DELL POWER SOLUTIONS • AUGUST 2005

DELL™

AUGUST 2005 • $12.95

POWER SOLUTIONS

THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS

with Dell Systems Management

Inside This Issue:

Technologies Behind the New Dell OpenManage IT Assistant Console

Guidelines for Assessing Data Center Power and Cooling Requirements

Approaches to Business Continuity for Microsoft Exchange

Take Control with Dell Systems Management

Gigabit Ethernet
Multi-Gigabit Ethernet

DELL™
POWER SOLUTIONS
THE MAGAZINE FOR DIRECT ENTERPRISE SOLUTIONS
AUGUST 2005

The Dell OpenManage suite, a comprehensive set of industry standards–based tools included with Dell PowerEdge systems, is designed to automate server and client management functions—thereby helping to simplify IT operations throughout the complete system life cycle. In addition, this Dell tool set works with a wide variety of integrated, standards-based management solutions from key Dell partners such as Altiris and Microsoft. Working together, these solutions can enable administrators to respond quickly and flexibly to fast-changing IT requirements, freeing valuable resources to focus on strategic business initiatives.

Simplifying data center operations with Dell's flexible, industry-standard systems management suite

TABLE OF CONTENTS

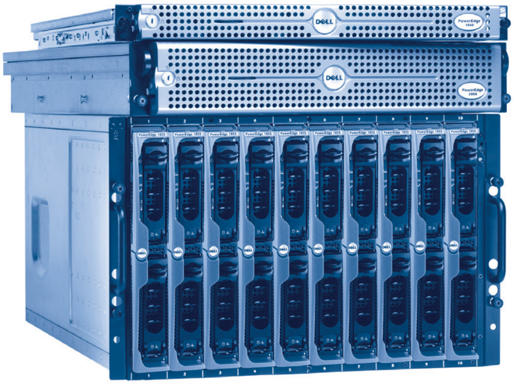

From top: Dell PowerEdge 1850, PowerEdge 2850, and PowerEdge 1855 servers

DATA CENTER TECHNOLOGY

BUSINESS CONTINUITY

STORAGE TECHNOLOGY

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

EDITORIAL

EDITOR-IN-CHIEF | Tom Kolnowski
MANAGING EDITOR | Debra McDonald
SENIOR EDITORS | Liza Graffeo, Vicki Van Ausdall
CONTRIBUTING AUTHORS | Tim Abels; Chris Auger; Kimber H. Barton; Bradley Branson; Bernard Briggs; Matthew Brisse; Alan Brumley; Michael Brundridge; Balasubramanian Chandrasekaran; Babu Chandrasekha; Swee Chew; Jamee Cushman-Ramsey; Narayan Devireddy; Puneet Dhawan; Lisa Filemyr; Roger Foreman; Jyeh Gan; Erin Geaney; Richard Golasky; Brian Goldstein; Alex Gorbansky; Saurabh Gupta; Landon Hale; Jordan Hargrave; Stephanie Hartley; Steve Heracleous; Ruoting Huang; Dave Jaffe, Ph.D.; Harish Jayakumar; Tom Johansmeyer; Nadine Latief; Qingsong Li; Kit Lou; Zafar Mahmood; Abhishek Mehta; Todd Mitchell; Krishna Mohan; David Moss; Chandra S. Mugunda; Todd Muirhead; Terry Myers; Vijay Nijhawan; Utpal Patel; Karthik Rajagopalan; Anusha Ragunathan; Scott Rosen; Paul Rubin; Tobin Rysenga; Bharat Sajnani; Rohit Sharma; Uday Datta Shet; Sudhir Shetty; Nik Simpson; Dennis Stultz; Arun Taneja; Ahmad Tawil; Teresa Taylor; Scott Thomas; Hector Valenzuela; Kevin Winert; and Man Xiong

ART

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

ONLINE

WEB PRODUCTION | Brad Klenzendorf

SUBSCRIPTION SERVICES

CIRCULATION | Amy Hargraves, Becky James

Subscriptions are free to qualified readers who complete the online subscription form or the subscription reply form on page 87 in this issue. 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, 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*, Mailstop 8456, 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.

August 2005

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.

THIS CHANGES EVERYTHING.

DELL OPENMANAGE™/ALTIRIS®
MANAGEMENT SUITE for DELL SERVERS

OS, Application and Hardware Management. Just one console.

Now, Dell™ PowerEdge™ administrators only need one console to deploy, manage, monitor, patch and update software and hardware for Microsoft® Windows® and Red Hat® Linux® environments. With all those features fully integrated, Dell OpenManage with Altiris Management Suite for Dell Servers helps get systems up and running fast, saving IT time and resources.

Take some time to see for yourself.

Visit dell.com/altiris6 today for a demonstration and whitepaper.

SERVER MANAGEMENT IN PROGRESS...

☑ APPLICATION UPDATE
☑ SECURITY PATCH
☑ OS INSTALLATION
☑ HARDWARE UPDATE
☑ PERFORMANCE REPORT

altiris®

GET MORE OUT OF CHANGE. GET MORE OUT OF NOW. DELL

Click www.dell.com/altiris6 Call 1.866.212.9344 toll free

Dell cannot be responsible for errors in typography or photography. Dell, PowerEdge, Dell OpenManage and the Dell logo are trademarks of Dell Inc. Altiris is a registered trademark and Server Management Suite is a trademark of Altiris, Inc. in the United States and other countries. Microsoft and Windows are registered trademarks of Microsoft Corporation. Red Hat is a registered trademark of Red Hat, Inc. Linux is a registered trademark of Linus Torvalds.

DATABASE TECHNOLOGY

VIRTUALIZATION TECHNOLOGY

SYSTEM ARCHITECTURE

ADVERTISER INDEX

ONLINE XTRA

DELL POWER SOLUTIONS MAGAZINE ARTICLES AVAILABLE ONLY AT WWW.DELL.COM/POWERSOLUTIONS

WWW.DELL.COM/ POWERSOLUTIONS

Unattended Management of the Dell OpenManage Server Administrator Life Cycle

By Swee Chew and Bernard Briggs

Beginning with Dell OpenManage 4.3, IT administrators can enhance management of their installed base of Dell OpenManage Server Administrator software by using the Microsoft Windows Installer Service and Windows Management Instrumentation Service. This article provides technical guidance on achieving a high degree of automation in the Dell OpenManage life cycle.

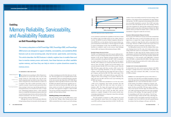

Enabling Memory Reliability, Availability, and Serviceability Features on Dell PowerEdge Servers

By Qingsong Li and Utpal Patel

The memory subsystems on Dell PowerEdge 1850, PowerEdge 2800, and PowerEdge 2850 servers are designed to support reliability, availability, and serviceability (RAS) features such as error-correcting code, chip fail correct, spare banks, and mirroring. This article describes the RAS features in detail and explains how they are enabled, how they can affect available system memory, and how they can help to minimize system downtime caused by memory errors.

A Technical Overview of the Dell PowerEdge 1855 Chassis and I/O Modules

By Michael Brundridge, Babu Chandrasekha, Jyeh Gan, and Abhishek Mehta

The Dell PowerEdge 1855 modular server system is a high-performance, highly integrated system. This article discusses various aspects of the shared components, their interconnections, redundancy, and the interfaces that can be used to configure the components.

Enabling Demand-Based Switching in Red Hat Enterprise Linux 4 on Dell PowerEdge Servers

By Jordan Hargrave

Demand-based switching (DBS) is a new technology that is designed to help minimize data center power and cooling requirements and ultimately lower IT costs. This article describes the DBS features and utilities available with Red Hat Enterprise Linux 4 to run on Intel Xeon processor–based Dell PowerEdge servers.

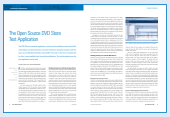

The Open Source DVD Store Test Application

By Dave Jaffe, Ph.D., and Todd Muirhead

The DVD Store e-commerce application, used as a test workload in many recent Dell white papers and demonstrations, has been released to the general public under the open source GNU General Public License (GPL). The code, in the form of compressed tar files, is now available from linux.dell.com/dvdstore. This article explains how the test application can be used.

UNSCRAMBLE THE PUZZLE ABOUT UNIX MIGRATION.

SIMPLIFIED SOLUTIONS. MAXIMIZED VALUE.

Work through the complexity of UNIX® migration with the powerful alliance of Dell and Microsoft. We build your migration solution to order, so you can choose the services, guidance and training best suited to your unique challenges. This customization delivers the results you need to solve old business problems and tackle new ones—with more power and greater productivity. In fact, migration might lower your TCO as much as 15-20%[1]. Find out how we put all the right pieces in place. Call 1-866-871-9881 or visit www.DELL.com/MSmigration and get a free business case analysis on migrating to Dell | Microsoft® Windows Server™ 2003 solutions.

visit www.DELL.com/MSmigration or call 1-866-871-9881
for your free UNIX migration business case analysis

DELL™ Microsoft®

[1] Based on independent research by Alinean, Inc., customer case studies from January 2004 through August 2004 modeled with the Alinean UNIX Migration ROI Analyst v3.0 software to quantify customer investment, TCO savings and business benefits of consolidation and migration from UNIX to Microsoft Windows® for specific workloads.

Relational Bliss

With hefty reference books at hand—the likes of *Software Engineering: A Practitioner's Approach* and *Database System Concepts*—the *Dell Power Solutions* editorial team contemplated various ways to deliver an accessible, easy-to-use content categorization feature with a repository that would require minimal maintenance time. In the end, after skimming heavy chapters on user interface design, database system architecture, entity relationship diagrams, and database normalization, we settled on a relational database approach. Following a relatively short design and development cycle, the Related Categories repository was born and hosted on a Dell™ Precision™ workstation.

In our Austin-based editorial offices, we then embarked on the task of categorizing an ever-growing library of *Dell Power Solutions* content. Effective with the May 2005 issue, all articles appearing in the print and online editions have been published with a Related Categories box located in the margin of the first page, populated with as many as a dozen related categories (see Figure 1). Our editors carefully screen each article to select relevant categories from an ever-expanding list of more than 200 IT-centric topics that cover the landscape of best practices, key technologies, standards, products and services, technical terms, vendor offerings, and much more.

In addition, we scanned our back-issue archives and categorized *Dell Power Solutions* content from the November 2003 issue onward. Then, ever true to relational database form, we cross-referenced Related Categories listings across hundreds of articles (and associated editorial content, such as technical posters) into a single entity: the *Dell Power Solutions* Related Categories Index. Essentially a downloadable PDF file, this index is organized alphabetically by category and includes details such as subject/content, date, print page number, and—most importantly—a clickable URL for instant access to the editorial content. To visit the latest Related Categories Index, download the file at www.dell.com/downloads/global/power/power_index.pdf and then click on the respective URL to access specific *Dell Power Solutions* content or an entire back issue.

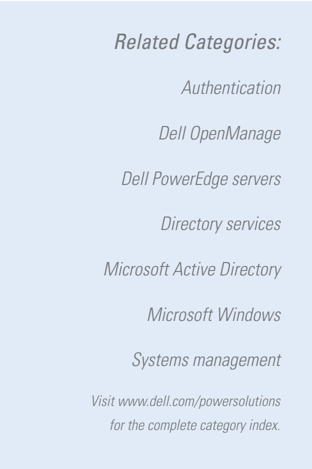

Related Categories:

Authentication

Dell OpenManage

Dell PowerEdge servers

Directory services

Microsoft Active Directory

Microsoft Windows

Systems management

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

Figure 1. Sample Related Categories box

In this August 2005 issue of *Dell Power Solutions,* we have added a healthy dose of content to the systems management category. Our cover story, "Take Control with Dell Systems Management," provides insight into Dell's standards-based approach to enterprise systems management, and four more supporting articles delve into the enhanced technology behind the new Dell OpenManage IT Assistant 7 management console. In all, 25 articles in the print edition and another five Web-exclusive articles feature in-depth technical coverage on systems management techniques, Microsoft® SQL Server and Oracle® databases, storage technology, the Linux® OS, and more. As you page through each issue of *Dell Power Solutions,* we hope the Related Categories Index will help you connect rapidly to the technical information you need most.

Tom Kolno

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

Take Control with Dell Systems Management

The Dell™ OpenManage™ suite, a comprehensive set of industry standards–based tools included with Dell PowerEdge™ systems, is designed to automate server and client management functions—thereby helping to simplify IT operations throughout the complete system life cycle. In addition, this Dell tool set works with a wide variety of integrated, standards-based management solutions from key Dell partners such as Altiris and Microsoft. Working together, these solutions can enable administrators to respond quickly and flexibly to fast-changing IT requirements, freeing valuable resources to focus on strategic business initiatives.

BY PAUL RUBIN AND TERRY MYERS

Related Categories:

Altiris

Baseboard management controller (BMC)

Change management

Dell OpenManage

Dell PowerEdge servers

Microsoft

Remote access controllers (RACs)

Remote management

Scalable enterprise

Standards

System deployment

Systems management

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

Systems management—and change management in particular—is becoming an unwieldy process for enterprises of all sizes. In many organizations, the IT infrastructure is tightly interwoven with business-critical systems that have been developed on an ad hoc basis over the years. To survive, administrators must simplify systems management tasks and enable fast, flexible response to diverse—and fast-changing—business conditions. Some enterprises have turned to expensive proprietary systems as a "magic bullet," while others have developed custom applications that meet immediate needs but can be difficult to scale as computing requirements grow.

In contrast, Dell's strategy for enabling the scalable enterprise is to provide cost-effective, industry-standard data center components that can be upgraded incrementally whenever and wherever necessary. And Dell's solution for the systems management dilemma is no different: to provide standards-based, interoperable systems management capabilities that enable administrators to choose tools that are tailored to their specific enterprise IT needs—and that scale easily to meet unpredictable growth requirements. Using a core set of tools offered by Dell, administrators can proactively manage basic functions of their servers and clients throughout the complete system life cycle. Augmenting the Dell tool set is a wide variety of integrated, industry-standard solutions that are designed to provide a comprehensive systems management framework.

This approach is designed to allow organizations to take charge of their IT infrastructure, laying the foundation for a standards-based IT environment in which system purchases no longer lock enterprises into specific management tools. The Dell systems management framework allows administrators the flexibility to use systems management capabilities that match their current IT requirements while enhancing the scalability of management environments to allow for growth and change. By implementing flexible, standards-based management solutions and enabling administrators to automate tasks and optimize IT resources, enterprises can free valuable resources to focus on strategic business initiatives. This article explains systems management challenges faced by today's IT organizations, various Dell products that can help address these challenges, and benefits of using the Dell open standards model.

Challenges of managing a complex IT environment

As today's IT infrastructures evolve and Intel® architecture–based systems perform much of the critical workload, a new set of management considerations is coming into play. As shown in Figure 1, a growing IT infrastructure presents major operational challenges: clients, servers, and storage systems must be deployed; changes and updates made; and systems regularly monitored for health and status. Emerging organizational needs must be supported through the reuse of existing IT resources, as well as the deployment of additional applications. Pressing concerns that administrators must contend with in today's dynamic IT environment include accelerated rate of change, increased scale and complexity, heightened security risks, and proliferation of tools and processes.

Accelerated rate of change. In a typical IT infrastructure, the one constant is usually change. Not surprisingly, change management has become a pain point for many IT organizations today. The velocity of change that must be managed can be characterized as a function of the number of clients, servers, and storage devices multiplied by the number of product updates and changes introduced by hardware and software providers. Given innumerable updates to hardware, operating systems, and applications on a regular basis—plus patches necessitated by security threats—many organizations have had to budget change management as a full-time job.

Increased scale and complexity. Most IT infrastructures must accommodate continual growth in the amount of data stored and managed as well as in the number and types of clients, servers, and storage platforms, operating systems, and applications. All too often, as more individual data center components are added to the computing environment, more unique management tools and processes are required to support that environment.

Heightened security risks. Ensuring IT security has become a daunting task. With new security issues a daily concern in operating systems and applications, administrators are faced with applying security updates to many systems to prevent system failures and intrusions to their networks. Securing IT infrastructure will likely remain a high-priority challenge, requiring topflight technologies and procedures.

Proliferation of tools and processes. As they contend with an ever-increasing number of clients, servers, storage systems, operating systems, and applications, IT organizations face the prospect of skyrocketing investments in staffing and training to support multiple management tools. This growth often calls for complex, tool-specific management processes that need to be communicated and maintained. Moreover, many organizations have implemented management tools to address specific functions and challenges, and some of these tools have overlapping capabilities so they are not fully utilized when combined—which can be another contributor to high overhead.

Many IT vendors take a proprietary approach to systems management—an approach based on their own hardware, software, and professional services. This approach may be beneficial to providers, but IT organizations are often left with complex, expensive multivendor environments and a variety of systems management tools. Such a situation can limit organizational response by forcing choices among specific vendors that may not address important enterprise IT and business challenges. To help simplify operations and advance best practices for enterprise-wide systems management, Dell promotes open standards that are designed to enhance flexibility and enable products to interoperate across complex, heterogeneous IT environments.

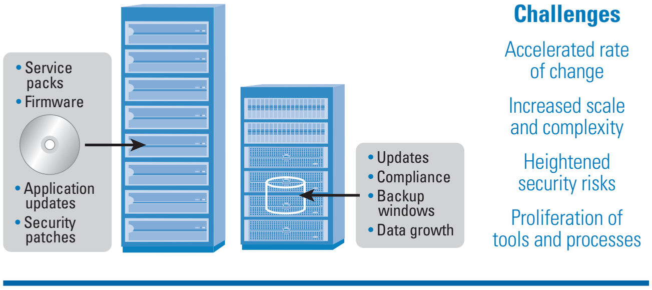

Figure 1. Operational challenges for growing IT infrastructures

Dell tools for managing the scalable enterprise

Dell has become a worldwide leader in high-volume, high-performance client, server, and storage systems through a consistent focus on industry standards. From active participation in standards-setting organizations to driving the development and integration of offerings from alliance partners, Dell has endorsed standards as the foundation of its scalable enterprise strategy. Industry-standard systems and software enable enterprises to respond quickly and flexibly to changing business requirements in cost-effective increments. In addition, industry standards can help maximize the IT investment by making it possible for older data center components to continue serving the infrastructure when new systems or upgrades are deployed. Dell leverages industry standards and the expertise of its leading systems management partners to provide a fully scalable framework for systems management solutions.

Systems management standards

To optimize existing resources and facilitate future growth, organizations need the capability to manage systems and processes efficiently and cost-effectively across a variety of hardware platforms, operating systems, and applications from multiple vendors. To that end, Dell's systems management strategy focuses on three goals that are designed to simplify IT operations through industry standards:

- Enable management of Dell's instrumented clients, servers, and storage platforms
- Champion systems management standards
- Support integrated, standards-based systems management solutions

Instrumented platforms can enable a complete life-cycle approach to managing systems and storage. Dell platforms are designed to provide the needed management information and control functions to support deployment, health status monitoring, fault recovery, and change management. Partnering with key systems management vendors such as Altiris and Microsoft, Dell can provide for an expanding range of solutions by offering toolkits designed to link Dell instrumentation with partner management software. This combination of Dell and partner technologies can lead to enhanced choices for systems management solutions.

Today's emerging standards are laying the groundwork for future management systems to deploy, monitor, and upgrade heterogeneous clients, servers, and storage systems from a central point of control (see Figure 2). Dell is helping to advance the standardization of client, server, and storage management by participating in numerous organizations, including the following:

- **Unified Extended Firmware Interface (UEFI) Forum:** Standardizing system firmware interfaces for client and server BIOSs
- **Intelligent Platform Management Interface (IPMI) Forum:** Developing standards for core server monitoring and control
- **Organization for the Advancement of Structured Information Standards (OASIS):** Driving the development, convergence, and adoption of e-business standards, including Web services for systems management
- **Storage Networking Industry Association (SNIA):** Standardizing the management of storage subsystems
- **Distributed Management Task Force (DMTF):** Creating standards for systems management information and access, including emerging standards for server management that are being developed by the Server Management Working Group (SMWG)

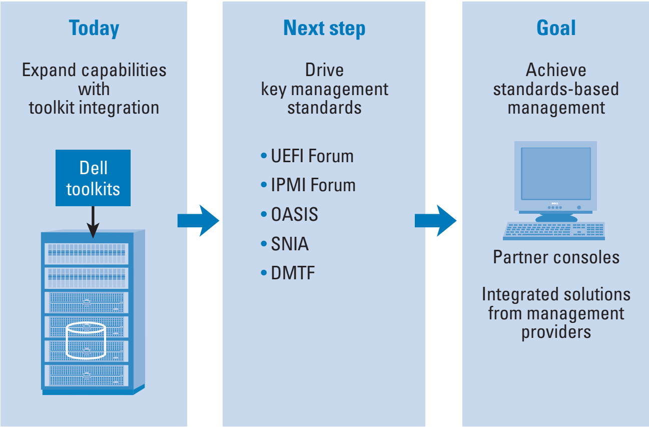
Today
Expand capabilities with toolkit integration
Dell toolkits
Next step
Drive key management standards
• UEFI Forum
• IPMI Forum
• OASIS
• SNIA
• DMTF
Goal
Achieve standards-based management
Partner consoles
Integrated solutions from management providers

Figure 2. Path to centralized hardware management

Active involvement in the preceding initiatives helps Dell to enable maximum standardization of systems and storage instrumentation in Dell data center components and to make provisions for a far-reaching, standards-based systems management framework.

In addition, standardization is expected to lead to broader, more versatile choices in systems management solutions, which can lead to a pivotal benefit: enabling providers to cost-effectively customize hardware and software products that are designed to meet the specific management needs of individual IT organizations. Broad use and acceptance of management standards may further develop the power of choice because an increasing number and variety of products typically creates competition that in turn can help lower cost, drive innovation, and enhance quality.

Dell OpenManage for Servers

By offering integrated hardware and software management capabilities, Dell OpenManage helps support server management solutions that are designed to provide organizations of all sizes with an advanced, reliable, and easy-to-manage IT infrastructure (see Figure 3). Dell OpenManage for Servers can simplify operations and help keep overall management costs low by enabling standardization and automation of server deployment, monitoring, and change-management tasks. Standards-based instrumentation built into the PowerEdge server BIOS, baseboard management controller (BMC), and remote access controller, as well as the Dell storage controllers, enables detailed inventorying, environmental monitoring, and server and storage health monitoring.

Dell PowerEdge servers are engineered for manageability, and Dell OpenManage uses PowerEdge server instrumentation that enables IT administrators to monitor and control server operations. The following sections describe Dell OpenManage components that address server deployment, day-to-day operations, and management of server change.

Server deployment with Dell OpenManage

Dell OpenManage facilitates automation of single-server deployment through Dell OpenManage Server Assistant, and automation of multi-server deployment through the Dell OpenManage Deployment Toolkit (DTK). The DTK is typically used in conjunction with deployment solutions provided by partners, including Microsoft and Altiris.

Dell OpenManage Server Assistant

Dell OpenManage Server Assistant (DSA) is a CD-based single-server deployment tool that is delivered with Dell PowerEdge servers on the Dell PowerEdge Installation and Server Management CD. DSA works in conjunction with OS media from Microsoft, Novell, and Red Hat to provide step-by-step guidance for server hardware configuration and OS installation. It includes tools and

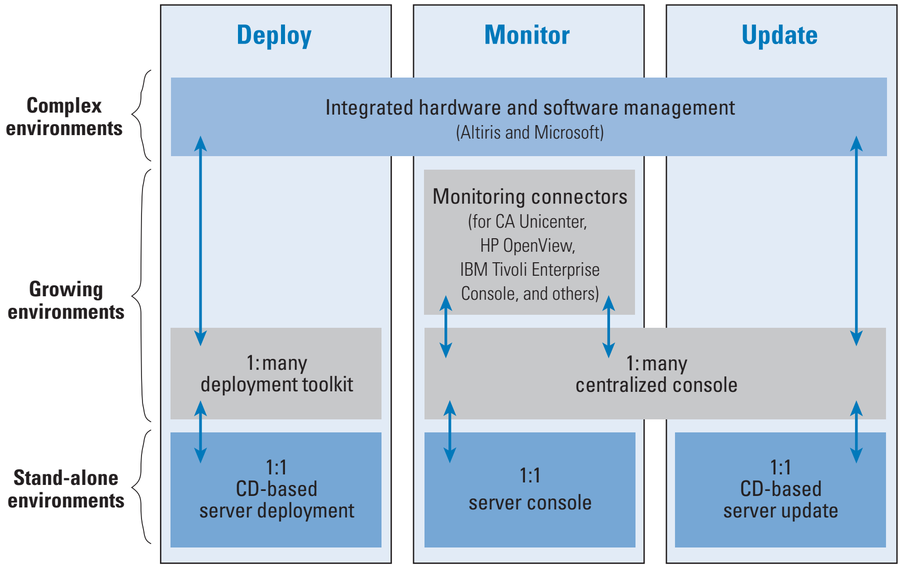

Figure 3. Dell OpenManage systems management framework

the latest drivers—including RAID and network interface card (NIC) drivers—to help speed the setup, configuration, and optimization of Dell PowerEdge systems.

Dell OpenManage Deployment Toolkit

The Dell OpenManage Deployment Toolkit is a set of utilities designed to configure, record, and replicate a pre-OS server configuration for the BIOS, BMC, and RAID storage controller. DTK utilities are booted from a floppy, CD, or the network at the start of a server installation. These utilities are often used in command scripting as part of an automated server deployment solution. DTK is designed to enable quick and easy configuration of multiple servers from a bare-metal state through OS and application setup. Servers typically can be up and running within 30 minutes.

In addition, Dell has partnered with Microsoft and Altiris to support automated multi-server deployments. DTK has been integrated with Microsoft® Advanced Deployment Services (ADS) and with Altiris® Server Provisioning and Altiris Blade Deployment Solution. These approaches use DTK to script the exact configuration of PowerEdge hardware before selected OS and application software is installed on the server.

Server monitoring with Dell OpenManage

Dell OpenManage supports ongoing monitoring of server status and enables server control to help administrators handle server failures and outages virtually anytime, from virtually anywhere. Dell OpenManage enables monitoring through a combination of hardware and software components that operate directly on the managed server and on a management console (see Figure 4).

Baseboard management controller

Dell PowerEdge servers—including the PowerEdge SC1425, PowerEdge 800, PowerEdge 1800, PowerEdge 1850, PowerEdge 1855, PowerEdge 2800, PowerEdge 2850, PowerEdge 6800, and PowerEdge 6850 systems—are equipped with a BMC that implements the IPMI standard. Connected to server sensors, the BMC is designed to proactively monitor server hardware, log server fault events, and alert administrators when server faults occur. The BMC also enables basic remote operations, independent of server state, with server text console redirection, power control, and reset control.

Dell Remote Access Controller

An optional feature for PowerEdge servers, the Dell Remote Access Controller (DRAC) is designed to provide high levels of remote control, enabling administrators to remotely operate the server through industry-standard protocols such as HTTP, Secure Sockets Layer (SSL), and Secure Shell (SSH). Working in conjunction with the BMC, the DRAC enables advanced remote management capabilities including a continuous video console, virtual media, text console connectivity, a dedicated network controller, and Microsoft Active Directory® user authentication and authorization.

The BMC and DRAC provide out-of-band management for PowerEdge servers—that is, they operate regardless of the state of the server or OS. The BMC and DRAC can power, control, and reset the server, making it possible for administrators to control a PowerEdge server across the network.

Dell OpenManage Server Administrator

Dell OpenManage Server Administrator (OMSA) is designed to simplify single-server management with a secure command-line

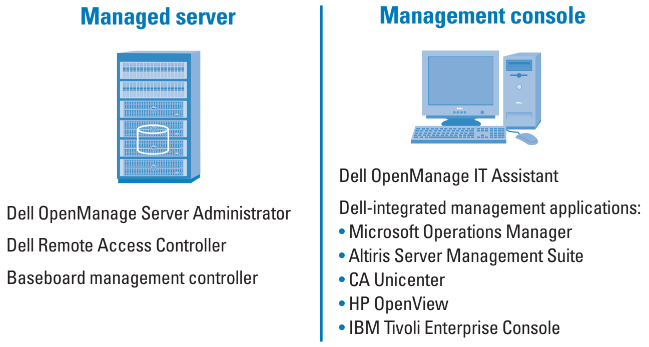
Managed server
Management console
Dell OpenManage Server Administrator
Dell Remote Access Controller
Baseboard management controller
Dell OpenManage IT Assistant
Dell-integrated management applications:
• Microsoft Operations Manager
• Altiris Server Management Suite
• CA Unicenter
• HP OpenView
• IBM Tivoli Enterprise Console

Figure 4. Overview of server status and control monitoring tools

interface (CLI) and Web-based graphical user interface (GUI). OMSA provides in-band management of PowerEdge servers, operating when the server OS is up and running normally. OMSA performs two major functions:

- **Standards-based instrumentation:** Monitors server subsystem health, reports inventory, enables server configuration, diagnoses the server, and alerts administrators about failures and warning conditions
- **Local server console:** Supports secure server operations via the CLI and GUI

OMSA subsystems include the Storage Management service for managing a server's RAID storage and Online Diagnostic services for diagnosing server hardware without taking the server offline. OMSA instrumentation provides management information, fault alerts, server configuration, and control functions for applications running on the management console. These applications can include Dell OpenManage IT Assistant, group management consoles such as Microsoft Operations Manager, and enterprise management solutions such as Computer Associates (CA) Unicenter. Dell provides versions of OMSA for leading server operating systems, including Microsoft Windows® 2000 Server and Windows Server™ 2003, Red Hat® Enterprise Linux®, and Novell® NetWare® platforms. OMSA is delivered on the Dell PowerEdge Installation and Server Management CD.

Dell works with industry leaders in systems management to enable group and enterprise management systems to inventory, monitor, and manage PowerEdge systems.

Dell OpenManage IT Assistant

Dell OpenManage IT Assistant (ITA) provides an integrated view of the comprehensive suite of Dell client and server monitoring and reporting tools. ITA discovers Dell clients and servers, gathers inventory, monitors system health status, and alerts administrators about system failures via e-mail and paging.

ITA also supports group change management for PowerEdge servers as detailed in the "Server change management with Dell OpenManage" section in this article. ITA is installed on a client or server running Microsoft Windows XP, Windows 2000 Server, or Windows Server 2003. Beginning with version 7, the ITA user interface is accessible from clients running Microsoft Internet Explorer or the Mozilla Web browser. ITA is delivered on the Dell Systems Management Consoles CD.

Integration with other management applications

Dell works with industry leaders in systems management to enable group and enterprise management systems to inventory, monitor, and manage PowerEdge systems. Integrated support is available for the following applications: Microsoft Operations Manager, Altiris Server Management Suite™ software, CA Unicenter, HP OpenView, and IBM® Tivoli Enterprise Console.

Server change management with Dell OpenManage

One of the biggest challenges for today's IT environments is change management: updating system, OS, and application software to keep it current and secure. The Dell OpenManage suite is designed to help maintain PowerEdge systems with single-server change management; multi-server change management; and integrated hardware, OS, and application change management. Components for server change management include the Dell Server Update Utility, ITA change-management services, and integrated change-management solutions from Altiris and Microsoft.

Dell Server Update Utility

To simplify the update of a single server, Dell provides a CD-based tool called the Server Update Utility (SUU). After SUU is loaded onto a server, it inventories the Dell firmware and drivers on that server; compares the installed firmware and drivers to the content of the SUU CD; and, if an update is required, automates the update process and any needed reboot of that server. SUU is designed to simplify single-server updates through features such as inventories, reports and recommendations, and checks for prerequisite conditions. SUU is delivered on the Dell PowerEdge Updates CD.

Dell OpenManage IT Assistant multi-server updates

For organizations that have deployed a significant number of PowerEdge servers, automation is essential to avoid performing the time-consuming task of single-server updates. To enable Dell firmware and driver updates for multiple servers, ITA 7 (or later) provides capabilities to inventory multiple Dell servers, identify servers that require updates, and update those servers from the ITA console. ITA uses the content of the Dell PowerEdge Updates CD to create a repository of firmware and drivers that it compares against the installed PowerEdge servers. ITA works in conjunction with OMSA to conduct the inventory and update the servers. *Note:* ITA change management focuses on managing only Dell firmware and drivers; change management of server operating systems and application software requires the use of an integrated change-management solution.

Integrated change-management solutions

To further standardization of the change-management process for server software, Dell provides a partner developer kit with features that enable Dell server updates: interfaces for inventorying server firmware and drivers, comparisons between inventory results and the latest updates, and mechanisms to install updated server firmware and drivers.

Dell has worked directly with industry leaders Microsoft and Altiris to enable their change-management solutions not only to update Dell server OS and application software, but also to help manage change for Dell server firmware and drivers. Microsoft and Altiris are delivering integrated change-management solutions that are designed to manage all software on Dell servers—enabling organizations to simplify IT operations by using a single process to maintain Dell server firmware, drivers, OS software, and applications.

Advanced systems management tools for the scalable enterprise

As enterprises continue to grow in scale and complexity, administrators must simplify operations to take control of the IT infrastructure and meet fast-changing business requirements. By understanding how to use highly instrumented clients, servers, and storage systems together with integrated, standards-based systems management solutions, administrators can proactively manage servers and clients throughout the complete system life cycle. The core Dell tool set discussed in this article, together with integrated systems management solutions from Dell partners such as Altiris and Microsoft, can help administrators automate tasks and optimize IT investments—freeing valuable resources to focus on strategic business initiatives.

Paul Rubin is a senior product planner in the Dell OpenManage marketing group. He is responsible for long-range planning of Dell systems management products including system instrumentation, base management, remote access, and standardization. Paul has a B.A. in Computer Science from The University of Texas at Austin.

Terry Myers is a senior marketing manager in the Dell outbound marketing group. He is responsible for global marketing programs and messaging in support of Dell systems management products. Terry has a bachelor's degree in business and an M.B.A.

FOR MORE INFORMATION

Dell scalable enterprise strategy:
www.dell.com/enterprise

Dell OpenManage systems management framework:
www.dell.com/openmanage

Dell standards-based solutions:
www.dell.com/standards

HOT
QLogic® Fibre Channel HBAs have officially smoked the competition.
#1
in FC HBAs

Introducing Dell OpenManage IT Assistant

Dell™ OpenManage™ IT Assistant (ITA) 7 introduces features that can significantly enhance an organization's capability to manage its systems from a centralized console. For example, ITA 7 enables administrators to obtain detailed hardware and software inventory information from a centralized database, run reports against that data, create and run scheduled tasks for selected devices, and create dynamic groups of systems based on inventoried hardware and software assets in the database.

BY SUDHIR SHETTY

Related Categories:

Change management

Dell OpenManage

Systems management

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

Dell OpenManage IT Assistant (ITA) 7 provides several key features not available in previous releases, including enhanced security and a redesigned user interface as well as additional capabilities for creating dynamic groups, managing tasks, controlling devices, launching applications, reporting device information, updating software, and troubleshooting. This article explores significant new features of ITA 7.

Enhanced security

Dell OpenManage IT Assistant 7 has significantly enhanced administrative security compared to previous versions of ITA. For example, ITA 7 employs HTTP over Secure Sockets Layer (HTTPS) communication between the user interface and the management station, authentication using OS credentials, and single sign-on capability to enable a seamless security login model.

Roles within IT Assistant

To enable seamless integration with the security model that is built into Microsoft® Windows® operating systems, ITA 7 uses the OS to authenticate users rather than defining a proprietary user access model. Each user is automatically assigned a role—User, Power User, or Administrator—based on that individual's privileges on the management station or role as configured in the Microsoft Active Directory® directory service. The *Dell OpenManage Version 4.3 Installation and Security User's Guide* describes the process for extending the Active Directory schema to enable user roles to be configured via Active Directory.[1]

The ITA user roles and privileges are as follows: User, read-only access; Administrator, full access to ITA operations; and Power User, access to several administrative functions except for the abilities to modify discovery configuration settings, modify alert management settings, delete tasks, and schedule tasks for later execution.

Single sign-on capability

To enable a seamless login, ITA 7 also supports single sign-on. This support means a user is not required to log in again after already logging in to the OS. To enable single sign-on support, administrators must navigate to Tools > Internet Options in the menu bar of Microsoft Internet Explorer. In the Internet Options pop-up window,

[1] The *Dell OpenManage Version 4.3 Installation and Security User's Guide* is available at support.dell.com/support/edocs/software/smsom/4.3/en/ug/index.htm.

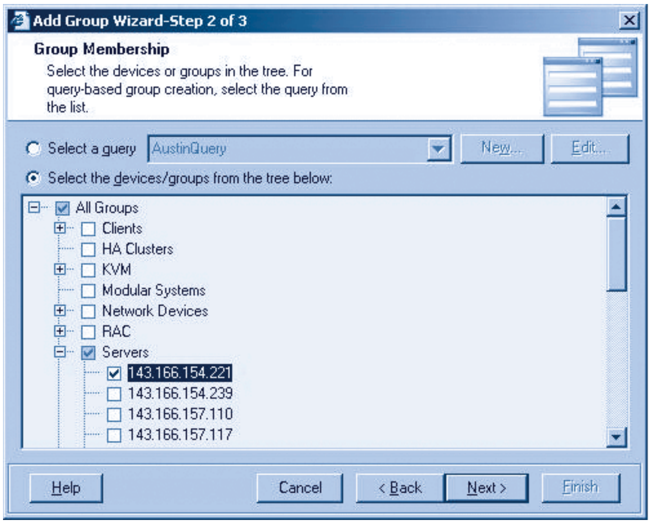

Add Group Wizard-Step 2 of 3

Group Membership

Select the devices or groups in the tree. For query-based group creation, select the query from the list.

○ Select a query: AustinQuery

● Select the devices/groups from the tree below:

Figure 1. Device selection pane within the Add Group Wizard

administrators should select the Security tab and then set the Web content zone in which the ITA management station resides; best practices recommend selecting the "Trusted sites" zone to avoid other security pop-ups. Next, administrators should click the Custom Level button and select the radio button next to "Automatic logon with current username and password." To enable remote login to a management station, administrators should use the following formats to define the URL that points to that management station:

- **Authentication using the local system's user credentials:** https://*machinename*:2607/?authType = ntlm&application = ita&locallogin = true
- **Authentication using Active Directory:** https://*machinename*:2607/?authType = ntlm&application = ita

Redesigned user interface

The standardized look-and-feel of the redesigned ITA graphical user interface (GUI) enables administrators to create objects through intuitive wizards and to easily invoke actions by right-clicking or using the Actions menu.

Because the GUI is Web based, remote administrators can manage their networks via a Web browser. ITA 7 supports Microsoft Internet Explorer on Windows operating systems and Mozilla 1.7.3 (or later) on Linux® operating systems.

Dynamic groups

Administrators can create dynamic groups based on dynamic queries or static selection of devices within the device tree. First, administrators must right-click on the All Groups node and then select "New Group." In the device selection pane of the Add Group Wizard (see Figure 1), administrators can either manually select the devices that constitute the group or define a query that determines group membership.

Reprinted from *Dell Power Solutions,* August 2005. Copyright © 2005 Dell Inc.
All rights reserved.

A query definition is based on the data that is inventoried from remote devices. To define a query, administrators should select the radio button next to "Select a query" and click the New button. Administrators can then define the device group based on certain criteria. For example, an administrator could define a group that comprises servers located in Austin, Texas, by leveraging the server naming convention the organization uses—including "aus" in the domain name. As shown in Figure 2, this example scenario uses the query "Name Contains 'aus' AND Device Type Is Server."

As servers are added or deleted, the group membership is modified dynamically to reflect those changes. The updated custom group appears in the device tree. This new custom group can be leveraged in other areas of Dell OpenManage IT Assistant that are dependent on device groups such as reporting, task management, and event filters.

