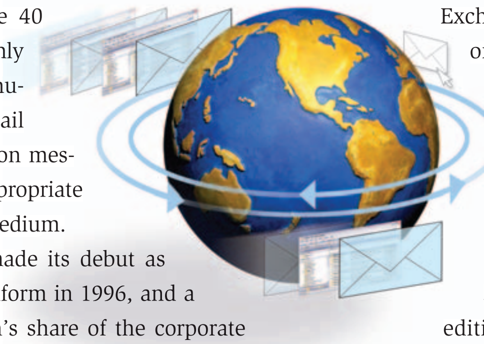

Microsoft® Exchange Server made its debut as an enterprise-ready messaging platform in 1996, and a decade later it commands the lion's share of the corporate e-mail market. Acknowledging the pervasiveness of e-mail and the prominence of Exchange as it reaches its 10th anniversary, this May 2006 issue of *Dell Power Solutions* features the latest guidance for Exchange-based messaging platforms. In our cover story, "Best Practices for Managing a Global Migration to Microsoft Exchange Server 2003," we offer a unique portal into the planning and execution of Dell's internal migration to the latest version of Exchange—as told by the Dell IT management team that orchestrated a worldwide migration for an active production environment with more than 100,000 mailboxes. Read about the Dell IT team's top 8 strategies honed during the migration, including their "pilot extensively, migrate aggressively" approach, and much more from an insider's perspective.

Of those approximately 160 billion messages estimated to be sent per day in 2006, approximately 67 percent are expected to be spam.[1] Two articles also featured in the Microsoft Exchange section in this issue address methods for securing, protecting, and archiving Exchange environments. In the first article, "Introducing Symantec Email Security and Availability for Microsoft Exchange," a layered hierarchy approach is prescribed for implementing e-mail security mechanisms at key points in the network. And the second article, "Providing Multi-Tiered Security for Microsoft Exchange Environments," delves into details on integrated solutions from Symantec and Dell that address pressing concerns about spam and virus-laden e-mail traffic.

As yet another extension of the message as medium, we are pleased to announce the availability of syndicated Web content via RSS (really simple syndication) feeds for the online edition of *Dell Power Solutions*. You can subscribe your favorite RSS reader to the new feeds by going to www.dell.com/powersolutions or www.dell.com/rss. As soon as we publish our quarterly print editions and timely Web-exclusive postings to our Online Extra section, you will have real-time access to *Dell Power Solutions* content.

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

[1] "Exchange Server Market Share Statistics, 2005" by the Radicati Group, July 19, 2005; www.microsoft.com/exchange/evaluation/compare/market_share.mspx.

Best Practices for Managing a Global Migration to

Microsoft Exchange Server 2003

The Dell IT operations team is the first to hear about lost revenue when the company's business-critical messaging infrastructure goes down. So over the years it has developed well-honed strategies for managing an enterprise-wide Microsoft® Exchange Server migration. This article shares Dell's insider perspectives on how to mitigate the risks associated with such a huge, worldwide undertaking—including best practices that show how enterprises can minimize the time and expense of their own migration to Microsoft Exchange Server 2003.

BY JESSE FREUND, TYRONE FREITAS, AND KATHRYN WHITE

Related Categories:

Best practices

Case study

Dell PowerEdge servers

Dell Services

Microsoft Active Directory

Microsoft Exchange

Microsoft Windows Server 2003

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

E-mail and messaging are indispensable to everyday business at Dell. So suffice it to say, when Dell decides to go through a Microsoft Exchange Server migration, it is a big deal. For a sense of the challenge, consider that Dell's worldwide messaging infrastructure includes more than 100,000 mailboxes and more than 250 Microsoft Exchange servers. This article provides an insider's view of how the Dell IT operations team managed a seamless, worldwide migration to Microsoft Exchange Server 2003—and shares best practices designed to help enterprises achieve such a huge undertaking without disrupting the flow of vital business operations.

To date, Dell has gone through several Exchange Server migrations. Most recently, the company began a worldwide migration from Exchange Server 2000 to Exchange Server 2003. Within six weeks, the Dell IT operations team successfully migrated all systems in the Americas region to the new messaging platform, the first phase of the worldwide migration. What's more, the project included a complete refresh of the hardware platform to the latest Dell™ PowerEdge™ 6800 servers, the Microsoft Windows Server™ 2003 OS, and application layers.

Some might wonder why Dell performed a complete hardware refresh instead of the seemingly easier in-place upgrade. Having gone through multiple Exchange Server migrations, the IT operations team at Dell has learned that migrating to a completely new messaging platform offers significant benefits. For example, a complete refresh helps enterprises to reduce downtime by cutting over to a preconfigured hardware infrastructure and rolling back the environment easily if something goes wrong.

Based on first-hand experience migrating to Exchange Server 2003 on Dell's business-critical messaging infrastructure, the Dell IT operations team has developed expertise that can help mitigate the risk associated with such a huge undertaking. This article shares that knowledge by offering well-honed strategies and best practices for managing a global Exchange Server migration (see Figure 1).

Strategy 1: Get buy-in from dependency groups

Before any piloting can take place, the leaders of the migration project must identify dependency groups and ensure buy-in. Before moving forward, the Dell migration team found it extremely helpful to explain the scope of the project and solicit involvement so participants did not feel left out or steamrolled. This step alone can prove crucial to the success of any application rollout—all the more when it involves something as complex as an enterprise-wide messaging environment.

At Dell, the IT operations team helped identify the hardware platform as well as the dependency groups that needed to support the Exchange Server migration—including the help desk, data center logistics, server operations, monitoring, and backup and recovery teams, as well as all the necessary regional IT managers. Then the project team met with all the dependency groups to ask for their help and involvement in making the migration a priority. Ultimately, much time was spent coordinating resources, but doing this work up front saved considerable time and money once the project got started.

While the process of garnering active support from the busy dependency groups at Dell headquarters in Austin, Texas, had its challenges, global participation was even trickier to coordinate. In Europe alone, Dell has 25 regional sites spread across several time zones. Plus, Dell schedules regular change moratoriums that vary from region to region. Although the task was daunting, the Dell migration team worked tirelessly to ensure worldwide interaction in order to devise an enterprise-wide project plan that addressed the requirements of all the dependency groups and regional offices. Why? The stakes were too high to do otherwise.

Strategy 2: Define primary business requirements

To help ensure that the project makes financial sense, organizations should define business requirements before the technical work begins. A slew of competing business requirements must be distilled into crucial, achievable needs that are stated in the simplest possible terms. For example, Dell had two primary business requirements for its Exchange Server 2003 migration. First, the environment had to be stable—at the very least, as stable as the Exchange 2000 Server environment. Second, the migration had to be seamless—users involved in vital business processes should not be aware they are being migrated.

To meet these requirements, the Dell IT operations team undertook several basic but important steps. First, it worked to define precisely what a stable environment would look like based upon its own historical data and the performance thresholds that Microsoft provided. The Dell team scheduled mailbox migrations from 1 A.M. to 5 A.M. local time, so the move would take place when most users were asleep. Finally, the migration started conservatively and scaled up over time. Exacting attention to detail enabled the Dell IT operations team to meet its business requirements cost-effectively and keep to its project time line.

Strategy 3: Leverage vendor relationships

Long before embarking on an extensive migration project, enterprises are well advised to cultivate a tight relationship with their messaging solution vendor and key third-party application partners because a test plan that includes insider input is likely to catch potential complications before they turn into problems. Dell benefited from its relationship with Microsoft in this way, utilizing the partnership to ease the challenges of its Exchange Server 2003 migration project.

Leading up to the migration, the Dell IT operations team met on a weekly basis to develop test plans, evaluate criteria, and establish performance baselines. Once the project started, the weekly meetings became sit-downs every other day, and frequent checkpoints uncovered hiccups before they could threaten the project. During this time, the Dell team also solicited feedback from Microsoft to help ensure that the test plan identified and resolved issues that otherwise might have interfered with smooth business operations.

Strategy 4: Profile the healthy pre-migration environment

Prior to the migration, Dell had a stable Exchange environment in place. So before the first pilot took place, the Dell IT operations team captured live data from the production environment during peak times to arrive at a baseline for healthy Exchange Server performance. Besides its internal analysis, the Dell team used the Microsoft Exchange Server Best Practices Analyzer (ExBPA) tool to help define what constituted a healthy Exchange Server environment.

ExBPA is designed to help administrators determine the overall health of their Exchange servers and topology. Before starting its first pilot, the Dell IT operations team used ExBPA to take a snapshot of the company's Exchange Server environment. In addition, the team used ExBPA to evaluate the environment during different stages of the migration project. ExBPA proved to be a crucial tool for enforcing standards and uncovering inconsistencies, as well as prioritizing potential problems according to criticality.

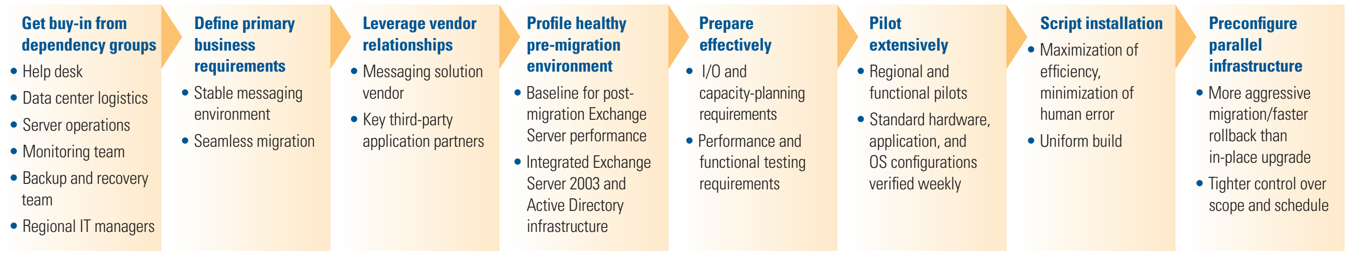

Get buy-in from dependency groups
- Help desk
- Data center logistics
- Server operations
- Monitoring team
- Backup and recovery team
- Regional IT managers

Define primary business requirements
- Stable messaging environment
- Seamless migration

Leverage vendor relationships
- Messaging solution vendor
- Key third-party application partners

Profile healthy pre-migration environment
- Baseline for post-migration Exchange Server performance
- Integrated Exchange Server 2003 and Active Directory infrastructure

Prepare effectively
- I/O and capacity-planning requirements
- Performance and functional testing requirements

Pilot extensively
- Regional and functional pilots
- Standard hardware, application, and OS configurations verified weekly

Script installation
- Maximization of efficiency, minimization of human error
- Uniform build

Preconfigure parallel infrastructure
- More aggressive migration/faster rollback than in-place upgrade
- Tighter control over scope and schedule

Figure 1. Best-practice considerations for managing a global Exchange Server migration

As part of its analysis, the Dell team identified stale data that did not need to be migrated to the Exchange Server 2003 environment. The migration team worked closely with local administrators to uncover swollen mailboxes that had not been logged in to for months and public folders that could be archived and removed. To avoid the hassle and expense of moving junk material to the new environment, the team took an aggressive approach toward consolidating and removing seldom-used mailboxes and public folders.

Also, due to the tight integration between the two systems, the Microsoft Active Directory® service is critical to the Exchange environment. For instance, part of the Exchange Server 2003 installation involved a schema extension that mandated changes to the Active Directory infrastructure. To help ensure smooth integration between Exchange Server 2003 and Active Directory, the Dell project team made sure to perform the necessary installations and groom the Active Directory data.

Strategy 5: Prepare to architect and pilot test effectively

Every time data is read from or written to Exchange, disk I/O is generated. One of the biggest challenges that organizations face when configuring Exchange Server is understanding the I/O requirements. So the migration team carefully assessed the speed, performance, and number of disk spindles needed in Dell's messaging environment. Furthermore, the team explored capacity-planning issues to help ensure that the system would be designed to accommodate future growth quickly and flexibly. Finally, the team developed specific item-retention policies to govern how long different types of data should be stored. Without retention policies, the size of Dell's Exchange databases could quickly double or triple.

Once the storage requirements had been determined, the migration team developed an extensive test plan. Again, vendor involvement played a key role, and Microsoft proved instrumental in helping Dell develop a comprehensive test methodology. In addition, the Dell team turned to a number of software tools to help ensure thorough functional and performance testing.

To perform functional testing, the Dell migration team relied on the Microsoft Exchange Server 2003 Load Simulator (LoadSim) stress-testing tool and the Microsoft MailStorm utility. Using LoadSim and MailStorm, the test team populated the Active Directory with user accounts and mailboxes. The tests sent multiple messaging requests to the Exchange server to simulate a typical mail load. This allowed the Dell team to evaluate how a server running Exchange Server 2003 responded to large e-mail loads. Ultimately, LoadSim and MailStorm proved to be valuable tools for right-sizing servers and validating the deployment plan.

To test performance and I/O requirements, the migration team turned to the Microsoft Exchange Server Jetstress Tool. Specifically, Jetstress created a simulated Exchange database and stressed the storage with virtual users in order to monitor performance. In the end, Jetstress helped ensure that the Exchange Server 2003 disk subsystem was adequately sized to meet the desired performance criteria. Plus, using Jetstress, the Dell migration team was able to determine the system breaking points before loading servers with actual users.

In addition to testing the Exchange Server infrastructure with Microsoft tools, the Dell migration team performed extensive third-party application testing to verify that the Exchange Server 2003 environment could support functionality in the existing Exchange 2000 Server environment. To accomplish this task, the Dell migration team identified the applications that integrated with Exchange—including mobile messaging, antivirus, electronic fax, backup and recovery, and Simple Mail Transfer Protocol (SMTP) gateway applications—and performed independent functional testing. In the end, the extensive testing of third-party applications paid off because it identified several applications that required reengineering or vendor involvement to provide seamless integration with the Exchange Server 2003 environment.

Strategy 6: Pilot extensively and migrate aggressively

With the test plan in hand, the migration team performed regional pilots to break the testing into phases, as shown in the Figure 2 schematic of a best-practices enterprise messaging architecture. The first test phase encompassed the Americas, the second phase correlated to the Asia-Pacific region, and the third phase included Europe, the Middle East, and Africa (EMEA). The idea behind this phased approach was to account for regional differences by having the local personnel drive the regional pilots. In each region, local administrators determined the location, size, and success criteria for each pilot.

In addition to regional pilots, the Dell team performed functional pilots for each type of Exchange server involved in the migration. The front-end servers, the mailbox servers, the public folder servers, and the bridgehead servers (which shuttle mail from region to region) were each isolated and tested individually. Again, the goal was to perform the pilots in manageable phases and help ensure success in each of the functional areas.

After the pilot phase, the Dell IT operations team invoked the mantra *pilot extensively and migrate aggressively* to guide the graduated cutover to Exchange Server 2003. During the first week, 800 users a night were migrated to the new platform. After everything went according to plan, 1,600 users a night were migrated during the second week of the project, and by the fourth week, the Dell IT operations team hit its maximum migration rate of 2,400 users per night.

The aggressive migration was made possible by the parallel hardware infrastructure Dell decided to install, as well as the multithreaded Move Mailbox tool in Exchange Server 2003. The Move Mailbox tool permitted the Dell team to schedule the moves during off-peak hours and it allowed administrators to migrate mailboxes from one server to many servers or from many servers to one server, which made it possible to break down the entire e-mail infrastructure at will. Finally, after migrating a certain number of users, the Dell team ran the ExBPA tool each week to verify that the hardware and software met the configuration standards.

Strategy 7: Script the installation

A scripted installation can pay big benefits in helping to reduce the project time line and the cost to correct errors. However, before a scripted installation can take place, detailed configuration standards must be developed and implemented. The configuration standards must include a locked-down hardware profile, a standard layout for applications, and a standard OS setup, including hot fixes and service packs.

When the time came to migrate to Exchange Server 2003, the Dell team needed to install 18 hot fixes. Some of those hot fixes would have taken two to three hours to load manually, so the ability to script the installation offered definite time savings. Aside from speeding the implementation, the goal of the scripted installation was to minimize the potential for human error. For example, a local administrator's typo could lead to an inconsistent environment.

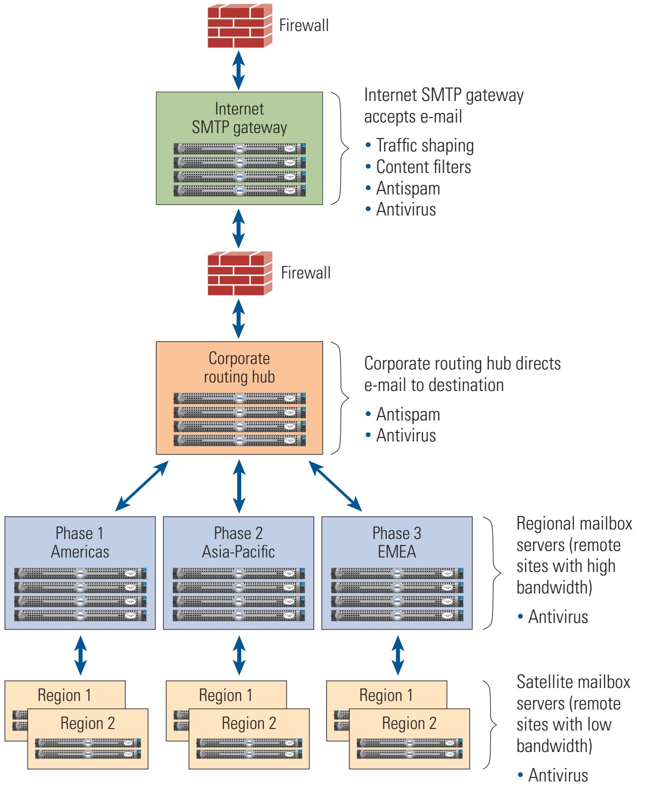

Figure 2. Enterprise messaging architecture for conducting regional pilots

The Dell installation script used the unattended installation feature built into Exchange Server 2003. The script performed the Exchange base installation with the service pack and then stopped the Exchange services to install the hot fixes. Then the script rolled through the installation in the order specified. Finally, the script configured the databases to help ensure standard naming conventions across the server and storage environment.

In the end, a scripted installation allows organizations to manage a large environment with minimal headcount. The script results in a standard build across all environments, which helps lower total cost of ownership. For example, an administrator in Limerick, Ireland, can support a server in Tokyo. Plus, the benefit of having a familiar, standardized environment helps lead to less administrative overhead and lower cost of ownership for organizations of all sizes.

Strategy 8: Consider a complete hardware refresh

Historically at Dell, the regional offices have selected their own hardware and developed their own configuration standards. Over time, in an effort to standardize the hardware environment, the Dell IT operations department listened to the problems the regional offices were having, solicited input from the regional teams, and developed hardware standards and configurations that could address the needs of offices around the globe.

When it came time to perform the Exchange migration, the Dell IT operations team found few decisions more rewarding than the decision to completely refresh the Exchange hardware foundation based upon a configuration that was standardized across the entire enterprise. Although some organizations may view the decision to deploy new servers up front as an extravagance, Dell's internal cost/benefits analysis determined that a new, standardized hardware platform could pay off handsomely by streamlining the IT change-management process—enabling fast, flexible business response anywhere in the world. In fact, establishing a standardized hardware configuration for the Exchange Server 2003 migration advanced Dell's own scalable enterprise framework, significantly enhancing the company's ability to simplify operations, improve resource utilization, and scale cost-effectively.

For example, Exchange Server 2003 boasts the multi-threaded Move Mailbox tool that enabled the migration team to move up to four mailboxes simultaneously. Thanks to the standardized hardware environment, it was possible to develop a standard configuration for the Move Mailbox tool, which allowed for multiple sessions to operate simultaneously with each session moving four mailboxes. Although it is not advisable to do so, given the parallel hardware infrastructure and the multi-threaded Move Mailbox tool, it would have been theoretically possible to move all of Dell's 100,000 mailboxes in a single night.

Using a preconfigured, parallel hardware infrastructure can help an organization to migrate on an aggressive time line. Performing a complete hardware refresh delivers another important benefit: If something goes wrong, it is easy to roll back to the previous environment—which is not the case with an in-place upgrade. Not only can the standard hardware configuration lead to a faster and more consistent rollout than an in-place upgrade, but it also offers a form of insurance should unforeseen problems arise during the migration.

Monitoring post-migration Exchange Server performance

To help ensure that the Exchange Server environment would continue to perform up to specification after the migration, the Dell IT operations team determined the areas that needed to be monitored and established performance and capacity thresholds. Then the team scripted a tool to monitor issues such as Remote Procedure Call latency, disk I/O thread counts, and log record stall. The script runs on a regularly scheduled basis, collecting data about performance thresholds during the peak hours between 7 A.M. and 3 P.M. local time. Analyzing this data helps the IT operations team stay on top of service-level agreements and address problems such as a bad hard disk or overloaded server.

To wit: capacity issues must be monitored closely. Through load testing, the migration team determined—due to varying usage patterns among users—that certain PowerEdge 6800 servers could handle fewer users before performance issues began to appear. So Dell keeps the user count for those servers at a level 20 to 30 percent below the norm, to allow excess capacity for spikes in the current load as well as for future growth. Similarly, the team keeps a close watch on database capacity. Dell strictly adheres to a 12 GB limit on Exchange Server databases and runs a capacity-monitoring program every eight hours that is specifically designed to let administrators know if a database has grown beyond its specified capacity threshold. As a best practice, Dell keeps its Exchange Server databases small to affect as few users as possible should an outage occur or a database become corrupted.

For more comprehensive systems management, the Dell IT operations team plans to replace its internally developed scripts with Microsoft Operations Manager (MOM), a comprehensive event and performance management suite for monitoring and controlling both hardware and software resources. In addition, MOM provides valuable trend analysis capabilities that enable predictive monitoring.

Mitigating risks in a global Exchange Server migration

Offering a Dell IT operations department perspective, this article aims to help enterprises avoid pitfalls typically associated with the migration to a new messaging infrastructure. By coordinating dependency groups, defining business requirements, and securing vendor involvement, enterprises of all sizes can properly prepare for an Exchange Server migration. Through meticulous architecting, extensive piloting, and aggressive migrating, enterprises can help make the transition a success. And by scripting the installation and refreshing the Exchange server hardware, organizations can minimize the time, expense, and risk associated with an Exchange Server migration as compared with an in-place upgrade.

Jesse Freund is a business and technology writer based in San Francisco. He has written about business and technology for leading publications, corporations, and organizations, including *Business 2.0* and *Wired* magazines. Jesse has a B.A. in History from the University of California, Berkeley.

Tyrone Freitas has been at Dell for more than eight years, and currently manages the Global Messaging and Directory Services group within the Dell Global IT Operations department. He attended Western Connecticut State University.

Kathryn White is the features editor for *Dell Power Solutions.* She has 25 years of development, communications, and marketing experience in the IT business. Kathryn has a B.S. in Mathematics from the University of South Carolina.

FOR MORE INFORMATION

Dell and Microsoft Exchange Server 2003:
www.dell.com/exchange

Microsoft Exchange Server 2003:
www.microsoft.com/exchange

Microsoft Exchange Server 2003 deployment tools:
www.microsoft.com/technet/prodtechnol/exchange/downloads/2003/tools.mspx

IT Executive Learning Series

By IT Leaders, for IT Leaders

Engage in an IT-to-IT Discussion

Dell IT ELS events offer a number of informative sessions hosted by Dell's top IT executives. These sessions give a behind-the-scenes look at Dell's IT organization and cover a range of topics including supply chain management, security, disaster recovery, data management, and much more*. These complimentary events provide an outstanding opportunity for you to connect with your peers and gain insight into several Dell IT best practices. The IT Executive Learning Series is an IT-to-IT event specifically designed and presented "By IT Leaders, For IT Leaders."

*Topics vary by event

What Your Peers Have Said About Dell IT Events

- "I found the series to be very informative, and more importantly, thought provoking."
- "Excellent, candid presentations…"
- "Substantially more relevant than going to a Gartner or META conference."
- "I appreciated the IT-to-IT discussion as opposed to a sales event."
- "Time well spent, and I look forward to coming back again."

Visit **www.dell.com/it** for dates, topics, and registration information

DELL™
Information Technology

Introducing Symantec Email Security and Availability

for Microsoft Exchange

The Symantec® Email Security and Availability solution for Microsoft Exchange is designed to protect both the systems and the information in Microsoft® Exchange e-mail infrastructures. This solution prescribes a layered hierarchy—implementing e-mail security mechanisms at key points within the network—to filter out unwanted messages and keep e-mail systems running efficiently.

BY WERNER ZURCHER AND GARRETT P. JONES

Related Categories:

E-mail technology

Microsoft Exchange

Security

Symantec

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

Enterprises that depend on e-mail for employee, customer, and partner communications require no-compromise security and high availability for their e-mail infrastructure as well as for the information that passes through the messaging system and is ultimately archived. However, security and availability are interdependent variables that are often achieved at the expense of one another: high security is often traded off in exchange for high availability, and vice versa.

The Symantec Email Security and Availability solution for Microsoft Exchange is designed to reduce the volume of spam e-mail, eliminate the risk of virus infection, automatically manage the e-mail life cycle through archiving, and keep an enterprise e-mail infrastructure resilient against failure. As a result, this integrated, multilayer approach can help reduce costs and simplify management of the e-mail environment and life cycle.

Figure 1 illustrates the components of the Symantec Email Security and Availability solution, which layers different types of protection at various levels of the e-mail architecture. This approach focuses on server products and does not encompass desktop protection options.

Increasing e-mail security

The first layer of the Symantec Email Security and Availability solution for Microsoft Exchange is designed to provide e-mail security by reducing incoming e-mail volume, securing the perimeter, and filtering e-mail internally.

Reducing e-mail volume

The first line of defense against unwanted e-mail content is deployed outside the messaging infrastructure—before the data can affect internal servers, including the Simple Mail Transfer Protocol (SMTP) mail gateways. This first line of defense is provided by the Symantec Mail Security 8160 appliance.

This appliance, which integrates Symantec software on a Dell™ PowerEdge™ 1850 server, helps prevent spam by evaluating sender reputation and using traffic shaping on the inbound SMTP stream. If a significant amount of incoming e-mail is spam, traffic shaping can help reduce the overall e-mail volume by blocking the transmission of SMTP traffic from known spammers and Internet hosts that they have commandeered, without affecting other e-mail transmissions.

Securing the perimeter

The network perimeter is a critical area for enhancing network security. As the second line of defense after the Symantec Mail Security 8160 appliance, the perimeter solution—incorporating Symantec software as well as Symantec Mail Security 8200 Series appliances—combines state-of-the-art spam and virus detection with turnkey operation. The following Symantec Mail Security 8200 Series appliances are available:

- **Mail Security 8220:** Built on a Dell OptiPlex™ desktop and designed for environments with less than 100 users
- **Mail Security 8240:** Built on a Dell PowerEdge 850 server and designed for environments with 100 to 1,000 users
- **Mail Security 8260:** Built on a Dell PowerEdge 1850 server and designed for environments with more than 1,000 users

Symantec's perimeter components include a mass-mailer cleanup capability to remove entire messages and prevent unnecessary virus notifications based on the presence of a mass-mailer worm; the ability to block e-mail based on customizable rules; the ability to process spam based on antispam engine verdicts (for example, deleting spam messages but quarantining suspected spam messages for further review); and a Web-based Spam Quarantine server, which removes spam messages from the messaging environment but makes them available for further processing and review. By blocking spam and other unwanted e-mail messages, Symantec's perimeter protection reduces the volume of e-mail that must be distributed and processed internally.

Filtering e-mail internally

While Symantec's perimeter protection plays a key role in minimizing the negative impact of Internet e-mail traffic, Symantec Mail Security for Microsoft Exchange is designed to keep internal message traffic free of malicious or inappropriate content. This software is tightly integrated with Exchange using Microsoft-supported application programming interfaces, helping to ensure maximum performance and minimum conflicts with the underlying messaging architecture. Similar to the perimeter protection components, Symantec Mail Security for Microsoft Exchange leverages the same core antivirus technology, updates, and response mechanism. In addition to core scanning services, Mail Security for Microsoft Exchange offers content inspection capabilities, such as subject-line and message-body filtering, attachment stripping, and restricted message size.

Archiving e-mail and increasing content accessibility

The second layer of the Symantec Email Security and Availability solution for Microsoft Exchange archives e-mail. Archiving helps ensure that e-mail content is accessible and available whenever it is needed. This layer utilizes VERITAS Enterprise Vault™ software to archive, index, search, and retrieve information. The archiving process is designed to be automatic and seamless. Enterprise Vault implements user-defined policies to automatically archive e-mail, file system content, instant messaging, and other content from operational storage locations to a cost-effective online vault—without affecting end-user access to the data. Users can access archived information directly from their e-mail clients or Web browsers and can access it while offline by using the Offline Vault option.

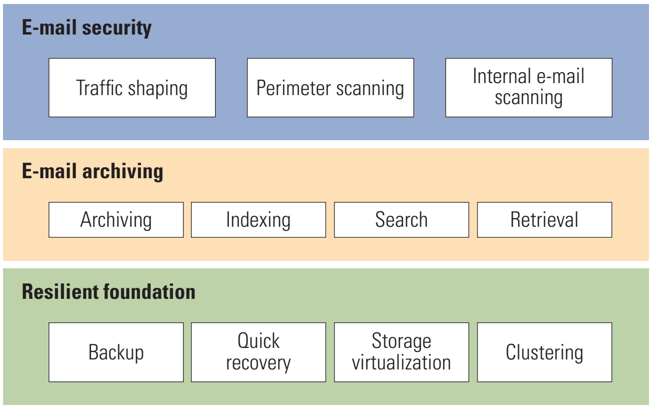

Figure 1. Components of the Symantec Email Security and Availability solution

IT administrators can automatically discover, collect, migrate, and eliminate Microsoft Personal Folders (.pst) files by moving the content to the vault. Enterprise Vault can also archive Exchange Journals and Public Folders, in addition to Microsoft Exchange mailboxes. Archived data is automatically compressed, duplicate copies are removed, and data is retained based upon business policies. Users, compliance departments, legal professionals, and corporate risk management functions can securely and easily search through messages, files, and attachments.

Message archiving using Enterprise Vault can provide benefits in three core areas:

- **Enhanced e-mail availability:** Enterprise Vault is designed to reduce the amount of data stored in primary messaging servers and file servers, helping to minimize corruption and performance problems that are observed when these servers reach capacity thresholds. By archiving data for long-term retention and providing search capabilities, Enterprise Vault can help maintain end-user access to data.
- **Minimized e-mail cost:** Enterprise Vault is designed to reduce primary storage costs throughout the e-mail environment by archiving outdated or infrequently accessed data to low-cost storage. This approach helps reduce backup costs significantly because archived data does not require frequent backups. Support and migration costs also can be minimized through elimination of e-mail quotas and .pst files and reduction in the amount of data to be moved during upgrades and server consolidation.

- **Compliance with e-mail retention policies and regulations:** Enterprise Vault is designed to facilitate e-mail retention by following defined business rules to meet legal discovery and corporate or regulatory requirements.

Enterprise Vault can be integrated with Symantec Mail Security appliances and software. If an enterprise is legally required to keep a copy of all the e-mail it receives, a Web-based Spam Quarantine server that is fed spam and other junk e-mail messages by Symantec Mail Security can deliver the junk e-mail to Enterprise Vault for journaling. In this way, Symantec Mail Security 8200 Series appliances or Symantec Mail Security software can forward all SMTP e-mail communication to Enterprise Vault servers for journaling.

Building a resilient foundation

The third layer of the Symantec Email Security and Availability solution for Microsoft Exchange is designed to enhance e-mail system availability. Symantec offers various products to match varying organizational needs for information availability. Symantec Backup Exec™ and VERITAS Storage Foundation™ software form the lower tiers of the Symantec availability hierarchy. Backup Exec and Storage Foundation are designed to enable near-instantaneous recovery from storage device failures and quick recovery for application logic or data corruption. Backup Exec can be used as a data backup management tool to send data to tape as usual, but it also can be used to create on-disk backups, on-disk snapshots, and backups that are staged on disk and then migrated to tape.

For enterprises that require comprehensive protection and fast recovery when failures occur, Symantec offers VERITAS Storage Foundation High Availability (HA) for Windows. This advanced solution works with existing hardware and infrastructure components to enable cost-effective clustering capabilities designed to provide high-availability disaster recovery and business continuity for business-critical applications and databases. Alternatively, organizations may consider deploying Microsoft Cluster Service (MSCS), a component of Microsoft Windows® server operating systems designed to provide high availability.

Figure 2 provides a view of the Symantec Email Security and Availability solution in relation to the overall network topology. In Figure 2, the various tiers—network boundary, gateway, mail server, and archive—are shown in relation to the Symantec products that can be deployed at each tier.

Enhancing management of Exchange environments

Symantec Email Security and Availability for Microsoft Exchange is a comprehensive e-mail system solution that is designed to help ensure the security, availability, and resilience of e-mail systems and information, while helping to reduce the total cost of maintenance of the e-mail infrastructure. This solution takes a multilayered approach to e-mail security, incorporating antivirus, antispam, archiving, backup and recovery, and storage management capabilities.

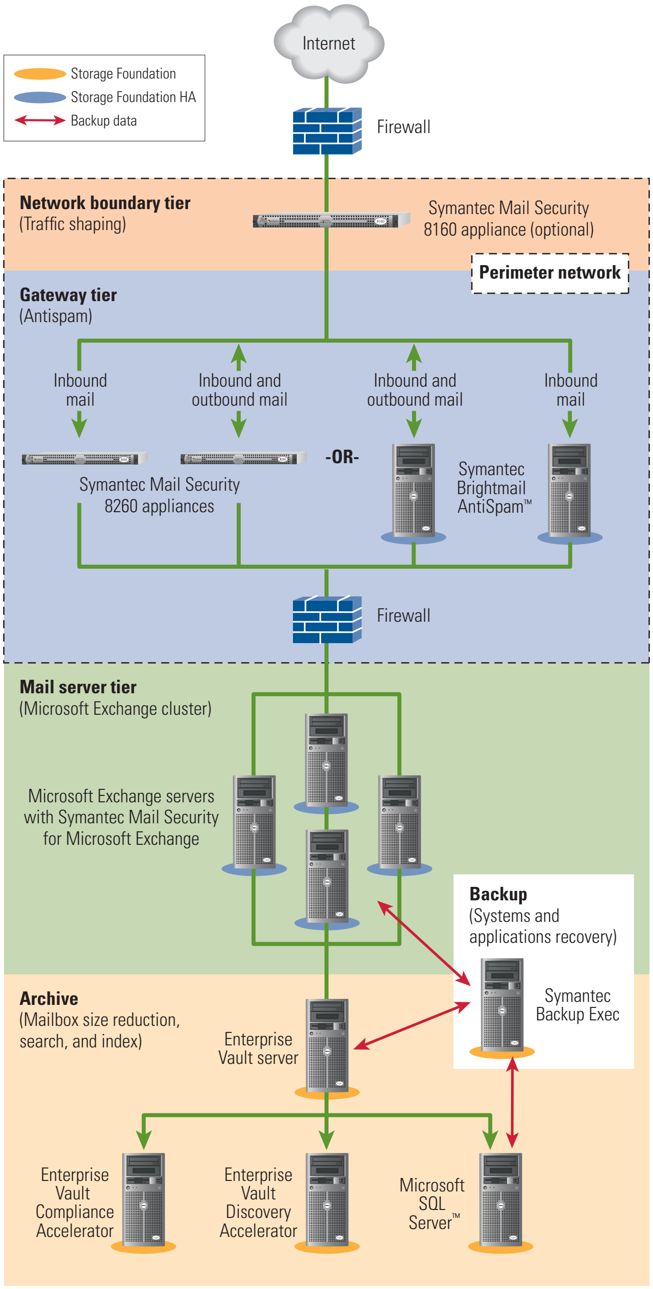

Figure 2. Network topology incorporating the Symantec Email Security and Availability solution for Microsoft Exchange

Werner Zurcher is the director of product management in the Global Solutions Group at Symantec. Werner has degrees in Electrical Engineering and Computer Science from Brown University.

Garrett P. Jones is the Symantec global alliance manager in the Dell Enterprise Product Group. Garrett has a B.A. in Business Economics from The University of Texas at Austin.

Providing Multi-Tiered Security

for Microsoft Exchange Environments

For most enterprise IT organizations, eradicating spam and virus-laden e-mail traffic can be a daunting task. By deploying a multi-tiered protection strategy involving the desktop, file server, mail server, gateway, and network boundary, administrators can provide a strong defense against attacks and disruptions. A select set of products from Symantec and Dell can help fortify enterprise messaging environments.

BY WERNER ZURCHER AND GARRETT P. JONES

Related Categories:

E-mail technology

Microsoft Exchange

Security

Symantec

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

The requirements for managing an e-mail infrastructure such as Microsoft® Exchange changed dramatically as e-mail evolved into a mission-critical application. At the same time, spam and viruses present new and constantly changing threats to e-mail security and availability. To ward off attacks and help reduce business risks, IT administrators must protect e-mail infrastructures with a combination of integrated, highly accurate antispam, antivirus, and content filtering technologies.

Preventing spam, viruses, and other unwanted content from reaching the Microsoft Exchange environment and end users can help significantly improve overall e-mail productivity, enhance network security, and reduce total cost of ownership. In addition, significant reductions of spam and viral content can also help reduce backup windows and speed data recovery.

To help ensure e-mail security and availability, Symantec best practices recommend that organizations implement a multi-tiered approach. Each tier can help reduce the potential downstream risk posed by security threats and spam. For example, Figure 1 shows the primary tiers of e-mail protection for client desktops and the Symantec® products that are available for securing each tier.

To secure e-mail systems and keep them available, organizations must be able to control and manage the flow of messaging information from start to finish. In functional terms, this means removing spam, viruses, and unwanted or unneeded content from the messaging infrastructure at the appropriate time. Multilayered defenses should complement each other by using multiple methodologies to complicate any attempts to attack. The multi-tiered strategy also can help reduce both security risks and e-mail volume while helping to ensure that messages are legitimate and "clean" before they pass to the next tier.

Network boundary tier

Large enterprise IT organizations that need to significantly reduce spam before it enters their networks may want to

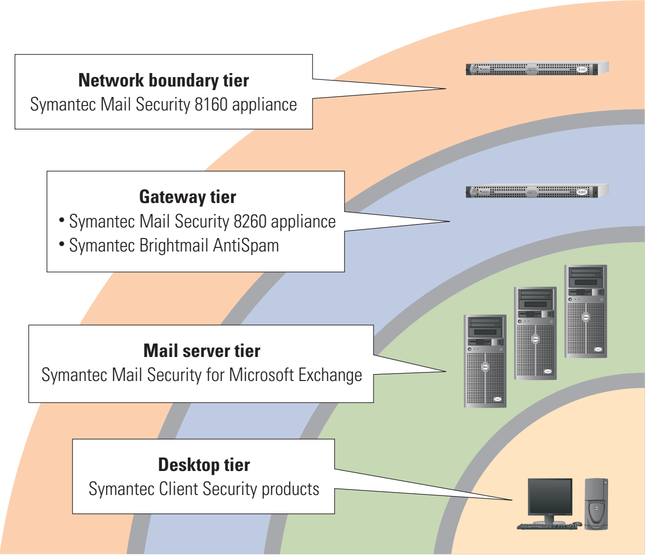
Network boundary tier
Symantec Mail Security 8160 appliance
Gateway tier
• Symantec Mail Security 8260 appliance
• Symantec Brightmail AntiSpam
Mail server tier
Symantec Mail Security for Microsoft Exchange
Desktop tier
Symantec Client Security products

Figure 1. Multi-tiered approach to e-mail protection

deploy antispam devices at the network boundary. One of the most effective ways this can be achieved is by using the Symantec Mail Security 8160 appliance. This is a network boundary device that acts as a router, not a message transfer agent (MTA), to inspect incoming Simple Mail Transfer Protocol (SMTP) traffic. The Mail Security 8160 is designed to significantly reduce the spam volume before it affects the internal network, including any SMTP gateway defenses.

Unlike traditional defenses, the Mail Security 8160 employs inspection and traffic shaping at the TCP/IP level by sampling and analyzing SMTP packets in real time based on their content, origin, and the sender's reputation. Over time, the appliance determines a sender's reputation based on cumulative history and reputation of the e-mail path itself. Once the reputation is established, incoming traffic can be "shaped" based on that reputation. Traffic shaping involves dynamically controlling the speed at which SMTP packets are accepted, and therefore controlling the number of e-mail messages that can be received from known spam senders. The net result is a significant reduction in the volume of spam reaching the mail server per unit of time, as well as a significant deterrent to spammers because spam jobs become backed up on their own servers.

Typically, traffic shaping is implemented only in large, high-volume environments (usually 2,000 or more e-mail users) that handle significant e-mail volumes. The Symantec Mail Security 8160 appliance is built using a Dell™ PowerEdge™ 1850 server.

Gateway tier

The two primary e-mail–borne threats and disruptions are viruses and spam. Several measures can be taken to prevent viruses and spam from reaching downstream servers, storage, archives, and e-mail users.

First, the most common virus content found in e-mail is the product of mass-mailer worms. These programs use e-mail addresses found on compromised systems and automatically generate e-mail messages to replicate and distribute their payload to unsuspecting users and systems. Because e-mail messages from mass-mailer worms have no intrinsic business value, they can be deleted automatically without fear of legitimate data loss. Gateway-based antivirus scanners should be able to identify and distinguish mass-mailer worms and allow administrators to delete them.

Second, spam content can be eliminated or removed from the internal mail streams to further reduce the burden on mail systems. Spam quarantines, generally housed on a server separate from the mail infrastructure, are ideal places to move unwanted spam content from active message stores (and consequently end-user mailboxes) to less-expensive media. Quarantine servers are easier to scale and maintain than mail servers because they have fewer functions. Antispam systems are not 100 percent accurate and businesses cannot risk the loss of legitimate e-mail, so spam quarantines provide a place to review spam-tagged messages. However, the reliability of the antispam system can play a significant role in reducing the amount of data that is held in quarantine and minimizing the amount of data requiring review.

Finally, an organization must not be perceived as a source of inappropriate or malicious content. All outbound e-mail should be scanned for viruses and inappropriate content. Also, organizations can put measures in place to stop unauthorized Internet e-mail (SMTP) traffic by defining network firewall rules that restrict outbound SMTP e-mail to only authorized e-mail servers. They can also establish desktop firewall rules that prevent the generation of SMTP e-mail protocol messages by end-user systems. Figure 2 shows how the multiple technologies provided with Symantec Mail Security 8200 Series

Challenge	Solution
Keeping spam and other unwanted e-mail from reaching mail servers	Brightmail AntiSpam technology uses more than 20 spam-prevention techniques to block spam. The embedded Symantec antivirus technology features real-time scanning. Virus protection capabilities also include mass-mailer cleanup, which automatically removes e-mail messages associated with mass-mailer worms.
Reducing e-mail infrastructure costs	E-mail firewall technologies, which include Directory Harvest Attack Prevention and Sender Reputation, are designed to restrict connections from spam-sending servers.
Controlling outbound content	Content-compliance features help administrators control outbound e-mail content. Besides controlling viruses, administrators can manage sensitive e-mail content and enforce content rules to conform with corporate and regulatory policies.

Figure 2. Security challenges and solutions for the gateway tier

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

> The continually changing landscape of threats—including spam, viruses, phishing, and spyware—requires tools that automatically keep up with the latest antispam and antivirus policies and rules.

appliances and Symantec Brightmail AntiSpam™ software can help address the challenges that IT organizations face when protecting the enterprise network.

A disadvantage of some antispam systems is that they provide high spam-detection rates at the expense of accuracy. The standard metrics for antispam reliability are the spam catch rate, which specifies how much spam was detected, and the accuracy rate, which shows what percentage of messages were correctly identified as legitimate (thus avoiding spam "false positives"). Symantec Mail Security 8200 Series appliances and Symantec Brightmail AntiSpam software are designed to provide highly effective and accurate antispam technology.

Furthermore, the continually changing landscape of threats—including spam, viruses, phishing, and spyware—requires tools that automatically keep up with the latest antispam and antivirus policies and rules. Symantec Mail Security 8200 Series appliances and Brightmail AntiSpam software include an integrated virus and spam signatures update mechanism that is frequently and automatically updated. Mail Security 8200 Series appliances deliver antispam, antivirus, content filtering, e-mail firewall, and quarantine capabilities and are available in the following models:

- **Mail Security 8220:** Built on a Dell OptiPlex™ desktop and designed for environments with less than 100 users
- **Mail Security 8240:** Built on a Dell PowerEdge 850 server and designed for environments with 100 to 1,000 users
- **Mail Security 8260:** Built on a Dell PowerEdge 1850 server and designed for environments with more than 1,000 users

Mail server tier

The mail server tier processes outbound e-mail and processes and stores inbound and internal e-mail. Even with solid perimeter protection in place, messaging environments require inspection of internal e-mail traffic and stored messages. This is necessary because viruses can enter through other vectors such as through personal, Web-based e-mail or removable media (for example, USB drives of computers with outdated virus definitions). Also, post-attack virus cleanup of message stores (after early-stage virus infestations) using the latest antivirus definitions is critical.

Symantec Mail Security for Exchange enables administrators to inspect content in real time as e-mail is being committed to and accessed from the Exchange Information Store—which comprises both the private and public Exchange information databases. Administrators also can conduct sweeps of Information Store content on a scheduled or on-demand basis using updated virus definitions or specific content rules that are designed to identify suspicious or inappropriate content. Figure 3 explains common challenges that IT departments face when inspecting internal traffic and how Symantec Mail Security for Exchange can help address those challenges.

Desktop tier

At the innermost tier of the network, desktop users interact with Microsoft Exchange and other inboxes. At this tier, security threats and viruses are often launched by users who remain unaware of malicious activity. Consequently, having protection at the desktop level is a critical component of a tiered defense strategy. The Symantec antivirus, anti-spyware, and personal firewall software tools—Symantec Antivirus Corporate Edition and Symantec Client Security—are designed to stop the launch of threats delivered through e-mail at the desktop tier.

Challenge	Solution
Scanning for viruses that enter the network by bypassing the network boundary and gateway tiers	Viruses can enter the network through personal, Web-based e-mail or removable media such as USB drives. Mail Security for Exchange can scan mail downstream of the gateway servers to help ensure that new threats are exposed and handled.
Ensuring redundancy in e-mail inspection	Although inbound e-mail is a common delivery mechanism, viruses can enter e-mail systems from other sources. Running defenses at the gateway can provide coverage of inbound e-mail, but not all threats can be detected and removed at that tier—virus detection and cleanup also should be performed at the mail server tier.
Preventing authorized content from being sent to unauthorized users	Typically, companies carefully secure internal Web sites from unauthorized individual or departmental access. However, information from a secured Web site can be downloaded to a desktop system and easily forwarded to virtually anybody. This possibility exposes data to unauthorized users both inside and outside the company. Mail Security for Exchange incorporates rules-based content filtering to help prevent unwanted content from entering—and confidential information from leaving—the network.
Enforcing e-mail usage policies	Companies enforce e-mail policies to prevent inappropriate language in e-mail and unwanted or oversized attachments (such as MP3 music files; AVI and other video file types; and file types commonly used for delivery of viruses, such as executables). Mail Security for Exchange is designed to enforce these policies at the mail server tier to help prevent internally introduced and inappropriate e-mail from propagating inside and outside the company.

Figure 3. Security challenges and solutions for the mail server tier

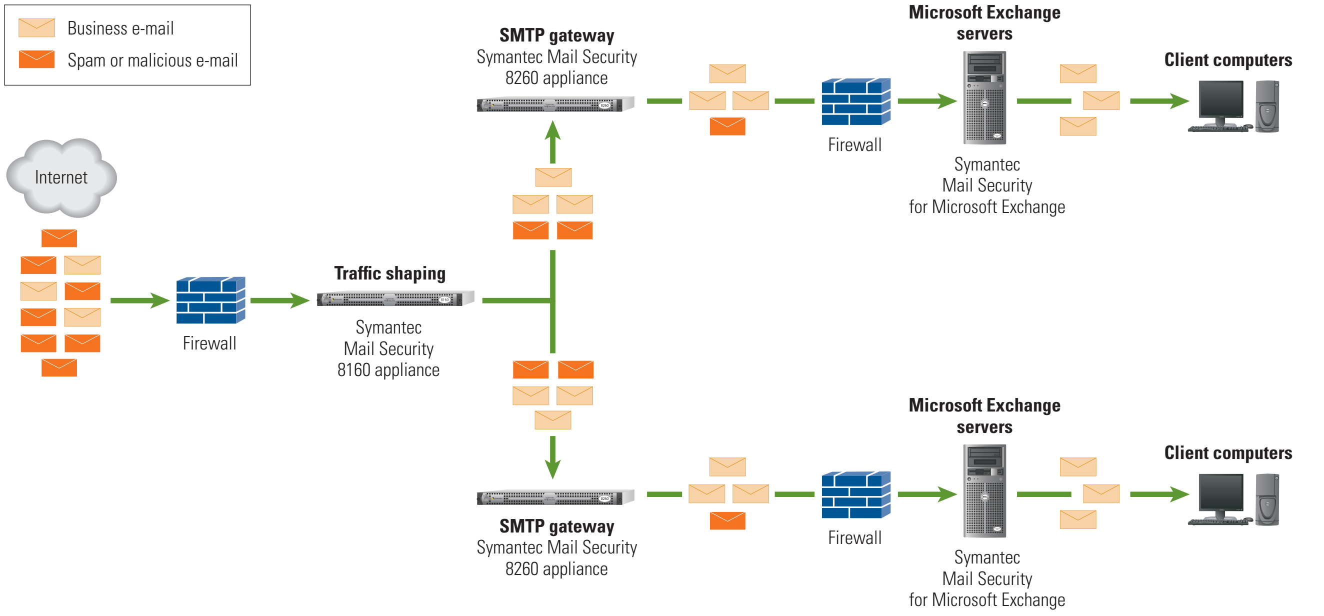
Business e-mail
Spam or malicious e-mail
Internet
Firewall
Traffic shaping
Symantec Mail Security 8160 appliance
SMTP gateway
Symantec Mail Security 8260 appliance
Firewall
Microsoft Exchange servers
Symantec Mail Security for Microsoft Exchange
Client computers
SMTP gateway
Symantec Mail Security 8260 appliance
Firewall
Microsoft Exchange servers
Symantec Mail Security for Microsoft Exchange
Client computers

Figure 4. Recommended architecture for multi-tiered e-mail security approach

Although desktop protection tools are highly customizable and individually effective, they cannot offer organization-wide protection because they protect individual desktop computers only. Comprehensive enterprise protection is possible only with a multi-tiered approach.

Architectural overview

For organizations supporting between 1,000 and 2,000 e-mail users, Symantec recommends implementing e-mail protection at the desktop, mail server, and gateway tiers. In high-volume e-mail environments of 2,000 or more users, additional protection should be implemented at the network boundary tier. Symantec and Dell recommend the following e-mail security products for each tier:

- **Desktop tier:** Symantec antivirus, anti-spyware, and personal firewall software
- **Mail server tier:** Symantec Mail Security for Microsoft Exchange
- **Gateway tier:** Symantec Mail Security 8260 appliance
- **Network boundary tier:** Symantec Mail Security 8160 appliance

Figure 4 shows the recommended architecture for a multi-tiered e-mail security approach for companies with more than 1,000 employees.

E-mail protection on multiple levels

E-mail plays a critical role in today's mission-critical enterprise applications. By filtering out spam and viruses at multiple levels of the e-mail infrastructure, enterprise IT organizations can help prevent disastrous security intrusions and help keep e-mail systems running efficiently. Symantec's offerings can be configured to address the security and availability needs of a multilayer e-mail infrastructure comprising the network boundary tier, the gateway tier, the mail server tier, and the desktop tier. This comprehensive approach provides appropriate tools to enhance e-mail security and enables high availability in mission-critical messaging environments.

Werner Zurcher is the director of product management in the Global Solutions Group at Symantec. Over the past 11 years, he has held various product management positions at VERITAS and Symantec, including Microsoft Windows® product line manager. Werner has degrees in Electrical Engineering and Computer Science from Brown University.

Garrett P. Jones is the Symantec global alliance manager in the Dell Enterprise Product Group and has also worked in the Dell Advanced Systems Group. Garrett has a B.A. in Business Economics from The University of Texas at Austin.

FOR MORE INFORMATION

Dell and Symantec:
www.dell.com/symantec

Symantec security products:
enterprisesecurity.symantec.com/esa

Microsoft Exchange Server 2003 Deployment Considerations

for Small and Medium Businesses

A Dell™ PowerEdge™ 2800 server can provide an effective platform for Microsoft® Exchange Server 2003. A team of Dell engineers tested Exchange Server 2003 performance and scalability for varying numbers of users on a PowerEdge 2800 server platform, in scenarios typical of those in a small or medium-sized organization.