Task management

Dell OpenManage IT Assistant provides a powerful task management infrastructure that allows administrators to create and schedule tasks. The following sections describe the different types and categories of tasks. For each task category, administrators can choose the target devices for the task as well as the schedule corresponding to the task. Figure 3 shows the task panel view.

Command-line tasks

The ITA 7 command-line interface (CLI) allows administrators to perform generic CLI activities, including Intelligent Platform Management Interface (IPMI) tasks, remote client instrumentation tasks, and remote Dell OpenManage Server Administrator tasks.

Generic CLI tasks. The CLI enables system administrators to run a custom, user-defined task against a specific set of devices. This provides the flexibility to schedule and run user-defined scripts or programs against a selected set of devices.

IPMI tasks. Administrators can run CLI tasks against the baseboard management controller (BMC) of a target server. For example, administrators can reset the device and clear the system event log. Administrators use the in-band Dell OpenManage Server Agent to discover the BMC. *Note:* This requires the BMC Management Utility to be installed on the management station running ITA.

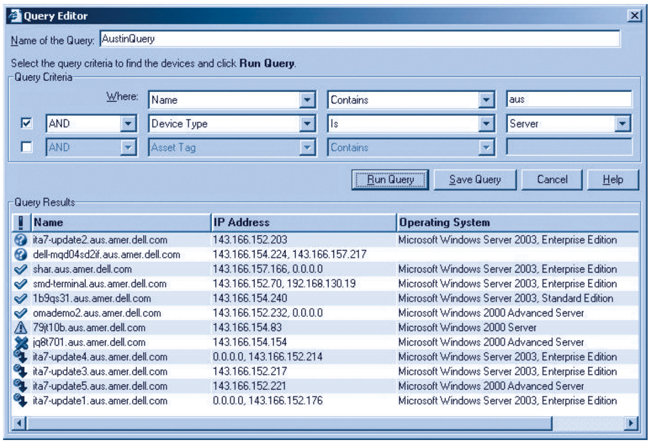

Name of the Query: AustinQuery

Select the query criteria to find the devices and click **Run Query**.

Query Criteria

	Where:	Name	Contains	aus
☑	AND	Device Type	Is	Server
□	AND	Asset Tag	Contains	

Query Results

!	Name	IP Address	Operating System
?	ita7-update2.aus.amer.dell.com	143.166.152.203	Microsoft Windows Server 2003, Enterprise Edition
?	dell-mqd04sd2if.aus.amer.dell.com	143.166.154.224, 143.166.157.217	
✓	shar.aus.amer.dell.com	143.166.157.166, 0.0.0.0	Microsoft Windows Server 2003, Enterprise Edition
✓	smd-terminal.aus.amer.dell.com	143.166.152.70, 192.168.130.19	Microsoft Windows Server 2003, Enterprise Edition
✓	1b9qs31.aus.amer.dell.com	143.166.154.240	Microsoft Windows Server 2003, Standard Edition
✓	omademo2.aus.amer.dell.com	143.166.152.232, 0.0.0.0	Microsoft Windows 2000 Advanced Server
⚠	79jt10b.aus.amer.dell.com	143.166.154.83	Microsoft Windows 2000 Server
✖	jq8t701.aus.amer.dell.com	143.166.154.154	Microsoft Windows 2000 Advanced Server
↓	ita7-update4.aus.amer.dell.com	0.0.0.0, 143.166.152.214	Microsoft Windows Server 2003, Enterprise Edition
↓	ita7-update3.aus.amer.dell.com	143.166.152.217	Microsoft Windows Server 2003, Enterprise Edition
↓	ita7-update5.aus.amer.dell.com	143.166.152.221	Microsoft Windows 2000 Advanced Server
↓	ita7-update1.aus.amer.dell.com	0.0.0.0, 143.166.152.176	Microsoft Windows Server 2003, Enterprise Edition

Figure 2. Query definition pane for creating a new device query

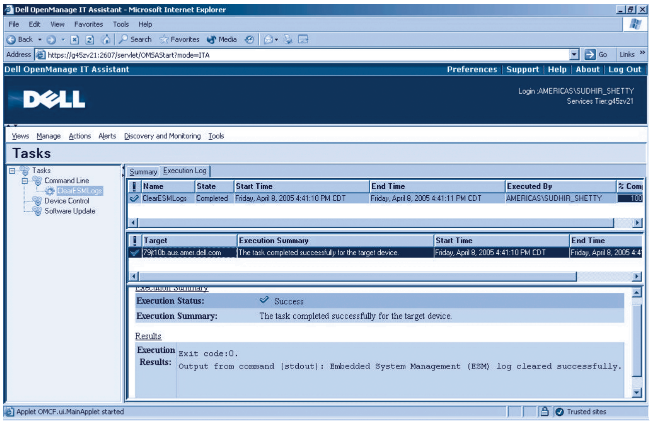

Address https://g45zv21:2607/servlet/OMSAStart?mode=ITA

Login :AMERICAS\SUDHIR_SHETTY
Services Tier:g45zv21

!	Name	State	Start Time	End Time	Executed By	% Com
✓	ClearESMLogs	Completed	Friday, April 8, 2005 4:41:10 PM CDT	Friday, April 8, 2005 4:41:11 PM CDT	AMERICAS\SUDHIR_SHETTY	100

!	Target	Execution Summary	Start Time	End Time
✓	79jt10b.aus.amer.dell.com	The task completed successfully for the target device.	Friday, April 8, 2005 4:41:10 PM CDT	Friday, April 8, 2005 4:4

Execution Summary

Execution Status: ✓ Success

Execution Summary: The task completed successfully for the target device.

Results

Execution Results:

```
Exit code:0.
Output from command (stdout): Embedded System Management (ESM) log cleared successfully.
```

Applet OMCF.ui.MainApplet started

Figure 3. Task panel view showing successful run of a Server Administrator CLI task to clear the ESM logs

Remote client instrumentation tasks. Administrators can perform remote Dell OpenManage Client Connector CLI tasks against client devices running Dell OpenManage Client Instrumentation 7.x or later. This enables administrators to perform BIOS updates for remote client devices and set values via Common Information Model (CIM) to update client BIOS configuration parameters. *Note:* This requires Client Connector 2.0 (or later) to be installed on the management station running ITA.

Remote Dell OpenManage Server Administrator tasks. Administrators can perform remote server administration CLI tasks, extending the powerful Dell OpenManage Server Administrator CLI to a one-to-many context. *Note:* Server Administrator 2.0 (or later) must be installed on the remote target device to enable this functionality. The Server Administrator CLI enables administrators to perform a wide variety of configuration and reporting tasks such as clearing Embedded Server Management (ESM) logs, reporting on storage attributes, and running diagnostics.

Device control

ITA 7 enables administrators to control devices using functions such as device wake-up and device shutdown over the network.

Device wake-up. Administrators can wake up devices on the network using the standard Wake-on-LAN (WOL) magic packet. *Note:* Network interface cards on the target devices must support this feature for this capability to work in a distributed environment. Also, the intermediate routers must be configured to enable directed broadcasting or subnet broadcasting.

Device shutdown. Administrators can shut down devices on the network, using either the Windows Management Instrumentation (WMI) protocol or Simple Network Management Protocol (SNMP). In addition, administrators can remotely reboot, power cycle (if supported), or power down the remote systems if they possess the appropriate administrative credentials.

Software update tasks

ITA 7 also enables administrators to perform remote server software updates and remote client software updates over the network.

Server software update. Administrators can perform a remote software update of a server. *Note:* This requires that Server Administrator 2.0 (or later) be installed on the remote device. Remote software updates can also be initiated after administrators run a compliance report against a specific Dell Update Package or a System Update Set.

Client software update. Administrators can perform a remote update of a client BIOS using the client software update mechanism. *Note:* To view this task type in the ITA 7 user interface, administrators must install Client Connector 2.0 (or later) on the management station running ITA.

For more information about updating software through ITA, see the "Software updates" section in this article.

Application launch

Dell OpenManage IT Assistant 7 provides a consolidated launch point for several individual device managers—including Dell OpenManage Server Administrator; Dell OpenManage Array Manager; the console for Dell remote access controllers; the console for Dell PowerConnect™ switches; and the console for digital keyboard, video, mouse (KVM) switches. ITA 7 can launch one-to-one element managers or applications for a specific device. The applications that are enabled for a particular device are typically determined by the applications that are installed on the management station, the applications that are installed at the user interface tier, and the capabilities of the target device.

Figure 4 shows an example user interface in which the administrator has the options of launching the Server Administrator, Array Manager, and Remote Desktop Connection applications against a specific device. Server Administrator and Array Manager were enabled for this device because the appropriate agent software was detected on the remote target device. Remote Desktop Connection is enabled for devices running a Windows OS.

Reporting

Dell OpenManage IT Assistant 7 is designed to gather data from a wide variety of devices and consolidate that information into a centralized database. The Report Wizard enables administrators to create custom reports that identify specific inventoried attributes for a selected set of devices. ITA 7 provides a rich reporting infrastructure that allows administrators to report on data attributes that are retrieved from

remote agents. For example, ITA 7 can report on data about storage, cost of ownership, hardware inventory, and software inventory. Reports can be executed and displayed as HTML, comma-separated value (CSV), or XML. CSV reports can be imported into a third-party tool such as Microsoft Excel to enable complex data charting and graphing capabilities.

Software updates

Dell OpenManage IT Assistant 7 enables administrators to view the Dell Update Packages and System Update Sets that reside on the Dell PowerEdge™ Updates CD, which is available as part of the Dell OpenManage Subscription Service. A System Update Set is a certified aggregate of Dell Update Packages that are specific to a hardware platform. Dell Update Packages are also available on the Dell support Web site at support.dell.com.

ITA 7 enables administrators to gather software inventory from target devices, compare it with a Dell Update Package or System Update Set, and report on devices that are not compliant with the updates. In addition, ITA 7 can schedule a remote software update for remote devices.

For example, a compliance report could be run against the System Update Set for the Dell PowerEdge 2800 server. The report might show two systems on the network that have an outdated BIOS and outdated BMC firmware. The administrator could then create a software update task that would run the System Update Set to bring the two target systems into compliance.[2]

Troubleshooting

Dell OpenManage IT Assistant includes a troubleshooting tool to help diagnose common networking problems for remote devices.

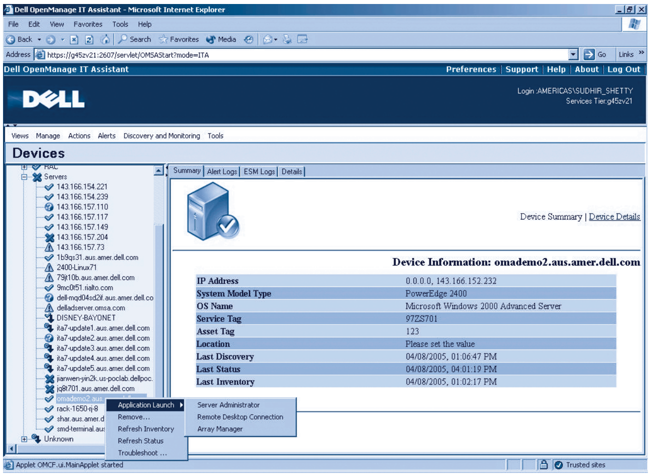
Dell OpenManage IT Assistant - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://g45zv21:2607/servlet/OMSAStart?mode=ITA
Go Links
Dell OpenManage IT Assistant
Preferences | Support | Help | About | Log Out
DELL
Login :AMERICAS\SUDHIR_SHETTY
Services Tier:g45zv21
Views Manage Actions Alerts Discovery and Monitoring Tools
Devices
RAC
Servers
143.166.154.221
143.166.154.239
143.166.157.110
143.166.157.117
143.166.157.149
143.166.157.204
143.166.157.73
1b9qs31.aus.amer.dell.com
2400-Linux71
79jt10b.aus.amer.dell.com
9mc0t51.rialto.com
dell-mqd04sd2if.aus.amer.dell.co
delladserver.omsa.com
DISNEY-BAYONET
ita7-update1.aus.amer.dell.com
ita7-update2.aus.amer.dell.com
ita7-update3.aus.amer.dell.com
ita7-update4.aus.amer.dell.com
ita7-update5.aus.amer.dell.com
jianwen-yin2k.us-poclab.dellpoc.
jq8t701.aus.amer.dell.com
omademo2.aus.
rack-1650-rj-8
shar.aus.amer.d
smd-terminal.aus
Unknown
Summary | Alert Logs | ESM Logs | Details
Device Summary | Device Details
Application Launch
Remove...
Refresh Inventory
Refresh Status
Troubleshoot ...
Server Administrator
Remote Desktop Connection
Array Manager
Applet OMCF.ui.MainApplet started
Trusted sites

Device Information: omademo2.aus.amer.dell.com

IP Address	0.0.0.0, 143.166.152.232
System Model Type	PowerEdge 2400
OS Name	Microsoft Windows 2000 Advanced Server
Service Tag	97ZS701
Asset Tag	123
Location	Please set the value
Last Discovery	04/08/2005, 01:06:47 PM
Last Status	04/08/2005, 04:01:19 PM
Last Inventory	04/08/2005, 01:02:17 PM

Figure 4. Application launch options for a single server device

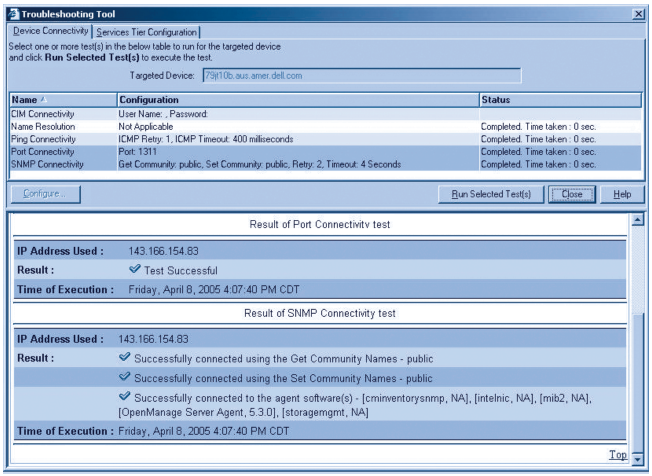

Select one or more test(s) in the below table to run for the targeted device
and click **Run Selected Test(s)** to execute the test.

Targeted Device: 79jt10b.aus.amer.dell.com

Name	Configuration	Status
CIM Connectivity	User Name: , Password:	
Name Resolution	Not Applicable	Completed. Time taken : 0 sec.
Ping Connectivity	ICMP Retry: 1, ICMP Timeout: 400 milliseconds	Completed. Time taken : 0 sec.
Port Connectivity	Port: 1311	Completed. Time taken : 0 sec.
SNMP Connectivity	Get Community: public, Set Community: public, Retry: 2, Timeout: 4 Seconds	Completed. Time taken : 0 sec.

Result of Port Connectivitv test

IP Address Used :	143.166.154.83
Result :	✓ Test Successful
Time of Execution :	Friday, April 8, 2005 4:07:40 PM CDT

Result of SNMP Connectivity test

IP Address Used :	143.166.154.83
Result :	✓ Successfully connected using the Get Community Names - public
	✓ Successfully connected using the Set Community Names - public
	✓ Successfully connected to the agent software(s) - [cminventorysnmp, NA], [intelnic, NA], [mib2, NA], [OpenManage Server Agent, 5.3.0], [storagemgmt, NA]
Time of Execution :	Friday, April 8, 2005 4:07:40 PM CDT

Figure 5. Results of the Port Connectivity and SNMP Connectivity tests run against a specific target device

Network administrators can use ITA 7 to run a set of test suites to verify appropriate device connectivity (see Figure 5). This enables network administrators to use SNMP or CIM to check device communication capability, port connectivity on a specific port, or Domain Name System (DNS) configuration. Troubleshooting tips are also provided to help administrators diagnose and correct network connectivity issues within their specific IT environments.

The ITA 7 services tier configuration tests help ensure that administrators have set up the e-mail configuration correctly. E-mail alerts are among the actions that can be configured for an incoming alert from a remote device.

Enhanced systems management across the data center

Dell OpenManage IT Assistant 7 introduces significant features that are designed to enhance enterprise systems management. ITA 7 provides a centralized console for viewing hardware and software asset inventory, performing tasks, and running reports on data gathered from remote target devices. Such capabilities can help streamline systems management across the data center.

Sudhir Shetty works in the Dell Systems Management Consoles Group responsible for defining IT Assistant software functionality. He has an M.S. in Computer Science from The University of Texas at Austin.

FOR MORE INFORMATION

Dell OpenManage:
www.dell.com/openmanage

[2] For more information about compliance reports and the software update process, see "Software Change Management Using Dell OpenManage IT Assistant" by Sudhir Shetty, Steve Heracleous, and Rohit Sharma in *Dell Power Solutions,* August 2005; www.dell.com/downloads/global/power/ps3q05-20050106-Sharma.pdf.

Guidelines for Deploying and Troubleshooting Remote Software Updates with Dell OpenManage IT Assistant

The remote update system in Dell™ OpenManage™ IT Assistant (ITA) 7 is designed to streamline the deployment and execution of software update packages. For systems running a Microsoft® Windows® OS, ITA 7 uses Windows Management Instrumentation (WMI) to perform updates. This article describes ITA 7's security and remote update capabilities, focusing on common problems encountered when updating over WMI and how to troubleshoot failed updates when such problems occur.

BY BRADLEY BRANSOM AND TOBIN RYSENGA

Related Categories:

Change management

Dell OpenManage

Microsoft Windows

Remote management

Systems management

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

Dell OpenManage IT Assistant (ITA) 7 is designed to deploy and update enterprise software efficiently across complex enterprise networks with complicated security policies. ITA will use different protocols to perform the update, depending on the OS installed on the remote target system. On systems running a Microsoft® Windows® OS, updates are performed over Windows Management Instrumentation (WMI). WMI allows ITA to connect to target systems, upload software packages to a temporary location, and execute the installation under the account matching the provided credentials. Alternatively, on systems running a Linux® OS, updates are performed over Secure Shell (SSH). Plink is the SSH client used to establish the SSH connection to the target system. SSH and Secure Copy (SCP) are used to create a secure environment for deploying and executing the remote updates. PuTTY SCP (PSCP) is used to deploy the files to the managed node.

Troubleshooting a failed WMI update

All too often, IT organizations managing complex network configurations, numerous security policies, and a large number of servers must contend with a high failure rate for remote software updates. This section explains the steps administrators can follow to help determine and fix many common problems that may occur when WMI updates fail. Figure 1 lists the error codes returned by the remote update system and some suggested next steps for solving the problem. If the steps provided in Figure 1 do not solve the problem, then general-purpose troubleshooting steps—described in the following sections—usually can help solve the problem.

Step 1: Connectivity

The first step with any failure in ITA is to verify that the system running ITA can reach the target system. Administrators can check this connectivity by pinging the target system—either through the ITA Troubleshooting Tool or the Windows command line. If the remote system cannot be reached, administrators must check various network components, such as cables, routers, and firewalls associated with that particular system.

Step 2: Wbemtest.exe

Wbemtest.exe is the WMI Tester from Microsoft that allows administrators to verify WMI commands. Running this application brings up the Windows Management Instrumentation

Error code	Error description	Suggested next step
3328	The deploy operation failed. The administrator is unable to copy the file to the remote system.	Verify that the account used to deploy the update has write access to the temp directory on the target machine.
3329	The administrator is unable to connect to the remote system.	See the "Step 1: Connectivity" section in this article.
3330	The administrator is unable to access the specified share on the remote system.	Verify that the account used to deploy the update has write access to the temp directory on the target machine.
3331	The administrator is unable to execute the WMI method to create the process on the remote system or unable to execute the SSH client.	For Windows targets, see the "Minimum WMI permissions" section in this article.
3332	The remote system did not execute the remote process.	For Windows targets, see the "Minimum WMI permissions" section in this article.
3335	The administrator is unable to create a temporary file to hold execution results on the remote system.	Verify that the account used to deploy the update has write access to the temp directory on the target machine.
3340	The administrator is unable to connect to the remote device.	If a firewall is configured, refer to the ITA readme.txt for additional information about configuring firewalls to enable remote connection.
3341	The administrator is unable to retrieve the results from the temporary file that contains the remote program execution output.	Verify that the account used to deploy the update has write access to the temp directory on the target machine.
3342	The administrator has invalid permissions to run the remote command.	For Windows targets, see the "Minimum WMI permissions" section in this article.
3353	An invalid username was specified.	Verify that the credentials provided correspond to an account that has administrative rights on the target system.
3354	And invalid password was specified.	Verify that the credentials provided correspond to an account that has administrative rights on the target system.
3360	An exception occurred during remote Dell OpenManage Server Administrator command-line interface (CLI) processing.	Choose the Trust option in ITA.
3365	Invalid network credentials were specified when connecting the share.	Verify that the credentials provided correspond to an account that has administrative rights on the target system.
3366	The administrator has invalid authentication rights.	Verify that the credentials provided correspond to an account that has administrative rights on the target system.
3367	There are insufficient privileges for WMI to execute processes on the remote system.	For Windows targets, see the "Minimum WMI permissions" section in this article. Check that the following items are configured correctly: the privilege "Replace a Process Level Token" in the Local Security Settings must grant access to the Local Service and the Network Service; and the Local Security Settings, WMI Security Settings, and any other security settings that deviate from the default settings must be set to default.
3369	The path of the remote Server Administrator CLI executable that is attempting to execute on the remote system is invalid.	Make sure that omexec.exe is in the default path of the remote system. A reboot after installation of Server Administrator should help ensure that the path is correct.

Figure 1. Error codes and troubleshooting tips for failed WMI updates

Tester dialog box. From this dialog box, administrators should click the Connect button to open the Connect dialog box, in which they can enter the details of the target system being tested (see Figure 2).

Administrators should enter "\\IP of remote system\root\cimv2" in the Namespace field and the proper username, password, and authority in the Credentials section—these should be the same credentials that are set up in ITA. The other fields on this screen can be left in the default state. Administrators should then click the Connect button to establish a connection to the target system.

If an error message occurs at this point, the provided credentials most likely do not have access to the target system. If the connection was successful, administrators will see the Wbemtest main screen. They should then click the Execute Method button, which will open a small dialog box prompting for an object path. Administrators should enter "win32_process" (making sure there are no spaces before or after this phrase) and click the OK button. This action brings up the Execute Method dialog box, in which administrators should use the Method drop-down menu to select the Create method and then click the Edit In Parameters button. The object editor, shown Figure 3, will then appear. Administrators should scroll down through the Properties list box, select the CommandLine property, and click the Edit Property button.

In the Property Editor dialog box, administrators need to change the Value section by setting the radio button to "Not NULL" and entering "Ping /t 127.0.0.1" in the text area (see Figure 4). Next,

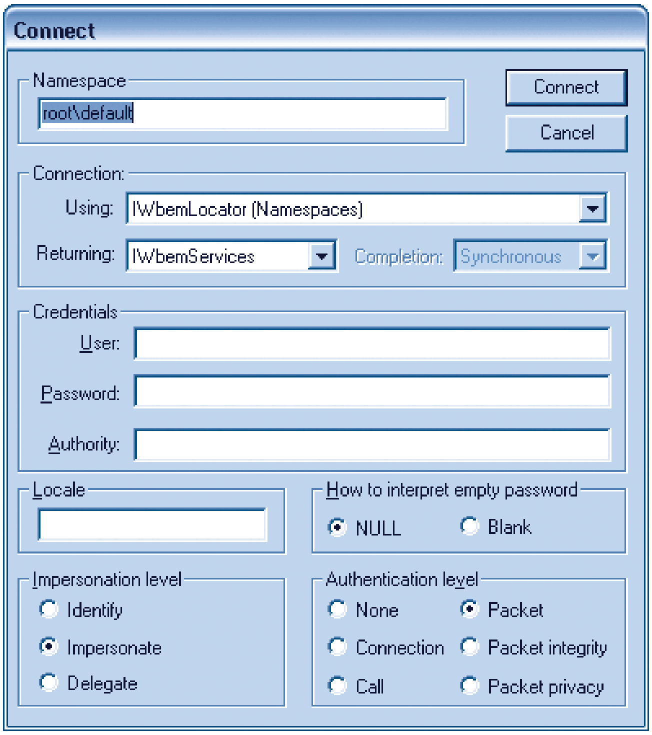

Namespace: root\default

Using: IWbemLocator (Namespaces)

Returning: IWbemServices   Completion: Synchronous

How to interpret empty password: ● NULL ○ Blank

Impersonation level: ○ Identify ● Impersonate ○ Delegate

Authentication level: ○ None ● Packet ○ Connection ○ Packet integrity ○ Call ○ Packet privacy

Figure 2. Wbemtest Connect dialog box

Dave is on vacation.

He's been notified of a problem with his company's servers.
He's able to manage those servers from anywhere.
He's not cutting his vacation short.

He's right here.

The Dell remote console switch provides the traditional functionality of KVM switching with the added benefit of **KVM Over IP**. 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 or across the globe your data center will always be within your reach.

Take **control** of your servers with the **Dell Remote Console Switch.**

Digital Availability. **Easy as DELL**

www.dell.com

Dell and the Dell logo are trademarks of Dell Inc.

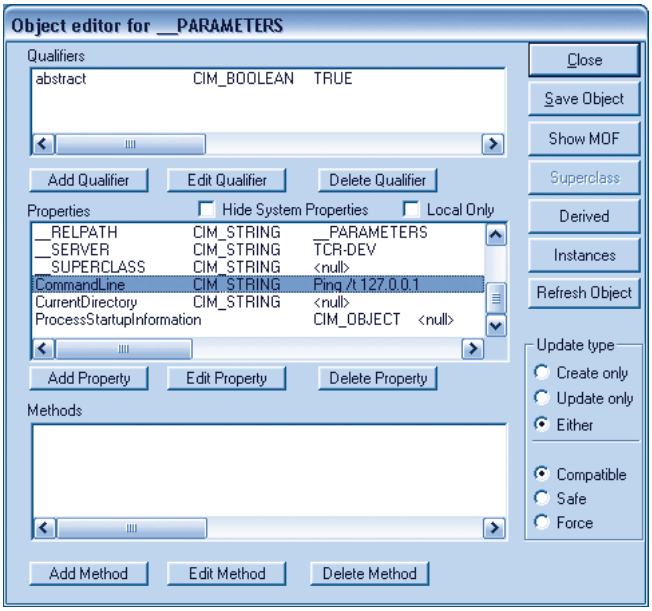

Object editor for __PARAMETERS

Qualifiers

```
abstract                CIM_BOOLEAN   TRUE
```

Properties □ Hide System Properties □ Local Only

```
__RELPATH                  CIM_STRING   __PARAMETERS
__SERVER                   CIM_STRING   TCR-DEV
__SUPERCLASS               CIM_STRING   <null>
CommandLine                CIM_STRING   Ping /t 127.0.0.1
CurrentDirectory           CIM_STRING   <null>
ProcessStartupInformation               CIM_OBJECT   <null>
```

Update type: ○ Create only ○ Update only ● Either

● Compatible ○ Safe ○ Force

Figure 3. Wbemtest object editor

administrators should click the Save Property button to return to the object editor. They should then click the Save Object button to save the changes, and click the Execute Method button to return to the Execute Method dialog box, where they can click the Execute! button. If this execution works, administrators will see a successful message pop-up window. They can click the Edit Out Parameters button in the Execute Method dialog box to bring up the output property editor screen (see Figure 5). Administrators should scroll down in the Properties list, select the ReturnValue property, and click the Edit Property button to read the return value from the method. The value of the ReturnValue property provides information about the WMI method. Figure 6 explains the return values. Once administrators have looked up the return value, they can cancel all the way out and click the Exit button to quit wbemtest.

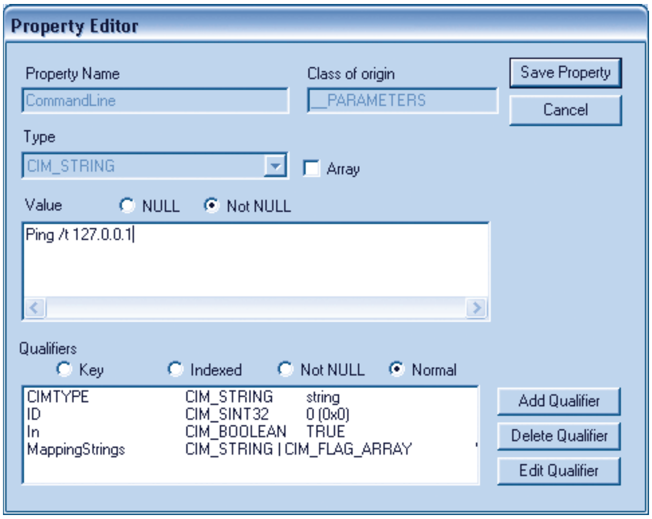

Property Editor

Property Name: CommandLine

Class of origin: __PARAMETERS

Type: CIM_STRING □ Array

Value ○ NULL ● Not NULL

Ping /t 127.0.0.1

Qualifiers ○ Key ○ Indexed ○ Not NULL ● Normal

```
CIMTYPE          CIM_STRING     string
ID               CIM_SINT32     0 (0x0)
In               CIM_BOOLEAN    TRUE
MappingStrings   CIM_STRING | CIM_FLAG_ARRAY
```

Figure 4. Wbemtest Property Editor dialog box

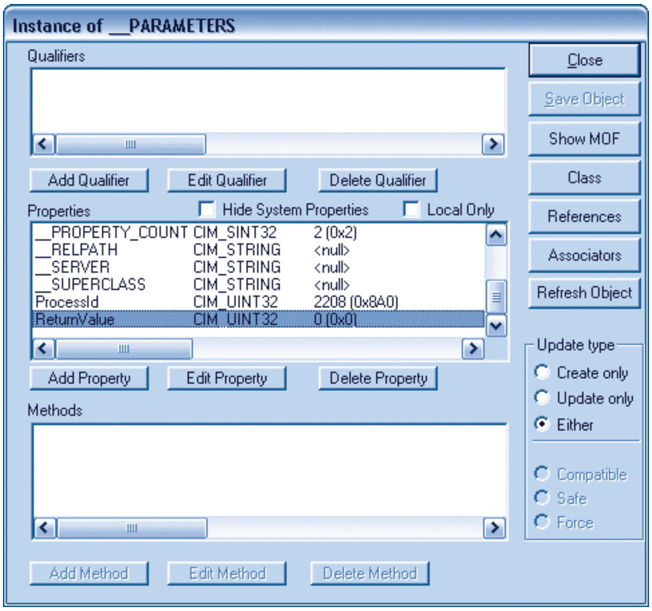

Instance of __PARAMETERS

Properties □ Hide System Properties □ Local Only

```
__PROPERTY_COUNT  CIM_SINT32    2 (0x2)
__RELPATH         CIM_STRING    <null>
__SERVER          CIM_STRING    <null>
__SUPERCLASS      CIM_STRING    <null>
ProcessId         CIM_UINT32    2208 (0x8A0)
ReturnValue       CIM_UINT32    0 (0x0)
```

Update type: ○ Create only ○ Update only ● Either

○ Compatible ○ Safe ○ Force

Figure 5. Wbemtest output property editor

Identifying symptoms and problem areas for WMI

This section discusses configuration and permissions that are required for successful remote updates.

Minimum WMI permissions

Before the remote update feature can run, Dell OpenManage IT Assistant 7 requires the capability to create a process on the target system. This process creation can be implemented through the Create method of the WMI win32_process class.

Process creation is performed using the CreateProcessAsUser application programming interface (API) while the provider (the WMI engine on the remote system) simulates the caller (the credentials the user provided). To be able to call win32_process.CreateProcess, the wmiprvse.exe account requires the following privileges:

Return value	Description
0*	The WMI command was successfully executed.
2	The administrator has invalid authentication rights. Please check that the person whose credentials are being used has at least administrative rights on the remote system.
3	There are insufficient privileges for WMI to execute processes on the remote system. Check that the following items are configured correctly: the privilege "Replace a Process Level Token" in the Local Security Settings must grant access to the Local Service and the Network Service; and the Local Security Settings, WMI Security Settings, and any other security settings that deviate from the default settings must be set to default.
8	An unknown error was returned from the Create method of the WMI win32_process.
9	The path of the remote Dell OpenManage Server Administrator CLI executable that is attempting to execute on the remote system is invalid. Make sure that omexec.exe is in the default path of the remote system.
21	An invalid parameter was passed into the Create method of the WMI win32_process.

*If the return value is 0, administrators should go to the remote system, verify that the ping is running in the task manager, and then stop the ping.

Figure 6. Wbemtest return values that help verify WMI commands

- **SeImpersonatePrivilege:** Simulate a client (Microsoft Windows Server 2003 or Windows XP with Service Pack 2)
- **SeAssignPrimaryTokenPrivilege:** Replace a process-level token
- **SeIncreaseQuotaPrivilege:** Adjust memory quotas for a process

To authorize WMI users and set permissions, administrators should access WMI Control. To do so, they must right-click the My Computer icon on the target system, select "Manage," and then select "Services and Applications" in the Computer Management tree. Next, they should right-click on "WMI Control" and select "Properties." From this window, administrators should click the Security tab, select the namespace for which they want to assign a user or group access, and then click the Security button.

For example, an administrator provides access to Root > CIMV2. In the Security for ROOT\CIMV2 dialog box, the administrator clicks the Add button; then the Select Users, Computers, or Groups dialog box appears, where the name of the object (user or group) to be added is entered. The administrator clicks the Check Names button to verify the entry and then clicks the OK button. Administrators may have to change the location or use the Advanced button to query for objects. The dialog box Help feature—which is accessed by right-clicking on any element and selecting "What is this?"—can provide more details.

At the top of this list, the caller account should have FILE_EXECUTE and FILE_READ_DATA access to the image it is about to launch. The account should also be able to bypass traverse checking (SeChangeNotifyPrivilege is enabled).

In the Security dialog box, under Permissions for Administrators, administrators should select the permissions to allow or deny the new user or group. Permission levels are as follows:

- **Execute Methods:** Allows methods exported from the WMI classes or instances to be run.
- **Full Write:** Allows full read, write, and delete access to all WMI objects, classes, and instances.
- **Partial Write:** Allows write access to static WMI objects.
- **Provider Write:** Allows write access to objects that are supplied by providers.
- **Enable Account:** Allows read access to WMI objects.
- **Remote Enable:** Allows remote access to the namespace.
- **Read Security:** Allows read-only access to WMI security information.
- **Edit Security:** Allows read and write access to WMI security information.

Note: To set WMI permissions, an administrator must be a member of the Administrators group on the local computer.

Administrators also can set permissions on a remote computer. To access a remote computer, they should right-click on "Computer Management" to connect to the other computer.

DNS name problems

Dell OpenManage IT Assistant offers administrators the ability to discover systems by IP address range or by host name. Once ITA has determined the host name for a discovered device, however, it will connect to that machine using the host name for remote updates. With a network that contains a mixture of domains and workgroups, more than one system may match the host name that ITA is using to access the target system. In such a case, several types of failures may occur because ITA is communicating with the wrong system. The only way to correct this situation is to identify the problem system(s) and change the host names to be unique across the network.

This duplicate name error also can occur if systems that are part of a domain are *ghosted*—that is, the hard drive image for a system was created from ghost software used to duplicate images of hard drives. All members of a domain should be uniquely named; otherwise, a system with a duplicate name will fail when attempting various network tasks. In this case, the system must be removed from the domain and added again with a unique name.

Streamlining remote software updates

Through WMI, Dell OpenManage IT Assistant 7 can enable efficient deployment and execution of remote software updates across Microsoft Windows–based systems. In particular, enterprises with complex networks, numerous security policies, and a large number of servers face a high probability of encountering problems when deploying remote software updates in a WMI environment. The troubleshooting guidelines discussed in this article can help administrators streamline systems management and the remote software update process using ITA 7.

Bradley Bransom is a senior developer on the Dell Server Administrator team. Prior to joining Dell, Bradley was employed by AMD, 3COM, and Texas Instruments. He has a B.S. in Computer Science with a minor in Mathematics from Texas A&M University.

Tobin Rysenga is a senior developer on the Dell OpenManage IT Assistant development team. Prior to joining Dell, Tobin worked for Maxis, Microprose, and Glass Eye and has worked as an independent consultant.

FOR MORE INFORMATION

Dell OpenManage IT Assistant User's Guide:
support.dell.com/support/edocs/software/smitasst/7.0

Shetty, Sudhir, Steve Heracleous, and Rohit Sharma. "Software Change Management Using Dell OpenManage IT Assistant." *Dell Power Solutions,* August 2005. www.dell.com/downloads/global/power/ps3q05-20050106-Sharma.pdf

Software Change Management

Using Dell OpenManage IT Assistant

Dell™ OpenManage™ IT Assistant 7 enables organizations to significantly enhance management of the software life-cycle process, allowing administrators to inventory software asset information (BIOS, firmware, and drivers) in a centralized database; compare that asset information with a certified Dell Update Package or System Update Set; and schedule software updates for remote systems. This article outlines the IT Assistant 7 software change-management architecture and feature set.

BY SUDHIR SHETTY, STEVE HERACLEOUS, AND ROHIT SHARMA

Related Categories:

Asset management

Change management

Dell OpenManage

Dell PowerEdge servers

Inventory management

Systems management

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

Dell OpenManage IT Assistant (ITA) 7 provides several key features that enable IT administrators to streamline management of the software life-cycle process. This feature set includes capabilities that enhance software inventory, software release management, compliance reporting, and software update processes. This article explains these features and provides step-by-step instructions for implementing software updates with ITA 7.

Figure 1 represents the Dell OpenManage IT Assistant change-management architecture. ITA requires the appropriate agent to be installed on the remote Dell PowerEdge™ server to enable the change-management infrastructure. The minimum supported agent is Dell OpenManage Server Administrator 2.0 (available on the Dell OpenManage 4.3 CD) running on a Microsoft® Windows® or Red Hat® Enterprise Linux® OS.

Software inventory

ITA collects software asset data from remote agents via Simple Network Management Protocol (SNMP) and Windows Management Instrumentation (WMI). This information is gathered from Dell OpenManage Server Administrator agents (version 4.3 or later) and centrally stored in the ITA database. This capability enables administrators to run software inventory reports via either ITA or third-party enterprise reporting tools.

ITA collects the software inventory information during an inventory cycle scheduled by the network monitoring service. The inventory cycle can be scheduled within ITA. Alternatively, the inventory cycle can run on demand for a specific device or for a specific range.

Software release management

ITA enables administrators to view the Dell Update Packages and System Update Sets that reside on the Dell PowerEdge Updates CD, which is available as part of the Dell OpenManage Subscription Service. A System Update Set is a certified aggregate of Dell Update Packages that are specific to a hardware platform. Dell Update Packages are also available on the Dell support Web site at support.dell.com.

In ITA, administrators can open a software repository via the Software Updates View. To open the repository located on the Dell PowerEdge Updates CD, administrators should right-click on the Software Update Repositories node in the tree and select “Open Repository (Update CD).”

EASIER TO GROW IS ABOUT TO GET EASIER TO AFFORD.

Accelerate ahead with the tested, validated, low-cost Oracle Database 10*g* Standard Edition and Oracle Real Application Clusters solution.

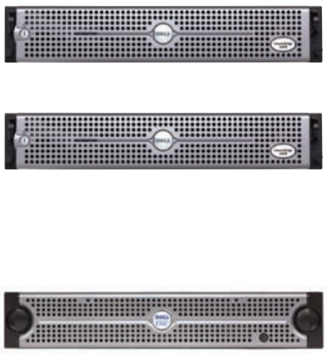

With Dell's supported configuration for Oracle Database 10*g* Standard Edition and Oracle Real Application Clusters, running on Microsoft® Windows® Operating System, you can reduce the price of success. Tested and validated for performance, the configuration lets you create a highly available, redundant, 2-node server solution. The database you build today scales to meet your needs tomorrow. Dell has teamed with Oracle to help you accelerate your move to 10*g*. It's a complete solution, at unheard-of value, from world-renowned providers. All with a single point of contact: Dell. Get enterprise-level server capability that's easy to buy, install and expand.

Which leaves just one question. How fast can you dial?

ORACLE® Microsoft®

GET MORE OUT OF NOW. DELL

Call 877-274-7375 or visit www.dell.com/oracle for more information.

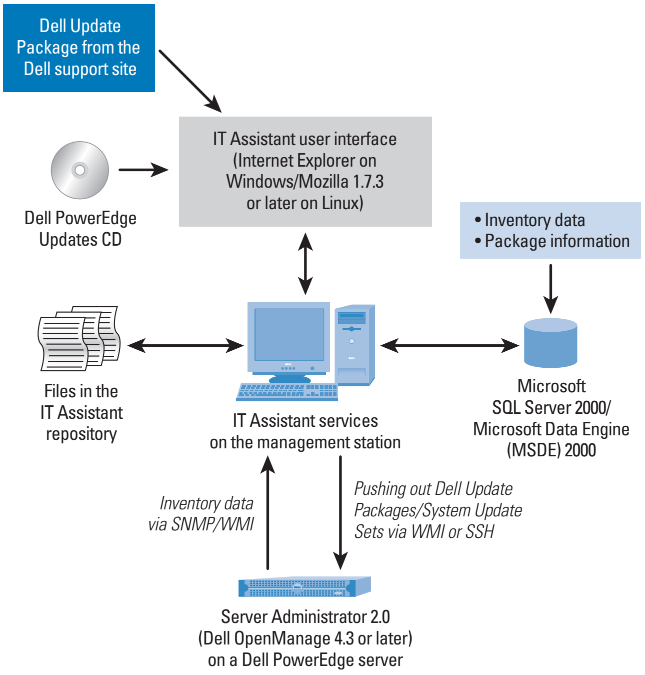
Dell Update Package from the Dell support site
Dell PowerEdge Updates CD
IT Assistant user interface (Internet Explorer on Windows/Mozilla 1.7.3 or later on Linux)
• Inventory data
• Package information
Files in the IT Assistant repository
IT Assistant services on the management station
Microsoft SQL Server 2000/ Microsoft Data Engine (MSDE) 2000
Inventory data via SNMP/WMI
Pushing out Dell Update Packages/System Update Sets via WMI or SSH
Server Administrator 2.0 (Dell OpenManage 4.3 or later) on a Dell PowerEdge server

Figure 1. Overview of the IT Assistant change-management architecture

Administrators can then locate and select the catalog.xml file in the repository directory.

After the repository is opened, specific packages or System Update Sets can be imported into the ITA repository by right-clicking on those packages and selecting “Import.” Packages can also be added directly to the ITA repository by right-clicking on “IT Assistant Repository” in the tree and selecting “Add.” This operation can be used for Dell Update Packages that have been downloaded individually from the Dell support Web site.

Compliance reports

ITA compliance reporting enables administrators to compare the software inventory of remote systems with a Dell Update Package or a System Update Set. After the import process is complete, administrators can perform a compliance check and determine which components on the remote target device are not compliant with the Dell Update Package or System Update Set.

Figure 2 displays an example compliance report that has been executed with a Dell PowerEdge 2800 System Update Set. The compliance has been executed against all devices in the enterprise. Administrators can pick specific devices in the device tree for the compliance report, or they can run the report against devices that meet certain query criteria. A system also must meet the following prerequisites to be listed in the compliance report:

- Running the appropriate Server Administrator agent (Dell OpenManage 4.3 or later).
- Containing inventory data in the ITA database.
- Complying with the update criteria for the specified Dell Update Package or System Update Set. ITA uses the prerequisite information in the Dell Update Package or System Update Set to determine applicability to a particular remote target system. Only the PowerEdge 2800 servers will show up in the example compliance report shown in Figure 2.

In the example compliance report, the system named bransom2800 is compliant with the System Update Set. Green checkmark icons in the compliance report indicate that components have the correct versions of the BIOS, baseboard management controller (BMC) firmware, Intel® connection driver, and backplane firmware. Note that the compliance report will not list components that are not present on the system. Hence, even if the System Update Set has packages corresponding to Broadcom network interface card (NIC) drivers, these packages will not be listed in the compliance report if a Broadcom NIC is not installed in the target system. Figure 3 explains the icons that appear in a compliance report.

Software update process

ITA can schedule remote software updates by deploying packages via WMI to a Microsoft Windows–based system or via Secure Shell (SSH) to a Linux-based system. The Server Administrator agent at the remote system applies the packages and performs the necessary reboots, if required.

After a compliance report is generated as described in the preceding section, administrators can proceed with updating devices that are not compliant with the Dell Update Package or System Update Set. If a target device is selected in the compliance report, the Update button in that panel becomes enabled (see Figure 4). After pressing the Update button, the administrator is guided by the Task Creation Wizard to create a software update task.

The wizard-guided process is similar to the process for creating a software update task via the task panel in ITA (invoked by going

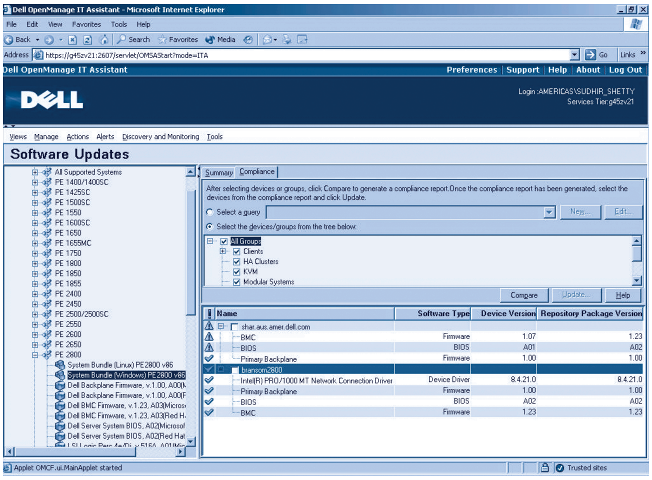
Dell OpenManage IT Assistant - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Address https://g45zv21:2607/servlet/OMSAStart?mode=ITA
Go Links
Dell OpenManage IT Assistant
Preferences | Support | Help | About | Log Out
DELL
Login :AMERICAS\SUDHIR_SHETTY
Services Tier:g45zv21
Views Manage Actions Alerts Discovery and Monitoring Tools
Software Updates
All Supported Systems
PE 1400/1400SC
PE 1425SC
PE 1500SC
PE 1550
PE 1600SC
PE 1650
PE 1655MC
PE 1750
PE 1800
PE 1850
PE 1855
PE 2400
PE 2450
PE 2500/2500SC
PE 2550
PE 2600
PE 2650
PE 2800
System Bundle (Linux) PE2800 v86
System Bundle (Windows) PE2800 v86
Dell Backplane Firmware, v.1.00, A00(M
Dell Backplane Firmware, v.1.00, A00(F
Dell BMC Firmware, v.1.23, A03(Micros
Dell BMC Firmware, v.1.23, A03(Red H.
Dell Server System BIOS, A02(Microsof
Dell Server System BIOS, A02(Red Hat
Summary Compliance
New... Edit...
Compare Update... Help
Applet OMCF.ui.MainApplet started
Trusted sites

After selecting devices or groups, click Compare to generate a compliance report.Once the compliance report has been generated, select the devices from the compliance report and click Update.

○ Select a query
● Select the devices/groups from the tree below:

☑ All Groups
☑ Clients
☑ HA Clusters
☑ KVM
☑ Modular Systems

!	Name	Software Type	Device Version	Repository Package Version
⚠	shar.aus.amer.dell.com			
⚠	BMC	Firmware	1.07	1.23
⚠	BIOS	BIOS	A01	A02
✓	Primary Backplane	Firmware	1.00	1.00
✓	bransom2800			
✓	Intel(R) PRO/1000 MT Network Connection Driver	Device Driver	8.4.21.0	8.4.21.0
✓	Primary Backplane	Firmware	1.00	1.00
✓	BIOS	BIOS	A02	A02
✓	BMC	Firmware	1.23	1.23

Figure 2. Example compliance report executed against all systems in an enterprise

Compliance report icons	
✓	The version of the software component on the remote target system is the same as the version of the component within the selected Dell Update Package or System Update Set
⚠	The version of the software component on the remote target system is older than the version within the selected Dell Update Package or System Update Set and thus the component needs to be updated.
✓↓	The version of the software component on the remote target system is newer than the version within the selected Dell Update Package or System Update Set. (*Note:* These components will not be altered unless the administrator forces a downgrade or selects the re-apply option when creating the software update task.)
✖↓	The version of the software component on the remote target system does not meet the minimum prerequisites to be remotely updated via ITA. This typically occurs because the software component is too old or requires a floppy-based mechanism to perform the software update. In this scenario, the option to perform a remote software update is disabled for that system.
Task execution status icons	
✓	The update was successful.
⚠	The target system must be rebooted for the update to take effect.
✖	The update task failed.
●	The update task is in progress.
Individual package summary icons	
✓	The package applied successfully or was not applicable for this device.
⚠	The package requires a reboot to take effect.
✖	The package failed to apply. To resolve this issue, administrators should try another method (such as a floppy-based image) to update this component.

Figure 3. Icon legend for Dell OpenManage IT Assistant 7

to Manage > Tasks on the menu bar). However, when creating a task via the task panel, administrators must select the specific Dell Update Package or System Update Set and also explicitly select the devices to be updated. To benefit from a compliance check before applying an update package, best practices recommend that administrators create update tasks using the Software Updates view.

Step 1: Create task

The first pane for creating a software update task enables administrators to enter a name and a description for the task. Context-sensitive help is available by clicking the Help button in each pane. If the administrator clicks the Cancel button, the task creation operation is canceled.

Step 2: Select options

The second pane for creating a software update task enables administrators to pick configuration options such as forcing a downgrade or re-apply, allowing a reboot, and generating a trusted key.

Force downgrade or re-apply. If applicable, this option will downgrade or re-apply the same version of the software package on the target device. If the component software on the target system is the same version or later of the package being installed, the default behavior is that the component software will not be updated.

Allow reboot. This option enables the software update task to reboot the system if necessary. Otherwise, a message will be added to the task execution log, indicating that a reboot may be required for the update to be successfully applied.

Generate trusted key. For software updates on Linux target systems, ITA uses SSH to communicate with the remote Linux device. During the SSH connection, if the host key or device identifier is not recognized, then the SSH client software issues a warning. This warning may also be issued if administrators connect to a device for the first time. To trust the host key automatically and ignore the warning, administrators should enable this option. If working from within a company network and protected from the Internet by a firewall, administrator may choose to trust the host key without checking it.[1]

Step 3: Select schedule

The third pane for creating a software update task enables administrators to schedule tasks at a specific time or on an hourly, daily, weekly, or monthly basis. For a software update task, administrators typically select the Run Now radio button to immediately execute the task, or they select "Run once" to pick the specific start date and time for running the task.

Step 4: Enter credentials

The fourth pane for creating a software update task enables administrators to specify the credentials required for task execution. Administrators must enter credentials that have administrative privileges on the target system. Windows OS user IDs should be specified as *domain\username* or *localhost\username*. On Linux target systems, the user ID is an administrative username for that system.

Step 5: Review task creation summary

Figure 5 displays the summary pane for the software update task, presenting pertinent information supplied during the task creation

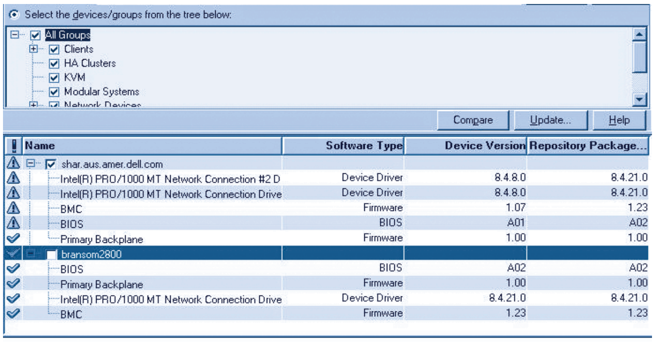

◉ Select the devices/groups from the tree below:

⊟ ☑ All Groups
- ⊞ ☑ Clients
- ☑ HA Clusters
- ☑ KVM
- ☑ Modular Systems
- ⊞ ☑ Network Devices

!	Name	Software Type	Device Version	Repository Package...
⚠	⊟ ☑ shar.aus.amer.dell.com			
⚠	Intel(R) PRO/1000 MT Network Connection #2 D	Device Driver	8.4.8.0	8.4.21.0
⚠	Intel(R) PRO/1000 MT Network Connection Drive	Device Driver	8.4.8.0	8.4.21.0
⚠	BMC	Firmware	1.07	1.23
⚠	BIOS	BIOS	A01	A02
✓	Primary Backplane	Firmware	1.00	1.00
✓	⊟ ☐ bransom2800			
✓	BIOS	BIOS	A02	A02
✓	Primary Backplane	Firmware	1.00	1.00
✓	Intel(R) PRO/1000 MT Network Connection Drive	Device Driver	8.4.21.0	8.4.21.0
✓	BMC	Firmware	1.23	1.23

Figure 4. Software update task creation for a device listed in the compliance report

[1] For additional information about configuring keys for SSH communication to a remote target Linux system, administrators should view the ITA readme.txt file in the IT Assistant folder of the installation directory.

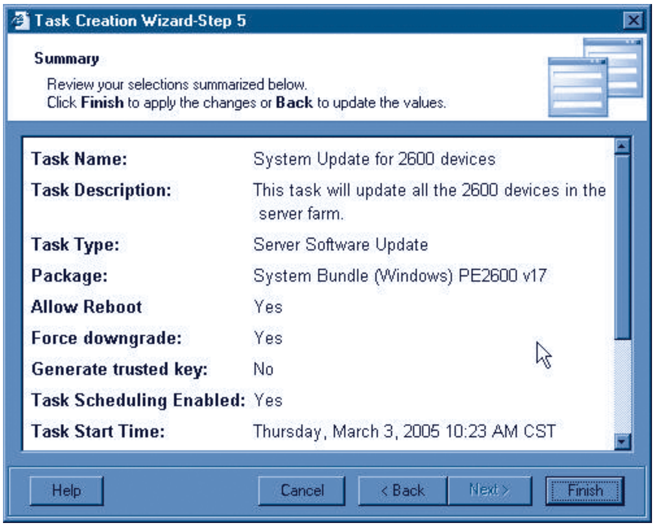

Task Creation Wizard-Step 5

Summary

Review your selections summarized below.
Click **Finish** to apply the changes or **Back** to update the values.

Task Name:	System Update for 2600 devices
Task Description:	This task will update all the 2600 devices in the server farm.
Task Type:	Server Software Update
Package:	System Bundle (Windows) PE2600 v17
Allow Reboot	Yes
Force downgrade:	Yes
Generate trusted key:	No
Task Scheduling Enabled:	Yes
Task Start Time:	Thursday, March 3, 2005 10:23 AM CST

Figure 5. Summary of the task creation process

process. At this stage, administrators can cancel the task, use the Back button to update data that was entered in previous panes, or click the Finish button to complete the task creation process.

Task execution

Once the task creation process is completed, ITA will attempt to persist the task information in the database. Errors are reported to the administrator. If the task was successfully created, the administrator's context is switched to the task panel and the summary information is displayed. The execution log can be viewed by clicking the Execution Log tab.

The task execution log provides detailed information about each run of a task. The bottom panel displays the overall task execution status for a particular target system as well as the individual execution log for each Dell Update Package in the System Update Set. The overall status is represented by the task execution icons explained in Figure 3.

Typically, a task will fail for one of the following reasons:

- **Improper credentials:** The credentials must have administrative privileges on the remote target system.
- **Networking problems:** ITA must be able to communicate with the remote target system. The troubleshooting tool can help diagnose network connectivity issues; the Ping Connectivity and CIM (Common Information Model) Connectivity tests can verify a successful connection to a target device.
- **Dell Update Package failures:** The detailed task execution log contains additional information about why an individual package failed to apply. The individual package summary icons are detailed in Figure 3.

Figure 6 describes additional troubleshooting issues.

Problem	Resolution
Administrator is unable to view the software inventory in the device details panel for a system in the device tree view.	• Verify that Server Administrator 2.0 (or later) is installed on the managed system. • Right-click on the device name in the device tree and select "Troubleshoot." Run the SNMP Connectivity or CIM Connectivity test, and verify that the system is connected. Via SNMP, verify that ITA can successfully connect to the cminventorysnmp agent on the target device. • Right-click on the device name and select "Refresh Inventory" or "Perform Inventory Now" from the ranges view. Note that the agent can take up to five minutes to successfully inventory all the software on the device.
Compliance report for a Dell Update Package or System Update Set indicates that no devices can be updated by the selected package.	• Follow the preceding steps to help ensure that ITA can successfully retrieve the inventory for a specific system. • Note that the package contains prerequisite information that helps determine which PowerEdge systems or target operating systems can be updated. (For example, if the administrator is implementing a PowerEdge 2800 Linux System Update Set, devices will appear in the compliance report only if a PowerEdge 2800 server is running Red Hat Enterprise Linux within that environment.)
Software update task failed to execute and has an ✖ icon in the task execution log.	• View the task execution results to obtain a description of the failure. • View the detailed package log for more information about why the task execution failed.

Figure 6. Resolution of common troubleshooting issues

Effective change management using IT Assistant 7

Dell OpenManage IT Assistant 7 includes powerful features that are designed to enable administrators to streamline management of the enterprise software life-cycle process. ITA provides a centralized console for viewing software asset inventory, checking version compliance, and scheduling remote software updates. These capabilities allow administrators to keep their enterprises updated with the latest certified software and to manage distributed environments from a central console.

Sudhir Shetty works in the Dell Systems Management Consoles Group responsible for defining IT Assistant software functionality. He has an M.S. in Computer Science from The University of Texas at Austin.

Steve Heracleous works in the Dell Systems Management Consoles Group responsible for developing the IT Assistant user interface. He has a bachelor's degree and a master's degree in Electrical Engineering from The University of Texas at Austin.

Rohit Sharma is a test engineer in the Dell OpenManage software development and test organization. He has a bachelor's degree in Computer Engineering from the University of Mumbai in India and a master's degree in Computer Science from North Carolina State University.

FOR MORE INFORMATION

Dell OpenManage:
www.dell.com/openmanage

Best Practices for Migrating

to Dell OpenManage IT Assistant

Dell™ OpenManage™ IT Assistant (ITA) is a centralized systems management console that allows administrators to monitor and manage network devices in various environments. This article describes the features introduced in ITA 7 and details best practices for migrating from previous versions of ITA. Administrators who plan to start using ITA as their systems management console can also benefit from this discussion.

BY KRISHNA MOHAN AND SCOTT THOMAS

Related Categories:

Best practices

Change management

Dell OpenManage

Systems management

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

In growing enterprises, systems management consoles must have the capability to scale quickly and flexibly to manage, monitor, and update heterogeneous devices on the network. Dell OpenManage IT Assistant (ITA) 7 is designed to centrally manage various devices, including Dell PowerEdge™ servers, Dell OptiPlex™ desktops, Dell Precision™ workstations, and Dell PowerConnect™ switches. Beginning with version 7, ITA includes enhanced features and capabilities designed to manage, monitor, and update several types of network devices in various enterprise environments. This article explains these functions and provides guidance for administrators planning to migrate to ITA 7 software.[1] To perform this migration, administrators can obtain ITA 7 from the Dell OpenManage Applications CD or download ITA 7 from the Dell support Web site at support.dell.com.

General overview of IT Assistant components

The Dell OpenManage IT Assistant systems management console is a multi-tiered application consisting of a services tier and a client user interface (UI) tier (see Figure 1), which comprise the following components:

- **ITA services tier:** IT Assistant Network Monitoring Service, IT Assistant Connection Service, and ITA database
- **ITA client UI:** Console view for managing and monitoring all devices; multiple ITA client UIs can connect to a single ITA services tier

In ITA versions prior to 6.5, the client UI is a Web application that can be accessed remotely via the Microsoft® Internet Explorer Web browser from any Microsoft® Windows® OS as long as the system on which the console was installed is running the client UI as a published Web application under Microsoft Internet Information Services (IIS). The client UI is implemented using the Microsoft Java Virtual Machine (JVM). In ITA versions 6.5 to 6.5.3, the client UI was developed using the Microsoft .NET framework because Microsoft JVM was no longer supported. The client UI console is an application and does not support Web access. Beginning with ITA 7, the console view can be remotely accessed via different Web browsers and can also be remotely accessed on systems running the Red Hat® Enterprise Linux® OS.

[1] Best practices discussed in this article assume a basic level of familiarity with the use of ITA in a networked environment. Many concepts discussed in this article build on the ITA online help, which can be accessed directly from the installed product as well as from the *Dell OpenManage IT Assistant User's Guide* (available at support.dell.com/support/edocs/software/smitasst/7.0).

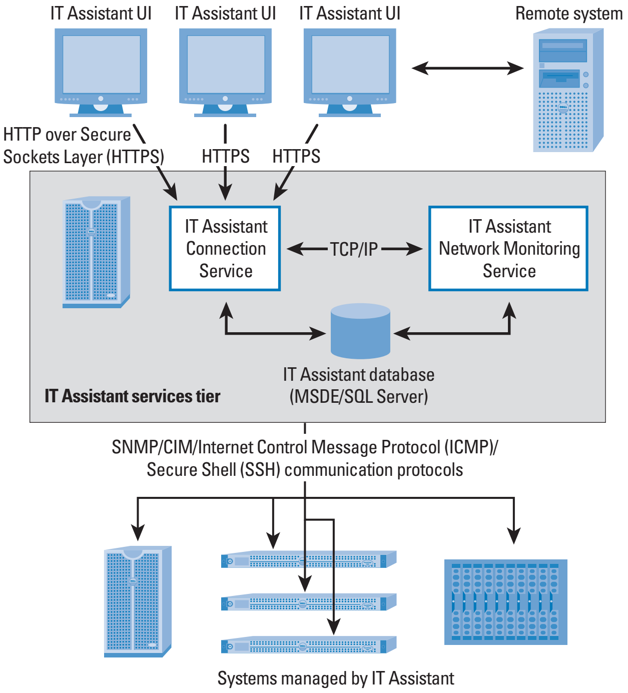
IT Assistant UI
IT Assistant UI
IT Assistant UI
Remote system
HTTP over Secure Sockets Layer (HTTPS)
HTTPS
HTTPS
IT Assistant Connection Service
TCP/IP
IT Assistant Network Monitoring Service
IT Assistant database (MSDE/SQL Server)
IT Assistant services tier
SNMP/CIM/Internet Control Message Protocol (ICMP)/ Secure Shell (SSH) communication protocols
Systems managed by IT Assistant

Figure 1. Overall component view within an ITA-managed environment

Features introduced in IT Assistant 7

The latest version of Dell OpenManage IT Assistant includes additional features and enhanced capabilities such as the following: native installation, user authentication, UI design and online help, remote UI access, topology view, phasing out of Desktop Management Interface (DMI) support, troubleshooting, task management, reporting, inventory cycles, dynamic groups, application launch, and single sign-on.

Native installation. The entire suite of Dell OpenManage systems management applications can now be installed using the native installation technology of the OS. For example, ITA uses Microsoft Windows Installer (MSI) technology to install on a Windows OS.

User authentication. ITA now uses OS-based or domain-based authentication; the ITA 6.*x* read/write password is no longer used. See the “Single sign-on” section in this article for additional information.

UI design and online help. The user-friendly ITA UI has been completely redesigned. The layout is arranged to enhance the functional orientation and now includes wizard-based dialogs for performing many standard tasks. The ITA menu bar items have changed. Comprehensive online help is now available, both from the “Help” link at the top right of the ITA window and from content-specific Help buttons within individual dialogs.

Remote UI access. The UI is exclusively Web based, uses Java technology, and now supports Linux systems.

Topology view. In the UI, administrators can go to Views > Topology on the menu bar to display a hierarchical graphical representation of the devices within a network. By double-clicking the icon for the group they want to view, administrators can move down through the hierarchy. In addition, they can display detailed device information by moving the cursor over each icon. In this view, administrators can also perform tasks on the devices such as application launching, inventory and status refreshing, and troubleshooting.

Phasing out DMI support. ITA no longer supports the DMI protocol. As a result, systems running DMI using Dell OpenManage Server Agent 4.5.1 (or earlier) and Dell OpenManage Client Instrumentation 6.0 (or earlier) will not be discovered by ITA.

Troubleshooting. ITA now provides a graphical troubleshooting tool available at Tools > Troubleshooting Tool on the menu bar. This tool can be used to diagnose and resolve discovery and configuration problems, including Simple Network Management Protocol (SNMP) and Common Information Model (CIM) issues. Administrators can also use the tool to test device and e-mail connectivity.

Task management. ITA now provides an updated tasking functionality that allows administrators to set up and remotely run certain tasks on enterprise systems, including device control (shut down and wake up) and command-line execution. To use the tasking functionality, administrators should go to Manage > Tasks on the menu bar.

Reporting. ITA now offers a customizable reporting feature that gathers data from the Microsoft Data Engine (MSDE) or Microsoft SQL Server database. The reports are based on the data collected in the last discovery or inventory cycle. The report interface wizard is designed to allow administrators to select actual fields in the ITA database and then create reports containing details about the hardware devices being managed by ITA—including servers, switches, and storage devices; BIOS, firmware, and driver versions; and other asset or cost-of-ownership details. The output format for these reports can be specified as HTML, XML, or comma-separated value (that is, for use in a spreadsheet tool such as Microsoft Excel). ITA saves the report definitions for later use and retrieval. To use the ITA report wizard, administrators should go to Views > Reports on the menu bar. These reports can be accessed remotely as well.

Inventory cycles. ITA collects inventory information such as software and firmware versions as well as device-related information about memory, processors, power supplies, Peripheral Component Interconnect (PCI) cards and embedded devices, and storage. This information is stored in the ITA database and used for generating reports.

Dynamic groups. ITA now allows for creating dynamic groups of devices to enhance the device management and monitoring process. This is accomplished by creating groups based on a complex query that is designed to dynamically build these groups during discovery.

Application launch. ITA provides a consolidated launch point for the following systems management applications and devices: Dell OpenManage Server Administrator (includes support for Dell OpenManage Server Administrator Storage Management), Dell OpenManage Array Manager, Dell OpenManage Client Connector, Dell remote access controllers (RACs), Dell PowerConnect switches, and Dell digital keyboard, video, mouse (KVM) switches.

Single sign-on. The Sign-On option on Windows systems enables all logged-in users to bypass the login page and access ITA by clicking on the IT Assistant icon on their desktop. To allow single sign-on for ITA, administrators should perform the following steps:

1. In Microsoft Internet Explorer, go to Tools > Internet Options on the menu bar.
2. In the Internet Options pop-up window, select the Security tab.
3. Select the "Trusted sites" security zone (the ITA system falls within this category).
4. Click the Custom Level button.
5. Under "User Authentication," select the "Automatic logon with current username and password" radio button.

Changes from previous versions of IT Assistant

This section explains the feature sets in ITA 7 that differ from previous versions of ITA. It also explains best practices for migrating previous versions of ITA to ITA 7 for the changed feature sets, including installation, database migration, custom groups and event filters, device view, application launch, management of discovered devices, browser setup, paging configuration, and group configuration.

Installation. The ITA installation program will no longer, by default, install the Microsoft database engine required by ITA if the engine is absent from the system. Rather, the prerequisite compliance program for the installer will report the absence of a default instance of SQL Server or MSDE and will provide a link, which will install a default instance of MSDE. ITA 7 is designed to install quickly, often in a fraction of the time required by previous versions of ITA.

Database migration. The database schema in ITA 7 is completely redesigned, which makes the database schema between previous ITA versions and ITA 7 incompatible. Hence, an upgrade to ITA 7 from previous versions will not preserve all the database settings. However, a few settings can be preserved during the upgrade to ITA 7:

- Global configuration settings
- Discovery range settings
- Configured event action settings from the Event Management system

Custom groups and event filters. Custom groups and event filters will not be preserved during an upgrade from a previous version of ITA. The discovered individual devices will not be visible via ITA 7 the first time; a discovery will need to be initiated to rediscover these devices. Custom groups will need to be reconfigured. The Event Management system will preserve only the configured event actions, while the event filters will need to be reconfigured.

Device view. The device view includes the Summary, Alert Logs, and ESM Logs (when applicable) tabs by default (see Figure 2). The Status, Asset Info, and Users tab are not included, unlike in previous ITA versions. The Logs tab in previous versions of ITA is now replaced with the Alerts Logs and ESM Logs tabs. The Details tab is available by going to Tools > User Preferences on the menu bar, selecting the View Preferences tab, and clicking the check box adjacent to "Enable Details tab in the Device Tree View."

Application launch. In earlier ITA versions, the application launch point existed on the status page of the device view under each of the element managers. ITA 7 provides a consolidated launch point for various systems management applications: Server Administrator for servers; Client Connector for clients; the PowerConnect console for switches; the console for digital KVM switches; the Array Manager console; and the RAC console for devices such as the Dell Remote Assistant Card II (DRAC II), Dell Remote Access Card III (DRAC III), Dell Remote Access Controller 4 (DRAC 4), Embedded Remote Access (ERA) controller, and the ERA/MC controller. It also provides a launch point for Remote Desktop Connection, SOL Proxy, and Telnet. These applications can be launched from within ITA using one of the following methods:

- Select a device in the device tree view and right-click to view the action menu. The action menu can also be accessed by selecting Actions > Application Launch on the menu bar.
- Select a device icon in the topology view and right-click to view the action menu. The action menu can also be accessed by selecting Actions > Application Launch on the menu bar.

Management of discovered devices. Because the status page is not available in ITA 7 for any discovered devices, administrators cannot reset configurable thresholds on temperature probes, fan probes, or voltage probes—nor can they change asset tags on the discovered devices in ITA via the status page. However, administrators can configure these thresholds using the following methods:

- To configure a single device, right-click on the device icon in the device tree view and select "Application Launch." The device tree view can also be accessed by selecting Actions > Application Launch on the menu bar. Select the

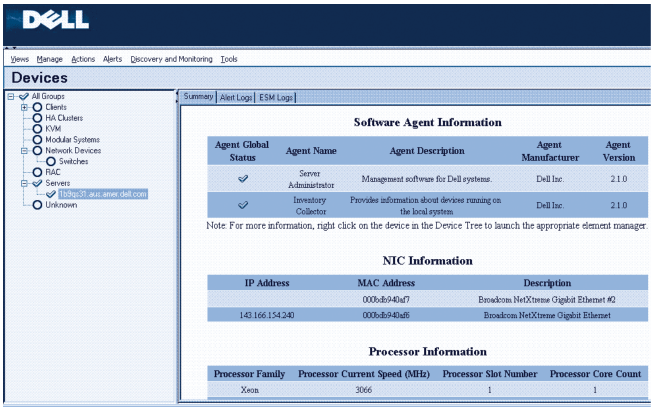

Software Agent Information

Agent Global Status	Agent Name	Agent Description	Agent Manufacturer	Agent Version
✓	Server Administrator	Management software for Dell systems.	Dell Inc.	2.1.0
✓	Inventory Collector	Provides information about devices running on the local system	Dell Inc.	2.1.0

Note: For more information, right click on the device in the Device Tree to launch the appropriate element manager.

NIC Information

IP Address	MAC Address	Description
	000bdb940af7	Broadcom NetXtreme Gigabit Ethernet #2
143.166.154.240	000bdb940af6	Broadcom NetXtreme Gigabit Ethernet

Processor Information

Processor Family	Processor Current Speed (MHz)	Processor Slot Number	Processor Core Count
Xeon	3066	1	1

Figure 2. Default ITA device view page

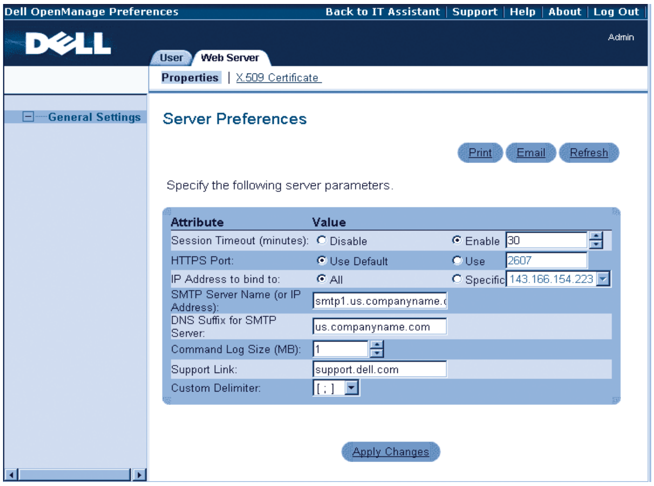
Dell OpenManage Preferences
Back to IT Assistant | Support | Help | About | Log Out
DELL
Admin
User
Web Server
Properties | X.509 Certificate
General Settings
Print
Email
Refresh
Apply Changes

Server Preferences

Specify the following server parameters.

Attribute	Value	
Session Timeout (minutes):	○ Disable	● Enable 30
HTTPS Port:	● Use Default	○ Use 2607
IP Address to bind to:	● All	○ Specific 143.166.154.223
SMTP Server Name (or IP Address):	smtp1.us.companyname.c	
DNS Suffix for SMTP Server:	us.companyname.com	
Command Log Size (MB):	1	
Support Link:	support.dell.com	
Custom Delimiter:	[;]	

Figure 3. Configuring e-mail alert settings on the ITA Server Preferences page

appropriate systems management application, and it will launch in context (for example, Server Administrator for servers or Client Connector for clients). Reset the configurable values within this application. Once the element application is launched, configure the appropriate element threshold.

- For configuring multiple devices like servers or clients (such as desktops, workstations, and notebooks), use the tasking functionality by selecting Manage > Tasks on the menu bar. Create a new command-line task. For the task type, select "Remote Server Administrator Command Line"[2] if configuring servers and "Remote Client Instrumentation Command Line"[3] if configuring clients.

Browser setup. Beginning with ITA 7, the UI (console view) can be remotely accessed via different Web browsers on both Microsoft Windows and Red Hat Enterprise Linux operating systems. The UI is no longer required to run as a published Web application under Microsoft IIS.

ITA supports the Mozilla 1.7.1 Web browser on Red Hat Enterprise Linux and Microsoft Windows operating systems. However, the Mozilla 1.7.3 Web browser is supported only on Red Hat Enterprise Linux.

Paging configuration. ITA no longer uses the WinBEEP paging application to send paging alert actions. Instead, it supports paging through e-mail alerts.

To configure e-mail alert actions, administrators must configure—via the Server Preferences page of ITA—the Simple Mail Transfer Protocol (SMTP) server name or IP address as well as the Domain Name System (DNS) suffix for the SMTP server. Administrators access this page by clicking the "Preferences" link on the title bar next to the "Support" link. They then select the Web Server tab once in the Server Preferences section and enter the SMTP server name and the DNS suffix for the SMTP server (see Figure 3). Finally, they click the Apply Changes button, which will restart the ITA services. Administrators can click the Email button to test the e-mail connectivity.

Group configuration. The group configuration utility is no longer supported in ITA 7 because the database schemas in this version differ drastically from previous versions. Administrators must recreate the custom groups, which can be accomplished using queries to rebuild the custom groups dynamically.

In the Add Group Wizard, administrators should select the "Select a query" radio button and then click the New button to create a new query. This will then take administrators to the query editor, where they can build a query using various system attributes. If a query is already built, then administrators can select the query from the drop-down menu and proceed.

Enhanced tool for managing the IT infrastructure

Dell OpenManage IT Assistant 7 offers several features that can help system administrators take control of their IT infrastructures. By understanding key differences between earlier versions of ITA and ITA 7, administrators can determine whether migrating to ITA 7 is appropriate for their overall enterprise IT environment and how such a migration might affect end users.

Krishna Mohan is a software engineer consultant in the Dell OpenManage Software Development Group. He has worked on several systems management solutions. Krishna has a B.S. in Engineering from Mangalore University in India and an M.S. in Electrical Engineering from The University of Texas at Austin.

Scott Thomas is a software engineer in the Dell OpenManage Software Development Group. He has a B.S. in Electrical Engineering from Florida A&M University and plans to enter the M.B.A. program at the University of Michigan this fall.

FOR MORE INFORMATION

Dell OpenManage:
www.dell.com/openmanage

[2] The "Remote Server Administrator Command Line" task can be executed only on servers running Dell OpenManage Server Administrator 4.3 (or later). During task creation, any servers in the device tree view that do not meet this criteria will not be selectable for executing the task.

[3] For the "Remote Client Instrumentation Command Line" task type to appear in the command-line task, Dell OpenManage Client Connector 7.0 needs to be installed on the management station running the ITA middle tier, and this software must be detected by ITA. A restart of the ITA services is required if Client Connector is installed after ITA.

Deploying Dell OpenManage Server Administrator with Altiris Deployment Solution

Altiris® Deployment Solution™ is designed to be an intuitive, cost-effective tool for deploying, configuring, and managing servers and software from a centralized location. This article examines how administrators can use Altiris Deployment Server to deploy Dell™ OpenManage™ Server Administrator 4.3 (or later), which leverages installation technologies that are native to the Microsoft® Windows® and Linux® operating systems.

BY BERNARD BRIGGS AND KIT LOU

Related Categories:

Altiris

Dell OpenManage

Linux

Microsoft Windows

Operating system deployment

Systems management

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

Altiris Deployment Solution is part of the Altiris Server Management Suite™ for IT life-cycle management, which enables administrators to cost-effectively deploy and manage servers from a centralized management console that helps streamline OS deployment, software deployment, and configuration tasks. Altiris Deployment Solution supports multiple operating systems, including Microsoft Windows and Linux. Because installers for Dell OpenManage Server Administrator 4.3 (or later) leverage installation technologies that are native to the Windows and Linux operating systems, these installers can be integrated easily into the Altiris Deployment Solution. This article explains methods for deploying Windows- and Linux-based Dell OpenManage 4.3 (or later) software installation packages using Altiris Deployment Solution.

Installing and setting up Altiris Deployment Solution

A simple setup of Altiris Deployment Solution includes the Altiris Deployment Server, a SQL database, and an optional Altiris Preboot Execution Environment (PXE) server. The Altiris Deployment Server itself runs on the Microsoft Windows 2000 Server or the Windows Server™ 2003 OS and requires a deployment database in the form of Microsoft Data Engine (MSDE) or Microsoft SQL Server. Managed servers can be Windows or Linux systems and require an agent to be installed; alternatively, servers can be provisioned from bare metal. The Altiris Deployment Server provides a user-friendly console to allow easy, remote management of the servers. This article focuses on a simple installation; for in-depth details, refer to the *Altiris Deployment Solution 6.1 Product Guide* at www.altiris.com/docs/support/deploymentserver/6.1/deployment.pdf.

For a simple installation, administrators must first install Altiris Deployment Server on a Windows 2000 or Windows Server 2003 system with MSDE or Microsoft SQL Server. MSDE is available from Microsoft. It is also available on the Dell Systems Management Consoles CD because it is a prerequisite for the installation of Dell OpenManage IT Assistant. On a Windows 2000 system, administrators also need to install Microsoft .NET Framework 1.1, which is available for download from Microsoft. Administrators who are installing Altiris Deployment Solution for the first time will also need specific DOS files available from a Windows 95 or Windows 98 CD. These DOS files are required for DOS-specific preboot tasks that leverage the Dell Deployment Toolkit bundled with the Altiris solution. After a successful installation of the Altiris Deployment Server, administrators will be able to see three entries in the Add/Remove Programs window: Altiris eXpress Deployment

Console, Altiris eXpress Deployment DataStore, and Altiris eXpress Deployment Server.