BY SUMAN KUMAR SINGH, ABHIJIT CHATTOPADHYAY, AND BHARATH VASUDEVAN

Related Categories:

Dell PowerEdge servers

E-mail technology

Microsoft Active Directory

Microsoft Exchange

Microsoft Windows Server 2003

Performance

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

Today, small and medium businesses (SMBs) increasingly rely on messaging solutions. E-mail has become a mission-critical application. Whether in the office or on the road, communicating internally or with customers, enterprises depend heavily upon e-mail availability. Performance and availability of the messaging infrastructure is critical. Therefore, enterprise IT organizations must consider hardware sizing and mailbox design before the initial deployment or redeployment after hardware or software upgrades are performed.

This article examines how to select the appropriate platform for deploying Microsoft Exchange Server 2003 in SMB environments[1] and why this platform can be well suited for messaging solutions for the SMB segment. Dell engineers conducted testing in November 2005 with the intent of answering the following frequently asked questions:

- How many mailboxes can be hosted on one server?
- How much memory is needed for the Exchange server?
- If server processor utilization is very low, why is e-mail access slow?
- Is storage performance too slow?

Selecting appropriate hardware for Microsoft Exchange Server 2003

Dell provides several server platforms that can be used to deploy Exchange Server 2003 for SMBs up to approximately 900 employees. Different types of workload can be imposed by clients—consuming a combination of memory, processor, I/O, and network resources on the server. Typical client workloads for Microsoft Exchange tend to stress the memory and I/O subsystems, while exercising the processor and network subsystems to a lesser extent. Single-socket servers may not allow future scalability because they are limited in their expansion capabilities. Microsoft Exchange Server 2003 is a 32-bit application, which can only address 4 GB of virtual memory. The current version of Exchange Server 2003 does not support 64-bit versions of the Microsoft Windows® OS. Given the 4 GB memory limit and the lighter processor workload,

[1] For guidance on sizing large Microsoft Exchange installations, refer to the Dell Exchange Advisor Tool at www.dell.com/exchange.

Server	Dell PowerEdge 2800
CPU	Two Intel® Xeon® processors at 2.8 GHz with 1 MB level 2 cache
Memory	4 GB of double data rate 2 (DDR2) error-correcting code (ECC) of RAM
NIC	Two dual-port Intel 8254NXX Gigabit Ethernet* adapters
RAID controller	PowerEdge RAID Controller 4, Extended Dual Channel integrated
Internal disks	Ten 73 GB, 15,000 rpm SCSI drives
OS	Microsoft Windows Server™ 2003, Enterprise Edition, with Service Pack 1 (SP1)
Messaging application	Microsoft Exchange Server 2003, Enterprise Edition, with SP1

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

Figure 1. Server configuration for Microsoft Exchange test environment

the processing power of a quad-socket server may not be required for hosting Exchange alone.

Ideally, a server should support 4 GB of memory and have scalable I/O options while providing the processor scalability and network bandwidth required to host Microsoft Exchange Server 2003. Dual-socket Dell PowerEdge servers, available in both tower and rack optimized form factors, can help meet such requirements. These general-purpose servers can provide the performance characteristics for Exchange as well as a large amount of internal storage. For future I/O scalability, using external storage is recommended. However, this article focuses on internal server storage only, which is typically used in small Exchange environments.

Configuring the Microsoft Exchange test environment

The Dell test team performed sizing and performance tests at the Dell Enterprise Solutions Engineering Labs. The purpose of the tests was to determine the processor, memory, network, and disk utilization while hosting between 600 to 1,050 Exchange mailboxes. Figure 1 describes the hardware components of the test environment. For more information about the test bed setup, the test tools, and the setup procedures, see the supplemental online section of this article at www.dell.com/powersolutions.

Analyzing the test results

Of the various factors affecting the performance of an Exchange server, the test team studied the effect of increasing the number of mailboxes on processor, memory, and network utilization as well as on disk response times. Using these results, the team analyzed the potential bottlenecks that can be encountered in an Exchange deployment. The tests were conducted for 600 to 1,050 simulated Exchange users.

Processor performance

To observe the effect of increasing the number of mailboxes on processor utilization, the test team installed LoadSim on each of the clients—which were Dell PowerEdge 750 servers—to simulate an Exchange workload. The test was run several times for different user counts. Figures 2 and 3 show the results obtained for processor measurement counters.

For 600 to 1,050 users, the average processor utilization ranged from 5 to 10 percent, while the maximum processor utilization ranged from 17 to 27 percent. Processor privileged time and user time followed the same trend as processor utilization. As shown in Figure 2, the average processor utilization reached only about 10 percent for 1,050 Messaging Application Programming Interface (MAPI) Messaging Benchmark 3 (MMB3) users. Typically, if processor utilization is consistently greater than 75 percent it may be considered a bottleneck. In this case, an average processor utilization of 10 percent shows that the PowerEdge 2800 had ample computing power available to handle 1,050 or more mailboxes. Thus, processing power should not be expected to become a bottleneck for this size of Exchange deployment. The spare computing power

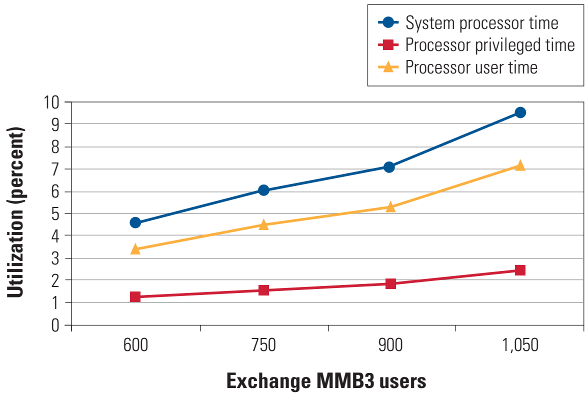

Figure 2. Average processor utilization for test environment

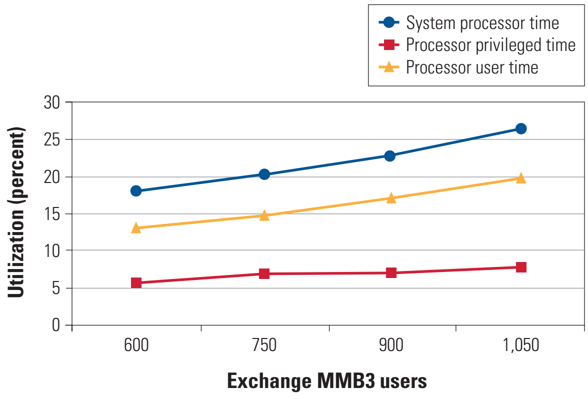

Figure 3. Maximum processor utilization for test environment

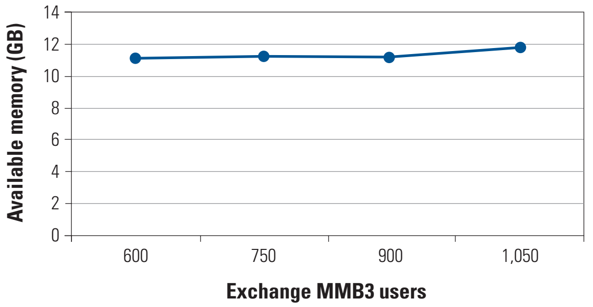

Figure 4. Memory utilization for test environment: Available memory

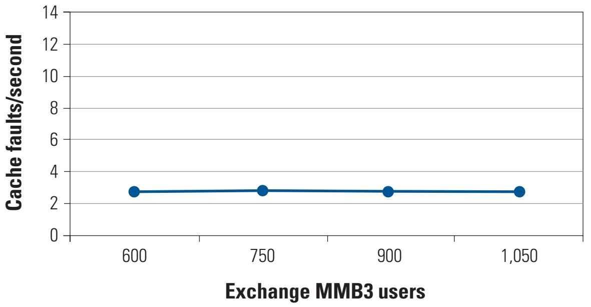

Figure 5. Memory utilization for test environment: Cache faults

available may be used by other applications such as antivirus, antispam, or backup software.

Memory performance

The tests used to study memory utilization were similar to those used for evaluating processor utilization. These tests measured available memory (see Figure 4) and cache faults (see Figure 5). The results show that available memory slightly declined and cache faults per second slightly increased as the number of mailboxes increased.

With the rapid increase in the speed of processors, memory, and system buses, the disk has become the slowest component of a server, making the disk seek process a very expensive operation in terms of performance. Exchange Server 2003 is an I/O- and memory-intensive application, so cache faults should be minimized as much as possible to limit data seek from disks and thus optimize performance. However, because Exchange is a 32-bit application, 4 GB of RAM is the maximum amount of memory that an Exchange server can efficiently use. Best practices recommend installing the maximum 4 GB of physical RAM.

Exchange memory optimization. By default, a 32-bit Windows OS allocates 2 GB of virtual address space for kernel mode (OS) and 2 GB for user mode (applications).[2] But 2 GB of virtual address space may not be enough for memory-intensive applications like Exchange. To solve this problem and make more memory space available for user mode, administrators can set the `/3GB` switch in the boot.ini file. The `/3GB` switch makes 3 GB (3,072 MB) of memory space available for user mode, and therefore more memory is available for Exchange services. Administrators should also set the `/USERVA=3030` parameter in the boot.ini file to allocate 3,030 MB of the 3 GB to Exchange and keep 42 MB available for dynamically allocating memory back to kernel mode if needed.

However, all client activity generated by a MAPI client causes updates to the Exchange Jet database, which produces random I/O activity. Installing the recommended 4 GB of RAM does not necessarily solve all memory issues nor guarantee that most I/Os will be cached. Using this switch may only help ensure that the Exchange components do not lack memory resources. For optimized performance, administrators should allocate an appropriate value of virtual memory to different Exchange components as well as to the OS. If the server is not a dedicated Exchange server and it hosts other applications such as antivirus, antispam, or backup software, additional RAM may be required.

As shown in Figures 4 and 5, increasing the number of users does not drastically increase the cache faults per second; neither does it significantly reduce the available memory. These results indicate that the Exchange components have sufficient memory and the deployment may be performing satisfactorily.

Network performance

Figure 6 shows network performance (kilobytes of data sent and received) for a varying number of users. These test results show that as the number of mailboxes increased, network traffic increased. Both the incoming and outgoing network traffic increased proportionally with the increasing number of mailboxes.

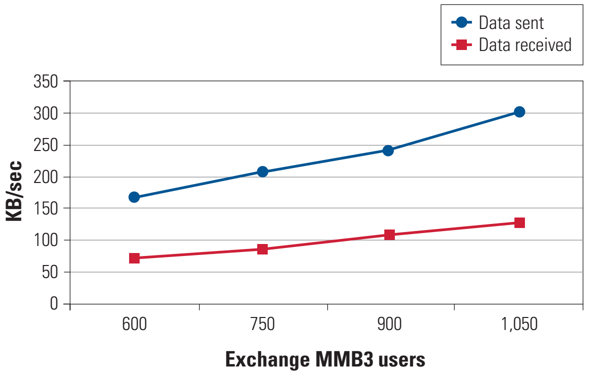

Figure 6. Network utilization for Microsoft Exchange test environment

[2] Microsoft Windows 2000 Advanced Server and 32-bit Windows Server 2003, Enterprise Edition, can support more than 4 GB of RAM by using Physical Address Extension (PAE). For more information about PAE and large memory support in Windows 2000 and Windows Server 2003, visit www.microsoft.com/whdc/system/platform/server/PAE/PAEdrv.mspx and support.microsoft.com/?kbid=283037.

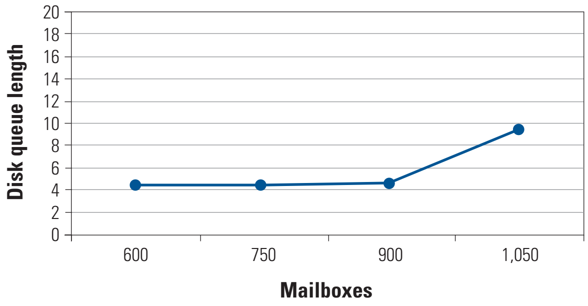

Figure 7. Storage performance for test environment: Disk queue length

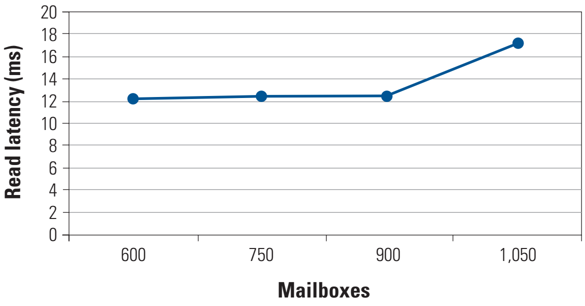

Figure 8. Storage performance for test environment: Read latency

From the magnitude of network traffic shown in Figure 6, the full-duplex 100 Mbps LAN connection appeared to perform sufficiently. However, these results do not include backup traffic or any other application traffic. If enterprise IT organizations plan to use the network for backups and restores or if other applications generate substantial network traffic, they should consider using a Gigabit Ethernet network.

Exchange uses the Microsoft Active Directory® directory service. Every Exchange mailbox must be an Active Directory user. Exchange servers and messaging clients access Active Directory in various situations, such as when logging on to the network and connecting to a mailbox or accessing server-based address lists. These activities create heavy network traffic between the servers, so there should be sufficient network bandwidth available between Exchange and the Active Directory server.

Storage performance

To observe the effect of increasing the number of mailboxes on storage performance, the test team used the Microsoft JetStress tool to simulate an Exchange I/O workload. The test was run several times for different mailbox counts with 50 MB mailbox size and 0.67 I/Os per second (IOPS) per user. Figures 7, 8, and 9 show the test results.

As shown in Figures 7, 8, and 9, disk queue length, read latency, and disk transfers per second rose steeply for more than 900 mailboxes. The increase in read latency was proportional to the increase in disk queue length. For 900 mailboxes and fewer, the read latency and disk queue length remained nearly constant.

Exchange is an I/O-intensive application and all client activity causes updates to the Exchange database, which produces I/O operations to disk. The disk subsystem should be able to meet these demands, and thus administrators should size the disks for performance and not just for capacity of the mailboxes.

For high performance, best practices recommend that the average read or write latency not exceed 20 milliseconds (ms) and the average disk queue length per spindle be less than 2. A disk queue length of less than 2 per spindle suggests that, while one I/O operation is being processed by the disk, another I/O is waiting in the queue. A value greater than 2 per spindle suggests that I/O requests are coming in at a higher rate than they can be processed by the disk. This may ultimately result in high disk latencies and therefore is not recommended. Best practices also recommend using fast drives rather than high-capacity drives. Fast drives can process more I/Os per second with low latency. If considering using SCSI drives, enterprise IT organizations should deploy 15,000 rpm Ulra320 drives. To further improve disk performance, IT administrators may also consider using the Microsoft DiskPar utility to verify that the disk tracks are sector aligned.

As shown in Figures 7, 8, and 9, the disk queue length, read latency, and disk transfers per second remained relatively consistent up to 900 mailboxes, even though the workload was increasing. However, these values increased sharply for 1,050 mailboxes. For example, the disk queue length for 1,050 mailboxes increased to 10. The database RAID group in this test environment comprised six spindles, so the disk queue length per spindle for 1,050 mailboxes was 1.67 (10 divided by 6), which is close to the maximum recommended limit of 2. Also, the read latency increased to 17 ms, which is approaching the limit of 20 ms.

These test results show that the disk performance parameters approached their maximum recommended limit for 1,050 mailboxes. To allow for spikes and periods of heavy load, best practices recommend no more than 900 mailboxes on the Exchange server.

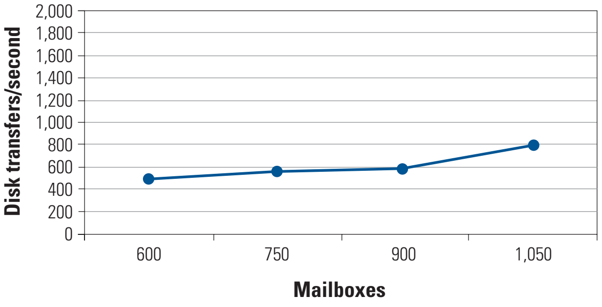

Figure 9. Storage performance for test environment: Disk transfers per second

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

To accommodate future growth, administrators should consider further reducing the number of hosted mailboxes. For example, to accommodate 25 percent future growth up to a maximum of 900 mailboxes, no more than 675 mailboxes should initially reside on the Exchange server.[3]

Available disk bandwidth can become a bottleneck as the number of mailboxes increases. To overcome this limitation and host large numbers of mailboxes, administrators may consider moving storage to external storage devices including direct attach storage (DAS) or storage area networks (SANs).

Considering other Microsoft Exchange configuration options

IT administrators deploying Microsoft Exchange in SMB environments may consider configurations such as high-availability (HA) clusters and front-end Microsoft Exchange servers.

High-availability clustering

The basic goal of HA clustering is to make sure that the physical server hosting an application is not a single point of failure by providing the ability for that application to be restarted on one of multiple servers in a cluster. If the server running the application fails, another designated server takes over the responsibility of running that application. Dell HA cluster implementations employ Microsoft Cluster Service (MSCS), and are designed and tested to help make sure that no single point of failure exists.[4] HA clustering using MSCS requires shared storage, because every node in the cluster needs access to the Exchange data. Therefore, to implement this type of configuration, administrators should consider moving storage to external devices such as DAS or SANs.

Front-end servers

Microsoft Exchange Server 2003 supports an Exchange architecture consisting of front-end and back-end servers. The front-end server accepts requests from clients and proxies them to the appropriate back-end server for processing. This architecture is usually recommended when the Exchange environment has multiple back-end servers. However, administrators may deploy a front-end server in an environment with a single back-end server because this configuration can provide the following benefits:

- **Single namespace:** Allows single namespace for all users; same URL for Microsoft Outlook® Web Access, Post Office Protocol 3 (POP3), and Internet Message Access Protocol 4 (IMAP4) clients; and consistent server name (does not change even if mailbox is moved or new servers are added)
- **Offload processing:** Uses front-end server to manage all encryption and decryption processing
- **Strong security:** Uses front-end server as a single point of access either on or behind a firewall; provides an additional layer of security for mailboxes; and does not require Remote Procedure Call (RPC) ports to be opened from perimeter network to internal network
- **Scalability:** Enables the number of front-end or back-end servers to be increased or decreased without disrupting users

Deploying Microsoft Exchange in an SMB environment

SMBs can range in size from 50 to 2,000 employees. A dual-socket Dell PowerEdge 2800 server can provide a suitable platform for SMBs to deploy messaging solutions. The results presented in this article show that, if used as a dedicated Microsoft Exchange server, the internal drives of a PowerEdge 2800 server can host up to 900 mailboxes.

Although the PowerEdge 2800 has enough processing power, memory, and network capacity to support much more than 900 mailboxes, the available disk bandwidth using the on-board hard drives can become a limiting scaling factor. To overcome this limitation and to host more than 900 mailboxes on a single server, enterprise IT administrators can move storage to external SCSI devices or host the Exchange mailboxes on a Dell/EMC Fibre Channel SAN. The best practices discussed in this article also can help administrators identify the bottlenecks in their deployments and take the appropriate actions to correct them.

Suman Kumar Singh is a systems engineer in the High-Availability Systems Group at Dell. He specializes in messaging systems architecture and sizing. His other interests include SANs, virtualization, and security. Suman has published several papers and presented at virtualization-related industry conferences.

Abhijit Chattopadhyay is a senior engineer in the High-Availability Solutions Engineering Group at the Dell Bangalore Development Center. Abhijit has a degree in Electronics and Communication Engineering from K.L.E. Society's College of Engineering and Technology in Karnataka, India. His current interests include high-availability clustering, SANs, and Microsoft Exchange.

Bharath Vasudevan currently manages the High-Availability Cluster Group at Dell. He has previously designed server hardware and served as a lead for multiple cluster releases. His current interests include application performance characterization and storage technologies. He has a master's degree in Electrical and Computer Engineering from Carnegie Mellon University.

[3] These storage recommendations are for Exchange users each generating 0.67 IOPS. Performance may vary depending on the nature of the Exchange users.

[4] For more information about deploying Microsoft Exchange Server 2003 on a Dell HA cluster, see "Microsoft Exchange Server 2003 Scale-Out Performance on a Dell PowerEdge High-Availability Cluster" by Arrian Mehis, Ananda Sankaran, and Scott Stanford in *Dell Power Solutions,* February 2005; www.dell.com/downloads/global/power/ps1q05-20040216-Stanford.pdf.

Backup Strategies for Microsoft Exchange Server 2003

Microsoft® Exchange Server 2003 installations require a solid backup and recovery architecture. This article presents common topologies and operational strategies for backing up Exchange in a variety of enterprise scenarios, including stand-alone servers, LANs, and storage area networks.

BY SUMAN KUMAR SINGH AND QUOCDAT NGUYEN

Related Categories:

Backup

Dell PowerEdge servers

Dell PowerVault storage

Dell/EMC storage

Disaster recovery

Microsoft Exchange

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

Enterprises typically translate e-mail downtime into lost revenue, lost productivity, or both. The size of an organization and the number of applications and service-level agreements that the data center must support are essential considerations when determining a suitable e-mail backup architecture. Enterprises must assess cost/benefits trade-offs that include the impact of the backup infrastructure on the availability and performance of business-critical applications—defining an acceptable duration for the backup and restore window, for example.

To protect e-mail data from potential disaster, the first line of defense is usually to back up critical information using tape or disk. To help administrators determine the backup method that is most appropriate for their specific enterprise requirements, this article describes three models for backing up the e-mail infrastructure: stand-alone server backup, LAN-based backup, and storage area network (SAN)–based backup.

Stand-alone server backup model

A stand-alone backup and restore scenario can be appropriate for small Microsoft Exchange environments that are hosted on a single stand-alone server. This approach typically locates the Exchange database on the server's internal storage or on direct attach SCSI or Fibre Channel storage. Even if the Exchange database is fairly large, the high storage capacities of advanced tape technologies may allow an organization to back up data onto a single tape. For example, the tape backup unit in the stand-alone server backup model can be a dedicated tape drive, such as the Dell™ PowerVault™ 110T Ultrium 3 Linear Tape-Open (LTO-3) tape drive, or an autoloader tape backup library, such as the Dell PowerVault 132T tape library (see Figure 1).

The advantages of this model are that it is simple, easy to deploy and implement, and cost-effective. However, there are a few limitations. For example, the management

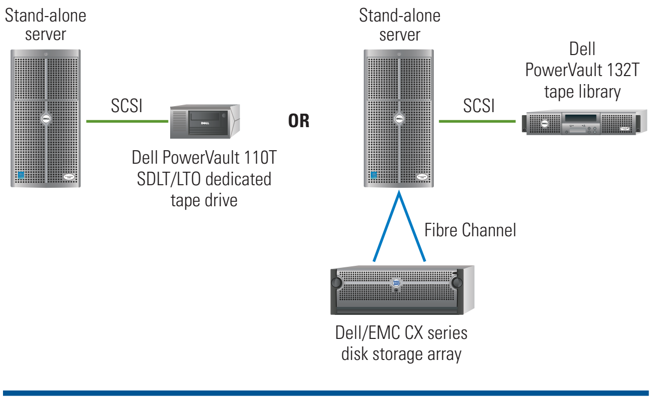
Stand-alone server
SCSI
Dell PowerVault 110T SDLT/LTO dedicated tape drive
OR
Stand-alone server
Dell PowerVault 132T tape library
SCSI
Fibre Channel
Dell/EMC CX series disk storage array

Figure 1. Stand-alone server backup model

of a stand-alone server backup model can be difficult if the data center is running multiple applications that must be backed up separately. In addition, this backup model offers limited scalability compared to LAN-based and SAN-based approaches. Because this configuration does not use a separate backup server, the application server's system resources must be shared and dedicated among different applications, including the backup task. As a result, backups may affect the performance of production applications.

LAN-based backup model

In the LAN-based scenario, the tape library attaches to a separate server known as the backup server or media server. The backup server connects to the application servers over the LAN, as shown in Figure 2. The LAN-based model may be suitable for environments that support multiple servers running multiple applications—for example, Oracle® database, Microsoft SQL Server™, and Microsoft Exchange applications. Configured either on stand-alone servers or on clustered servers, the applications can share the same backup tape library across the LAN. The master backup server initiates backup tasks and provides a centralized location housing the catalog database as well as the logical/physical management tree for the entire backup organization. The backup application agent must be installed on all application client nodes to facilitate application backups via the backup server. The backup server has a physical SCSI- or Fibre Channel–based interface to the tape library. The data from the Exchange server or other applications is sent over the LAN to the backup server and then to the tape library. The data flow is shown in Figure 2.

Compared to the stand-alone server backup model, the LAN-based model offers enhanced scalability to meet future needs for expanded storage and backup capacity.

LAN-based backup has several advantages over the stand-alone server backup model. In this model, different application servers can share a single tape library over the LAN. This centralized approach helps simplify backup administration. Compared to the stand-alone server backup model, the LAN-based model offers enhanced scalability to meet future needs for expanded storage and backup capacity. However, the LAN approach can also incur significant performance penalties because backup is performed over the network. To help avoid contention between the application traffic and backups, best practices recommend configuring a separate, isolated subnet dedicated to backup traffic. This limitation can be addressed in a SAN-based backup model.

SAN-based backup model

The SAN-based scenario is similar to the LAN-based approach—it also streamlines administration through centralized backups. The network topology for SAN-based backups is designed to improve application performance because it is routed over a high-speed Fibre Channel network. As a result, SAN-based backups enable the following benefits:

- Enhanced performance of transaction-intense applications using a high-bandwidth network interface
- High reliability and availability
- Ability of heterogeneous servers and operating systems to coexist and share the same tape library across the network

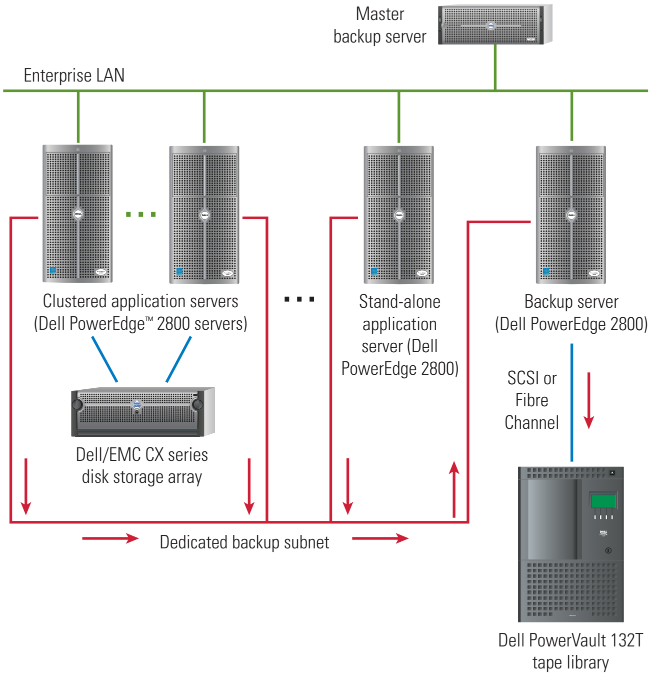
Master backup server
Enterprise LAN
Clustered application servers (Dell PowerEdge™ 2800 servers)
Stand-alone application server (Dell PowerEdge 2800)
Backup server (Dell PowerEdge 2800)
SCSI or Fibre Channel
Dell/EMC CX series disk storage array
Dedicated backup subnet
Dell PowerVault 132T tape library

Figure 2. LAN-based backup model

As shown in Figure 3, application servers, a master backup server, a storage system, and a tape library are connected across the Fibre Channel fabric. Data traffic can be routed from the application servers through a high-speed Fibre Channel switch and written directly to the tape library. As in the LAN-based model, the master backup server controls the backup tasks.

At press time, the front-end Fibre Channel interface for Dell-based SANs is designed to support data transfer rates of up to 4 Gbps (up to 800 MB/sec full duplex). Moreover, the latest multi-drive autoloader tape libraries or modular tape libraries—such as the Dell PowerVault 136T and PowerVault ML6000 series, respectively—enable administrators to back up several applications concurrently. Of the three backup topologies described in this article, the SAN-based model is designed to provide the highest I/O throughput for backup. Consequently, the SAN approach has the potential to drastically reduce backup and restore windows compared to LAN and stand-alone server models—thereby helping to improve performance for application service-level agreements.

The SAN approach has the potential to drastically reduce backup and restore windows compared to LAN and stand-alone server models—thereby helping to improve performance for application service-level agreements.

In addition to these benefits, the SAN-based model can take advantage of storage software and backup techniques such as EMC® SnapView™ snapshot software, EMC MirrorView™ SAN-based mirroring, and the Microsoft Windows®–based Volume Shadow Copy Service (VSS). The SAN-based approach also enables cluster-aware backups, in which a backup job can be failed over to another available node in the event of a failure during backup operations.

Backup strategies

For all three scenarios described in this article, administrators must ensure that backup software running on the backup infrastructure is compatible with Microsoft Exchange Server 2003. To take advantage of online Exchange backup, backup software must support the Exchange Server 2003 Backup and Restore application programming interface (API) or the Windows VSS writer.

It is equally important to back up all the data required to restore applications running on an organization's server to a previous known good state. Along with the applications, support software and management scripts must be backed up. For Exchange Server 2003, backing up the contents of mailboxes, public folders, and requisite configuration data for the Exchange environment is critical. In addition, best practices recommend that Exchange data be backed up separately—not together with Windows or with the full server backup operation.

Administrators should ensure that backups include the following:

- Microsoft Windows OS
- Backup and systems management software
- Management scripts
- Microsoft Active Directory® data
- System state, including the Microsoft Internet Information Services (IIS) metabase
- Cluster quorum (if Exchange uses clusters)
- Certification services (if applicable)
- Exchange databases and log files
- Exchange message-tracking logs

After hardware and software components are configured properly in the backup infrastructure and critical backup data is identified, administrators must implement a backup strategy. Exchange works with one or a combination of the following methods: full backup, differential backup, incremental backup, and mirror backup.

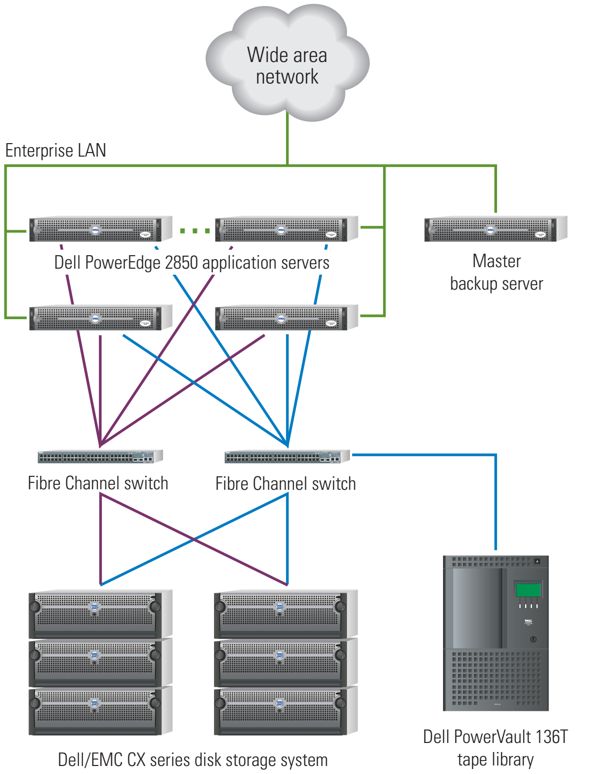
Wide area network
Enterprise LAN
Dell PowerEdge 2850 application servers
Master backup server
Fibre Channel switch
Fibre Channel switch
Dell/EMC CX series disk storage system
Dell PowerVault 136T tape library

Figure 3. SAN-based backup model

Full backup. A full backup is designed to store all data, including Exchange database files and transaction logs. This approach helps simplify the recovery process because it saves all the data files and transaction log files in a single backup session. However, a full backup operation consumes the most network bandwidth and requires the most storage space compared to differential, incremental, and mirror backups. For that reason, best practices recommend a full backup operation be performed at regular intervals, in rotation with other backup strategies.

Differential backup. A differential backup contains only the Exchange transaction log files that have changed since the last full backup; the database files are not copied. Because all the transaction logs since the last full backup are required for a restore operation, circular logging cannot be enabled during a differential backup. Recovery requires both the last full backup and the last differential backup. Best practices recommend that a full backup be performed at regular intervals and supplemented with daily differential backups.

Incremental backup. An incremental backup contains the Exchange transaction log files that have changed since the last full, differential, or incremental backup. Of these three types, incremental is the fastest backup method and may be suitable for large Exchange databases with a high volume of daily activity. The drawback to the incremental approach is that recovery requires the last full backup and all subsequent incremental backups. Best practices recommend that a full backup be performed at regular intervals and supplemented with daily incremental backups.

Key considerations when determining the e-mail infrastructure include the time necessary to perform the backup, the number of tape backup sets required for the restore, the time necessary to complete the restore, and the system resources available to perform the restore.

Mirror backup. A mirror backup is similar to a full backup except that no file marking is performed. Mirror backup is not ordinarily used for recovery purposes. This method can be used to make a full copy of the Exchange database without disrupting any incremental or differential backup procedures.

Microsoft Exchange Server 2003 supports online backup, which can be performed using a program such as Microsoft Windows NT® Backup (Ntbackup) or a third-party backup utility that supports the Exchange backup API. During the online backup process, Exchange services typically continue to run normally and users experience no downtime. Online backups can be performed for full, differential, incremental, and mirror backup strategies.

Practical considerations

For enterprises of all sizes, e-mail has become a mission-critical application. When determining a suitable backup architecture for the e-mail infrastructure, administrators must factor in the number and variety of applications and service-level agreements that the data center must support. Of particular concern is the impact of the e-mail infrastructure on the availability and performance of business-critical applications. Key considerations when determining the e-mail infrastructure include the time necessary to perform the backup, the number of tape backup sets required for the restore, the time necessary to complete the restore, and the system resources available to perform the restore. Last but not least, administrators must ensure that backup media is stored in a secure location.

Last but not least, administrators must ensure that backup media is stored in a secure location.

The three strategies described in this article for backing up Microsoft Exchange Server 2003—stand-alone server backup, LAN-based backup, and SAN-based backup—offer general guidelines for Exchange backup operations. Actual implementations may differ based on the specific requirements of individual Exchange organizations.

Suman Kumar Singh is a systems engineer in the High-Availability Systems Group at Dell. He specializes in messaging systems architecture and sizing. His other interests include SANs, virtualization, and security. Suman has published several papers and presented at enterprise computing–related industry conferences.

QuocDat Nguyen is a systems engineer in the High-Availability Cluster Development Group at Dell. His responsibilities include developing SAN-based, high-availability clustering products that comprise Dell servers and Dell/EMC Fibre Channel storage systems. QuocDat has a B.S. in Electrical Engineering from the University of Houston.

FOR MORE INFORMATION

Microsoft Exchange Server 2003 upgrades on Dell platforms:
www.dell.com/exchange

Dell storage:
www.dell.com/storage

SQL Server 2005 on a Dell Scalable Enterprise Foundation

Dell's vision for the scalable enterprise is based on the standardization of core elements of the data center to provide superior value, and encompasses the core tenets of simplified operations, improved utilization, and cost-effective scaling. Using Microsoft® SQL Server™ 2005 as a representative workload on Dell™ PowerEdge™ servers, the Dell Scalable Enterprise Technology Center Labs team describes high-level considerations for building a scalable enterprise foundation that can be customized to support aggressive growth.

BY TIM ABELS AND TODD MUIRHEAD

Related Categories:

Dell PowerEdge servers

Dell Scalable Enterprise Technology Center

Microsoft Operations Manager (MOM)

Microsoft SQL Server 2005

Microsoft Systems Management Server (SMS)

Microsoft Windows Server 2003

Scalable enterprise

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

The Dell Scalable Enterprise Technology Center Labs are dedicated to developing and testing representative scalable enterprise architectures and workloads. This first article in an ongoing series examines the advantages of using industry-standard platforms as building blocks to create a foundation for managing an example Microsoft SQL Server 2005 application on Dell PowerEdge servers. Upcoming articles will detail best practices and key decision points regarding specific areas such as deployment, management, virtualization, and scalability of the example Microsoft SQL Server 2005 application, based on actual configurations demonstrated in the Dell Scalable Enterprise Technology Center Labs test bed.

This article illustrates how organizations can lay the groundwork for growing enterprise applications from the smallest to the largest possible scale without changing the underlying IT infrastructure. In fact, the Dell Scalable Enterprise Technology Center Labs team reuses many infrastructure components from project to project to show how different applications and OS environments can run on the same basic foundation of industry-standard data center platforms.

Standards are the most important part of a scalable enterprise architecture. Dell's scalable enterprise architecture relies on standards-based software and hardware, such as Microsoft SQL Server 2005 on Dell PowerEdge 2850 servers, to enable organizations to scale quickly and flexibly in response to changing business requirements.[1] Dell PowerEdge servers are built on industry-standard Intel® Xeon® processors that provide

[1] The Dell Scalable Enterprise Reference Architecture addresses solutions that are available today and looks ahead to future possibilities, considering how the data center may function as technology continues to mature and additional standards are defined. For a detailed definition of the Dell Scalable Enterprise Reference Architecture, see "Dell Scalable Enterprise Architecture" by Jimmy Pike and Tim Abels, Dell Inc., August 2005; www.dell.com/downloads/global/vectors/2005_scalable_enterprise.pdf.

a common, cost-effective platform for high-performance, high-availability SQL Server 2005 applications.

This automated and standardized data center vision of Dell's scalable enterprise strategy is not fully realized today. This is because some high-level data center functions have yet to be standardized and management tools will need to be created around those standards. However, organizations can advance toward this goal by laying the groundwork with industry-standard data center components that can be managed from a central console—achieving tangible benefits today from simplified operations, improved resource utilization, and cost-effective scaling. By implementing progressive levels of data center automation in pragmatic, planned phases, organizations can equip their IT infrastructure with components that are consistent with their current business goals and practices, and be prepared to take advantage of the larger scalable enterprise vision as high-level standardization in the data center matures. (For details, see the supplemental online section of this article, "Progressive levels of automation for implementing the scalable enterprise," at www.dell.com/powersolutions.)

Building a scalable enterprise foundation

The Dell Scalable Enterprise Technology Center Labs team built a scalable enterprise architecture focusing on the example SQL Server 2005 database platform (Figure 1). This example architecture was used to demonstrate how a highly scalable enterprise architecture can be built using industry-standard hardware and software to help achieve the three primary benefits of Dell's scalable enterprise strategy: simplified operations, improved resource utilization, and cost-effective scaling.

One centralized enterprise management console simplifies monitoring and operations to a great extent. Other management consoles may be allowed to perform one-time tasks, but whenever possible, all monitoring and alerting should be done through a single management console. The integration of Dell PowerEdge servers into Microsoft Operations Manager (MOM) 2005 and Microsoft Systems Management Server (SMS) 2003 provides consolidated tools for managing, monitoring, and deploying both software and hardware. As a result, administrators can manage hardware monitoring, hardware inventory, and patching with the same tools they use to monitor and deploy SQL Server 2005 software.

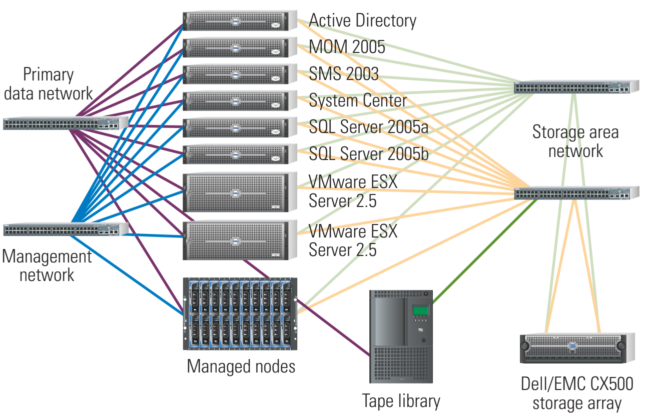

Figure 1. Example SQL Server 2005 configuration implemented in the Dell Scalable Enterprise Technology Center Labs

This example architecture was used to demonstrate how a highly scalable enterprise architecture can be built using industry-standard hardware and software to help achieve the three primary benefits of Dell's scalable enterprise strategy: simplified operations, improved resource utilization, and cost-effective scaling.

Virtualization is a key data center platform that contributes significantly to improved resource utilization. Not only does virtualization software enable servers to be provisioned, deployed, and maintained without incurring downtime, but it also provides a fast, flexible way to consolidate underutilized servers. The ability to reprovision servers immediately to meet daily business fluctuations also helps increase the utilization of existing resources. Dell PowerEdge servers running Microsoft Virtual Server 2005, VMware® ESX Server™, or VMware GSX Server™ software enable a virtualization platform that allows a single physical server to run many virtual machines (VMs) at the same time—with each VM running a separate OS and applications environment—while still being managed through a central console.

Industry-standard building blocks of hardware and software support the growth of scalable enterprise applications because high-volume, industry-standard data center components are typically less expensive to acquire and maintain than proprietary and custom solutions—and they typically attract a larger community of support than proprietary and custom solutions. Advances in performance enabled by 64-bit support and dual-core capability of Intel Xeon processors, combined with scalable standards-based software, help provide powerful, cost-effective enterprise solutions today and anticipate continued improvement in the design and performance of integrated, cost-effective, industry-standard platforms.

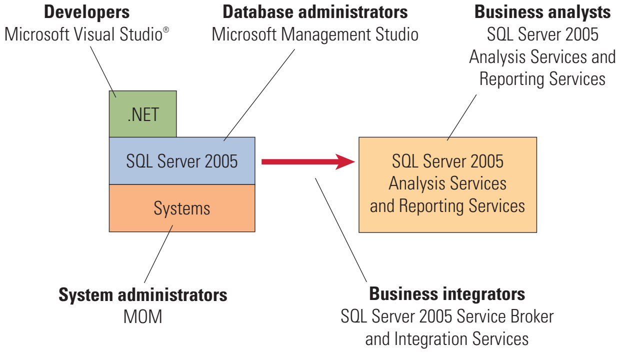
Developers
Microsoft Visual Studio®
Database administrators
Microsoft Management Studio
Business analysts
SQL Server 2005 Analysis Services and Reporting Services
.NET
SQL Server 2005
Systems
SQL Server 2005 Analysis Services and Reporting Services
System administrators
MOM
Business integrators
SQL Server 2005 Service Broker and Integration Services

Figure 2. Tools available to different types of enterprise users in the integrated Dell–SQL Server 2005 implementation

Integrating resource management

The Microsoft SQL Server 2005 enterprise database platform introduces significant changes in features, tools, and interfaces that enable many advances for enterprises building large-scale, industry-standard database solutions. The integrated Dell and Microsoft SQL Server 2005 stack configured in the Dell Scalable Enterprise Technology Center Labs shows how a single solution can be designed to address the needs of developers, system administrators, database administrators, business analysts, and business integrators. Figure 2 shows the tools that are available to various SQL Server 2005 users in the integrated Microsoft and Dell implementation.

MOM 2005 and management packs

Tight integration with MOM 2005 enables system administrators to manage SQL Server 2005 database operations in a simplified way. By providing a central operations console that can monitor and manage Microsoft SQL Server 2005, Active Directory® services, Windows Server™ operating systems as well as Dell PowerEdge server hardware, MOM significantly simplifies data center operations. MOM is extended to manage these components with management packs (MPs). Each MP contains the definitions that MOM needs to be able to monitor and manage the software or hardware. Additionally, MOM allows administrators to extend these MPs with custom definitions and thresholds to tailor MOM to their own environments.

Microsoft's SQL Server MP for MOM provides predefined enterprise-level availability and configuration monitoring; performance data collection; and default thresholds designed to proactively increase the security, availability, and performance of a SQL Server infrastructure. MOM 2005 can monitor a variety of SQL Server 2005 configurations, including failover clusters, named instances, log shipping, replication clusters, data mirrors, partitioned clusters, 32-bit and 64-bit processors, multi-core processors, and Data Protection Manager for hot backup. For free downloads of the SQL Server MP, the Dell MP, and related MPs for server operating systems, clusters, Virtual Server, Active Directory, and Data Protection Manager, visit www.microsoft.com/mom/managementpacks.

The Dell MP, when loaded into the MOM 2005 console, imports predefined computer groups, processing rules, computer attributes, scripts, tasks, the Dell Knowledge Base, and public views for managing Dell-specific applications and hardware. The Dell MP provides the connections for MOM to be able to interact with Dell OpenManage™ Server Administrator (OMSA) running on Dell servers. This includes links in a Dell MP–generated alert to launch the OMSA console for systems experiencing a problem condition. OMSA is a Web-based console that is included with Dell PowerEdge servers. It runs on each server to provide detailed monitoring, alerting, and diagnostic tools for the Dell server hardware. The Dell MP is designed to create tight integration of the Dell OpenManage hardware-level monitoring with MOM's enterprise-level console. The Dell MP is available as a free download from support.dell.com.

Figure 3 shows how a MOM server can manage the SQL Server 2005 platform. *Note:* The support servers may assist MOM with one-time support services, but MOM provides continual event monitoring, management, and unified reporting.

SMS 2003 and Dell PowerEdge servers

SMS 2003 Service Pack 1 supports the Inventory Tool for Dell Update, which enables the SMS agent to inventory Dell PowerEdge servers for the relevant Dell hardware patches. Once the inventory is complete, SMS then applies the needed patches in the correct order. SMS can also pull updates down automatically from the Internet to procure and apply current patches for Dell servers.

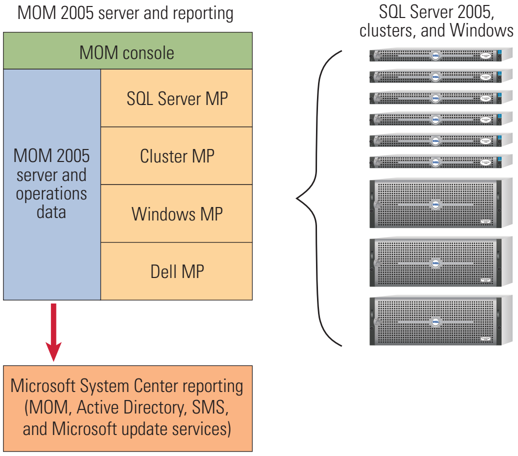
MOM 2005 server and reporting
SQL Server 2005, clusters, and Windows
MOM console
MOM 2005 server and operations data
SQL Server MP
Cluster MP
Windows MP
Dell MP
Microsoft System Center reporting (MOM, Active Directory, SMS, and Microsoft update services)

Figure 3. MOM 2005 server managing SQL Server 2005 systems

To accomplish this integration, Microsoft used the Dell Partner Development Kit (PDK) to enable SMS to interact with Dell servers using the appropriate interfaces. SMS can be configured to update Dell PDK components automatically as Dell updates them, or administrators can decide to upgrade manually when needed.

When Dell PowerEdge servers running SQL Server 2005 need to be updated and patched, using Microsoft SMS with the integrated Dell tools helps further simplify the management of these systems. The ability to automatically and systematically update the server software and hardware is essential to a highly manageable infrastructure.

Virtualization and consolidation

Many data center servers and storage systems are underutilized because they are built out to support peak capacity of several departments and multiple locations. Virtualization platforms for the Intel x86 architecture play a principal role in enabling server consolidation and can lead to improved server utilization. SQL Server 2005 on Dell PowerEdge servers provides an excellent example of how servers can be consolidated using standard hardware, proper sizing, and virtualization software.

Light SQL Server 2005 applications, including most developer or test databases, can be consolidated using server virtualization. This approach allows many VMs to run on a single physical server, isolating each VM with its own OS and application(s). By allowing many virtual servers to be consolidated onto a single physical server, virtualization enables servers to operate close to the maximum specified processing capacity—and corrective reprovisioning can be accomplished quickly when necessary. In this way, many underutilized servers that are each running a single application can be replaced with a single, highly utilized server running many VMs. Through server consolidation, virtualization can lead to cost savings in reduced hardware acquisition and maintenance costs, reduced data center space requirements, and simplified management through a common virtualization layer. MOM 2005 MPs are available for Microsoft Virtual Server, VMware ESX Server, and VMware GSX Server products, enabling management of these virtualization environments as part of the simplified operations design.

Dynamic reprovisioning and workload balancing

Consolidated hardware and virtualization software can be used as flexible building blocks to enable dynamic reprovisioning and workload balancing. Standardizing and consolidating on Dell PowerEdge servers helps provide a single hardware platform with rich management tools that can be supported with a single system image. A single enterprise manager, managing a single image, is key to reprovisioning across a range of servers.

When dynamically reprovisioning systems, enterprises must be careful to provide additional capacity beyond the anticipated workload. Reprovisioning and virtualization techniques enable enterprises to rebalance loads across VMs running on consolidated hardware platforms so they use resources efficiently, enabling a dynamic response to changing business conditions. Dell PowerEdge servers provide a hardware platform that supports a common OS image, which enables easy reprovisioning of servers running the common system image to support various tasks. For example, Dell PowerEdge 1850, PowerEdge 2800, and PowerEdge 2850 servers share a common system image, whether they are based on single-core or dual-core processors.

A major restriction when scaling 32-bit SQL Server 2000 databases was the 4 GB limit for directly addressable memory. Address Windowing Extensions (AWE) allowed 32-bit Windows to address up to 32 GB of memory for SQL Server 2000 Enterprise Edition, but it was not practical for many memory-intensive, high-performance enterprise applications due to AWE address translation overhead, and the 4 GB limit could be extended only for a single SQL service and cache. In comparison, 64-bit Windows Server 2003 with 64-bit SQL Server 2005 running on 64-bit Dell PowerEdge servers is designed to directly address all SQL services up to the maximum amount of memory supported by the hardware, with no overhead penalties.

> Consolidated hardware and virtualization software can be used as flexible building blocks to enable dynamic reprovisioning and workload balancing. Standardizing and consolidating on Dell PowerEdge servers helps provide a single hardware platform with rich management tools that can be supported with a single system image.

Leading price/performance examples

Dell has worked with Microsoft to provide a cost-effective, scalable database platform integrating Dell PowerEdge servers and SQL Server 2005 applications. Dell PowerEdge servers are available with high-performance components, including dual-core Intel Xeon processors, PCI Express expansion slots, double data rate 2 (DDR2) memory, and on-board management controllers.

For example, a Dell PowerEdge 2800 with a dual-core Intel Xeon processor running in Extended Memory 64 Technology (EM64T) 64-bit mode broke the $1 cost per transaction barrier on the TPC-C

System	tpmC	Price/tpmC (US$)	Processor
Dell PowerEdge 2800	38,622	0.99	1 dual-core Intel Xeon processor at 2.8 GHz, 2×2 MB cache
Dell PowerEdge 2800	28,244	1.29	1 single-core Intel Xeon processor at 2.8 GHz, 2 MB cache
Dell PowerEdge 2800	28,122	1.40	1 single-core Intel Xeon processor at 3.4 GHz, 2 MB cache
Dell PowerEdge 2850	26,410	1.53	1 single-core Intel Xeon processor at 3.4 GHz, 1 MB cache

Source: Top Ten TPC-C by Price/Performance Version 5 Results as of April 13, 2006; www.tpc.org/tpcc/results/tpcc_price_perf_results.asp. For detailed configurations, visit the preceding TPC-C results page and click on each system, and then click "Executive Summary."

Figure 4. Four top TPC-C results by price/performance

benchmark[2] using a configuration that included SQL Server 2005. The second, third, and fourth best price/performance TPC-C results were also Dell PowerEdge servers, with configurations that included SQL Server 2000 (see Figure 4). Based on these TPC-C results, the upgrade to dual-core and SQL Server 2005 improved the price/performance by 30 to 50 percent. These results show the potential for improved performance across a wide range of SQL applications and servers in the data center.