For each managed server on which administrators intend to install Dell OpenManage Server Administrator, they must first install the Altiris agent software. For Windows systems, the simplest way to install the agent is to click the Remote Agent Install icon in the task bar of the Deployment Server Console. For systems running Red Hat® Enterprise Linux 2.1 (or later), administrators must install the Linux agent manually—push installation is not available out-of-the-box for Linux servers, although many administrators deploy Linux images with the Altiris agent already installed. When a simple installation of the Deployment Server is performed, the Altiris Linux agent is available in the c:\Program Files\Altiris\eXpress\Deployment Server directory. The name of the agent is altiris-adlagent-*x.y-b*.i386.rpm—for example, altiris-adlagent-2.2-14.i386.rpm. Administrators must transfer this Red Hat Package Manager (RPM™) file locally to the Linux system and run a simple RPM installation. After the installation, they can configure the agent by running the configure utility under the /opt/altiris/deployment/adlagent/bin directory. Once the Altiris agent is installed, an icon representing the server is created in the Altiris console. From that point forward, Linux servers are managed through many of the same features used to manage Windows servers.

With a simple installation as described in this article, administrators can launch the Deployment Server Console, view the Windows and Linux managed servers, and begin setting up jobs to remotely deploy Server Administrator onto these remote managed systems. Altiris Deployment Solution provides a sample job for installing Server Administrator on Dell servers. This article describes some advanced concepts beyond the scope of the sample job. Figure 1 shows the Deployment Server Console with two active computers: a Dell PowerEdge™ 2800 server running Windows Server 2003 and a PowerEdge 2600 server running Linux. The Jobs pane of this console shows several Altiris Deployment Solution jobs—described in the following sections—that can be executed on these servers.

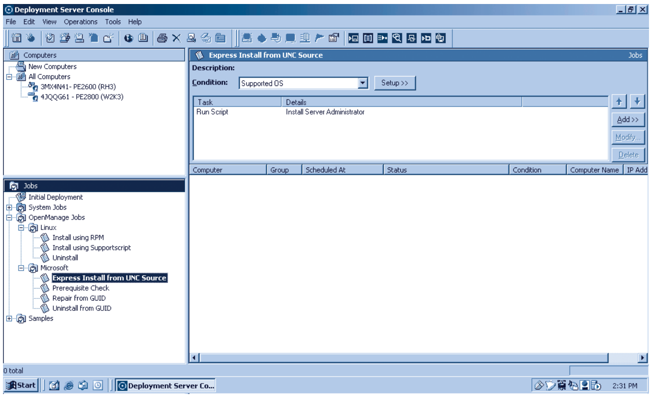

Description:

Condition: Supported OS

Task	Details
Run Script	Install Server Administrator

Figure 1. Altiris Deployment Server Console

Deploying Server Administrator on Windows

This section describes the various Altiris Deployment Solution jobs that can be used to deploy Dell OpenManage Server Administrator onto Windows-based servers.

Prerequisite checking

The Prerequisite Checker is a utility independent of the Server Administrator installer. It should be used in tandem with the installer to help ensure that the server meets the specified requirements prior to deploying and executing the installer via msiexec.exe.

The Prerequisite Checker utility is located at /srvadmin/windows/PrerequisiteChecker/runprereqchecks.exe on the Dell Installation and Server Management CD. Administrators can script the execution of this utility with an `/s` parameter that invokes the utility silently. This invocation returns one of several return or error codes. For more details, see the Run Script task in Figure 2. In addition to the return code, administrators can query the Windows Registry for feature-level detail regarding which features will be affected by a prerequisite warning or error.

The installer package expects the Prerequisite Checker to have been executed prior to the installation. Upon launch, it will query the registry for the prerequisite check results and then deselect default features that indicate a warning and disable features that indicate an error. If the prerequisites are not run, the installer will apply the features selected for installation and Server Administrator may, as a result, run in a degraded state.

Administrators should use the Altiris Job Wizard and create a job to execute the Prerequisite Checker. Next, they should add a Run Script task to the job and embed the script shown in Figure 2. Administrators should run this job on all servers to which they plan to deploy Server Administrator and correct any failed prerequisites.

Express installation

Several methods are available to install Server Administrator. The Express Install job applies the default features available in the Microsoft Windows Installer (MSI) package. Administrators can execute an express installation by issuing the following command:

```
msiexec /qa /i /srvadmin/windows/SysMgmt/
    SysMgmt.msi /l*v c:/sysmgmt.log
```

The `/qa` parameter invokes the installation silently while the `/l*v` parameter logs the progress of the installation to the specified file. Several return codes could result from executing the installation command. These codes should be evaluated to determine whether the installation was successful. Figure 3 shows common error codes for an express installation on a Windows system.

Note that Altiris jobs can be configured to watch for specific return codes and to reroute job functions based on specific values.

```
@ECHO OFF
REM  Run Prerequisite Checker
REM
\\altiris_1\srvadmin\windows\PreReqChecker\
   RunPreReqChecks.exe /s
SET ERRORLEV=%ERRORLEVEL%
REM Prereq Error
IF %ERRORLEV%==3 GOTO ErrorLabel
REM Prereq Warning
IF %ERRORLEV%==2 GOTO WarningLabel
REM Prereq Information
IF %ERRORLEV%==1 GOTO SuccessLabel
REM Prereq Success
IF %ERRORLEV%==0 GOTO SuccessLabel
REM Other errors
GOTO OtherLabel

:SuccessLabel
EXIT 0

:FailureLabel
EXIT %ERRORLEV%

:WarningLabel
REM Prerequisite Check reported a Warning
EXIT 0

:OtherLabel
REM
REM  process return code if desired
REM  - Prerequisite checker did not run.
REM
REM  -1  Windows Host Scripting Error
REM  -2  Operating system not supported
REM  -3  User lacks Administrator privileges
REM  -4  Unused
REM  -5  Failed to change working directory to %TEMP%
REM  -6  Destination directory does not exist
REM  -7  Internal error
REM  -8  An instance is already running
REM  -9  Windows Host Scripting Error: wrong
REM      version, corrupted or not installed
REM -10  Error with scripting environment

EXIT %ERRORLEV%
```

Figure 2. Prerequisite Checker sample script for Windows systems

For example, the default setting for a return code of zero will allow the job to advance to the next sequenced task included in the job. For a return code of 1603, administrators may want to use the Altiris LogEvent utility to write a specific message back to the Altiris console and store it in the database. In the case of a 1602 code, administrators may want the console to display "The user canceled the installation."[1]

Administrators can use the Altiris Job Wizard to create a Software Deployment Job. When the wizard prompts, administrators should browse to the /srvadmin/windows/SysMgmt/SysMgmt.msi installation package on the Dell Installation and Server Management CD. In the Package Distribution Options pane associated with the package, administrators should select "Copy all files and subdirectories" and add the logging parameter to the optional parameters field as discussed earlier. Administrators should run this job on all servers on which they want Server Administrator installed.

Note: Best practices recommend that the installation package is deployed only to servers that have successfully executed the Prerequisite Checker job with no warnings or errors. Administrators can enforce this recommendation by defining specific return code actions. For more information, see the *Altiris Deployment Solution 6.1 Product Guide* at www.altiris.com/docs/support/deploymentserver/6.1/deployment.pdf.

For sample scripts to repair installation and uninstall Dell OpenManage Server Administrator, visit *Dell Power Solutions* online at www.dell.com/powersolutions.

Value	Description
0	The action completed successfully.
1601	The Windows Installer service could not be accessed.
1602	The user canceled the installation.
1603	A fatal error occurred during the installation.
1618	Another installation is already in progress.
1619	The installation package could not be opened. The administrator should verify the package exists and that it is accessible.
1638	Another version of this product is already installed.
1639	An invalid command-line argument was provided.
1641	The installer has started a reboot after a successful installation.

Figure 3. Common return codes after an express installation on a Windows system

[1] For more information about using the Altiris LogEvent and WLogEvent utilities within custom code, see page 170 in the *Altiris Deployment Solution 6.1 Product Guide* at www.altiris.com/docs/support/deploymentserver/6.1/deployment.pdf.

Deploying Server Administrator on Linux

This section describes the various Altiris Deployment Solution jobs that can be used to deploy Dell OpenManage Server Administrator onto Linux-based servers.

Installation using RPM

To deploy Server Administrator onto a Linux system using the express setup, administrators should use the Altiris Job Wizard to create a job with three tasks. The first task copies source RPM files to a temporary location on the system; the second task performs the actual installation from the temporary location and starts the related services; and the third task cleans up the files in the temporary location.

```
rpm -ivh /tmp/srvadmin/* 2>&1 | tee
    -a /tmp/ominstall.log
ecode=$?
case "$ecode" in
0)
/opt/altiris/deployment/adlagent/bin/logevent
    -c:0 -l:1 -ss:"Install Succeeded."
sh /opt/dell/srvadmin/omil/supportscripts/
    srvadmin-services.sh start
rcode=0;;
*)
/opt/altiris/deployment/adlagent/bin/logevent
    -c:$ecode -l:3 -ss:"Install Failed."
rcode=255;;
esac
exit $rcode
```

Figure 4. RPM sample script for installing Server Administrator on Linux systems

```
sh /tmp/srvadmin/supportscripts/srvadmin-install.sh
    --express 2>&1 | tee -a /tmp/ominstall.log
ecode=$?
case "$ecode" in
0)
/opt/altiris/deployment/adlagent/bin/logevent
    -c:0 -l:1 -ss:"Install Succeeded."
sh /opt/dell/srvadmin/omil/supportscripts/
    srvadmin-services.sh start
rcode=0;;
*)
/opt/altiris/deployment/adlagent/bin/logevent
    -c:$ecode -l:3 -ss:"Install Failed."
rcode=255;;
esac
exit $rcode
```

Figure 5. Supportscript sample script for installing Server Administrator on Linux systems

Administrators add a Copy File task to copy the desired directory and files to the remote server. If deploying Server Administrator to an eighth-generation PowerEdge server with a Dell Remote Access Controller 4 (DRAC 4), such as the PowerEdge 2800 server, administrators should use the /srvadmin/linux/express-install-with-RAC4 directory on the Dell Installation and Server Management CD. If deploying to a PowerEdge system with a DRAC 3, such as the PowerEdge 2600 server, administrators should use the /srvadmin/linux/express-install-with-RAC3 directory on the Dell Installation and Server Management CD. To help ensure that all files within the directory are copied, administrators should choose the “Copy directory” option and set the destination path to /tmp/srvadmin. Next, they should add a Run Script task, embed the script shown in Figure 4, and select Linux as the OS.

Administrators can use Altiris Deployment Solution to deploy software onto multiple Dell servers efficiently and cost-effectively from an intuitive, centralized management console.

Next, administrators should add another Run Script task, embed the following script, and choose Linux as the OS:

```
rm -rf /tmp/srvadmin
```

This installation job can then be deployed to the appropriate Server Administrator servers.

Note: The Prerequisite Checker utility is not available on Linux.

Installation using supportscript

To deploy Server Administrator onto a Linux system using the srvadmin-install support script, administrators should use the Altiris Job Wizard to create a job with four tasks. The first task copies source RPM files from the RPM System (RPMS) directory to a temporary location on the system; the second task copies source files from the supportscripts directory to a temporary location on the system; the third task performs the actual installation from the temporary location by launching the srvadmin-install support script and starts the associated services; and the fourth task cleans up the files in the temporary location.

Administrators add a Copy File task to copy the RPMS directory under the srvadmin/linux folder on the Dell Installation and Server Management CD. To help ensure that all files within the directory are copied, administrators should choose the “Copy directory” option and set the destination path to /tmp/srvadmin/RPMS.

Next, administrators should add a second Copy File task to copy the supportscripts directory under the srvadmin/linux folder on the Dell Installation and Server Management CD. They should then choose the “Copy directory” option to help ensure that all files within the directory are copied and set the destination path to /tmp/srvadmin/supportscripts. Finally, they should add a Run Script task, embed the script shown in Figure 5, and select Linux as the OS.

Administrators should add another Run Script task, embed the following script, and choose Linux as the OS:

```
rm -rf /tmp/srvadmin
```

This installation job can then be deployed to the appropriate Server Administrator servers.

Administrators can customize the parameters to install specific components using the srvadmin-install.sh script. Figure 6 describes the silent installation options of the srvadmin-install.sh script. These options can be combined, such as `sh srvadmin-install.sh --diags --web` or `sh srvadmin-install.sh -dw`.

Uninstalling Server Administrator

To remove Server Administrator from a Linux system, administrators can use the Altiris Job Wizard to create a job and add a Run Script task. They should embed the script shown in Figure 7 in the task and select Linux as the OS. Administrators can then deploy the uninstall job to the appropriate Server Administrator servers.

Deploying Server Administrator with the Altiris Deployment Solution

The procedures presented in this article can enable administrators to use Altiris Deployment Solution to deploy software onto multiple Dell servers efficiently and cost-effectively from an intuitive, centralized management console. Once a job is built, administrators can use it again and again for different Dell servers within their organizations. This article describes sample Altiris jobs for deploying Dell OpenManage Server Administrator onto Windows- and Linux-based Dell servers. These jobs can help administrators install and uninstall Server Administrator. On Windows-based systems, these jobs also can perform prerequisite checking and repairs. By understanding how to take advantage of such capabilities, administrators can use Altiris Deployment Solution to help simplify systems management in Dell server environments.

Option	Description
`-x	--express`	Installs all components, including remote access controller (RAC) components; any other options passed are ignored
`-b	--base`	Installs base components
`-d	--diags`	Installs diagnostic components, including base components
`-s	--storage`	Installs storage components, including base components
`-r	--rac`	Installs applicable RAC components, including base components
`-w	--web`	Installs Web server components, including base components

Figure 6. Silent installation options of the svradmin-install.sh script for Linux systems

```
rpm -e 'rpm -qa | grep srvadmin' 2>&1 | tee
   /tmp/ominstall.log
ecode=$?
case "$ecode" in
0)
/opt/altiris/deployment/adlagent/bin/logevent
   -c:0 -l:1 -ss:"Uninstall Succeeded."
rcode=0;;
*)
/opt/altiris/deployment/adlagent/bin/logevent
   -c:$ecode -l:3 -ss:"Uninstall Failed."
rcode=255;;
esac
exit $rcode
```

Figure 7. Uninstall sample script for Linux systems

Bernard Briggs is the engineering manager of the Dell OpenManage Install Development Group. He has more than seven years of experience with Dell. Bernard has a bachelor’s degree in Computer Science from The University of Texas at Austin.

Kit Lou is a senior engineer in the Dell OpenManage Install Development Group and is focused on various software installation and deployment technologies. He has a B.S. in Computer Science from Loyola University, New Orleans, and an M.S. in Computer Science from The University of Texas at Austin.

FOR MORE INFORMATION

Altiris Deployment Solution:
www.altiris.com

Dell OpenManage:
www.dell.com/openmanage

Advancing the Update Process for Dell Server Components Using

Altiris Patch Management and Deployment Tools

Altiris and Dell have worked closely together to offer IT administrators two advanced, integrated tools for updating Dell server components: Altiris® Patch Management Solution™ for Dell Servers and Altiris Deployment Solution™ for Dell Servers. Each approach is optimized to address the update needs of Dell server hardware at two specific points in the server life cycle: the provisioning or repurposing stage and the production stage.

BY TODD MITCHELL, HECTOR VALENZUELA, KEVIN WINERT, AND LANDON HALE

Related Categories:

Altiris

Change management

Dell PowerEdge servers

Linux

Microsoft Windows

Operating system deployment

System deployment

Systems management

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

Altiris Deployment Solution for Dell Servers helps automate the update process, enabling Dell™ server components in both Microsoft® Windows® and Linux® OS environments to be updated as part of the server build or rebuild process—essentially out-of-the-box without further administrative intervention. Altiris Patch Management Solution for Dell Servers is designed to address the needs of servers already in production. It offers a policy-driven approach for detecting required hardware updates and then automates their download and distribution.

In addition, Altiris Patch Management Solution for Dell Servers integrates with Altiris Patch Management Solution for Windows, providing a single console for both flash-updating Dell hardware components and patching a Windows OS. Figure 1 shows the Altiris console in which a drop-down list allows administrators to filter available updates by vendor before launching the patch wizard to build a distribution policy.

Altiris Deployment Solution for Dell Servers, Altiris Patch Management Solution for Dell Servers, and Altiris Patch Management Solution for Windows are all included in a cost-effective suite: the Altiris Server Management Suite for Dell Servers. This article compares the two Altiris packages that are designed to update Dell server hardware and examines how their feature sets complement each other.

Using Altiris Deployment Solution for Dell Servers to help automate server builds

Altiris Deployment Solution for Dell Servers provides comprehensive functionality for rapidly provisioning and repurposing Dell servers from a bare-metal state. This package includes the following prebuilt sets of server management tasks, or *jobs* (see Figure 2):

- **Pre-OS hardware configuration:** Leverages the Dell OpenManage™ Deployment Toolkit to provide BIOS, Dell Remote Access Controller (DRAC), RAID, and baseboard management controller (BMC) configuration; additional provided jobs can be used to create the Dell Utility Partition, clear the Master Boot Record, and so forth
- **Scripted or image-based installation:** Supports Windows and Linux operating systems with Dell-specific drivers

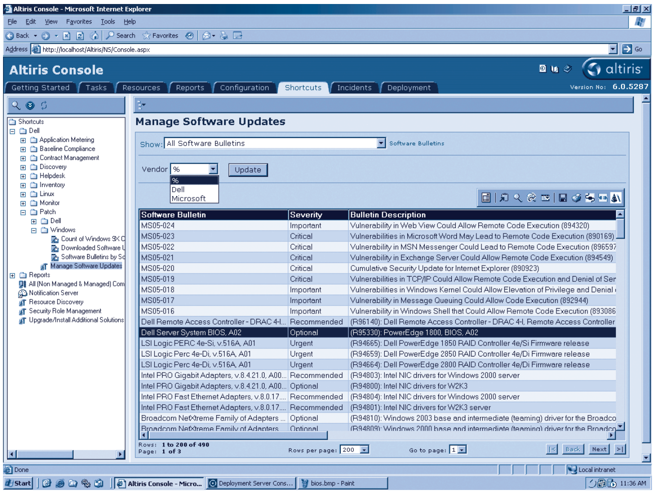

Manage Software Updates

Show: All Software Bulletins — Software Bulletins

Vendor: % — Update

%
Dell
Microsoft

Software Bulletin	Severity	Bulletin Description
MS05-024	Important	Vulnerability in Web View Could Allow Remote Code Execution (894320)
MS05-023	Critical	Vulnerabilities in Microsoft Word May Lead to Remote Code Execution (890169)
MS05-022	Critical	Vulnerability in MSN Messenger Could Lead to Remote Code Execution (896597
MS05-021	Critical	Vulnerability in Exchange Server Could Allow Remote Code Execution (894549)
MS05-020	Critical	Cumulative Security Update for Internet Explorer (890923)
MS05-019	Critical	Vulnerabilities in TCP/IP Could Allow Remote Code Execution and Denial of Ser
MS05-018	Important	Vulnerabilities in Windows Kernel Could Allow Elevation of Privilege and Denial
MS05-017	Important	Vulnerability in Message Queuing Could Allow Code Execution (892944)
MS05-016	Important	Vulnerability in Windows Shell that Could Allow Remote Code Execution (893086
Dell Remote Access Controller - DRAC 4-I...	Recommended	(R96140): Dell Remote Access Controller - DRAC 4-I, Remote Access Controller
Dell Server System BIOS, A02	Optional	(R95330): PowerEdge 1800, BIOS, A02
LSI Logic PERC 4e-Si, v.516A, A01	Urgent	(R94665): Dell PowerEdge 1850 RAID Controller 4e/Si Firmware release
LSI Logic Perc 4e-Di, v.516A, A01	Urgent	(R94659): Dell PowerEdge 2850 RAID Controller 4e/Di Firmware release
LSI Logic Perc 4e-Di, v.516A, A01	Urgent	(R94664): Dell PowerEdge 2800 RAID Controller 4e/Di Firmware release
Intel PRO Gigabit Adapters, v.8.4.21.0, A00...	Recommended	(R94803): Intel NIC drivers for Windows 2000 server
Intel PRO Gigabit Adapters, v.8.4.21.0, A00...	Optional	(R94800): Intel NIC drivers for W2K3
Intel PRO Fast Ethernet Adapters, v.8.0.17....	Recommended	(R94804): Intel NIC drivers for Windows 2000 server
Intel PRO Fast Ethernet Adapters, v.8.0.17....	Recommended	(R94801): Intel NIC drivers for W2K3 server
Broadcom NetXtreme Family of Adapters ...	Optional	(R94810): Windows 2003 base and intermediate (teaming) driver for the Broadco
Broadcom NetXtreme Family of Adapters	Optional	(R94809): Windows 2000 base and intermediate (teaming) driver for the Broadco

Rows: 1 to 200 of 490
Page: 1 of 3
Rows per page: 200
Go to page: 1
|< Back Next >|

Figure 1. Altiris console drop-down list to filter available updates by vendor before launching the Altiris patch wizard

- **Application distribution and installation:** Includes example jobs for popular applications such as Dell OpenManage Server Administrator, Microsoft SQL Server, the Apache Web server, and so forth

Altiris jobs package the preceding capabilities into customized sets of ordered tasks that display as a single drag-and-drop event in the Altiris console. One job can be created for each server configuration an organization requires. For example, an organization may have three configurations for a Web server and four configurations for a database server—with Altiris Deployment Solution, every configuration becomes a single build job.

Altiris Deployment Solution for Dell Servers also provides sample jobs for updating server components using Dell Update Packages. These sample jobs will install Dell Update Packages in the order recommended by Dell, and they can be included within more comprehensive server build jobs to help ensure that all needed Dell updates are applied as part of the build process. This capability can be valuable because it provides administrators with the flexibility to adapt to a variety of provisioning scenarios. For example, before some server applications can be installed or configured, they may require a Dell server to first adhere to a specific BIOS level, Embedded Server Management (ESM) firmware revision, BMC firmware revision, DRAC firmware revision, RAID driver version or firmware revision, a network interface card (NIC) driver version, or a combination of such specifications.

Altiris Deployment Solution can address such scenarios while streamlining the server deployment process. Updating server hardware components as an integrated function of the server's build can eliminate the need for additional agents to install, configure, determine which updates are needed, and then apply them. Altiris offers a unified approach that is designed to address comprehensive server provisioning needs through all aspects of hardware configuration and updates, OS installation, and application installation. In any environment in which fast server builds or rapid repurposing is required, the capability of Altiris Deployment Solution to apply Dell hardware updates as an integrated part of a customized server build can help organizations meet business requirements quickly and efficiently.

Employing Altiris Patch Management Solution for Dell Servers to update servers

Altiris Patch Management Solution for Dell Servers is designed to address the ongoing need to assess and maintain component revision levels on Dell servers already in production. These servers may have been provisioned many months or even years ago, but they still must be periodically updated to maintain the latest hardware component versions from Dell. The needs for applying updates in production vary significantly from those in the provisioning or repurposing stage of a server's life cycle.

Organizations that depend on a large number of servers often struggle to maintain an awareness of the various updates that Dell makes available and the different servers to which these updates apply. As a result, many IT organizations choose to update server components only when there is an associated problem. This approach obviously increases risks and is not recommended.

Altiris and Dell have streamlined the update process by enabling periodic downloads of Dell Update Packages directly from Dell and then scanning an organization's IT environment to determine existing component revision levels. By correlating the results into a series of easy-to-read reports as shown in Figure 3, Altiris Patch Management Solution for Dell Servers helps administrators quickly determine which servers lack the requisite updates and then automates the download and distribution of updates to the appropriate systems.

Administrators use a simple wizard in the Altiris software to build policies that distribute Dell Update Packages. The wizard

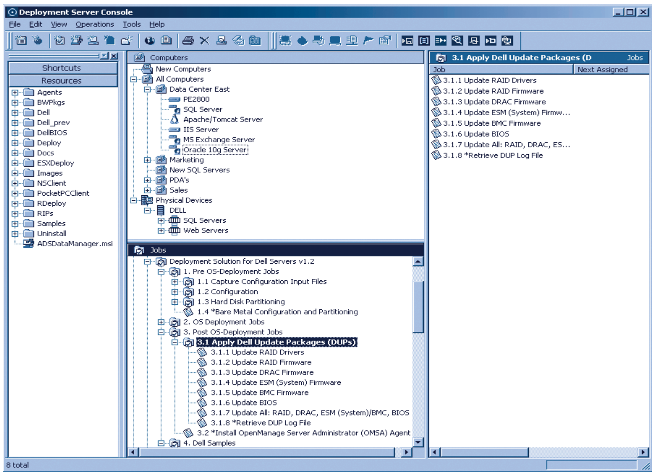

Jobs

Deployment Solution for Dell Servers v1.2
1. Pre OS-Deployment Jobs
1.1 Capture Configuration Input Files
1.2 Configuration
1.3 Hard Disk Partitioning
1.4 *Bare Metal Configuration and Partitioning
2. OS Deployment Jobs
3. Post OS-Deployment Jobs
3.1 Apply Dell Update Packages (DUPs)
3.1.1 Update RAID Drivers
3.1.2 Update RAID Firmware
3.1.3 Update DRAC Firmware
3.1.4 Update ESM (System) Firmware
3.1.5 Update BMC Firmware
3.1.6 Update BIOS
3.1.7 Update All: RAID, DRAC, ESM (System)/BMC, BIOS
3.1.8 *Retrieve DUP Log File
3.2 *Install OpenManage Server Administrator (OMSA) Agent
4. Dell Samples

3.1 Apply Dell Update Packages (D — Jobs

Job	Next Assigned
3.1.1 Update RAID Drivers	
3.1.2 Update RAID Firmware	
3.1.3 Update DRAC Firmware	
3.1.4 Update ESM (System) Firmw...	
3.1.5 Update BMC Firmware	
3.1.6 Update BIOS	
3.1.7 Update All: RAID, DRAC, ES...	
3.1.8 *Retrieve DUP Log File	

Figure 2. Prebuilt Altiris Deployment Solution jobs that can update Dell server components during the build process

allows administrators to accept or modify default suggestions for target collections, reboot schedules, multicast settings, and so forth. Administrators can also choose whether updates should be deployed on a one-time basis or whether a given policy is dynamic—that is, whether it will also automatically deploy updates in the future to servers that meet predefined criteria. Soon after a server receives the Altiris agent, these dynamic policies are designed to scan the system and then automatically download needed updates without administrator involvement.

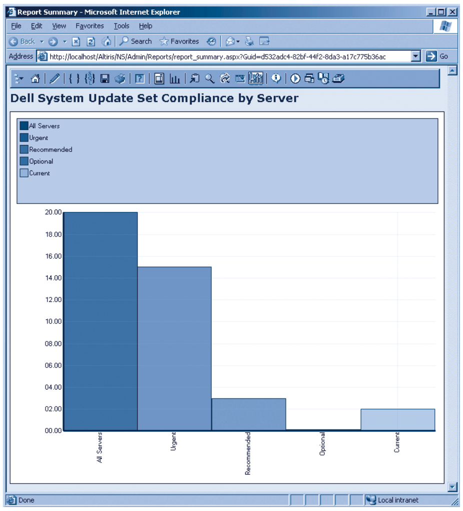
Report Summary - Microsoft Internet Explorer

Address http://localhost/Altiris/NS/Admin/Reports/report_summary.aspx?Guid=d532adc4-82bf-44f2-8da3-a17c775b36ac

Dell System Update Set Compliance by Server

Figure 3. Altiris Patch Management Solution compliance report displaying server update needs by severity

Altiris Patch Management Solution for Dell Servers can also display scan results in easy-to-read graphs. Figure 3 shows the Dell System Update Set Compliance by Server display, which summarizes update needs by severity—that is, how many servers lack urgent, recommended, or optional updates versus how many servers are current. In this graph, administrators can click each bar to view a detailed list of server names and applicable updates. From there, the Dell Update Wizard can be launched to distribute the updates to prebuilt collections of servers.

In addition, Altiris Notification Policies can be configured to automatically send e-mail messages that alert administrators when new updates become available from Dell and indicate which updates are lacking on specific servers (see Figure 4).

The capability of Altiris Patch Management Solution to integrate hardware updates and OS patches into a single console can enhance an IT organization's efficiency by helping to minimize the number of tools administrators need to learn and maintain. At the same time, this approach enables organizations to leverage the efficiencies of a single infrastructure and agent to optimize role and scope security, update delivery, policy creation, and so forth.

Finally, Altiris Patch Management Solution for Dell Servers employs an agent-side rules engine that can queue separate policy requests for updates and reorder them based on Dell best practices. For example, if a given policy is created for a BIOS update in the morning by Administrator A and then Administrator B creates a separate policy in the afternoon to apply an ESM update, both policies will be reserved for execution until the off-peak reboot schedule permits. At that time, the Altiris Patch Management Solution rules engine will examine which updates are required and will re-sequence them if necessary—for example, to apply the ESM update before the BIOS update. Although Dell Update Packages contain logic designed to ensure that servers meet pre-execution requirements, this Altiris safeguard can help ensure the proper update sequence when custom command-line options, or other nonstandard conventions, may be involved.

Applying Dell updates throughout the server life cycle

Although administrators can use either Altiris package discussed in this article to update Dell server hardware, each of the two tools is specifically tailored to address the needs of a key point in the server life cycle: Altiris Deployment Solution focuses on the provisioning or repurposing stage, whereas Altiris Patch Management Solution focuses on the live production stage.

Altiris Patch Management Solution for Dell Servers differs from Altiris Deployment Solution in that it focuses on inventorying server components and comparing existing revision levels to the latest catalog of available Dell updates. Altiris Deployment Solution for Dell Servers does not consider the original server configuration and component levels because they are typically overwritten as part of the automated server build process. The objective of Altiris Deployment Solution is to quickly and completely provision a server—through all phases of hardware, OS, and application deployment and configuration—to meet the build standards as defined by an IT organization. Typically, organizations have little need for up-front scanning and reporting when using Altiris Deployment Solution because the focus is on what the server configuration will be, not on what it was.

Therefore, Altiris Deployment Solution for Dell Servers is best utilized when:

- Server component updates should occur at a precise point in the server build or repurposing workflow.

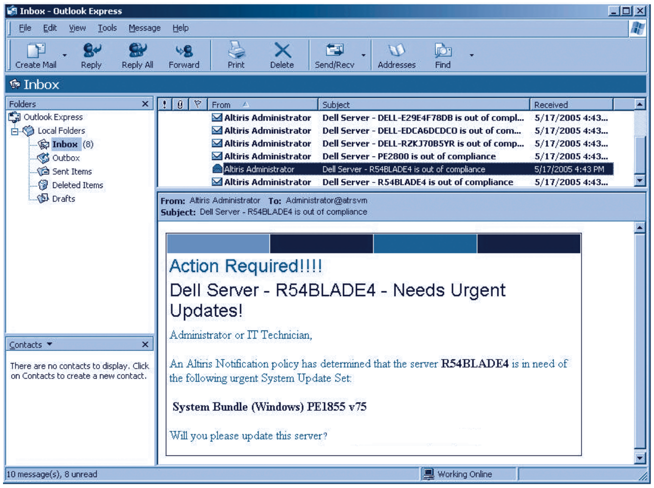
Inbox - Outlook Express

From	Subject	Received
Altiris Administrator	Dell Server - DELL-E29E4F78DB is out of compl...	5/17/2005 4:43...
Altiris Administrator	Dell Server - DELL-EDCA6DCDCO is out of com...	5/17/2005 4:43...
Altiris Administrator	Dell Server - DELL-RZKJ70B5YR is out of comp...	5/17/2005 4:43...
Altiris Administrator	Dell Server - PE2800 is out of compliance	5/17/2005 4:43...
Altiris Administrator	Dell Server - R54BLADE4 is out of compliance	5/17/2005 4:43 PM
Altiris Administrator	Dell Server - R54BLADE4 is out of compliance	5/17/2005 4:43...

From: Altiris Administrator To: Administrator@atrsvm
Subject: Dell Server - R54BLADE4 is out of compliance

Action Required!!!!

Dell Server - R54BLADE4 - Needs Urgent Updates!

Administrator or IT Technician,

An Altiris Notification policy has determined that the server **R54BLADE4** is in need of the following urgent System Update Set:

System Bundle (Windows) PE1855 v75

Will you please update this server?

There are no contacts to display. Click on Contacts to create a new contact.

10 message(s), 8 unread

Figure 4. Altiris Notification Policies e-mail messages, which are sent when specific update conditions exist

- An organization needs to include hardware updates as an immediate and seamless part of the server build process—without relying on or waiting for the installation of other management agents to determine update applicability, download the needed updates, and execute the updates on a predetermined reboot schedule.
- Security or other concerns require Dell updates to be applied before any OS is installed. Although Altiris Deployment Solution for Dell Servers provides comprehensive sample jobs for applying Dell Update Packages in post-OS Windows or Linux environments, some organizations need to apply Dell updates before any OS is installed. For that purpose, Altiris Deployment Solution for Dell Servers provides a sample job that uses Dell tools to apply a BIOS update in a preboot environment, before the installation of an OS.

Altiris Patch Management Solution for Dell Servers is best utilized when:

- Automated applicability assessments are needed to identify which updates are required. The value of this automatic assessment increases with the number of servers in production and the number of different models and configurations being used.
- Dynamic policies are required to reduce the time administrators must spend downloading and applying updates. In production environments, organizations value the accuracy and convenience of automated downloads and update distribution. As server configurations change in the data center, dynamic policies can provide a safety net for detecting needed updates in predefined scenarios and applying them without manual administrative intervention.
- Comprehensive reporting and notification alerts are needed. Altiris Patch Management Solution for Dell Servers provides graphical reporting and various automated notification options to help administrators track the exact revision levels on existing servers. Alerts can trigger e-mail messages, generate help-desk tickets, execute custom processes, forward Simple Network Management Protocol (SNMP) traps, and so forth.
- Flexible scheduling capabilities are needed to apply patches during off-peak hours. Updating production servers often requires a tightly controlled schedule that supports management windows for periodic reboots during off-peak hours.
- Advanced software delivery options including bandwidth throttling, checkpoint recovery, and update multicasting are required. Environments in which remote production servers must be updated across slower-bandwidth links are well suited for the strengths of Altiris Patch Management Solution for Dell Servers. In comparison, Altiris Deployment Solution for Dell Servers offers checkpoint recovery but generally assumes high-bandwidth connections given the nature of server provisioning tasks—such as sending very large image files, scripted installations, and server applications across the network.

Altiris tools are specifically designed to advance the update process at different points throughout the server management life cycle. Because both packages for updating Dell server hardware are offered as part of the Altiris Server Management Suite for Dell Servers, organizations can use them together alongside several other management tools—through a single Altiris console.

Streamlining management of Dell server environments

The integration of Altiris patch management and deployment tools with the Dell OpenManage 4 infrastructure enables highly automated solutions for updating Dell server components. Based on feedback from server administrators, Altiris has designed the Altiris Server Management Suite for Dell Servers to help streamline management across the data center. Two components of this suite—Altiris Deployment Solution for Dell Servers and Altiris Patch Management Solution for Dell Servers—enhance efficiency by targeting the unique needs of Dell environments.

Todd Mitchell is the Dell alliance technical director at Altiris. Todd has a bachelor's degree from Brigham Young University.

Hector Valenzuela is a solutions architect in the Custom Solutions Engineering Group at Dell. Hector has a bachelor's degree in Electrical Engineering from The University of Texas at El Paso.

Kevin Winert is a product marketing manager in the Dell OpenManage Marketing Group at Dell. Kevin has a B.S. in Computer Science from Brigham Young University and an M.B.A. from The University of Houston.

Landon Hale manages Dell's relationship with Altiris within the Dell Global Alliances team. Landon has a B.A. in Political Science from Carleton College and an M.B.A. from the Marshall School of Business at the University of Southern California.

FOR MORE INFORMATION

Altiris Patch Management Solution for Dell Servers:
www.altiris.com/products/dellpatch

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

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

Time-Savings Validation for Dell Server Deployment with Altiris Deployment Solution

Altiris® Deployment Solution™ is designed to reduce administrative time and management costs by enabling a complete server build to be captured as a simple drag-and-drop job. Once created, the job can be easily executed again and again simply by dragging and dropping it onto one or more servers in the Altiris console—with no further involvement by the administrator. This article examines independent testing conducted by KeyLabs to verify that Altiris Deployment Solution can help reduce the time administrators spend deploying servers from hours to minutes.

BY TODD MITCHELL AND LANDON HALE

Related Categories:

Altiris

Dell OpenManage

Operating system deployment

System deployment

Systems management

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

Altiris and Dell jointly commissioned KeyLabs in November 2004 to empirically substantiate that Altiris Deployment Solution can help dramatically reduce the time that an administrator spends deploying and repurposing servers. To evaluate this claim, KeyLabs compared the administrative time required to deploy Dell™ servers using three popular build methods: manual server deployment, the Dell OpenManage™ Server Assistant (DSA) CD, and Altiris Deployment Solution for Dell Servers.

Altiris Deployment Solution is designed to automate the server build process, so an administrator can start or schedule the build process and then move on to other tasks—without having to physically attend installation functions. The objective of the KeyLabs study was to quantify how much deployment time Altiris Deployment Solution could eliminate for server administrators, compared with traditional server build methods. This article summarizes the methodology and findings of that study; the full-length version of the study documentation is available at www.dell.com/downloads/global/solutions/Deployment%20Comparison%20for%20Dell%20PowerEdge%20Servers.pdf.

KeyLabs test methodology

This section briefly describes the test methodology KeyLabs employed, including the types of server builds that were deployed and the techniques used to deploy these builds.

Tested builds

Dell PowerEdge™ 2650 servers were used as the hardware platform for this study because this type of server is one of the most popular PowerEdge server models. Using this Dell server hardware, the KeyLabs team defined two server builds for testing:

- **Server Build 1:** Microsoft® Windows® 2000 Advanced Server with Microsoft SQL Server 2000
- **Server Build 2:** Red Hat® Enterprise Linux® AS 3 with LAMPP (Linux, Apache, MySQL, PHP, and Perl)

For each build, KeyLabs testers adhered to the following configuration steps:

1. Configure BIOS, Dell Remote Access Controller (DRAC), and RAID components (set a RAID-5 configuration).
2. Install the OS and apply any OS updates or patches; apply any Dell-specific drivers to support RAID and network interface card (NIC) components.
3. Install other applications and service packs (Microsoft SQL Server 2000 or LAMPP).
4. Install Dell OpenManage Server Administrator (OMSA).
5. Apply Dell Update Packages to update server components, including the BIOS, Embedded Server Management (ESM) firmware, DRAC, and so forth.

Deployment methods

KeyLabs used three deployment techniques and measured the administrative time required for each technique.

Manual server deployment. Administrators cycle through CDs, software installation screens, BIOS configuration screens, and so forth as needed to install and configure each server. No efficiencies are gained as additional servers are deployed—that is, each server deployed requires fundamentally the same amount of administrative time because administrators must go through the same installation steps as the previous server builds, even if previous servers are the same model of server.

Deployment using the Dell OpenManage Server Assistant CD. The DSA CD is a bootable CD that ships with every Dell PowerEdge server and helps automate OS deployment by automatically scripting the OS installation to include Dell-specific drivers. This CD does not provide setup assistance for the BIOS, DRAC, and baseboard management controller (BMC) components, nor does it help with post-OS tasks such as applying hardware updates or installing applications. This approach helps reduce OS installation times somewhat, but administrators must repeat the same tasks for each server deployed.

Deployment using Altiris Deployment Solution. This method is designed to automate server builds. Using Altiris Deployment Solution requires an initial time investment to install the Altiris environment and to create and test server build jobs. After that process is completed, however, Altiris Deployment Solution enables hands-free, remote provisioning of one or more servers—including hardware, OS, and application configuration.

Special testing considerations for Altiris Deployment Solution

Altiris Deployment Solution requires a specific infrastructure: a dedicated Altiris server and a back-end Microsoft SQL Server 2000 database. To facilitate hands-free functionality, many advanced deployment features such as pre-provisioning (the ability to assign a server build to a Dell Service Tag for a server that has not yet been physically added to the LAN) and rip-and-replace (the ability to dynamically provision blade servers based on their location in a chassis) require the use of an Altiris Preboot Execution Environment (PXE) server.

The installation and configuration of Altiris Deployment Solution—including the PXE server configuration—took KeyLabs approximately 45 minutes. In addition to this installation time, KeyLabs built two jobs, one for Server Build 1 and one for Server Build 2, as defined earlier in this article. Unlike the manual and DSA CD installation methods, Altiris Deployment Solution required an initial investment of approximately six to seven hours before the two server builds in the study could be deployed. Figure 1 conveys the breakdown for this time requirement.

Test environment

Data was compiled for server deployments of various sizes, including deployments of 1 server, 25 servers, 50 servers, and 100 servers. The KeyLabs team isolated physical-layer setup concerns from the study—that is, the time required to unbox, cable, and rack-mount the servers was not recorded. Administrative time was recorded only for provisioning tasks that servers have in common regardless of form factor, including hardware component configuration, OS deployment, and application installation.

A standard Fast Ethernet (100 Mbps) network infrastructure was used—the study did not determine the relative impact of network latencies across the various test cases. The servers tested did not rely on remote storage such as storage area networks or network attached storage. Even though Altiris Deployment Solution can be configured to support such storage systems, the tested servers were configured with internal disk arrays.

Results of the KeyLabs study

KeyLabs testers determined that Altiris Deployment Solution can enable significant reductions in the administrative time required to deploy servers when compared with manual server deployments or server

Task	Time (hours:minutes)
Altiris Deployment Solution installation	0:15
Altiris Deployment Solution configuration and testing (for example, PXE server and Boot Disk Creator)	0:30
Creation and testing of Server Build 1 (Windows) job	3:00
Creation and testing of Server Build 2 (Linux) job	3:00
Total setup time*	6:45

**The KeyLabs engineer who performed the Altiris testing had completed basic Altiris product training. Time spent on the initial server build may be slightly longer or shorter for other organizations, depending on the staff's level of Altiris training and experience.*

Figure 1. Altiris Deployment Solution setup time for both server builds

deployments that leverage the DSA CD. KeyLabs found that, once the initial time was invested to set up the Altiris Deployment Solution infrastructure, administrative time-savings could be compounded exponentially. Simply dragging and dropping server build jobs onto groups of 5, 10, 25, 50, or more servers in the Altiris console took the administrator only a few seconds and the jobs could execute completely unattended.

As shown in Figures 2 and 3, the Altiris server deployment methodology began to generate significant time-savings after the deployment of the third server when compared with the manual installation process; it generated time-savings after the deployment of the fourth server when compared with the DSA CD. This reduction in time was similar for both Windows and Linux environments.

Both Figures 2 and 3 show that the administrative time for the manual and DSA deployment methods increased with each server build, whereas the time for Altiris Deployment Solution decreased after the first server build. Figures 4 and 5 show the actual time spent for each tested server build and deployment method; Figure 4 compares time spent deploying one server, while Figure 5 compares the time spent deploying 25 servers. In both Figures 4 and 5, comparative data for Altiris Deployment Solution shows the administrative time required after the initial server build was completed, as described in Figure 1.

As previously described, Altiris Deployment Solution can require up to three hours to create and test a comprehensive server deployment job. Once that job is created, however, KeyLabs testers found that it requires only five seconds to execute automatically and consistently on one or more servers.

Actual build time for an Altiris job can vary dramatically depending on how the job is built. For example, instead of executing application installations for Dell OMSA and Microsoft SQL Server directly from a network share, the KeyLabs team built jobs to first copy the needed installation files to each server. This step added significant time that slowed deployment speed for each server—contributing more than 30 minutes to the Altiris Deployment Solution

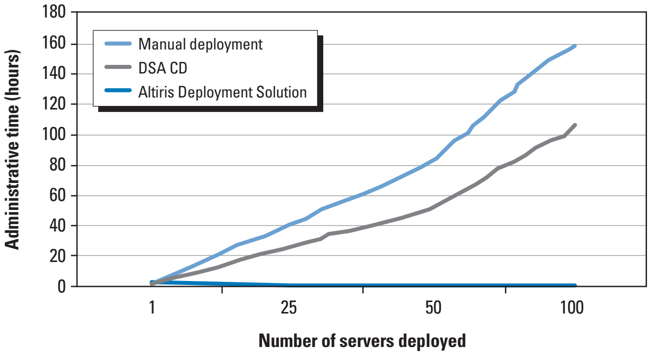

Figure 2. Comparing administrative time for three types of server deployment using Server Build 1 (Windows environment)

BENEFITS OF ALTIRIS DEPLOYMENT SOLUTION

Looking beyond the scope of the study described in this article, the KeyLabs team noted several enterprise IT benefits enabled by Altiris Deployment Solution:

- **Build consistency:** In a manual server build, every keystroke represents a potential difference that can translate into an improperly configured server and, ultimately, downtime. The automation provided by Altiris Deployment Solution is designed to eliminate build discrepancies by helping ensure that every server is built exactly the same way—leading to high availability and easy troubleshooting if a problem occurs.
- **Detailed server history:** Altiris maintains a complete build history of each server's deployment and subsequent management. All status messages returning from the managed server are recorded—including status messages from custom Linux shell scripts, VBScript, and batch files. After a server is deployed, all subsequent management is recorded in the job history, including such tasks as software updates, installations, and configuration tasks (IP address changes, server renaming, and so forth). The manual and DSA CD installation methods do not provide this benefit.
- **Cross-platform support:** The efficiencies discussed in this article are gained in both Microsoft Windows and Red Hat Enterprise Linux server environments—both are managed using the same conventions from the Altiris console.
- **Role and scope security:** The security features of Altiris Deployment Solution can limit access to servers and management capabilities based on a user's role. Altiris Deployment Solution can also provide a complete audit trail detailing who configured what and when.
- **Scheduled execution:** Servers can be tagged for immediate provisioning or scheduled for later deployment during off-peak periods, days or months in advance. As a result, deployment jobs can be assigned and set to execute long before the server hardware arrives on site.
- **Quick recovery:** The recovery process for each server can be significantly reduced by leveraging the same jobs that are used to deploy servers. Jobs can be quickly and easily modified to help recover failed servers or quickly provision additional servers with the same build configuration as a failed server.
- **Resource archiving and reuse:** By building and maintaining Altiris jobs, administrators are enabled to collect a robust archive of images, scripts, software packages, and so forth that can be easily reused. This resource pool can be leveraged by jobs and other projects beyond the initial server deployment. No such archive is compiled for the manual or the DSA CD installation method.

build times listed in Figure 4—but it also helped ensure that a failed network connection would not cause the installation to hang. If the file copy failed, the Altiris console would indicate the failure without starting the application installation.

Once a job begins executing, build status is indicated in real time on the Altiris console. Administrators are not required to be physically present at any servers. Jobs can be configured to send an e-mail message or call into other messaging systems for problem notifications or status updates. Administrators can view the console to determine when server deployments are complete, and they can remotely control servers to verify manually that servers are ready for production. To further reduce manual verification times, Altiris server deployment jobs can be configured to automatically run customer-defined test scripts as a final task. These scripts can be used to exercise a server or validate its configuration. Altiris utilities can be called within customer scripts to send custom status messages to the Altiris console and build logs.

Server build	Manual installation (hours:minutes)	Dell OpenManage Server Assistant CD (hours:minutes)	Altiris Deployment Solution*	
			Administrative time (hours:minutes)	Total build time (hours:minutes)
Server Build 1	2:28	2:06	0:05	1:35
Server Build 2	1:43	1:35	0:05	1:15

** This method also required an initial setup time of 6 hours and 45 minutes to install Altiris Deployment Solution and create Server Build 1 and Server Build 2.*

Figure 4. Comparing administrative deployment time for one Dell PowerEdge 2650 server

Server build	Manual installation (hours:minutes)	Dell OpenManage Server Assistant CD (hours:minutes)	Altiris Deployment Solution*	
			Administrative time (hours:minutes)	Total build time for all 25 servers (hours:minutes)
Server Build 1	41:15	26:15	0:05	4:15
Server Build 2	30:25	16:32	0:05	3:55

** This method also required an initial setup time of 6 hours and 45 minutes to install Altiris Deployment Solution and create Server Build 1 and Server Build 2.*

Figure 5. Comparing administrative deployment time for 25 Dell PowerEdge 2650 servers

In this study, once the initial server build was completed as described in Figure 1, the 25-server deployment using Altiris Deployment Solution was launched with only a few seconds of administrative time, and all 25 servers were ready for production about four hours later—without additional involvement by the administrator. As tested by KeyLabs, a Server Build 1 deployment to 25 Dell PowerEdge 2650 servers using Altiris Deployment Solution reduced administrative time by approximately 90 percent when compared with the manual installation method (see Figure 5).

Improved deployment efficiency with Altiris Deployment Solution

The KeyLabs study described in this article, commissioned jointly by Altiris and Dell, verified that Altiris Deployment Solution can effectively reduce server provisioning time. In addition, advanced Altiris features such as pre-provisioning and blade server rip-and-replace can further extend the time-savings enabled by Altiris Deployment Solution. By automating the server build process, Altiris Deployment Solution streamlines the installation and configuration of Dell servers—significantly enhancing the efficiency of routine systems management tasks and freeing administrators to support business-critical applications.

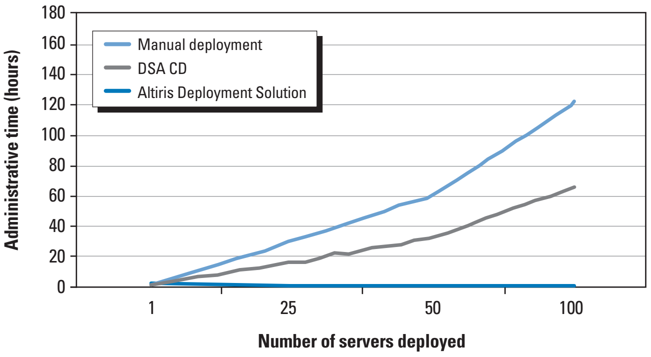

Figure 3. Comparing administrative time for three types of server deployment using Server Build 2 (Linux environment)

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.

Landon Hale manages Dell's relationship with Altiris within the Dell Global Alliances team. Previously, he worked in various product marketing, sales, and sales management roles at Dell and at Sea-Land Service. Landon has a B.A. in Political Science from Carleton College and an M.B.A. from the Marshall School of Business at the University of Southern California.

FOR MORE INFORMATION

Deployment Comparison for Dell PowerEdge Servers:
www.dell.com/downloads/global/solutions/Deployment%20Comparison%20for%20Dell%20PowerEdge%20Servers.pdf

LANDesk Server Manager:

Enhancing Standards-Based Manageability for Dell Servers

Server management technologies built into the hardware level, such as Intelligent Platform Management Interface 1.5–compliant baseboard management controllers and the Dell Remote Access Controller 4 (DRAC 4), can be enhanced by systems management software such as LANDesk® Server Manager. LANDesk Server Manager helps provide extended access and control by enabling administrators to access various built-in management modules on multiple servers through a single, unified console. This approach is designed to enable optimal server performance in high-availability computing environments.

BY ROGER FOREMAN, LANDON HALE, AND KIMBER H. BARTON

Related Categories:

Intelligent Platform Management Interface (IPMI)

LANDesk

Remote access controllers (RACs)

Remote management

Systems management

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

Dell servers offer advanced, built-in manageability at the hardware level. Eighth-generation platforms such as the Dell™ PowerEdge™ 1850 and PowerEdge 2850 servers feature Intelligent Platform Management Interface (IPMI) 1.5–compliant baseboard management controllers (BMCs) to enable standards-based manageability and out-of-band recoverability.

With the recent release of the IPMI 2.0 specification, future Dell servers are expected to take advantage of extended authentication, virtual LAN (VLAN) support, enhanced security, and more. The Dell Remote Access Controller 4 (DRAC 4) currently provides many such advanced capabilities for both IPMI-enabled and non-IPMI-enabled servers. When combined with advanced server management software such as LANDesk Server Manager (see Figure 1), advanced hardware management capabilities can enhance the manageability of a wide range of components—including IPMI-compliant BMCs, DRAC 4 devices, and standards-based host bus adapters (HBAs) such as SMBus Controller—on multiple servers through a single, unified console. Using an approach such as the unified LANDesk console, administrators can be enabled to access advanced features of standards-based hardware management interfaces regardless of the underlying hardware component.

LANDesk Server Manager features and capabilities

LANDesk Server Manager provides advanced manageability for Dell PowerEdge servers through a unified management environment. Key features and capabilities include the following:

- Remote recoverability with support for both IPMI 1.5– and IPMI 2.0–compliant hardware as well as support for both the DRAC III and DRAC 4
- Extended device discovery, including both in-band and out-of-band IPMI discovery
- OS vulnerability detection and patch management

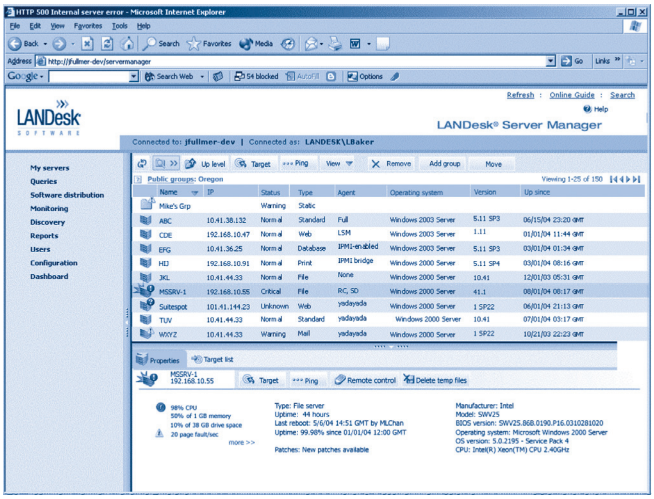

HTTP 500 Internal server error - Microsoft Internet Explorer

Address: http://jfullmer-dev/servermanager

Connected to: jfullmer-dev | Connected as: LANDESK\LBaker

Public groups: Oregon

Viewing 1-25 of 150

Name	IP	Status	Type	Agent	Operating system	Version	Up since
Mike's Grp		Warning	Static				
ABC	10.41.38.132	Normal	Standard	Full	Windows 2003 Server	5.11 SP3	06/15/04 23:20 GMT
CDE	192.168.10.47	Normal	Web	LSM	Windows 2003 Server	1.11	01/01/04 11:44 GMT
EFG	10.41.36.25	Normal	Database	IPMI-enabled	Windows 2003 Server	5.11 SP3	03/01/04 01:34 GMT
HIJ	192.168.10.91	Normal	Print	IPMI bridge	Windows 2000 Server	5.11 SP4	03/01/04 08:16 GMT
JKL	10.41.44.33	Normal	File	None	Windows 2000 Server	10.41	12/01/03 05:31 GMT
MSSRV-1	192.168.10.55	Critical	File	RC, SD	Windows 2000 Server	41.1	08/01/04 08:17 GMT
Suitespot	101.41.144.23	Unknown	Web	yadayada	Windows 2000 Server	1 SP22	06/01/04 21:13 GMT
TUV	10.41.44.33	Normal	Standard	yadayada	Windows 2000 Server	10.41	07/01/04 03:17 GMT
WXYZ	10.41.44.33	Warning	Mail	yadayada	Windows 2000 Server	1 SP22	10/21/03 22:23 GMT

MSSRV-1
192.168.10.55

98% CPU
50% of 1 GB memory
10% of 38 GB drive space
20 page fault/sec

Type: File server
Uptime: 44 hours
Last reboot: 5/6/04 14:51 GMT by MLChan
Uptime: 99.98% since 01/01/04 12:00 GMT

Patches: New patches available

Manufacturer: Intel
Model: SWV25
BIOS version: SWV25.86B.0190.P16.0310281020
Operating system: Microsoft Windows 2000 Server
OS version: 5.0.2195 - Service Pack 4
CPU: Intel(R) Xeon(TM) CPU 2.40GHz

Figure 1. LANDesk Server Manager GUI

- Recoverability of crashed IPMI-enabled servers over the network using remote power-control and problem-resolution tools
- Robust hardware and software inventory
- In-band remote control using low-impact, on-demand technology as well as out-of-band remote access and console redirection through IPMI or DRACs
- Predictive failure analysis and repair through active performance monitoring, alerting, and automated response
- Historical performance logging and extended event logging, including access to out-of-band system event logs (SELs)—providing administrators with extended out-of-band root-cause analysis capabilities for component failures in nonresponsive servers
- Easy scalability designed to manage thousands of servers through a single management server
- A single, unified console to access server management features and functions

Standards-based management through IPMI

The IPMI specification is designed to address the need for a consolidated remote hardware management standard. IPMI defines a common platform instrumentation interface that helps enable interoperability between the baseboard and the chassis, between the baseboard and the server management software, and between different physical servers. IPMI uses intelligent hardware that allows administrators to monitor and access platform instrumentation even when a server is powered down or when the OS is nonresponsive.

Standardized interfaces help provide an extensible management framework that is designed to adapt automatically as system components are added or removed (see Figure 2). At the hardware interface level, data abstracted from embedded sensors and control elements can travel either in-band through the host or out-of-band through a hardware management controller. In an effective management system, the specific interface should be irrelevant—normalized data is simply passed to the next layer. In this way, a comprehensive server management application can mask the underlying communications through a user interface that presents data and control elements in a consistent manner regardless of the specific hardware interface. Each layer adds unique value that contributes to the overall effectiveness of the system.

By specifying automatic alerting, system shutdown and restart, power control, and asset tracking capabilities, the IPMI standard enables IT administrators to monitor the health of their servers regardless of which OS the server uses—and regardless of the state of the OS. In an out-of-band state, for example, IPMI is designed to provide autonomous management data even when a server is powered down.

Centralized control

Many Dell servers are equipped with more than 100 on-board sensor devices and chips. Using LANDesk Server Manager, administrators can view data from these sensors, receive sensor alerts, and configure the BMC through the LANDesk console.

LANDesk Server Manager implements the IPMI standard through the following features and capabilities:

- **Discovery:** LANDesk Server Manager discovers IPMI-enabled systems both in-band and out-of-band. The software's Node Discovery Engine (NDE) takes a range of addresses and sends an Alerting Standards Forum (ASF) ping to the nodes on the address block. IPMI-enabled systems will respond with an ASF response. For each system that responds, LANDesk Server Manager establishes an out-of-band session and scans the Sensor Data Records (SDRs) for a valid globally unique identifier (GUID). This GUID is a unique ID that was created during installation, when the BMC was configured. The GUID is then sent back to the NDE, which enters it into the LANDesk Server Manager database.

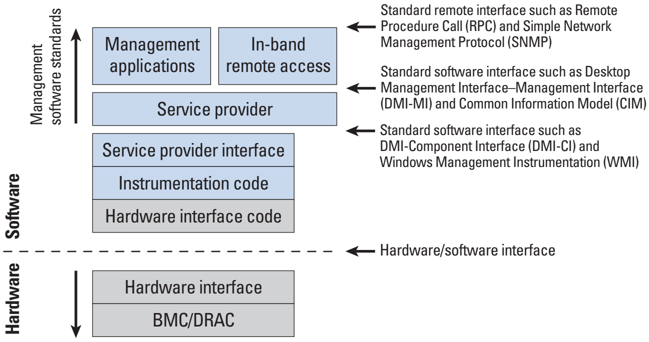
Management software standards
Software
Hardware
Management applications
In-band remote access
Service provider
Service provider interface
Instrumentation code
Hardware interface code
Hardware interface
BMC/DRAC
Standard remote interface such as Remote Procedure Call (RPC) and Simple Network Management Protocol (SNMP)
Standard software interface such as Desktop Management Interface–Management Interface (DMI-MI) and Common Information Model (CIM)
Standard software interface such as DMI-Component Interface (DMI-CI) and Windows Management Instrumentation (WMI)
Hardware/software interface

Figure 2. Layered hardware and software management stack

- **SEL event creation:** LANDesk Server Manager provides localized descriptions for events defined in the IPMI specification. The IPMI specification provides several ways for administrators to add their own SEL entries. Using the LANDesk Server Manager software developer's kit (SDK), administrators can add their own localized descriptions for customized IPMI events and sensors.
- **Caching:** Instead of accessing information from the BMC each time an administrator submits a request to LANDesk Server Manager, the software caches several types of information—streamlining information retrieval to help improve performance.
- **Watchdog timer:** LANDesk Server Manager provides support for the watchdog timer, a service present on IPMI-enabled systems. This service is configured during installation. The watchdog starts a timer on the BMC, then resets it after a specified amount of time. If the system hangs, the watchdog will time out, causing the hardware to reboot the system. In this way, the watchdog timer enables remote reboot.

LANDesk Server Manager supports both IPMI 1.5 and IPMI 2.0 specifications; IPMI 2.0 is backward compatible with IPMI 1.5. IPMI 2.0 includes integrated Serial Over LAN (SOL) management; low-cost BMCs (sometimes referred to as mini-BMCs); Remote Management Control Protocol+ (RMCP+) packet formats; enhanced security through authentication and encryption algorithms; LAN session enhancements such as discovery and support for new payload types under IPMI; modular extensions such as field replaceable unit (FRU) identification, node replacement, and redundant management bus monitoring; serial redirection over LAN as well as extended terminal mode; and continued support for original equipment manufacturer (OEM) value-added feature integration.

The IPMI specification is designed to address the need for a consolidated remote hardware management standard.

During installation, LANDesk Server Manager enables and configures the BMC out-of-band channels for communication. After installation, the LANDesk software detects the current IP configuration of the OS and hardware and synchronizes the BMC out-of-band channels with that configuration.

LANDesk Server Manager then continuously monitors the hardware and OS for changes in the IP configuration and resynchronizes the BMC as necessary. When a server that has LANDesk Server Manager installed is discovered and managed by the administrative console, the administrator may configure the BMC alert policies and Platform Event Filters (PEFs) for the desired alerting.

A layered management approach

Manageability features built into the server hardware layer—including IPMI-based components and the DRAC 4—can be enhanced when used in conjunction with systems management software running under the OS, such as LANDesk Server Manager. Together, such systems management software and the OS can help provide sophisticated control, error handling, and alerting.

Through a single console, LANDesk Server Manager enables integrated systems management for Dell PowerEdge servers using an IPMI-compliant BMC or mini-BMC, DRAC, HBA, or a combination of these hardware interfaces. This approach enables administrators to combine extended in-band management with robust out-of-band management to take advantage of advanced hardware management technologies—helping to provide efficient, effective systems management for high-availability servers.

Roger Foreman is a senior product manager on the Dell OpenManage Marketing team. He previously served as a senior consultant in the Dell Technology Showcase and a senior manager of e-infrastructure services for Dell Technology Consulting. Before that, he spent 19 years in sales, software development, and consulting at IBM. Roger has a B.S. and an M.S. in Electrical Engineering from Iowa State University, as well as an M.S. in Computer Science from the University of Arizona.

Landon Hale manages Dell's systems management alliance ecosystem on the Dell Global Alliances team. Previously, he worked in various product marketing, sales, and sales management roles at Dell and at Sea-Land Service. Landon has a B.A. in Political Science from Carleton College and an M.B.A. from the Marshall School of Business at the University of Southern California.

Kimber H. Barton is a technical director in the Strategic Alliance group at LANDesk Software Inc. Over the past 20 years, he has held a wide range of industry positions including technical marketing engineer, product solution specialist, and systems engineer for several leading hardware and software technology companies such as Intel and LANDesk. Kimber has an A.S. in Electrical Automation and Robotics Technology from Utah Valley State College and a B.S. in Business Administration from the University of Phoenix, as well as many key industry-standard certifications.

FOR MORE INFORMATION

LANDesk Server Manager:
www.landesk.com/Products/Server

Dell PowerEdge servers:
www.dell.com/servers

Quick and Easy Hardware Diagnostics

for Dell PowerEdge and PowerEdge SC Servers

Dell™ PowerEdge™ Diagnostics is a lightweight, stand-alone program designed to perform comprehensive, cross-platform troubleshooting. This article describes how administrators can install and use the Dell PowerEdge Diagnostics application to enhance server administration.

BY HARISH JAYAKUMAR AND KARTHIK RAJAGOPALAN

Related Categories:

Dell PowerEdge SC servers

Dell PowerEdge servers

Diagnostics

Linux

Microsoft Windows

Systems management

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

Troubleshooting hardware to determine the cause of failures and problem conditions is critical for ensuring proper server maintenance and enabling a high level of availability. An efficient, effective hardware diagnostic tool can help organizations isolate faults and resolve problems quickly, minimizing system downtime considerably. Dell provides two approaches to hardware diagnostics for Microsoft® Windows® OS–based and Linux® OS–based server environments: the Dell OpenManage™ suite and the Dell PowerEdge Diagnostics tool.

The Dell OpenManage suite is a comprehensive software suite that includes an integrated set of rich systems management features, of which diagnostics is one component. Dell OpenManage, included with Dell PowerEdge servers (but not with PowerEdge SC servers), is installed on a system if the administrator selects the Diagnostic Service option when setting up Dell OpenManage software on the server.

Alternatively, Dell PowerEdge Diagnostics software provides the capability to run diagnostics on PowerEdge and PowerEdge SC servers. Thus, it is well suited for organizations that own PowerEdge SC servers or simply do not want to install the entire Dell OpenManage suite. Dell PowerEdge Diagnostics is a small (approximately 5 MB) application that comprises a suite of diagnostic programs, or *test modules,* that run locally on the target system. After running the test modules, administrators can easily delete the Dell PowerEdge Diagnostics application from the system if desired. When organizations need diagnostic capabilities but do not require a complete set of management tools or prefer not to permanently install a diagnostic application on their servers, Dell PowerEdge Diagnostics can be an excellent choice.

Downloading Dell PowerEdge Diagnostics

Administrators can obtain the Dell PowerEdge Diagnostics application from the Dell support Web site or from a Dell CD. For PowerEdge servers, the application is located on the Dell PowerEdge Service and Diagnostic Utilities CD. For PowerEdge SC servers, the Dell PowerEdge Diagnostics application is located on the Dell OpenManage Server Assistant CD (in Service Mode). Administrators can download the application from either CD as follows:

1. Select the appropriate eighth-generation PowerEdge or PowerEdge SC server model (Dell PowerEdge Diagnostics does not support earlier-generation servers) from the Select Server menu on the Welcome screen, and then select the server's OS from the Select Drivers/Utilities menu. Only Windows and Linux are supported.
2. Click "Continue." A screen showing details of the server configuration will appear.
3. Click "PowerEdge Diagnostics" to download an executable file containing the utility.

Administrators can download the application from the Dell support Web site as follows:

1. Go to support.dell.com/support/downloads, select “PowerEdge” from the Product Model menu, and then select the appropriate server.
2. Click the blue arrow next to the server. The Select Criteria screen will appear.
3. Select “Diagnostic Utilities” from the Select Your Download Category menu.
4. Select the server’s OS from the Select an Operating System menu.
5. Select the server’s OS language from the Select an Operating System Language menu.
6. Click “Submit” and select “Diagnostic Utilities” on the next screen.
7. Click “Dell PowerEdge Diagnostics” to download an executable file containing the utility.

Installing Dell PowerEdge Diagnostics

Once administrators have obtained the Dell PowerEdge Diagnostics package, they can install the application on either Windows- or Linux-based servers.

Installation on Windows-based servers. On servers running a Windows OS, administrators should double-click the executable file that they downloaded (see the preceding section in this article, “Downloading Dell PowerEdge Diagnostics”) and then follow the steps in the Dell PowerEdge Diagnostics installation wizard. By default, the wizard extracts the files required to run the application into a folder named “Dell” on the desktop, and then launches the application. The StartDiags shortcut within the Dell directory can be used to launch the application later. After the diagnostics are completed, administrators may delete the Dell folder. *Note:* A user must be logged in with Administrator access to run Dell PowerEdge Diagnostics.

Installation on Linux-based servers. On servers running a Linux OS, the executable file for the application is available for download as a .tar.gz file. This file must first be moved to the desired directory by the root user before the files within it are extracted; extracting the files creates a subdirectory within that directory named “Dell.” The startDiag.sh command-line script, which is one of the files extracted from the .tar.gz file, can be used to launch the application under the X Window System, a nonproprietary standardized set of display-handling routines.[1]

In the following example, a directory named DellPEDiagnostics is created in /root. After extraction of the .tar.gz file, a folder called “Dell” appears. The startDiag.sh script resides inside the Dell folder; administrators can invoke this script to launch the application as follows:

```
[root@server5 DellPEDiagnostics]#  ls
Dell
[root@server5 DellPEDiagnostics]#  cd  Dell
[root@server5 Dell]#  ls
oldiags readme.txt  startDiags.sh
[root@server5 Dell]#  ./startDiags.sh
```

Administrators do not need to reboot the server to run Dell PowerEdge Diagnostics.

Understanding the program interface

After it is launched, the Dell PowerEdge Diagnostics application is designed to automatically discover the devices in the system. *Note:* The application must be restarted to discover hardware changes implemented after the device discovery process.

When the device discovery process is complete, the application presents Diagnostic Information and Diagnostic Selection panes (see Figure 1) for administrators to view as follows:

- **Diagnostic Information:** Displays system configuration and operational status
- **Diagnostic Selection:** Lists available tests and the devices on which these tests can be run

The Diagnostic Selection pane offers three viewing options: Group by Connection, Group by Device, and Group by Test.

Group by Connection. The Group by Connection view organizes the devices based on how they are connected to each other.

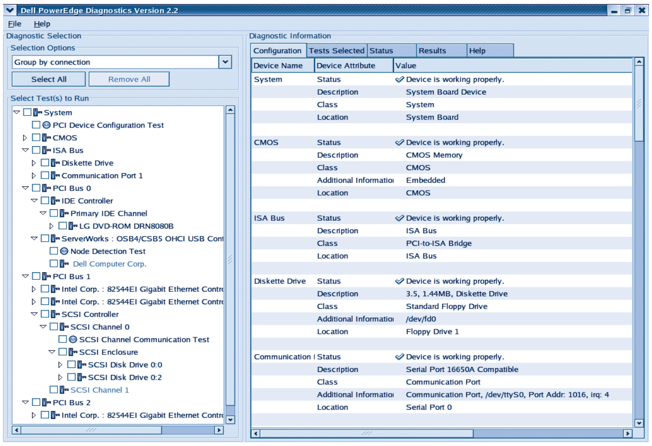

Group by connection

□ System
□ PCI Device Configuration Test
□ CMOS
□ ISA Bus
□ Diskette Drive
□ Communication Port 1
□ PCI Bus 0
□ IDE Controller
□ Primary IDE Channel
□ LG DVD-ROM DRN8080B
□ ServerWorks : OSB4/CSB5 OHCI USB Cont
□ Node Detection Test
□ Dell Computer Corp.
□ PCI Bus 1
□ Intel Corp. : 82544EI Gigabit Ethernet Contr
□ Intel Corp. : 82544EI Gigabit Ethernet Contr
□ SCSI Controller
□ SCSI Channel 0
□ SCSI Channel Communication Test
□ SCSI Enclosure
□ SCSI Disk Drive 0:0
□ SCSI Disk Drive 0:2
□ SCSI Channel 1
□ PCI Bus 2
□ Intel Corp. : 82544EI Gigabit Ethernet Contr

Device Name	Device Attribute	Value
System	Status	✓ Device is working properly.
	Description	System Board Device
	Class	System
	Location	System Board
CMOS	Status	✓ Device is working properly.
	Description	CMOS Memory
	Class	CMOS
	Additional Informatio	Embedded
	Location	CMOS
ISA Bus	Status	✓ Device is working properly.
	Description	ISA Bus
	Class	PCI-to-ISA Bridge
	Location	ISA Bus
Diskette Drive	Status	✓ Device is working properly.
	Description	3.5, 1.44MB, Diskette Drive
	Class	Standard Floppy Drive
	Additional Informatio	/dev/fd0
	Location	Floppy Drive 1
Communication I	Status	✓ Device is working properly.
	Description	Serial Port 16650A Compatible
	Class	Communication Port
	Additional Informatio	Communication Port, /dev/ttyS0, Port Addr: 1016, irq: 4
	Location	Serial Port 0

Figure 1. Dell PowerEdge Diagnostics interface

[1] For high-quality graphics, administrators can run the Dell PowerEdge Diagnostics application under GNU Network Object Model Environment (GNOME), a Windows-like desktop system that runs on Linux systems. For more information about the open source GNOME utility, visit www.yolinux.com/TUTORIALS/GNOME.html.

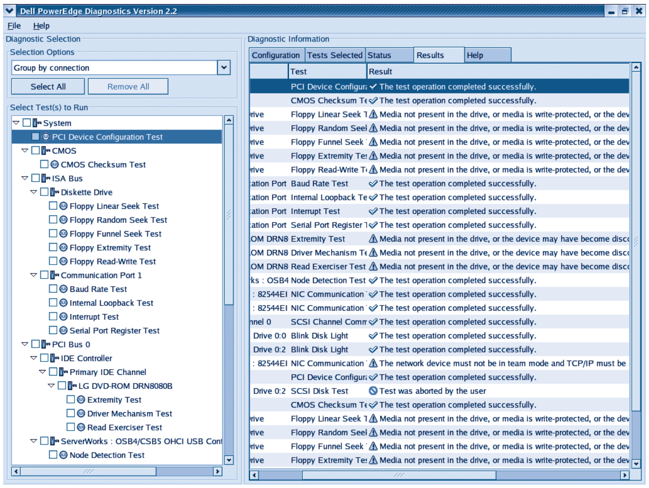

Group by connection

- ▽ □ System
  - □ PCI Device Configuration Test
  - ▽ □ CMOS
    - □ CMOS Checksum Test
  - ▽ □ ISA Bus
    - ▽ □ Diskette Drive
      - □ Floppy Linear Seek Test
      - □ Floppy Random Seek Test
      - □ Floppy Funnel Seek Test
      - □ Floppy Extremity Test
      - □ Floppy Read-Write Test
    - ▽ □ Communication Port 1
      - □ Baud Rate Test
      - □ Internal Loopback Test
      - □ Interrupt Test
      - □ Serial Port Register Test
  - ▽ □ PCI Bus 0
    - ▽ □ IDE Controller
      - ▽ □ Primary IDE Channel
        - ▽ □ LG DVD-ROM DRN8080B
          - □ Extremity Test
          - □ Driver Mechanism Test
          - □ Read Exerciser Test
    - ▽ □ ServerWorks : OSB4/CSB5 OHCI USB Cont
      - □ Node Detection Test

	Test	Result
	PCI Device Configur	✓ The test operation completed successfully.
	CMOS Checksum Te	✓ The test operation completed successfully.
rive	Floppy Linear Seek T	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Random Seel	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Funnel Seek	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Extremity Tes	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Read-Write T	⚠ Media not present in the drive, or media is write-protected, or the dev
ation Port	Baud Rate Test	✓ The test operation completed successfully.
ation Port	Internal Loopback Te	✓ The test operation completed successfully.
ation Port	Interrupt Test	✓ The test operation completed successfully.
ation Port	Serial Port Register T	✓ The test operation completed successfully.
OM DRN8	Extremity Test	⚠ Media not present in the drive, or the device may have become disc
OM DRN8	Driver Mechanism Te	⚠ Media not present in the drive, or the device may have become disc
OM DRN8	Read Exerciser Test	⚠ Media not present in the drive, or the device may have become disc
ks : OSB4	Node Detection Test	✓ The test operation completed successfully.
: 82544EI	NIC Communication	✓ The test operation completed successfully.
: 82544EI	NIC Communication	✓ The test operation completed successfully.
nnel 0	SCSI Channel Comm	✓ The test operation completed successfully.
Drive 0:0	Blink Disk Light	✓ The test operation completed successfully.
Drive 0:2	Blink Disk Light	✓ The test operation completed successfully.
: 82544EI	NIC Communication	⚠ The network device must not be in team mode and TCP/IP must be
	PCI Device Configur	✓ The test operation completed successfully.
Drive 0:2	SCSI Disk Test	⦸ Test was aborted by the user
	CMOS Checksum Te	✓ The test operation completed successfully.
rive	Floppy Linear Seek T	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Random Seel	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Funnel Seek	⚠ Media not present in the drive, or media is write-protected, or the dev
rive	Floppy Extremity Tes	⚠ Media not present in the drive, or media is write-protected, or the dev

Figure 2. Results tab showing completed tests, aborted tests, and tests that encountered non-critical errors

This view is comparable to the Devices by Connection viewing option in Device Manager, which is a function of the System Properties utility in Windows operating systems. However, the devices that Dell PowerEdge Diagnostics displays are a subset of the devices shown in Device Manager. Devices that the Dell PowerEdge Diagnostics application does not support are not shown.[2]

Administrators may want to use the Group by Connection view when they are familiar with the available tests and want to run specific tests on specific devices.

Group by Device. The Group by Device view organizes the devices as a set of flat nodes—that is, the devices are listed alphabetically, with no reference to how they are connected. The tests that can be run on a device are shown when the administrator expands the node corresponding to that device. No tests are available for devices that are grayed out.

Administrators may want to use the Group by Device view when they need to run all available tests on a specific device (for instance, to troubleshoot a device that has been functioning poorly).

Group by Test. The Group by Test view organizes all available tests as a set of flat nodes. The devices that can be tested by a given test are shown when the administrator expands the node corresponding to that test.

Administrators may want to use the Group by Test view when they are familiar with the available tests and want to run a specific test on all the devices for which that test is available.

Running a quick health check on a system

Once administrators are familiar with the program interface, they can use the Dell PowerEdge Diagnostics application to run quick diagnostic tests to discover problem conditions using the following procedure:

1. In the program interface, click "Select All" in the Diagnostic Selection pane to select all available tests. The Diagnostic Information pane will then automatically shift to the Tests Selected tab.
2. On the Tests Selected tab, select the Quick Test box, and click "Run Tests." *Note:* On systems that have high-capacity hard disks, the diagnostics may take a while to complete even though the Quick Test diagnostic option is selected. However, the tests would take even longer without the Quick Test option selected.
3. Wait while the Dell PowerEdge Diagnostics application adds the tests to a queue. This too may take some time, depending on the number of devices in the system. Administrators can view the progress of the tests under the Status tab.
4. While tests are running, administrators can view results as they are generated under the Results tab of the Diagnostic Information pane (see Figure 2). Administrators can also suspend or abort tests while they are running. *Note:* Not all tests support suspend and abort functionality. For a list of tests that do not support suspend and abort features, administrators should refer to the *Dell PowerEdge Diagnostics Version 2.2 User's Guide* (support.dell.com/support/edocs/software/smpediag/2.2/index.htm).

The results shown in the Results tab are represented by icons followed by detailed descriptions of the test results. Based on results obtained using the Quick Test option, administrators can choose to de-select Quick Test and run one or more tests to check exhaustively for problem conditions in a specific device or devices.

Finding a convenient, easy approach to server diagnostics

For organizations that want to add quick, lightweight diagnostics to their server maintenance toolkit, Dell PowerEdge Diagnostics can be a cost-effective, easy-to-use application designed with the flexibility to support both Windows- and Linux-based servers. Even for systems not initially configured with diagnostic software, this application can be easily deployed—without complex installations, rebooting, or extended downtime—to enhance systems management functions.

Harish Jayakumar is a test engineer in the Dell OpenManage software development and test organization. He has an M.S. in Computer Science from Arizona State University and a B.S. in Computer Science from the University of Madras in India.