Licensing considerations

Dell offers the SQL Server 2005 license as an option that can be purchased with a PowerEdge server. Microsoft Windows and SQL Server 2005 licensing enable a dramatic improvement in price/performance because they are both based on sockets instead of cores. For example, SQL Server 2005 on PowerEdge 6800 or PowerEdge 6850 servers with dual-core processors and Intel Hyper-Threading Technology provides eight cores and 16 threads for the same licensing fee as SQL Server 2000 on single-core processors, which provides only four cores and eight threads.[3]

Similarly, Windows Server 2003 Release 2 (R2), Enterprise Edition, improves licensing by enabling up to four instances of SQL Server in VMs on a single physical server. The VMs can be hosted with Microsoft Virtual Server, VMware ESX Server, or VMware GSX Server and can include instances of Windows Server 2003 R2, Standard Edition or Enterprise Edition. Additionally, the license limit of four VMs applies to running instances only, allowing for many inactive system images on a storage area network or file servers.[4]

Advancing scalable enterprise goals

Standardization and integration are at the heart of Dell's scalable enterprise architecture. By deploying industry-standard data center components such as Dell PowerEdge servers and SQL Server 2005 database applications together with virtualization technology, any size organization—from a small firm starting out with a stand-alone server environment to a huge global enterprise with an established infrastructure in place—can build a highly manageable, scalable foundation. Microsoft Operations Manager 2005 and Microsoft Systems Management Server 2003 are instrumental in simplifying operations, improving resource utilization, and scaling cost-effectively because they enable administrators to monitor and manage the integrated environment—including SQL Server 2005, Windows Server 2003, and Dell PowerEdge servers—from a central console.

By demonstrating how to provision, deploy, and manage representative scalable enterprise architectures, the Dell Scalable Enterprise Technology Center shares best practices and identifies key decision points that can help organizations benefit today and lay the groundwork to take full advantage of the scalable enterprise approach as high-level standards for data center automation mature.

Tim Abels is a senior software architect on the Dell Scalable Enterprise Technology Center Labs team. Tim has an M.S. in Computer Science from Purdue University.

Todd Muirhead is a senior engineering consultant on the Dell Scalable Enterprise Technology Center Labs team. Todd has a B.A. in Computer Science from the University of North Texas.

FOR MORE INFORMATION

Dell scalable enterprise:
www.dell.com/enterprise

Microsoft SQL Server on Dell and SQL Server services:
www.dell.com/sql

Microsoft Windows and server management:
www.dell.com/openmanage

MOM 2005 and management packs:
www.microsoft.com/mom

SMS 2003 and Inventory Tool for Dell Update:
www.microsoft.com/smserver

Virtualization resources:
www.vmware.com/dell

[2] For more information about TPC-C, tpmC, and the Transaction Processing Performance Council, visit www.tpc.org.

[3] For details about Microsoft's licensing policy for multi-core processors, visit www.microsoft.com/licensing/highlights/multicore.mspx.

[4] For details about Microsoft licensing for Windows Server 2003 R2 in a virtualized environment, visit www.microsoft.com/windowsserver2003/howtobuy/licensingr2/overview.mspx.

Best Practices: **Enterprise Test Management**

Testing products before deployment is a critical step for enterprises. Those administering the tests must determine when a product is ready to be deployed to production. But before this decision is made, the entire testing process should be carefully planned, managed, and reviewed. With effective test management in place, enterprises can help ensure that products receive comprehensive, well-documented testing and are truly ready for production.

BY CYNTHIA LOVIN AND TONY YAPTANGCO

Related Categories:

Application development

Enterprise management

Enterprise testing

Performance

Planning

Project management

Scalable enterprise

System deployment

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

Successful test management requires development of a comprehensive testing strategy. Decisions must be reached about who will manage the test project, what measurement tools will be needed, and where and by whom the testing will be performed. Enterprises can apply a few general rules when developing and managing a testing strategy:

- Time durations for each test phase vary according to the complexity and stability of the product during the course of the test execution.
- Test cases or test scenarios are developed to focus on the testing objective for each test phase. Depending on the stability that the product demonstrates during the test phases, these tests may be run multiple times.
- Changes in the product design or additions to the scope of the product late in the cycle can affect an enterprise's test strategy as well as its resources and budget.

This article is the second in a series of articles examining best practices in enterprise testing. The previous article, "Best Practices: Enterprise Testing Fundamentals," which appeared in the February 2006 issue of *Dell Power Solutions,* focused on the fundamentals of unit, product, and system testing in a phased approach.[1] This article focuses on test management.

Developing a test strategy

A test strategy provides a framework for the testing effort. A comprehensive strategy goes beyond the technical requirements of the product or system under test. Figure 1 discusses some factors enterprises should consider when building a test strategy.

Enterprises may also want to consider industry best practices to help them develop their test strategy and processes:

- Industry experts recommend that the frequency of defects discovered during each phase of the test cycle fall within the following ranges: 60 to 65 percent during unit testing; 30 to 35 percent during product testing; and 5 to 10 percent during system testing.

[1] For more information, see "Best Practices: Enterprise Testing Fundamentals" by Cynthia Lovin and Tony Yaptangco in *Dell Power Solutions,* February 2006; www.dell.com/downloads/global/power/ps1q06-20050111-Lovin.pdf.

Factor to consider	Impact on testing strategy
Is the product under test a new architecture, first release, or sustaining effort?	Determines where to focus test coverage and technical strategy—for example, new feature validation, integration testing, or regression testing
Was the product under test developed in house, provided by a supplier, or jointly developed?	Determines the level of knowledge about the system's internal structure or coding and whether low-level or "white box" testing is possible; drives decisions about who performs testing as well as when and where to test
Will future, similar test efforts be needed?	Drives decisions about staffing and test automation
Will testing be performed by staff, a contract test house, or a supplier?	Affects the test costs, training, and retention of technical knowledge; drives the audit strategy (degree of revalidation to perform on testing conducted outside the enterprise)
Will testing be performed in a single geographic region or in multiple regions?	Determines the need to evaluate certain factors: ownership of and accountability for overall test effort; a clear delineation of responsibilities and tasks; the impact of working in different time zones; communication strategies, methods, and frequencies; and cultural factors such as trust, communication protocols, and holiday scheduling
How many configurations and prototype units are budgeted and available for testing?	Affects appropriate staffing levels; determines whether configuration variations can be tested in parallel or serially; helps establish a sound technical test strategy
Which configurations will be used most in deployment or pose the greatest risk to successful deployment?	Determines the test depth and frequency of test repetition; affects the test automation strategy

Figure 1. Factors that affect a test strategy

- Identifying and fixing defects during unit testing is significantly less expensive than doing so during the other test phases. System testing and deployment are the most expensive phases during which to identify defects.

Managing resources

Before testing begins, enterprise test teams should ensure that they and their testing partners (suppliers, contract houses, and so forth) have adequate capacity and resources for the job. Test personnel are needed in a variety of roles:

- A test project engineer or manager leads the test planning and execution.
- Test engineers develop and validate the necessary test cases (including automation).
- Test technicians set up the systems under test, run the test cases, and report any defects.

All necessary prototypes and test tools must be procured, shipped to the testing location(s), configured, and made ready to run the required tests. Test teams must make sure that they have a plan to upgrade hardware as needed during the project. Defect tracking and measurement systems must be ready. All members of the test team must understand the reporting requirements.

Example test phase entrance criteria	Example test phase exit criteria
• All test cases are available. • All test resources are available (both personnel and hardware). • All technical specifications and user documentation are available for this phase of the test effort. • Correct revision levels of hardware and software are installed. • A defect tracking system is in place.	• All test cases were attempted. • Test pass rates were high enough to warrant transition to the next phase. • All defect fixes were verified and regression on adjacent areas was successfully performed. • No unresolved high-severity defects exist that would adversely affect transition to the next phase.

Figure 2. Example test phase entrance and exit criteria

Managing the test project

Test project management involves planning the test project, establishing criteria that must be met before starting and concluding each test phase, supervising the test process, documenting the test procedures and results, and analyzing the test results.

Project planning

Managing a project begins well in advance of the start date. Stakeholders should have enough time to review the test strategy and test cases, and the test team should begin physical preparation of the test tools, configurations, and infrastructure.

Entrance and exit criteria

Test entrance and exit criteria should be defined for each phase of the test cycle. The test team must document criteria for entrance to and exit from the unit, product, and system test phases. Entrance criteria focus on component availability and other readiness measurements to start the test phase, whereas exit criteria are used to assess the maturity or progress of the product and readiness to move out of that phase. Figure 2 shows examples of entrance and exit criteria.

Test execution

Managing the execution phase consists primarily of keeping the test effort on track. This means that the test team must make sure that the tests are validating the planned functional coverage or use, defect fixes are being verified, and regression testing is being conducted within the allocated schedule and resources. Test teams often must achieve these goals even when the product design, scope, or implementation details change during the test cycle. Successfully managing test execution requires flexibility, sustained attention to test progress metrics (see the "Measuring success" section in this article), and an ability to make necessary last-minute adjustments.

Test documentation

Test documentation at a minimum should include a test plan, test cases, and a test summary report. Suggested contents for each type of document are as follows:

- **Test plan:** Includes what is being tested; what is excluded from testing; the features to be verified; the configurations to be used; and an estimate of the resources and time required for testing.
- **Test cases:** Include the goal of each test; a detailed description of the test procedures; the required environment for test execution; whether testing will be automated or manual; and the criteria by which to judge whether testing passed or failed.
- **Test summary report:** Includes the dates and revisions for test execution; whether entrance and exit criteria were met; the test case pass/fail rates; how many defects were detected, during which test phase they were detected, and their severity level; whether any defects remained in the product when it was shipped or deployed; and a description of those defects.

Post-project assessment

Soon after testing is completed, the test team should review what went well and what went poorly during the project. In particular, they should capture the historical metrics for the project, as described in the "Measuring success" section in this article and in Figure 3. The results of this assessment should be used to improve planning and execution of future testing.

Post-project assessment may include how long testing took compared to the test plan, how many resources were employed compared to the test plan, and a comparison of the test effort with previous similar efforts.

Measuring success

In general, test measurements are sound only insofar as they are balanced. A balanced set of test measurements includes assessments of productivity, financials, quality, and timeliness.

As part of the test management strategy, test teams should establish measurement systems with three distinct objectives: operational measurements to track real-time progress through test execution (assessing the quality and release readiness of the product); historical measurements to assess the project (comparing the actual results of the test project with the test plan); and business measurements to drive organizational improvement over time (comparing performance to that of similar past projects or to industry benchmarks). Figure 3 provides example metrics that can be used.

Learning from test projects

At times, the complexity of managing a test project can rival the complexity of the systems or products being tested. Enterprises must develop a strategy, a test plan, and test cases in advance to meet the unique requirements of their environment. Furthermore, test teams must keep their projects on track through technical difficulties, product changes, and resource constraints. By taking balanced measurements of their progress, test processes, and product quality, test teams can help their enterprises create successful products—and the lessons learned from each test effort can help improve planning and execution for future projects.

Type of metric	Example metrics
Operational	• Execution progress for test cases • Percentage of completed tests • Number of tests that passed, failed, or were unable to progress because of blocking defects • Number of outstanding defects (measured by severity and by component) • Rate of defect discovery (measured by severity and by component) • Rate of defects being fixed compared to the discovery rate • Number of defect fixes requiring regression testing
Historical (post-project assessment)	• Number of defects discovered per test hour • Number of defects that escaped discovery in a previous test phase • Number and percentage of valid defects discovered • Number and type of defects reported by users in the first 90 days after the product was introduced to the user community • Comparison of the project's planned execution (time, headcount, and so forth) with its actual execution
Business	• Duration of effort compared to a previous product of similar size or complexity • Resource consumption compared to a previous product of similar size or complexity • Number of defects found, phase during which the defects were found, severity of the defects, and comparison of these metrics with previous test efforts • Number of defects that escaped discovery in previous test phases and comparison of this metric with previous test efforts

Figure 3. Types and examples of testing metrics

Cynthia Lovin is a senior consultant test engineer in Dell's Product Group Global Test Department, with 16 years of test- and quality-engineering experience in companies ranging from startups to Fortune 500 enterprises. She has Six-Sigma Green Belt certification and a B.A. in Business Administration from The University of Texas at Austin.

Tony Yaptangco is the director of the system test group within the Dell Product Group Global Test Department. He has 25 years of experience in various software and test environments, including 16 years in engineering management positions. Tony has a B.S. in Computer Science from San Diego State University and an M.S. in Engineering Management from National Technological University.

Automated Deployment

of Novell SUSE Linux Enterprise Server 9 on Dell PowerEdge Servers

Administrators of large enterprises can save time when they install and configure systems that use the Novell® SUSE® Linux® Enterprise Server 9 (SLES 9) OS by implementing automated installation. This article introduces the concepts and step-by-step methods for automating the installation of SLES 9.

BY ANOOP K. AND JOHN HULL

Related Categories:

Linux

Linux engineering

Novell SUSE

OS deployment

System deployment

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

System deployment in an enterprise can be very time-consuming, especially if it involves a large number of servers. The traditional method of booting to the OS installation CD, manually answering configuration questions, and then changing CDs when prompted can be inefficient for data centers in which multiple servers must be deployed quickly and easily.

To deploy the Novell SUSE Linux Enterprise Server 9 (SLES 9) OS on a large number of servers in an enterprise, best practices recommend using an automated installation method that provides the following capabilities:

- Remote boot of the OS installation kernel
- Remote access to OS installation media
- Unattended installation of the OS

By harnessing the power of the SUSE Linux installer, administrators can easily install SLES 9 throughout the enterprise.

SUSE Linux Enterprise Server 9 automated installation package

The SLES 9 installer package comprises two programs: linuxrc and YaST (which stands for "Yet another Setup Tool"). The sections that follow describe these two programs and explain how to use them for automated installation of SLES 9 on Dell™ PowerEdge™ servers.

The linuxrc program

After the installation media boots, the SUSE Linux installation kernel and ramdisk are loaded, after which a small program called linuxrc starts. The linuxrc program supplies the necessary functions to prepare the installation environment, such as analyzing the system and loading kernel modules. It will prompt the administrator for language, keyboard layout, and installation media location—all of which are needed for the administrator to interact with YaST, the installer used in SUSE Linux.

The YaST program

After linuxrc has completed its tasks, it loads YaST and hands off control to this program. YaST then prompts the administrator for partitioning schemes, software selection, time zone, and other configuration options. After the administrator has provided all the required information, the installation proceeds. The system is partitioned and the requested software is installed.

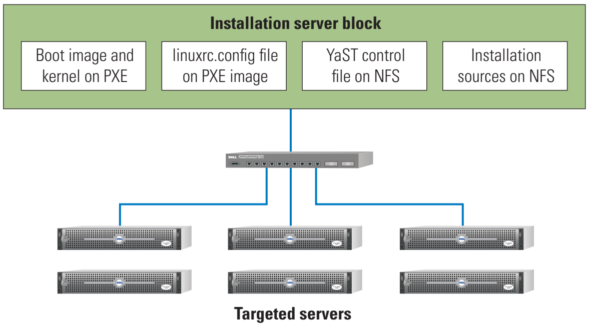
Installation server block
Boot image and kernel on PXE
linuxrc.config file on PXE image
YaST control file on NFS
Installation sources on NFS
Targeted servers

Figure 1. Framework for automated installation of Novell SUSE Linux Enterprise Server 9

During this process, YaST prompts the administrator multiple times to change the OS installation CDs. After the installation is complete, YaST prompts the administrator for post-installation configuration information such as display properties, network configuration, and a root password.

Framework for automated installation

The automated installation framework presented in this article specifies three steps:

- Setting up an AutoYaST control file
- Configuring a network server containing SLES 9 installation sources
- Creating the SLES 9 installation option on a server that supports the Preboot Execution Environment (PXE)

The AutoYaST control file, which is similar to a Red Hat® Enterprise Linux kickstart file, contains configuration information for the system to be installed. The network server contains the SLES 9 installation media configured for access over the network by YaST. The PXE server contains the SLES 9 installation kernel and the initial RAM disk (initrd) and allows them to be downloaded over the network for a network boot. Figure 1 shows the components involved in the automated installation framework.

The automated installation process follows this sequence:

1. The target installation server downloads the SLES 9 installation kernel and initrd to the target system using the PXE boot process.
2. The target system loads the Linux kernel and initrd and starts the linuxrc program.
3. Linuxrc loads YaST and the AutoYaST control file from the network.

After these steps complete, the OS installation and configuration proceed and eventually complete.

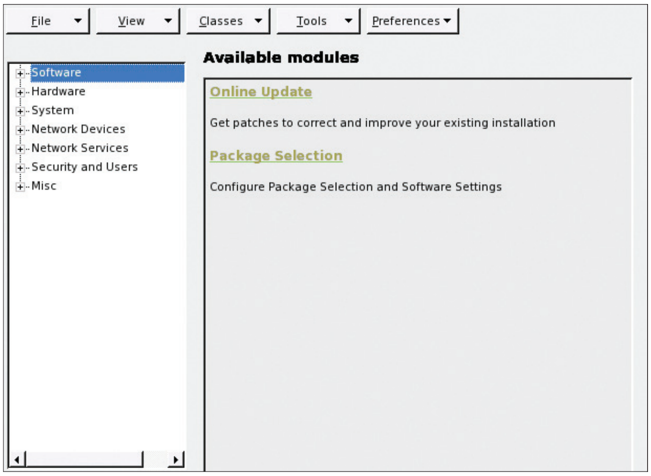

Available modules

Online Update

Get patches to correct and improve your existing installation

Package Selection

Configure Package Selection and Software Settings

Figure 2. The AutoYaST module screen

Preparing an AutoYaST control file

The AutoYaST control file contains all the configuration details required by the YaST installer program during installation. This file is most easily created using the YaST AutoYaST module, which provides a user-friendly graphical interface that guides the administrator through the control-file creation process. The AutoYaST module can be started by entering `yast2 autoyast` at a command prompt within the X Window System. Figure 2 shows the AutoYaST module screen.

Installation options can be configured through AutoYaST. Submenus are presented within each category and can be used to configure items such as software to be installed, partitions, security options, and network configuration. Administrators should complete all the necessary configuration choices and then save this file as "autoinst.xml." Figure 3 shows an example autoinst.xml file.

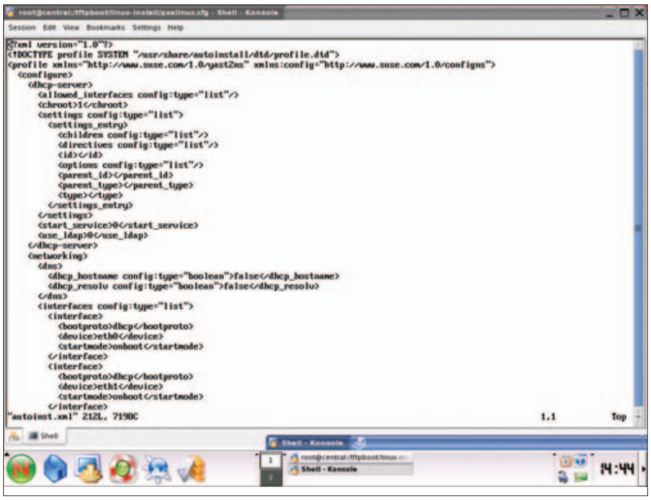
root@central:/tftpboot/linux-install/pxelinux.cfg - Shell - Konsole
Session
Edit
View
Bookmarks
Settings
Help
Shell
Shell - Konsole
14:44

```
<?xml version="1.0"?>
<!DOCTYPE profile SYSTEM "/usr/share/autoinstall/dtd/profile.dtd">
<profile xmlns="http://www.suse.com/1.0/yast2ns" xmlns:config="http://www.suse.com/1.0/configns">
  <configure>
    <dhcp-server>
      <allowed_interfaces config:type="list"/>
      <chroot>1</chroot>
      <settings config:type="list">
        <settings_entry>
          <children config:type="list"/>
          <directives config:type="list"/>
          <id></id>
          <options config:type="list"/>
          <parent_id></parent_id>
          <parent_type></parent_type>
          <type></type>
        </settings_entry>
      </settings>
      <start_service>0</start_service>
      <use_ldap>0</use_ldap>
    </dhcp-server>
    <networking>
      <dns>
        <dhcp_hostname config:type="boolean">false</dhcp_hostname>
        <dhcp_resolv config:type="boolean">false</dhcp_resolv>
      </dns>
      <interfaces config:type="list">
        <interface>
          <bootproto>dhcp</bootproto>
          <device>eth0</device>
          <startmode>onboot</startmode>
        </interface>
        <interface>
          <bootproto>dhcp</bootproto>
          <device>eth1</device>
          <startmode>onboot</startmode>
        </interface>
"autoinst.xml" 212L, 7198C                                    1,1           Top
```

Figure 3. Example autoinst.xml file viewed as plain text

SLES CD 1	*Installation directory*/SLES/CD1/
SLES CD 2	*Installation directory*/CORE/CD1/
SLES CD 3	*Installation directory*/CORE/CD2/
SLES CD 4	*Installation directory*/CORE/CD3/
SLES CD 5	*Installation directory*/CORE/CD4/
SLES CD 6	*Installation directory*/CORE/CD5/
Service Pack CD 1	*Installation directory*/SP/CD1/
Service Pack CD 2	*Installation directory*/SP/CD2/

Figure 4. Installation directories for Novell SUSE Linux Enterprise Server 9 CDs

Configuring the installation server

The installation target system can use network protocols such as Network File System (NFS), HTTP, or FTP to access the SLES 9 installation sources over the network, rather than obtaining them from a CD. To access installation sources, a network server must be configured to make the sources available and must have those sources configured as described in this section.

Setting up the installation repository. In SLES 9, setting up an installation repository is not straightforward as it often is in other Linux distributions. The SLES CD set comprises one SLES personality CD, five CORE CDs, and two Service Pack (SP) CDs. The first CD of the SLES CD set is the boot CD, and the second through sixth CDs are the CORE CDs.

Administrators must choose a network protocol for sharing the installation sources over the network—the example in this article uses NFS—and a directory from which to share the sources. Administrators can follow the directory format shown in Figure 4 for creating the SUSE installation source. Administrators must create the SLES, CORE, and SP subdirectories, and then copy the eight SLES CDs to the installation directory in the sequence shown in Figure 4.

Creating symbolic links. After all the CDs are copied, symbolic links to the files must be created to complete the process. These links should reside in the installation directory:

```
# cd  Installation directory
# ln -s  SLES/CD1/boot
# ln -s  SLES/CD1/control.xml
# ln -s  SLES/CD1/media.1
```

Administrators including a Service Pack in the installation must create one more link for the /SP/CD1/driverupdate file in the installation directory. This will help the installer to use the latest driver files in the installation:

```
# ln -s SP/CD1/driverupdate
```

If administrators do not plan to use a Service Pack during installation, there is no need to include the SP directories.

Creating the yast directory. After creating the symbolic links in the directory as detailed in the preceding section, administrators must create a directory and name it “yast.” This directory will contain files that control the installation order of the SLES CD directories. Administrators should create a file named “order” in the YaST directory, and add the following lines to the file (press Tab to separate the entries):

```
/SP/CD1      /SP/CD1
/SLES/CD1    /SLES/CD1
/CORE/CD1    /CORE/CD1
```

Administrators should then create a file named “instorder” and add the following entries:

```
/SP/CD1
/SLES/CD1
/CORE/CD1
```

Once these steps are complete, the installation server has been configured.

Setting up PXE boot

After configuring the AutoYaST control file and the installation sources, administrators should set up a PXE server to allow automated installation.[1] To enable PXE booting, administrators must perform the following steps:

1. Place the AutoYaST control file in the root level of the installation sources directory that was previously configured.
2. Locate the files named “linux” (kernel) and “initrd” in the /boot/loader directory on the first SP CD (or SLES CD 1 if a Service Pack is not employed), and copy them to the correct location in the /tftpboot directory on the PXE server.
3. Modify the linuxrc.config file, which is located in the initrd file, to define the location on the network of the installation sources and the AutoYaST control file. This file sets variables for the linuxrc program to use during initialization of the installation environment.

[1] For more information about setting up a PXE server using PXELINUX, see “Harnessing PXE Boot Services in Linux Environments” by John Hull, Robert Hentosh, and Rogelio Noriega in *Dell Power Solutions,* June 2004; www.dell.com/downloads/global/power/ps2q04-028.pdf.

```
KernelPCMCIA:   1
UseUSBSCSI:     1
Product:        SUSE LINUX Enterprise Server 9
UpdateDir:      /linux/suse/x86_64_sles9
MemLoadImage    163840
MinMemory:      256
MemYaST:        64000
MemYaSTText:    64000
ModuleDisks:    0
Language:       en_US
Keytable:       english-us
Install:        nfs://172.16.64.1/pub/nfs/suse/
                sles/9/sp2/x86_64/install
InstMode:       nfs
Insmod:         e1000
Netdevice       eth0
Netconfig:      dhcp
Textmode:       1
AutoYaST:       nfs://172.16.64.1/pub/nfs/
                autoinst.xml
```

Figure 5. Example contents of a /mnt/tmp/linuxrc.config file

To modify the linuxrc.config file, complete the following steps as the root administrator:

```
# mv initrd initrd.gz
# gunzip initrd
# mount -o loop initrd /mnt/tmp
```

4. Open the /mnt/tmp/linuxrc.config file for editing. An example of the linuxrc.config file is shown in Figure 5. Add the following entries to the this file:

   ```
   Install: network path to installation sources
   AutoYaST: network path to AutoYaST file
   ```

5. Modify the existing /mnt/tmp/linuxrc.config file values with the following values to start the installation program:

   - `Language`: Default language in which to start the installer
   - `Keytable`: Keyboard layout
   - `Install`: Installation method
   - `InstMode`: Installation mode—for example, NFS or FTP
   - `Insmod`: Additional module to load other than auto-detected modules
   - `Netdevice`: Device to use to start the installation
   - `Netconfig`: Mode of IP configuration—for example, DHCP
   - `Textmode`: 1 or 0 for enabling text-mode installation
   - `AutoYaST`: Path of AutoYaST control file—for example, floppy or NFS path

6. After completing the linuxrc.config file edit, recompress the initrd file to its original state:

   ```
   # umount /mnt/tmp
   # gzip -9 initrd
   # mv initrd.gz initrd
   ```

7. Finally, create an entry in one of the PXELINUX configuration files to allow a PXE-based boot of the kernel and initrd.

Integrating the system

After the entire setup is done, administrators can connect the target server to the installation server network. The target server should boot to PXE using the linuxrc.config and initrd files hosted at the installation server. This process will use the parameters embedded in the linuxrc.config program to start YaST.

Before starting YaST, the linuxrc program should be able to locate the SLES 9 CD repository. After YaST is started, it will search for the AutoYaST control file and use the parameters found there to continue the installation. After installation is complete, the system will reboot in a fully configured, ready-to-use state.

A productivity-enhancing installation tool

Implementing an unattended installation of Novell SUSE Linux Enterprise Server 9 can be helpful to system administrators who must quickly deploy and configure operating systems. Any enterprise can reap the benefits of this approach, but it will be especially useful for data centers or high-performance computing cluster environments that contain numerous servers to configure and manage.

Anoop K. is a senior engineer on the Linux engineering team at the Dell Bangalore Development Center. He has six years of industrial experience and is presently attaining his B.S. in Information Systems from Birla Institute of Technology and Science (BITS) in Pilani, India.

John Hull is a software engineer at Dell and is the lead for the SUSE Linux program. He has a B.S. in Mechanical Engineering from the University of Pennsylvania and an M.S. in Mechanical Engineering from the Massachusetts Institute of Technology.

FOR MORE INFORMATION

Novell Linux Desktop: AutoYaST Installation Tool:
www.novell.com/products/desktop/features/autoyast.html

Novell Linux Desktop Cool Solutions:
www.novell.com/coolsolutions/nld

VMware ESX Server Performance

on Dual-Core Dell PowerEdge 2850 Servers

Several Dell™ PowerEdge™ server models can be equipped with dual-core Intel® Xeon® processors. Dell engineers tested the Dell PowerEdge 2850 server with both single-core and dual-core processors, and the results showed a 28 to 51 percent performance gain when moving from single-core to dual-core processors.

BY TODD MUIRHEAD AND DAVE JAFFE, PH.D.

Related Categories:

Characterization

Multi-core technology

Performance

Virtualization

VMware

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

Multi-core processors are designed to improve hardware performance compared to traditional single-core processors. Dual-core processors from Intel are now available in eighth-generation Dell PowerEdge server models, including the Dell PowerEdge 2850. To determine the benefits of dual-core processors in a VMware ESX Server software–based virtual computing environment, a team of Dell engineers tested two Dell PowerEdge 2850 servers, one equipped with two single-core processors and another equipped with two dual-core processors. On these servers, the test team deployed VMware ESX Server software and created multiple virtual machines (VMs) to run three different types of enterprise applications: a Microsoft® SQL Server™ database, the NetBench benchmark (simulating a file server), and the LAMP (Linux®, Apache, MySQL, PHP) Web platform.

Hardware configuration and setup

The Dell PowerEdge 2850 server is a 2U dual-processor system with an 800 MHz frontside bus and up to 16 GB of RAM. Both PowerEdge 2850 servers used for testing were configured with 8 GB of RAM. The PowerEdge 2850 server can be equipped with either single-core or dual-core Intel Xeon processors. In the test environment, the single-core processors were 3.6 GHz and the dual-core processors were 2.8 GHz. The PowerEdge 2850 server supports up to six internal disks, offers three PCI slots, and includes dual on-board Gigabit Ethernet[1] network interface cards (NICs). In each test server, two of the PCI slots were used for QLogic QLA2340 Fibre Channel host bus adapters (HBAs) to provide connectivity to the storage area network (SAN). An Intel PRO 1000XT Gigabit Ethernet NIC was used in the third PCI slot. By equipping each server with three NICs, testers were able to dedicate one NIC to the VMware ESX Server service console, one to the VMs, and one to VMware VMotion™ VM migrations. Figure 1 summarizes the configuration of the PowerEdge 2850 servers tested.

Both servers were connected to the SAN via the QLogic QLA2340 HBAs. The dual HBAs in each server enabled ESX Server software to provide failover across multiple paths to the logical units (LUNs). SAN storage was provided by a Dell/EMC CX700 array. The VMs used in the test were spread across ten 73 GB, 10,000 rpm Fibre Channel disks located on the CX700 storage array. These disks were divided into two five-disk (4 + 1) RAID-5 LUNs. The VMs running Microsoft SQL Server, NetBench, and LAMP were evenly divided on the two LUNs so that half of each workload type were on each LUN. Figure 2 summarizes the SAN configuration used in the test environment.

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

Virtualization software	VMware ESX Server 2.5.2
CPU	Two single-core Intel Xeon processors DP at 3.6 GHz with 2 MB L2 cache or two dual-core Intel Xeon processors DP at 2.8 GHz with 2 MB L2 cache per core
Memory	8 GB
Internal disks	Two 73 GB drives
NICs	Two 10/100/1,000 Mbps internal NICs and one Intel PRO 1000XT Gigabit Ethernet NIC
Disk controller	Dell PowerEdge RAID Controller 4, embedded integrated
Fibre Channel HBA	QLogic QLA2340 HBA
Form factor	2U (3.5 inches)

Figure 1. Configuration for PowerEdge 2850 servers used in test environment

VMware ESX Server software for virtualized platform

The Dell team installed VMware ESX Server 2.5.2 virtual infrastructure software on both Dell PowerEdge 2850 servers used in testing. VMware ESX Server allows several operating systems and applications to run on the same physical server simultaneously by providing a virtualization layer that resides just above the hardware layer. ESX Server software creates VM "containers" on the physical server, enabling each VM to run its own OS, which in turn governs its own set of applications and services.

Because ESX Server isolates each VM from other VMs residing on the same physical server just as physical systems are isolated from one other, administrators have a great amount of flexibility in using ESX Server to run different types of applications and operating systems at the same time. Each VM can be rebooted or powered down without affecting other VMs running on the same physical server. This capability allows administrators to patch or upgrade and then reboot an application on one VM without incurring downtime for other VMs running on that physical system.

Test workloads to simulate multiple applications

To simulate how enterprises typically run applications on VMs using ESX Server, in October 2005 the test team increased the number of VMs until CPU utilization for the entire physical server reached 85 percent. This represents a reasonably high level of usage for a production server, but it is well below the maximum 100 percent utilization that is used by many industry-standard benchmarks—and a 100 percent utilization level is not typically reached in production.

To compare relative performance differences between single-core and dual-core processors in the PowerEdge 2850 server, the test team ran three different types of enterprise workload: Microsoft SQL Server 2000 with an online transaction processing (OLTP) workload, Novell® SUSE® Linux with a LAMP stack, and Microsoft Windows Server™ 2003 with NetBench 7.03.[2]

Each workload ran on multiple VMs and under the same load at the same time. By keeping all settings on the VM and driver systems identical and then observing how many VMs could be run simultaneously, the test team was able to measure how many VMs the physical server could support. Figure 3 shows the configuration of the VMs used in the test environment.

Microsoft SQL Server 2000. On the SQL Server 2000 VMs, the test team installed Microsoft Windows Server 2003 and SQL Server 2000 with Service Pack 4 (SP4). The SQL Server version of the Dell DVD Store database was loaded into SQL Server 2000 using the scripts provided with the DVD Store application. The complete DVD Store application code, including the SQL Server version and a LAMP version, is freely available for public use under the GNU General Public License (GPL) at linux.dell.com/dvdstore. The DVD Store database simulates the database back end of a simple Web-based storefront. The database is small, approximately 1 GB, and representative of a database used for development or testing.

To simulate a load against the VM, the test team used the DVD Store driver program, which is included as part of the DVD Store download. Each SQL Server 2000 VM was driven by a single thread of the driver application with a 20-millisecond delay.

SUSE LAMP. For the LAMP workload, the test team installed Novell SUSE Linux Enterprise Server 9, Apache 2, and MySQL 5 on a VM. The MySQL version of the Dell DVD Store application was loaded into MySQL 5, and the PHP version of the DVD Store application was

Controller	Dell/EMC CX700 storage array
Disk enclosure	Dell/EMC DAE2 disk array enclosure
Disks	Ten 73 GB, 10,000 rpm drives
LUNs	Two five-disk RAID-5 LUNs
Software	EMC Navisphere® Manager and EMC Access Logix™

Figure 2. Configuration for SAN used in test environment

Workload	RAM	Disk	Virtual NIC type	Number of virtual CPUs
SQL Server 2000	512 MB	10 GB	Vmxnet	1
SUSE LAMP	1,024 MB	10 GB	Vlance	1
NetBench	512 MB	10 GB	Vmxnet	1

Figure 3. Configuration for VMs used in test environment

[2] The three workloads used for this test were also used in a previous study. For more information about that study, see "VMware ESX Server Performance on Dell PowerEdge 2850 and PowerEdge 6850 Servers" by Todd Muirhead; Dave Jaffe, Ph.D.; and Scott Stanford in *Dell Power Solutions*, February 2006; www.dell.com/downloads/global/power/ps1q06-20050312-Muirhead.pdf.

set up on Apache. In this setup, the Web tier and the database tier ran on the same VM to have a complete LAMP stack.[3]

The driver for the LAMP stack differs from the driver used in the SQL Server testing in that it emits HTTP requests and receives HTML code returned from the Apache/PHP layer, whereas the SQL Server driver communicates directly with the database. However, the test team measured the same parameters for this workload: total orders per minute (OPM) handled by the application and average response time as experienced by simulated customers. Each SUSE LAMP VM was driven by a single thread of the driver program with a 30-millisecond delay in this test.

	PowerEdge 2850 with single-core Intel Xeon processors at 3.6 GHz			PowerEdge 2850 with dual-core Intel Xeon processors at 2.8 GHz			
Workload	**Number of VMs**	**Performance**	**CPU utilization for ESX Server host**	**Number of VMs**	**Performance**	**CPU utilization for ESX Server host**	**Performance gain compared to single-core processors**
SQL Server 2000	10	1,094 OPM	86%	15	1,654 OPM	86%	51%
SUSE LAMP	13	1,685 OPM	86%	16	2,156 OPM	86%	28%
NetBench	11	295 MB/sec	85%	18	443 MB/sec	85%	50%

Figure 4. Multiple-workload test results for PowerEdge 2850 configured with single-core and dual-core processors

NetBench 7.03. NetBench 7.03, developed by *PC Magazine,* is a benchmark tool that is designed to simulate a file server workload. The program creates and accesses a set of files according to predefined scripts. NetBench is typically run with an increasing number of client engines running against a single server to measure how much throughput (in megabytes per second) can be achieved with a given number of client connections.

To determine how many VMs could run on an ESX Server host, the Dell test team increased the number of VMs and the number of client engines at the same rate, until the CPU utilization on the ESX Server host reached 85 percent. NetBench 7.03, with the included standard DiskMix script, was used with a 0.6-second think time to connect two client engines to each VM. This simulates multiple file servers hosted on the same ESX Server system, similar to a file server consolidation scenario. The driver systems on which the client engines ran had mapped drives to all of the VMs in the test. In NetBench, the test directories path file was modified so that, as successive client engines were added, they would use the next drive letter, which corresponded to the next VM.

Results of single-core and dual-core performance tests

First, the VMs used in testing were run on the PowerEdge 2850 with single-core Intel Xeon processors at 3.6 GHz in successive tests, adding VMs in each round until the total CPU utilization on the system reached 85 percent. Then, using VMware VMotion software, the VMs were moved to the PowerEdge 2850 with dual-core Intel Xeon processors at 2.8 GHz and the process was repeated. The difference in the number of VMs and the associated performance metric—OPM for SQL Server and LAMP and megabytes per second for NetBench—indicated the relative difference in performance. The test team calculated the percent performance gain by using these two metrics.

In this test study, the SQL Server 2000, SUSE LAMP, and NetBench workloads all demonstrated a significant increase in performance when moved from single-core to dual-core Intel Xeon processors. Performance gains of 51 percent for the SQL Server 2000 workload and 50 percent for the NetBench workload were achieved on the dual-core PowerEdge 2850. The LAMP workload showed a 28 percent increase in performance when running on the dual-core system. These test results demonstrate that performance gains range depending on the type of workload. Figure 4 summarizes the test results for all three workloads running on the PowerEdge 2850 in both single-core and dual-core configurations.

Multi-core technology in virtual environments

Dual-core Intel Xeon processors are designed to provide a significant boost in performance over single-core Intel Xeon processors. In the Dell tests described in this article, a VMware ESX Server–based virtual computing environment demonstrated performance gains of up to 51 percent when workloads were moved from a single-core to a dual-core platform. The test results also indicate that performance increases vary depending on the type of work being performed by the VMs. In addition, these test results show that more VMs can be hosted on a physical server using dual-core processors compared to a physical server using single-core processors—and thus a greater overall throughput is possible with dual-core processors.

Todd Muirhead is a senior engineering consultant on the Dell Scalable Enterprise Technology Center team. Todd has a B.A. in Computer Science from the University of North Texas.

Dave Jaffe, Ph.D., is a senior consultant on the Dell Scalable Enterprise Technology Center team. He has a B.S. in Chemistry from Yale University and a Ph.D. in Chemistry from the University of California, San Diego.

[3] The LAMP stack has been fully documented in the Dell Enterprise Product Group white paper, "MySQL Network and the Dell PowerEdge 2800: Capacity Sizing and Performance Tuning Guide for Transactional Applications" by Todd Muirhead, Dave Jaffe, and Nicolas Pujol; www.dell.com/downloads/global/solutions/mysql_network_2800.pdf.

Managing Dell PowerEdge Servers Using the

Dell Management Pack for Microsoft Operations Manager

The Dell™ Management Pack for Microsoft® Operations Manager is a software module designed to provide specific information about Dell applications and hardware. It can enhance the ability of administrators to manage Dell PowerEdge™ servers with Microsoft Operations Manager 2005. In addition, the Dell Management Pack can be customized to suit specific environments.

BY BALASUBRAMANIAM J.

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Microsoft Operations Manager (MOM)

Microsoft Windows

Systems management

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

To help manage hardware and software resources across a data center, IT organizations can deploy Microsoft Operations Manager (MOM) 2005. This comprehensive tool provides event management, proactive monitoring and alerting, and system and application knowledge to help reduce IT costs and improve systems availability. The Dell Management Pack (Dell MP) for Microsoft Operations Manager is a software module that can be easily integrated into MOM 2005. Management packs are predefined software packages for managing and monitoring specific environments.

The Dell MP provides ready-to-use knowledge for managing and monitoring Dell applications and hardware. It enables MOM to categorize Dell systems into a Dell-specific computer group and to accurately depict the status of each system on the network. The status monitoring of Dell systems includes Dell-specific alerts and pre-failure alerts that enable administrators to assess, respond to, and help ensure the availability of Dell systems being monitored in the MOM environment. The alerts also provide a link to launch Dell OpenManage™ Server Administrator (OMSA) and the Dell Remote Access Controller (DRAC) to further help manage and monitor Dell systems.

Importing the Dell Management Pack

Microsoft Operations Manager 2005 Workgroup Edition is designed for Microsoft Windows Server™ OS–based environments comprising 10 or fewer servers. When obtained from Dell, the MOM 2005 Workgroup Edition installation CD contains the Dell Management Pack, which is automatically imported during the MOM installation. Alternatively, enterprise IT organizations that have purchased the full edition of MOM 2005 or MOM 2005 Workgroup Edition can download the Dell MP from support.dell.com. In that

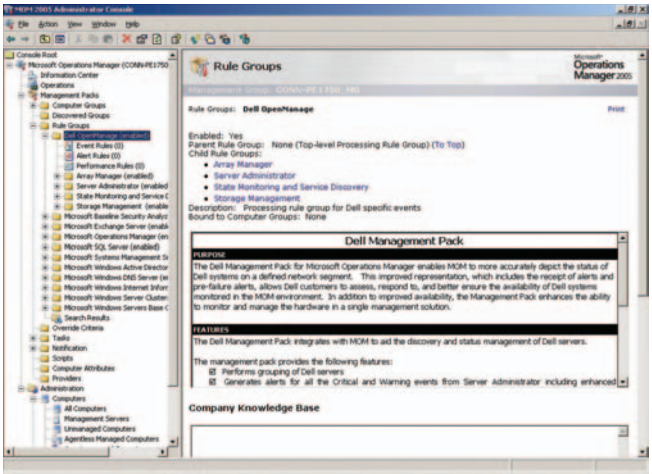
MOM 2005 Administrator Console
File
Action
View
Window
Help
Console Root
Microsoft Operations Manager (CONN-PE1750
Information Center
Operations
Management Packs
Computer Groups
Discovered Groups
Rule Groups
Dell OpenManage (enabled)
Event Rules (0)
Alert Rules (0)
Performance Rules (0)
Array Manager (enabled)
Server Administrator (enabled
State Monitoring and Service D
Storage Management (enable
Microsoft Baseline Security Analyz
Microsoft Exchange Server (enabl
Microsoft Operations Manager (en
Microsoft SQL Server (enabled)
Microsoft Systems Management S
Microsoft Windows Active Director
Microsoft Windows DNS Server (e
Microsoft Windows Internet Inforr
Microsoft Windows Server Cluster
Microsoft Windows Servers Base C
Search Results
Override Criteria
Tasks
Notification
Scripts
Computer Attributes
Providers
Administration
Computers
All Computers
Management Servers
Unmanaged Computers
Agentless Managed Computers
Rule Groups
Microsoft Operations Manager 2005
Management Group: CONN-PE1750_MG
Print
Company Knowledge Base

Rule Groups: **Dell OpenManage**

Enabled: Yes
Parent Rule Group: None (Top-level Processing Rule Group) (To Top)
Child Rule Groups:

- Array Manager
- Server Administrator
- State Monitoring and Service Discovery
- Storage Management

Description: Processing rule group for Dell specific events
Bound to Computer Groups: None

Dell Management Pack

PURPOSE

The Dell Management Pack for Microsoft Operations Manager enables MOM to more accurately depict the status of Dell systems on a defined network segment. This improved representation, which includes the receipt of alerts and pre-failure alerts, allows Dell customers to assess, respond to, and better ensure the availability of Dell systems monitored in the MOM environment. In addition to improved availability, the Management Pack enhances the ability to monitor and manage the hardware in a single management solution.

FEATURES

The Dell Management Pack integrates with MOM to aid the discovery and status management of Dell servers.

The management pack provides the following features:

☒ Performs grouping of Dell servers
☒ Generates alerts for all the Critical and Warning events from Server Administrator including enhanced

Figure 1. Dell OpenManage processing rule groups

case, administrators can import the Dell MP using the Management Pack Import Export wizard.

When imported into the MOM Administrator console, the Dell MP provides predefined computer groups, processing rule groups, computer attributes, scripts, tasks, the Dell Knowledge Base, and public views for managing Dell-specific applications and hardware. At the MOM Administrator console, administrators can create, import, or export management packs as well as configure MOM settings, discover systems, deploy agents, and create and maintain user privileges.

For effective management of Dell servers once the Dell MP has been imported into the MOM Administrator console, the Dell servers must be grouped into the Dell Computers group.

The Dell OpenManage processing rule group (PRG) is the "parent" rule group that consists of four "children" PRGs: Server Administrator, Array Manager, Storage Management, and State Monitoring and Service Discovery. Dell PRGs are associated with the Dell Computers group and are enabled by default. Figure 1 shows Dell PRGs displayed in the MOM Administrator console.

The Server Administrator, Array Manager, and Storage Management PRGs contain the event-processing rules for critical and warning events occurring in the Dell OpenManage Server Administrator, Dell OpenManage Array Manager, and Dell OpenManage Server Administrator Storage Management services, respectively. These event-processing rules generate an alert when a Dell event is detected and the rule criteria are met. Each of these PRGs also includes an alert-processing rule. The alert-processing rule sends out a notification message to the hardware support group that is provided in the MOM Administrator console on receipt of a critical or warning alert. *Note:* The Dell PRGs do not process informational events.

The State Monitoring and Service Discovery PRG contains the event-processing rules for monitoring the server and storage components in the Dell state view and also contains rules for service discovery. The state monitoring rules generate alerts that update the Dell state view. These state monitoring alerts are logged in the Dell State Monitoring Alerts view.

Dell Computers group. Dell Computers is the group defined to group all Dell servers. The criteria for grouping Dell systems into the Dell Computers group are based on the Dell computer attribute. The Dell processing rules are associated with this computer group, helping to define the computers on which the Dell processing rules must be deployed.

Dell Knowledge Base. The Dell MP provides a knowledge base for all Dell processing rules and Dell PRGs. The Dell Knowledge Base provides a brief summary of each event, its cause, and the recommended resolution (see Figure 2).

Discovering Dell servers

For effective management of Dell servers once the Dell MP has been imported into the MOM Administrator console, the Dell servers must be grouped into the Dell Computers group. For this to occur, all the Dell servers must have OMSA installed and must be listed in the Agent-Managed Computers view, which is located in the Administration > Computers section of the MOM Administrator console.

To discover the computers over the network, administrators can create a computer discovery rule and initiate a computer scan by right-clicking on "Computer Discovery Rules" and then

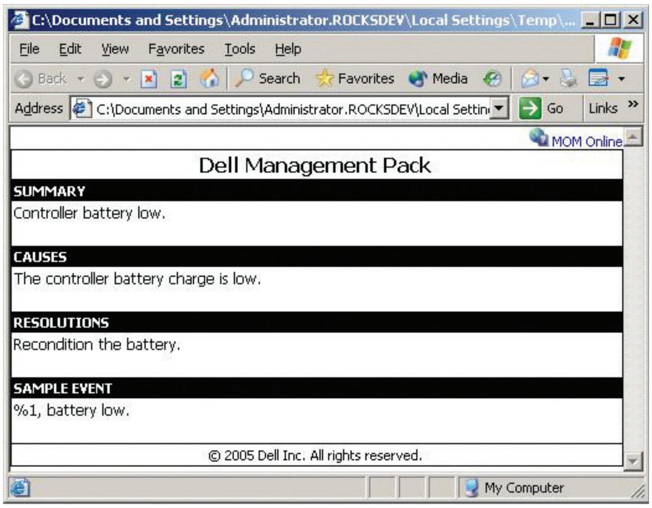
C:\Documents and Settings\Administrator.ROCKSDEV\Local Settings\Temp\...
File
Edit
View
Favorites
Tools
Help
Back
Search
Favorites
Media
Address
C:\Documents and Settings\Administrator.ROCKSDEV\Local Settin
Go
Links
MOM Online
My Computer

Dell Management Pack

SUMMARY

Controller battery low.

CAUSES

The controller battery charge is low.

RESOLUTIONS

Recondition the battery.

SAMPLE EVENT

%1, battery low.

© 2005 Dell Inc. All rights reserved.

Figure 2. Dell Knowledge Base information sample

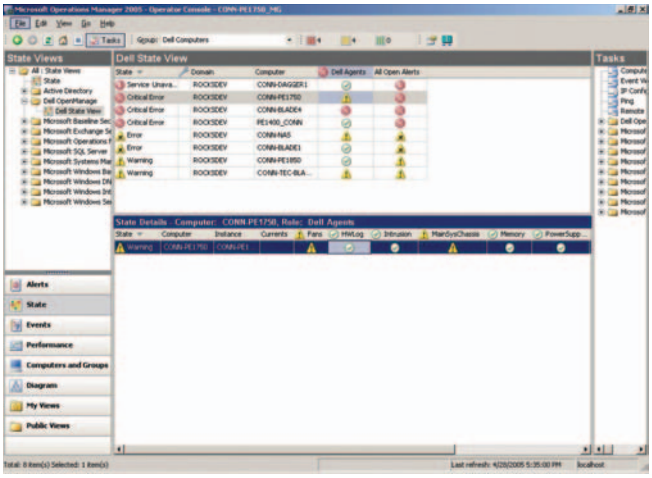

Dell State View

State	Domain	Computer	Dell Agents	All Open Alerts
Service Unava...	ROCKSDEV	CONN-DAGGER1	✓	●
Critical Error	ROCKSDEV	CONN-PE1750	⚠	●
Critical Error	ROCKSDEV	CONN-BLADE4	●	●
Critical Error	ROCKSDEV	PE1400_CONN	✓	●
Error	ROCKSDEV	CONN-NAS	⚠	⚠
Error	ROCKSDEV	CONN-BLADE1	✓	⚠
Warning	ROCKSDEV	CONN-PE1850	✓	⚠
Warning	ROCKSDEV	CONN-TEC-BLA...	⚠	⚠

State Details - Computer: CONN-PE1750, Role: Dell Agents

State	Computer	Instance	Currents	Fans	HWLog	Intrusion	MainSysChassis	Memory	PowerSupp...
Warning	CONN-PE1750	CONN-PE1		⚠	✓	✓	⚠	✓	✓

Total: 8 item(s) Selected: 1 item(s) Last refresh: 4/28/2005 5:35:00 PM localhost

Figure 3. Dell state view

selecting "Run Computer Discovery Now." For more information about creating computer discovery rules, refer to MOM documentation at www.microsoft.com/mom/techinfo/productdoc. Once the scan is completed, all systems with OMSA installed are grouped under the Dell Computers group.

Monitoring and managing Dell servers

The Dell MP monitors and manages the Dell servers that are grouped in the Dell Computers group. The management pack provides a list of public views to help administrators monitor and manage the Dell systems in the network.

Dell state view

The Dell MP provides a Dell-specific state view reflecting the health of each Dell server managed by MOM. To view the Dell state view, administrators should open the MOM Operator console, which is used to monitor system health, view problems, and obtain recommended resolutions. In the MOM Operator console, administrators can click the State tab in the bottom left pane. This shows the state views of all imported management packs. Expanding "Dell OpenManage" in the State Views tree shows the Dell state view (see Figure 3).

The Dell state view shows the status of each Dell hardware component. State changes are triggered by the MOM alert infrastructure.[1] The component status is designated by colors—red, yellow, and green—that represent the level of alert severity. The health of a component is derived by reviewing its non-resolved alerts. The status becomes the severity level of the most severe, non-resolved alert that has an active problem. If a component is not present in a specific server, its status is shown as white. To set the state for the Dell MP, the Dell service discovery script should run at least once and discovery instances of Dell servers should exist.

The Dell state view is based on roles and components. Dell agents represent the role, and the hardware components are the parts of the role that sum up the global health for that particular role. The health of any Dell server is dependent on its components' global health.

Hardware component status is updated whenever an event is sent by the managed node instrumentation service to the Windows Event Log. Server status can also be updated by administrative action. To update the status, administrators can select a server and click the "Update status" task that appears on the right pane of the MOM Operator console.

Updating the Dell state view based on events. The Dell state view is updated whenever a Dell event occurs. The Dell agents generate an event when a hardware component changes its status. The Dell MP has an event-processing rule to update the status of the components in the state view. This event-processing rule calls a Dell script, which is then executed locally on the managed node using the local system account. The script queries the global status of the individual components. Based on the status, the script generates an event for each component. Another event-processing rule processes this script-generated event and generates the state-based alert for updating the state view. Figure 4 shows the steps taken to update the state view when an event is generated.