Karthik Rajagopalan is a developer in the Dell OpenManage software development organization. He has a master's degree in Electrical Engineering from the Indian Institute of Technology, Madras, and a bachelor's degree in Instrumentation from the Birla Institute of Technology and Science in Pilani, India.

[2] For a comprehensive list of devices supported by Dell PowerEdge Diagnostics, see the *Dell PowerEdge Diagnostics Version 2.2 User's Guide* at support.dell.com/support/edocs/software/smpediag/2.2/index.htm.

Using the DRAC 4 and Dell OpenManage DTK in

Remote Deployments Without a PXE Server

The Dell™ OpenManage™ Deployment Toolkit (DTK) allows administrators to create or enhance a framework for rapidly deploying system images on Dell servers. A deployment environment can be as simple as a bootable CD or as complex as a remote network boot environment. This article explains how to extend a network-based deployment framework using the Dell Remote Access Controller 4 (DRAC 4).

BY ALAN BRUMLEY AND ANUSHA RAGUNATHAN

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Dell Remote Access Controller (DRAC)

Remote management

System deployment

Systems management

Virtual media

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

Deployment frameworks require target servers to boot into a deployment OS, which is designed to enable administrators to control the installation of system images. Using current technology, remote deployment of systems without keyboard, video, and mouse (KVM) necessitates the use of Preboot Execution Environment (PXE). PXE is a convenient and efficient method for booting a server on a trusted and secure network. However, PXE has three potential drawbacks that may lead administrators to opt for a different approach to deployment: security, the use of User Datagram Protocol (UDP), and the need for a Dynamic Host Configuration Protocol (DHCP) server.

- **Security:** A PXE server processes DHCP requests and passes along a file name and file server address to the target server when assigning an IP address to the target server during the PXE boot process. The target server then uses Trivial FTP (TFTP) to connect to the PXE server, downloads a bootstrap loader file from the PXE server, and executes the file. However, security concerns may arise from unauthorized PXE servers and clients because of the open and "trusting" nature of the PXE boot process.
- **UDP:** Another potential drawback of PXE is its heavy use of UDP. The UDP transport does not readily provide a mechanism for detecting and resending lost packets, which can lead to disruptions in data transmission.
- **DHCP:** Because PXE requires a DHCP server, some common network topologies can create situations in which PXE may not be not a good fit. For example, a remote branch office that is equipped with only one or two servers may be too small to cost-justify its own PXE server. If the remote office is not connected by a virtual private network (VPN) to the corporate office—or if the VPN is not configured to relay DHCP requests—PXE booting may not be possible without a local server.

The primary advantage of PXE is its capability to boot multiple servers from the network into a deployment OS without requiring a specific configuration of the PXE server for each client system. However, in scenarios such as that of a small branch office network where the pool of servers is small and redeployments are infrequent, this advantage may not be significant. In such cases, the security, UDP, and DHCP drawbacks inherent in PXE may outweigh the advantage of the hands-off deployments that PXE enables. Alternatively, administrators may consider using the Dell Remote Access Controller 4 (DRAC 4) as a boot device by leveraging its virtualization features.

Method	Advantages	Disadvantages
PXE	• Suited for one-to-many booting • No PXE server reconfiguration required for each target server • Central repository for all boot images	• Potential security issues • UDP transport used • DHCP server required
DRAC 4	• TCP-based connection • Authentication on the DRAC 4 helps prevent unauthorized booting • No additional DHCP/PXE servers needed for network booting	• Interaction with target server required • Additional network connection required

Figure 1. Advantages and disadvantages of PXE versus the DRAC 4 for remote booting and remote deployment

Advantages of booting using the DRAC 4 versus PXE

In eighth-generation Dell PowerEdge™ and PowerEdge SC systems and in planned future Dell servers, the DRAC 4 is designed to provide dependable boot functionality that allows administrators to reuse most of their DOS-booting PXE deployment environment. Virtualization features built into the DRAC 4 address each of the three drawbacks of PXE (see Figure 1).

First, management connections to the DRAC 4 can be created over a secure Web interface, helping address security concerns about booting servers located at remote sites. Furthermore, a username and password are required to enable virtualization and to connect physical media—such as floppy disk drives—on the target system. Such precautions can make it difficult for an unauthorized user to remotely access boot media. In contrast, systems booting with PXE cannot determine whether the server communicating with the client is the “official” PXE server. The first server that responds to a client will boot that client.

Second, the connection from the DRAC 4 to the management station that contains the boot media is maintained over a TCP connection. In contrast to a UDP connection, a TCP connection makes it difficult for network conditions to cause a disruption because handshaking on both ends helps ensure correct and dependable delivery of the boot image.

Third, the DRAC 4 does not require a DHCP server. If an organization’s existing network does not have a DHCP server, the DRAC 4 can be configured to use a fixed IP address. Furthermore, to enable booting, the DRAC 4 need not even reside on the same network as the corporate DHCP server. *Note:* Once the system is booted and network drivers are loaded, the target server’s network interface cards (NICs)—not the DRAC 4’s NIC—are used to connect to network drives and resources to continue deployment.

Limitations of DRAC 4 virtualization in DOS

When using DRAC 4 virtualization features from within DOS, administrators should be aware of two caveats. First, DRAC 4 virtualization features offer a choice of boot devices: either a floppy disk or a CD. Unfortunately, administrators cannot access both devices from DOS at the same time; only the device that initiates the boot is available to DOS. *Note:* This limitation does not necessarily apply to systems that use a Microsoft® Windows® or Linux® OS with the appropriate virtualization drivers installed.

Second, administrators may experience problems when booting a generic floppy disk, because any CD drivers that load will attempt to use the IDE CD that may be installed on the target server. Most DOS CDs begin booting by using a process known as boot floppy emulation, in which the floppy disk image stored on the CD loads CD drivers, mounts the rest of the data on the CD, and then continues deployment. If a DOS CD that uses boot floppy emulation is used in conjunction with DRAC 4 virtualization, the target server will typically respond in one of the following ways:

- Boot as normal but display an error message stating that the IDE CD is not present
- Halt and display an error message that the CD drive is empty (that is, assuming no CD in the IDE CD drive on the target server)

To avoid these issues, Dell recommends using boot floppy disk virtualization only when remotely booting the target server into the deployment framework. *Note:* This limitation applies only to DOS. For instructions on remote booting and installing other operating systems, such as Windows and Linux, please consult the DRAC 4 documentation.

Provided that a boot floppy disk does not attempt to load CD drivers, few if any changes are typically necessary to make the floppy disk work properly in the remote system. Consequently, the PXE boot floppy disk is a natural candidate to be used as the DRAC 4 virtual floppy disk because most PXE boot floppy disks do not detect and configure CDs on the target server; instead, they focus on getting a network stack loaded and connected.

DRAC 4 virtual media capabilities for deployment

The virtual media capabilities of the DRAC 4 are designed to provide organizations with the capability to access remote media such as floppy disks or CDs as though they were physically present on the system. The virtual media feature can be used in conjunction with the Dell OpenManage Deployment Toolkit (DTK) to help overcome hurdles in deployment frameworks.[1]

The DRAC 4 is primarily advantageous as a deployment tool to manage remote servers on a dedicated network in which no PXE server is available. In the Figure 2 example, a corporate PXE server manages the management station, and the management station is connected to the remote target server over a dedicated IP-based VPN. When the target server is not on the same network as the PXE server, deployment can be a challenge. A strategy that uses the virtual media capabilities of the DRAC 4 can avoid the need for a PXE server on the target server’s network by booting the target server to an environment running DTK on DOS.

[1] DTK is not supported on PowerEdge SC servers, so any remote deployment that is set up on these servers will not include DTK components.

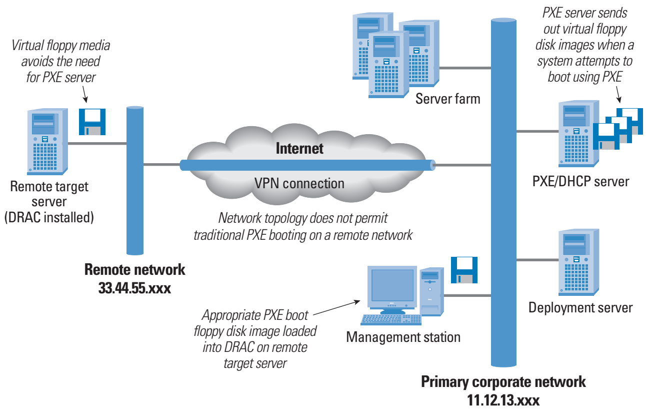
Virtual floppy media avoids the need for PXE server
Remote target server (DRAC installed)
Remote network 33.44.55.xxx
Internet
VPN connection
Network topology does not permit traditional PXE booting on a remote network
Server farm
PXE server sends out virtual floppy disk images when a system attempts to boot using PXE
PXE/DHCP server
Deployment server
Appropriate PXE boot floppy disk image loaded into DRAC on remote target server
Management station
Primary corporate network 11.12.13.xxx

Figure 2. Booting a remote server to an environment running DTK on DOS

Network stack considerations

When using the DRAC 4's virtual media capabilities, administrators must take into account network stack considerations. When using PXE to remotely boot systems into DOS, administrators often use a generic DOS floppy disk image that they capture and store on the PXE server. Whenever a server attempts to boot via PXE, the PXE server transmits this virtual floppy disk image to the target server, which boots as if this image were on a floppy disk inserted into the target server's floppy disk drive. Normally, a network connection is necessary and the virtual floppy disk image must contain drivers for the target system's NIC as well as a network stack.

Some network administrators may employ a generic driver that uses the Universal Network Driver Interface (UNDI) and PXE base code infrastructure. The generic driver allows use of the network stack without the need for a driver that is specific to a particular type of NIC.

When a server boots from the DRAC 4 virtual floppy disk, the PXE network stack may not be present (because PXE may be turned off) or the PXE network stack may not be in the state that the driver would expect because the system did not boot via PXE. Administrators who experience problems when booting from the DRAC 4 virtual floppy disk should update the boot floppy disk image to use DOS network drivers specific to the NIC instead of the generic UNDI or PXE base code drivers.

Finally, administrators should keep in mind that the NIC integrated into the DRAC 4 cannot be used for OS network traffic. This means that administrators need to use the embedded NICs on the target server or an add-in card to access the network from within the deployment OS. The DRAC 4 NIC is used only to communicate with the DRAC 4 and to perform management functions.

Preparation of the virtual image

Best practices recommend that administrators include network drivers for the target servers on the DOS image that is booted.[2] Network drivers allow the target server to access the file share that contains the DTK tools and scripts. *Note:* Administrators must have a valid license for DOS 6.22 (on Windows 95 and Windows 98 operating systems) before deploying servers in their network.

Administrators should edit autoexec.bat to enable network access to the file server share that contains the DTK tools and scripts before creating the network boot floppy disk. Once created, the image from the boot floppy disk can be used as a virtual floppy disk or a virtual CD.

Although the virtual image strategy avoids the need for PXE, deploying multiple servers using the virtual image method can be difficult and time-consuming because this process requires administrators to access each DRAC 4 one by one. To conveniently deploy to multiple servers, administrators can use the remote Racadm command-line utility from either Windows or Linux on the management station.

To use the Racadm utility, administrators should follow this procedure:

1. During the power-on self-test (POST) of the target server, use the DRAC 4 Ctrl-E boot utility to set the DRAC 4 IP address to either static or DHCP. *Note:* This step requires physical access to the target server.
2. Use the F2 BIOS setup utility to set the boot sequence on the target server so that the virtual floppy disk or CD is the first device in the boot sequence, followed by the primary hard disk (or IDE CD if desired). *Note:* Administrators can use DRAC 4 console redirection to accomplish this step if the target server is in a remote location.

   The target server's boot sequence should be left in this state during general use. Each time the server is rebooted, it will not boot to the virtual floppy disk because, without a connected management workstation, the drive will appear empty to the BIOS.
3. When the system needs to be deployed or redeployed, insert the boot floppy disk image into the floppy disk drive of the management station and use a Web browser on the management station to connect to the target server's DRAC 4.
4. Use the DRAC 4 power control feature or traditional OS methods to reboot the target server. The target server will boot and start the deployment using DTK scripts.
5. At the appropriate stage of a given DTK deployment script, the Racadm utility can be used to "eject" the virtual floppy disk by sending a `vmdisconnect` command to the DRAC 4, causing the DRAC 4 to sever connections to the management station that sent the DRAC 4 the boot disk. After this, the target server will not see a bootable virtual floppy disk image and instead will fail over to the hard disk.

[2] For information about how to create a DOS floppy disk that includes network drivers, visit www.nu2.nu/bootdisk/network.

Note: Part of the deployment process includes adding the virtualization drivers to the deployment image so that the OS can access the virtual media after the OS is installed. Please consult the OS user's guide for information about scripted driver installations.

Powerful out-of-band management

The DRAC 4 is a powerful management device that can help reduce cost of ownership. In addition to the convenience of remote out-of-band access to target servers, virtual media capabilities allow the DRAC 4 to be used as a remote boot device. In this way, the DRAC 4 can help IT organizations that support small, remote branch offices to save time and money by deploying servers without traveling to the target servers' physical locations—and without installing and supporting a PXE server at each remote location.

Alan Brumley is the lead engineer for the Dell OpenManage Deployment Toolkit. He has been at Dell for more than five years. Alan has a B.S. in Computer Engineering from the University of South Carolina.

Anusha Ragunathan is a software engineer in the Dell Product Group and is on a team that develops deployment tools for PowerEdge servers as part of Dell OpenManage systems management offerings. Anusha has a B.S. in Computer Science Engineering from Bharathiyar University in India and an M.S. in Computer Science Engineering from Arizona State University.

FOR MORE INFORMATION

Remote management using DRAC 4:
Pan, Weimin, and Gang Liu. "Remote Management with Virtual Media in the DRAC 4." *Dell Power Solutions,* October 2004. www.dell.com/downloads/global/power/ps4q04-20040105-Pan.pdf.

Remote OS deployment using the DRAC 4:
Brown, Michael E., Manoj Gujarathi, and Gong Wang. "Remote OS Deployment Using Dell OpenManage Server Assistant 8 and DRAC 4." *Dell Power Solutions,* February 2005. www.dell.com/downloads/global/power/ps1q05-20040170-Gujarathi.pdf.

Updating the Dell PowerEdge 1855 Blade Server Chassis and Server Blades

Many hardware components in the Dell™ PowerEdge™ 1855 blade server have been designed with a firmware upgrade option—including the Dell Remote Access Controller/Modular Chassis (DRAC/MC) management module; server baseboard management controllers; I/O modules; and keyboard, video, mouse modules. This article discusses standards-based frameworks for updating individual server blades and available software utilities as well as recommended methods and best practices for updating chassis component firmware.

BY NARAYAN DEVIREDDY AND RUOTING HUANG

Related Categories:

Blade servers

Dell PowerEdge blade servers

New-generation server technology

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

Equipped with 10 server blades and several I/O modules and management modules, the Dell PowerEdge 1855 blade server contains numerous system software components to update during the change-management process. Although best practices dictate careful planning and methodical implementation when deploying system software updates for the PowerEdge 1855 blade server, such change management need not be difficult. As is the case with other PowerEdge servers, administrators can manage software updates for the PowerEdge 1855 blade server using standards-based Dell OpenManage™ tools, third-party enterprise change-management tools, or both.

Updating server blades

The process for updating system software components on PowerEdge 1855 server blades is designed to be the same as on other PowerEdge servers. This similarity allows administrators to leverage existing change-management processes in their IT organization to manage updates to PowerEdge 1855 server blades. On PowerEdge 1855 server blades, administrators can typically update firmware on the following system software components:

- Server system BIOS
- Baseboard management controller (BMC)
- PowerEdge Expandable RAID Controller 4, Integrated Mirroring (PERC 4/IM)
- Adaptec SCSI Card 39160
- SCSI hard disk drives
- OS-level drivers

As part of the change-management planning process, administrators must decide which update method best fits their environment. System software components on a PowerEdge 1855 server blade can be updated using one of two methods: offline updates or online updates.

Offline updates

To perform an offline update, administrators must schedule server downtime. Most offline update utilities boot into a pre-OS environment such as DOS to perform the update. These update utilities may also require multiple system reboots to complete the update process. Although the offline update process is time-consuming and requires several manual steps, this is a safe method to update system software on virtually any server because the network is not involved. In contrast, network traffic, Trivial FTP (TFTP) server response time, packet loss, and link loss can affect firmware updates that are performed over a network.

Although the offline update process is time-consuming and requires several manual steps, this is a safe method to update system software on virtually any server because the network is not involved.

On a PowerEdge 1855 server blade, administrators can update BIOS, BMC firmware, RAID firmware, SCSI hard disk drive firmware, SCSI controller firmware, and OS-level drivers using DOS-based update utilities found on the Dell support Web site (support.dell.com). *Note:* The PERC 4/IM is integrated into the PowerEdge 1855 server blade motherboard, so PERC 4/IM firmware is included as part of the system BIOS update package.

Online updates

Unlike an offline update, the online update process minimizes server downtime by allowing the process to run in a normal OS environment. The greatest advantage of the online approach is its capability to automate the update process. Online updates enable organizations to update multiple server blades from a central console using Dell OpenManage change-management tools or third-party change-management frameworks.

Dell OpenManage change-management tools for the PowerEdge 1855 blade server comprise the following:

- Dell Update Packages
- Dell OpenManage Server Update Utility (SUU)
- Dell OpenManage IT Assistant (ITA) 7

Dell Update Packages. Dell Update Packages are self-contained, easy-to-use programs; each package updates a single system software component on the server on which it is executed. Each Dell Update Package contains the logic to verify that the update will work on a given system. Dell Update Packages support both a graphical user interface (GUI) and a command-line interface (CLI). BIOS and BMC firmware Dell Update Packages are available for PowerEdge 1855 server blades. Dell Update Packages are supported on Microsoft® Windows Server™ 2003 and Red Hat® Enterprise Linux® operating systems. In addition, Dell Update Packages can be integrated into third-party or custom software distribution application frameworks such as Microsoft Systems Management Server (SMS) 2003, Altiris® Deployment Solution, and others.[1]

Dell OpenManage Server Update Utility. SUU is a comprehensive system update utility that provides a mechanism to update several system software components at once. It is a CD-based application that can identify systems and apply the appropriate updates. For example, on a PowerEdge 1855 blade server, administrators can use the SUU CD to update all system software on a server blade. The SUU update process requires administrators to update one server blade at a time.

Dell OpenManage IT Assistant. ITA 7 provides centralized software update capability for Dell servers. IT Assistant allows administrators to load Dell Update Packages and Dell System Update Sets into a central repository and then compare the packages to the software versions currently running on PowerEdge systems. Administrators can then decide whether to update systems that are not in compliance, either immediately or according to an administrator-defined schedule.

BMC flash updates

The BMC does not disable the power button on the front of an individual PowerEdge 1855 server blade during a flash update. If an administrator presses the power button during a flash update, the server blade will power off and leave the BMC in an unknown state. To recover the server blade, an administrator should remove the server blade from the PowerEdge blade server chassis, wait five seconds, reinsert the server blade, and allow the server blade to boot. Then the flash update can be performed again.

Note: Keyboard, video, mouse (KVM) hot-key keyboard sequences are not supported during a BMC flash update. If an administrator attempts a hot-key sequence from a server blade that is performing a BMC flash update, the flash update may fail. To recover, the administrator must execute the flash update again.

Updating chassis components

Because individual server blades in the PowerEdge 1855 blade server share common chassis infrastructure components, keeping the system software of the chassis components up-to-date can play a critical role in

[1] For more information about how Dell Update Packages are designed to work in third-party software distribution applications, see "Scripting Dell Update Packages on Windows and Linux" by Manoj Gujarathi, Pritesh Prabhu, and Subbu Ganesan in *Dell Power Solutions,* October 2004, www.dell.com/downloads/global/power/ps4q04-20040125-Gujarathi.pdf; and "Deploying Dell Update Packages Using Microsoft Systems Management Server 2003" by Sandeep Karandikar and Manoj Gujarathi in *Dell Power Solutions,* February 2005, www.dell.com/downloads/global/power/ps1q05-20040111-Gujarathi.pdf.

the overall PowerEdge 1855 blade server change-management process. However, the frameworks that are used to update individual server blades do not scale to manage the updates of chassis components. This section describes how to update components in the PowerEdge 1855 blade server chassis that have administrator-upgradeable system software. These firmware components include:

- Management modules
- KVM modules
- I/O modules

Administrators can update the PowerEdge 1855 blade server chassis module firmware using either the GUI or the CLI of the Dell Remote Access Controller/Modular Chassis (DRAC/MC). Both interfaces require administrators to download the firmware image from a TFTP or an FTP server. The updated firmware image should be made available in a designated directory on the TFTP or FTP server. *Note:* The *Dell Remote Access Controller/Modular Chassis User's Guide* provides specific instructions on how to set up the TFTP or FTP server for a firmware update at support.dell.com/support/edocs/software/smdrac3/dracmc.

Management module firmware updates

The DRAC/MC is the management module that allows administrators to monitor and manage chassis components in a Dell PowerEdge 1855 blade server. The PowerEdge 1855 blade server offers two different configurations for management module firmware:

- Single DRAC/MC
- Redundant DRAC/MC

Single DRAC/MC firmware. DRAC/MC firmware can be upgraded using the GUI, CLI, or DRAC/MC firmware recovery console.[2] The DRAC/MC firmware update is a TFTP-based update process. For all three DRAC/MC firmware update methods, administrators must complete the following setup procedures before starting the firmware update process:

1. Set up a TFTP server and copy the firmware image to the root of the TFTP server.
2. Record the IP address of the TFTP server and the file name of the updated firmware image.
3. Log in to the DRAC/MC, using either the GUI or CLI.

To update the firmware using the GUI, administrators should navigate to the Update tab and select "Firmware Update." Next, they should enter the TFTP IP address of the firmware image file name and start the DRAC/MC firmware update process by clicking "Update Firmware."

To update the firmware using the CLI, administrators should enter the following Racadm command at the DRAC/MC console prompt:

```
DRAC/MC: racadm fwupdate -a TFTP_IP_ADDRESS
    -d mgmt.bin
```

where `TFTP_IP_ADDRESS` is the IP address of the TFTP server and `mgmt.bin` is the name of the firmware image file.

Note: This Racadm command can be entered at the DRAC/MC serial console or at a Telnet session. Best practices recommend using the serial console because, if a TFTP download fails, administrators will lose Web access and network connections such as Telnet. DRAC/MC will boot to the DRAC/MC firmware recovery console, which can be accessed only at the serial console. The firmware recovery console provides the following options:

- Upgrade firmware from the serial port
- Upgrade firmware from the network
- Configure network parameters

Using the DRAC/MC firmware recovery console, administrators can restart the firmware update process either via serial port or the network.

Note: While in recovery mode, the DRAC/MC does not monitor chassis components of the PowerEdge 1855 blade server. For that reason, administrators should take extra care to minimize the amount of time that the DRAC/MC spends in recovery mode.

Redundant DRAC/MC firmware. In a redundant configuration, two separate DRAC/MC modules are installed in a chassis:

- A primary DRAC/MC module, which actively monitors the chassis
- A standby DRAC/MC module, which monitors the active signal from the primary DRAC/MC module (If a failure in the primary DRAC/MC module occurs for more than five seconds, the standby DRAC/MC module is designed to become the active, primary DRAC/MC module.)

The PowerEdge 1855 blade server supports redundant DRAC/MC mode if the DRAC/MC is running firmware version 1.1 or higher. Although redundant DRAC/MC modules can be updated with a single firmware package, the DRAC/MC goes through the following steps to complete the firmware update process once the administrator

[2] For more information about how to access the DRAC/MC GUI and the CLI (also known as the Racadm command-line utility), refer to the *Dell Remote Access Controller/Modular Chassis User's Guide* at support.dell.com/support/edocs/software/smdrac3/dracmc.

initiates a firmware update task, using either the GUI or the `racadm fwupdate` command from the CLI:

1. The primary DRAC/MC module starts the TFTP firmware update.
2. The standby DRAC/MC module monitors the chassis while the primary DRAC/MC module is updated. At this time, neither DRAC/MC is accessible, either through Telnet or the GUI.
3. When the primary DRAC/MC module completes the TFTP update, the TFTP update on the standby DRAC/MC module begins. The primary DRAC/MC module continues to monitor the chassis while the standby module is updating the firmware. At this time, neither DRAC/MC is accessible, either through Telnet or the GUI.
4. When the standby DRAC/MC module completes the firmware update process, the primary DRAC/MC module is available for network access. Telnet and the GUI become available.

KVM module firmware updates

The KVM module enables administrators to access server blades in the PowerEdge 1855 blade server by providing keyboard, monitor, and mouse functions as if the administrator were directly connected to the module. The PowerEdge 1855 blade server provides a built-in analog KVM module and an optional digital KVM module. Both KVM modules are flash-upgradeable.

Analog or digital KVM module firmware can also be updated using the DRAC/MC GUI or CLI. For both firmware update methods, administrators must complete the following setup procedures before starting the firmware update process:

1. Set up a TFTP server, and copy the firmware image to the root of the TFTP server.
2. Record the IP address of the TFTP server and the file name of the new firmware image.
3. Log in to the DRAC/MC, using either the GUI or CLI.

To update the firmware using the GUI, administrators should navigate to the Update tab and select "KVM Firmware Update." They should then enter the TFTP server IP address of the firmware image file name, and start the KVM firmware update process by clicking "Update Firmware." The TFTP download and firmware update process may take up to six minutes. After the update completes, the KVM will reset.

To update the firmware using the CLI, administrators should enter the following Racadm command at the DRAC/MC console prompt:

```
DRAC/MC: racadm fwupdate -a TFTP_IP_ADDRESS
    -d kvm_firmware_name -mkvm
```

I/O module firmware updates

The PowerEdge 1855 blade server chassis provides extensible I/O functionality such as networking, Fibre Channel, or InfiniBand connectivity. The McDATA 4314 Fibre Channel switch and Brocade Silkworm 3014 Fibre Channel switch provide Fibre Channel connectivity. The Dell PowerConnect™ 5316M Ethernet switch, a managed Layer 2 network switch, provides network functionality.

To update the firmware of the preceding I/O modules, administrators need to procure the IP addresses of the switches. As part of the installation of the McDATA, Brocade, and PowerConnect switches, administrators must configure the IP address using the corresponding switch configuration application. However, administrators can obtain the IP address of the Brocade switch using the DRAC/MC CLI as follows:

1. Log in to the DRAC/MC and connect to the switch using the `DRAC/MC: connect switch-N` command.
2. Log in to the switch with the username "admin" and password "password."
3. Enter the `ipaddrshow` command to obtain the IP address.

Unlike the TFTP-based DRAC/MC firmware update process, the McDATA and Brocade Fibre Channel switch module firmware update process is FTP based. Administrators must complete the following setup procedures before starting the firmware update process:

1. Set up an FTP server on the management station, and unzip the firmware in a local directory.
2. Record the IP address of the switch and the FTP server.
3. Ensure that the switch is in normal operation mode by inspecting the status LEDs.

McDATA 4314 Fibre Channel switch firmware. To provide consistent performance throughout the fabric, administrators should ensure that all switch modules are running the same version of firmware. Installing updated firmware requires a switch reset. A stable fabric is required to successfully activate the firmware on a switch without disrupting traffic. Therefore, administrators must ensure that no administrative changes are in progress anywhere in the fabric before installing the Fibre Channel switch firmware.

McDATA provides management station software called Enterprise Fabric Connectivity Manager (EFCM), which provides a GUI to update the switch firmware. Detailed instructions on how to use EFCM can be found in the McDATA 4314 switch documentation.[3]

[3] For more information about EFCM, refer to the *EFCM Management Guide* on the CD that ships with the McDATA 4314 Fibre Channel switch.

For the firmware update process, the McDATA management interface requires administrators to log in to the switch using the "admin" account and access the advanced configuration console mode using the `admin start` command.[4]

Brocade Silkworm 3014 Fibre Channel switch firmware. Brocade Fibre Channel switch firmware can be upgraded using either a Web-based GUI or a CLI. To update the fabric OS in command-line mode, administrators should execute the `firmwaredownload` command from an FTP server or from a local Network File System (NFS) directory while in `admin` mode:

```
firmwaredownload options host_or_IP,user,
   /path/to/the/pfile,passwd
```

The updated firmware is in the form of Red Hat® Package Manager (RPM™) packages with names defined in pfile, a binary file that contains specific firmware information and the names of firmware packages to be downloaded.

In dual-domain systems, the `firmwaredownload` command downloads the firmware image by default to both control processors (CPs) in rollover mode, which helps prevent disruption to application services. This operation depends on support for the High-Availability (HA) feature, which can be enabled through the `haenable` command in the switch CLI. If HA support is not available, administrators can still upgrade the CPs one at a time, using the `-s` option.[5]

To provide consistent performance throughout the fabric, administrators should ensure that all switch modules are running the same version of firmware.

Systems supported by the Brocade firmware have two partitions of nonvolatile storage areas—a primary and a secondary partition—to store two firmware images. The `firmwaredownload` command loads the updated image into the secondary partition and swaps the secondary partition to be the primary partition. The command then reboots the CP and activates the updated image. Finally, it performs the `firmwarecommit` procedure automatically to copy the updated image to the other partition (unless the `-n` option is used).

To update the firmware using the GUI, administrators should launch the Brocade Web console by entering the IP address of the switch in the browser address line. They should then log in as "admin" and navigate to the Firmware tab. On the Firmware page, administrators should enter the FTP server's IP address and the path to the firmware image file, then begin the firmware update process.

At the switch console, administrators can use the `firmwaredownloadstatus` command to monitor the download process. After the download is finished, administrators can enter the `firmwareshow` command to verify that the firmware update completed successfully.

Dell PowerConnect 5316M Ethernet switch firmware. Two firmware images can be stored in the flash memory of the PowerConnect 5316M switch module. The images are called active and nonactive, depending on which image the switch is currently running. The switch also supports two protocols to download the images: network-based TFTP and serial port–based xmodem.

To use the TFTP method, administrators must complete the following setup procedures before starting the firmware update process:

1. Set up a TFTP server.
2. Install the updated firmware image on the TFTP server.
3. Log in to the switch, and enter the privileged EXEC mode.[6]

After logging in, administrators can execute the following command in privileged EXEC mode to copy the file named "image" to the nonactive image file:

```
console# copy tftp://hostname/path/to/the/
systemimage flash
```

After the flash update is complete, the switch can be instructed to boot from either of the two images by executing the following command in privileged EXEC mode:

```
console# boot system {image1 | image2}
```

Administrators should enter the following command to verify whether the switch successfully booted into the updated system image:

```
console# show version
```

Although administrators may never need to upgrade the switch boot image, they can do so by executing the following command:

[4] For more information about the CLI-based firmware update process for the McDATA 4314 switch, refer to the *McDATA 4314 Command Line Interface Guide* on the CD that ships with the McDATA switch.

[5] For more information about command options for Brocade Silkworm 3014 Fibre Channel switch firmware, refer to the *Brocade Fabric Operating System (FOS) Reference Manual* on the CD that ships with the Brocade switch.

[6] For more information about the operation modes of the PowerConnect 5316M switch and how to configure the system identity, refer to the *Dell PowerConnect 5316M Ethernet Switch Module User's Guide* and the *Dell PowerConnect CLI 5316M Reference Guide* at support.dell.com/support/edocs/network/PC5316M/en.

```
console# copy tftp://hostname/path/to/the/
    bootimage boot
```

Note: Best practices recommend saving extra copies of the switch configurations on a TFTP server, especially before a firmware upgrade.

The PowerConnect 5316M Ethernet switch module used in the PowerEdge 1855 blade server chassis does not come equipped with its own serial console port. Instead, the serial console can be accessed as a module to which the DRAC/MC is connected. Therefore, to use xmodem as the source for the management station where the firmware image is stored, administrators must perform the following steps from the DRAC/MC:

1. Ensure that the current shell interface is already at the DRAC/MC command prompt. If not, switch back to the context of the DRAC/MC command prompt by pressing the Enter key, the tilde key, and the period key. *Note:* Press the Shift key if the tilde character is located in the upper register of the keyboard, and then press the period key.
2. At the DRAC/MC command prompt, issue the following command:

```
DRAC/MC: racadm config -g cfgSerial
    -o cfgSerialConsoleIdleTimeout 0x3000
```

3. Redirect the DRAC/MC serial console to the internal serial console interface of the PowerConnect 5316M Ethernet switch module in binary mode by entering the following command:

```
DRAC/MC: connect -b switch-N
```

where *N* is the chassis I/O module bay number in which the PowerConnect 5316M Ethernet switch module is inserted. Press the Enter key several times to ensure that the terminal connection is successfully established and that the Ethernet switch module prompt appears.

Note: To terminate the binary mode connection to the PowerConnect 5316M Ethernet switch module's serial console, disconnect the current session of the terminal.[7]

Keeping blade server system software components up-to-date

The modular design of blade servers, such as the Dell PowerEdge 1855 blade server, brings an added dimension to the traditional change-management process. Today's industry-leading change-management and software distribution frameworks such as Altiris Patch Management Solution and Microsoft SMS are designed to provide robust automated tools to manage updates on modules that reside on server blades. Because these frameworks do not scale to manage the updates of chassis management modules and I/O modules, updating blade server chassis modules requires careful planning and deployment such as the approach described in this article.

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, and Computer Associates in different capacities. He has an M.S. in Computer Science from Alabama A&M University.

Ruoting Huang is a development engineer in the Dell Enterprise Systems Management Software organization. He focuses on parallel processing and internetworking. Ruoting has an M.S. in Computer Science from the Asian Institute of Technology.

FOR MORE INFORMATION

Dell PowerEdge 1855 Systems User's Guide:
support.dell.com/support/edocs/systems/pe1855

Dell OpenManage:
www.dell.com/openmanage

Dell servers:
www.dell.com/servers

[7] For more information about configuring and using the DRAC/MC, refer to the *Dell Remote Access Controller/Modular Chassis User's Guide* at support.dell.com/support/edocs/software/smdrac3/dracmc.

Guidelines for Assessing Power and Cooling Requirements in the Data Center

As computing deployments continue to increase in density, cooling dynamics have become a paramount consideration in many data center environments. This article discusses steps that administrators can take to become familiar with intensifying power and cooling requirements, particularly for rack-dense blade servers—and explores tactics to help optimize the deployment of computing components throughout the data center.

BY DAVID MOSS

Related Categories:

Blade servers

Data center density

Data center technology

Planning

Services

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

The advent of blade server technology per se did not create a new challenge in data center power and cooling requirements, but it did exacerbate existing conditions. Administrators no longer can rely simply on HVAC (heating, ventilation, and air-conditioning) systems to cool data centers. They must move past the mindset of cooling the room into a mindset of cooling each rack.

Servers generate heat—in fact, 100 percent of the input power is released as heat, usually measured in kilowatts (kW). Most servers are cooled with internally generated airflow, commonly expressed as cubic feet per minute (CFM). In general, the more heat a server produces, the more airflow it is likely to consume to maintain the temperature requirements for its internal components. Unless data center administrators are satisfied with racks operating at a 2 to 3 kW load, they should pay special attention to airflow consumption within each rack. For example, in the late 1990s a typical 4U, two-processor server consumed about 40 to 50 CFM to cool a maximum load of approximately 400 watts (W)—or a nominal load of around 200 W. A full 42U rack of these servers would typically operate on 2 to 4 kW of power and require 400 to 500 CFM for cooling. In contrast, eighth-generation Dell™ PowerEdge™ 1855 blade servers consume approximately 2,000 CFM to cool 15 to 25 kW.

Moreover, as computing density increases in the data center, a corresponding increase in equipment exhaust temperatures occurs. Inside each rack, an inadequate supply of chilled air may result in the consumption of used air; some equipment may consume air from its own exhaust traveling over the top of the rack (see Figure 1). Exhaust temperatures can severely limit deployment potential if the warm air is allowed to recirculate to the intake of any servers. This is a fundamental challenge in

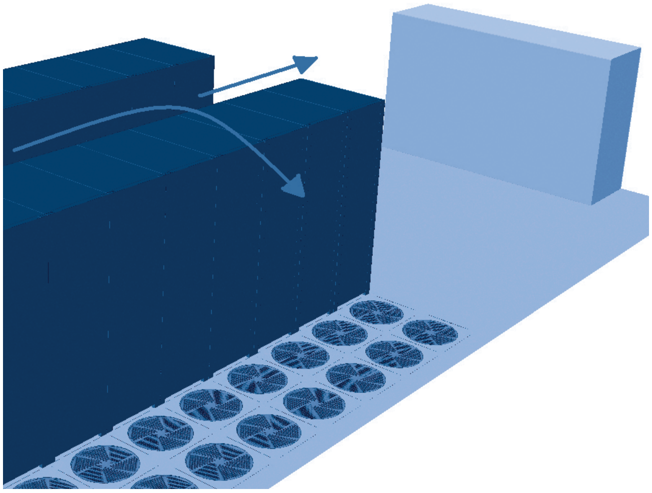

Figure 1. Equipment exhaust recirculated through the rack

today's data centers. The server intake aisle must fill up with chilled air about as quickly as the servers consume the air. Otherwise, hot exhaust air may recirculate through the equipment and result in overheating.

Because of the open nature of a raised-floor environment, administrators cannot be certain that the air they supply directly in front of each server actually gets to that server. Administrators must understand and plan around the actual airflow dynamics of the data center. The airflow typically includes areas of weakness, which can result in hot spots, as well as areas of strength, which can be cool enough to support more equipment. The challenge is to understand and take advantage of the strengths and weaknesses in data center locations or to intelligently manipulate those strengths and weaknesses to normalize the airflow.

Following standard best practices

The first step for data center administrators is to understand basic best practices. Wherever possible, administrators should ensure strict adherence to the *hot aisle/cold aisle* approach of alternating intake aisles and exhaust aisles—that is, rack intakes should face each other and rack exhausts should face other. This approach accomplishes two important goals: It helps keep the exhaust from one row feeding the next, and it provides for a concentration of chilled air at the server's intake. Administrators should avoid introducing chilled air to the exhaust aisle—both intentional introduction through hot-aisle vent tiles and seepage from floor cable openings. Even a small cable opening can equal or exceed the amount of air that flows through a cold-aisle vent tile. Cable grommets can be used to help block the intrusion of chilled air into the hot aisle, as shown in Figure 2.

Although it may be tempting to cool a hot aisle so the temperature is more comfortable for administrators, the exhaust aisle should remain hot to maximize the HVAC return temperature, which helps the HVAC to operate as efficiently as possible. Gaps between equipment should be avoided—including gaps between racks and gaps within racks. Blanking panels can be used within racks. Such basic best practices lead to a common goal: creating the coldest possible intake aisle with the highest possible volume of chilled-air delivery and the warmest possible exhaust aisle.

Estimating airflow rates and cooling capacity

To move beyond the basics, administrators should develop an understanding of data center dynamics, including airflow delivery capability as well as the equipment's airflow consumption rate. The delivery capability is simple—administrators simply measure tile-flow rates. The facilities department within an organization may have a flow-hood, such as the one shown in Figure 3, to help take such measurements. In addition, consulting services such as Dell™ Data Center Environment Assessment (DCEA) are available not only to measure airflow rates but also to analyze that data and make specific recommendations.

The first step for data center administrators is to understand basic best practices.

The other side of the equation is consumption, which may be trickier to calculate because hardware vendors do not always provide airflow consumption rates. Dell provides consumption rates for most of its rack-mount equipment through the Dell Product Configuration Calculator, which is available at www.dell.com/calc.

Dell servers use variable-speed fans controlled by algorithms that use ambient and component temperature sensors. Airflow rate has a high dependence on inlet ambient temperature. Cooler inlet temperatures equate to lower airflow rates. In the absence of measured values, administrators can estimate equipment airflow rate by

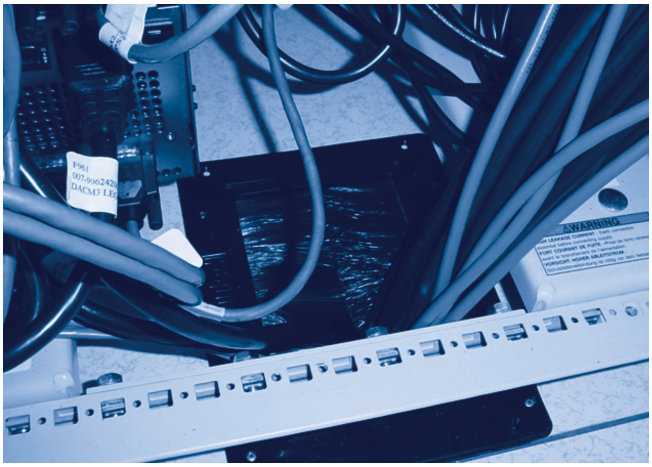

Figure 2. Cable grommets to minimize loss of chilled air

using the temperature difference between the intake and exhaust as well as the equipment's power draw. The airflow rate can be approximated using one of the following formulas:

$$\text{CFM} = 1.78 \times \frac{\text{equipment power (W)}}{\text{temperature difference (°C)}}$$

$$\text{CFM} = 3.2 \times \frac{\text{equipment power (W)}}{\text{temperature difference (°F)}}$$

If the equipment power is known but the temperature difference is not, administrators can approximate the airflow rate at 9 CFM per 100 W, which assumes a temperature difference of about 20°C or 36°F, depending on the formula used. Administrators should be careful when estimating airflow rates with either method. Because the airflow rate is directly proportional to the equipment power, administrators can easily overestimate airflow rates if using conservative power numbers. Best practices recommend using measured values over estimates for equipment power whenever possible.

This information can then be used to calculate cooling capacity. At minimum, airflow rate balances should be studied. Administrators should compare cumulative rack consumption rates versus cumulative delivery rates, which are best compared on a row-by-row basis. Also, the total airflow rate supplied to an aisle should be examined in comparison to the consumption of the aisle's two flanking equipment rows. This process should then be applied to the rest of the data center. For example, if vent tiles are allowing some areas to be overprovisioned with cooling in comparison to consumption, administrators may ameliorate the situation simply by moving vent tiles away from areas that are overcooled to areas where more cooling is needed.

Administrators can determine the effect on an area in which the chilled-air delivery is reduced by studying the exhaust temperatures and top intake temperatures in that area. If the intake of the highest mounted systems is still significantly below equipment specifications—which is generally around 35°C (95°F) but may be de-rated at higher altitudes—then administrators may be able to move some of the cooling from that area. If the exhaust temperatures in the aisles flanking that cold aisle are high and approach the equipment operating (intake) temperature limits, then the removal of the chilled-air supply in that area could result in a deficiency and subsequent recirculation of exhaust into the intake aisle. A simple test of blocking off the airflow from a tile or two may help administrators understand whether chilled air can be reprovisioned from one area to another. If no significant change in the inlet temperatures of several surrounding top-mounted servers ensues, the air could be reprovisioned to an area of greater need. *Note:* The methods described in this section involve a substantial amount of trial and error, which can affect mission-critical equipment. A more scientific approach is explained in the next section.

Predicting data center cooling needs with CFD

Numerical methods for computer modeling have been used effectively to predict the behavior of many types of dynamic systems. Whether for the prediction of automobile aerodynamics, the strength of a building structure, or the weather patterns for a five-day forecast, computer modeling helps solve a variety of engineering problems, and the data center is an excellent candidate for this approach. *Computational fluid dynamics* (CFD) is the term generally applied to the analytical modeling of fluid systems or air systems, such as in the case of a data center. Several software products can help administrators create a thermal/airflow model of their data centers to perform virtual trial-and-error scenarios. *Note:* Although this type of analysis is best performed by engineers with the proper expertise, data center administrators can also learn to perform this type of environmental study effectively.

Organizations also have the option of hiring CFD consultants, including those from Dell. Dell plans to offer this type of analysis as an option of the Dell DCEA service. Through the DCEA service, an organization can have its data center measured and analyzed, including a CFD

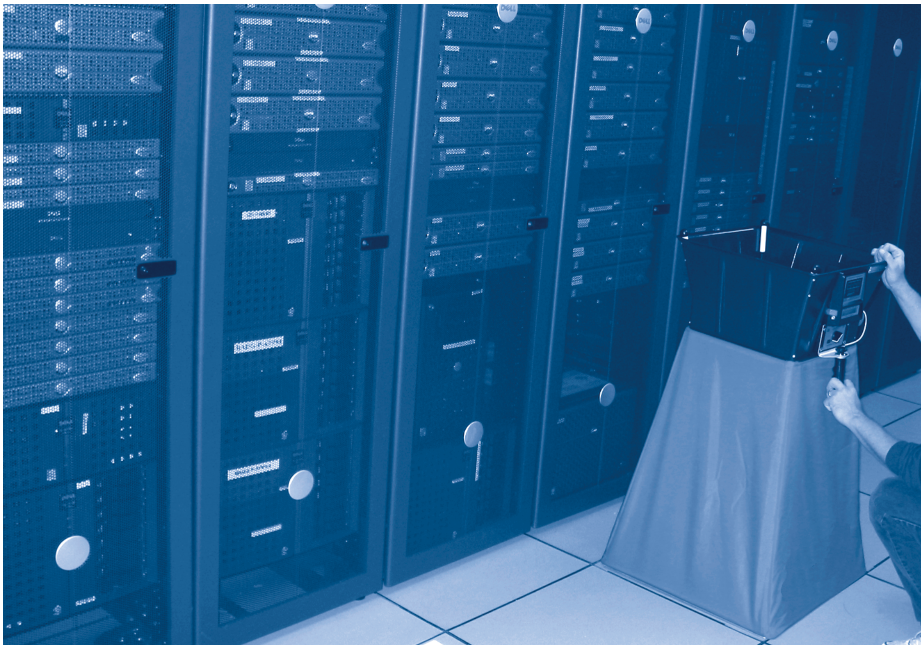

Figure 3. Flow-hood to measure the airflow rate of an entire tile

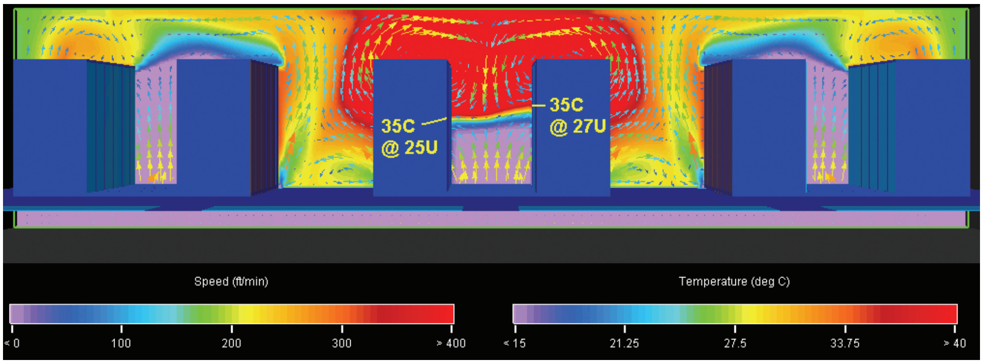

Figure 4. Temperature and air vector plot for example CFD analysis

analysis, with a resulting confidence level that the infrastructure will support future needs.

To understand how a CFD analysis can help, administrators can consider the example CFD model depicted in Figure 4, which was created using Flotherm, a leading thermal analysis application from Flomerics. This cross-section through six rows of racks shows airflow patterns superimposed onto a temperature profile. The rows are arranged according to the hot aisle/cold aisle approach. In this particular example, the CFD analysis proposed that equipment with a much higher air consumption and heat load be deployed in the two middle racks. Airflow is distributed uniformly to each of the three cold aisles. However, the cool temperature (depicted in lavender in Figure 4) completely fills the two outside cold aisles but does not fill the middle aisle. In fact, the cool temperature billows out of these two outside aisles. In the middle aisle, the cool air travels only about halfway up the rack before it is depleted by system consumption. The air pattern shows hot air being re-circulated over the top of the rack at temperatures that exceed equipment specifications.

The plot in Figure 4 is set up to represent in red any air temperatures greater than 40°C (104°F), which is five degrees higher than most equipment's maximum inlet temperatures. The analysis ultimately shows that, without changes, this middle aisle cannot handle the proposed deployment of new equipment because equipment deployment above the 25U rack location will receive unacceptably high inlet temperatures. However, the overabundance of cold air in the outer two aisles indicates that a significant amount of air could be diverted from these aisles to the middle aisle to help satisfy its cooling needs. Using CFD analysis, the administrator in this example scenario could model various configurations to help determine the best way to adequately divert the excess cold air—without incurring the risk of physically reconfiguring the actual data center equipment for numerous trial-and-error comparisons. In this particular example, subsequent numeric iterations suggested removal of 20 percent of the vent tiles from the outer cold aisles and the insertion of several grates in the center cold aisle. The CFD analysis indicated that grates, at 60 percent open, could offer substantially more airflow than the standard 25 percent open, perforated tiles that had previously been in use.

By taking advantage of CFD analysis, Dell has helped many organizations optimize their data center cooling capabilities—often without the need for additional HVAC infrastructure. While such optimizations can be determined through the trial-and-error approach, using CFD analysis software as a predictive, iterative tool enables data center administrators to find a quick, efficient path to resolution.

Once the delivery of chilled air has been optimized, data center administrators who need the capability to increase the deployment should consider alternate sources of chilled air. For example, supplemental cooling systems can offer chilled-air generation at the rack level, and refrigerant-based systems are designed to blow air down into the cold aisle to supplement the raised-floor delivery. Other types of systems can fit in line with racks to deliver supplemental or primary chilled air. In addition, some self-contained racks can generate chilled air inside the rack for the enclosed equipment.

Balancing power and cooling needs in the data center

In today's world of high-density data center equipment and chilled-air delivery challenges, administrators must understand the supply limitations of the data center and the demand of the equipment. In short, power and cooling needs are a supply and demand problem. By learning how to properly assess data center power and cooling requirements and procuring the appropriate tools to help resolve such environmental considerations, administrators can be enabled to design the optimal enterprise computing environment—and to deploy rack-dense servers with confidence.

David Moss is a senior thermal/mechanical architect with Dell and has over 20 years of experience in electronics packaging design. In his role at Dell, he holds 18 U.S. patents. David has a B.S. in Mechanical Engineering from Texas Tech University.

FOR MORE INFORMATION

Dell Product Configuration Calculator:
www.dell.com/calc

Information Life-Cycle Management

in a Virtual Data Center

Growing acceptance of commodity, standards-based hardware and server virtualization is leading many organizations to consider dramatic changes in their IT infrastructure. These changes are expected to transform the enterprise data center into well-defined building blocks of virtualized compute, storage, and network resources. In the resulting IT environment, business unit managers should be able to establish performance, availability, compliance, and data classification requirements for enterprise applications. The intelligent IT infrastructure will be designed to automatically provision the necessary hardware and software resources once business requirements have been defined, without further administrative intervention.

BY NIK SIMPSON AND MATTHEW BRISSE

Related Categories:

Data center technology

Information life-cycle management (ILM)

Planning

Storage

Systems management

Virtual data center

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

Despite advances in many areas of data center technology, IT organizations are still grappling with growing complexity and cost. At the same time, administrators are under pressure to control total cost of ownership for the IT infrastructure. One promising solution is a data center infrastructure that is designed to adapt to changing business and application requirements by automatically provisioning resources—both hardware and software—as demand requires. The goal is to develop a data center that is built upon well-defined, standards-based building blocks of virtualized compute, storage, and network resources.

In such a virtual data center, when business unit managers roll out new applications they will establish business requirements for performance, availability, compliance, and data retention. In response, the virtual data center will have the capability to provision the necessary hardware and software resources to meet business requirements without further administrative intervention once those business requirements have been defined. The virtual data center will be designed to monitor each application to ensure that it meets the established business requirements (or updated requirements caused by changing business conditions)—dynamically adjusting data center resources as required. The result is envisioned as a closed-looped management system that will have the capability to become increasingly self-managed over time.[1]

The concept of the virtual data center is powerful, but to exploit its full capability, enterprises require a strategy for information life-cycle management (ILM). The need for ILM derives from a basic tenet of the virtual

[1] For more information about the virtual data center, see "Progressive Degrees of Automation Toward the Virtual Data Center" by Jimmy Pike and Tim Abels in *Dell Power Solutions,* February 2005; www.dell.com/downloads/global/power/ps1q05-20040228-Pike.pdf.

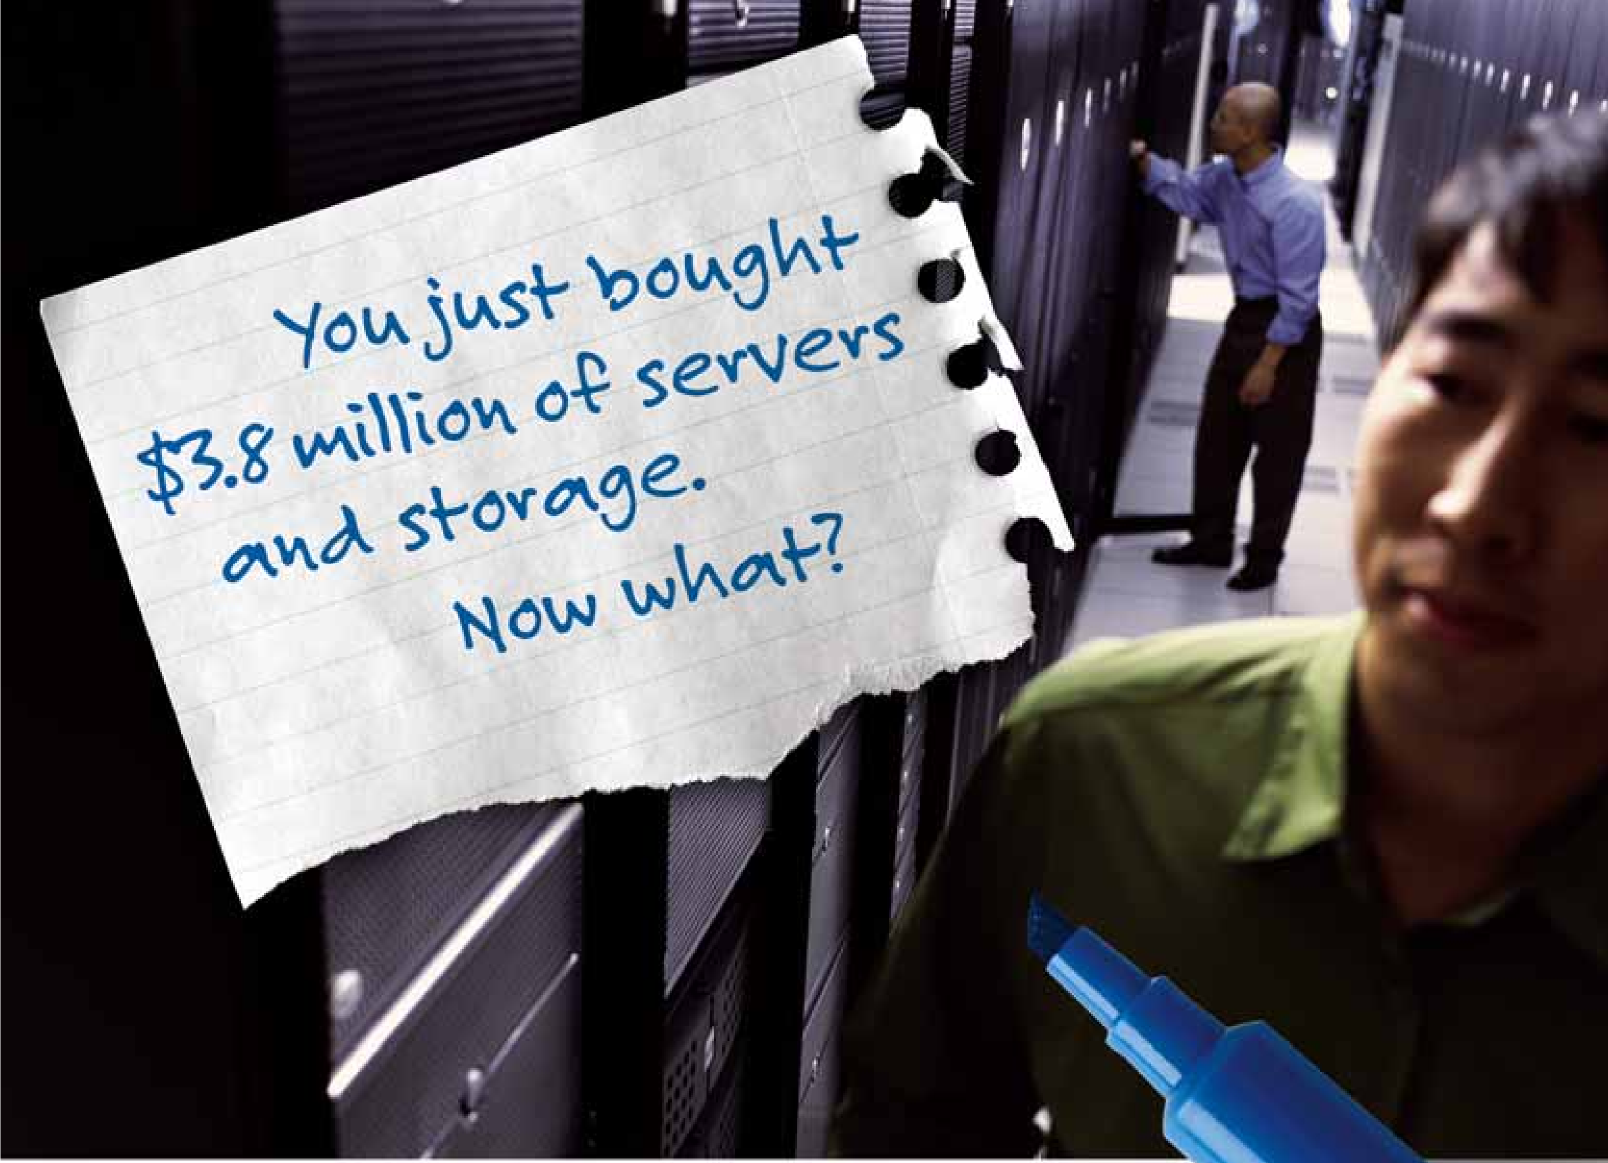
You just bought
$3.8 million of servers
and storage.
Now what?

data center: installing an application is a logical task for which the administrator enters basic requirements—or *service-level objectives* (SLOs)—and the infrastructure management layer provisions the resources to meet those requirements. In the virtual data center context, ILM requirements must be reflected in the SLOs when an application is deployed. Examples of ILM requirements that might be specified include:

- **Downtime:** The amount of downtime that is acceptable for the application. This metric helps the infrastructure layer determine which disaster recovery options are available.
- **Time to zero value:** The length of time for the value of the information to reach zero. For example, if the system is producing weather forecasts, then the time to zero value is one day—yesterday's forecast has historical significance but little value. This metric can be used to determine when to purge old information or move old information to less-expensive storage as its value declines.
- **Regulatory compliance:** Regulations, if any, that affect the information created by the application. For example, the information may need to meet the retention requirements of the Sarbanes-Oxley Act. This metric could be used to set up off-site file archives.

Providing ILM requirements as part of the application deployment process helps align information management actions with the value of the application from the moment it is installed.

Understanding the information management problem

Although the basic concept of ILM has been a standard operating procedure in the mainframe world for several years, outside the mainframe world ILM is far less mature. In the open systems world, most IT organizations have some procedures in place (often in the form of custom scripts and manual administration tasks) to manage the information produced by applications and users. Unfortunately, these procedures typically suffer from the following problems:

- **Lack of integration:** Few IT organizations have a holistic approach to information management; procedures tend to be application- and operation-specific, using different tools with different management interfaces for each task. This heterogeneity can make it difficult to document procedures and ensure that the right procedures are in place for every application.
- **Insufficient distinction between information types:** Not all information is of equal value, and the procedures for tasks such as backup and disaster recovery should be driven by the value of the information. Today's data centers typically implement one of two classes of ILM: a combination of off-site disaster recovery, extensive archiving, and comprehensive backup; or a bare-minimum tape backup. This binary approach can lead to excessive IT costs because administrators tend to err on the side of caution and overprotect information that may have little residual value.
- **Inability to scale:** An ad hoc approach to information management is feasible as long the number of servers and quantity of information remains small. However, even in midsize companies, existing ILM approaches can be difficult to scale as the number of servers and the quantity of stored information continue to increase.

Until recently, such problems have been marginalized while efforts were focused on physical infrastructure management. Today, management tools for physical infrastructure have matured and the time is fast approaching for organizations to address information management in a consistent, flexible, and scalable way. Reasons for a renewed focus on ILM include the following considerations:

- **Growing complexity in data and information management:** As organizations have come to rely on computing infrastructures for a growing number of day-to-day functions, effective information management has become increasingly important.
- **Regulatory compliance:** The problem of complexity is compounded by the increasing number of government regulations—such as the Health Insurance Portability and Accountability Act (HIPAA) and the Sarbanes-Oxley Act in the United States, and the Data Privacy Directive in the European Union—that affect the way organizations manage data.
- **Growth of digital information storage:** New applications, new regulations, and increased use of existing applications are leading to rapid growth in capacity requirements.
- **Performance problems caused by information overload:** As the amount of information stored by an application grows, the performance of the application typically degrades and administrative operations such as backups and disaster recovery also take longer.
- **Increasing maturity in infrastructure management:** In the past few years, manufacturers and industry standards bodies have cooperated to produce standards, such as the Storage Network Industry Association (SNIA) initiatives to simplify the management of storage infrastructure.

Altogether, these issues have caused information management to become a high priority for many organizations. Centralized control of how information is managed can be a key first step toward more ambitious goals such as creating a virtualized data center. Perhaps the biggest roadblock for many organizations is a lack of standards.

The need for open standards

A standards-based approach to ILM procedures will require broad cooperation from the entire IT industry. To that end, the SNIA is currently focusing on information management and the underlying storage infrastructure through two projects:

- **Storage Management Initiative (SMI):** The SMI aims to create a standard set of technologies for the management of storage—whether in a storage area network (SAN), network attached storage (NAS), or a locally attached disk—and its associated infrastructure. With support from major vendors in the storage industry, the SMI has made great strides in creating a set of interfaces that are designed to manage storage infrastructure components regardless of vendor.
- **Data Management Forum (DMF):** The DMF is charged with two tasks: defining a common language and a set of SLOs for information management; and working with the SMI to define a road map for the development of ILM technologies and how those ILM technologies will interact with other management interfaces like the SMI specification.

The goal of the SMI and DMF efforts is to create the set of service layers shown in Figure 1. These service layers will be designed to communicate through open interfaces so that high-level SLOs at the ILM layer can be translated into low-level data management and storage management services that have the capability to provision and manage the storage infrastructure.

For example, an ILM SLO could require certain application data to be available 24/7 with a two-hour maximum acceptable recovery time. This would cause the data management layer to request a locally mirrored volume (for maximum reliability under normal operation) as well as synchronous mirroring to a remote data center for two-hour recovery in the event of catastrophic failures. As with other elements of the virtual data center, this capability is far from being realized. However, the DMF has defined a set of steps that can help organizations to develop the fully interoperable tool sets required for broad adoption of ILM technologies.

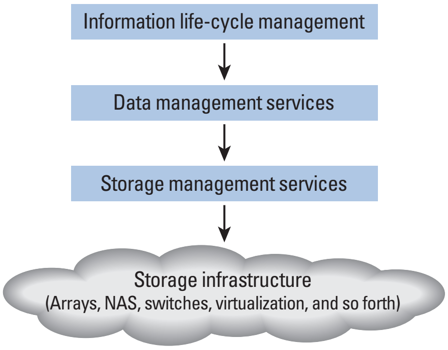
Information life-cycle management
Data management services
Storage management services
Storage infrastructure
(Arrays, NAS, switches, virtualization, and so forth)

Figure 1. A layered approach to the virtual data center

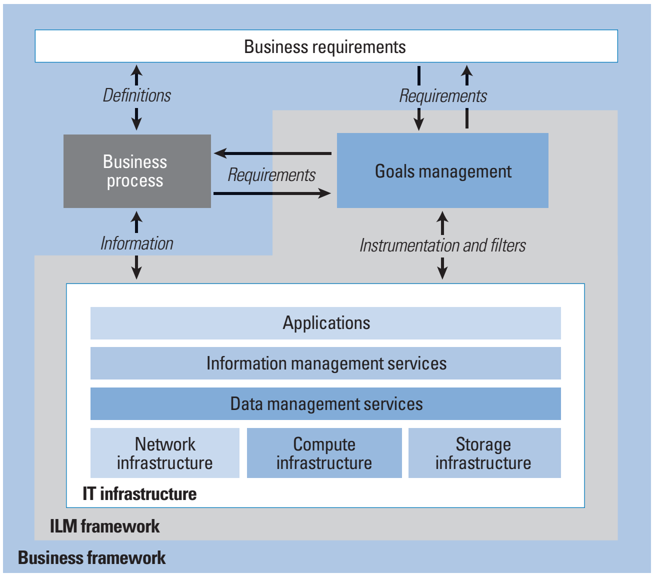

Figure 2. DMF framework for ILM in the virtual data center

Employing the DMF ILM model

Understanding the DMF's model for ILM (see Figure 2) is critical to understanding how ILM may be implemented in the future. The ILM model reflects a pragmatic approach to data management, in which management and control of the IT infrastructure is based on needs defined by the business framework. Business requirements coming from the business framework drive the management of applications and information. Within the ILM framework, the goals management layer transforms business requirements into policies that are enforced in the service and infrastructure layers. In turn, the goals management layer provides feedback to the business framework regarding cost, risk, and status.

Based on this ILM model, the system administrator would register an application with the mapping and control layer of the virtual data center management framework, which is detailed in Figure 3. The administrator would then specify a level of service and performance or a compliance requirement such as Sarbanes-Oxley. The virtual data center model would reference the ILM requirements to deploy the application and provision the necessary resources and methodologies to meet the SLO—without requiring the administrator to provision resources, tune the application, or define a backup and recovery strategy.

Implementing intelligent, flexible data management

Because the information management needs of an application change over time, managed information should be continually reevaluated to help ensure that it is managed appropriately. Achieving this goal requires a flexible and intelligent data classification system capable of analyzing the information and applying data management rules that match the needs of the application and the value of the information. However, intelligent classification is only half of the solution. Once information has been classified, intelligent data placement is required to direct that information to the appropriate tier of storage—or to arrange operations such as backup and replication to help ensure that the services provided to the application meet the SLO.

The requirement for dynamic information management is spurring a new generation of software designed to automate data classification and use the results from classification to drive intelligent data placement. The combination of intelligent data classification and placement can allow software to execute simple tasks such as ensuring that adequate capacity is provided for the data created by an application. In addition, intelligent data classification and placement can enable the development of software designed to execute complex tasks that have the capability to meet the high-level requirements of Sarbanes-Oxley regulations. Creating such software will require the development of standard interfaces for every aspect of an application's environment—allowing data to be moved, copied, purged, archived, or otherwise manipulated to meet the stipulated data management requirements.

Such comprehensive ILM requirements indicate that next-generation systems must be designed to support the following features and capabilities:

- **All data types:** To meet complex information management requirements, systems should be capable of working with all types of data—regardless of which application produces the data. Platforms that support a single data type or application will not suffice.
- **Any platform:** Corporate information is stored on many types of platforms, including UNIX®, Microsoft® Windows®, and Linux® operating systems. A comprehensive system should be agnostic about OS platforms and infrastructure choices such as NAS versus SAN.
- **Flexible classification:** The system must be able to combine a variety of techniques for classifying information to determine value. Conventional methods, such as age or size, should be combined with emerging methods, such as enterprise search engines, to provide the flexibility needed for truly intelligent data placement.

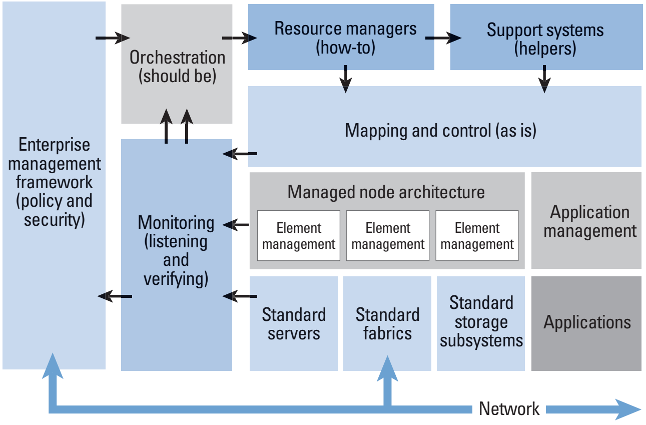
Enterprise management framework (policy and security)
Orchestration (should be)
Resource managers (how-to)
Support systems (helpers)
Mapping and control (as is)
Monitoring (listening and verifying)
Managed node architecture
Element management
Element management
Element management
Application management
Standard servers
Standard fabrics
Standard storage subsystems
Applications
Network

Figure 3. Virtual data center management framework

Evolving into the virtual data center

Well-defined, standards-based data center components for compute, storage, and network resources can enable enterprises to transition to a virtual data center infrastructure that has the potential to automatically provision hardware and software resources in response to changing business and application requirements. The key advantage of information life-cycle management is its ability to help solve the growing problem of data center complexity. By streamlining business processes, ILM can help reduce complexity and enable substantial cost reductions for IT infrastructure. Given its focus on business requirements and emerging standards, the ILM approach explored in this article is well suited to help automate the data center of the future.

Nik Simpson is director of marketing at Scentric Inc., an independent software vendor that designs next-generation ILM technologies. Nik also co-chairs the DMF's Information Lifecycle Management Initiative Technical Liaison Group and co-authored the SNIA's *Guide to Storage Virtualization*.

Matthew Brisse is a technology strategist in the office of the Dell CTO and is vice chair of the SNIA board of directors. At Dell, Matthew chairs the Standards Review Board and is a member of the Systems Management Architecture team.

FOR MORE INFORMATION

Storage Network Industry Association:
www.snia.org

Business Continuity for Exchange:

Protecting the E-mail Infrastructure

Safeguarding e-mail applications and developing a comprehensive business continuity strategy are critical considerations for the enterprise IT infrastructure. This article explains how administrators can build an effective business continuity plan to protect their Microsoft® Exchange e-mail environments. In addition, this article explains how to implement that plan using Dell software and support services running on Dell™ PowerEdge™ servers, Dell/EMC Fibre Channel storage arrays, and Dell PowerVault™ tape backup.

BY ARUN TANEJA AND ALEX GORBANSKY

Related Categories:

Business continuity

Dell/EMC storage

Disaster recovery

Microsoft Exchange

Planning

Storage

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

E-mail has become indispensable to nearly every organization because virtually everyone uses e-mail to communicate throughout the day about ongoing business operations. Whether for sales, marketing, or manufacturing, many core business processes now depend on e-mail as a critical internal communications platform. In addition, e-mail has become an important repository of business information—to such an extent that many organizations now use e-mail as a substitute for file systems or document management systems. For these reasons, e-mail performance and availability can have an enormous impact on business processes throughout the enterprise.

One of the most commonly deployed software platforms for e-mail systems is Microsoft® Exchange, which leverages features and benefits of the prevalent Microsoft Windows® OS. This article explores the challenges of protecting Exchange-based e-mail environments from failure, steps to building an effective business continuity plan for Exchange systems, and Dell products and services that are designed to enable business continuity in Exchange environments.

Challenges of managing Exchange-based environments

Administrators with direct responsibility for e-mail systems face many different types of challenges in dealing with an application that has enormous visibility to the organization at large and to every employee in the organization. First, administrators must proactively manage rapid e-mail growth, which entails keeping the system up and running while mailboxes are added for a continuing influx of new online users.

Second, administrators must develop a contingency plan to help protect the e-mail environment against failures. What if the e-mail server experiences a serious system problem or mailboxes become corrupted? Will the e-mail platform be adequately protected? Administrators should consider such issues, especially if planning to upgrade the e-mail system. The implications of not doing so may be catastrophic to the organization.

Unfortunately, current business continuity practices in Exchange environments typically do not address business and operational realities. For example, many Exchange e-mail systems use only off-site tape vaulting for business continuity. In such environments, data may not be recoverable for hours or days—and in some cases, not at all.

Of course, no business continuity plan can guarantee that an organization will be protected from an outage or a data-loss scenario. An administrator's responsibility is to understand the expected service levels and recovery objectives of the organization and then to design the appropriate approach to minimize the risk and impact of various failures that could occur.

Key layers in a business continuity strategy

One of the most difficult challenges for e-mail administrators can be determining how to create a business continuity plan. No one strategy will work across the board for all Exchange environments. Also, the business importance of e-mail systems will vary significantly among organizations and even within a specific organization.

To enable administrators to meet a diverse set of application and business requirements, best practices recommend that a business continuity plan be organized into layers. From those layers, administrators can determine products and services that are appropriate for the specific environments that require protection. Figure 1 lists the key protection layers that can help guide a business continuity strategy.

Layer 1: Building a foundation with platform availability

A viable business continuity strategy depends on a solid foundation, which is represented by the platform layer. This layer consists of internal high-availability system features in the server and storage nodes that comprise the Exchange infrastructure. Administrators should make sure that each system is designed with the appropriate level of redundancy and fault tolerance, enabling it to survive one or more internal component failures. That requires redundant, hot-swappable fans and power supplies as well as mirrored disk drives for booting.

Moreover, administrators should validate that existing platforms have the appropriate levels of internal fault tolerance to support uptime requirements. All too often, organizations invest in higher-layer technologies such as remote replication without solidifying the foundation. Without a solid platform layer, organizations may end up relying on costly disaster recovery systems rather than helping to prevent disasters in the first place. Disaster recovery procedures can exact a heavy toll in terms of time and expense, especially when recovering business-critical systems. Starting with a solid system platform and building protection in logical layers can help avoid triggering disaster recovery procedures.

Protection layer	Helps safeguard against	Technology
1	Platform downtime	High-availability system features
2	Data loss and corruption	RAID, snapshots, and tape or disk backup
3	Application and server failure	Clustering
4	Site failure	Remote data replication

Figure 1. Layers of protection in the business continuity strategy

Layer 2: Safeguarding Exchange from data loss and corruption

The data layer can help provide capabilities to safeguard Exchange from both data loss and data corruption. First, RAID technology is designed to provide a primary level of data protection from storage media (disk drive) failures, helping to prevent data loss when an individual drive fails. However, RAID does not address the issue of data corruption—and the corruption of individual mailboxes represents one of the most common operational challenges in Exchange environments.

Point-in-time copies, or *snapshots,* are designed to address data corruption issues. Snapshots enable administrators to seamlessly and automatically roll back the state of a mailbox to the last "good copy" with data integrity. This approach allows administrators to take a series of snapshots during the course of the day and keep the snapshots on rotation for recovery purposes. Thus, if a particular mailbox or Microsoft Personal Folders (.pst) file becomes corrupted, administrators can roll back to the latest consistent version of the data using the appropriate snapshot. Because the rollback is from disk rather than tape, recovery can be almost instantaneous and highly reliable.

Finally, administrators can configure a tape- or disk-based backup system. If all else fails, a backup system can be used to restore application data that is stored on tape or on disk. Restoring from tape can be time-consuming and is prone to tape media errors. If relying on tape alone as a backup medium, administrators may have to accept relatively slow recovery-time objectives (RTOs) and significant gaps in recovery-point objectives (RPOs) due to the sequential nature of the tape medium. Disk-based backup systems can help significantly improve recovery times and recovery granularity compared to tape-based backups, but disk-based backups tend to be costlier than tape.

As a baseline, administrators need some form of RAID protection and tape backup. If the organization's backup window is shrinking and administrators cannot consistently meet RTOs and RPOs with tape alone, a disk-based backup system may be worth considering. Finally, if data corruption is a concern and an organization needs to minimize data loss, a snapshot approach may be appropriate.

Layer 3: Maintaining application and server availability

The first two layers of the business continuity strategy focus on the hardware infrastructure of the Exchange environment, while the third layer addresses OS and application availability and potential software failures. When an application fails, the organization could

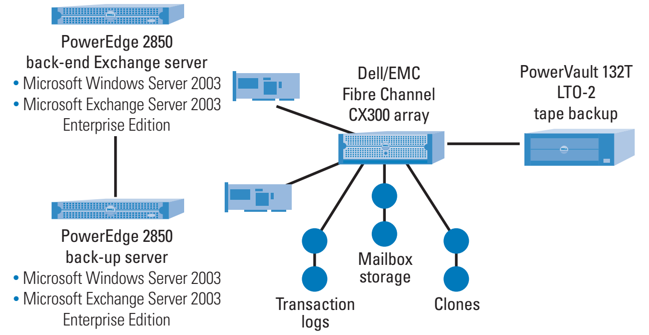

Figure 2. Exchange environment incorporating Dell PowerEdge servers as well as Dell/EMC storage arrays and Dell PowerVault tape backup

incur downtime of hours or even days. For many organizations, the requisite Exchange availability to be maintained in the event of a server or software failure is specified in an internal service-level agreement. If the organization cannot tolerate the potential downtime associated with restarting Exchange—or in some cases, procuring and provisioning a new Exchange server—a cluster configuration should be considered.

High-availability (HA) server clustering technologies are designed to recover from software failures. An HA cluster typically consists of a primary server and a secondary server connected—via a private LAN—to a shared storage pool. If the primary server fails, the secondary server is designed to transparently take over the application workload and access the same data set that was being accessed by the primary server, without disrupting end-user applications. In addition, clustering can enable administrators to perform software updates without interrupting service to end-user applications.

Layer 4: Preventing total site failure

An entire site outage typically represents the ultimate disaster, but it is nevertheless a plausible scenario for which administrators should develop robust recovery strategies. The capability to protect the organization from days of application downtime and data loss requires that data be replicated in a consistent manner to a remote standby site or a hot disaster recovery site.

Organizations can design and configure a replication site in-house, or they can use a third party to set up, host, and provide mobile access to the replication site. Determining how far a replication site should be from the central data center depends on geography and business risk factors, and organizations can work with insurance companies to help analyze such factors.

Several types of remote replication technologies exist. Host-based replication software runs on the application server and is platform- and OS-dependent. Array-based replication software runs on the storage array and is independent of the application server. More recently, replication software that runs on an appliance or switch in the network has become available. Each type of replication software offers it own set of advantages and disadvantages in terms of cost, performance, and flexibility.

Of course, every Exchange environment does not require this level of protection. Organizations will ultimately have to weigh all the costs associated with replication sites—including hardware, bandwidth, and administrative resources—against the benefits of risk reduction. If the data center cannot tolerate the time and potential data-loss window for recovering Exchange data from tape in the event of a site outage, remote replication may provide an extra layer of resiliency to help protect the e-mail environment.

A flexible Microsoft Exchange environment for business continuity

Dell's hardware, software, and support services are designed to help organizations build a complete Microsoft Exchange environment and effectively address the key layers of business continuity. Figure 2 shows an example Dell-based Exchange environment for approximately 2,000 users.