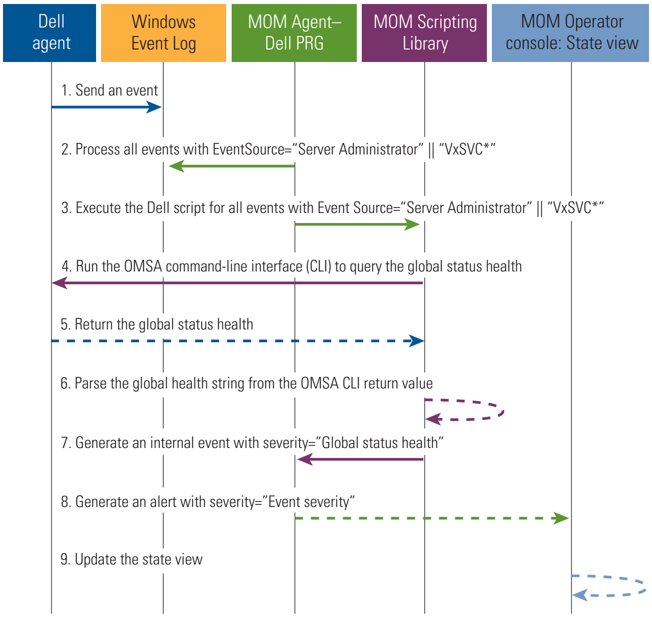

1. Send an event
2. Process all events with EventSource="Server Administrator" || "VxSVC*"
3. Execute the Dell script for all events with Event Source="Server Administrator" || "VxSVC*"
4. Run the OMSA command-line interface (CLI) to query the global status health
5. Return the global status health
6. Parse the global health string from the OMSA CLI return value
7. Generate an internal event with severity="Global status health"
8. Generate an alert with severity="Event severity"
9. Update the state view

Figure 4. Updating the Dell state view using events

[1] The Dell instrumentation (OMSA, Array Manager, and OMSA Storage Management) generates events, and in turn, the Dell MP issues alerts upon receipt of Dell events.

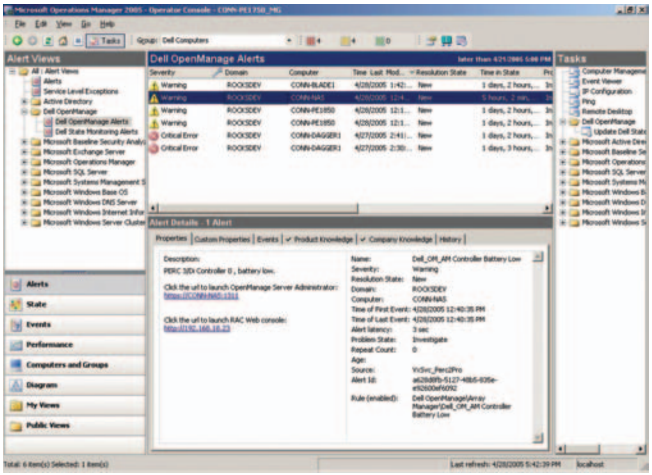

Dell OpenManage Alerts

Severity	Domain	Computer	Time Last Mod...	Resolution State	Time in State
Warning	ROCKSDEV	CONN-BLADE1	4/28/2005 1:42:...	New	1 days, 2 hours,...
Warning	ROCKSDEV	CONN-NAS	4/28/2005 12:4...	New	5 hours, 2 min,...
Warning	ROCKSDEV	CONN-PE1850	4/28/2005 12:1...	New	1 days, 2 hours,...
Warning	ROCKSDEV	CONN-PE1850	4/28/2005 12:1...	New	1 days, 2 hours,...
Critical Error	ROCKSDEV	CONN-DAGGER1	4/27/2005 2:41:...	New	1 days, 2 hours,...
Critical Error	ROCKSDEV	CONN-DAGGER1	4/27/2005 2:30:...	New	1 days, 3 hours,...

Alert Details - 1 Alert

Description:
PERC 3/Di Controller 0 , battery low.

Click the url to launch OpenManage Server Administrator:
https://CONN-NAS:1311

Click the url to launch RAC Web console:
https://192.168.18.23

Name: Dell_OM_AM Controller Battery Low
Severity: Warning
Resolution State: New
Domain: ROCKSDEV
Computer: CONN-NAS
Time of First Event: 4/28/2005 12:40:35 PM
Time of Last Event: 4/28/2005 12:40:35 PM
Alert latency: 3 sec
Problem State: Investigate
Repeat Count: 0
Age:
Source: VxSvc_Perc2Pro
Alert Id: a628d0fb-5127-40b5-805e-e92600ef6092
Rule (enabled): Dell OpenManage\Array Manager\Dell_OM_AM Controller Battery Low

Total: 6 item(s) Selected: 1 item(s)
Last refresh: 4/28/2005 5:42:39 PM

Figure 5. Dell OpenManage Alerts view

Updating the Dell state view based on an administrative request. Administrators can update the Dell state view by selecting the system from the state view for which the status needs to be updated and clicking the "Update status" task in the right-hand task pane of the MOM Operator console. This task launches the same script that queries the global health of the individual components. The status of all the components in the state view is updated when the task is executed.

Alert views

The Dell MP provides two alert views: Dell OpenManage Alerts view and Dell State Monitoring Alerts view. To locate these views, administrators can launch the MOM Operator console, click the Alerts tab in the bottom left pane, and expand "Dell OpenManage" in the Alert Views tree. These two views provide the following information:

- **Dell OpenManage Alerts view:** Displays warning and critical alerts generated by OMSA (including OMSA Storage Management) and Array Manager events
- **Dell State Monitoring Alerts view:** Displays warning and critical alerts that change the state of the components in the Dell state view

The Alert Details pane provides a detailed description of each alert (see Figure 5). To obtain more information about the alert, administrators can select the Product Knowledge tab. The information includes a description of the alert, possible causes for it, and any action plans to resolve the issue. If the alert is a Dell alert, the OMSA launch point and the DRAC launch point are provided in the alert description on the Properties tab of the Alert Details pane. If administrators click these Web addresses, they can launch the Web interface for OMSA or DRAC and retrieve more information about the system that generated the alert.

Diagram views

The Dell MP performs a service discovery of the hardware properties on computers in addition to discovering the roles and components for each computer. It offers the following diagram views that represent the Dell servers in the network:

- **Dell Computer Group Listing:** Displays the Dell Computers group and the Dell systems that are part of this group, along with system attributes and status (see Figure 6)
- **Dell Computer Listing:** Displays all Dell systems in the network, along with their attributes and status

These diagram views list all discovered Dell servers in the network and their status. The status is cumulative for all open alerts generated from the respective computers. A Dell computer icon represents all these servers. Dell-specific information and other attributes are shown in this view when a mouse cursor moves over the computer icons. Figure 7 shows which information is provided in this view.

The Dell asset tag value can be set from OMSA by entering a value for the Chassis Asset Tag field in the Information section of the Main System Chassis page of the OMSA Web interface. For more information about launching OMSA from the MOM Operator console, see the "Integrating MOM with Dell OpenManage Server Administrator for one-to-one management" section in this article.

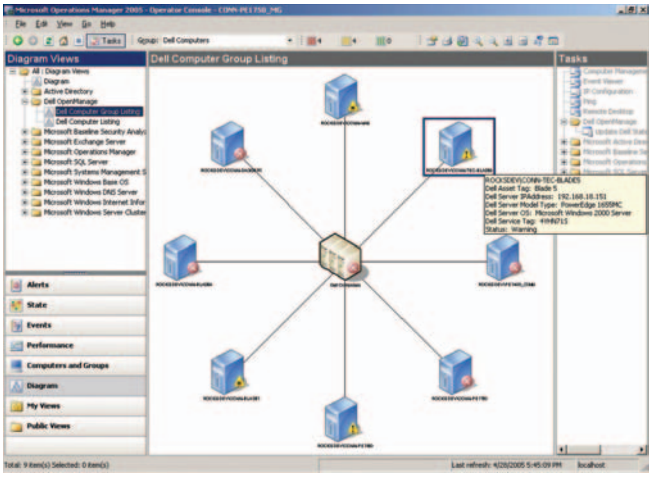
Microsoft Operations Manager 2005 - Operator Console - CONN-PE1750_MG
File Edit View Go Help
Tasks
Group: Dell Computers
Diagram Views
All : Diagram Views
Diagram
Active Directory
Dell OpenManage
Dell Computer Group Listing
Dell Computer Listing
Microsoft Baseline Security Analy
Microsoft Exchange Server
Microsoft Operations Manager
Microsoft SQL Server
Microsoft Systems Management S
Microsoft Windows Base OS
Microsoft Windows DNS Server
Microsoft Windows Internet Infor
Microsoft Windows Server Cluster
Alerts
State
Events
Performance
Computers and Groups
Diagram
My Views
Public Views
Dell Computer Group Listing
Tasks
Computer Manageme
Event Viewer
IP Configuration
Ping
Remote Desktop
Dell OpenManage
Update Dell State
Dell Computers

ROCKSDEV\CONN-TEC-BLADES
Dell Asset Tag: Blade 5
Dell Server IPAddress: 192.168.18.151
Dell Server Model Type: PowerEdge 1655MC
Dell Server OS: Microsoft Windows 2000 Server
Dell Service Tag: 4HN715
Status: Warning

Total: 9 item(s) Selected: 0 item(s)
Last refresh: 4/28/2005 5:45:09 PM

Figure 6. Dell Computer Group Listing diagram view

Event views

The Dell event views display all the Dell events that are processed by the Dell MP. The Dell MP supports events from OMSA—including the OMSA Storage Management service—and Dell OpenManage Array Manager. To monitor and manage the status of Dell servers, the MOM Agent must be installed and running with the latest rules on the Dell servers. Whenever a critical or a warning event is generated by a Dell server, the matching event-processing rule generates an alert. This alert is displayed in the Dell OpenManage Alerts view created by the Dell MP. Alerts are listed in descending order of severity.

The Dell Remote Access Controller is designed to provide remote management capabilities, crashed system recovery, and power-control functions for Dell PowerEdge servers.

Once the alert has been resolved, the state of the alert in the MOM Operator console must be resolved manually. If the alert state is not set to "resolved," then the status of the server is not updated.

While resolving an alert, administrators also can add company-specific knowledge base information to the rule. If the knowledge base information needs to be updated for this alert, administrators can go to the Company Knowledge tab in the Alert Details pane and click the Edit button before resolving the alert. They can then enter the knowledge base information specific to that alert and click the OK button.

The Resolution State field for the alert in the MOM Operator console should be set to "Resolved." Once the alert is resolved, it is removed from the Dell OpenManage Alerts view and the updated status is propagated to the parent nodes.

Integrating MOM with Dell OpenManage Server Administrator for one-to-one management

Dell OpenManage Server Administrator must first be installed and running on all the Dell servers that are managed by MOM through the Dell MP. Dell servers running OMSA can log events in the Windows Event Viewer. Events that match the criteria of the Dell PRGs generate alerts that are forwarded to the management station running MOM. Each Dell event-processing rule is associated with the VBScript-based DellOMSALaunch custom script. This script is executed when an alert is received from the managed system. The script appends to the alert description the URL for launching the OMSA Web interface. Administrators can click this link to launch OMSA for the server that generated the alert (see Figure 8). They can then log in to OMSA for one-on-one management of that server. Once the issue has been resolved on the Dell server, administrators can resolve the alert in the MOM Operator console.

Dell service tag	An alphanumeric serial number that uniquely identifies each Dell system
Dell asset tag	An individual code assigned to a system, usually by a system administrator, for security or tracking purposes
Dell system IP address	The IP address of the Dell system in the network
Dell server OS	The OS installed on the Dell system
Dell model type	The Dell PowerEdge or PowerVault™ model name

Figure 7. Types of information provided in Dell diagram views within the MOM Operator console

The OMSA URL can also be found in the alert description of Dell state monitoring alerts. Administrators need not explicitly resolve state monitoring alerts. When the issue is resolved, the alert state is automatically set to inactive. The state is set back to active only if the component's status changes to a critical or warning state.

Integrating MOM with the Dell Remote Access Controller for remote management

The Dell Remote Access Controller is designed to provide remote management capabilities, crashed system recovery, and power-control functions for Dell PowerEdge servers. The Dell MP supports third- and fourth-generation DRACs. In the Dell alert description, the Dell MP provides a launch point for the DRAC Web interface (see Figure 9), which can be used to remotely manage the Dell server.

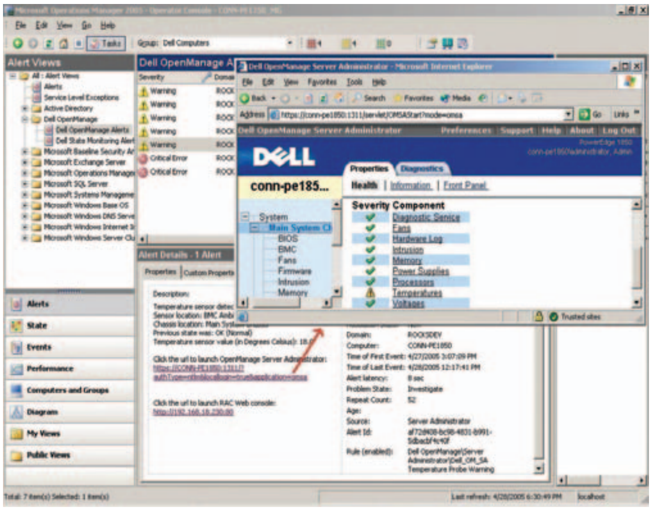

Address: https://conn-pe1850:1311/servlet/OMSAStart?mode=omsa

Description:
Temperature sensor detec
Sensor location: BMC Ambi
Chassis location: Main Syst
Previous state was: OK (Normal)
Temperature sensor value (in Degrees Celsius): 18.0

Click the url to launch OpenManage Server Administrator:
https://CONN-PE1850:1311/?authType=ntlm&localogin=true&application=omsa

Click the url to launch RAC Web console:
http://192.168.18.230:80

Domain:	ROCKSDEV
Computer:	CONN-PE1850
Time of First Event:	4/27/2005 3:07:09 PM
Time of Last Event:	4/28/2005 12:17:41 PM
Alert latency:	8 sec
Problem State:	Investigate
Repeat Count:	52
Age:	
Source:	Server Administrator
Alert Id:	af72d408-bc98-4831-8991-5dbacbf4c40f
Rule (enabled):	Dell OpenManage\Server Administrator\Dell_OM_SA Temperature Probe Warning

Total: 7 item(s) Selected: 1 item(s)

Last refresh: 4/28/2005 6:30:49 PM localhost

Figure 8. Launching Dell OpenManage Server Administrator from MOM 2005

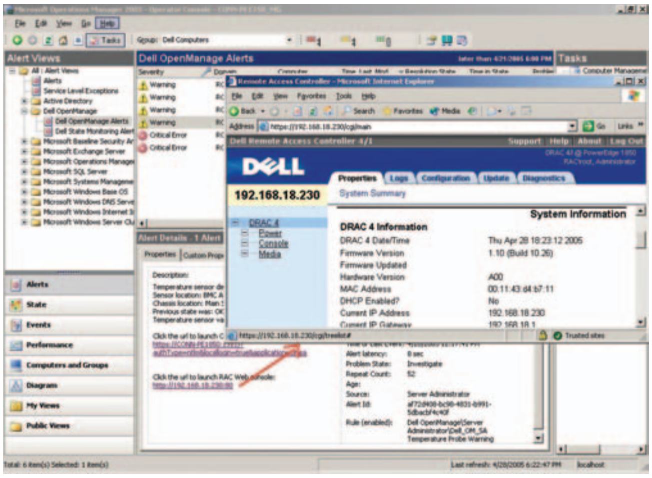

Address: https://192.168.18.230/cgi/main

192.168.18.230

System Information

DRAC 4 Information

DRAC 4 Date/Time Thu Apr 28 18:23:12 2005
Firmware Version 1.10 (Build 10.26)
Firmware Updated
Hardware Version A00
MAC Address 00:11:43:d4:b7:11
DHCP Enabled? No
Current IP Address 192.168.18.230
Current IP Gateway 192.168.18.1

https://192.168.18.230/cgi/treelist#

Description:
Temperature sensor de
Sensor location: BMC A
Chassis location: Main S
Previous state was: OK
Temperature sensor va

Click the url to launch C
https://CONN-PE1850-2W2/
authType=ntlm&login=true&application=...

Click the url to launch RAC Web console:
http://192.168.18.230:80

Alert latency: 0 sec
Problem State: Investigate
Repeat Count: 52
Age:
Source: Server Administrator
Alert Id: af72d408-bc98-4831-b991-5dbacbf4c40f
Rule (enabled): Dell OpenManage\Server Administrator\Dell_OM_SA Temperature Probe Warning

Total: 6 item(s) Selected: 1 item(s)

Last refresh: 4/28/2005 6:22:47 PM localhost

Figure 9. Launching the DRAC Web interface from MOM 2005

For a DRAC launch point to be present, a DRAC card and the DRAC agent must be installed and running on all the Dell servers that are managed by MOM. If a DRAC is installed, administrators must install the DRAC agent using the Dell OpenManage Systems Management CD for OMSA versions 1.6 to 1.9 or the Dell PowerEdge Installation and Server Management CD for OMSA versions 2.0 and later. During installation, administrators should select Managed Node > DRAC Agent.

> The Dell Management Pack can help enterprise IT organizations using MOM 2005 enhance the monitoring and management of their Dell systems.

Each Dell event-processing rule is associated with the VBScript-based Dell RAC Console Launch custom script. This script queries Dell managed systems for the remote access card URL and appends the URL to the alert description. Administrators can click the URL to launch the DRAC Web interface, then log in to the DRAC console and remotely manage the server.[2]

Sending notifications to the hardware support group

For all critical and warning alerts generated by Dell OpenManage Server Administrator (including OMSA Storage Management) and Array Manager events, a notification message is sent to the hardware support notification group. Information events are not processed and hence the hardware support group is not notified about the informational events that are issued from Dell servers. System administrators must be added as operators to this hardware support group to be notified by either e-mail or pager whenever a critical or warning alert is generated.

Customizing the Dell Management Pack

The Dell MP can be customized to suit a managed environment. This section describes two example scenarios for customizing the Dell MP: creating a Dell OpenManage task to launch the OMSA Web interface and customizing the Dell state view to display the status of MOM Agents and the MOM Server.

Scenario 1: Creating a Dell OpenManage task

Administrators can easily launch the OMSA Web interface by directly launching it as a task. To do so, they must create a new task under Dell OpenManage:

1. Open the MOM Administrator console.
2. Go to Management Packs > Tasks > Dell OpenManage.
3. Right-click on "Dell OpenManage" and select "Create Task" from the pop-up menu. The Create Task wizard should launch.
4. After the welcome screen of the Create Task wizard appears, click the Next button.
5. Select "Operator Console" as the "Run" location and "Command-line" as the task type. Click the Next button.
6. Select "Computers" as the view type (this enables the task to be executed when any computer is selected from the MOM Operator console).
7. Under "Task command line," enter the binary path of the active Web browser installed on the system. For example, if

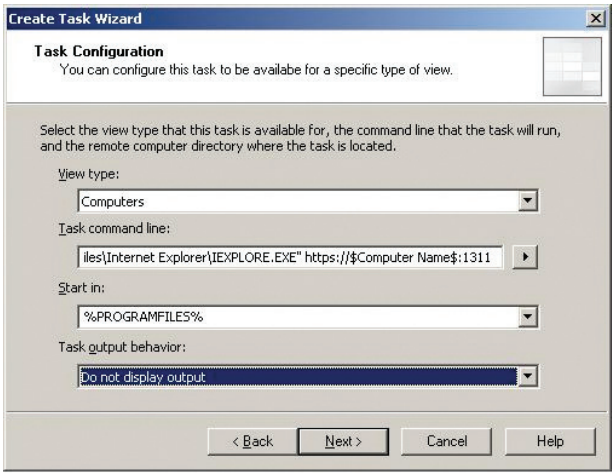

Task Configuration

You can configure this task to be availabe for a specific type of view.

Select the view type that this task is available for, the command line that the task will run, and the remote computer directory where the task is located.

View type: Computers

Task command line: iles\Internet Explorer\IEXPLORE.EXE" https://$Computer Name$:1311

Start in: %PROGRAMFILES%

Task output behavior: Do not display output

Figure 10. Creating a task to launch Dell OpenManage Server Administrator

[2] For more information about remotely managing servers using Dell Remote Access Controllers, refer to the DRAC user guides available at support.dell.com/support/edocs/software/smdrac3.

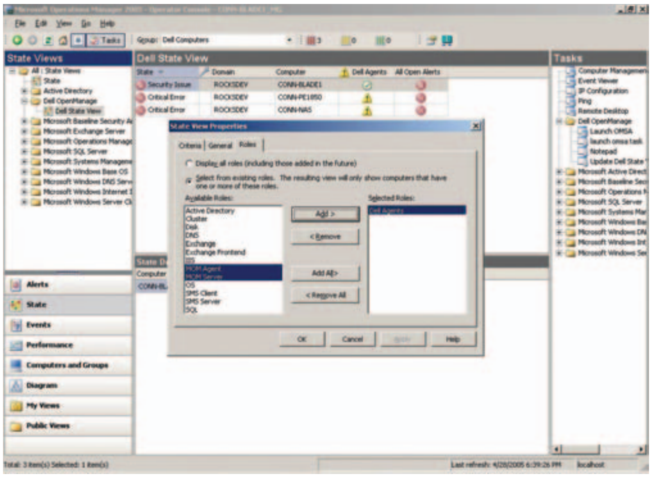

State	Domain	Computer	Dell Agents	All Open Alerts
Security Issue	ROCKSDEV	CONN-BLADE1		
Critical Error	ROCKSDEV	CONN-PE1850		
Critical Error	ROCKSDEV	CONN-NAS		

○ Display all roles (including those added in the future)

● Select from existing roles. The resulting view will only show computers that have one or more of these roles.

Available Roles: Active Directory, Cluster, Disk, DNS, Exchange, Exchange Frontend, IIS, MOM Agent, MOM Server, OS, SMS Client, SMS Server, SQL

Selected Roles: Dell Agents

Figure 11. Customizing the Dell state view

Microsoft Internet Explorer is the active Web browser, enter "C:\Program Files\Internet Explorer\IEXPLORE.EXE" in this field. As shown in Figure 10, enter the following information as a parameter to the Web browser in this field: https://$Computer Name$:1311. The limitation for this task is that all servers managed by the Dell MP must have the OMSA HTTP port number of 1311 (this is the default port number for OMSA). If the port number is changed for any server, the Web interface cannot be launched from this task. However, the updated URL is still reflected in the Dell alert description.

8. Select "%PROGRAMFILES%" as the "Start in" location.
9. Select "Do not display output" for the task output behavior and click the Next button.
10. Enter the name for the task as "Launch OMSA" and provide a description for the task.
11. Optionally, configure a shortcut key to execute this task directly from the keyboard without using the mouse. Select CTRL + SHIFT and enter "O" as the shortcut key to launch the OMSA console.
12. Click the Finish button.

To execute this task, administrators should close the MOM Operator console (if opened) and reopen it. The far right pane shows all the tasks. Administrators can expand "Dell OpenManage" to see the Launch OMSA task that was just created. Next, they should select a Dell server from any of the Dell-specific views (alerts, events, diagram, state, or computers), and then use the shortcut key entered for the task (CTRL + SHIFT + O) or click on the Launch OMSA task. Either action opens up the OMSA Web interface for the selected server.

Scenario 2: Customizing the Dell state view

The Dell state view can be customized to add new roles such as MOM Agents and the MOM Server. This customization can help administrators to monitor the status of these newly added roles. To add the MOM Agent and MOM Server roles into the Dell state view, administrators can take the following steps:

1. Open the MOM Operator console.
2. Select the State tab in the bottom left pane.
3. In the State Views tree, go to Dell OpenManage > Dell State View.
4. Right-click on "Dell State View" and select "Properties" from the pop-up menu.
5. Select the Roles tab in the State View Properties window.
6. Select the roles ("MOM Agent" and "MOM Server") to be monitored from the Available Roles list box and click the Add button (see Figure 11).
7. Click the Apply button and then click the OK button to close the properties dialog box. The Dell state view should then have the new roles added to it.

Enhancing MOM 2005 capabilities for managing Dell systems

By supporting Dell-specific software and hardware components within the managed environment, the Dell Management Pack for Microsoft Operations Manager is designed to extend the capabilities of MOM 2005. The Dell MP can help enterprise IT organizations using MOM 2005 enhance the monitoring and management of their Dell systems. Furthermore, IT administrators can customize the Dell MP to suit a data center's specific needs.

Balasubramaniam J. is a senior engineering analyst at Dell. With more than five years of experience in systems management applications, Bala currently works as a lead engineer for Dell OpenManage Connection. He has a bachelor's degree in Computer Science and Engineering from Madras University in India and a master's degree in Software Systems from the Birla Institute of Technology in Pilani, India.

FOR MORE INFORMATION

Dell OpenManage:
www.dell.com/openmanage

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

Managing Dell Client Hardware Using the

Dell Client Manager from Altiris

As a component of both Dell™ OpenManage™ Client Administrator 3.0 and Altiris® Client Management Suite™ software, the Dell Client Manager from Altiris is designed to provide effective management capabilities for enterprise IT organizations using Dell client hardware. The Dell Client Manager can enable automated, centralized management of Dell Precision™ workstations, OptiPlex™ desktops, and Latitude™ notebooks.

BY JORDAN GARDNER AND TODD MITCHELL

Related Categories:

Altiris

Dell OpenManage

Systems management

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

Dell OpenManage Client Administrator (OMCA) is a suite of integrated products based on Altiris technology that enables IT administrators to manage Dell Precision workstations, Dell OptiPlex desktop computers, and Dell Latitude notebooks. With OMCA, administrators can create, edit, copy, and remotely deploy system images as well as remotely distribute software applications, updates, and service packs over a network. OMCA also enables administrators to remotely migrate old PC software and settings to a new OS as part of a hardware refresh—all from a central management console.

In addition to deployment and migration features, Dell OMCA offers tools for system inventory and software delivery. Specifically, version 3.0 of OMCA includes the following Altiris products:

- **Altiris Deployment Solution:** Deploys operating systems via imaging or scripted installations, moves user data and settings between systems, performs remote control operations, and provides software packaging functionality
- **Altiris Software Delivery Solution:** Provides policy-based software delivery over a wide area network or LAN
- **Altiris Inventory Solution:** Performs robust hardware, software, and user inventory
- **Dell Client Manager:** Includes Dell-specific hardware and BIOS inventory information, helps configure and update BIOS, and provides hardware health monitoring

Altiris Deployment Solution™, Altiris Software Delivery Solution™, and Altiris Inventory Solution® software is designed to manage any standards-based x86 hardware, not just Dell hardware. In addition, the Dell Client Manager[1] from Altiris provides advanced management functions specifically for Dell systems. In addition to being a component of OMCA 3.0, the Dell Client Manager can be added to Altiris Client Management Suite—extending the Altiris service-oriented architecture (SOA) to include Dell-specific hardware properties.

[1] To download a trial version of the Dell Client Manager, visit www.altiris.com/eval/dell and select "Dell OpenManage Client Administrator 3.0" from the drop-down menu.

Altiris Client Management Suite includes the capabilities of OMCA 3.0 and adds integrated tools for patching Microsoft® Windows® operating systems, application metering (usage monitoring and denial of unauthorized applications), advanced remote control, and application management. While OMCA addresses fundamental client management needs, Client Management Suite offers a comprehensive approach by including additional tools and advanced, policy-based options to fully automate desktop management.

Whether as part of Dell OMCA 3.0 or Altiris Client Management Suite, the Dell Client Manager extends either product to include options for managing Dell systems at the hardware level. This article examines the architecture and key functions of the Dell Client Manager.

Building upon the Altiris architecture

The Dell Client Manager employs the same architecture used by all Altiris software tools. This architecture is designed to simplify management tasks by integrating installed Altiris solutions into a central console that allows data and management functions to be shared across individual solutions (see Figure 1).

This architecture enables multiple Altiris solutions to work together. For example, Altiris Server Management Suite™ software, OMCA 3.0, and Altiris Handheld Management Suite™ software can be installed on the same back-end Altiris server to provide a single management console across an entire enterprise for server, desktop, and mobile device management. The Altiris role-and-scope security engine also works across installed Altiris tools. By leveraging the same infrastructure—policy engine, role/scope security engine, and client/server communication model—Altiris solutions are designed to provide efficiency and scalability.

Altiris Notification Server

The key component of the Altiris architecture is Altiris Notification Server™. Notification Server is available at no cost and can be installed independently of any Altiris solutions. It is the engine that manages communication with the remote Altiris agents and the SQL database. The database can be collocated on the Altiris server or installed on a remote server. Notification Server manages the Altiris Web console, Altiris connectors into third-party products, and system notification policies. For more information about Altiris Notification Server, see the "Altiris Notification Server communication architecture" sidebar in this article.

Altiris solutions built on the Notification Server framework use the same Altiris Agent. When additional Altiris solutions are installed, the disk and memory footprints for this agent are dynamically extended with sub-agents as needed. The typical disk footprint for the Altiris Agent is approximately 5 MB to 12 MB, depending on which solutions are installed and how many policies the Altiris administrator has created.

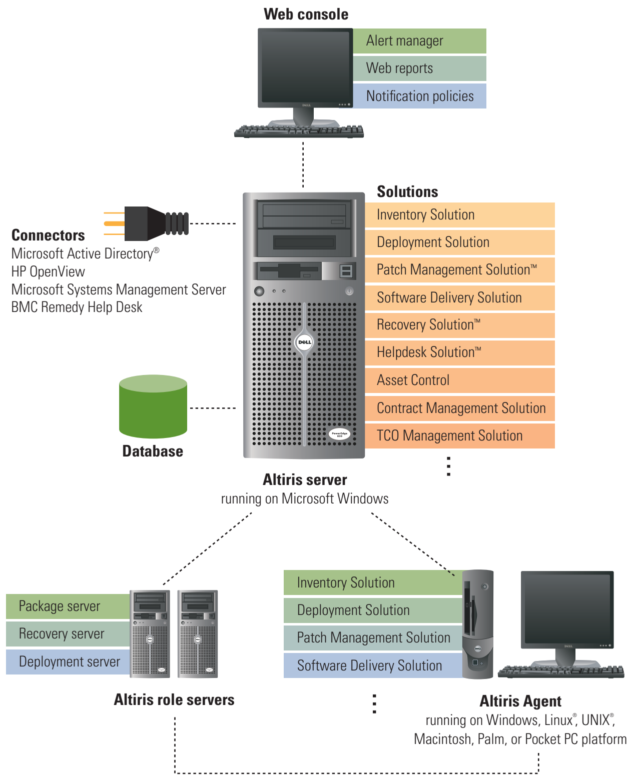

Figure 1. Typical configuration for IT environment managed by Altiris software

Fundamentally, Altiris Notification Server uses policies to associate tasks and software packages with collections of systems. Computers can belong to multiple collections, and collections can be either static or dynamic. When an administrator puts a system in a static collection, it stays in that collection until the administrator explicitly removes it. Membership in a dynamic collection is based on the properties of the system—as properties change, the system automatically moves into and out of dynamic collections. For example, an administrator can create a dynamic collection that consists of desktops running the Microsoft Windows XP OS with Service Pack 2 and belonging to a specific domain. If either of these properties changes for any system in the collection, Notification Server automatically removes the system from that collection, and the system is then disassociated with any policies assigned to manage that collection.

Dynamic collections can be a powerful mechanism for automating systems management. For example, a policy can be created to deliver, on an ongoing basis, a specific Dell BIOS update to any system that requires it. If a system is added to the LAN after the policy is created, that system joins all the predefined collections it qualifies for shortly after the Altiris Agent is installed on the system. Any policies assigned to those collections become effective for the added system—it automatically receives the BIOS update, if needed, and executes any other tasks associated with the assigned policies.

Dell management function	Required Dell OpenManage tool	Included in the Dell Client Manager?
Hardware inventory	ITA/OMCI	Yes
BIOS settings inventory	DCCU or ITA/OMCI	Yes
Policy-based, remote BIOS configuration	None	Yes
BIOS updates	DCCU, ITA/OMCI, or OMCC/OMCI	Yes
Hardware health monitoring	ITA/OMCI	Yes

Figure 2. Dell management functions included in the Dell Client Manager

Dell Client Manager sub-agent

The Dell Client Manager sub-agent contains the same code that is available from Dell as Dell OpenManage Client Instrumentation (OMCI), which is commonly installed on any Dell system managed by Dell OpenManage Client Connector (OMCC) or Dell OpenManage IT Assistant (ITA). As part of its functionality, OMCI publishes hardware-specific values in the Dell Windows Management Instrumentation (WMI) namespace, allowing for detailed hardware and BIOS inventory to be taken from each client. When the Dell Client Manager agent is deployed on client systems, it automatically uninstalls any 6.*x* versions of OMCI before installing the latest OMCI version. It also upgrades any 7.*x* versions of OMCI to the latest version of OMCI as needed. If the Dell Client Manager sub-agent is uninstalled, administrators have the option of allowing the OMCI WMI component to remain on the client system.

Another Dell OpenManage tool, the Dell Client Configuration Utility (DCCU), can be used to remotely retrieve or configure BIOS settings and update the BIOS version. The Dell Client Manager sub-agent includes code from DCCU to provide these functions.

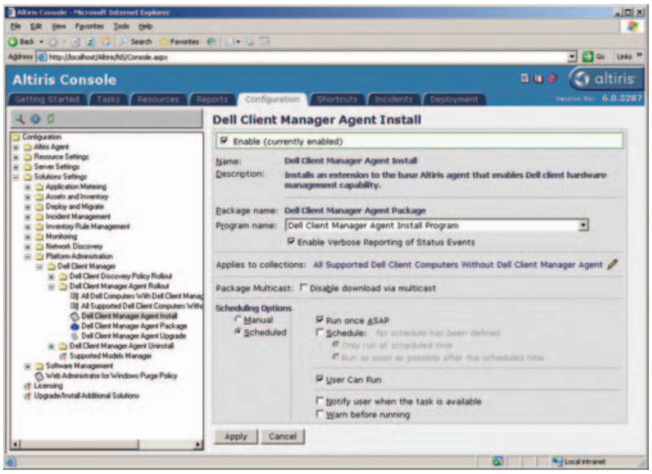
Altiris Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://localhost/Altiris/NS/Console.aspx
Go Links
Altiris Console
altiris
Getting Started
Tasks
Resources
Reports
Configuration
Shortcuts
Incidents
Deployment
Version No: 6.0.5287
Configuration
Altiris Agent
Resource Settings
Server Settings
Solutions Settings
Application Metering
Assets and Inventory
Deploy and Migrate
Incident Management
Inventory Rule Management
Monitoring
Network Discovery
Platform Administration
Dell Client Manager
Dell Client Discovery Policy Rollout
Dell Client Manager Agent Rollout
All Dell Computers With Dell Client Manag
All Supported Dell Client Computers With
Dell Client Manager Agent Install
Dell Client Manager Agent Package
Dell Client Manager Agent Upgrade
Dell Client Manager Agent Uninstall
Supported Models Manager
Software Management
Web Administrator for Windows Purge Policy
Licensing
Upgrade/Install Additional Solutions
Apply
Cancel
Local intranet

Dell Client Manager Agent Install

☑ Enable (currently enabled)

Name: Dell Client Manager Agent Install
Description: Installs an extension to the base Altiris agent that enables Dell client hardware management capability.

Package name: Dell Client Manager Agent Package
Program name: Dell Client Manager Agent Install Program
☑ Enable Verbose Reporting of Status Events

Applies to collections: All Supported Dell Client Computers Without Dell Client Manager Agent

Package Multicast: ☐ Disable download via multicast

Scheduling Options
○ Manual
● Scheduled
☑ Run once ASAP
☐ Schedule: No schedule has been defined
● Only run at scheduled time
○ Run as soon as possible after the scheduled time
☑ User Can Run
☐ Notify user when the task is available
☐ Warn before running

Figure 3. Distributing the Dell Client Manager sub-agent by enabling a single policy

When the Altiris Agent is installed, necessary code from the OMCI and DCCU tools are included with the Dell Client Manager sub-agent. Administrators do not need to download and install additional tools from Dell because the Dell Client Manager provides the functionality required to manage the Dell system (see Figure 2). Enabling a single policy from the Altiris Web console extends the Altiris Agent with the Dell OpenManage code needed to manage Dell hardware (see Figure 3). By using the Dell Client Manager, administrators have a single tool and console to manage their Dell client hardware—avoiding the need for individual Dell OpenManage tools such as ITA, DCCU, OMCC, or OMCI (see Figure 4).

Exploring key features of the Dell Client Manager

The Dell Client Manager provides key functions unique to Dell OpenManage software, including many enhancements over previous Dell systems management software releases. In particular, the Dell Client Manager enables detailed Dell-specific hardware inventory integrated into the Altiris Resource Manager; remote BIOS updates; remote, policy-based BIOS configuration (model agnostic); and remote hardware health monitoring.

Detailed Dell-specific hardware and BIOS inventory

The Dell Client Manager is designed to provide detailed hardware and BIOS inventory scans on Dell hardware. Dell hardware and BIOS properties are read from the Dell WMI namespace, and then published once the Dell Client Manager sub-agent is present. This inventory is forwarded to the Altiris server using the client/server communication model (for more information about this process, see the "Altiris Notification Server communication architecture" sidebar in this article). The inventory gathered through the Dell Client Manager sub-agent is displayed in data classes identified with an "OMCA" prefix in the Altiris Resource Manager (see Figure 5). These data classes are displayed alongside other Altiris inventory data classes, allowing a single view of all information known for a particular Dell device.

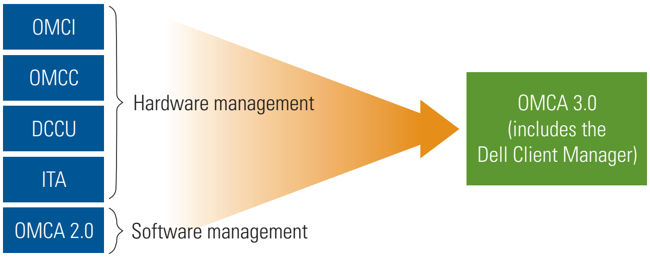

Figure 4. OMCA 3.0 provides capabilities of various Dell tools

Now you can be in two places at the same time. Or, if you're feeling ambitious, 200.

From a single interface you can remotely manage all your company's client hardware and software – with Dell™ OpenManage™ Client Administrator 3.0 powered by Altiris.® Systems administrators have to work hard to stay on top of things. Now there's a whole new way of working smart, thanks to a single solution that lets you automate your entire client infrastructure via remote configuration, management, and monitoring of hardware, O/S, and applications. It's the new Dell OpenManage Client Administrator 3.0 powered by Altiris, and it's the best way to remotely manage hardware and software from a single toolset. ***Find out more today at dell.com/Altiris, and let Dell and Altiris work harder – and smarter – for you.***

altiris®

The power of control.

DELL

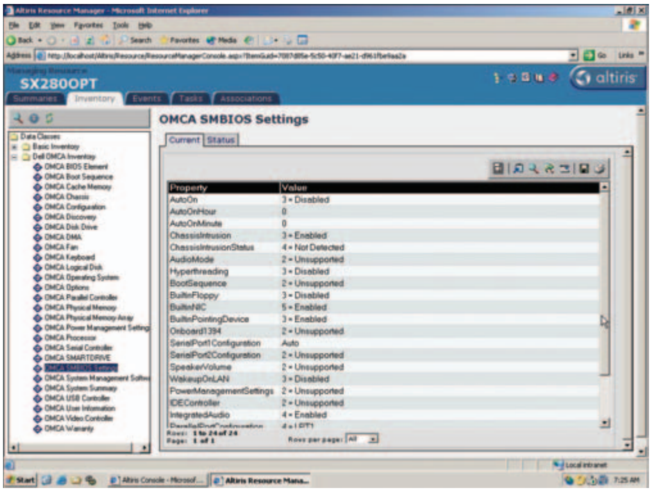

OMCA SMBIOS Settings

Property	Value
AutoOn	3 = Disabled
AutoOnHour	0
AutoOnMinute	0
ChassisIntrusion	3 = Enabled
ChassisIntrusionStatus	4 = Not Detected
AudioMode	2 = Unsupported
Hyperthreading	3 = Disabled
BootSequence	2 = Unsupported
BuiltinFloppy	3 = Disabled
BuiltinNIC	5 = Enabled
BuiltinPointingDevice	3 = Enabled
Onboard1394	2 = Unsupported
SerialPort1Configuration	Auto
SerialPort2Configuration	2 = Unsupported
SpeakerVolume	2 = Unsupported
WakeupOnLAN	3 = Disabled
PowerManagementSettings	2 = Unsupported
IDEController	2 = Unsupported
IntegratedAudio	4 = Enabled
ParallelPortConfiguration	4 = LPT1

Rows: 1 to 24 of 24
Page: 1 of 1

Figure 5. Dell OpenManage inventory displayed in the Altiris Resource Manager

The Altiris Extensible Management Architecture™ (EMA™) allows Dell-specific properties to be published to other components and solutions installed on the Altiris server. This data sharing can be used to create comprehensive management policies that automate management functions based on Dell hardware properties such as BIOS revisions and model numbers. The Altiris architecture allows Dell properties to drive Altiris notification policies, collection definitions, and reports.

Administrators can define the frequency of hardware and BIOS scans as well as the collections of systems for which the scan schedules apply. A collection can be defined using any hardware properties inventoried by Altiris software. Another option is to wake up powered-down systems to perform these scans. These settings are configured via two policies provided with the Dell Client Manager: the BIOS Inventory Scan Policy and the Hardware Inventory Scan Policy. These inventory scans use the same Dell Client Manager sub-agent; however, the policies differ in the WMI properties that they collect.

These policies exist on managed Dell clients as two small .xml files:

- **actions.xml:** Defines which WMI properties should be collected
- **schema.xml:** Formats the output of the scan so that Altiris Notification Server can recognize the data and successfully import it once it is posted to the Altiris server

These .xml files are communicated to the client prior to the first inventory scan and are periodically updated as additional Dell client models are released and additional BIOS properties become available with updated BIOS versions.

Administrators also can use Altiris Notification Server to throttle bandwidth to prohibit BIOS or hardware inventory from being sent over the network unless certain bandwidth thresholds are satisfied. These capabilities are particularly valuable for large environments or infrastructures with low-bandwidth connections.

Remote BIOS updates

The Dell Client Manager allows administrators to easily update the BIOS on Dell client hardware via automated policies. Administrators can download a Dell client BIOS update from support.dell.com to the Altiris server and then use the Altiris Web console to assign the update to a collection of Dell clients (see Figure 6). Most often, BIOS updates are assigned to a collection of systems that are the same Dell model and do not already have the latest BIOS update.

For example, administrators can build a policy to assign a BIOS update to all the Dell Latitude D600 notebooks within the marketing department that do not already have the latest A16 BIOS revision. The policy can define when the update is applied and force a Wake-on-LAN if necessary. If the policy is created as a dynamic collection, when a Latitude D600 system that does not have the latest A16 BIOS revision appears on the LAN, the update is automatically applied—without requiring administrator intervention.

The BIOS update policy uses the .hdr file included in the BIOS package to perform the upgrade. The .hdr file contains the BIOS image and metadata for that image. This package is delivered to the client system using the Altiris Agent's ability to deliver a software package that is communicated via HTTP, HTTP over Secure Sockets Layer (HTTPS), or TCP/IP (all configurable by the administrator).

Once the package has been delivered to the client, the .hdr file is extracted and its metadata is used to perform required checks. These checks can include verifying a valid system ID (to help ensure that the BIOS update is a valid update for the target system) and whether a downgrade is permitted (if desired). Once all the checks have passed, the BIOS image is taken from the .hdr file and copied

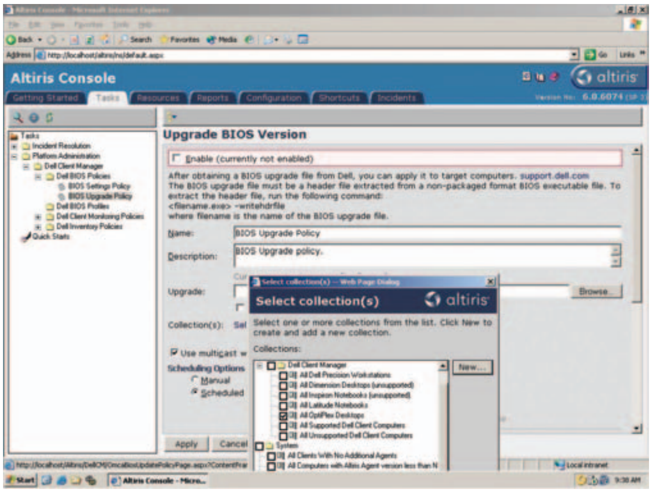
Altiris Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address http://localhost/altiris/ns/default.aspx
Go Links
Altiris Console
altiris
Getting Started Tasks Resources Reports Configuration Shortcuts Incidents
Version No: 6.0.6074 (SP 3)
Tasks
Incident Resolution
Platform Administration
Dell Client Manager
Dell BIOS Policies
BIOS Settings Policy
BIOS Upgrade Policy
Dell BIOS Profiles
Dell Client Monitoring Policies
Dell Inventory Policies
Quick Starts
Browse...
Apply Cancel
Select collection(s) -- Web Page Dialog
Select collection(s)
altiris
New...
http://localhost/Altiris/DellCM/OmcaBiosUpdatePolicyPage.aspx?ContentFrar
Local intranet
Start Altiris Console - Micro... 9:30 AM

Upgrade BIOS Version

☐ Enable (currently not enabled)

After obtaining a BIOS upgrade file from Dell, you can apply it to target computers. support.dell.com
The BIOS upgrade file must be a header file extracted from a non-packaged format BIOS executable file. To extract the header file, run the following command:
<filename.exe> -writehdrfile
where filename is the name of the BIOS upgrade file.

Name: BIOS Upgrade Policy
Description: BIOS Upgrade policy.
Upgrade:
Collection(s):
☑ Use multicast
Scheduling Options
○ Manual
● Scheduled

Select one or more collections from the list. Click New to create and add a new collection.

Collections:
☐ Dell Client Manager
☐ All Dell Precision Workstations
☐ All Dimension Desktops (unsupported)
☐ All Inspiron Notebooks (unsupported)
☐ All Latitude Notebooks
☑ All OptiPlex Desktops
☐ All Supported Dell Client Computers
☐ All Unsupported Dell Client Computers
☐ System
☐ All Clients With No Additional Agents
☐ All Computers with Altiris Agent version less than N

Figure 6. Creating a BIOS upgrade policy with the Dell Client Manager

to a locked, contiguous memory buffer. In the client's shutdown process, the BIOS detects a valid BIOS update image in the system memory buffer and uses the image data to reprogram the system's flash memory. When the client powers up, the updated BIOS takes effect. A primary benefit of the Dell Client Manager is the ability to configure the installation schedule, reboot time, and deferment of BIOS update packages.

Remote configuration of BIOS settings

Dell has traditionally provided BIOS updates as part of the Dell Client Configuration Utility, which uses BIOS setup packages. Typically, one DCCU package is created for each Dell model to be configured, and administrators use a third-party tool to deliver the DCCU packages to the appropriate Dell systems.

The Dell Client Manager is designed to provide several advantages over DCCU by enabling the following features:

- Capability to create model-independent BIOS profiles
- Enforcement of remote BIOS configurations via policies

New profiles can be created from the Altiris Web console (profiles created this way can support all possible Dell client BIOS settings), or they can be captured from a reference system and then modified. BIOS profiles are defined in the central Altiris Web console and saved as templates that can be applied over and over again. Altiris administrators can add, edit, clone, and delete BIOS profiles (see Figure 7).

Policies can be created to apply BIOS profiles across any supported Dell client hardware regardless of the model. Settings that are not available in the BIOSs of legacy systems are simply ignored. Each managed Dell client with the Dell Client Manager sub-agent installed applies the BIOS profile settings by reading from the BIOS profile configuration file. This configuration file defines the values to be applied to the writeable BIOS WMI properties published by the Dell Client Manager sub-agent.

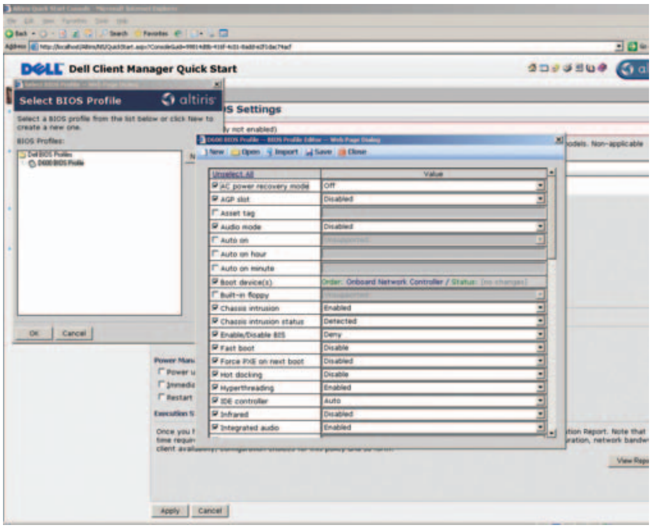
Dell Client Manager Quick Start

Select BIOS Profile

Select a BIOS profile from the list below or click New to create a new one.

BIOS Profiles:

Dell BIOS Profiles
D600 BIOS Profile

Unselect All	Value
☑ AC power recovery mode	Off
☑ AGP slot	Disabled
☐ Asset tag	
☑ Audio mode	Disabled
☐ Auto on	Unsupported
☐ Auto on hour	
☐ Auto on minute	
☑ Boot device(s)	Order: Onboard Network Controller / Status: [no changes]
☐ Built-in floppy	Unsupported
☑ Chassis intrusion	Enabled
☑ Chassis intrusion status	Detected
☑ Enable/Disable BIS	Deny
☑ Fast boot	Disable
☑ Force PXE on next boot	Disabled
☑ Hot docking	Disable
☑ Hyperthreading	Enabled
☑ IDE controller	Auto
☑ Infrared	Disabled
☑ Integrated audio	Enabled

Figure 7. Using the Dell Client Manager to obtain BIOS setting profiles

Remote hardware health and BIOS configuration monitoring

The Dell Client Manager enables administrators to monitor hardware components of Dell client systems. Several properties can be monitored, such as low disk space or chassis intrusion. Hardware health monitoring can be used in conjunction with Altiris notification policies and event handlers to quickly inform administrators of problem conditions and potential system failure.

ALTIRIS NOTIFICATION SERVER COMMUNICATION ARCHITECTURE

Altiris Notification Server incorporates a pull agent model. By default, the Altiris Agent requests a policy update from Notification Server every hour. In production environments, this interval is typically extended to be every 6 to 12 hours, although it may be longer or shorter. During this update, the Altiris Agent asks Notification Server to determine which new policies are applicable to the system hosting the agent.

If at least one policy addition or modification exists, the agent downloads a new policy configuration XML file, which informs the Altiris Agent of what work it should perform. For example, the file may inform the agent that it should run a software inventory scan every 12 hours and a hardware inventory scan once a week; it should deny access to any unauthorized software programs (such as games or instant messengers) during the working hours of 8 A.M. to 5 P.M.; and it should download and execute the most recent Microsoft patches or Dell BIOS updates from the Altiris server immediately. The entire process of downloading the configuration policy typically generates less than 2 KB of traffic per agent.

Communication between the Altiris Agent and Altiris Notification Server fundamentally consists of XML files that are compressed and transferred via HTTP (port 80) or HTTPS (port 443). As inventory scans are performed or other agent events are triggered, that data is communicated to Notification Server as an XML file with an *.nse extension (Notification Server Event).

The Dell Client System Monitoring Policy allows administrators to select which events and metrics should be monitored. This policy can also define the corresponding actions that should be triggered whenever those preset events or warnings occur (see Figure 8).

Dell Client Manager events are generated two ways:

- **Using WMI event consumers:** Whenever a change occurs to certain properties within the Dell namespace, registered WMI event consumers (contained in the Dell Client Manager sub-agent) pick up these changes and trigger an event. These events are read by the Altiris Agent and their information is forwarded to Altiris Notification Server for potential follow-up action.
- **Monitoring the status of WMI properties:** Selected property values are polled regularly, and change information is sent to the Altiris server as soon as it is detected. No registered WMI event consumers are used.

The Altiris server allows automated actions for responding to monitored events. These automated actions can include the following:

- Logging each event in the NT event log
- Sending notification e-mails to system administrators
- Creating or editing help-desk tickets
- Launching any custom script or executable on the Altiris server

Support for future Dell hardware models

The Dell Client Manager is extensible in that it is designed to support Dell BIOS updates and system models as they become available. Through the Supported Models Manager, administrators can view which models are currently supported, or they can update the Dell Client Manager to include additional Dell client models as they are released (see Figure 9).

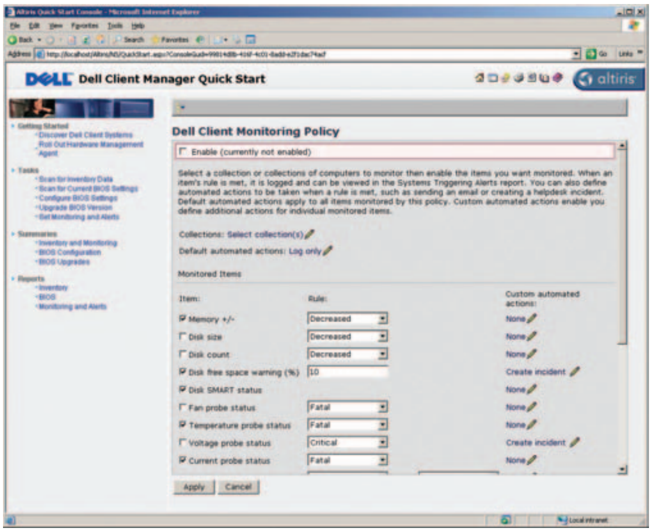
Altiris Quick Start Console - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address http://localhost/Altiris/NS/QuickStart.aspx?ConsoleGuid=99814d8b-416f-4c01-8add-e2f1dac74af
DELL Dell Client Manager Quick Start
altiris
Getting Started
Discover Dell Client Systems
Roll Out Hardware Management Agent
Tasks
Scan for Inventory Data
Scan for Current BIOS Settings
Configure BIOS Settings
Upgrade BIOS Version
Set Monitoring and Alerts
Summaries
Inventory and Monitoring
BIOS Configuration
BIOS Upgrades
Reports
Inventory
BIOS
Monitoring and Alerts
Apply
Cancel
Local intranet

Dell Client Monitoring Policy

□ Enable (currently not enabled)

Select a collection or collections of computers to monitor then enable the items you want monitored. When an item's rule is met, it is logged and can be viewed in the Systems Triggering Alerts report. You can also define automated actions to be taken when a rule is met, such as sending an email or creating a helpdesk incident. Default automated actions apply to all items monitored by this policy. Custom automated actions enable you define additional actions for individual monitored items.

Collections: Select collection(s)

Default automated actions: Log only

Monitored Items

Item:	Rule:	Custom automated actions:
☑ Memory +/-	Decreased	None
□ Disk size	Decreased	None
□ Disk count	Decreased	None
☑ Disk free space warning (%)	10	Create incident
☑ Disk SMART status		None
□ Fan probe status	Fatal	None
☑ Temperature probe status	Fatal	None
□ Voltage probe status	Critical	Create incident
☑ Current probe status	Fatal	None

Figure 8. Configuring actions corresponding to specific events for monitored clients

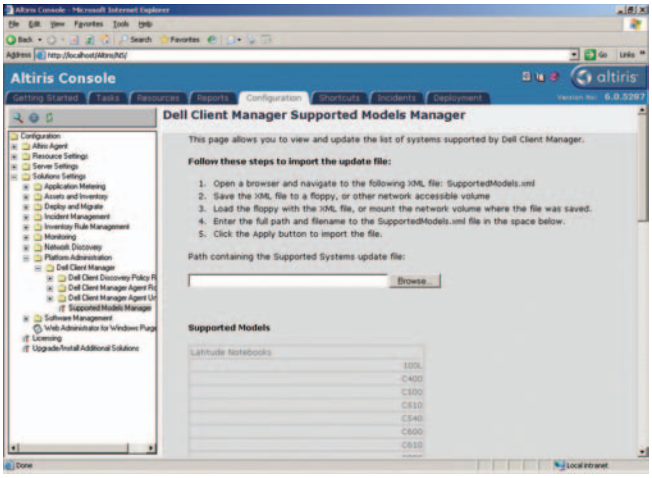

Dell Client Manager Supported Models Manager

This page allows you to view and update the list of systems supported by Dell Client Manager.

Follow these steps to import the update file:

1. Open a browser and navigate to the following XML file: SupportedModels.xml
2. Save the XML file to a floppy, or other network accessible volume
3. Load the floppy with the XML file, or mount the network volume where the file was saved.
4. Enter the full path and filename to the SupportedModels.xml file in the space below.
5. Click the Apply button to import the file.

Path containing the Supported Systems update file:

Supported Models

Latitude Notebooks
100L
C400
C500
C510
C540
C600
C610

Figure 9. Supported Models Manager within the Dell Client Manager

Managing Dell clients from a single console

The Dell Client Manager from Altiris offers a simple, robust tool for one-to-many management of Dell workstations, desktops, and notebooks. It bundles Dell OpenManage code and is designed to improve upon individual Dell OpenManage tools, such as DCCU, OMCI, OMCC, and ITA. A key advantage of the Dell Client Manager is its ability to bring Dell-specific functions and hardware properties into the comprehensive Altiris SOA. This approach enables a level of automation, control, and extensibility previously unavailable for Dell client hardware. The Dell Client Manager allows enterprise IT organizations to benefit from the capabilities of Dell OpenManage software while using the Altiris infrastructure—all from a single, easy-to-use console.

Jordan Gardner is a technical strategist on the Dell alliance team at Altiris. He works closely with Dell and Altiris customers to understand the Dell-specific functionality that Altiris solutions provide. Jordan has a bachelor's degree in Computer Science from Brigham Young University.

Todd Mitchell is the Dell alliance technical director at Altiris. He has worked with numerous Altiris customers to support Dell-specific implementations and management needs. Todd has a bachelor's degree from Brigham Young University.

FOR MORE INFORMATION

Dell Client Manager:
www.altiris.com/omca

Managing UNIX and Linux Platforms in a Windows World

Altiris® systems management software can be used to manage heterogeneous IT environments. Servers and clients running UNIX®, Linux®, and Microsoft® Windows® operating systems can be managed from a central Altiris console. Such multiplatform support allows enterprise IT organizations to gradually integrate UNIX/Linux-based systems into Windows-centric environments.

BY PATRICK BOURKE, TODD MITCHELL, AND RICH LACEY

Related Categories:

Altiris

Linux

Microsoft Windows

Operating system (OS)

Systems management

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

The Linux platform continues to gain significant ground in the IT market. And while more and more large corporate IT organizations are implementing Linux, the small-to-medium business market is faced with an interesting dilemma: Although these organizations may want to deploy Linux platforms, they do not want the expense and added operational burden of another OS.

Regardless of the distribution, Linux is a viable software alternative to Microsoft Windows—and not simply because of its low price. The economic benefits of reduced or absent licensing fees are obvious, but for many organizations, the value goes beyond that: Linux can be a driver for rapid growth, an effective solution for addressing seasonal fluctuations, or even a core component for solving strategic needs. Furthermore, as decision makers continue to feel increasing pressure to justify IT investments with special considerations for security, flexibility, and overall cost, implementing Linux can be easy to justify.

Implementing the Linux platform is not an all-or-nothing strategy. Many IT organizations opt for a gradual migration approach that allows them to begin introducing Linux into current operations while simultaneously improving efficiencies across their existing infrastructure—including Microsoft Windows and UNIX platforms.

Altiris systems management software can help IT organizations manage the existing Windows- and UNIX-based systems they have while expanding the Linux footprint in their data centers. The Altiris console abstracts many of the low-level differences between Windows and UNIX/Linux—enabling administrators to quickly become effective at managing diverse environments on both small and large scales.

Understanding the Altiris platform

Altiris provides a single-console solution with native agents for a variety of operating systems, including Windows, UNIX, Linux, and others. Although implementing a single management tool for heterogeneous environments offers several benefits, UNIX/Linux administrators may be reluctant to use a tool that is not open source or one that interoperates with Windows. However, CIOs and IT directors—who are responsible for implementing technology to generate strategic value for the organizations they serve—very often recognize and endorse the need to leverage efficiencies across disparate environments via a single point of control.

Altiris software is built upon the Extensible Management Architecture™ (EMA™) platform that allows the majority of Altiris products (including those for UNIX/Linux) to plug into a common back-end Altiris server. This framework leverages efficiencies across Altiris products by providing a common-role security engine for access to features and tasks, a Configuration Management Database (CMDB), a Web-based console, a consolidated client/server communication model, and an Altiris Agent that individual solutions extend to add features (as solutions are added into the console, the agent footprint automatically grows to accommodate new features).

Most Altiris software products are organized into suites, but they can also be purchased individually. This extensibility allows organizations to add functions and features as needs and budgets evolve. Altiris software is designed to work seamlessly with other Altiris tools to help future-proof a management tool set.

Altiris component	Supported UNIX and Linux distributions
Altiris Notification Server agent (required for most Altiris solutions)	IBM AIX 4.3.3, 5.1, 5.2, and 5.3
	HP-UX 11, 11i, and 11iv2
	Red Hat Linux 7.2, 7.3, 8, and 9; Red Hat Enterprise Linux AS 2.1; Red Hat Enterprise Linux 3 and 4
	Sun Solaris 7, 8, 9, and 10
	Novell® SUSE® Linux 8.0-8.1 and 9.0-9.3; SUSE Linux Enterprise Server (SLES) 8 and 9
	United Linux 1.0
Altiris Deployment Solution agent	Debian 3.1
	Fedora Core 3
	Red Hat Linux 7.3, 8, and 9; Red Hat Enterprise Linux AS 2.1; Red Hat Enterprise Linux 3 and 4
	Sun Solaris 8 and 9 (management agent only)
	Novell SUSE Linux 9.x; SLES 9
	VMware ESX Server 2.1 and 2.5
Altiris Monitor Solution	Red Hat Linux 7.2, 7.3, 8, and 9; Red Hat Enterprise Linux AS 2.1; Red Hat Enterprise Linux 3 and 4
	Sun Solaris 7, 8, and 9
	Novell SUSE Linux 8.0-8.1 and 9.0; SLES 8
Altiris security and compliance solutions	IBM AIX 4.3.3, 5.1, and 5.2
	HP-UX 11 and 11i
	Red Hat Linux 8 and 9; Red Hat Enterprise Linux 3
	Sun Solaris 8 and 9

Figure 1. Supported platforms for various Altiris software products

Altiris also organizes its software products into a maturity model that suggests increasing value at different points in an administrator's experience with the software. IT administrators that have not used Altiris software can benefit from products in level 1 of a suite. As they implement and learn those tools, then the products in level 2 become appropriate. For example, level 1 of Altiris Server Management Suite™ includes deployment, inventory, and software delivery tools among others. These tools are the foundational blocks that generate manageability gains for most IT organizations. After administrators become familiar with these products, the tools in levels 2 and 3 of Altiris Server Management Suite become relevant and valuable. The suites and levels suggest a proven starting place and a migration path for taking advantage of the Altiris platform.

Installing the Altiris Agent

Altiris software supports a variety of UNIX and Linux distributions (see Figure 1). This heterogeneity offers numerous benefits, including a single point of control, familiarity and usability, minimized licensing costs, and cross-trained administrators adept at managing a variety of operating systems within a single tool set. Another significant benefit administrators can realize by using Altiris software is the flexibility to focus on which applications are suited for critical IT projects as opposed to having to narrow user or project requirements to satisfy the limitations of a homogeneous management solution.

Installing the Altiris Agent is a prerequisite for implementing Altiris tools within the data center. To install the Altiris Agent onto a UNIX/Linux-based system, administrators need to use the root account or an account with root-equivalent privileges. Root access is required to perform many of the administrative tasks managed by the Altiris infrastructure, such as system inventory and software delivery.

The Altiris Agent can be installed using various methods, including pushing the agent remotely from the Altiris console (see Figure 2). When the agent is pushed, Secure Shell (SSH) is initially used to connect to the target computers to start the installation process. Altiris supports SSH versions 1 and 2 (SSH1 and SSH2, respectively) using username and password authentication or public and private key authorization. The push installation process assumes that SSH is properly configured and running on the target machine. Once the SSH connection is established, the implementation forces the selection of encryption algorithms in the following order:

- An SSH2 connection first tries to negotiate Advanced Encryption Standard (AES), and if that fails, Triple Data Encryption Standard (DES) is used. If Triple DES is not available, Blowfish is selected; if that is not available, the connection is dropped. DES is not negotiated because most servers no longer support it (primarily because it is not secure).

- An SSH1 connection first tries to negotiate the Triple DES algorithm, and if that fails, it tries Blowfish. If Blowfish is not available (via a remote server), the connection falls back to DES.
- If SSH fails, the Altiris Agent can be configured to resort to the less-secure Telnet protocol.

To push the agent, administrators can simply select the target host systems from a list of discovered systems in the Altiris console. Altiris software provides several methods for discovering systems on a network such as IP sweeping, TCP port scanning, and circular Domain Name System (DNS) resolution. Administrators also can manually enter the IP addresses of target systems into the agent rollout dialog or import a spreadsheet of previously defined systems. Once the agent is installed, HTTP over Secure Sockets Layer (HTTPS) is the default communication protocol between the Altiris Agent and the Altiris Notification Server.

A significant benefit administrators can realize by using Altiris software is the flexibility to focus on which applications are suited for critical IT projects as opposed to having to narrow user or project requirements to satisfy the limitations of a homogeneous management solution.

In addition to a push installation, the Altiris Agent can be installed by manually pulling it from a central URL or integrating the agent package into the imaging process. Integration into the imaging process helps guarantee that the agent is installed and configured prior to deployment in the environment. To help decrease the management footprint operationally, Altiris software offers several features to help reduce network utilization for agent communication, including checkpoint/restart and bandwidth throttling.

Using Altiris software in UNIX/Linux environments

Altiris offers several software products designed to provide comprehensive life-cycle management for UNIX/Linux-based systems. These products can provide the following capabilities:

- Provisioning
- Inventory
- Software delivery
- Monitoring
- Security auditing and vulnerability scanning
- Network discovery and topology mapping

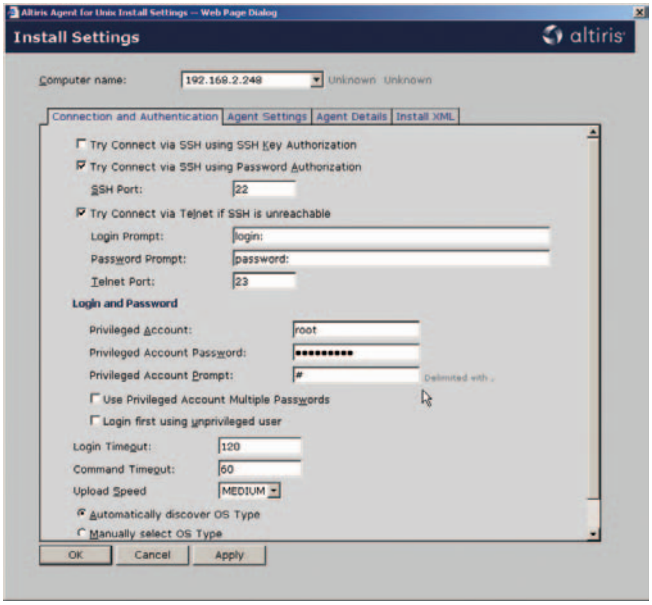
Altiris Agent for Unix Install Settings -- Web Page Dialog
Install Settings
altiris
Connection and Authentication
Agent Settings
Agent Details
Install XML
OK
Cancel
Apply

Computer name: 192.168.2.248 Unknown Unknown

☐ Try Connect via SSH using SSH Key Authorization
☑ Try Connect via SSH using Password Authorization
SSH Port: 22
☑ Try Connect via Telnet if SSH is unreachable
Login Prompt: login:
Password Prompt: password:
Telnet Port: 23

Login and Password
Privileged Account: root
Privileged Account Password: ●●●●●●●●●●
Privileged Account Prompt: # Delimited with ,
☐ Use Privileged Account Multiple Passwords
☐ Login first using unprivileged user
Login Timeout: 120
Command Timeout: 60
Upload Speed MEDIUM
● Automatically discover OS Type
○ Manually select OS Type

Figure 2. Pushing the Altiris Agent to UNIX/Linux-based systems

Altiris Deployment Solution

Altiris Deployment Solution™ software is designed to deploy and manage remote servers, desktops, and notebooks. This solution can create and distribute Linux images (with support for ext2 and ext3 file systems); perform a scripted Linux OS installation (using KickStart); remotely change configuration settings (such as host name and IP address); execute Linux shell scripts (such as bash, sh, csh, ksh, and perl); copy files from the Altiris server to a managed system; and remotely power-control machines. Furthermore, an add-on package provides support specifically for provisioning Dell™ PowerEdge™ servers. With Altiris Deployment Solution for Dell Servers, administrators can perform the following tasks:

- Update hardware components with Dell Update Packages
- Deploy the hidden 32 MB file allocation table (FAT)–formatted Dell Utility Partition
- Configure a server's Dell Remote Access Controller, baseboard management controller, and BIOS
- Set a RAID configuration

For more information about Altiris Deployment Solution for Dell Servers, visit www.dell.com/altiris.

Altiris Deployment Solution includes several sample jobs to provide examples for remotely installing applications such as Apache, Oracle® 10*g*, and VMware® ESX Server™ virtualization software (see www.altiris.com/vmware for more information). This product also supports DOS, Windows Preinstallation Environment (WinPE), and Linux preboot environments—giving administrators

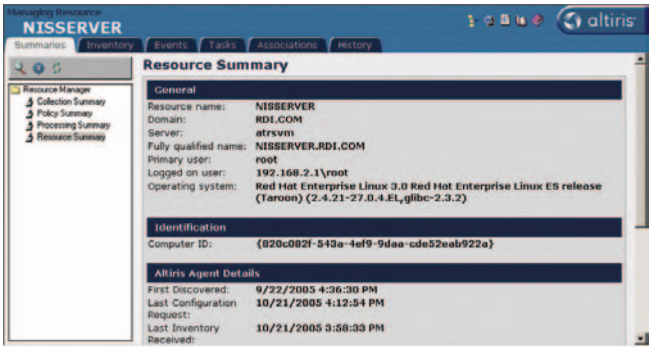

Resource Summary

General

Resource name: NISSERVER
Domain: RDI.COM
Server: atrsvm
Fully qualified name: NISSERVER.RDI.COM
Primary user: root
Logged on user: 192.168.2.1\root
Operating system: Red Hat Enterprise Linux 3.0 Red Hat Enterprise Linux ES release (Taroon) (2.4.21-27.0.4.EL,glibc-2.3.2)

Identification

Computer ID: {020c002f-543a-4ef9-9daa-cde52eab922a}

Altiris Agent Details

First Discovered: 9/22/2005 4:36:30 PM
Last Configuration Request: 10/21/2005 4:12:54 PM
Last Inventory Received: 10/21/2005 3:58:33 PM

Figure 3. Partial Linux inventory from the Altiris Resource Summary

tremendous flexibility to provision target systems. Administrators can even switch between any of these preboot environments as needed within a single deployment job.

Furthermore, Altiris Deployment Solution can deploy the following package file types:

- **.rpm:** Red Hat® Package Manager (RPM™) file
- **.bin:** Binary file
- **.gz:** Compressed file package
- **.tar:** Collection of files in a package
- **.tgz:** Compressed collection of files in a package
- **.bz2:** Compression file
- **.shar:** Tar file with a shell script as a package
- **.deb:** Debian package file
- **.pkg:** Solaris package file

The value of Altiris Deployment Solution is centered in its ability to capture a complete sequence of management tasks (including workflow with conditional logic) as a simple drag-and-drop job in the Altiris console. For example, a single job can be built to provision a server from bare metal (including low-level BIOS and RAID configuration) through OS deployment and application installation. Once a job is built, it can be executed over and over again simply by dragging and dropping it onto the icons representing managed systems in the Altiris console. This functionality can be used to create automated jobs for many of the activities administrators currently perform manually.

Administrators can use Altiris Deployment Solution to help their IT department define, standardize, and automate deployment processes for Linux-based servers. Deployment standardization helps ensure Linux-based configurations are reliable and consistent. Altiris Deployment Solution also can help significantly reduce the administrative time required to provision Linux-based servers. In fact, in a KeyLabs study that was jointly commissioned by Altiris and Dell in November 2004, 25 Dell PowerEdge 2650 servers running Red Hat Enterprise Linux AS 3 were deployed using various installation methods. Deployment using Altiris Deployment Solution was 87 percent faster than a manual deployment.[1] Dell IT also used Altiris Deployment Solution to deploy its own servers and was able to reduce deployment of Windows- and Linux-based servers from an average of 6 hours down to just 20 minutes per server.[2]

Altiris Inventory Solution

Altiris Inventory Solution® software is designed to gather hardware, software, and OS data from each UNIX/Linux-based system based on Altiris policies (see Figure 3). A policy implements a user-defined, recurring schedule for collecting data from target systems. Policies can provide comprehensive data collection or limit inventory scans to a subset of data specified by an administrator.

Administrators can also define custom inventory scans to collect additional data values not gathered by the default Altiris scan. Administrators can create shell scripts to pull data from a variety of sources, store the data in the Altiris CMDB, and display it in the Altiris console just like data returned from a standard scan. This data is also available for Altiris reports, policies, and collections in the same way that data from a default Altiris scan is made available.

The value of Altiris Deployment Solution is centered in its ability to capture a complete sequence of management tasks as a simple drag-and-drop job in the Altiris console.

The first time a system is scanned, all inventory data is aggregated into a single XML file (typically 300 KB in size) on the target system and then forwarded to the central Altiris server via HTTP or HTTPS. To minimize bandwidth, subsequent scans typically send only delta information (that is, data that has changed since the last scan)—typically 15 KB to 25 KB in size.

From the central console, administrators can view the Altiris-provided reports, create their own reports, or build notification policies. Notification policies scan incoming data for predefined values and immediately notify administrators of problem

[1] For a synopsis of the methodology and findings of this study, see "Time-Savings Validation for Dell Server Deployment with Altiris Deployment Solution" by Todd Mitchell and Landon Hale, *Dell Power Solutions,* August 2005; www.dell.com/downloads/global/power/ps3q05-20050221-Altiris.pdf. The complete study documentation can be found at www.dell.com/downloads/global/solutions/Deployment%20Comparison%20for%20Dell%20PowerEdge%20Servers.pdf.

[2] For more information, see the Dell success story at www.dell.com/downloads/global/casestudies/2005_altiris.pdf.

conditions via e-mail, help-desk tickets, reports, or launching of any user-defined action such as a custom script or application. Altiris Inventory Solution can provide a robust view of the IT environment, allowing administrators to accurately forecast growth requirements, track software installations, plan equipment upgrades or replacements, and assist with server consolidation planning and technology migrations.

Altiris policies assign tasks to collections. Collections are groups of managed assets that are either explicitly defined by an administrator or dynamically generated by a query against a set of properties (for example, all servers running Red Hat Enterprise Linux or all servers running Apache software). Administrators can use dynamic collections to automate management functions. As properties of UNIX/Linux-based systems change over time, the Altiris server can automatically move systems into and out of dynamic collections, thereby changing the tasks that apply to them without administrator involvement. If a system is added to a collection, any policies that apply to that collection automatically become effective for that system. For example, if a Linux-based server is reprovisioned to be an Apache Web server instead of a file server, the Altiris Agent collects new inventory information about the server and forwards it to the Altiris CMDB. The Altiris server then automatically removes the system from any previous collections that no longer apply and adds it to any new collections that have been defined for Apache servers.

Altiris Software Delivery Solution

The Altiris Software Delivery Solution™ tool allows administrators to install software packages, patches and patch bundles, and data or configuration files on remote UNIX/Linux-based systems via automated policies. In addition, this tool enables administrators to deliver and execute any type of shell script as an Altiris software delivery package.

Altiris Software Delivery Solution includes enhanced capabilities such as suggested install, uninstall, and rollback commands based on the success (exit status) of installation processes. A manifest file preserves UNIX/Linux permissions on packages located on Windows-based package servers; however, native UNIX/Linux-based servers can also act as package servers for software delivery. Altiris Software Delivery Solution also incorporates advanced options for distributing software in bandwidth-sensitive topologies, including checkpoint recovery and bandwidth throttling.

Many UNIX/Linux administrators manage their environment using a variety of vendor tools, but primarily with shell scripts. Software delivery allows centralized management and deployment of these scripts using a one-to-many deployment methodology. Administrators can store these scripts centrally and use Altiris Software Delivery Solution to execute them on multiple systems—even across different operating systems. Using a combination of Altiris Inventory Solution to track script versions and Altiris Software Delivery Solution to distribute scripts can help administrators greatly simplify script management.

Additionally, most UNIX/Linux administrators have skills particular to one type of UNIX or Linux distribution. These administrators may expect their skill sets to translate to other operating systems when, in reality, the commands to perform day-to-day functions can be very different. For example, the Solaris command for installing a software package is `pgkadd`; on Red Hat Enterprise Linux, this command is `rpm`; and on AIX, it is `installp`. Each of these installation commands also employs unique command-line options that further complicate the process. Altiris Software Delivery Solution is designed to help simplify this complexity by automatically suggesting the installation command for a given software package based on the file types it contains. Often, this can help administrators quickly learn how to install software on an unfamiliar OS.

As properties of UNIX/Linux-based systems change over time, the Altiris server can automatically move systems into and out of dynamic collections, thereby changing the tasks that apply to them without administrator involvement.

Altiris Monitor Solution

Altiris Monitor Solution™ software enables UNIX/Linux administrators to use the Altiris-provided monitor pack with predefined metrics in several categories or to create monitor packs using a variety of data providers, including UNIX/Linux commands, compound commands (scripts), port checks, and log file parsing. Out-of-the-box metrics include disk, memory, ports, printers, processor, and security. Common daemon metrics provided include Dynamic Host Configuration Protocol (DHCP), WU-FTPD (a replacement ftp daemon for UNIX), XINETD (a replacement daemon for inetd, the Internet services daemon), Simple Mail Transfer Protocol (SMTP), and the HTTP daemon (HTTPD).

Rules evaluate data metrics to determine whether predefined actions should be executed. Triggered rules can create a help-desk ticket, generate an e-mail to notify administrators of a system in a critical state, execute custom scripts or applications, and generate reports. Predefined reports can help administrators analyze data, and Altiris software supports creation of user-defined custom reports.

Altiris Monitor Solution offers both real-time and historical views of monitored metrics (see Figure 4) and can maintain an

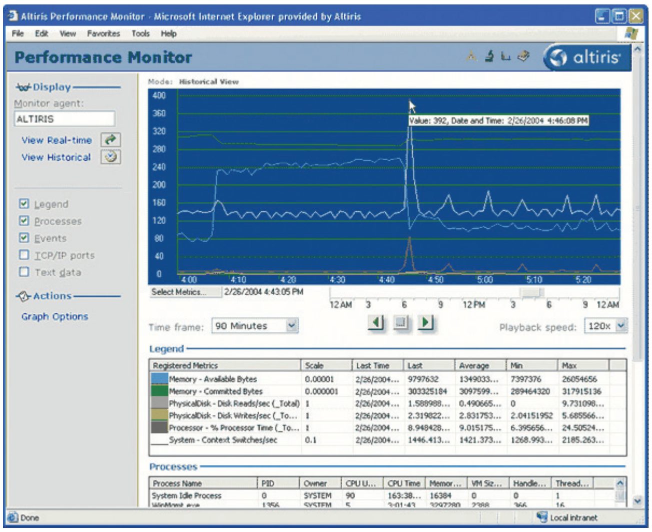

Monitor agent: ALTIRIS

Mode: Historical View

Value: 392, Date and Time: 2/26/2004 4:46:08 PM

2/26/2004 4:43:05 PM

Time frame: 90 Minutes

Playback speed: 120x

Registered Metrics	Scale	Last Time	Last	Average	Min	Max
Memory - Available Bytes	0.00001	2/26/2004...	9797632	1349033...	7397376	26054656
Memory - Committed Bytes	0.000001	2/26/2004...	303325184	3097599...	289464320	317915136
PhysicalDisk - Disk Reads/sec (_Total)	1	2/26/2004...	1.588988...	0.490665...	0	9.731098...
PhysicalDisk - Disk Writes/sec (_To...	1	2/26/2004...	2.319822...	2.831753...	2.04151952	5.685566...
Processor - % Processor Time (_To...	1	2/26/2004...	8.948428...	9.015175...	6.395656...	24.50524...
System - Context Switches/sec	0.1	2/26/2004...	1446.413...	1421.373...	1268.993...	2185.263...

Process Name	PID	Owner	CPU U...	CPU Time	Memor...	VM Siz...	Handle...	Thread...
System Idle Process	0	SYSTEM	90	163:38...	16384	0	0	1
WinMgmt.exe	1356	SYSTEM	5	3:01:43	3297280	2388	366	16

Figure 4. Altiris Performance Monitor

extensive history of low-level information for long-term trend analysis via the Web-based console. Administrators can analyze real-time and trended data together for an extremely accurate picture of data center performance.

A central monitor dashboard provides a holistic view of the state of all monitored machines, making it easy to identify problem areas. From the dashboard, administrators can drill down into detailed system logs and information provided by a variety of Altiris tools.

Another valuable feature of Altiris Monitor Solution is the periodic capture (every 60 seconds, by default) of the current process list. Servers often crash when no administrator is readily available—for example, on nights, weekends, or holidays. Determining what caused the failure typically involves a tedious search through numerous system and application log files. Because the Altiris Monitor Solution captures a variety of monitoring metrics or data points (including the process list), it can be a powerful tool to aid in tracing the root cause of a machine failure. Administrators can easily launch the historical monitoring view on the failed machine and view metric details up to the time of failure. The process list snapshot often reveals valuable insight about how applications were utilizing system resources immediately prior to system failure, enabling administrators to identify the problematic service and the circumstances surrounding it.

Altiris security auditing

IT administrators who are responsible for securing data and applications have two broad strategic goals: preventing unauthorized access to IT resources and maintaining IT services. Altiris software is designed to enhance security by automating vulnerability audits and leveraging best-of-breed remediation tools.

For administrators who want a comprehensive UNIX/Linux security audit (see Figure 5) based on predefined templates, Altiris tools offer a vulnerability assessment in seven areas:

- Antivirus status
- Security patch status
- Industry-known vulnerabilities
- Personal firewall status
- System security configuration settings
- Unauthorized software
- Unauthorized hardware

Altiris security tools can even provide agent-less support for UNIX/Linux environments using SSH. Agent-less auditing is often the preferred method for auditing desktop and server systems. Audit credentials can be assigned to the security audit team, but not revealed to the local desktop user. Full-time agent support is also available for Windows, UNIX, and Linux environments. No administrative credentials are required when the agent is installed locally.

Using Altiris security management software, IT administrators can perform the following tasks:

- Secure the environment according to best practices, while implementing mitigation and shielding techniques to prevent common breaches.
- Audit the environment against a database of known vulnerabilities and integrate the information with common configuration management components.
- Fix discovered vulnerabilities against a standard risk prioritization strategy to control or eliminate the root causes.

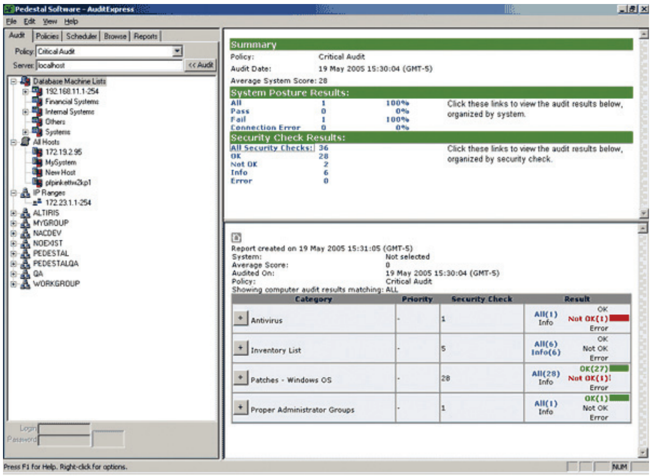

Policy: Critical Audit

Server: localhost

Summary

Policy: Critical Audit

Audit Date: 19 May 2005 15:30:04 (GMT-5)

Average System Score: 28

System Posture Results:

All	1	100%
Pass	0	0%
Fail	1	100%
Connection Error	0	0%

Click these links to view the audit results below, organized by system.

Security Check Results:

All Security Checks:	36
OK	28
Not OK	2
Info	6
Error	0

Click these links to view the audit results below, organized by security check.

Report created on 19 May 2005 15:31:05 (GMT-5)
System: Not selected
Average Score: 0
Audited On: 19 May 2005 15:30:04 (GMT-5)
Policy: Critical Audit
Showing computer audit results matching: ALL

Category	Priority	Security Check	Result	
Antivirus	-	1	All(1) Info	OK Not OK(1) 🟥 Error

Inventory List	-	5	All(6)
			Info(6)
				OK
				Not OK
				Error

Patches - Windows OS	-	28	All(28) Info	OK(27) 🟩 Not OK(1) Error
Proper Administrator Groups	-	1	All(1) Info	OK(1) 🟩 Not OK Error

Figure 5. View of a Linux-based server in Altiris AuditExpress

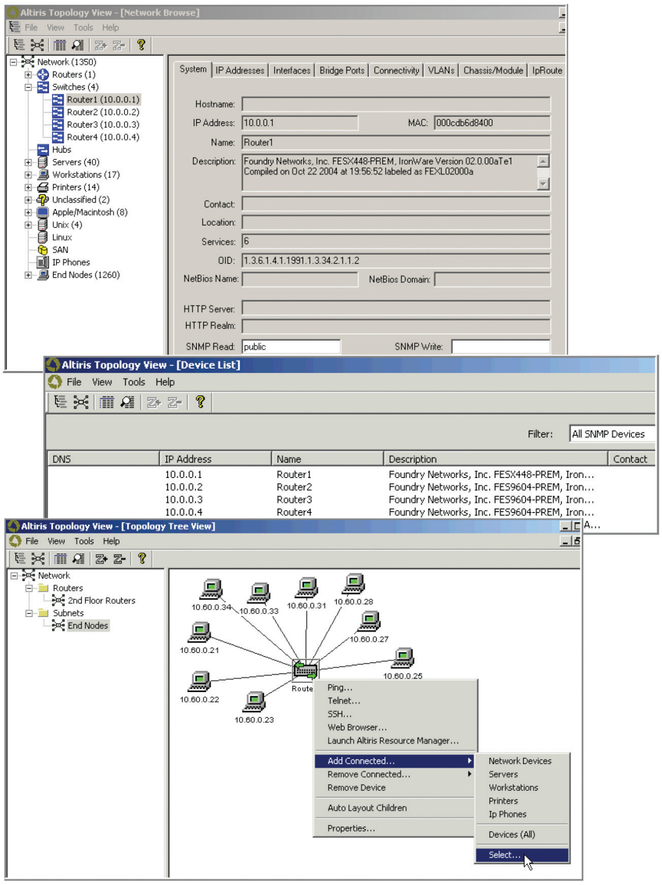

Hostname:
IP Address: 10.0.0.1 MAC: 000cdb6d8400
Name: Router1
Description: Foundry Networks, Inc. FESX448-PREM, IronWare Version 02.0.00aTe1
Compiled on Oct 22 2004 at 19:56:52 labeled as FEXL02000a
Contact:
Location:
Services: 6
OID: 1.3.6.1.4.1.1991.1.3.34.2.1.1.2
NetBios Name: NetBios Domain:
HTTP Server:
HTTP Realm:
SNMP Read: public SNMP Write:

Filter: All SNMP Devices

DNS	IP Address	Name	Description	Contact
	10.0.0.1	Router1	Foundry Networks, Inc. FESX448-PREM, Iron…	
	10.0.0.2	Router2	Foundry Networks, Inc. FES9604-PREM, Iron…	
	10.0.0.3	Router3	Foundry Networks, Inc. FES9604-PREM, Iron…	
	10.0.0.4	Router4	Foundry Networks, Inc. FES9604-PREM, Iron…	A…

Figure 6. Automated network discovery and management

- Enforce best practices through standardized configuration management, policy-based system maintenance, and change-control procedures.

Altiris network discovery and topology mapping

Documenting and tracking changes in network topology can consume a tremendous amount of time and effort. Typically, enough of the network has changed by the time such a documentation project ends that the resulting topology maps are outdated.

Altiris offers a variety of software products for standardizing the deployment and configuration of network devices and virtual LAN (VLAN) settings. Additionally, Altiris tools can generate a Layer 2 connectivity map based on a quick scan of the environment to eliminate inaccuracies and save time. An assortment of maps can be created using various filters and editing tools (see Figure 6). Furthermore, once topology maps are finalized, they can easily be exported to a Microsoft Visio® file for further editing.

These Altiris tools also can be used to locate rogue devices and quarantine them by moving the devices from the production network to a non-production VLAN. Needed system patches and configuration changes can be performed in the quarantine VLAN before moving a noncompliant device back to the production LAN.

Centralizing systems management for heterogeneous environments

Altiris provides a comprehensive set of software products for managing UNIX/Linux-based systems via a single console and management infrastructure. Altiris software can help dramatically improve an IT organization's ability to "do more with less," especially within heterogeneous environments. For Windows-centric organizations, Altiris tools not only can help to efficiently manage systems already in place, but they also can help position IT organizations to simply and easily begin building out a Linux strategy.

By providing policy-driven, one-to-many management tools for diverse environments, Altiris software is designed to automate several common management functions. Additionally, Altiris tools can provide a helpful layer of abstraction that helps minimize the differences between operating systems for low-level, one-to-one management tasks and utilities—enabling administrators to quickly become effective at managing diverse environments.

Patrick Bourke is a senior technical support engineer for Akibia, Inc. He has assisted numerous Dell customers, along with Akibia's internal IT department, in designing comprehensive Altiris management infrastructures. Prior to joining Akibia, Patrick spent many years as a network engineer and systems management consultant specializing in the design, implementation, and long-term support of complex, multivendor data center environments.

Todd Mitchell is the Dell alliance technical director at Altiris. He has worked with numerous Altiris customers to support Dell-specific implementations and management needs. Todd has a bachelor's degree from Brigham Young University.

Rich Lacey is the product line manager for server management and the UNIX/Linux advocate at Altiris. Prior to joining Altiris, Rich spent several years as a consultant and engineer specializing in the architecture, implementation, and integration of operational support systems in the telecommunications industry.

FOR MORE INFORMATION

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

Comparing DRAC 4 Serial Console Redirection Methods: COM2 Versus Video

The Dell™ Remote Access Controller 4 offers a rich feature set that allows administrators to configure and manage a Dell PowerEdge™ server from a remote location, whether it is across the hall or halfway around the world. This article describes Dell's implementations of the connect com2 and connect video methods, including how to configure servers for serial console redirection. It also examines the movement toward an industry-standard model for serial access to remote host consoles, enabling out-of-band management when in-band tools are not available.

BY CARL KAGY AND JON McGARY

Related Categories:

Dell PowerEdge servers

Dell Remote Access Controller (DRAC)

Remote management

Systems management

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

The Dell Remote Access Controller 4 (DRAC 4) offers a rich feature set, including serial console redirection for configuring and managing remote Dell PowerEdge servers. The DRAC serial interface lets administrators monitor a remote server and interact with the remote server's OS console through the server's serial port. From a serial session, administrators can power up, power down, power cycle, or reset a remote managed system; view logs; monitor sensor status; configure the DRAC; issue `racadm` commands at the command-line interface; and redirect the managed system's text console.

This article explains how the serial interface can redirect the managed system's text console and describes how Dell implements serial console redirection—exploring how the connect com2 and the connect video methods work, along with the limitations of each approach. In addition, this article discusses movement toward an industry-standard remote access model for serial console management, which is a critical component for system administrators who need to monitor and recover servers when they cannot access the LAN.

Understanding serial console redirection

Increased use of Linux® OS–based servers spurred the introduction of serial console redirection. The intent of this feature was to make the remote server console on the serial port (COM1 or COM2) available through the DRAC; the DRAC would internally connect to one of the server's serial ports and act as a proxy for that port on a Telnet or serial session. Dell developed two architectures to implement serial console redirection: the connect com2 method creates a serial connection to one of the system's COM ports while the connect video method pulls data from the video hardware.

The `connect com2` command establishes a serial session between the DRAC 4 and the managed system's COM port (see Figure 1). This approach allows the DRAC to redirect serial data from the host, with the benefit that the system's hardware controls the data flow for the connection. Although the connect com2 method works well on most Dell servers, exceptions have occurred when the server hardware is not designed to support an internal connection directly between the DRAC and the COM port. Simply to address that exception, the connect video method for serial console redirection was designed. However, to maintain a consistent feature set, the connect video approach is enabled in all Dell servers.

Although the console output is substantially the same, the connect video feature operates with entirely different input data than the connect com2 feature. Connect video architecture is different from connect com2 console redirection in that connect video requires the DRAC to access the managed server's video controller hardware and PCI bus. To redirect the serial console, the DRAC routinely polls the video frame buffer for its data. When a difference between frames is detected, the DRAC converts that difference into characters and VT100 sequences for output to the remote session.

Limitations of the serial console redirection methods

Connect com2 and connect video offer different approaches to serial console redirection. Both methods possess inherent limitations, such as performance trade-offs, OS crash history, data integrity risks, and server design constraints.

Performance trade-offs. The connect video method can suffer performance degradation because it requires the DRAC to control all of the data flow for the serial console redirection session. In contrast, connect com2 lets the managed system's hardware control all of the data flow for the session. As a result, connect com2 tends to perform faster and to use fewer DRAC processing cycles than connect video.

OS crash history. Because connect video architecture pulls data from the video frame buffer, it can retrieve a snapshot of the last output from a failed or crashed OS. Early versions of connect com2 could not provide this functionality. Until recently, connect com2 could show output from the OS only after the administrator had initiated a serial session and issued the serial console redirection command.

Versions 1.3 and later of DRAC 4 firmware maintain a serial history buffer of character output from the server's serial port, even when a DRAC serial session is not active. Subsequent sessions of connect com2 can then use the `-h` history option. The history buffer size is configurable, possessing a default maximum limit of 8,192 characters—more information than can be captured from the final video screen.

Data integrity risks. The connect video method of comparing snapshots of video data possesses an inherent disadvantage. When rapidly changing information is compared, some data can be missed between snapshots in time, which is equivalent to dropping characters. This problem is unlikely using the connect com2 method because connect com2 architecture establishes a connection directly to the managed server's serial port, which uses hardware flow control. The hardware controls the commencement and termination of all data sent over the serial line, which helps ensure that all characters are captured for presentation to the remote system.

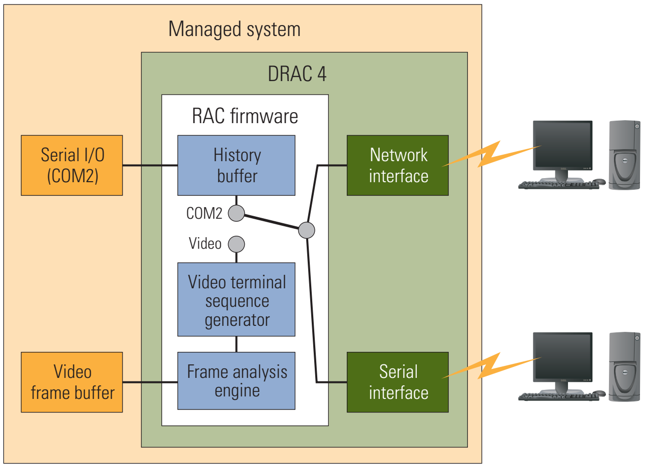

Figure 1. Connect com2 method combined with a history buffer

Server design constraints. The connect video method constrains server hardware designs because it requires the DRAC to be on the PCI system bus and tightly coupled with the video hardware. A second limiting factor is that connect video architecture does not extend to modular blade server systems in which multiple CPU blades reside in the same enclosure. For example, in the Dell Modular Server Enclosure, the DRAC hardware does not connect to the PCI bus or to the individual server blade module's video hardware. Consequently, of the two approaches, connect com2 provides the only viable way to redirect the serial console to an individual server blade that resides in a modular system enclosure.

Moving toward a standardized remote access model

Serial console management is a critical component enabling administrators to monitor and recover servers when the LAN is not available. For a step-by-step example of how to configure Dell PowerEdge servers (running a Microsoft® Windows® or Linux OS) for serial console redirection, see the "Tutorial: How to configure serial console redirection on a DRAC 4–based server" sidebar in this article.

The industry-standard Intelligent Platform Management Interface (IPMI) defines a Serial Over LAN (SOL) approach that enables servers to be monitored and controlled remotely. The capabilities and behavior of the IPMI SOL method for serial console redirection are similar to the Dell connect com2 approach. Both IPMI SOL and connect com2 provide output similar to that of a serial terminal concentrator, which is commonly used as the basis for systems

TUTORIAL: HOW TO CONFIGURE SERIAL CONSOLE REDIRECTION ON A DRAC 4–BASED SERVER

This scenario provides a step-by-step example of how administrators can configure a Dell PowerEdge server for serial console redirection. This example starts with a rack of Windows- or Linux-based servers, each configured with an up-to-date DRAC 4. The configuration is designed to enable the IT administrator to connect to a DRAC 4–equipped server and observe the recent activity on that server's console—especially if that server has experienced a failure. To help accomplish this objective, the IT administrator should perform the following tasks.

Enable serial and Telnet features within the DRAC 4

First, from the local Windows- or Linux-based system, enter the following commands:

```
racadm config -g cfgSerial -o
   cfgSerialConsoleEnable 1

racadm config -g cfgSerial -o
   cfgSerialTelnetEnable 1
```

Then, enter the following commands from the remote console:

```
racadm -u username -p password -r DRAC 4
   IP address config -g cfgSerial -o
   cfgSerialConsoleEnable 1

racadm -u username -p password -r DRAC 4
   IP address config -g cfgSerial -o
   cfgSerialTelnetEnable 1
```

Configure the system setup program on the managed system

Next, set up the BIOS to route the serial port signals to the DRAC 4 using the following procedures for both Windows- and Linux-based servers:

1. Turn on or restart the system.
2. Press F2 immediately after "<F2> = System Setup" appears on the display.
3. Scroll down, select "Integrated Devices," and press Enter.
4. In the submenu, scroll down to Serial Port 1 and select "RAC."
5. Scroll down and select "Console Redirection."
6. On the Console Redirection screen, select the following settings:
   - Console Redirection: Serial Port 1
   - Redirection After Boot: Disabled
7. Press Esc to exit the system setup program and complete the configuration.

At this point, the Windows-based server configuration is complete but Linux-based servers require a few additional steps.

Configure the Linux system files

Configuring Linux to use a serial console requires the modification of three startup files within the server's Linux system image.[1]

1. **Modify the Grand Unified Bootloader (GRUB) configuration.** In the Linux configuration file /etc/grub.conf, add the following two lines in the general settings section of the file:

   ```
   serial --unit=0 --speed=57600

   terminal --timeout=10 serial
   ```

   Next, append the `kernel` and `console=ttyS0,57600` options to the kernel line. For example:

   ```
   kernel /vmlinuz-2.4.20-8smp ro root=LABEL=/
      console=ttyS0,57600
   ```

   If a splashimage directive such as `splashimage=(hd0,2)/grub/splash.xpm.gz` appears, comment it out.

2. **Enable login to the console after boot.** In the Linux configuration file /etc/inittab, add a line to configure a getty on the COM1 serial port as follows:

   ```
   co:2345:respawn:/sbin/agetty -h -L 57600 ttyS0
      vt100
   ```

3. **Grant permission to initiate the session.** In the Linux configuration file /etc/securetty, add ttyS0 to the list of supported ports.

Establish a Telnet or serial session

After the initial configuration steps have been completed, the administrator can connect to the DRAC 4 using terminal emulation software such as Hilgraeve HyperTerminal or minicom; Telnet; or Secure Shell (SSH). SSH offers an interface text window similar to Telnet but with encryption for greater security.

After connecting to the DRAC 4 and authenticating with a username and password, the administrator can issue the command `connect -h com2` within the session to display the most recent output from the managed system's serial console. Typing "help" should display a list of valid serial session commands.

[1] For more information, visit the *Dell Remote Access Controller 4 User's Guide* at support.dell.com/support/edocs/software/smdrac3/drac4/1.1/en/UG/racugc3.htm#wp56214.

management connectivity in place of a DRAC or baseboard management controller within the server.

Serial console management is not limited to Linux-based servers. Both Linux and Windows communities are moving toward serial connectivity—not as an optional interface but as a systems management requirement. Microsoft Windows Server™ 2003 introduced the Emergency Management Services suite of features, which supports remote management and system recovery through the managed server's serial port. The Emergency Management Services suite runs while the system is booting and the Special Administration Console runs after the graphical interface is active; both use the server's serial port and thus can be accessed using serial console redirection.

Enabling reliable, cost-effective remote access

The server industry is moving toward a remote access model for both Windows- and Linux-based systems, in which connection to a host console is achieved through a serial port. Out-of-band access is not intended to replace traditional in-band management solutions. Instead, out-of-band access offers an alternative for administrators to monitor and control remote managed systems when standard in-band management tools are not available.

Today, the server industry is moving away from the connect video architecture, which has inherent drawbacks that limit performance and data integrity and restrict hardware design for future server generations and modular systems. The remote access model presented in this article enables out-of-band, serial access through the DRAC 4 to the managed system using the connect com2 method combined with a history buffer. In this way, administrators can access managed systems using serial console redirection features that are designed to provide a reliable, cost-effective remote connection mechanism.

Carl Kagy is a senior software developer in the Dell OpenManage Remote Management Group. Prior to joining Dell, Carl was employed by NCR, Tandem Computers, and IBM, and specialized in remote management of fault-tolerant computers. He has a B.S. from Case Western Reserve University.

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

FOR MORE INFORMATION

McGary, Jon and Carl Kagy. "Remote Configuration of Serial and Telnet Interfaces to the DRAC 4." *Dell Power Solutions,* October 2004. www.dell.com/downloads/global/power/ps4q04-20040102-McGary.pdf

FREE Subscription Request

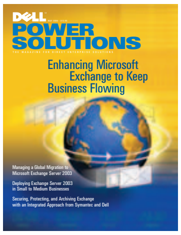
DELL
POWER SOLUTIONS
Enhancing Microsoft Exchange to Keep Business Flowing
Managing a Global Migration to Microsoft Exchange Server 2003
Deploying Exchange Server 2003 in Small to Medium Businesses
Securing, Protecting, and Archiving Exchange with an Integrated Approach from Symantec and Dell

❑ **Yes!** I want to receive *Dell Power Solutions* **Free.**

Check one:

❑ New subscription ❑ Renew subscription

❑ Address change ❑ Cancel subscription

Current subscriber ID (from mailing label): ______

First name: ______

Last name: ______

Company name: ______

Address 1: ______

Address 2: ______

City: ______ State, province, or territory: ______

ZIP/postal code: ______ Country: ______

Telephone: ______

E-mail address: ______

Please complete the questions listed below to receive your free subscription to *Dell Power Solutions*. You must answer all questions to qualify for a free subscription.

1. Which of the following best describes your job function?

❑ Management (CxO, VP, director)
❑ IT manager or supervisor
❑ Systems engineer/ technical consultant
❑ Systems analyst
❑ System administrator
❑ Network administrator
❑ Project manager
❑ Marketing/sales
❑ Other

2. How large is your company, in terms of annual sales volume?

❑ Less than $5 million
❑ $5–$9 million
❑ $10–$49 million
❑ $50–$99 million
❑ Greater than $100 million
❑ Greater than $1 billion

3. How large is your company, in terms of employees?

❑ Less than 200
❑ 200–500
❑ 500–1,000
❑ 1,000–5,000
❑ Greater than 5,000

4. What is the principal industry served by your company?

❑ Education
❑ Financial services and banking
❑ Government and military
❑ Healthcare
❑ Hospitality
❑ Internet or Internet service provider
❑ Manufacturing
❑ Retail
❑ Telecommunications
❑ Utilities
❑ Other

5. What Dell products does your company use?

❑ Desktops or notebooks
❑ Servers or storage
❑ All of the above
❑ None of the above

6. What operating systems does your company use?

❑ Windows
❑ Novell
❑ UNIX
❑ Linux
❑ Mixed
❑ Other

05/06

Subscriptions are free to qualified readers who complete the online subscription form or submit this subscription reply form. To sign up as a new subscriber, renew an existing subscription, change your address, or cancel your subscription, submit the online subscription form at www.dell.com/powersolutions_subscribe; fax this subscription reply form to +1 512.283.0363; or return this subscription reply form by surface mail to *Dell Power Solutions*, Mail Stop RR5-03, Dell Inc., One Dell Way, Round Rock, TX 78682, U.S.A. For subscription services, please e-mail us_power_solutions@dell.com.

Subscribe online at www.dell.com/powersolutions_subscribe

Improving Real-Time Access to Data Center Servers

with Dell 2161DS-2 and 4161DS Remote Console Switches

The Dell™ 2161DS-2 and 4161DS Remote Console Switches are designed to provide KVM (keyboard, video, mouse) control of up to 256 attached servers with expanded access and improved performance. Incorporating KVM over IP™ technology from Avocent, the rack-mountable Dell 2161DS-2 and 4161DS switches can give administrators flexible, centralized control of servers over a single IP connection.

BY ROBERT BERNSTEIN AND MAX A. BENHAM

Related Categories:

Avocent

Dell PowerConnect switches

Keyboard, video, mouse (KVM)

Systems management

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

With integrated KVM (keyboard, video, mouse) technology, Dell Remote Console Switches enable administrators to manage servers regardless of location. Like the previous-generation Dell 2161DS switch, the Dell 2161DS-2 allows for simultaneous access by two remote users and one local user. The Dell 4161DS switch, however, allows for simultaneous access by up to four remote users, or three remote users and one local user (see Figure 1).

The Dell 2161DS-2 and 4161DS switches both feature Digital Share Mode, which allows multiple users to access the same server simultaneously. In addition, the 2161DS-2 and 4161DS include performance enhancements over previous-generation Dell Remote Console Switches. These switches use an algorithm developed by Avocent called Dambrackas Video Compression™ (DVC). Specifically optimized for interactive user sessions of KVM over IP, DVC provides extremely high compression rates for video data. This technology enables a near–real-time experience by improving color depth (15-bit rather than 7-bit) and by minimizing mouse response latency.

Encryption and compression algorithms are also embedded in the switch hardware. Video quality through local and remote sessions can achieve resolutions of up to 1,024 × 768. A higher resolution of 1,280 × 1,024 can be obtained with specialized adapters for local sessions. With the Dell 2161DS-2 and 4161DS switches, interfacing with remote systems can be just like using the keyboard, video, and mouse that are directly attached to the computer.

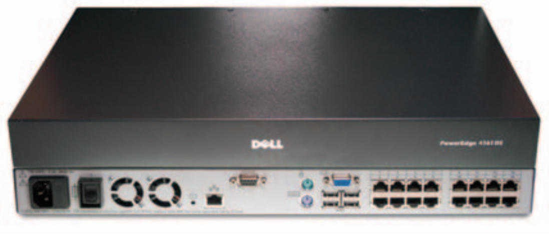

Figure 1. Dell 4161DS Remote Console Switch

Benefiting from advanced control

Like the previous-generation Dell 2161DS switch, the Dell 2161DS-2 and 4161DS switches provide secure analog (local) and digital (remote) connectivity to all major operating systems, server platforms, and serial devices. Each switch features 16 Category 5 (Cat 5) RJ-45 ports for server connections. By using inexpensive Cat 5 cables, administrators can install multiple servers in a single rack without the typical cable bulk.

The 2161DS-2 and 4161DS switches support local USB components in addition to PS/2 devices. Administrators can connect a USB or PS/2-enabled mouse and keyboard. In addition, these switches provide a 10/100/1,000 Ethernet interface for digital remote connections. This forward-looking Gigabit Ethernet support means switch functionality can scale over time. And because the 2161DS-2 and 4161DS switches are flash-upgradeable, they are designed to accommodate future hardware compatibility.

Administrators can access the switches locally or remotely through the Dell Remote Console Software (RCS) Java-based client interface (see Figure 2). RCS is a cross-platform management application that gives administrators a single point of access for the entire system. This easy-to-use software lets administrators configure the switch and select the servers to control.

Additionally, the 2161DS-2 and 4161DS switches include the Avocent On-Screen Configuration and Activity Reporting (OSCAR®) overlaid menu system. This interface provides analog control from the local workstation.

Supporting and expanding existing KVM configurations

With the Dell 2161DS-2 and 4161DS switches, administrators can choose from two types of server interface pods (SIPs)—USB or PS/2—to fit any server need (see Figure 3). The SIPs replace standard KVM cables by converting a server's KVM signals, then sending the converted signals through a single Cat 5 cable. Because the name of the server configured by OSCAR or RCS is stored in the SIP, whenever a server is moved or a rack is re-cabled, the name of that server automatically follows it to the new switch or port location.

The SIPs get their power from the server through the USB or PS/2 connections. This helps ensure that the server works with or without connectivity to the 2161DS-2 or 4161DS switch. By using this keep-alive technology, the SIP is designed to provide continuous keyboard and mouse emulation.

The Dell 2161DS-2 and 4161DS switches can support dense data center environments. With a port expansion module, each of the 16 ports on the switch can connect to up to eight servers—for a total of up to 128 servers per switch. Beyond that, each port on the 2161DS-2 or 4161DS can also connect to an existing analog Dell switch attached to as many as 16 servers, enabling administrators to scale out a single 2161DS-2 or 4161DS switch to connect to up to 256 servers. However, with these types of expansion scenarios, the same limits on simultaneous access by remote and local users still apply.

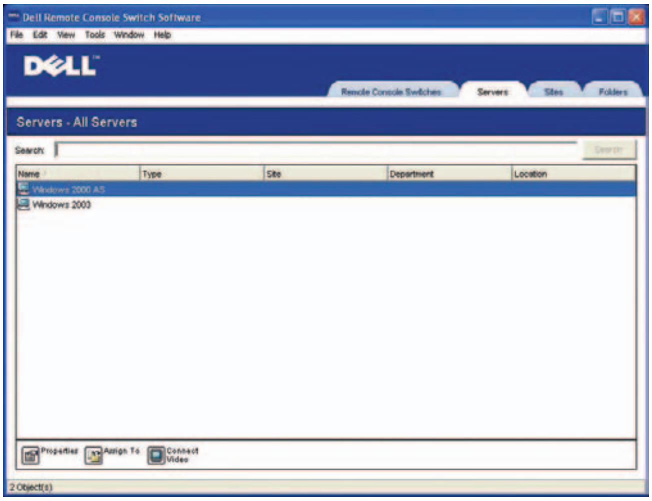

Name	Type	Site	Department	Location
Windows 2000 AS				
Windows 2003				

Figure 2. Dell Remote Console Software client interface

Identifying various usage models

From classroom environments to dense data centers, the Dell 2161DS-2 and 4161DS switches can help administrators gain secure remote access to any server at any time. The scenarios discussed in this section represent common uses for these Dell Remote Console Switches.

Server rooms

In enterprises of all sizes, server rooms are often in inconvenient locations for IT administrators. These administrators can waste valuable time making their way to a failed server in another building or across a corporate campus, which takes a toll on server downtime and user productivity. The Dell 2161DS-2 and 4161DS switches can help keep downtime minimal and productivity high by providing remote access to servers, no matter where they are located.

The need for remote access is especially apparent in branch-office environments, where the server room is not simply a few

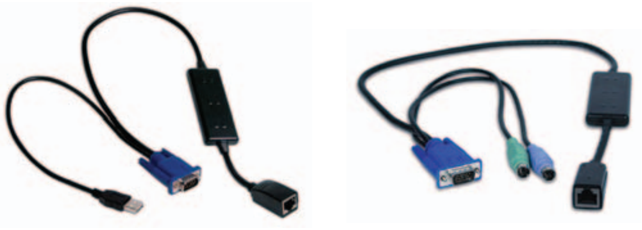

USB SIP

PS/2 SIP

Figure 3. Server interface pods for the Dell 2161DS-2 and 4161DS switches

buildings away, but rather a few cities—or countries—away. With the Dell 2161DS-2 and 4161DS switches, server administrators can perform the necessary managerial tasks as if they were sitting at the server, even though they could be halfway around the world.

As demands for computing power escalate, enterprises are adding more and more computers to the server room, making the space quite dense. With available space at a premium, IT administrators can benefit from the local access provided by the Dell 2161DS-2 and 4161DS switches. These switches also provide local control of multiple servers and other devices from a single keyboard, monitor, and mouse. Even if the systems are not homogeneous, nearly any type of server can be controlled from the same console.

The Dell 2161DS-2 and 4161DS Remote Console Switches enable IT administrators to respond quickly and securely to server issues, no matter where the servers or the administrators are located.

Secure data centers and network operations centers

Data centers and network operations centers often have physical security restrictions to consider. To prevent unauthorized access to equipment, these centers sometimes limit physical access to systems—even limiting those charged with maintaining the systems.

The Dell 2161DS-2 and 4161DS switches can help enterprises maintain secure access to key systems by letting administrators securely manage and control data center servers from remote consoles. Security is provided through multilevel password protection and 128-bit Secure Sockets Layer encryption, so the servers stay physically and virtually secure while administrators gain the access they need.

Co-location facilities

For advanced power and network redundancy, many enterprises place their servers—particularly Web servers—at co-location facilities. These centers typically house hundreds or even thousands of servers, providing a superfast Internet connection and a sustainable power supply. What these facilities do not provide, however, is maintenance of the servers.

Remote KVM over IP access can help administrators effectively manage servers housed in such facilities. The 2161DS-2 and 4161DS switches enable comprehensive remote operation with unfettered access by multiple users, while preventing access by unauthorized users.

Training environments

Because the Dell 2161DS-2 and 4161DS switches feature Digital Share Mode, multiple users can share a remote session on the same server. This feature can be beneficial in training environments, either within a classroom or with trainees spread all over the world.

In Digital Share Mode, the Dell 2161DS-2 or 4161DS switch lets up to 12 users view what a trainer is showing on-screen, regardless of where the users are located. The 12 users have view-only access, so they can see what is happening on the remote server's screen, but they cannot interfere with the actions of the trainer.

This capability can be used to train IT administrators at various remote locations simultaneously. For example, a specialist based at an enterprise's headquarters could guide administrators at several remote offices through a new backup procedure or server upgrade. The administrators would be able to view the screen as if they were sitting at the server console, learning the task without interfering with the procedure itself.

Securely managing remote servers in real time

The Dell 2161DS-2 and 4161DS Remote Console Switches enable IT administrators to respond quickly and securely to server issues, no matter where the servers or the administrators are located. Features such as Digital Share Mode and DVC as well as support for USB devices and Gigabit Ethernet are designed to significantly improve performance compared to previous-generation Dell Remote Console Switches. With the Dell 2161DS-2 and 4161DS switches, administrators have the ability to connect to, monitor, and troubleshoot servers in real time, whether they are across the hall or halfway around the world.

Robert Bernstein is a product marketing manager for racks and rack peripherals in the Dell Enterprise Systems Group. Previously, he assisted with advanced system sales in the Dell Small and Medium Business Group for eight years. Robert has a B.S. in Communications from The University of Texas at Austin and is a Microsoft Certified Systems Engineer.

Max A. Benham is the Dell appliance account manager at Avocent. Previously, he led the Avocent original equipment manufacturer (OEM) Program Management organization. Max has a B.S. in Economics and a B.A. in Slavic Languages and Literature from the University of Washington.

FOR MORE INFORMATION

Dell 2161DS-2 and 4161DS Remote Console Switches:
www.dell.com/downloads/global/products/pedge/en/2161DS-2_4161DS.pdf

Access the data center from the comfort of anywhere.

The Dell™ 2161DS-2/4161DS remote console switches give today's server administrators the power to monitor and control servers from virtually anywhere. For maximum flexibility, the 4161DS provides **access for up to four simultaneous users and by using Dell's remote console software, simply point-and-click to take control of your servers. Whether you're at the rack, in your office, across the globe, or in your living room, your data center can always be within your reach.**

Digital Availability. Easy as DELL

www.dell.com

Cascading the Avocent Digital Access KVM Switch to a KVM Infrastructure

As a component within the Dell Modular Server Enclosure, the Avocent Digital Access KVM (keyboard, video, mouse) switch provides remote access capabilities to Dell™ PowerEdge™ 1855 and PowerEdge 1955 blade servers. Digital Access KVM switches can be cascaded to Dell external KVM devices, such as the Dell 2161DS-2 and 4161DS Remote Console Switches, and can be managed using the On-Screen Configuration and Activity Reporting (OSCAR®) interface.

BY BABU CHANDRASEKHAR AND JAKE DINER

Related Categories:

Avocent

Keyboard, video, mouse (KVM)

Systems management

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

With Avocent Digital Access KVM (keyboard, video, mouse) switches, IT administrators can remotely monitor and control Dell PowerEdge 1855 and PowerEdge 1955 blade servers. The Digital Access KVM switch consolidates KVM connections and provides a single local terminal connection to the blade server chassis, known as the Dell Modular Server Enclosure. The Digital Access KVM switch has a 10/100 Mbps network connection that can be used to remotely connect to the individual server blades within the Dell Modular Server Enclosure to take advantage of virtual media and remote KVM features.

The Avocent Digital Access KVM can be cascaded to external remote console KVM switches, such as the Dell 2161DS-2 and 4161DS Remote Console Switches, or to analog console switches such as the Dell 180AS, Dell 2160AS, and other legacy KVM switches. The external KVM switch can then be connected to an enterprise's KVM infrastructure—interconnecting several servers and KVM switches over an enterprise-wide network and enabling administrators to manage multiple servers from a single console using Dell Remote Console Software (RCS).

Understanding the Avocent Digital Access KVM switch

The Avocent Digital Access KVM switch provides one RJ-45 network connector and a custom KVM dongle with two PS/2 ports and one video port (see Figure 1). The Dell Modular Server Enclosure can support an Avocent Digital Access KVM switch or an Avocent Analog KVM switch. In the analog switch, the RJ-45 connector is an Analog Console Interface (ACI) port that can be used to cascade to an external KVM switch. The RJ-45 network connector in the Digital Access KVM switch is not a direct console interface and therefore it cannot

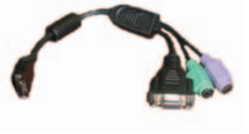

Figure 1. Dell PS/2 KVM dongle

be used for KVM cascading. Instead, the KVM dongle must be connected to a PS/2 server interface pod (SIP) to cascade the Digital Access KVM switch to an external KVM switch (analog or digital). On one end of the SIP are two PS/2 connectors and one video connector; on the other end is an RJ-45 tiering connector (see Figure 2). The RJ-45 connector on the SIP is an ACI port that can be connected to an external KVM switch using a standard Category 5 (Cat 5) cable, but it is not an Ethernet network port. The corresponding RJ-45 connector on an external KVM switch is called an Analog Rack Interface (ARI) port.

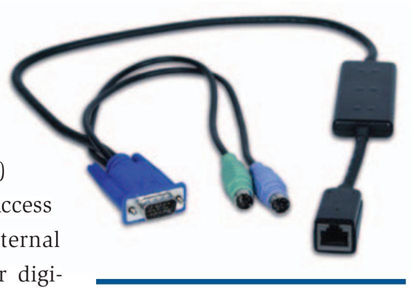
Figure 2. Dell PS/2 server interface pod

The Digital Access KVM switch supports two user interfaces. The network configuration, virtual media, and remote KVM settings can be configured from the Dell Remote Access Controller/Modular Chassis (DRAC/MC) user interface. The display interface for server selection is the On-Screen Configuration and Activity Reporting (OSCAR) menu, which is available in all Dell KVM switches. SIP module firmware can be updated through the OSCAR interface; Digital Access KVM firmware must be updated through the DRAC/MC interface.

Cascading to an external KVM switch

Dell external KVM switches such as the 2161DS-2 and 4161DS have ARI ports, which accept Cat 5 cabling from an ACI port. To cascade KVM switches, the KVM dongle from the Avocent Digital Access KVM switch should be connected to the SIP, and then the SIP's RJ-45 connector (ACI port) should be cabled to one of the external switch's ARI ports (see Figure 3). The Dell Modular Server Enclosure with the Digital Access KVM switch must be powered up before administrators can make the cascading connection to the external KVM switch.

The Digital Access KVM switch's local security settings such as the password and screen saver must be disabled before cascading the KVM switches. Administrators can disable these security settings from the local OSCAR menu by connecting a local keyboard and monitor to the Digital Access KVM dongle.

Dell KVM switches support only two levels of cascading; therefore, the Digital Access KVM switch can be tiered under only one switch. It cannot be connected to a port expansion module.

After setting up the cascading connection, administrators can use the external KVM switch's OSCAR menu to access the server blades within the Dell Modular Server Enclosure. Best practices recommend resetting the SIP module in order for the external KVM switch to replicate the Digital Access KVM switch's settings. The SIP can be reset from the OSCAR menu by selecting the following options in order: Commands, Display Versions, SIP-Number, Version, and Reset. After the reset, the SIP presents itself to the external KVM switch as a 10-port KVM switch.

If the external KVM switch is managed from a remote workstation using RCS, then the RCS database must be resynchronized with all of the connected KVM switches. The database resynchronization is initiated from RCS by selecting the Management Properties menu option and then clicking the Resync button in the Server category. To resynchronize multiple remote client workstations, administrators can save the resynchronized local database on the first workstation and then load it onto the other client workstations.

The Digital Access KVM switch consolidates KVM connections and provides a single local terminal connection to the blade server chassis, known as the Dell Modular Server Enclosure.

Connecting to external KVM switches from other vendors

The Avocent Digital Access KVM switch can be connected to non-Dell external KVM switches by connecting the KVM dongle directly to the PS/2 ports and video ports of the external KVM switch. In this type of configuration, known as non-seamless tiering, administrators first select the Digital Access KVM port from the menu of the external KVM switch's on-screen display menu, and then press the Print Screen key to access the OSCAR menu for the Digital Access KVM switch. Non-seamless tiering

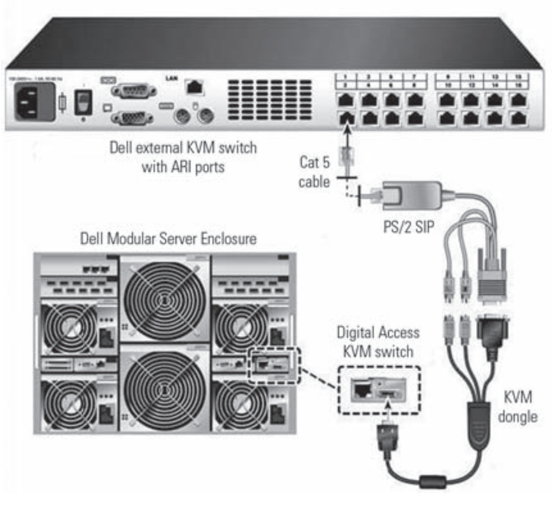

Figure 3. Cascading a Digital Access KVM switch to a Dell external KVM switch through Cat 5 cabling

is not supported by Dell external KVM switches.

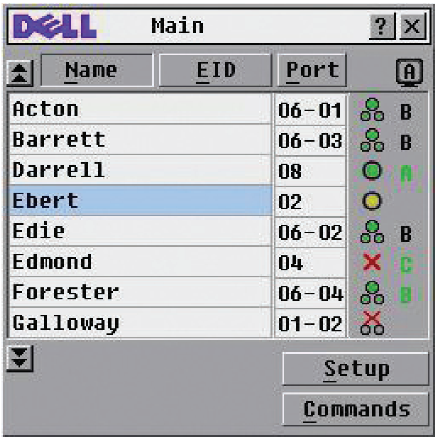

Name	EID	Port		A
Acton		06-01	🟢∴	B
Barrett		06-03	🟢∴	B
Darrell		08	🟢	A
Ebert		02	🟡	
Edie		06-02	🟢∴	B
Edmond		04	✗	C
Forester		06-04	🟢∴	B
Galloway		01-02	✗∴	

Figure 4. OSCAR interface showing SIP online status

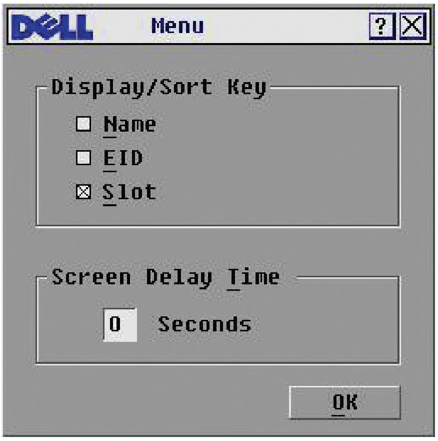

Display/Sort Key

☐ Name

☐ EID

☒ Slot

Screen Delay Time

0 Seconds

Figure 5. OSCAR interface to set screen delay

Configuring OSCAR settings

The OSCAR interface for the Avocent Digital Access KVM switch is a subset of the external KVM switch's user interface. The OSCAR interface displays the list of servers sorted by name, electronic identifier (EID), or port number. The EID option is available only on external switches. Once switches are cascaded, the external switch's OSCAR menu becomes the user interface for the Digital Access KVM switch. The server blades in the Dell Modular Server Enclosure are labeled in the external KVM switch's OSCAR menu as *x*-01 to *x*-10, where *x* is the port number to which the Digital Access KVM switch is connected.

A cascaded KVM connection handles only server selection logic and does not share custom settings from a KVM switch in a lower tier. Settings in the Digital Access KVM switch such as custom server names, language, security settings, and so forth are not propagated to the external KVM switch. The Digital Access KVM switch does not provide a language selection option because the firmware image differs for each supported language.

Avocent Digital Access KVM switches give server administrators the convenience of remotely monitoring and controlling Dell PowerEdge server blades within a Dell Modular Server Enclosure.

Channel status symbols are located to the right of the port number in the OSCAR interface (see Figure 4). A green selection mark indicates that the SIP is online, a yellow mark indicates that the SIP is updating, and a red cross mark indicates that the SIP is offline. When the SIP is connected to a switch such as the Digital Access KVM switch, the channel status symbol is a combination of three circles. If a user is connected to the channel, a letter is displayed to the right of the channel status indicating which user is connected.

Dell KVM switches allow server selection changes without launching the OSCAR interface—a feature known as soft-switching. To use soft-switching, administrators can select the Setup button on the OSCAR interface. Then, on the Menu page, they can select "Slot" for the Display/Sort Key setting and set the Screen Delay Time to be one or more seconds (see Figure 5). The Screen Delay Time setting is ignored if the password option is enabled in the external KVM switch.

Enabling remote access capabilities with Digital Access KVM switches

Avocent Digital Access KVM switches give server administrators the convenience of remotely monitoring and controlling Dell PowerEdge server blades within a Dell Modular Server Enclosure. By cascading a Digital Access KVM switch to external KVM switches and connecting it to the overall KVM infrastructure, administrators can interconnect multiple servers and KVM switches across an enterprise network. Using Dell Remote Console Software and OSCAR, administrators can manage these servers efficiently from a single console.

Babu Chandrasekhar is a lead software engineer in the Dell Enterprise Server Group. Before joining Dell, he worked as a software engineer for Digital Equipment Corporation, Intel Corporation, and the Bhabha Atomic Research Centre. He has a B.S. in Computer Science and Engineering from the University of Kerala in India.

Jake Diner is a software engineer in the Dell Enterprise Systems Management Software organization. His interests include public speaking and wireless communication. Jake has a B.S. in Computer Science from Michigan State University.

FOR MORE INFORMATION

Dell blade servers:
www.dell.com/blades

Understanding Redundancy

in Dell PowerEdge Blade Servers

When administrators plan blade server deployment, they should carefully review redundant system configuration options. Depending on the data center's design and IT policies, blade server systems may require redundancy at different levels. Dell™ PowerEdge™ blade servers offer redundancy options for various functions—ranging from power redundancy to chassis-management redundancy.

BY NARAYAN DEVIREDDY AND SANJEEV S. SINGH

Related Categories:

Blade servers

Dell PowerEdge blade servers

Systems management

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

By consolidating servers, infrastructure components, and management within a single chassis, Dell PowerEdge blade servers can help achieve high efficiency in the data center and provide an optimized rack environment. Blade server deployment requires careful planning of data center resources such as power, networking, infrastructure fabric, cooling, and management access—and the blade server architecture differs from monolithic servers in its possible redundant configurations.

The system chassis used in today's monolithic servers hosts a single compute node with one or more CPUs, local storage, network, and other infrastructure components such as management access, power, and cooling. Redundancy options for monolithic servers range from cooling to network redundancy, thereby eliminating single points of failure. Several Dell PowerEdge servers offer redundant hot-pluggable cooling fans and power supplies, redundant memory with failover memory banks for memory mirroring, RAID controllers for redundant storage configuration, and dual-port integrated network interface cards (NICs) with network teaming software for redundant network configuration. When ordered with power, cooling, or memory redundancy options, Dell PowerEdge servers arrive fully configured and these options do not require any additional configuration at the deployment site. Network and storage redundancy options, however, require additional configuration using management applications.

The system chassis of Dell PowerEdge blade servers, known as the Dell Modular Server Enclosure, can host up to 10 compute nodes (server blades), plus shared infrastructure components. The common infrastructure components include cooling fans, power supplies, network switches, I/O modules, a KVM (keyboard, video, mouse) module, and chassis management in the form of the Dell Remote Access Controller/Modular Chassis (DRAC/MC). These components are shared through the midplane, which is passive to help ensure high reliability—that is, the midplane contains no active logic, just connectors and traces. Each server blade in a modular server chassis typically offers the same redundancy options as a monolithic server.

The primary reason for implementing redundant configurations within servers is to avoid single points of failure. Depending on the type of device and the design of the redundancy algorithm, setting up a redundant configuration can simply require installing two modules and powering up the system—or it may require using a management interface to configure the redundancy options. In a redundant setup, when one device fails, the second device assumes control and service continues uninterrupted. When a failover occurs, most devices typically transmit informational events that alert IT management about the redundancy failover. The Dell Modular Server Enclosure offers redundant configurations for power supplies, fans, I/O modules, network connections, and management modules.

Power supply redundancy

In the Dell Modular Server Enclosure, the DRAC/MC is responsible for system power budgeting and management. Power supply redundancy allows uninterrupted system operation in the event of failure of one or two power supplies in the chassis. This means that the DRAC/MC allows the blade server system to be powered up only if the required power is less than the available power based on the redundancy policy selection. The DRAC/MC lets the administrators select the redundancy configuration that best meets their requirements.

The following redundancy policies are available with DRAC/MC firmware version 1.3 (see Figure 1):

- **No redundancy:** In this mode, the total available power from the installed power supplies is used to power up the server blades and chassis components. Failure of one power supply may cause the chassis to power down based on the power consumption and available power at the time of failure.
- **3 + 1 redundancy:** In this mode, the power capacity of one power supply is kept in reserve while powering up the server blades; failure of any one power supply does not cause the chassis to power down.
- **2 + 2 redundancy:** In this mode, the capacity of two power supplies is kept in reserve while powering up the server blades; failure of any two power supplies does not cause the chassis to power down.

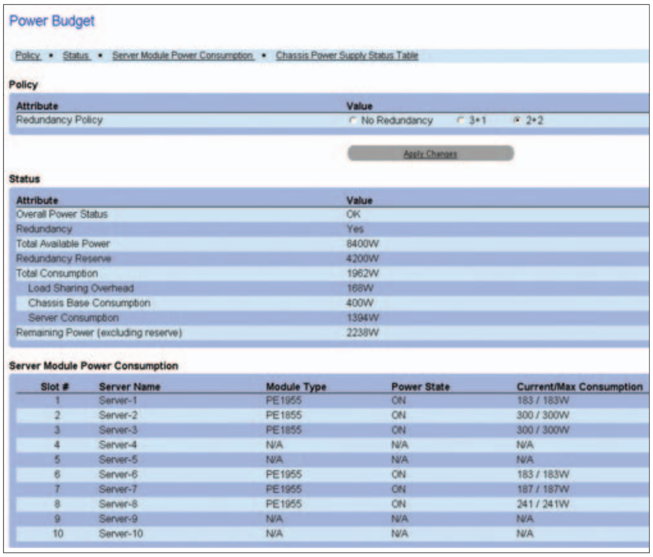

Power Budget

Policy • Status • Server Module Power Consumption • Chassis Power Supply Status Table

Policy

Attribute	Value
Redundancy Policy	○ No Redundancy ○ 3+1 ● 2+2

Status

Attribute	Value
Overall Power Status	OK
Redundancy	Yes
Total Available Power	8400W
Redundancy Reserve	4200W
Total Consumption	1962W
Load Sharing Overhead	168W
Chassis Base Consumption	400W
Server Consumption	1394W
Remaining Power (excluding reserve)	2238W

Server Module Power Consumption

Slot #	Server Name	Module Type	Power State	Current/Max Consumption
1	Server-1	PE1955	ON	183 / 183W
2	Server-2	PE1855	ON	300 / 300W
3	Server-3	PE1855	ON	300 / 300W
4	Server-4	N/A	N/A	N/A
5	Server-5	N/A	N/A	N/A
6	Server-6	PE1955	ON	183 / 183W
7	Server-7	PE1955	ON	187 / 187W
8	Server-8	PE1955	ON	241 / 241W
9	Server-9	N/A	N/A	N/A
10	Server-10	N/A	N/A	N/A

Figure 2. DRAC/MC Power Budget screen

Note: If administrators have already deployed a Dell Modular Server Enclosure with the DRAC/MC installed, they can upgrade the firmware to version 1.3. For more information about obtaining DRAC/MC firmware version 1.3, visit www.support.dell.com.

The desired power supply redundancy policy can be selected from the Power Budget page of the DRAC/MC Web interface (see Figure 2). The same selections can be made through the DRAC/MC command-line interface using the following command:

```
racadm config -g cfgChassisPower -o
  cfgChassisRedundancyPolicy value
```

Valid values for the `cfgChassisRedundancyPolicy` option are `0` (no redundancy), `1` (3 + 1 redundancy), and `2` (2 + 2 redundancy).

The redundancy policy selection is available only if four 2,100-watt power supplies are installed.[1] If a power supply is not installed, fails, or is removed, the redundancy policy selection is grayed out on the Power Budget page of the DRAC/MC Web interface.

Redundancy policy	Power supply wattage	Support for power supply redundancy	
		If all four power supplies are installed and working	**If one or more power supply fails**
3+1	1,200	Yes if only single-core Dell PowerEdge 1855 server blades are installed (If dual-core PowerEdge 1855 or PowerEdge 1955 server blades are installed, redundancy is not supported.)*	No
	2,100	Yes	No
2+2	1,200	No	No
	2,100	Yes	Yes if one power supply fails; no if two or more power supplies fail

**Dual-core Dell PowerEdge 1855 and Dell PowerEdge 1955 blade server systems do not support 1,200 watt power supplies*

Figure 1. Redundancy support for various power supply configurations in Dell PowerEdge blade servers

[1] The exception to this is that the 3+1 redundancy policy is supported if four 1,200-watt power supplies and only single-core Dell PowerEdge 1855 blades are installed.

AC versus DC redundancy

For AC redundancy, power supplies must be connected to different AC grids so that failure of one AC grid does not cause a total loss of power to the chassis. Thus, in a configuration with four 2,100-watt power supplies, two power supplies should be connected to one AC circuit, while the other two power supplies should be connected to a different AC circuit. If all four power supplies are connected to the same AC circuit, then the system has DC redundancy because it is protected against power supply failures but not against AC failure.

For maximum benefit, Dell best practices recommend that administrators use redundant AC power connections when selecting the 2 + 2 redundancy policy. This means that administrators should configure a separate AC connection for each pair of power supplies. In this configuration, the system can continue to function if one AC panel fails.

Fan module redundancy

The Dell Modular Server Enclosure provides two cooling fan modules. These are located in the middle of the rear of the chassis and provide redundant cooling to the chassis. Each fan module contains two fans that provide redundancy within the individual modules. In addition to these fan modules, the chassis's power supplies contain fans to cool the chassis. Administrators should install dummy power supplies in any empty power supply bay and not leave any bay unoccupied. Leaving a bay unoccupied affects cooling and can cause the blades to be throttled.

When a fan failure occurs in one of the modules, the status is communicated in two ways:

- The fan's fault indicator light turns amber.
- An entry is made in the DRAC/MC system event log, and if configured to do so, the system sends alerts to the appropriate management consoles and e-mail accounts.

The fans require no special configuration to be set up for redundancy. Administrators can monitor the fan status through the DRAC/MC interface. For more information, see the *Dell Remote Access Controller/ Modular Chassis User's Guide* at support.dell.com/ support/edocs/software/smdrac3/dracmc.

I/O module redundancy

The Dell Modular Server Enclosure supports a maximum of four I/O modules. I/O bays 1 and 2 can support only Ethernet-based I/O modules, whereas bays 3 and 4 can support any type of I/O module provided that I/O configuration rules are followed.[2]

Installing the I/O modules in pairs allows the modules to be configured redundantly. The redundancy is achieved through proper software configuration of the daughtercards. Installing a second Fibre Channel module in the Dell Modular Server Enclosure can create data-path redundancy and double performance by providing two paths for data to travel to and from the system. This redundant data-path configuration helps ensure server-to-storage connectivity if one of the data path fails.

Network connection redundancy

Each server blade has two embedded network LAN on Motherboards (LOMs) that have a dedicated circuit to the internal ports of the integrated switches or pass-through modules. The LOMs provide dedicated 1,000 Mbps full-duplex connections. LOM 1 on each server blade connects to an internal port of switch 1 or pass-through module 1, and LOM 2 on each server blade connects to the counterpart port of switch 2 or pass-through module 2 (see Figure 3).

Installing a second switch or pass-through module is optional. However, two installed switches or pass-through modules can enable additional connectivity or network redundancy and fault tolerance.

An important distinction between a blade server and other types of servers is that the connection between the LOM and the internal ports of the integrated I/O module (switch or pass-through) is hardwired through the midplane. This design enables the link between the LOM and the integrated I/O module to remain in a connected state—unless either a LOM or the I/O port fails. The link remains active even in the absence of a network connection between the external uplink ports on the integrated switch and the external network. Because of this, failures of cables or switch ports outside the enclosure do not trigger failover events in a network-teaming scenario in which an integrated switch is used. To overcome this limitation, administrators can use the `nic-redundancy` command available in Dell PowerConnect™ 5316M switch firmware version

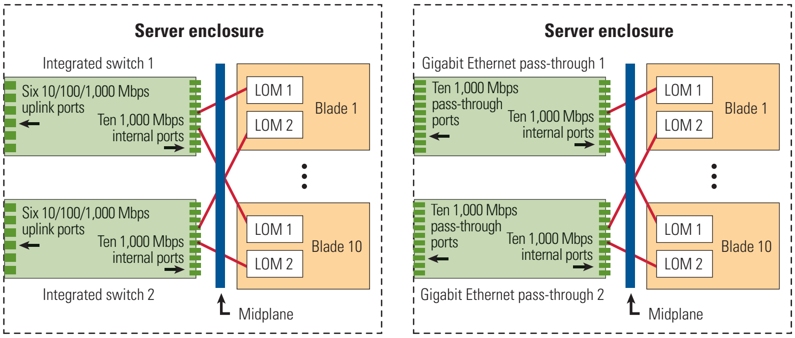
Server enclosure
Integrated switch 1
Six 10/100/1,000 Mbps uplink ports
Ten 1,000 Mbps internal ports
LOM 1
LOM 2
Blade 1
Six 10/100/1,000 Mbps uplink ports
Ten 1,000 Mbps internal ports
LOM 1
LOM 2
Blade 10
Integrated switch 2
Midplane
Server enclosure
Gigabit Ethernet pass-through 1
Ten 1,000 Mbps pass-through ports
Ten 1,000 Mbps internal ports
LOM 1
LOM 2
Blade 1
Ten 1,000 Mbps pass-through ports
Ten 1,000 Mbps internal ports
LOM 1
LOM 2
Blade 10
Gigabit Ethernet pass-through 2
Midplane

Figure 3. Dell PowerEdge blade server architecture showing integrated switches and pass-through modules

[2] For more information, see "A Technical Overview of the Dell Modular Server Enclosure and I/O Modules" by Michael Brundridge, Babu Chandrasekhar, Jyeh Gan, and Abhishek Mehta in *Dell Power Solutions,* August 2005; www.dell.com/downloads/global/power/ps3q05-20050163-Brundridge-OE.pdf.

1.0.0.35. This global configuration command enables the switch to support NIC teaming.

By default, NIC redundancy is disabled. Enabling the `nic-redundancy` option triggers a failover in teaming software if an external link on the integrated switch fails. The switch brings down the internal port links if there is no external port link and brings them back up once an external link is reestablished. The following command enables the `nic-redundancy` feature:

```
    console(config)# nic-redundancy
```

The following command disables the `nic-redundancy` feature:

```
    console(config)# no nic-redundancy
```

The scenario is different if a pass-through module is used. With a pass-through module, the link is in a connected state only if a network connection exists between the external ports on the pass-through module and the switch port outside the enclosure—as is true for a stand-alone server. Thus, for the pass-through module, the teaming software triggers a failover event if the LOM, pass-through port, cable, or external switch port fails.[3]

Management module redundancy

IT administrators can build chassis management redundancy into Dell PowerEdge blade servers by installing a redundant DRAC/MC in the chassis to serve as a standby for the primary DRAC/MC. The primary DRAC/MC actively monitors the chassis, while the standby DRAC/MC monitors the active signal from the primary DRAC/MC module. The standby DRAC/MC becomes the active primary DRAC/MC if a failure endures for more than five seconds.

Failures can be caused by any of the following conditions:

- The primary DRAC/MC network connection is broken. For example, the network cable has been disconnected or is broken.
- An administrator removes the primary DRAC/MC from the chassis.
- The primary DRAC/MC is rebooting or an administrator initiates a DRAC/MC reset.
- The primary DRAC/MC is nonresponsive and fails to exchange a heartbeat signal with the standby DRAC/MC.
- The firmware is being updated, causing a temporary failover. In this case, because the primary and standby DRAC/MCs use the same IP address, the console, Telnet, and user interface are rendered inactive. During a firmware update, the standby DRAC/MC monitors the chassis while the primary DRAC/MC updates its firmware. When the primary DRAC/MC completes its firmware update, the standby DRAC/MC continues its Trivial FTP (TFTP) update. The DRAC/MC network interface is unavailable until the firmware update is complete.

When a failover occurs, the primary and standby DRAC/MC share the same IP address and the same Media Access Control (MAC) addresses. At any point in time, only the current primary DRAC/MC provides chassis management and responds to an administrator's request from the network; the standby DRAC/MC does not perform any chassis management and does not respond to any administrator request from the network.

Designating the primary DRAC/MC

When the chassis powers up for the first time, the DRAC/MC that is located above power supply 1 becomes the primary module. The chassis orientation assumes that the administrator is viewing the chassis from the back. Viewed from the back, the primary DRAC/MC is the one located on the right side of the chassis during initial power-up.

High availability for blade server environments

The Dell Modular Server Enclosure provides multiple redundancy options for Dell PowerEdge blade servers. Configuring these options properly can help reduce downtime by increasing the availability and performance of Dell PowerEdge blade servers. Each component within the enclosure should be individually considered and configured to achieve optimal results.

Narayan Devireddy is a development manager in the Dell Enterprise Systems Management Software organization. He has 14 years of systems management product development experience. Before joining Dell, Narayan worked for Novell, Compaq, Cheyenne Software, and Computer Associates in different capacities. He has an M.S. in Computer Sciences from Alabama A&M University.

Sanjeev S. Singh is a software engineer in the Dell Enterprise Systems Management Software organization. Previously, he was a software engineer at Hewlett-Packard and NCR. He has a B.S. in Electrical Engineering, and he has an M.S. in Computer Engineering from North Carolina State University.

FOR MORE INFORMATION

Dell PowerEdge blade servers:
www.dell.com/blades

[3] For more information, see "Enhancing Network Availability and Performance on the Dell PowerEdge 1855 Blade Server Using Network Teaming" by Mike J. Roberts, Doug Wallingford, and Balaji Mittapalli in *Dell Power Solutions,* February 2005; www.dell.com/downloads/global/power/ps1q05-20040274-Roberts.pdf.

Streamlining Large-Scale Java Development Projects with SAP NetWeaver Application Server

SAP NetWeaver® Application Server provides tools for model-driven, service-oriented Java development as well as the SAP NetWeaver Development Infrastructure. This infrastructure is designed to reduce overall development costs and provide high reliability and flexibility for application development landscape management—including deployment into runtime systems and change-management processes.

BY WOLF HENGEVOSS AND CHRISTOPHER HEARN

Related Categories:

Application development

Application servers

Enterprise resource planning (ERP)

SAP

Scalable enterprise

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

Experienced application developers may have no trouble writing effective Java code or deploying their applications to a server. But working with large, geographically dispersed development teams on tasks such as application software updates can present challenges to even the most experienced developers.

When evaluating Java development tools, developers often focus on productivity for writing code. Although this is clearly an important feature, they should not underestimate the importance of managing the entire life cycle of an application—from setting up the development environment to managing software delivery and maintaining the application. The larger the project, the more critical it is from both efficiency and financial points of view.

Over the past 30 years, SAP has gathered in-depth experience in developing applications with large, geographically dispersed development teams. The delivery of software to customers, management of software upgrades, and synchronization of updates with customized code at the customer site are all issues that SAP has addressed. Now SAP is bringing the same high-quality approach to Java-based application development.

As a key component of the SAP NetWeaver composition platform, SAP NetWeaver Application Server (SAP NetWeaver AS) provides tools for model-driven, service-oriented Java development as well as an infrastructure that fully supports all phases of the application life cycle. This infrastructure—the SAP NetWeaver Development Infrastructure (NWDI)—is tightly integrated with the SAP NetWeaver Developer Studio tool delivered with SAP NetWeaver AS (see Figure 1). NWDI is designed to reduce overall development costs and provide outstanding reliability and flexibility for the application deployment and change-management processes.

Synchronized team development efforts

NWDI enables application developers to use Java in their local environment, while synchronizing the entire team's work via a central development environment. Within NWDI, the Design Time Repository (DTR) handles file versioning to help ensure that all developers are working from the same set of code. Developers access the central service via the SAP NetWeaver Developer Studio tool, check out files, produce new versions in the local file system, and check them back in after successful local testing. From the DTR, developers can compare versions and check version or revision history. During development, they can synchronize repository and local file systems at any time.

Using the DTR, developers can manage different versions of a development object in the same repository; multiple states of a software component (development and consolidation of several releases); and multiple users making modifications to the same development object (with conflict detection).

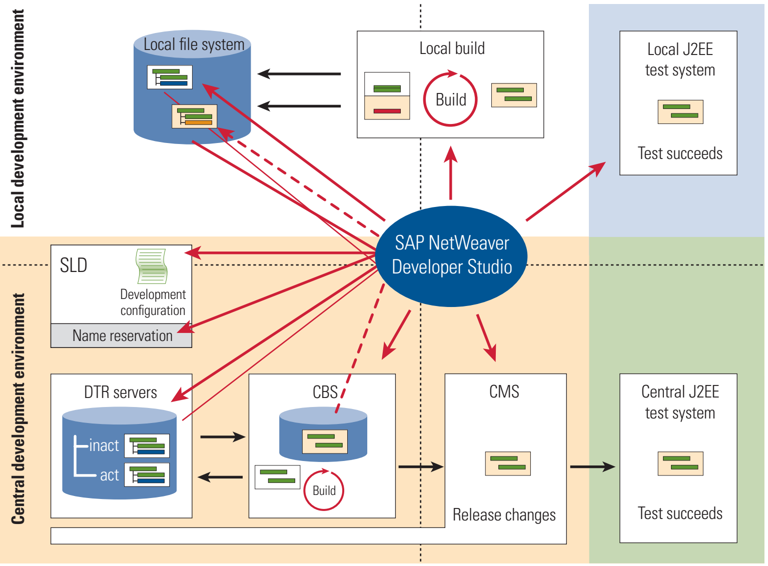

Figure 1. Developing Java-based applications within the SAP NetWeaver Development Infrastructure

Each state of software component development is represented in one workspace. The information about the state of a workspace can be propagated to other workspaces so the work of development teams can be synchronized using various instances of the DTR.

The DTR provides change-management features for distributed, multi-user development environments. As older versions of files are replaced with newer ones, the DTR handles this process centrally and keeps the version history. For more complex projects, additional capabilities are provided. For example, if modifications are occurring directly in end users' systems, various versions need to be developed in parallel and multiple DTRs must be in place. The version history is always transported along with the files for *global version history* of the DTR. As a result, versions created in parallel are detected automatically across repository boundaries. However, during code modifications, developers may not want an earlier version to be automatically overwritten by updates. In such cases, the DTR supports merging two versions, allowing developers to combine the advantages of the newer version with modifications.

Furthermore, to help reduce the maintenance effort as much as possible, developers need to integrate bug fixes from older releases into new releases from the maintenance cycle. The DTR supports this capability because changes are always deployed as an entire set of versions, instead of individually. This approach to change management means that the results are unaffected by the sequence in which the changes are applied. With its innovative approach to distributed development, the DTR is designed to enhance productivity and reduce development costs throughout the application life cycle.

Efficient component-based development

NWDI takes a component-based approach to development. A component hierarchy distinguishes multiple levels of granularity:

- **Development objects:** This is the finest level of granularity and includes tables, Java classes, and project files that are stored as versioned files.
- **Development components (DCs):** These are the units of the development and build process. DCs group development objects into larger components and define reusable components, working on the granularity of projects.
- **Software components:** These are the deployment and installation units that are made up of DCs. They represent larger building blocks of the product.
- **Product:** Although not part of the component model, products gather—and reuse—software components into applications that fulfill business requirements.

DCs are central to application development and follow a simple principle—only the parts declared to be public are visible to other DCs. That is, to use one DC from another DC, developers need to explicitly declare this usage. This explicit declaration of dependencies between DCs and precise relationships between objects allows encapsulation of functionality that leads to a fine-tuned, highly efficient build process in the Component Build Service.

Speed up build cycles with the Component Build Service

Like the DTR, the Component Build Service (CBS) is a Java 2 Platform, Enterprise Edition (J2EE) application that uses a database. It hosts all Java archives needed or produced during software development. For each software state, a *buildspace* is set up to contain these archives. Developers can trigger a central build of their components in the CBS at any time. Central builds apply only to modified DCs, along with any DCs that have dependencies with the changed archives.

This DC build approach allows developers to correct errors in small groups, dramatically helping to reduce bug-fix cycle times. A failed build process does not affect the build process of any other DC.

In addition, the CBS also provides J2EE cluster support for high-performance, automated build scripts for Java development, and automatic rebuilds of dependent DCs after changes to objects.

After a successful build, the CBS automatically makes the sources and archives available for use by other developers. Because all archives are centrally stored and up-to-date in the CBS, it is the source for consistently retrieving and updating used libraries in a local file system. Fast build cycles and a current build environment significantly help reduce development costs, time, and errors, especially for large projects.

Tools for managing the development environment

The development process starts with the definition of a product in the System Landscape Directory (SLD). From here, developers define which software components are used in the product and the dependencies between them.

This information is imported into the Change Management Service (CMS), where the environment for the developers is defined. Workspaces and buildspaces are created by defining a *track,* which is used to describe the development environment for one software component state. Administrators fill buildspaces with all the libraries required for a development configuration. The definition of a development configuration specifies access to project-specific sources and archives; this definition is imported as an XML file into the SAP NetWeaver Developer Studio.

After development and successful testing, developers release components to the CMS again. Here, the quality manager triggers the import into the consolidation system. After central testing, the final version is created by assembling all components and then approved for production use. For the next release, it is not necessary to physically set up a new system—developers simply define a new track. This allows for comprehensive control of the application life cycle from product definition to application patches.

Figure 2 shows how all transport steps are managed in the CMS Transport Studio—from check-in of required objects through imports during development and consolidation to deployment into the production system. Figure 3 shows how the landscape definition in the CMS is reflected in the SAP NetWeaver Developer Studio. Developers can view a software component and the DCs it contains.

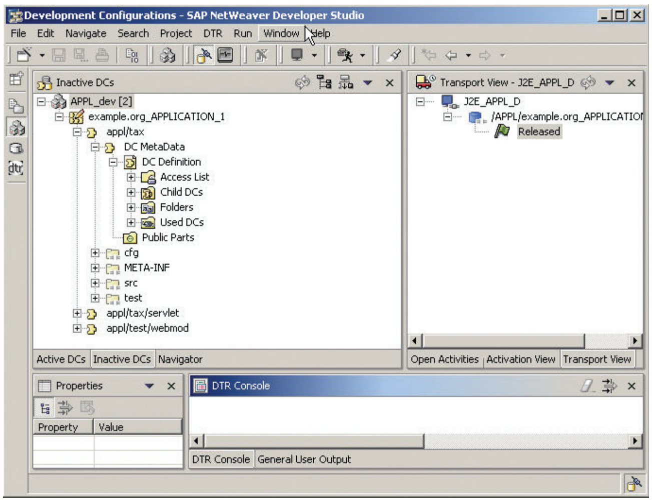

Figure 3. Viewing DCs from the SAP NetWeaver Developer Studio

Effective, efficient application development

SAP NetWeaver Development Infrastructure provides developers with a consistent development infrastructure centrally managed in the CMS. All sources are stored centrally in the DTR, which provides project-specific access to sources plus the reliability of a database. The DTR is tightly integrated with the CBS: For the central build process, the sources are retrieved from the DTR and the resulting archives are stored in the CBS centrally—helping to ensure that the latest versions of sources are used and that files are protected. Distributed versioning and concurrency control allow developers to manage large development projects taking place in different locations. The DTR's versioning capabilities help ensure that modifications are not overwritten during updates, but rather can be integrated into the new version. This is the basis of a safe modification process for Java-based applications.

In addition, NWDI allows developers to clearly define dependencies and encapsulate functions, facilitating the reuse and maintenance of development components. In short, SAP NetWeaver and its development infrastructure can help streamline Java development projects, enabling developers to work productively and efficiently.

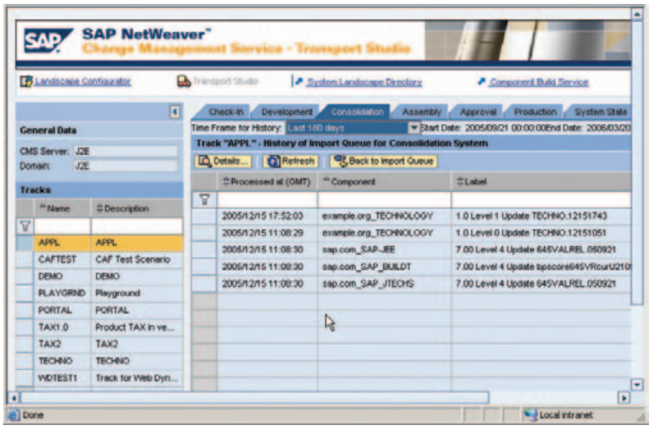

General Data

CMS Server: J2E
Domain: J2E

Time Frame for History: Last 180 days Start Date: 2005/09/21 00:00:00End Date: 2006/03/20

Tracks

Name	Description
APPL	APPL
CAFTEST	CAF Test Scenario
DEMO	DEMO
PLAYGRND	Playground
PORTAL	PORTAL
TAX1.0	Product TAX in ve...
TAX2	TAX2
TECHNO	TECHNO
WDTEST1	Track for Web Dyn...

Track "APPL" - History of Import Queue for Consolidation System

Processed at (GMT)	Component	Label
2005/12/15 17:52:03	example.org_TECHNOLOGY	1.0 Level 1 Update TECHNO.12151743
2005/12/15 11:08:29	example.org_TECHNOLOGY	1.0 Level 0 Update TECHNO.12151051
2005/12/15 11:08:30	sap.com_SAP-JEE	7.00 Level 4 Update 645VALREL.050921
2005/12/15 11:08:30	sap.com_SAP_BUILDT	7.00 Level 4 Update bpscore645VRcurU210
2005/12/15 11:08:30	sap.com_SAP_JTECHS	7.00 Level 4 Update 645VALREL.050921

Figure 2. Managing transport steps in the CMS Transport Studio

Wolf Hengevoss is a member of the SAP NetWeaver product management team and focuses on the roll-out of NWDI.

Christopher Hearn is the product marketing director for SAP NetWeaver. Educated at Oxford University in England, Chris has 20 years of experience in a variety of IT roles.

FOR MORE INFORMATION

SAP Developer Network–SAP NetWeaver Application Server knowledge center and forums for Java:
www.sdn.sap.com/irj/sdn/developerareas/java

Exploring the Distributed File System

in Microsoft Windows Server 2003 R2

The Distributed File System (DFS) in Microsoft® Windows Server™ 2003 Release 2 (R2) introduces DFS Namespaces and DFS Replication. These two enhancements are designed to help administrators manage distributed file server resources efficiently while enabling fast, fault-tolerant access with low-bandwidth replication.

BY MIN-JOHN LEE AND MAHESH VELLORE

Related Categories:

Microsoft Windows Server 2003

System deployment

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

Microsoft Windows Server 2003 Release 2 (R2) is the first release update for Windows Server 2003, and its features are designed to integrate seamlessly into the Windows Server 2003 environment. At the core of Windows Server 2003 R2 is the Distributed File System (DFS), which has been enhanced with the goal of providing increased performance and improved availability while reducing issues commonly linked to branch-office server deployments.

For example, users may have trouble locating the files they need on distributed file servers. In addition, administrators often contend with low-bandwidth connections for providing file replication and file availability over wide area networks when servers at the central or branch offices fail. To help resolve such concerns, Windows Server 2003 R2 provides DFS Namespaces and DFS Replication, which together enable fault-tolerant access to distributed files with low-bandwidth replication:

- **DFS Namespaces:** Formerly known as the Distributed File System, the DFS Namespaces service allows administrators to group shared folders that are located on different servers. Folders are presented as a virtual tree called a namespace, so users no longer have to remember physical file locations.
- **DFS Replication:** A follow-up to the File Replication Service that was introduced in the Microsoft Windows® 2000 OS, the DFS Replication service is designed to address low network bandwidth using the remote differential compression (RDC) algorithm. RDC enables DFS Replication to

transfer only the changes that have been made since the last file update. Another feature, cross-file RDC, identifies files that are similar to the one being replicated. DFS Replication can then use portions of those similar files to replicate the file, helping reduce the amount of data transferred over the network.

Planning a DFS deployment

Administrators must consider a wide variety of factors when planning a DFS deployment. Major concerns include deployment of the Microsoft Active Directory® service, connection availability between branch offices, bandwidth availability between branch offices, file-share access in heterogeneous network environments, cost-effective backup procedures, capacity planning and scalability requirements, regional failover, and file types. By weighing alternatives carefully, administrators can determine a suitable deployment strategy for their particular enterprise requirements.

Active Directory considerations. Some DFS features exist only in the Active Directory environment. DFS Namespaces can be stand-alone or domain based. Stand-alone DFS Namespaces should be selected if an organization does not use the Active Directory service or if it needs to host more than 5,000 folders within a single namespace. Because of performance and directory service synchronization requirements, Microsoft best practices recommend using domain-based DFS Namespaces to host up to 5,000 folders. If using a server cluster for high availability, an organization must choose stand-alone DFS Namespaces.

By weighing alternatives carefully, administrators can determine a suitable deployment strategy for their particular enterprise requirements.

The DFS Replication service must be deployed in an Active Directory environment because the RDC technology can function only in an Active Directory domain. Active Directory enables DFS Replication to store configuration objects and to delegate user rights precisely.

Connection availability between branch offices. In an environment where the network connectivity could be either unavailable due to maintenance or unreliable due to the quality of the Internet service provider, DFS allows administrators to configure multiple target file servers for a given namespace folder. In this way, administrators can configure target priority so that data access requests fail over to a desired next-available target when the primary target cannot be reached. This feature enables zero-downtime server maintenance without interrupting data availability. In addition, once the primary target comes back online, data access requests can automatically route back to the primary DFS server by enabling client failback.[1]

Bandwidth availability between branch offices. In many circumstances, branch offices are connected to the enterprise's hub file server using on-demand or permanent virtual private network (VPN) connections over the Internet. In such network environments, data synchronization incurs a certain minimum latency due to network bandwidth and availability. In Windows Server 2003 R2, the DFS Replication service's RDC algorithm detects changes made by users in the replicated files and sends only changes that have been made since the last update. This approach consumes considerably less network bandwidth than the full file replication used in previous versions of DFS, and the reduced network overhead in turn helps improve the reliability of file replication in a low-bandwidth environment.

File-share access in heterogeneous network environments. For networks supporting both Windows and Linux® or UNIX® operating systems, Windows Server 2003 R2 provides Network File System (NFS) services and a user-mapping service that allows non-Windows platforms to access file shares on systems running Windows Server 2003. The combination of the NFS Client and NFS Server components in Windows Server 2003 R2 with DFS Namespaces and DFS Replication allows Linux- and UNIX-based systems to take advantage of the RDC algorithm and benefit from the high availability these services are designed to provide.

Cost-effective backup procedures. A comprehensive DFS implementation should include backup considerations. In a distributed environment, file updates are scattered across multiple office locations. Careful planning of DFS Replication procedures can help reduce staff requirements at branch locations. When data is replicated to a central DFS server, IT staff is needed only where the enterprise's hub file server is located to perform administrative data backup tasks. Windows Server 2003 R2 DFS Replication allows organizations to establish replication groups and replication topology and to set up filters to help prevent synchronizing unnecessary content such as MP3 files.

Capacity planning and scalability requirements. When planning for storage services, an organization's two major tasks are meeting immediate storage requirements and planning for future growth. Because Windows Server 2003 R2 DFS allows organizations to configure multiple target servers for a given folder, administrators can upgrade the storage capacity of a major DFS namespace server while data access requests fail over to the secondary namespace

[1] To enable support for client failback on a Microsoft Windows XP system, administrators must install Windows XP Service Pack 2 or later as well as the hot fix described in Microsoft Knowledge Base article KB898900. For more information, visit www.microsoft.com/downloads/details.aspx?FamilyID=7d3f51e3-2d33-48c4-8b5f-fe2345b0a35e&DisplayLang=en.

server. As a result, storage capacity can be resized quickly and flexibly in response to immediate business requirements, and administrators can scale enterprise storage capacity in cost-effective increments as future business needs materialize.

Regional failover. DFS in Windows Server 2003 R2 can help organizations provide data redundancy in different geographic areas—for example, replicating to several remote locations as well as to the enterprise's hub file server to maintain high availability and safeguard business continuity despite regional disasters. The RDC algorithm enables data on DFS namespace servers to be synchronized reliably between remote locations, even in low-bandwidth network environments.

File types. Windows Server 2003 R2 DFS is designed to work effectively with file types that are not affected by latency due to network bandwidth and availability issues. For example, file types that require real-time data synchronization, such as transaction-oriented databases, are not particularly well suited for DFS deployments.

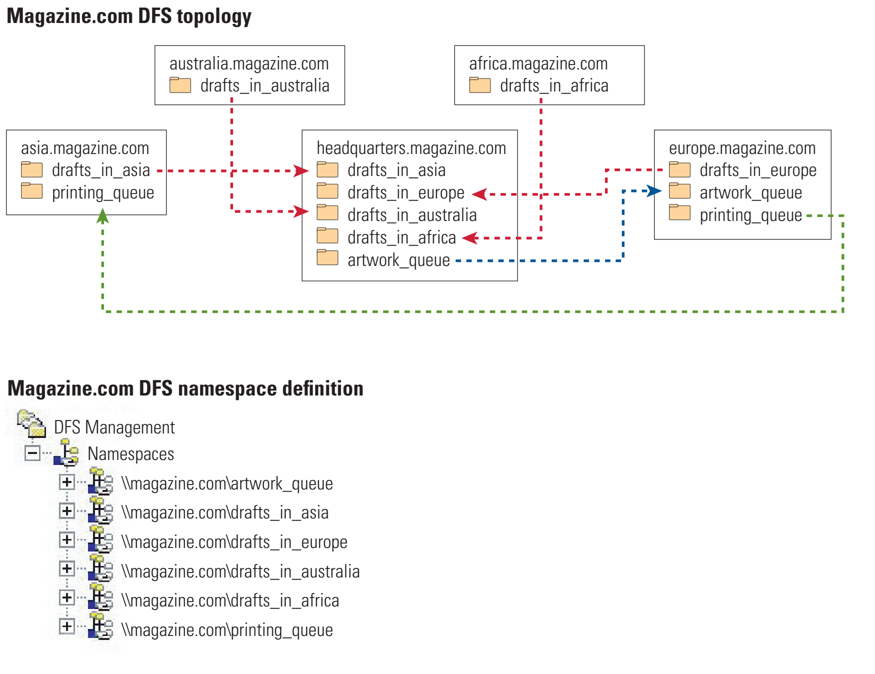

Figure 1. Global document workflow process for example scenarios

Customizing DFS deployments

The three DFS deployment scenarios described in this section demonstrate common branch-office requirements. Each scenario involves a global magazine company that has writers located around the world, article reviewers at the company's headquarters in North America, an artwork team in Europe, and a printing team in East Asia. DFS servers are deployed on each continent to transfer data around the world.

Scenario 1: Implementing business processes

The magazine's remote DFS servers can collect article drafts from writers around the world. A writer in Asia uses a notebook computer to establish a connection to the company's network and maps the DFS namespace—\\magazine.com\drafts_in_asia—in Microsoft Windows Explorer. The writer saves draft articles in the namespace as soon as they are ready for review. On the back end, the draft article is actually saved in a share folder hosted by a regional DFS server called \\asia.magazine.com\drafts_in_asia. This server is defined as a target folder in the \\magazine.com\drafts_in_asia namespace. By doing this, the writer does not have to remember all the DFS server names—only a company domain name. This setup can be duplicated on each continent.

At the North American headquarters, a reviewer uses Windows Explorer to map a drive to the same \\magazine.com\drafts_in_asia namespace. When this connection is established, the DFS Namespaces service on the back end actually establishes a connection to a local target folder hosted in a DFS server located at headquarters. The local target folder resides at \\headquarters.magazine.com\drafts_in_asia. The reviewer can view articles submitted by writers from Asia because a DFS Replication group is defined such that all files in \\asia.magazine.com\drafts_in_asia are replicated to \\headquarters.magazine.com\drafts_in_asia. Organizations can duplicate such a setup to implement a sophisticated document workflow process—for example, a process for legal review.

Once a reviewer approves an article, the article is moved to another namespace called \\magazine.com\artwork_queue. This namespace has target folders at \\headquarters.magazine.com\artwork_queue and \\europe.magazine.com\artwork_queue. The artwork team in Europe maps a drive to the same \\magazine.com\artwork_queue namespace and can view the same content as reviewers by using a DFS Replication group between the two target folders. Once the artwork for an article is completed, the artwork team moves the final version of the article to the \\magazine.com\printing_queue namespace. With a replication group set up between the Europe branch server and the Asia branch server, the printing team in Asia can obtain the final version of the article. Figure 1 shows the DFS topology and namespace definition for this global document workflow process.

Scenario 2: Collecting data for central backup

DFS Replication can be configured to perform data collection from different branch-office locations so that the backup is centralized on the enterprise's hub file server. The IT department at the magazine's headquarters can set up DFS Replication groups so that the data in the shared folders—drafts, artwork_queue, and printing_queue—is replicated from the server at each branch office (data source) to the hub file server (data destination). This approach helps reduce operational costs at branch offices by consolidating hardware for backup and related management tasks. The data collection topology also allows data to be replicated from the hub server to each branch server if data recovery at a branch server is needed. Either setting up shared folder permissions or disabling the hub-to-branch connection in the replication group can prevent data from changing on the hub file server.

> Replicating data to a branch-office server enables the branch office to provide fault-tolerant access by establishing failover to an office that is located in a different geographic region.

Scenario 3: Providing regional failover for high availability

The DFS Management snap-in within the Microsoft Management Console (MMC) can be used to configure DFS Replication–specific tasks such as creating replication groups, specifying replication group schedules, managing replication for specific connections, and managing replication filters. Implementing DFS Replication with DFS Namespaces can help achieve high data availability. If the branch server asia.magazine.com reaches its storage capacity, branch clients can fail over to the DFS server at headquarters while the Asia branch server is taken down to let administrators add storage space. With a DFS Namespaces enhancement known as client failback, branch clients can fail back to the branch server after the branch server is back online. Both failover and failback are transparent to end users, thus providing high data availability and business continuity. *Note:* DFS Replication for data collection is not recommended for database files or any other type of file that is held open for long periods of time because such files are replicated only after they are closed.

Enterprises with offices scattered across the globe must provide access to files that may be located at any remote office location. Replicating data to a branch-office server enables the branch office to provide fault-tolerant access by establishing failover to an office that is located in a different geographic region. This approach also helps provide an additional level of data protection if disaster strikes and distant offices are unaffected.

Extending DFS deployments with Storage Manager for SANs

A major requirement for branch offices is that storage be easily configured, provisioned, and managed. Windows Server 2003 R2 introduces the Storage Manager for Storage Area Networks (SANs) component, which helps simplify storage provisioning on external storage arrays. Using Storage Manager for SANs, administrators can host DFS replicated files on Fibre Channel and Internet SCSI (iSCSI) storage subsystems that support Virtual Disk Service (VDS).[2] This approach enables smooth integration into the Windows-based server, because Storage Manager for SANs is controlled through an MMC snap-in.

Enabling fast, fault-tolerant file access across the enterprise

Microsoft Windows Server 2003 R2 has enhanced the Distributed File System to address emerging needs for heightened performance, availability, and reliability of far-flung branch-office file servers while minimizing issues that are commonly linked to branch-office deployments, such as management overhead and limited connectivity. Two key capabilities introduced in Windows Server 2003 R2—DFS Namespaces and DFS Replication—are designed to simplify the management of distributed file server resources while providing easy and fault-tolerant access for remote branch-office locations. The three real-world scenarios described in this article show how DFS technologies enable enterprises to implement processes that help improve productivity while helping ensure high data availability and business continuity.

Min-John Lee is a software engineering consultant in the Server Operating Systems Engineering department in the Dell Product Group–Enterprise Software Development. Min-John has an M.S. in Electrical and Computer Engineering from Northwestern University.

Mahesh Vellore is a lead engineer in the Server Operating Systems Engineering Group at the Dell Bangalore Development Center. He has a diploma in Electronics Engineering and has been at Dell for more than two years.

FOR MORE INFORMATION

Microsoft Windows Server 2003 R2:
www.microsoft.com/windowsserver2003

[2] For more information about Storage Manager for SANs and VDS, visit www.microsoft.com/windowsserver2003/R2/storage/default.mspx.

Using Yosemite Backup Virtual Tape Libraries

to Accelerate Backup and Restore Operations

Virtual tape libraries combine the best of both worlds, using disk as the primary backup medium and tape for long-term archiving. Using backup software from Yosemite Technologies, IT organizations can combine the strengths of Dell™ PowerEdge™ servers and Dell PowerVault™ disk and tape hardware to create virtual tape libraries designed to improve the performance, reliability, flexibility, and scalability of network backup and restore operations.

BY NEIL MacLEAN, MIKE WEIMANN, JAKE SWIM, AND JIM LEE

Related Categories:

Backup

Dell PowerVault storage

Disaster recovery

Storage

Yosemite Technologies

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

Tape-based backups have commonly been used in data centers, but the problems with tape as a backup medium are well known: spotty reliability, high operational costs, and slow performance. For example, typical problems that plague tape-based backups include medium or cartridge failure during backups or restores; unreported failures or incomplete backups; misplaced or mislabeled media; long backup windows and verifications; and slow serial access times.

Meanwhile, the cost of disk-based storage has dramatically decreased, and in many scenarios backing up and restoring data to and from hard disks can be achieved at a much faster rate than possible with tape technologies, enabling shorter backup windows with disk-based storage. Growing dissatisfaction with traditional tape-based backups, combined with the viability of disk as the primary target for backups, has made disk-based backups a popular data protection strategy.

Although tape is still preferable for long-term archiving because of its portability and off-line security, many enterprise IT organizations are finding that a combination of the two technologies helps provide excellent, cost-effective data protection (see Figure 1). Yosemite Backup™ storage software can be used in conjunction with Dell disk and tape hardware to create such hybrid backup environments.

Enterprise-class performance for backups and restores

Disk-based backups can address some of the problems of tape-based data storage, but they can still leave critical data at risk. However, a disk-based backup can be effective when implemented in a multi-tiered storage hierarchy—for example, a disk-to-disk-to-tape (D2D2T) backup architecture with an automated tape loader or library. When implemented with tape and disk hardware,

Feature	Tape	Disk	Disk and tape
Administration	Good	Best	Better
Efficiency	Good	Better	Best
Expense	Best	Better	Good
Flexibility	Good	Better	Best
Performance	Good	Best	Better
Portability	Best	Good	Better
Reliability	Good	Better	Best
Scalability	Good	Better	Best

Figure 1. Comparison of typical backup and recovery media

Yosemite Backup software offers several features to enhance enterprise data storage.

Virtual tape libraries and Disk-to-Disk-to-Any option

One promising technology that is quickly gaining acceptance is the virtual tape library (VTL). This technology employs a disk array that functions as a virtual tape device or library, emulating the key characteristics of tape backup while providing the benefits and convenience of disk backup. External VTL appliances typically rely on proprietary hardware and software to implement the VTL function, which can increase costs and, in some cases, limit scalability. In contrast, Yosemite Backup software embeds the VTL function as an integrated feature, thereby providing the benefits of an external VTL appliance while avoiding drawbacks inherent in proprietary VTL technology.

Yosemite Backup software embeds the VTL function as an integrated feature, thereby providing the benefits of an external VTL appliance while avoiding drawbacks inherent in proprietary VTL technology.

The disk-to-disk technology in Yosemite Backup software offers the flexibility of two levels of disk-based backup architecture. The first level is disk-to-disk backup, in which the backup process uses designated disk volumes as the backup target—offering the performance, reliability, and easy scalability of disk-based backups. The second level is Yosemite Backup's Disk-to-Disk-to-Any (D2D2N$^{e\text{TM}}$) option. D2D2N^{e} allows for true management of data through a hierarchy of storage devices (see Figure 2). For example, a backup is performed from the source disk to a target disk (configured as a VTL). D2D2N^{e} then manages the movement of the data from the secondary disk to tertiary storage (magnetic disk, tape, or optical disc) using predefined criteria, such as the VTL reaching a capacity limit or an established time schedule.

D2D2N^{e} provides the flexibility to create multiple storage hierarchies that can consist of other virtual libraries. Data can be moved to other disk systems, and from there to tape for off-site archiving. Because Yosemite Backup is designed to provide true disk-to-disk-to-any functionality, it is not restricted to a rigid D2D2T configuration. In addition, all data is tracked through the Yosemite Backup catalog, freeing backup administrators from manually copying data to tape or keeping track of where data is going.

D2D2N^{e} is designed to maximize the capabilities of the storage infrastructure cost-effectively, enabling administrators to create VTLs from underutilized disk capacity in similar or heterogeneous environments. The storage folders that are created for use by D2D2N^{e} can be spread across any combination of storage and on any supported platform. As a result, storage can be pooled from Novell® NetWare®, Linux®, and Microsoft® Windows® environments to create the VTL. Yosemite Backup does not restrict IT environments to one platform, nor does it require proprietary hardware or additional third-party software to function as a VTL. In addition, Yosemite Backup software can be managed from a central console and is designed to restore data from any target medium in the hierarchy without staging to a VTL first.

As a result, Yosemite Backup's integrated VTL approach is designed to provide numerous advantages over traditional disk-only or tape-only backups, including the following benefits:

- A single, simple console for operation and administration
- Fast and highly reliable backups and restores with multiple drives or volumes

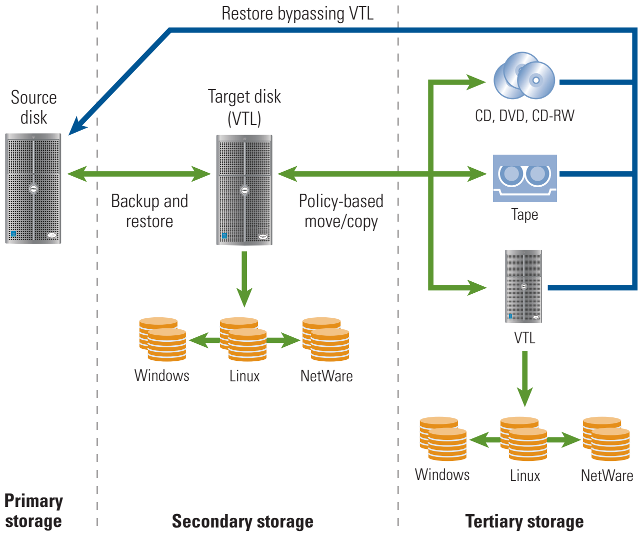

Figure 2. Hierarchy of storage devices allowed by the Yosemite Backup Disk-to-Disk-to-Any option

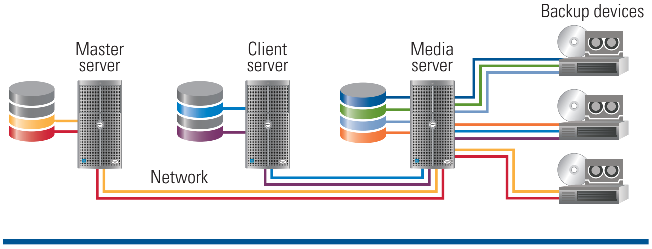

Figure 3. Intelligent data streams created by Yosemite Backup Self-Tuning Logic

- Simultaneous backup and restore operations
- Utilization and pooling of storage from multiple, heterogeneous systems
- Direct restores to the point of origin from any target—magnetic disk, tape, or optical disc
- "Spill-over" pools that allow heterogeneous servers to utilize disk volumes and thereby help ensure that backups complete successfully by not running out of disk space
- Creation of multiple backup hierarchies, such as D2D2D2T and D2D2D2O
- Management of multiple copies of the same backup on different media
- Ability to maintain established tape-based policies and processes, including rotation-type backup schedules

Self-Tuning Logic technology

The primary advantage of a disk-based backup is its potential to dramatically reduce the backup window. Yosemite Backup software maximizes this advantage by offering intelligent data streaming, which helps minimize the impact of backups on systems. The Self-Tuning Logic™ technology included in Yosemite Backup is designed to create an unlimited number of data streams to and from the backup source and target system for fast backup and restore operations (see Figure 3).

Self-Tuning Logic is designed to automatically configure the backup streams by detecting how the backup sources are configured, helping provide optimal performance. For example, if a physical disk is configured as multiple logical volumes, Self-Tuning Logic detects this configuration and sets up the backup appropriately.

Flexible, fast restores

In tape-based backup environments, storage software tools often require a restore from tape to a secondary disk and then back to the primary disk, effectively staging the data administrators wish to restore and increasing the restore time. In contrast, Yosemite Backup gives administrators the option of restoring directly from any backup target (magnetic disk, tape, or optical disc) if available on that medium.

Restore performance can be as important as backup performance. Because Yosemite Backup tracks all data in a central catalog, administrators can restore a single file or even a particular version of a single file, rather than a complete disk image—enabling fast, simple data restores that require minimal administrative time.

Dell hardware and Yosemite Backup: Multi-tiered storage architecture

Dell PowerVault and PowerEdge hardware enable enterprise IT organizations to easily add disk-based backup capabilities to their existing backup infrastructures. When combined with Yosemite Backup software's integrated VTL and D2D2N^e option, Dell hardware can be implemented in a plug-and-play manner to minimize the disruption to existing backup operations.

For example, consider an enterprise scenario in which an IT organization must back up data from eight servers and 130 client PCs. Microsoft SQL Server™ database software and Microsoft Exchange run on two of the eight servers. Data from all the servers and clients is backed up to a Dell PowerVault 122T DLT VS80 autoloader, which provides up to eight media slots (see Figure 4). A full backup is performed every Friday, and incremental backups are performed Saturday through Thursday. One month's worth of backups is stored at any given time. Data changes by about 20 percent each day. The backup window is eight hours, but the IT organization has difficulty backing up approximately 400 GB of data within this timeframe.

To resolve this problem, the IT organization in this scenario could upgrade the tape devices from digital linear tape (DLT) to Ultrium 3 Linear Tape-Open (LTO-3) technology. Such an upgrade would increase the backup speed because LTO-3 has a specified backup rate of 288 GB/hour compared to the specified backup rate of approximately 10.8 GB/hour for DLT VS80. However, the LTO-3 enhancement still would not overcome issues inherent in tape-only environments such as inefficiency, inflexibility, and limited scalability.

To help overcome such limitations, the organization could implement Dell disk hardware, configuring the environment for D2D2T backups. For instance, a Dell Storage Server (PowerEdge 830)

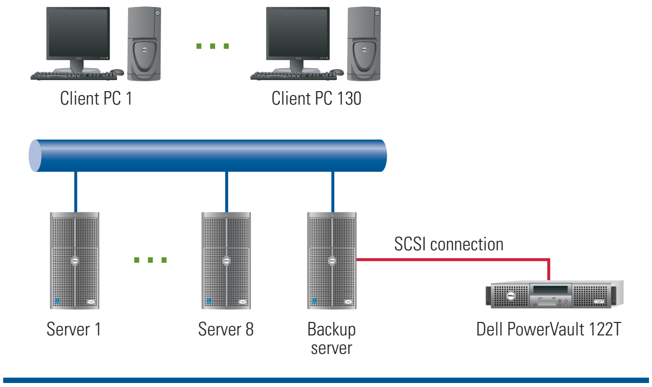

Figure 4. Example storage infrastructure configured for tape-only backups

running Microsoft Windows Storage Server 2003 Release 2 (R2) can provide a network attached storage (NAS) platform. In this scenario, the Dell Storage Server can act as both the Yosemite Backup application server and the disk storage platform, using Yosemite Backup software's integrated VTL to drive the D2D2T operation. Advantages of this implementation include the following:

- A backup server is not required when the Yosemite Backup VTL is integrated with the Dell Storage Server.
- Backups and restores can be completed more quickly and more reliably than possible with tape-only backups.
- Restores can occur directly from tape or disk back to the point of origin.
- The Yosemite Backup VTL can utilize available disk capacity from Windows, Linux, and NetWare servers, providing heterogeneous "spill-over" capability.
- Concurrent backups and restores can take place using the Yosemite Backup VTL architecture, allowing for multiple virtual drives.

Figure 5 shows the architecture for the example scenario with the Dell Storage Server added for D2D2T backups. This configuration allows the IT organization to keep the existing tape hardware while helping to improve backup and restore performance. Cost for the added hardware can be offset by enhanced scalability, performance, reliability, and availability.

For further scalability, a Dell PowerVault MD1000 system can be added (see Figure 6). This Serial Attached SCSI (SAS)/Serial ATA (SATA) device is attached behind the Dell server and is designed to provide up to 3.75 TB of additional raw disk storage capacity when fully populated with fifteen 250 GB SATA II hard disk drives. Furthermore, up to two more PowerVault MD1000 systems can be attached for up to 11.25 TB of additional raw disk storage.

The PowerVault 122T tape autoloader also can be replaced with the PowerVault 124T tape autoloader to provide a capacity and performance boost in the Figure 4, 5, and 6 scenarios. This tape autoloader provides up to 16 media slots and either an LTO-2-L drive or an LTO-3 drive, which are designed to achieve a native backup rate of 86.4 GB/hour or 288 GB/hour, respectively. This configuration enables disk-to-disk and disk-to-tape backups to complete well within the eight-hour backup window.

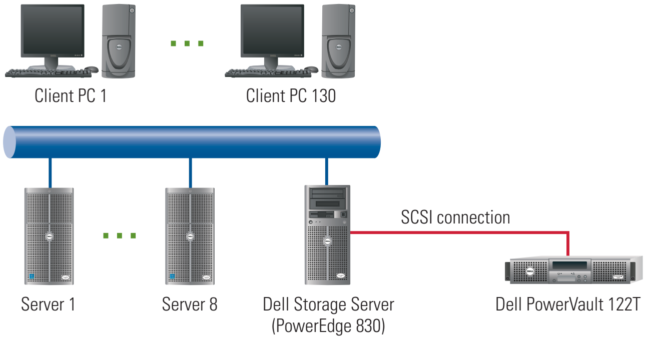

Figure 5. Example storage infrastructure configured for D2D2T backups

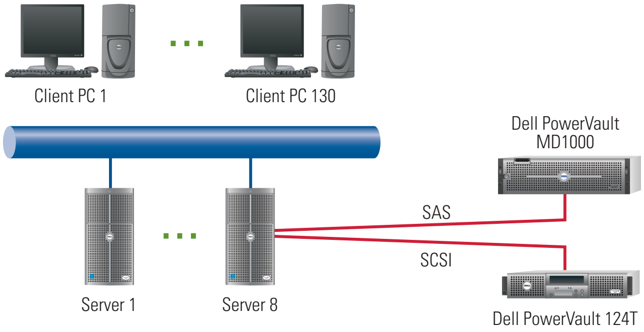

Figure 6. Example storage infrastructure configured to scale flexibly for additional disk capacity

Enhanced storage performance and reliability

Enterprise IT organizations can significantly enhance backup and recovery performance by migrating to D2D2T configurations such as those provided by Yosemite Backup software and Dell PowerEdge and PowerVault disk and tape systems. IT organizations can help reduce costs by integrating a disk system as a VTL—for example, deploying a Dell Storage Server to serve as a VTL can cost less than upgrading existing tape autoloaders and media. This savings can allow IT organizations to take advantage of investments in existing hardware and media while gaining the improvements in performance, reliability, scalability, and flexibility provided by disk-based backups.

Neil MacLean is a storage architect within the Office of the CTO at Dell.

Mike Weimann is a Dell global account manager at Yosemite Technologies.

Jake Swim is a Dell sales engineer at Yosemite Technologies.

Jim Lee is director of marketing at Yosemite Technologies.

FOR MORE INFORMATION

Dell storage:
www.dell.com/storage

Dell and Yosemite:
www.dell.com/yosemite

CommVault Galaxy Enhances Data Protection

for VMware ESX Server Virtual Machines

CommVault® Galaxy® software is a powerful suite of Unified Data Management™ capabilities optimized for large-scale, enterprise deployment. This suite includes capabilities for data backup, recovery, replication, snapshots, and archiving. These capabilities can be used to help manage data in VMware® ESX Server™ virtualized environments.

BY KELLY HARRIMAN-POLANSKI

Related Categories:

CommVault

Storage

Storage software

Virtualization

VMware

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

Backup, restore, and disaster recovery are crucial considerations when managing data center operations. This is no less true for virtualized servers. Traditional backup and recovery methods can be used to protect data in virtualized environments. However, IT organizations should also consider using additional methods for protecting data on virtual machines (VMs) such as those provided by CommVault Galaxy software.

Data protection needs of VMware-based environments

VMware software enables enterprise IT organizations to virtualize their computing, storage, and networking systems and manage them centrally through the creation of enterprise-class VMs. These virtual environments can help increase physical server utilization, performance, and system uptime cost-effectively. When deployed on Dell™ PowerEdge™ servers, VMware ESX Server virtual infrastructure software enables IT administrators to create multiple VMs on a single physical server, each of which can run a separate OS and applications without interfering with other VMs on the physical server. VMware VirtualCenter management software provides a central point of control for a data center's virtual computing resources. This highly scalable, cost-effective VMware virtualization platform is designed to provide advanced resource management capabilities for today's enterprise IT environments.

CommVault software can provide additional data protection and management capabilities in a VMware-based environment. By integrating data backup and recovery with replication, snapshots, and archiving—all of which can be managed from a single console—CommVault software offers cost-effective, scalable, and easy-to-use data management with a range of options for reliably protecting VMware-based environments. CommVault Galaxy software can help protect and manage data residing on VMware-based systems by providing the following capabilities:

- Object-level backup and recovery of file system data residing in guest-hosted VM environments
- Granular backup and recovery of application data residing in guest-hosted VM environments
- Deployment options for placement of the backup server components—within a VM and on a separate server
- Crash-consistent protection of VMs

Disk structure of VMware ESX Server

VMware ESX Server is a data center–class virtual infrastructure suitable for mission-critical environments. It boots from a version of the Linux 2.4 kernel on the x86 chipset. Because ESX Server runs directly on the server hardware, it is highly scalable and capable of running many more VMs than VMware GSX Server™ software, which runs on top of a host OS. The ESX Server system can guest-host VM environments running various operating systems, including Microsoft® Windows®, Red Hat® Enterprise Linux®, Novell® NetWare®, Novell SUSE® Linux, Sun Solaris, and other OS platforms (see Figure 1).

VMware provides its own file system—VMware File System (VMFS)—for storage of VMs. VMFS is optimized to store large files containing virtual disk images and memory images of suspended VMs. VMFS-2, used by ESX Server 2, may exceed 2 GB in size and can span multiple disk partitions across one or multiple logical units (LUNs) or physical disks.

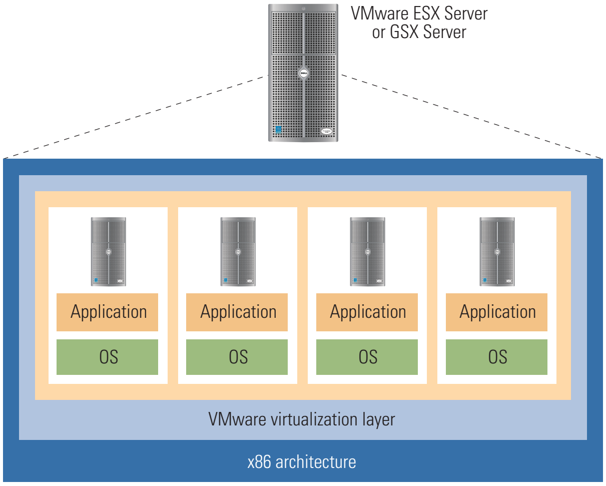

Figure 1. ESX Server environment with guest-hosted operating systems running as VMs

ESX Server captures and preserves the data residing on VMs in VMFS disk files. These files, which have .dsk or .vmdk file extensions, can be used to back up the entire VM. While the VM is in operation, these files remain open and in use. VMware provides various methods for gaining access to these files so that they are readily available for backup protection. One method for protecting the virtual disk file is to stop the VM while the file is being backed up. VMware also provides a procedure for releasing the disk file by first creating a redo file (log) to capture any changes, freezing the disk file, and then exporting a snapshot of the disk file. This snapshot image is then available for backup.

To help ensure effective recovery of an ESX Server system, administrators should protect three main components:

- Virtual disks—which contain the guest-hosted OS, applications, and the application data
- VM configuration files
- Physical machine configuration data

Options for protecting VMware ESX Server VMs with CommVault software

CommVault Galaxy backup and recovery software is flexible and modular—so enterprise IT organizations can customize the level of protection for their VMware-based environment. Administrators have two options for protecting data on VMs: treating the VMs as physical servers or treating them as files on an ESX Server system. For both options, Galaxy iDataAgent™ modules (iDAs) are used to implement data protection. These modules are designed to protect file systems in virtualized environments.

Protection option 1: Treating VMs as physical servers

IT administrators can deploy CommVault Galaxy software within VMware VMs just as if the VMs were separate physical servers. For example, administrators can install Galaxy iDAs within the VMs and then use a Galaxy CommServe™ module deployed on a separate physical system to protect those VM environments by sending backup data across the network. An advantage of this approach is the ability to run traditional incremental, differential, and synthetic full backups in addition to full backups. However, IT organizations cannot protect the entire encapsulated VM environment—that is, the VM configuration information in addition to the data residing on the VMs—for easy recovery.

Figure 2 depicts an ESX Server system with four VMs, each of which are protected by a Galaxy backup server running on a separate physical machine. The file systems and applications running within the VMs are protected via the Galaxy file system and application iDAs. Backup and recovery is performed using the normal Galaxy process, with each VM presented as a unique client system within the Galaxy unified console.

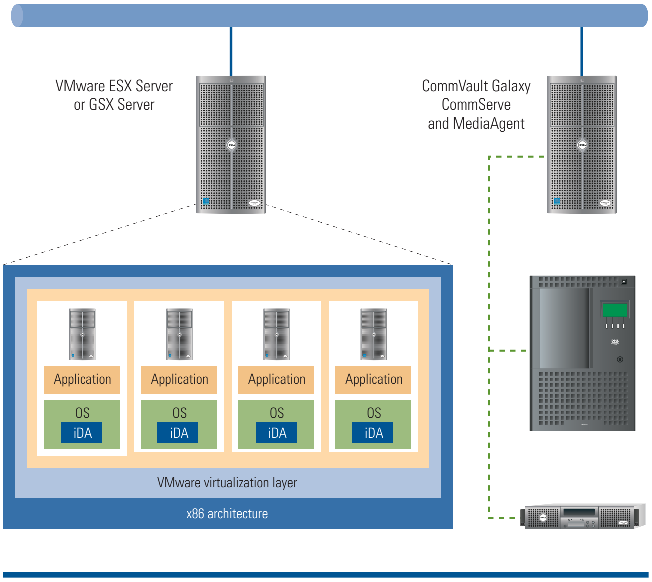

Figure 2. CommVault Galaxy software protecting an ESX Server VM environment

In this scenario, Galaxy software enables administrators to configure data movement from the virtual server environment to the backup server using standard Galaxy configuration options over the LAN. Backup devices, whether disk or tape, are configured and managed from the Galaxy MediaAgent™ component running on the backup server and can be shared among all VMs as well as other systems. These drives can also be specifically allocated to protect specific VMs, again using standard Galaxy configuration options.

Protect VM system state. In most methods for backing up VM data in a VMware-based environment, treating VMs as physical systems does not provide the added benefit of being able to protect and recover the VM configuration itself. Being able to protect and recover the encapsulated VM configuration is, therefore, described as a unique advantage of using virtual disk files to protect VM file system and application data, which is the other protection option (described in the next section). CommVault Galaxy provides full system state protection for both Windows and Linux file systems along with individual file protection and recovery—because CommVault builds system state protection into the file system iDA. Administrators do not need to protect the virtual disk files themselves to recover VMs at crash-consistent states.

Avoid sending backup copies across the network. IT administrators can avoid sending backup copies of VMs on an ESX Server system across the LAN. To do this, they can select a VM to deploy the Galaxy MediaAgent, which manages a backup device. Best practices recommend using a shared backup device for all VMs running on the ESX Server system. Alternatively, a MediaAgent can be run on more than one VM attached to more than one backup device, depending on factors such as how much data needs to be protected and how much load particular VMs need to support.

Protection option 2: Treating VMs as files on an ESX Server system

This option takes advantage of how ESX Server creates a file per VM disk device and keeps track of changes using redo logs. Backup and recovery using this method is simplified to the physical machine level and creates a crash-consistent recovery to a specific point-in-time.

For example, IT administrators can deploy a single Galaxy iDA for Linux file system protection to protect the entire ESX Server system and all of its guest-hosted VM environments. Then, when IT administrators need to recover a specific VM, the VM OS, file system data, applications, and application data are all recoverable to a consistent point in time.

Using this option means that IT administrators must plan to protect and recover very large files (typically more than 2 GB in size). In cases in which administrators need to recover a single file, they must first recover the entire disk file—which can take dramatically longer than recovering a single file, depending on the size of the virtual disk file.

Therefore, this option is well suited for small ESX Server implementations with file systems and no or few applications running. By definition, virtual disk files cannot be protected while the VM is running. To protect a virtual disk file, administrators must add a redo log, which makes the virtual disk file static and therefore available for backup. Administrators can also take a snapshot of the VM, creating multiple points in time—each of which can be backed up for crash-consistent recovery.

Effective backups for virtualized servers

As data needs continue to grow because of compliance regulations and end-user requirements, backup, restore, and disaster recovery are becoming crucial issues in today's data center operations. Virtualized servers such as VMware ESX Server systems require the same attention to backup strategies as physical systems. Together, CommVault and VMware can provide various options for effectively backing up and protecting the data in virtualized environments.

Kelly Harriman-Polanski is the director of product marketing for CommVault. Kelly recently joined CommVault and has worked in the storage software market for nearly 10 years, most recently at EMC/Legato. Her interests include data and information management, data archive, retrieval, compliance, data classification, and integrated snapshot and replication management. Kelly graduated magna cum laude and Phi Beta Kappa from Augustana College in Illinois.

FOR MORE INFORMATION

CommVault Galaxy:
www.commvault.com/backup_and_recovery.asp

Dell/EMC AX150 and AX150i Networked Storage Systems: Consolidation Without Complexity

Increased drive capacity, expanded host connectivity, and improved system availability mark the second generation of the world's easiest SAN storage array

Dell/EMC AX150i

Digital information is growing every which way, and small organizations need simple, cost-effective ways to manage that information. To keep employees working productively, meet regulatory requirements, and satisfy customers, organizations also must keep information secure and available. For organizations that cannot add staff and do not want the time and trouble of costly, complex storage-consolidation projects, Dell/EMC AX150 and AX150i networked storage arrays can be the answer. These systems, combined with either a Fibre Channel or IP network, are designed to make setting up and managing an entry-level storage area network (SAN) easy and affordable. And they feature safeguards to help keep information available.

Next-generation value

The AX150 and AX150i are the next generation of the Dell/EMC AX series disk arrays, replacing the Fibre Channel–attached AX100 and the Internet SCSI (iSCSI)–attached AX100i systems, respectively. The new-generation systems provide storage for up to 10 host servers, which can be redundantly connected either through a Fibre Channel switch on the AX150 or via a traditional Ethernet network on the iSCSI AX150i. These arrays can house up to twelve 3.5-inch Serial ATA II (SATA II) drives–available in 250 GB and 500 GB models–in a single 2U* rack enclosure, delivering 750 GB to 6 TB of raw disk capacity. For enhanced availability, the AX150 and AX150i systems can be configured with dual storage processors. Enhancements to power supplies, compared to previous-generation systems, provide extra redundancy to prevent potentially disruptive hardware failures.

Consolidation without complexity

Consolidated storage can help an organization lower total cost of ownership by simplifying operations and improving the utilization of storage resources. To help IT organizations configure and manage an entry-level SAN, the Dell/EMC AX150 and AX150i include setup wizards and intuitive, easy-to-use software. In addition, the iSCSI option enhances consolidation by allowing for simple connection over an existing LAN or dedicated IP network.

All software required for setup and day-to-day operation of an AX series–based SAN is included with the array. EMC Navisphere Express is specifically designed to make system setup and management easy. Administrators can configure a AX150 or AX150i in as little as four steps by using a graphical interface and wizards. They can provision the exact amount of storage necessary for each server and assign it–all from a single console. For fast data backup and recovery and increased application availability, EMC SnapView Express can create up to eight point-in-time copies of production data in seconds. EMC PowerPath software provides automatic load balancing and failover for redundant path connections.

Safeguards against disruption to data availability

Redundant hardware architecture helps make the AX150 and AX150i systems resilient to hardware malfunctions and data corruption. If a power failure occurs in dual-processor configurations, data in cache is written to disk, where it can be retrieved long after a battery-backed cache would have been exhausted. Fully redundant power supplies help keep systems up and running even if a power supply fails. And SATA II disks are designed to withstand the rigors of the small enterprise. These features, along with the integrated SnapView Express and PowerPath software, can help mitigate the risk of disruption to data availability.

Getting a grasp on proliferating business information can be easy and affordable with the Dell/EMC AX150 and AX150i networked storage systems. For small organizations that want high availability without high cost or complexity, the AX150 and AX150i set new standards of value and availability for SAN storage arrays.

For more information:

Dell storage: **www.dell.com/storage**

*Systems with dual storage processors require an additional uninterruptible power supply, which is 1U in height.

When information
comes together,
business just keeps
getting better.

DELL | EMC²

ALL THE RIGHT CONSOLIDATION, BACKUP AND ARCHIVE SOLUTIONS

Whether you need fast backup and complete protection or scalable and easy-to-manage storage consolidation for your mid-size enterprise, DellIEMC brings you solutions that are high on results – and simple to use. That's because it's easier than ever to put premium software, robust storage, and world-class technical support to work solving your business's critical IT challenges.

BUSINESS SOLUTIONS FOR MIDSIZE ENTERPRISES

Entry SAN Solution

- DellIEMC AX150 Storage Platform
- iSCSI or Fibre Channel Connectivity
- EMC® Navisphere® SAN Management Software

SAN Windows Backup Solution

- DellIEMC CX300 Storage Platform
- EMC® Navisphere® SAN Management Software
- EMC Snapview™ and EMC Replication Manager SE Software
- EMC SAN Copy™ Software

Data Archiving Solution

- EMC Centera™ Storage Platform
- Windows File System Archive Edition with EMC DiskXtender® Software
- Governance Edition with EMC EmailXtender® and EMC DiskXtender® Software

CALL 800.999.3355 toll free **www.dell.com/emc**

Data Recovery Outside the Core Data Center

Information availability is critical to an enterprise. Most enterprise IT departments have procedures in place to help ensure data availability, but if a move to a remote site becomes necessary, significant downtime may be incurred. To enhance disaster recovery capabilities and mitigate downtime, enterprises can build dual data centers. This article discusses secondary-site strategies and provides guidance on their deployment.

BY MICHAEL KIMBLE

Related Categories:

Business continuity

Disaster recovery

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

High availability of data is vital to enterprises. Without constant access to electronic information, services can be disrupted and enterprises can suffer far-reaching business and financial consequences. Minimizing this possibility is the goal of most enterprise IT departments. Although IT departments typically have procedures in place to meet availability requirements for localized problems such as data corruption or system failures, they can be unprepared to maintain information availability in the event of a disaster or catastrophic systems failure. Reasons for this may include the following:

- No single disaster recovery implementation meets every requirement.
- Business expectations are not aligned with IT processes.
- Thorough documentation and details for recovery are not up-to-date.

Recent events, including terrorist attacks and natural disasters, have demonstrated that an enterprise's need for adequate disaster recovery planning is greater than ever before. Although the odds of needing to transfer IT operations to a secondary data center are small, the consequences of not being prepared to do so are enormous. Knowing this, many enterprises are choosing to build secondary data center sites as a way of mitigating downtime risk. However, the time and expense required to do this means that they must consider all available options. This article reviews strategies for setting up secondary sites and provides guidance for estimating recoverability requirements using two criteria: allowable data loss, which is known as the recovery point objective (RPO); and the time required to restore operations, which is known as the recovery time objective (RTO).

Selecting the level of remote recoverability

For the past few decades, one of the most common data recovery methods has been to back up data to tape and send it off-site to be retrieved if necessary. In some cases, tapes are sent to secondary locations where the recovery process begins. This still holds true for the majority of today's enterprises—those that regard an RTO of one day or longer as acceptable—and probably will for the foreseeable future.

However, some enterprises prefer to have dual data centers that can provide transparent failover of applications with little interruption to end users. This has resulted largely from advances in data replication software and the reduced costs of reliable hardware. Although this is

Guideline	Class A	Class B	Class C	Class D
RPO	No data loss	Less than 10 minutes	4 hours	Best effort
RTO (internal)	Less than 30 minutes	Less than 4 hours	Less than 24 hours	Less than 48 hours
RTO (external)	Less than 24 hours	Less than 48 hours	3 days	1 week
Planned downtime	Less than 12 hours per year	Less than 24 hours per year	Less than 24 hours per year	Best effort
Remote technology used	Synchronous mirror	Asynchronous mirror	Host-based replication	Tape only
Local technology used	Point-in-time copy	Point-in-time copy	Backup to disk or tape	Backup to tape

Figure 1. Example recovery guidelines for various data classes

a legitimate goal, many enterprises lack the resources, budget, or justification for the expense and labor required for such an undertaking. For those reasons, enterprises should set realizable goals when planning a secondary site for disaster recovery.

Aligning business requirements with strategies

Just as not all data is created equal, the resources and requirements that must be committed to that data are not equal. A clear understanding of which applications and data affect an organization the most is the first step to defining a disaster recovery strategy. This understanding comes from performing a thorough business-impact analysis of the environment and documenting the findings in a data-classification report. The business-impact analysis estimates potential costs in lost revenue and productivity as well as damage to an enterprise's reputation and business relationships should a lengthy downtime occur.

Data classification involves many facets of an organization—from IT to business units to executive management. Each group often views the importance of an application and its data very differently. For example, e-mail might be the most important application to end users, while the enterprise's leaders might view an order management application as the most important because the enterprise cannot generate revenue and process orders without it. These differing views demonstrate the need to classify applications and data. Figure 1 shows an example classification matrix wherein each application or data pool is assigned to a defined class.

Examining secondary-site deployment options

To fully realize its goals for remote recoverability, an enterprise must examine its secondary-site data recovery options. Commercial hot sites and internally funded dual data centers are among the most popular choices, but other available options may be of interest. Determining the most suitable option should be based on the guidelines set forth in each enterprise's data classification and recovery decisions. Figure 2 summarizes secondary-site data recovery options.

Budget considerations

One of the primary factors in choosing a secondary-site strategy is the cost required to implement and maintain it compared to the return for potential use.

Hot site. With commercial hot sites, enterprises pay recurring fees to maintain a facility and keep equipment and resources available should they need them. These fees can vary widely based on the established RTO and can range from US$100,000 for recovery over several days to upwards of US$1 million for recovery in hours. An advantage of this option is that administrators can test preparations regularly and help ensure that processes and procedures are in order.

Cold site. With cold-site facilities, an enterprise can own or hold a lease on the space it requires and can then purchase hardware, software, and other items as needed. Although this can be a significant expense, an enterprise has to spend money to furnish and operate the secondary data center only if a disastrous failure occurs. This can be an appealing alternative to both hot-site and dual data center approaches. However, data recoverability using a cold site is largely based on the ability of hardware and software providers to deliver their products in a timely fashion.

Dual data centers. This approach to data recoverability requires that enterprises construct secondary facilities to house their emergency processing needs. Although this option provides the fastest recoverability, it can cost up to three times as much as

Type of recovery site	Description	Typical recovery time
Commercial hot site	Commercial hot-site vendors offer fully staffed facilities ready to support an enterprise's processing needs should the need arise. They offer not only processing hardware, but also workspace, telephones, and network connectivity.	24 to 48 hours
Cold site (shell site) and replacement systems	A cold-site facility is prepared to receive replacement hardware and technologies to resume processing. Hardware (and software, if not vaulted with tapes) is procured from a primary vendor or third-party contract providers.	3 to 7 days
Mobile shell site	A mobile shell site is a mobile processing facility that can be delivered to a destination of choice. The facility contains hardware and limited workspace.	1 to 3 days
Reciprocal backup agreement	A reciprocal backup agreement is an agreement between two noncompeting enterprises to share hardware resources in the event of failure.	12 to 36 hours
Dual data center	A dual data center is an internally funded data center equipped with the same type of hardware as the primary data center. It may be ready for processing immediately, depending on the vaulting strategy.	Instantaneous to 12 hours

Figure 2. Types of secondary sites for data recovery

the primary data center—and it may never be used. This additional cost can be attributed to obtaining the facility; purchasing hardware and software; and retrofitting the infrastructure to meet HVAC, telecommunications, security, and workspace requirements.

Because executive management often has a hard time justifying such a large expense for something to sit idle, organizations often try to use the secondary site for some primary processing. However, a thorough review of the impact of a catastrophic failure on an enterprise can often justify the need for a dual data center. For example, if an enterprise stands to lose US$1 million for every hour that its primary applications are down, the cost of a standby data center can be acceptable.

A further consideration is that although the primary- and secondary-site configurations are initially sound and the data is mirrored, the two sites often lose congruity as time progresses. This can occur when an enterprise adds new hardware and functionality to its primary site but overlooks updating the secondary site because of cost, staffing shortages, or other reasons. This problem has led to the adoption of virtualization technologies at the remote site to keep physical hardware costs low. Most disaster recovery situations do not require all operations to function at full capacity, so virtualization technology can be used to consolidate operations onto fewer physical servers than exist at the primary site—which can lead to savings in hardware costs.

Implementing a dual data center

After investigating all the available options, an enterprise may decide that a dual data center approach is the best choice. In that event, the enterprise should carefully choose an appropriate location and create a comprehensive plan for hardware, staff, and data connectivity.

Location and facility considerations

Choosing a location for a secondary data center can be difficult. Consideration should be given to distance, need for site renovation, availability of utilities, and staff proximity, among other factors. Distance from the primary site typically depends on where the facility is located and the type of threats to which it is exposed. For example, in Florida and the Gulf Coast of the United States, the wide path of hurricanes dictates the need for distance separations of 100 miles or more, whereas in a tornado zone, a much smaller distance—for instance, seven miles—would suffice.

These factors also affect the enterprise's recovery objectives and the type of replication it employs. An RPO of zero data loss, for example, requires the use of synchronous replication software, which is often limited to approximately 60 miles. Going beyond that distance may require asynchronous tools that increase RPO exposure.

Power grid considerations

After a geographic location has been identified, the task of selecting a specific site begins. The most important factor is determining which power grid the facility operates on and which other agencies use that grid. Typically, first-responder facilities such as hospitals, police stations, fire stations, and city functions are the first to be restored in a power grid. After power is restored to a data center, data communications services and connectivity soon follow. Although an enterprise's data center may have power and data available, it provides no advantage unless users can access it both locally and remotely. Therefore, enterprises should ensure they have processes in place for users to gain connectivity to the secondary data center.

Another important factor to consider in choosing a site is the nature and extent of the resources that are required to make the facility an acceptable data center. Enterprises should determine whether a costly retrofit—such as a raised floor, HVAC, or security—is necessary and whether the site can provide adequate workspace for staff operations.

Understanding factors that can inhibit disaster recovery

Personnel and data are two key components in any successful disaster recovery. However, other factors can greatly affect an enterprise's success and its ability to meet recovery objectives:

- **Applications:** Although data may be backed up and protected, current copies of an enterprise's applications also should be stored off-site.
- **Legacy hardware:** The need for legacy hardware such as specialized cards and readers should not be overlooked.
- **Documents and forms:** Preprinted forms such as checks and statement forms should be stored off-site or arrangements should be made to procure them quickly.
- **Local user data:** A significant amount of data may be stored on local users' hard drives, and this data must be backed up.

Finding the right fit

Selecting a disaster recovery strategy requires a significant amount of research and time. However, the time and research invested in studying an enterprise's environment and defining realizable recovery goals can help enterprises determine their choices and reach sound decisions. If enterprises choose to implement a secondary site for disaster recovery, they can benefit from the business continuity and high availability that this disaster recovery option provides.

Michael Kimble is an enterprise technologist in the Advanced Systems Group at Dell. Working with Dell consultants and customers, he helps design storage implementations for disaster recovery and business continuity. Michael has a B.S. in Finance and Economics from the University of Central Florida.

FOR MORE INFORMATION

Dell business continuity:
www.dell.com/businesscontinuity

Phoenix Recover Pro 6

Helps Close the PC Disaster Recovery Gap

Desktop and mobile PC users must be able to recover quickly from viruses, OS crashes, and user errors that threaten to disrupt business continuity, hamper productivity, and increase IT support costs. Server-based recovery solutions need access to storage and require the Microsoft® Windows® OS to be operational. Enterprises can implement a PC endpoint restore solution to help protect data and applications on PCs and enable immediate recovery regardless of whether users can boot into Windows or access the network.

BY SCOTT CHAMBERS

Related Categories:

Backup

Desktop software

Disaster recovery

Microsoft Windows

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

Enterprises spend a significant amount of money each year to address PC disaster recovery. Even with safeguards such as advanced firewalls, virus protection, intrusion prevention, access control technologies, and regular backups of key databases, a disaster recovery gap can exist. This occurs because desktop and notebook PCs, which most enterprise employees rely on for day-to-day business operations, typically lack built-in OS and data restore capabilities.

It is no secret that a significant percentage of business-critical data resides on employee PCs, not on enterprise servers. This creates a disaster recovery gap that is expected to widen as PC *endpoints*—that is, any device with an IP address including Web-enabled devices such as smart phones, kiosks, ATMs, and remote testing equipment—become more pervasive and subject to infection. However, IT management can help protect enterprise data residing on employee PCs by considering an advanced PC endpoint restore solution such as Phoenix® Recover Pro™ 6 software.

Identifying a PC endpoint restore solution

When evaluating PC endpoint restore solutions, enterprises should consider several criteria, including the capability for immediate recovery even when Windows does not boot, self-contained recovery, a protected recovery environment, ease of use, mean time to recovery, and affordability.

Immediate recovery even when Windows does not boot. A PC endpoint restore solution should be able to provide recovery from a Microsoft Windows OS crash regardless of the cause. A best-in-class recovery solution does not depend on the correct functioning of the Windows OS, but rather can restore the PC endpoint device to its state before the disaster occurred.

Self-contained recovery. Most disaster recovery solutions rely upon network access and recovery media when Windows does not boot. This approach significantly complicates and lengthens the restore process. An advanced approach to disaster recovery enables a CD-free recovery environment that can be instantly activated even in the event of a total Windows failure.

Protected recovery environment. Effective PC recovery solutions are immune to virus infection and always available to end users. The key to this immunization is to house the recovery mechanism and data in a protected area of the hard disk that is designed to be impervious to attacks on Windows and to remain available even if Windows crashes.

Ease of use. An optimal PC recovery solution does not require IT interaction. Nontechnical users should be able to restore their own systems without IT involvement.

Mean time to recovery. A disaster recovery plan should minimize user downtime and practically eliminate IT staff involvement. In a typical enterprise environment, users must wait several hours for the IT team to recover a failed desktop or notebook system and restore all the user's data files using traditional methods such as CDs and other media. Even then, traditional recovery methods rarely restore a user's system completely to the state that existed prior to the disaster. An optimal recovery solution restores a system in minutes versus hours and is capable of restoring system files, registries, applications, and data files to the exact state that existed prior to the disaster.

Affordability. To enable viable PC endpoint disaster recovery throughout organizations of all sizes, the solution must be affordable—making universal installation on each device cost-effective.

Closing the PC disaster recovery gap with Recover Pro 6

Phoenix Recover Pro 6 software enables the full restoration of PC endpoint data and applications in minutes—without IT intervention, network connection, or the need for external recovery media. A self-contained application housed in a protected area of the hard drive, Recover Pro software is designed to work even when Windows does not boot. In this way, Recover Pro helps solve a growing problem faced by today's enterprises: how to enable desktop and mobile PC users throughout an organization to recover quickly from worms, viruses, and other malicious attacks or OS crashes that can disrupt business continuity, hamper employee productivity, and increase IT support costs.

Scott Chambers is a director of product marketing at Phoenix Technologies and is responsible for the company's recovery and imaging product lines. Prior to Phoenix, he worked at Adaptec, Apple Computer, and Arthur Anderson & Co. His product experience includes storage systems and components, consumer electronics, and software. Scott has a B.S. in Accounting from Brigham Young University and an M.B.A. from Santa Clara University.

FOR MORE INFORMATION

Phoenix Recover Pro 6:
www.phoenix.com

Migrating an Oracle9*i* RAC Database to Oracle 10*g*

A vastly enhanced feature set and proven scalability across multiple nodes coupled with automated management capabilities have motivated many enterprises to upgrade their systems to Oracle® Database 10*g* Real Application Clusters. This article focuses on a tool available in Oracle Database 10*g* that helps simplify the upgrade process significantly—the Database Upgrade Assistant utility.

BY SANJEET SINGH, ANTHONY FERNANDEZ, AND UDAY SHET

Related Categories:

Best practices

Clustering

Database

Dell PowerEdge servers

Oracle

Scripting

System deployment

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

An outstanding feature set with enhanced support for clustered applications and automated management capabilities enable Oracle Database 10*g* to provide a highly available and scalable platform for enterprise database applications. The Dell™ PowerEdge™ line of dual- and quad-processor servers combined with the Real Application Clusters (RAC) architecture of Oracle Database 10*g* can provide a desirable price/performance advantage to enterprise customers. This combination enables a highly scalable architecture that allows an IT infrastructure to grow in cost-effective increments.

Oracle RAC 10*g* introduces features such as Automatic Workload Repository, Automatic Storage Management, Automatic Database Diagnostic Monitor, and Oracle Data Pump, plus enhancements to Oracle Enterprise Manager, tablespace control, and self-tuning capabilities. Added to a standardized and modular set of hardware components, this set of enhanced management and self-tuning features can significantly reduce total cost of ownership. Dell provides a validated and tested stack of Oracle RAC 10*g* database solutions. This can reduce the time required for integration and deployment significantly in addition to providing a proven solution.

These factors prompt many enterprises to migrate their database systems to Oracle RAC 10*g*. The upgrade can be performed manually by running the SQL-based upgrade scripts provided by Oracle 10*g*. Alternatively, administrators can use the Database Upgrade Assistant (DBUA) utility, which is the Dell-recommended method. The DBUA is an automated wizard introduced in Oracle Database 10*g* that supports upgrades from Oracle8*i* (8.1.7), Oracle9*i* (9.0.1), and Oracle9*i* Release 2. Administrators upgrading from earlier versions such as Oracle 8.1.6 must first upgrade to Oracle 8.1.7 before they can use the Oracle Database 10*g* DBUA to upgrade their version of the database.

Cluster upgrade considerations

The database upgrade procedure involves installing the new version of the database—Oracle Database 10*g* binaries—and upgrading the existing database. Dell recommends that administrators do a test run of the upgrade on a test database to verify that the application data and custom objects can upgrade successfully before upgrading the production system.

To help ensure minimal or zero downtime, Dell recommends that administrators have a logical standby database. Oracle Database 10*g* RAC has support for basic rolling upgrades—the RAC instances on different nodes can be upgraded in a sequential manner and hence require minimal downtime. However, rolling upgrades are not supported for database versions prior to Oracle Database 10*g*.

The Oracle Database 10*g* upgrade installation can be performed on various configurations that have been tested and validated by Dell and Oracle.[1] However, additional factors must be considered when upgrading a preexisting Oracle9*i* cluster to Oracle Database 10*g* RAC:

- Oracle Database 10*g* should be installed from an ORACLE_ HOME directory that is different from the preexisting Oracle9*i* database ORACLE_HOME. For example, if the Oracle9*i* ORACLE_HOME is /opt/oracle/product/9.2.0, then Oracle Database 10*g* should be installed under /opt/oracle/product/10.1.0.
- The Oracle9*i* Global Services Daemon (GSD) must be stopped before Cluster Ready Services (CRS) is installed. However, the Oracle9*i* Cluster Manager (oracm) should be running.
- If the Oracle9*i* server management database file (srvm.dbf) is located on an Oracle Cluster File System (OCFS) partition residing on shared storage, the server management configuration file (/var/opt/oracle/srvConfig.loc) should be moved to a temporary location for the duration of the CRS installation; this step helps ensure successful execution of the CRS Oracle Universal Installer (OUI). Administrators should return /var/opt/oracle/srvConfig.loc to the proper OCFS partition after the CRS installation has completed.
- During the CRS installation process, the OUI prompts the administrator to log in as the root user and run the scripts located in ORA_CRS_HOME/root.sh. To upgrade the environment variable for the CRS, the Oracle Cluster Registry must point to the Shared Oracle9*i* server management database file (srvm.dbf). The administrator must edit ORA_CRS_HOME/root.sh on all nodes and update the CRS_OCR_LOCATION to indicate the location of the srvm.dbf file. The administrator must also edit the /etc/oracle/ocr.log file with the updated path to the srvm.dbf file.

After CRS is started on a cluster, Oracle Database 10*g* software can be installed according to OS-specific directions that can be found at www.dell.com/10g. However, administrators should be aware of the following considerations:

- When Oracle Database 10*g* software is being installed, CRS and the Oracle9*i* oracm stack must be running. However, the Oracle9*i* GSD should not be running.
- The OUI for Oracle Database 10*g* detects any existing database and displays a page from which the administrator can select the older database to be upgraded.

Procedural upgrade enhancements

All too often, the Oracle database upgrade process has been cumbersome and error-prone. With releases before Oracle Database 10*g*, the administrator performing the upgrade was required to manually enter the initialization parameters and system resources. For example, the administrator had to ensure that adequate space was provided in the system tablespace and that rollback segments were properly sized.

When upgrading to Oracle Database 10*g*, administrators can opt to run the upgrade scripts manually or use the DBUA wizard. The DBUA is an interactive graphical tool that guides the administrator through the upgrade process, calling upon other tools in the background to perform individual tasks as required.

Pre-Upgrade Information Tool

The Pre-Upgrade Information Tool (utlu101i.sql) is invoked by the DBUA. It analyzes the database to verify that system requirements for upgrading to Oracle Database 10*g* have been met before the upgrade begins. For example, the Pre-Upgrade Information Tool checks the following parameters and system resources:

- The current version of the database is 8.0.6; 8.1.7; 9.0.1; or 9.2.
- The parameter `compatible` is 9.2.0 or greater.
- Tablespaces meet the minimum recommended sizes.
- Redo logs are at least 4 MB.
- The parameter `pga_aggregate_target` is at least 24 MB.
- The parameters `shared_pool_size`, `large_pool_size`, and `java_pool_size` are set to at least 96 MB, 8 MB, and 48 MB, respectively.
- The size of the sysaux table is adequate for Oracle Database 10*g*.

In addition, the Pre-Upgrade Information Tool checks for updated, deprecated, and obsolete initialization parameters, and estimates the time to perform the upgrade.

By default, the DBUA turns off archiving to improve performance. However, if the database being upgraded is part of an Oracle Data Guard data protection or disaster recovery environment, turning off archiving may cause problems. To avoid related complications, the administrator can enable the archiving mode during the upgrade process by setting the parameter `DisableArchiveLogMode` to “false” in the Oracle Database 10*g* _oracle_home/rdbms/admin/utlu101x.sql file.

[1] For detailed information about Oracle Database 10*g* support and product offerings—including specific configurations that have been tested and validated by Dell and Oracle—visit www.dell.com/10g.

Storage of auxiliary database metadata

The sysaux tablespace is a required tablespace in Oracle Database 10*g* that provides storage of non-"sys"-related tables and indexes that resided in the system tablespace of earlier versions of Oracle databases. This system-owned tablespace provides a centralized location for all auxiliary database metadata that does not reside in the system tablespace. For example, the objects related to the Oracle Recovery Manager (RMAN) catalog, online analytical processing (OLAP), and Oracle Text now reside in sysaux instead of the system tablespace. The sysaux tablespace also reduces the number of tablespaces created by default, and it has the same security profile as the system tablespace.

Before the upgrade begins, the DBUA asks for the location of the sysaux tablespace and recommends a size for it. The administrator has the option of changing its size before the tablespace is created. Then the DBUA creates a sysaux tablespace of the determined size.

Database backup option

The Oracle Database 10*g* DBUA also prompts administrators to create a backup of the database before the upgrade begins. This backup option enables administrators to revert to a previous version if a problem occurs during the upgrade. Best practices advise database administrators to back up production databases and perform a test upgrade on a clone of the existing database. This trial run can help ensure that the upgraded production platform meets enterprise requirements.

Post-Upgrade Status Tool

After completing the necessary pre-upgrade steps described in the preceding sections, the DBUA proceeds to upgrade the database.[2] In earlier versions of the Oracle database, it could be difficult for administrators to determine whether an upgrade was successful. With Oracle Database 10*g*, after an upgrade is complete, the DBUA invokes the Post-Upgrade Status Tool. This tool is designed to check the status of the upgrade and display a confirmation message if all components in the database have successfully completed the upgrade. If it finds a problem, the DBUA lists components that are invalid or do not reflect the correct version.

The Post-Upgrade Status Tool works by querying dba_server_registry and listing the status and version of each component. If a component has not been upgraded, it provides a brief description of the corrective action that the administrator needs to take to manually upgrade the component. The DBUA provides the option to recompile invalid objects immediately after the upgrade or at runtime on a when-needed basis. When-needed recompiling can shorten upgrade time, but it may also degrade database performance when recompiles are performed at runtime. If an upgrade is not satisfactory, the DBUA allows the administrator to restore the saved database. At that point, all changes performed during the upgrade are reversed.

Streamlined migration to Oracle Database 10*g* clusters

Oracle Database 10*g* can help enterprises improve the availability and performance of clustered database applications by providing a significantly enhanced feature set and automated management capabilities. In addition, the DBUA utility helps simplify the process of upgrading from Oracle9*i* to Oracle Database 10*g* clusters, enhancing the efficiency and productivity of database administrators. As a result, enterprises running earlier versions of the Oracle database need not shy away from the benefits of migrating to this feature-rich, highly scalable database.

References

Graves, Thomas. "The Self-Managing Database: Easy Upgrade." Oracle OpenWorld Convention, September 7–11, 2003. www.oracle.com/openworld/archive/sf2003/solutions_management.html

Oracle Corporation. "Generic Linux Upgrade from Oracle9*i* Real Application Clusters to Oracle 10*g*R1 Real Application Clusters Using Database Upgrade Assistant (DBUA)." Oracle Corporation, February 2005. www.oracle.com/technology/products/database/clustering/pdf/9i_to_10g_RAC_upgrade_on_Linux.pdf

Soorma, Gavin. "Using the Oracle 10*g* DBUA to Make the Giant Leap Forward." Oracle OpenWorld Convention, December 5–9, 2004. download-east.oracle.com/oowsf2004/1127.pdf

Sanjeet Singh is a software engineer in the Dell Database and Application Engineering Group. Sanjeet has a B.S. in Electrical Engineering and an M.S. in Computer Engineering from Purdue University.

Anthony Fernandez is a senior analyst with the Dell Database and Applications Team of Enterprise Solutions Engineering in the Dell Product Group, where he focuses on database optimization and performance. He has a B.S. in Computer Science from Florida International University.

Uday Shet is a senior engineering analyst in the Dell Database Solutions Group. Uday has a B.E. in Computer Science and is also an Oracle Certified Professional (OCP).

FOR MORE INFORMATION

Oracle Database 10*g*:
www.oracle.com/database

Oracle database backup and recovery:
www1.us.dell.com/content/topics/global.aspx/power/en/ps1q03_singh?c=us&cs=555&l=en&s=biz

[2] For more information, including a detailed, step-by-step Oracle database upgrade procedure, visit www.dell.com/oracle, locate the "Learning Center" section, and click on "Migration and Upgrade Documents."

Scaling an Oracle RAC System with

Disk-based Backup and Recovery

Data growth is inevitable in most IT environments, and yesterday's cutting-edge system can become today's bottleneck. Tested and validated Dell/Oracle industry-standard database components and Dell/EMC storage platforms can help enterprises expand their IT capacity to accommodate their increasing data. One way to scale is to add storage for disk-based backup and recovery, which can help minimize backup times, speed database recovery, add recovery points, and help meet service-level agreements.

BY MAHMOUD AHMADIAN, UJJWAL RAJBHANDARI, DAVID MAR, AND CHETHAN KUMAR

Related Categories:

Backup

Backup to disk

Business continuity

Clustering

Dell/EMC storage

Disaster recovery

Microsoft Windows Server 2003

Oracle

Red Hat Enterprise Linux

Scalable enterprise

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

Dell best practices strongly recommend deploying tested and validated configurations for databases because of the complexity and business-critical nature of such environments. Administrators can deploy these configurations with confidence that they have been thoroughly tested and stressed.

Although database configurations can be custom-tailored to an enterprise's needs, such configurations are typically unique in their combination of software, hardware, and driver versions. Consequently, they may not have been thoroughly tested or stressed in that particular combination. To alleviate such issues and to provide comprehensive support for the enterprise, Dell defines and tests stacks of database components, helping to identify interoperability issues before they reach the enterprise. These tested and validated Dell™ configurations are defined by a configuration matrix that includes hardware and software components and the process by which to deploy these components (see Figure 1).

Deploying tested and validated Dell/Oracle database configurations can hold several advantages for an enterprise. Such an approach can help eliminate deployment problems and enable system support and fast deployment. A deployment CD is provided with Dell/Oracle configurations to help ensure that drivers, kernel configurations, and Oracle prerequisites are optimally fulfilled for proper implementation. Furthermore, Dell promotes a scalable enterprise strategy, which calls for a phased deployment approach that can scale to meet the needs of any size enterprise. Dell/Oracle configurations use industry-standard components that can easily be upgraded to meet business requirements—providing a predictable foundation on which enterprises can later expand.

Achieving high data availability with Dell/Oracle configurations

Enterprise databases may range from a single-instance database consisting of a few gigabytes of data to multiple

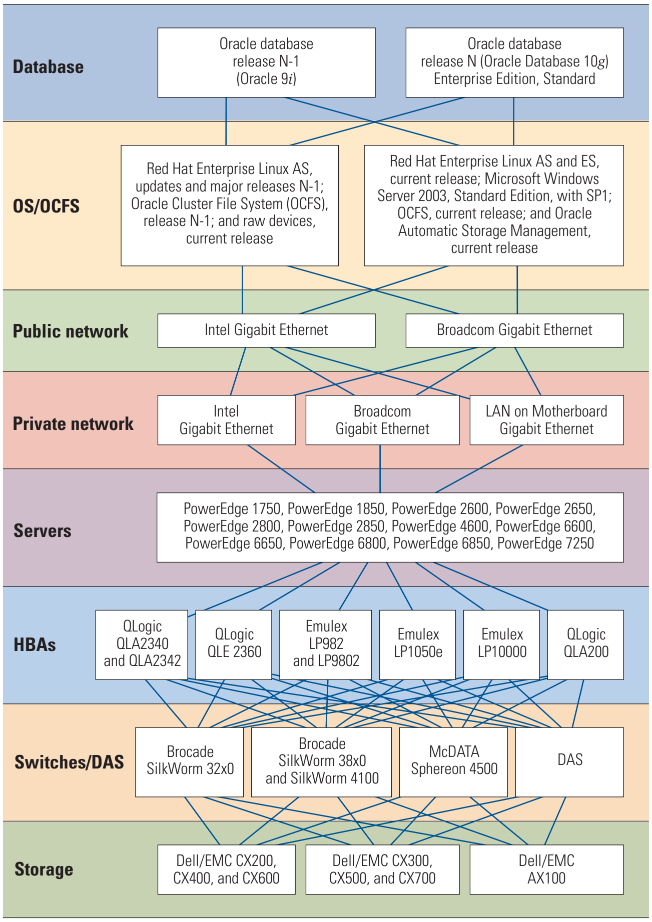
Database
Oracle database release N-1 (Oracle 9*i*)
Oracle database release N (Oracle Database 10*g*) Enterprise Edition, Standard
OS/OCFS
Red Hat Enterprise Linux AS, updates and major releases N-1; Oracle Cluster File System (OCFS), release N-1; and raw devices, current release
Red Hat Enterprise Linux AS and ES, current release; Microsoft Windows Server 2003, Standard Edition, with SP1; OCFS, current release; and Oracle Automatic Storage Management, current release
Public network
Intel Gigabit Ethernet
Broadcom Gigabit Ethernet
Private network
Intel Gigabit Ethernet
Broadcom Gigabit Ethernet
LAN on Motherboard Gigabit Ethernet
Servers
PowerEdge 1750, PowerEdge 1850, PowerEdge 2600, PowerEdge 2650, PowerEdge 2800, PowerEdge 2850, PowerEdge 4600, PowerEdge 6600, PowerEdge 6650, PowerEdge 6800, PowerEdge 6850, PowerEdge 7250
HBAs
QLogic QLA2340 and QLA2342
QLogic QLE 2360
Emulex LP982 and LP9802
Emulex LP1050e
Emulex LP10000
QLogic QLA200
Switches/DAS
Brocade SilkWorm 32x0
Brocade SilkWorm 38x0 and SilkWorm 4100
McDATA Sphereon 4500
DAS
Storage
Dell/EMC CX200, CX400, and CX600
Dell/EMC CX300, CX500, and CX700
Dell/EMC AX100

Figure 1. Dell component test matrix

databases storing hundreds of terabytes of data. Regardless of their size, databases require high data availability to support mission-critical deployments. Integrating the Dell/EMC AX100 storage array into an Oracle® automatic disk-based backup and recovery system can enable enterprises to achieve online data recovery at a reasonable cost.

Configuring Oracle's flash recovery area on an economical storage unit such as the AX100 can help ease management through retention policies and enable automated management of backup space, redo logs, and other recovery-related files on disk-based storage. This article describes two example configurations that combine Dell/EMC storage arrays with existing data storage systems to support the Oracle Database 10*g* flashback area, thereby achieving cost-effective online data backup. The first configuration combines Oracle Database 10*g* Standard Edition (SE) Real Application Clusters (RAC), the Microsoft® Windows Server™ 2003 OS, and a Dell/EMC AX100 storage array; the second configuration invokes a larger, enterprise-class system with a Dell/EMC CX300 storage array, Oracle Database 10*g* Enterprise Edition, and the Red Hat® Enterprise Linux® AS 3 Quarterly Update 4 (QU4) OS.

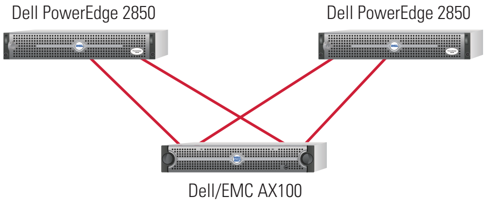

Figure 2. Dell/EMC AX100 in a DAS configuration

Dell/Oracle configuration 1: entry-level database

The Dell/EMC AX100 is designed to provide easy-to-use, low-cost storage in a direct attach storage (DAS) configuration, as shown in Figure 2, or in a storage area network (SAN) configuration. The AX100 is well suited for small workgroups, medium-sized businesses, and branch offices of large corporations. Tested and validated Dell configurations using the AX100 and a Microsoft Windows® OS can be excellent for small and medium-sized databases. Additionally, enterprises starting with these solutions can achieve low cost per unit of storage capacity for a cluster.

For small-to-medium businesses and budget-conscious enterprises, Dell has introduced a tested and validated configuration for Oracle Database 10*g* SE RAC on Windows Server 2003, Standard Edition, with Service Pack 1 (SP1).[1] The configuration uses two Dell PowerEdge™ 2850 servers and the Dell/EMC AX100 (see Figure 3). This configuration was selected for the test environment because it characterizes an entry-level offering wherein cost and ease of use are paramount. Although this configuration represents an entry-level scenario, the existing hardware can be integrated into a more complex, scaled design.

Servers	Two Dell PowerEdge 2850 servers
OS	Microsoft Windows Server 2003, Standard Edition, with SP1
Oracle database software	Oracle Database 10*g* SE, Release 1 base (10.1.0.2) plus the 10.1.0.4 patch set
Fibre Channel switches	None (servers directly attached to storage)
HBAs	Two QLogic QLA2340 Fibre Channel adapters
Storage	Dell/EMC AX100

Figure 3. Components of configuration 1

[1] For more information about this configuration, visit www.dell.com/oracle and click "Dell Supported Configurations." Then select the "Oracle Database 10*g*" tab, and in the "Microsoft Windows 2003 SP1" section, click "Oracle Database 10*g* Standard Edition with Real Application Clusters."

Dell/Oracle configuration 2: database cluster

Combining Oracle Database 10*g* Enterprise Edition, Red Hat Enterprise Linux AS 3, and the Dell/EMC CX300 storage array can yield a tested and validated configuration suitable for a larger enterprise than that envisioned for the first configuration. For its database, this configuration relies on a Dell/EMC CX300 storage unit (see Figure 4). Designed for slightly larger and more performance-critical databases than those for which the AX100 is intended, the Dell/EMC CX300 is an entry-level RAID storage system with 2 Gbps Fibre Channel host interfaces and capacity for as many as 60 Fibre Channel disk drives. Capable of operating as DAS, in a SAN, or attached to a Dell PowerVault™ network attached storage (NAS) system, the CX300 can provide up to 18 TB of FC2 raw storage capacity or up to 27 TB of Serial ATA (SATA) storage capacity using Dell/EMC 2 GB Disk Array Enclosures.

Tested and validated Dell configurations using the AX100 and a Microsoft Windows OS can be excellent for small and medium-sized databases. Additionally, enterprises starting with these solutions can achieve low cost per unit of storage capacity for a cluster.

Scaling the tested and validated Dell/Oracle configurations

Configurations 1 and 2 differ significantly. Although the end goal involves a similar configuration of Dell/Oracle and Dell/EMC products, the motivation to scale these configurations is dissimilar.

Scaling configuration 1: enhanced performance

Configuration 1 represents an entry-level Oracle database that is easy to deploy and to which storage can be added at a relatively low cost. However, improving Oracle cluster performance requires hardware additions. Consequently, to scale configuration 1, best practices recommend maximizing the use of initial hardware while seeking to improve performance and response times.

To accomplish these requirements, a Dell/EMX CX300 was added to the initial hardware configuration. The database that originally resided on the AX100 was then moved onto the CX300, while the flashback area remained on the AX100. This modification improved database performance because the CX300 provided greater storage speed and cache size compared to the AX100, and the initial hardware could then be used for disk-based recovery features.

If an enterprise currently uses the QLogic QLA200 host bus adapter (HBA) card in a fabric environment, Dell best practices recommend that administrators employ an HBA architecture that can offload tasks from the host node. They should also select a card with higher buffer credit levels than the QLA200 can provide—although that setup was not part of these test designs. Often the AX100 is coupled with QLA200 cards, which are designed for entry-level systems. Although moving to new HBAs requires initial hardware repurposing, replacing the entry-level HBAs can bring worthwhile performance improvement.

Scaling configuration 2: optimized cost

Configuration 2 assumes that a Dell/EMC CX300 storage system has already been integrated into an Oracle cluster. Enterprises can help reduce storage costs by adding an AX100 array to the cluster for disk-based backup and recovery.

Although an extra disk array enclosure (DAE) can be added to the cluster to increase storage capacity, this test design chose the goal of cost optimization. To minimize costs and add disk-based recovery, Dell engineers designed configuration 2 to leverage existing components by keeping the database logical units (LUNs) on the CX300. In contrast, the flashback and recovery area was moved to the AX100. This modification anticipated database performance at levels comparable to those of the initial configuration 2 while leveraging the added AX100 for disk-based recovery features, thereby dedicating disk space and storage processor cycles entirely to the database files.

The cost trade-offs for adding a DAE to the CX300 versus adding an AX100 were simple to analyze:

- Cost of DAE + (number of drives × cost of drives for DAE)
- Cost of Dell/EMC AX100 + (number of drives × cost of SATA drives for AX100)

Configuration 2 represents a scenario in which the addition of an AX100 can provide cost advantages. Because the Oracle flashback

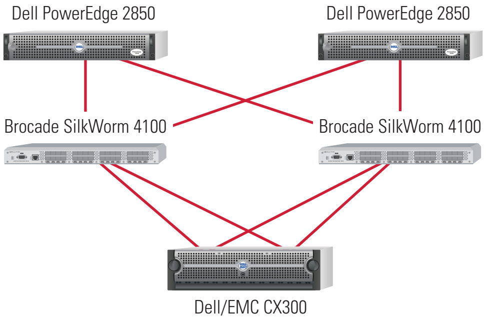

Figure 4. Dell/EMC CX300 in a SAN configuration

Servers	Two Dell PowerEdge 2850 servers
OS	Red Hat Enterprise Linux AS 3 QU4
Oracle database software	Oracle Database 10*g*, Release 1 base (10.1.0.2) plus the 10.1.0.4 patch set
Fibre Channel switches	Two Brocade SilkWorm 4100 switches running at 2 Gbps
HBAs	Two QLogic QLA2340 Fibre Channel adapters per node
Additional software	EMC PowerPath software
Storage	Dell/EMC AX100 and CX300

Figure 5. Components of configuration 2 scaled for cost and performance

recovery and backup functions do not require high random I/O reads and writes, backup and flashback storage can be placed on the AX100, which is slower than the CX300, without performance degradation to the database. Additionally, the cache on the CX300 can be leveraged to greater potential than in its initial configuration because sequential reads and writes do not fill up the cache.

In addition, Oracle enhanced the disk-based backup and recovery capability of its database software in Oracle 10*g*. Oracle added several features to take advantage of disk-system economics. These features allow leveraging additional cost-effective storage to help reduce expensive system outage time and associated costs.

Analyzing the optimal Dell/Oracle configuration

To optimize both cost and performance, a model configuration would involve both an AX100 and a CX300. In such a configuration, the CX300 stores the data files while the AX100 houses the flashback recovery area. To accomplish this, the Dell engineers placed the database disk groups on the CX300 LUNs. To enable flashback, they allocated a flash recovery area on the AX100. In tests performed by Dell engineers in October 2005, the AX100 was exclusively designated for the retention and backup components of the data file images, redo logs, and control files.

The scaled configuration comprised three layers: the server layer with two Dell PowerEdge 2850 servers; the switch layer with two Fibre Channel Brocade SilkWorm 4100 switches running in 2 Gbps mode; and the storage layer with both switches connected to the CX300 and AX100 arrays (see Figure 5).

High-availability design principles

Dell's Best Practices Program for Oracle encourages high-availability design principles. Consequently, the scaled design described in this article used redundant components at each level. Because the initial configuration already employed two PowerEdge 2850 servers, no additional servers were required. The SAN fabric environment used multiple Fibre Channel switches that allowed the use of more than one back-end storage unit while adding redundancy. Additionally, the presence of EMC PowerPath® software and multiple QLogic QLA2340 adapters in each PowerEdge 2850 server helped ensure that, if an HBA failed, alternate paths would be available from each node.

Figure 6 shows a two-node Oracle RAC cluster deployed in a high-availability configuration. Such a configuration should minimally contain the following components:

- Two servers
- Two HBAs per server
- Two Fibre Channel switches
- Multiple storage systems

Similar to a DAS configuration, each node's HBA in this high-availability configuration can access both storage processors; preconfigured zoning allows the HBAs to do so. In addition, with redundant switches, each server can still reach virtual disks if a switch fails because both storage processors are connected to alternate switches.

Configuration 1 versus the optimized scaled implementation

The scalability goals of the first configuration included reuse of initial hardware coupled with increased performance. Because the database was moved to the CX300, performance of the scaled configuration was significantly better than that of the AX100 alone, as test results show.[2] The Dell test team employed a benchmark that was similar to TPC-C from the Transaction Processing Performance Council (TPC), and these benchmark results revealed a significant performance increase for the Oracle cluster after the CX300 was added.[3]

Configuration 2 versus the optimized scaled implementation

In configuration 2, adding an AX100 easily achieved the goal of cost optimization for expanded storage capacity. However, the

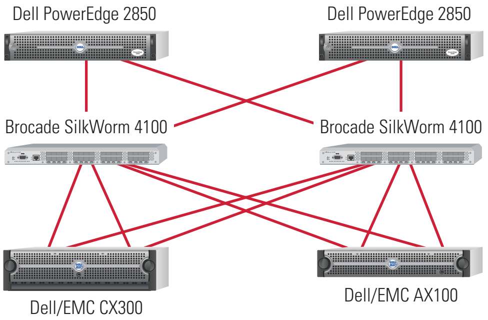

Figure 6. Dell/EMC CX300 and AX100 in a scaled configuration

[2] For more information and a detailed comparative performance analysis of the two systems, please see the associated white paper on this topic at www.dell.com/oracle. Search for "Dell Technology White Papers," and then click on the link. From the list of articles, click on "A Detailed Look into Scaling an Oracle RAC System with Disk-Based Backup and Recovery."

[3] See the associated white paper for test results.

performance of the scaled configuration needed to be comparable to that of the CX300 alone. To verify this, the Dell engineers first tested the CX300 alone and then tested the AX100 and CX300 combined configuration.

For the benchmark test of configuration 2, four LUNs were created in the CX300 and shared between the cluster nodes. A 1 GB LUN was assigned for storing Oracle Cluster Recovery (OCR) and voting-disk information. A 133 GB LUN was used for storing database files. Two 20 GB LUNs were used for flash recovery. The same storage processor was assigned as the default owner of the database and the flash recovery LUNs (to analyze the effect of offloading backup and flash recovery cycles to the AX100). Tests results showed that the CX300 performed relatively well by itself.

After the initial configuration 2 system was benchmarked, it was migrated to the optimized scaled configuration, which included both the AX100 and CX300. Two LUNs (1 GB and 133 GB) were created in the CX300 and two LUNs (20 GB) were created in the AX100 for the cluster nodes. The 1 GB LUN in the CX300 was assigned for storing OCR and voting-disk information; the 133 GB LUN was used for storing database files. The 20 GB LUNs in the AX100 were used for flash recovery.

Next, the Dell testing team subjected the scaled configuration 2 to a TPC-C benchmark–like workload—a series of tests using 200, 400, and 600 users. The tests showed that the AX100 and CX300 combination performed better in some cases than the CX300 alone, and that no serious degradation resulted from introducing the AX100 system to a cluster with a CX300.

The slight performance increase may have been attributable to the CX300 cache. The CX300 divides the cache into two components—one area is designated as a share cache for all LUNs, and a second area is broken into smaller segments wherein each segment is dedicated to a specific LUN. Because the flashback LUN was removed from the CX300, the entire cache was dedicated to the database. Consequently, small performance improvements might have resulted until the SAN cache was filled.

Balancing performance and cost in Dell/Oracle environments

The trade-off between speed and cost-effectiveness of disk-based backup and recovery is a factor that every enterprise must constantly weigh. Dell engineers analyzed two Dell/Oracle tested and validated configurations and scaled them to a common solution configuration. The goal in scaling the first configuration was to increase the solution's transaction rate. The goal in the second configuration—where database transaction rate was acceptable—was to add cost-effective disk-based recovery without sacrificing performance. In both cases, a common goal was to retain as much of the initial hardware as possible.

An ideal storage system provides high performance (like that provided by the Dell/EMC CX300) at a low price point (like that of the AX100). The final scaled database approach chosen for both initial configurations was to leverage two distinct storage units. Enterprises that originally deploy the AX100 alone and need to scale out can add a CX300 and continue to use the AX100. In the Dell tests, this configuration proved to add significant performance to the database with minimal hardware reprovisioning. Conversely, enterprise deployments that start with a CX300 can augment storage for disk-based recovery at a low price point with the addition of an AX100. This scaled solution demonstrated the ability to maintain the performance level of a CX300 system while enjoying the cost advantages of the entry-level AX100 storage system.

By combining an AX100 and a CX300 within an IT environment, enterprises can scale tested and validated Dell/Oracle configurations without sacrificing performance or cost. This cost-effective, phased scalability strategy can be instrumental to an enterprise's growth.

Mahmoud Ahmadian is an engineering consultant with the Database and Application Engineering Department of the Dell Product Group. Mahmoud has an M.S. in Computer Science from The University of Houston, Clear Lake.

Ujjwal Rajbhandari is a systems engineer in the Database and Application Engineering Department of the Dell Product Group. He is responsible for validating Oracle RAC solutions on Dell PowerEdge servers and Dell/EMC Fibre Channel storage. 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.

David Mar is a senior software engineer in the Database and Application Engineering Department of the Dell Product Group. His principal focus is developing deployment strategies for Oracle-based solutions on Dell PowerEdge servers, PowerVault storage units, and Dell/EMC storage units. David has a B.S. in Computer Engineering and an M.S. in Computer Science from Texas A&M University.

Chethan Kumar is a systems engineer and advisor in the Database and Applications Group at Dell. He has an M.S. in Computer Science and Engineering from The University of Texas at Arlington.

FOR MORE INFORMATION

Dell/Oracle:
www.dell.com/oracle

Dell/EMC storage arrays:
www.dell.com/storage

Sizing Dell PowerEdge Servers for **SAP Business One**

Dell™ PowerEdge™ servers can be reliable, robust platforms for the SAP® Business One application. However, when these platforms support other applications in addition to SAP Business One, performance can degrade. This article examines tests performed in two PowerEdge server–based environments to help determine how many users could adequately be supported while providing high performance for all the applications.

BY ALEXANDER ARTHUR

Related Categories:

Application servers

Characterization

Microsoft SQL Server

Microsoft Windows

SAP

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

SAP Business One is an integrated, affordable business management application designed specifically for small and midsize businesses. It provides enterprises with a single system to automate critical operations, including sales, finance, purchasing, inventory, and manufacturing, and offers an accurate, up-to-the-minute picture of their business.

To derive a series of recommended configurations for deploying SAP Business One in a real-world production environment, a team from Summit Business Solutions, an SAP Business One partner, tested this software on Dell PowerEdge servers in January and February 2006. The SAP Business One application by itself performs well on Terminal Services, a component of Microsoft® Windows® operating systems. However, performance problems may arise when additional applications that integrate into SAP Business One run on the same server as SAP Business One. Although the issue of application stability has been largely addressed by Service Pack 1 (SP1) of SAP Business One 2005, this test project addressed the issue of sizing a suitable server to the appropriate number of users and functions. Using Dell PowerEdge servers, the test team was able to get a sizeable number of users running SAP Business One with multiple applications integrated.

Configuring the test environments

The team tested two environments: SAP Business One running in a Terminal Services environment in an average-size deployment with 10 to 20 users; and SAP Business One running on Microsoft Small Business Server 2003 Premium Edition in a small business environment with 3 users.[1]

SAP Business One in an average-sized environment

In the first scenario, the test team ran the applications in a Terminal Services environment on two Dell PowerEdge 1850 servers, each with a dual-core Intel® Xeon® processor with up to 8 GB of RAM, dual 1 GB network interface cards, and two 73 GB SCSI hard drives. The team added the two servers to a Microsoft Active Directory® domain as member servers with access to a network printer and other standard business accessories typically connected to business workstations.

One server ran Microsoft Windows Server™ 2003 and Microsoft SQL Server™ 2000 with all the security patches

[1] For more information about the test configurations, visit www.dell.com/sapb1. A Webinar series for SAP Business One using Dell and Microsoft SQL Server and a demo copy of SAP Business One are also available at this Web site.

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

SIMPLIFYING AND ENHANCING RETAIL OPERATIONS WITH CITIXSYS iVEND RETAIL, SAP BUSINESS ONE, AND DELL HARDWARE

The speed and pace of today's retail environment demand systems and processes that help evaluate, plan, and respond quickly to client needs. A retail management system must accurately record sales and maximize the customer experience at the front end, while providing quality data behind the scenes for making quick, informed decisions about ordering, replenishment, and logistics.

The iVend Retail point-of-sale (POS) and retail management software from CitiXsys Technologies is designed to simplify retail operations and enhance the information available for store management and business planning. The iVend Retail software system powered by SAP Business One has been optimized for Microsoft SQL Server and Dell PowerEdge servers operating at both store and head-office levels. This optimization is critical because moving updated information between stores and the head office quickly and efficiently is a key factor for decision making in retail operations. Too often there is a disconnect between what is happening at the POS and what management needs to know to manage inventory, security, customer satisfaction, and profitability.

iVend Retail provides a comprehensive solution that combines easy-to-learn, easy-to-use functionality at the POS terminal with a full range of dynamically updated operations information conveyed back to the head office. Inventory, security, promotions, commissions, and many other store-specific requirements are set up by the head office to maintain control, but managed at the store level through a central store server. The iVend Retail head office administration is built into SAP Business One rather than into a separate head office system, which cuts down one level of synchronization and replication.

In terms of architecture, each POS terminal links to a local store server, which in turn links to the head office through various replication options—wide area network, FTP, dial-up, or CDs. The replication frequency schedule is determined by the administrator and can range from real time to once each day. Each POS terminal can continue to work even if the store server goes down; when the server is functional again, the data collected at the POS is transmitted at the next replication event. Standard options for backup mechanisms include CD-RW drives, database backup possible at every POS, and SQL backup possible at each store location

Each POS database is uploaded to the store server database, and the store server databases are uploaded to the FTP repository server through the Common, Head Office, Stores, and Status folders. The Status folder is used to signal the status of the synchronization process. The FTP repository server then downloads to the iVend Retail management server which, because it is built inside the accounting application, is also the server running SAP Business One. Data originating at the head office is uploaded to the FTP repository server, and files move to the Synchronized folder automatically upon completion of the synchronization process. The files in the Synchronized folder are periodically deleted through a simple batch process. Data is downloaded onto the store server from the FTP repository server. iVend Retail data replication provides centralized control from the head office over master data, support for tracing the synchronization process at the record level for every destination site through a detailed error log, and a set of distribution rules that support the store-specific transfer of data.

Dell PowerEdge servers deployed for the application/database and FTP repository servers must have dual Intel Pentium or Intel Xeon processors at 1.6 GHz or higher and 2 GB of error-correcting code (ECC) RAM. Depending on customer requirements, the hard drive configuration should be RAID-0, RAID-1, or RAID-5, with a minimum of one system disk and three RAID disks. The OS platform must be Microsoft Windows 2000 Server, Standard Edition, or later. Required software includes Microsoft SQL Server 2000 with mixed mode authentication (SQL Server and Windows), Microsoft Data Access Components 2.7, and Microsoft .NET Framework 1.1. The FTP repository server also requires Microsoft Internet Information Services 5.0. For the POS terminals, minimum requirements include one Intel Pentium III processor at 600 MHz, 128 MB of RAM, Windows XP Professional with SP1or later, and .NET Framework 1.1. A variety of Dell POS monitors, printers, scanners, keyboards, and cash drawers are supported as well.

iVend Retail from CitiXsys Technologies is designed to work in a wide range of small and medium-sized enterprises, from one store with two terminals to 100 stores or more. By using iVend Retail with industry-standard Dell servers and POS system technology and the SAP Business One application, enterprises can streamline store operations and refine business processes—helping to reduce inventory risk, securely manage employees and purchase transactions, and improve analysis and planning.

loaded. This server, called Dell-SAP1, acted as a database, file, and Web server. The test team also installed the license for SAP Business One as well as the server components of SAP Business One. The second server, Dell-SAP2, had Microsoft Windows Server 2003 with Terminal Services activated. It also had all the relevant security patches for Windows Server 2003 loaded.

For remote access and application publishing, the test team used the Provision Management Framework–Enterprise Edition from Provision Networks. This tool let the team publish and load balance SAP Business One, with instant delivery over Remote Desktop Protocol (RDP), using features such as Seamless Windows and Session-Sharing. The Provision Networks tool used Microsoft SQL Server on Dell-SAP1 as the database engine as well as the Web access component for the terminal server. Finally, the team installed iBOLT Special Edition for SAP Business One from Magic Software on Dell-SAP1 to perform the functions of recalculating a quote price and inserting the new calculated price in the sales order module.

On the second server, Dell-SAP2, which had Terminal Services activated, the test team installed the client components of SAP Business One and other standard office applications such as Microsoft Office XP, Adobe Acrobat Reader, and a desktop fax client, allowing faxing of documents from the desktop through the network to a fax machine. This server also had Fixed Asset and Outlook Integration for SAP Business One installed. In addition, XL Reporter was loaded and all the printers defined on the server were mapped to a print server outside the terminal server environment—hence, no local spooling.

After successfully testing this environment, the test team replaced the memory on both servers with 4 GB of double data rate 2 (DDR2) memory sticks and tested this environment again.

SAP Business One in a small business environment

For the small business environment, the test team used a Dell PowerEdge SC1420 server—one of Dell's entry-level servers configured with 2 GB of RAM and adequate hard drive space—which ran Microsoft Small Business Server 2003 Premium Edition. The team installed SAP Business One as well as all the applications used in the average-sized test environment to test up to three users. For the desktops, the team used standard Dell Dimension™ desktops with Intel Pentium® 4 processors and 512 MB of RAM. They also tested the desktops configured with 1 GB of RAM to identify any performance differences.

On the PowerEdge SC1420 server, the test team increased the memory to 8 GB and configured Microsoft SQL Server to use a maximum of 4 GB of RAM. The team then installed iBOLT, XL Reporter, Microsoft Office XP, Adobe Reader, and the server components of SAP Business One. All of the software had the latest patches, including patch 8 on 2005A for SAP Business One.

The functions performed in the test environment included generating quotes, invoice processing, and receiving payment. The test team also performed data maintenance, such as changing addresses and telephone numbers. They built enough data to be able to run manufacturing resource planning (MRP) on fairly sizeable data and run manufacturing processes as well. On the accounting side, they entered typical office journal entries using scripts learned in the training class for SAP Business One.

Finally, the test team installed Provision Networks Web-IT component and published the application over the Internet, allowing remote users access to the applications as well as the ability to print reports to a local printer defined as a slave of their workstations. For Internet bandwidth, broadband (DSL or cable) was assumed to be sufficient for the test.

Testing performance of an average-sized deployment of SAP Business One

The team began testing the average-sized environment by simulating 10 concurrent users on Dell-SAP2. The Terminal Services technology helped improve navigation and the responsiveness of the SAP Business One application because only one copy of the application is loaded into memory. All of the other applications installed also performed much better with the Terminal Services technology than without it. The test team then activated the Max-IT module from Provision Networks to improve any remaining performance issues and to address some of the environment-related issues that were encountered. Max-IT helped resolve many of these issues, and as a result, XL Reporter and Outlook Integration were more responsive than before Max-IT was activated. Figure 1 shows virtual memory usage and savings for each application on the Dell-SAP2 server.

The test team then created an entire set of quotes for iBOLT to recalculate the sales order prices by dividing by 10 and inserting the new calculated value as a sales order in SAP Business One. While this was occurring, users were accessing SAP Business One and either creating, printing, or querying orders. No performance degradation was observed. The test team started to print reports through

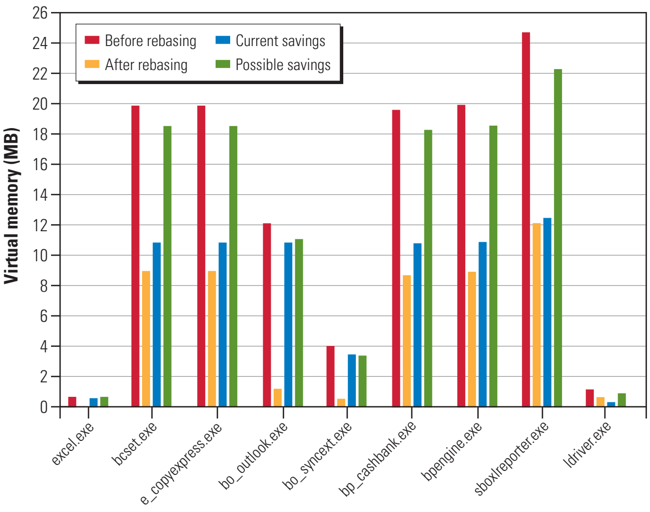

Figure 1. Per-application memory usage and savings in an average-sized terminal server environment for SAP Business One

the in-house print server and again there was no notable performance issue. The server and the complete environment remained up and performed satisfactorily. This proved that the Dell server (Dell-SAP1) was more than capable of handling the responsibilities of a file server, a license server, a Web server, and a database SQL server as well as that of an application server.

With Max-IT running, the test team decided to increase the workload by simulating up to 15 concurrent users, adding one user at a time. Performance remained acceptable at 15 users. As the team increased the number of users to 20, they began to notice some performance changes. Around the 17-user level, the change in application response was noticeable. XL Reporter and Outlook Integration, in particular, started showing very noticeable performance changes. However, the environment continued to function. For remote users, the only noticeable change was a screen refresh on some of the graphics.

In general, the logon process was ahead of the SAP splash logo but that was not an issue with users. Some of the graphs from XL Reporter took longer to refresh than they had with fewer simulated users. Data navigation was not affected as the number of users increased, and the system continued performing well even when performing computations or posting.

Running MRP

In all test scenarios, the test team noticed a significant change in performance when MRP ran. However, most enterprises typically do not run MRP in the middle of the day—they either run it at the beginning of the day to plan their day or at the end of the day to plan the next day. So this effect on performance should not necessarily be viewed as a concern. In addition, with just one user running MRP, the performance was fast.

Comparing performance expectations with test results

The average-sized deployment performed as the test team expected. XL Reporter, however, did not perform as expected—it consumed a significant amount of resources. IT organizations should consider this fact when sizing a Dell PowerEdge server and the number of XL Reporter users. In contrast, the Dell servers and SAP Business One kept up with demands. The Provision Networks software allowed the test team to get that extra margin of performance from the hardware, hence reducing total cost of ownership.

Testing performance of a small business deployment of SAP Business One

The Dell PowerEdge SC1420 server handled the workload of SAP Business One and the other applications surprisingly well—especially considering that Microsoft Small Business Server Premium Edition was used to run the entire network, including other workstations and servers on the network. Performance was somewhat sluggish but acceptable for a three-user environment. Screen refresh sometimes took an extra second, especially the posting and computing screens. For data navigation and other simple tasks, the PowerEdge SC1420 performed as well as needed to browse through data intelligently. Upgrading the client memory to 1 GB enhanced the performance of the applications; in particular, the performance of XL Reporter was much better than the results achieved with the 512 MB configuration.

Identifying best practices for sizing deployments of SAP Business One

The average-sized server environment is well suited for most deployments of SAP Business One. Although the Dell PowerEdge 1850 server with the maximum specifications proved capable of handling the workload, placing all applications on one server is not recommended. Instead, best practices recommend using two PowerEdge 1850 servers with average specifications—one to act as the database and license server and the other to act as the terminal server deploying the applications to the workstations. In addition, instead of having all 20 users on one terminal server, IT administrators can deploy multiple small terminal servers and load balance them with tools such as those from Provision Networks. This can then be scaled up for large environments, with all the terminal servers connecting to an enterprise server on which the databases are housed. The PowerEdge 1850 is well suited for such deployments. Microsoft SQL Server can house SQL databases but should not perform other functions on the network. A PowerEdge SC server can be used to act as an authenticating server to run the network and perform other functions such as file serving.

With the Microsoft Small Business Server environment, the results showed that the entry-level Dell PowerEdge SC1420 server can be a reliable platform for a robust deployment of SAP Business One. Small businesses can deploy SAP Business One without spending a significant amount on infrastructure. However, in this case, workstations should be upgraded and administrators must be creative as to when to run certain tasks such as MRP. Administrators should set appropriate expectations for performance, but for IT organizations with tight budgets, this deployment can be cost-effective.

Alexander Arthur is the vice president of professional services for Summit Business Solutions (www.summitt-biz.com), an SAP Business One partner. He has been deploying server-based computing solutions globally for more than 10 years.

FOR MORE INFORMATION

Dell and SAP Business One:
www.dell.com/sapb1

SAP Business One:
www.sap.com/solutions/sme/businessone

SAP
BUSINESS
ONE

ORACLE
DATABASE
STANDARD EDITION ONE
ORACLE