Achieve high availability with servers and storage arrays

Dell PowerEdge servers have been designed to enable a high degree of system reliability and availability, including the following features and capabilities:

- Redundant, hot-swappable disk drives, fans, and power supplies
- Redundant, integrated network interface cards and host bus adapter slots
- RAM protection using a combination of memory mirroring and error-correcting code (ECC)

Dell has also partnered with EMC to deliver storage arrays that are designed to enable high availability and reliability, including the following features and capabilities:

- Redundant, hot-swappable disk drives, fans, and power supplies
- A designated global hot-spare drive for each array
- Redundant storage processors for front-end server connectivity
- Mirrored, battery-backed write caches for each storage processor

Prevent data loss and corruption with RAID, backup systems, and point-in-time copies

Dell/EMC storage arrays provide a range of RAID options designed to protect applications from data loss in the event of hard drive failures. For this purpose, several Dell PowerVault backup systems as well as Dell/EMC disk subsystems can provide low-cost, high-capacity Serial ATA (SATA) drives. In addition, Dell has partnered with EMC Legato, CommVault, and VERITAS to provide backup applications such as Legato NetWorker® software, CommVault Galaxy, and VERITAS NetBackup and Backup Exec.

Galaxy and Backup Exec include Exchange suites that are particularly tailored to Dell-based environments.

For point-in-time copies, EMC® SnapView™ software offers an array-based snapshot approach that comes as an option with the Dell/EMC product line. SnapView enables administrators to create multiple, space-efficient point-in-time snapshots or addressable business continuity volumes of Exchange data. This approach enables administrators to keep several snapshots from the same day on rotation for recovery purposes.

Legato® PowerSnap™ software enables strong support for heterogeneous storage environments that use a variety of snapshot technologies. PowerSnap provides a common and open management framework that allows administrators to manage snapshots across various applications, operating systems, and storage platforms.

Gain application and server availability with clustering

Administrators can implement Microsoft Cluster Service, a standard Windows-based clustering product, on Dell platforms. Dell provides the option of using SCSI or Fibre Channel (FC) storage on the back end, making clustering a viable approach for both small and midsize enterprises.

Protect the site with remote replication

Dell offers remote data replication software to help protect Exchange environments. For example, EMC SAN Copy™ software provides an array-based approach to data movement that allows data to be replicated between Dell/EMC arrays at the same site or between remote sites. Data can be replicated using FC and FC over IP. EMC Replication Manager SE (RMSE) serves as the management interface, giving administrators a single point of control for local and remote replication operations. Using RMSE and SAN Copy, administrators can replicate a copy of the Exchange database and log files created by Microsoft Volume Shadow Copy Service (VSS) to a secondary disk array. The use of VSS helps ensure that a consistent point-in-time copy of Exchange is taken and replicated for extended business continuity.

EMC RepliStor® software provides an asynchronous, file-based approach to replication that can be configured to copy data from source to target systems (and vice versa) across a LAN, metropolitan area network (MAN), or wide area network (WAN). RepliStor runs directly on a server platform and allows clients to be redirected to access data from the target in the event of a failure at the source. RepliStor is well suited for small or branch-office Exchange environments using the Dell/EMC CX300 or AX100 storage array.

Dell business continuity services

After examining the various types of Dell hardware and supported software, organizations face the challenge of selecting the appropriate products and integrating them into their Exchange environments for maximum availability and resiliency. Dell business continuity professional services can help perform these steps—identifying objectives, suggesting products and technologies, and implementing systems that are designed to meet stipulated Exchange data availability and recovery requirements for organizations of all sizes.

Dell can provide end-to-end business continuity services such as conducting Exchange migrations, designing and deploying HA and Exchange clusters, and implementing data recovery and storage systems that use tools such as SAN Copy and SnapView. Dell can also provide disaster recovery and business continuity planning and services through joint offerings with SunGard Availability Services. In addition, Dell's partnerships with EMC and Microsoft can provide administrators with access to underlying technologies that enable Exchange environments to benefit from leading-edge features and capabilities without the operational headaches of managing numerous vendors.

Strategy for enabling a reliable e-mail environment

Organizations that have not fully assessed the business continuity risks to their Exchange environments may be exposed to significant vulnerabilities. Faced with the need to upgrade to Exchange Server 2003 while meeting stringent compliance regulations, many administrators are under increasing internal pressure to develop a robust business continuity plan. Creating such a plan means understanding the organization's tolerance for risk—whether it be data loss or downtime—and designing the appropriate approach based on those requirements. The layers of business continuity protection defined in this article can provide a good starting point for creating a business continuity plan.

Arun Taneja is the founder, president, and consulting analyst of the Taneja Group, an analyst and consulting group focused on storage and storage-centric server technologies. He has 25 years of experience in the IT industry, specifically in the areas of servers, operating systems, file systems, storage area networks, network attached storage, FC, Internet SCSI (iSCSI) and InfiniBand, clustering, and storage management software. He specializes in identifying technologies that have disruptive market potential and assisting companies with market positioning and strategy. Arun has a bachelor's degree in Technology from the Indian Institute of Technology, Delhi, and an M.S. in Electrical Engineering and an M.B.A. from the University of New Hampshire.

Alex Gorbansky is a senior analyst at the Taneja Group. Prior to the Taneja Group, Alex worked at EMC, where he was responsible for the product definition and launching of new enterprise software products; and at Loudcloud, a managed services provider, where he managed the storage and backup services business. Alex has a B.S. in Financial Engineering from Stanford University.

Enhancing Microsoft Exchange Migrations with VERITAS Enterprise Vault from Symantec

As organizations prepare to replace their legacy e-mail systems with the latest versions of Microsoft® Exchange, IT administrators must find a way to migrate these business-critical systems smoothly and efficiently. Most enterprises must address three major issues during the transition: considerable added infrastructure and resource costs when legacy e-mail and Exchange Server 2003 systems must coexist; extended changeover time, which can affect data availability; and heightened business risks given the mission-critical nature of the systems being moved. This article describes how VERITAS® Enterprise Vault® software can help organizations overcome such challenges.

BY SCOTT ROSEN

Related Categories:

Data archiving

Data center technology

E-mail technology

Microsoft Exchange

Storage

Storage software

VERITAS

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

Now that e-mail has become a business-critical application for many enterprises, finding ways to help reduce the risks associated with e-mail migration is a paramount concern. Administrators must manage several major areas of risk awareness during the transition, including data availability and the potential downtime of the core e-mail business system if a failure occurs during the changeover from a legacy e-mail system to Microsoft Exchange Server 2003.

Many organizations are planning to migrate to Exchange Server 2003, which is designed to streamline management and minimize infrastructure costs while leveraging the power of the Microsoft platform. But first, enterprises must plan a cost-effective, efficient way to contend with three significant challenges: considerable added infrastructure and resource costs when legacy e-mail and Exchange Server 2003 systems must coexist; extended changeover time, which can affect data availability; and heightened business risks given the mission-critical nature of the systems being moved. VERITAS Enterprise Vault software from Symantec can help administrators mitigate the risks inherent in an e-mail system migration, whether an organization is migrating from a legacy e-mail system or upgrading from an older version of Exchange to Exchange Server 2003.

Basic guidelines for migrating to Exchange Server 2003

A significant proportion of the time, effort, and costs associated with a migration project can be attributed to the physical volume of e-mail that must be transitioned. Decreasing the physical volume of data to be migrated can help minimize overall business risks as well as the coexistence time when both the old and new e-mail systems are running together, which can impose a major load on administrative and support resources.

When migrating from a legacy e-mail system, best practices advise administrators to focus on five main components:

- Mailbox profile
- Mailbox content
- Personal folder content
- Public folder content
- Address books, both personal and corporate

While the overall project can typically be managed using standard Microsoft Exchange or third-party migration tools, nearly all such tools can increase the volume of storage required. A typical scenario involves running parallel mailboxes in the legacy system and in Exchange Server 2003, which can double the amount of e-mail storage.

Even after the migration is complete, the amount of storage consumed is likely to be significantly higher as a result of the loss of *single-instancing,* which is often referred to as Single Instance Storage (SIS). Migration tools operate largely on a Messaging Application Programming Interface (MAPI) basis that includes no provision for single-instancing, which is usually provided through the Exchange message transfer agent (MTA). Migrated messages become unique, and as a result, the new e-mail environment typically requires greater storage capacity—in fact, some e-mail systems may require as much as two to three times more capacity than the previous e-mail system.[1]

Addressing the three core principles of controlling storage, minimizing administrative resources, and maintaining user transparency can provide a solid foundation for successful technology deployment.

Using Microsoft tools when migrating from a non-Exchange system to Exchange cannot solve this problem. Administrators can avoid an explosion of storage when migrating to Exchange Server 2003 by performing an in-place upgrade of the existing Exchange system. However, this approach requires system downtime and can be risky because all mailboxes must be converted at the same time. Consequently, administrators cannot adopt a phased approach; if a problem occurs, the whole process must be abandoned and the entire system reinstated.

Of course, administrators must consider the needs of the end users throughout the migration process. End users need uninterrupted access to the e-mail system and complete access to their personal e-mail knowledge base, and they should have a single point of access with no need to run parallel systems.

An overriding consideration for any migration or upgrade is to deliver the benefits of the new technology without introducing undue risk and unnecessary ongoing costs. Addressing the three core principles of controlling storage, minimizing administrative resources, and maintaining user transparency can provide a solid foundation for successful technology deployment.

Benefits of using VERITAS Enterprise Vault in Exchange migrations

Of the five main components administrators should focus on, as identified in the preceding section, VERITAS Enterprise Vault can benefit organizations most by helping to minimize the amount of mailbox content to be moved. By keeping the amount of mailbox content as low as possible, VERITAS Enterprise Vault can also help minimize the time required to perform migrations and enable administrators to keep the overall storage requirement under control, both for the migration phase and for the ongoing Exchange Server 2003 infrastructure. During this process, administrators should survey the amount of space that personal folder file content consumes and determine the impact that this file storage could have on the migration. Additionally, if an organization is considering transitioning from a non-Exchange system, Enterprise Vault can be used to help minimize storage capacity requirements during the migration to Exchange Server 2003, when both the non-Exchange system and Exchange Server 2003 must run together to help maintain system availability.

When migrating between different versions of Exchange, organizations can use Enterprise Vault before, during, and after the migration process to help minimize storage costs, migration time, and project risk. Enterprise Vault is designed to reduce the size of the Exchange message store before the physical migration occurs. By moving older items into a separate Enterprise Vault repository—which is Exchange-version independent and provides its own method of single-instancing and compression—administrators can minimize the amount of the content to be moved. Once in Enterprise Vault, data does not need to be converted when the organization moves to Exchange Server 2003. The data remains accessible to end users in the same seamless way it was before the migration. If required, the data can be restored to Exchange in the correct native format.

Enterprise Vault migration methods

Administrators can choose from four basic approaches to an Enterprise Vault–assisted migration: implementing Exchange Server 2003 without moving mailbox content, minimizing mailbox content to be moved, protecting the investment in Exchange Server 2003,

[1] For more information, see "Whatever Happened to Single-Instance Storage in Exchange?" by Jerry Cochran in *Windows IT Pro,* January, 18, 2002; www.storageadmin.com/Articles/Index.cfm?ArticleID=23819&pg=1&show=1373.

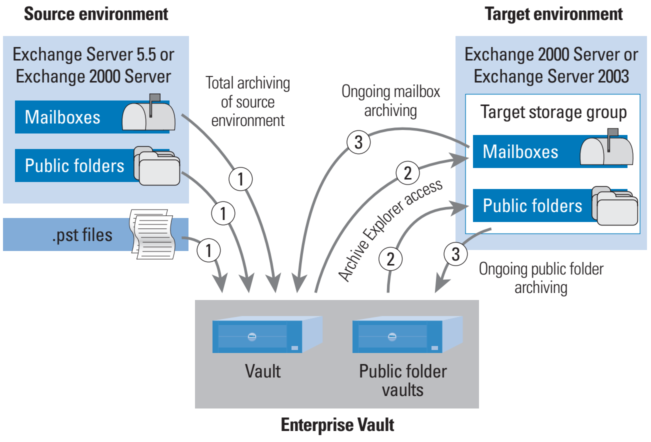

Figure 1. Migrating to Exchange Server 2003 without moving mailbox content

and using Enterprise Vault after migration. The choice of approach depends on how an organization views the role of Enterprise Vault—that is, as part of the migration project itself or as a separate project on its own. Each method is designed to significantly reduce the time, effort, risk, and cost involved because migrating the data represents a substantial portion of the migration project and Enterprise Vault helps minimize the amount of data that must be moved. This section examines the merits and considerations of these four methods.

Implement Exchange Server 2003 without moving mailbox content

Figure 1 shows an example scenario in which Enterprise Vault is deployed in a source environment running an older version of Exchange—for example, Exchange Server 5.5 or Exchange 2000 Server. In this environment, Enterprise Vault is used to archive all the e-mail messages from both public and private mailbox stores. At the same time, Enterprise Vault is also deployed in the target environment, which helps limit the migration project to primarily personal address books and mailbox profiles. The migration tasks include the following:

- Archive all content, including Microsoft Personal Folders (.pst) files, from the source environment. Migrate mailbox profiles and address books to the target environment.
- Provide access to archived mailbox and .pst content via Enterprise Vault Archive Explorer.
- Configure ongoing archiving in the target environment, with access to archived content via both Archive Explorer and shortcuts in mailboxes.

Cost savings can be achieved by providing end users with uninterrupted access to archived mail without requiring administrators to move that mail into Exchange Server 2003.

Minimize mailbox content to be moved

The most common way Enterprise Vault can be used in a migration is to minimize the amount of mailbox content that is physically moved across the two environments. Figure 2 depicts Enterprise Vault deployed in both the source environment and the target environment. In this scenario, Enterprise Vault can be used before migration to archive content from the mailboxes into the Enterprise Vault repository. Unlike the first method—which avoids moving mailbox content—in this scenario either all or a percentage of the content is archived from the source environment and replaced with seamless shortcut links in the mailboxes and public folders. This approach enables administrators to focus the data migration effort on moving the residual shortcuts and any content left behind.

The settings applied to this approach commonly archive any content older than 30 days. Residual shortcuts are left behind for the archived content or for any content that meets custom-defined criteria (for example, content that is up to one year old). Migrations previously performed by VERITAS Professional Services have shown that, on average, these policies can reduce the source mailbox and public folder content by as much as 80 percent and thus can significantly streamline the data migration effort—with the added benefit of providing seamless access to content archived from the source environment in the target mailboxes.

Protect the investment in Exchange Server 2003

If an organization has already begun its Exchange migration project or is migrating content from legacy e-mail systems, it may not be possible or appropriate to introduce a different technology into the legacy environment. In this case, Enterprise Vault can be introduced solely into the Exchange Server 2003 environment to help ensure best-practices mailbox management.

As noted in the section “Basic guidelines for migrating to Exchange Server 2003,” a significant side effect of migration to Exchange Server 2003 can be the loss of single-instancing,

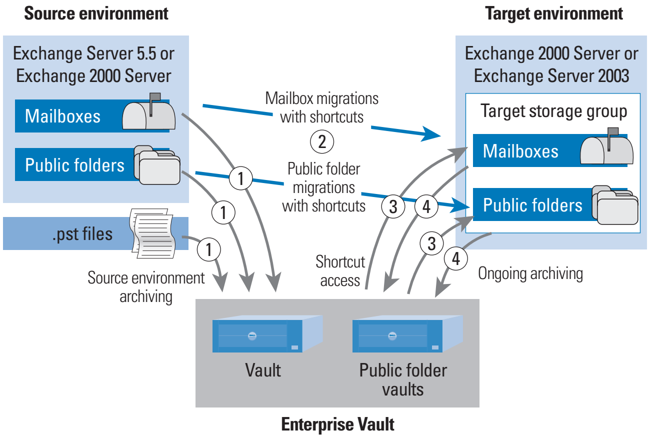

Figure 2. Migrating to Exchange Server 2003 by minimizing mailbox content

which can lead to a situation in which migrated data consumes more physical storage in the target environment than necessary. Enterprise Vault is designed to reduce the physical requirements for storing this data through archiving and the re-creation of lost single-instancing. The process can be seamless to end users because their original e-mail messages are replaced with shortcuts. Moreover, because this approach helps keep physical storage requirements low, it also helps minimize the associated costs of managing migrated content.

Minimize the size of Exchange databases after migration

Finally, Enterprise Vault can help organizations that have already completed their Exchange migration projects and, as a result, have larger private and public databases than before the migration, along with a corresponding increase in backup and recovery times.

In this case, the primary concern is to reduce the size of the Exchange databases quickly and, if necessary, limit their size to control unnecessary growth. The desired outcome is a defined service-level agreement for Exchange performance, a predictable backup and recovery strategy, and a reduction in ongoing storage and storage management costs. Exchange-initiated mailbox quotas may be used to limit mailbox sizes, but this approach can lead end users to create .pst files or delete information, which may increase the risk that important content such as corporate records will be lost. By introducing an archiving policy in conjunction with a mailbox quota, administrators can keep Exchange growth under control in a way that is unobtrusive to end users, preserving long-term access to important Exchange content. An example archiving policy using this model might constrain mailbox sizes by archiving at 75 percent of a mailbox quota of 100 MB, thus effectively limiting the Exchange database to 75 MB multiplied by the number of mailboxes, with an effective mailbox size governed by the amount of storage allocated to a mailbox archive.

Best practices for Enterprise Vault–assisted Exchange migrations

Determining the appropriate method for Enterprise Vault–assisted Exchange and legacy e-mail system migration may be governed by factors such as the following:

- Degree of data availability and system downtime the organization finds acceptable for such migration projects
- Availability of storage to address migrated e-mail content
- Availability of backup technology to address migrated e-mail content
- Time available to perform the migration process
- Status in terms of the migration project (not started, in progress, or concluded)
- Budget available for resources and software tools needed to perform the migration

When an organization implements VERITAS Enterprise Vault early in the planning stage of a migration project, the benefits of Enterprise Vault help justify the project expense and storage costs because they can enable a significant reduction in resource and management costs and a general reduction in overall project risk. The later that VERITAS Enterprise Vault is used in a migration project, the more its benefits derive from storage cost-savings—that is, an organization can move away from expensive local disks to cost-effective network attached storage, storage area networks, or other Dell and Dell/EMC storage options. The most common methods for using Enterprise Vault to assist in the Exchange migration process are to avoid moving mailbox content and to minimize the mailbox content being moved, as discussed in the section "Enterprise Vault migration methods."

Regardless of an organization's stage in a migration project, one type of mail content can benefit significantly from the use of Enterprise Vault: Microsoft .pst files. Enterprise Vault provides functionality to archive .pst content and help eradicate .pst files from the organization in an end-to-end process. This approach is designed to reduce the cost of the migration process by helping ensure that each stage is managed and controlled; to automatically determine .pst content ownership, which helps minimize the time taken to complete the process and also enables flexible security maintenance; and to allow a choice of server-based pull migration or client-based push migration—or a combination of both. Enterprise Vault also is designed to reduce risk and enable organizations to save time and money by migrating, repatriating, and consolidating .pst file content into an archive that is seamlessly accessible by Microsoft Windows® OS users.

An effective tool to help organizations move forward

Planning for a successful and error-free transition to Microsoft Exchange Server 2003 depends on several factors, but the actual migration process will not be a risk-free operation. Using VERITAS Enterprise Vault to assist in the management of Exchange content can help minimize infrastructure and resource costs, migration time and the migration's impact on data availability, and the risk of downtime that prevents access to business-critical e-mail systems. Enterprise Vault is designed to be a powerful tool to help organizations migrate to and take advantage of the next generation of e-mail systems.

Scott Rosen manages the Global Dell Relationship and Appliance Systems for the VERITAS Enterprise Vault suite, a product from Symantec Corporation. His focus is on building solutions with Dell technology. He has a degree in Organizational Psychology and Finance from the University of Michigan.

FOR MORE INFORMATION

VERITAS Enterprise Vault:
www.veritas.com/kvs

Deploying the McDATA 4314 Fibre Channel Switch Module for the Dell PowerEdge 1855 Blade Server

By integrating the McDATA® 4314 Fibre Channel switch module into the Dell™ Modular Server Enclosure, data center administrators can create a Fibre Channel storage area network (SAN) for the Dell PowerEdge™ 1855 blade server or expand their current SAN. Understanding the user-friendly features and monitoring capabilities of the McDATA 4314 Fibre Channel switch can help administrators simplify storage management in SAN environments.

BY RICHARD GOLASKY AND STEPHANIE HARTLEY

Related Categories:

Blade servers

Dell PowerEdge blade servers

Fibre Channel switches

Fibre Channel

McDATA

Storage

Storage architecture

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

The Dell PowerEdge 1855 blade server and the Dell Modular Server Enclosure offer many features that can benefit data center administrators. Given the capability to combine servers, Ethernet switches, systems management interfaces, and even a storage area network (SAN) fabric inside a single chassis, administrators no longer need to allocate large amounts of floor space for such data center components. The addition of a blade-sized Fibre Channel SAN interface—such as the McDATA 4314 Fibre Channel switch module—to a blade server enclosure enables administrators to easily integrate this SAN component into Fibre Channel infrastructures. Using the integrated McDATA 4314 in a Dell Modular Server Enclosure allows organizations to reduce the per-port cost of Fibre Channel connectivity compared to external Fibre Channel alternatives.

An internal Fibre Channel switch such as the McDATA 4314 is typically preferred over a Fibre Channel pass-through module interconnected to an external switch, except when a minimal number of server blades in a chassis are connecting to a SAN or when existing switches have numerous available ports without any potential future use. The internal switch is preferred over the pass-through module primarily because, when a pass-through module is used to connect to an external switch, the number of ports needed on the external switch is dependent on the number of server blades installed in the enclosure. When a system uses the switch module, a maximum of four ports can be used on the interconnected external switch, regardless of the number of server blades installed in the chassis.

The Dell Modular Server Enclosure supports the McDATA 4314 Fibre Channel switch module (FCSM). The McDATA 4314 is a 14-port switch that installs in module bay 3 of the enclosure for nonredundant fabrics, or in module bays 3 and 4 for redundant fabrics. An important cost-saving benefit is that the McDATA 4314 does not require separate licenses to enable additional Fibre

Channel ports; all 14 Fibre Channel ports are enabled for use when the switch is installed. Once the McDATA 4314 is installed in the enclosure and the Dell 2342 Fibre Channel host bus adapters are installed in the server blades, administrators can connect Fibre Channel storage devices such as the Dell PowerVault™ 136T tape library and Dell/EMC storage arrays to the PowerEdge 1855 blade server. This article explores the features of the McDATA 4314 FCSM when deployed in the Dell Modular Server Enclosure.

Operating mode

When introducing a McDATA Fibre Channel switch to a SAN environment, administrators must decide whether the switch should operate in Open Fabric mode or in McDATA Fabric mode. The McDATA 4314 is designed to operate in either mode; however, only the McDATA Fabric operating mode will be supported in the initial phase. The switch will default to McDATA Fabric mode, allowing for easy integration into an existing McDATA environment. The McDATA Fabric mode provides options—such as zoning methods—that are not available in the Open Fabric mode. When operating in the McDATA Fabric mode, the McDATA 4314 can join fabrics that consist of exclusively McDATA switches. The McDATA 4314 is also well suited for administrators who are building their first SANs.

Management interface

The McDATA 4314 provides three management interfaces to monitor switch components:

- Enterprise Fabric Connectivity Manager (EFCM) Basic (a storage network management application)
- Command-line interface (CLI) shell
- Embedded Web server

These interfaces are critical for managing day-to-day operations on a SAN. The primary management interface to the McDATA 4314 Fibre Channel switch is the EFCM Basic application. The EFCM Basic graphical user interface (GUI) provides user-friendly utilities to monitor and configure switch components, including:

- **Temperature:** Continuously monitors switch temperature
- **Performance:** Displays performance throughput for all 14 ports
- **Link and port statistics:** Displays information about port and link errors
- **Simple Network Management Protocol (SNMP) alerts:** Provides monitoring and trap functions (SNMP v1 and v2)
- **Alert thresholds:** Triggers alerts when ports exceed predefined error thresholds
- **Security:** Provides user account and authentication support using Secure Sockets Layer (SSL) and Secure Shell (SSH) protocols

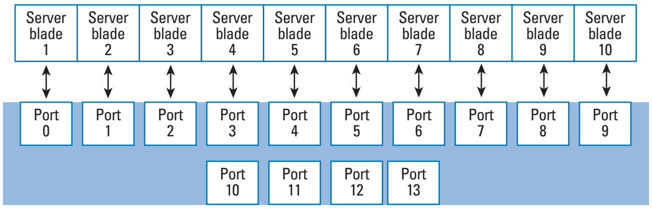

Figure 1. Server blade and switch port mapping

For systems on which EFCM Basic is not installed, the McDATA 4314 provides an embedded Web server that allows management from any Web browser. The embedded Web server provides the same functionality as EFCM Basic except for performance monitoring and the zoning wizard.

The McDATA 4314 also provides a CLI to manage the switch. Users can telnet into the switch, and then engage the CLI to invoke actions that could otherwise be done using EFCM Basic.

One key feature of the McDATA 4314 switch is McDATA's Hot Code Activation Technology (HotCAT®) feature, which provides the mechanism to invoke a switch reset without interrupting I/O operations. This procedure is also known as a Non Disruptive Code Load and Activation (NDCLA) and is invoked during and after a firmware upgrade. The McDATA 4314 switch's HotCAT feature is designed to upgrade firmware on the switch and then perform a hot reset without affecting data movement or incurring downtime.

> One key feature of the McDATA 4314 switch is the HotCAT feature, which provides the mechanism to invoke a switch reset without interrupting I/O operations.

To monitor link performance, EFCM Basic provides a powerful performance viewer applet to monitor the transmit links and the receive links on each Fibre Channel port. The applet can be easily customized for viewer preference.

Topology configuration

The 14 Fibre Channel ports on the McDATA 4314 include 10 internal ports for connection to the 10 server blades in the Dell Modular Server Enclosure and four external ports for connectivity to the SAN infrastructure. Because each of the 10 internal ports connects directly to a server blade via the enclosure's passive midplane, administrators cannot directly access these ports. However, administrators may configure the speed and topology settings by using EFCM Basic. Figure 1 shows the physical mapping between each server blade and the internal Fibre Channel ports of the McDATA 4314.

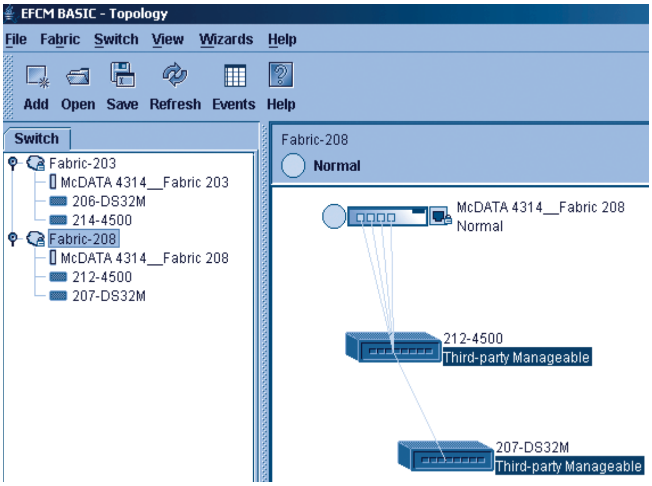

Figure 2. Using the EFCM Basic GUI to display ISLs between switches

The four external Fibre Channel ports provide direct connectivity to target devices such as Dell/EMC storage arrays and tape devices, or connectivity to other McDATA Fibre Channel switches operating in the McDATA Fabric mode.

Although the McDATA 4314 does not include an integrated trunking feature, it does support Inter-Switch Links (ISLs). Using the four external ports, an administrator can connect up to four ISLs to other McDATA switches. The ISLs help provide fault tolerance as well as load balancing between the switches. A loss on one of the ISLs will force data from the failed link to be routed to the remaining links. When connecting switches in Dell/EMC storage environments, administrators are limited to using a maximum of four interconnect switches, or *hops,* between the blade server and the target device, such as a Dell/EMC storage array. The four hops refer to the McDATA 4314 and three additional switch interconnects.

In Figure 2, the EFCM Basic GUI displays the physical layout of a fabric involving Fibre Channel switches that are linked together through ISLs. In this example, the four external ports of the McDATA 4314 are all linked to a McDATA Sphereon 4500 switch via ISLs. The Sphereon 4500 is then linked, via ISLs, to an EMC DS-32M switch. The four links between the McDATA 4314 and the Sphereon 4500 are green, to indicate four active ISL connections. If any of the links on the McDATA 4314 fail, then the link color will change from green to red and an alert will be generated.

Each Fibre Channel port on the McDATA 4314 can be configured as the following port types at 1 Gbps or 2 Gbps:

- G_Port
- GL_Port
- F_Port
- FL_Port
- E_Port (used for expansion ports)

Fibre Channel standards compliance

The McDATA 4314 is compliant with the following Fibre Channel standards:

- Fibre Channel Physical and Signaling Interface (FC-PH) Revision 4.3
- FC-PH-2
- FC-PH-3
- Fibre Channel Fabric Generic Requirements (FC-FG)
- Fibre Channel Switch Fabric 2 (FC-SW-2)
- Fibre Channel Generic Services (FC-GS)
- FC-GS-2
- FC-GS-3
- Fibre Channel Arbitrated Loop (FC-AL) Revision 4.6
- FC-AL-2 Revision 7
- Fibre Channel Fabric Loop Attachment (FC-FLA)
- FC-TAPE
- Fibre Channel Virtual Interface (FC-VI)

In addition, the McDATA 4314 adheres to the Fibre Channel Element MIB (management information base) Specification (FE-MIB) and the Fibre Alliance MIB Specification (FA-MIB). It also supports Class 2, Class 3, and Class F service.

Fabric zoning

A key component of the McDATA 4314 is its user-friendly interface for zone configuration. Zoning allows ports (initiators and targets) to be grouped together or masked from each other. In a fabric connecting hundreds or thousands of devices, the task of grouping devices can become tedious. To zone devices, administrators must know the World Wide Name (WWN) of each device and then group the device objects into a zone. Using the McDATA 4314 EFCM Basic SAN management application, administrators can add device objects to a zone simply by dragging and dropping. Once zone sets are created, administrators can run an error-checking procedure to verify that no conflicts have been created by newly created zones.

By using aliases, administrators can group several ports together as a single object, thus reducing the number of times they have to drag and drop ports onto a zone.

To help simplify zoning methods, the fabric zoning interface enables administrators to create aliases and nicknames. By using aliases, administrators can group several ports together as a single object, thus reducing the number of times they have to drag and drop ports onto a zone.

Maximum zone sets	256
Maximum zones	2,000
Maximum aliases	2,500
Maximum number of members per zone	2,000

Figure 3. Zoning database limits for the McDATA 4314 switch

Nicknames can be created to help simplify viewing of port devices. By using nicknames, administrators can associate a “friendly” name with a port’s WWN instead of having to refer to a device by its “unfriendly” WWN.

The zoning database limits the number of entries that can exist. Figure 3 defines these categories and limits.

In Figure 4, the left-hand panel shows the zones and aliases that were created by dragging and dropping individual objects from the right-hand pane. The right-hand pane contains objects associated with targets and initiators. Ports 0 through 9 represent the 10 server blades, and the remaining ports represent the four external ports.

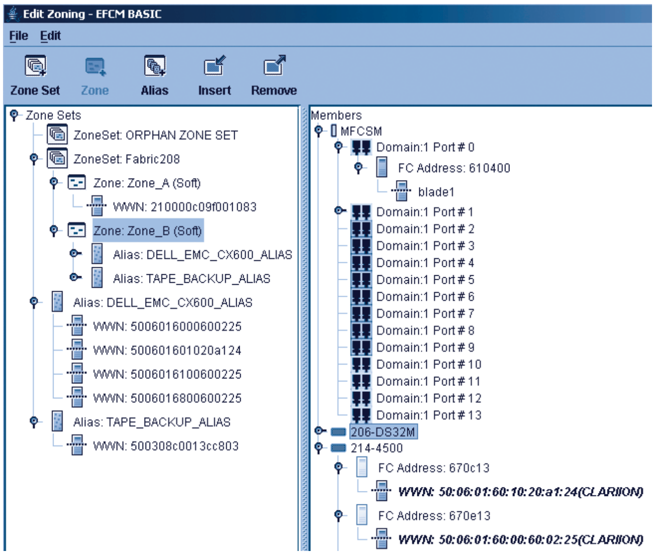

```
Zone Sets
    ZoneSet: ORPHAN ZONE SET
    ZoneSet: Fabric208
        Zone: Zone_A (Soft)
            WWN: 210000c09f001083
        Zone: Zone_B (Soft)
            Alias: DELL_EMC_CX600_ALIAS
            Alias: TAPE_BACKUP_ALIAS
    Alias: DELL_EMC_CX600_ALIAS
        WWN: 5006016000600225
        WWN: 500601601020a124
        WWN: 5006016100600225
        WWN: 5006016800600225
    Alias: TAPE_BACKUP_ALIAS
        WWN: 500308c0013cc803
```

```
Members
MFCSM
    Domain:1 Port # 0
        FC Address: 610400
            blade1
    Domain:1 Port # 1
    Domain:1 Port # 2
    Domain:1 Port # 3
    Domain:1 Port # 4
    Domain:1 Port # 5
    Domain:1 Port # 6
    Domain:1 Port # 7
    Domain:1 Port # 8
    Domain:1 Port # 9
    Domain:1 Port # 10
    Domain:1 Port # 11
    Domain:1 Port # 12
    Domain:1 Port # 13
206-DS32M
214-4500
    FC Address: 670c13
        WWN: 50:06:01:60:10:20:a1:24(CLARIION)
    FC Address: 670e13
        WWN: 50:06:01:60:00:60:02:25(CLARIION)
```

Figure 4. Creating zones and aliases for components within a McDATA 4314 switch fabric

Extended credits

Each Fibre Channel port on the McDATA 4314 switch supports a data buffer limit of 16 credits, with each credit supporting a maximum frame size of 2,148 bytes. In most Fibre Channel environments—at speeds of 1 Gbps and 2 Gbps—a credit limit of 16 is sufficient for the distances that data will travel. However, when optical cable distances exceed 500 meters, performance degradation can occur, thus decreasing the speed and efficiency at which data travels. The McDATA 4314 switch module is designed to compensate for these conditions by allowing each Fibre Channel port to “donate” a maximum of 15 credits to a pool from which needed ports can “borrow.” Increasing the buffer credits on a port allows additional data frames to be sent without requiring the switch to wait for acknowledgments from the end device—thereby helping to minimize the impact on long-distance performance.

> When incorporated into a Dell Modular Server Enclosure, the McDATA 4314 Fibre Channel switch module enables administrators to integrate server blades quickly and easily into a SAN environment.

Integrated SAN management

When incorporated into a Dell Modular Server Enclosure, the McDATA 4314 Fibre Channel switch module enables administrators to integrate server blades quickly and easily into a SAN environment. By using the EFCM Basic storage network management application provided with the McDATA 4314 to configure and monitor switch operations, administrators can be informed when a failure occurs within the switch. In addition, the easy-to-use GUI makes zoning a convenient process by allowing nicknames, aliases, and drag-and-drop methods. Target devices such as Dell tape libraries and Dell/EMC storage arrays can be easily connected directly to the McDATA 4314 switch or to an external switch linked via ISLs. This ability to integrate into existing Dell server and storage infrastructures—as well as easy-to-use management capabilities—are designed to make the McDATA 4314 Fibre Channel switch a powerful addition to SAN environments.

Richard Golasky is a development engineer senior consultant in the Dell/EMC Enterprise Storage Group. His responsibilities include all areas of tape backup, including Fibre Channel storage, high-availability clusters, network attached storage, and performance analysis. He has a B.S. in Electrical Engineering from Florida Atlantic University.

Stephanie Hartley is a consultant development engineer in the Dell/EMC Enterprise Storage Group. Her responsibilities include Dell PowerConnect™ switch development, Fibre Channel storage, and program management.

FOR MORE INFORMATION

McDATA switches:
www.mcdata.com

Dell networking devices:
www.dell.com/networking

Streamlining Server Management

with Boot-from-SAN Implementations

Booting servers from storage area networks (SANs) can provide significant benefits for today's complex data center environments. One of the driving forces behind SANs is the need to deliver mission-critical data quickly—at any time, without interruptions or delays. A key determinant in meeting that requirement is how quickly administrators can replace a failed server in a SAN environment. To that end, this article discusses systems management benefits and complexities associated with boot-from-SAN implementations.

BY MATTHEW BRISSE AND AHMAD TAWIL

Related Categories:

Dell/EMC storage

Fibre Channel

Servers

Storage

Storage architecture

Systems management

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

A storage area network (SAN) is designed to attach storage devices, such as intelligent storage arrays and tape libraries, to servers. Until recently, SANs were usually reserved for Fibre Channel networks, but with the introduction of Internet SCSI (iSCSI), SAN usage has now expanded to include IP networks as well as Fibre Channel networks. This article focuses on Fibre Channel–based SANs; however, the methods and challenges discussed apply to any fabric-based SAN.

Although network servers traditionally boot their OS from a local disk, today's rapidly evolving technologies enable servers to boot from a storage array located on a SAN. As a result, if a server needs to be taken offline or replaced, the network administrator can simply replace that server in the SAN with another server that has the same configuration, direct the storage volume to the replacement server, and boot the replacement server from the external storage array.

An important distinction must be made between booting from a SAN and using diskless servers. The two techniques present different capabilities and challenges. For example, a server booting from a SAN could contain disks for the page and swap files. Conversely, a diskless server could boot from a locally attached storage array.

Exploring the benefits of booting from SANs

Booting from SANs can help streamline systems management. Separating the boot image from the server enables administrators to leverage the high availability, data integrity, and storage management of the storage network. Booting servers from SANs enables far-reaching enterprise benefits including enhanced disaster tolerance, centralized administration, minimized cost of ownership, high-availability storage, enhanced business continuance, rapid server repurposing, and image management consolidation.

Enhanced disaster tolerance. If a server becomes faulty, unavailable, or compromised in a boot-from-SAN environment, it can be swapped out for a replacement server with the same configuration right away. This approach enables administrators to deploy the replacement server without first being required to reconstruct the booting, operating, and application environments. Administrators do not have to swap out hard drives, reconfigure arrays, or restore data and applications from backup—saving valuable time during a mission-critical restore. The capability to replace and add servers in minutes can enhance return on investment (ROI) by helping to prevent lengthy downtime of core servers—also helping enterprises to avoid service interruptions in the mission-critical applications that run on those servers.

Centralized administration. Transitioning from a distributed server environment, where the boot device is local to the server, to one or more storage arrays on a SAN helps reduce the management overhead associated with traditional distributed environments. Centralized systems management can streamline administrative workflow and typically requires fewer administrators to maintain the environment.

Minimized cost of ownership. Booting from a SAN enables the deployment of diskless servers. Depending on the network topology and number of servers, organizations can gain significant savings by using diskless servers. For example, 100 servers accessing redundant 250 GB drives can share a single storage array, using a fraction of the number of disks that would be required to support the same number of servers if each server were equipped with its own local disk storage.

High-availability storage. Booting from a SAN usually involves booting from an external RAID controller. Typically, external RAID controllers enable high availability. Adding the multipathing features of a Fibre Channel host bus adapter (HBA), which is required when booting from a SAN, helps provide even higher availability through redundancy. This benefit also enhances ROI because RAID controller costs can be amortized over multiple servers.

Enhanced business continuance. Data protection features inherent in SANs—such as backup and restore, data migration, data replication, and disk-capacity expansion features—can be effectively used for the boot drive on the RAID controller without incurring additional cost.

Rapid server repurposing. By storing the boot image on a SAN, administrators can create an application instance that includes the logical grouping and associations of a server model, installed applications, the OS image, the actual boot logical unit (LUN), and the data LUNs. A preconfigured application instance can be used to rapidly repurpose a server to perform a specific job at a scheduled time (for example, a backup server) or dynamically repurpose the server to sustain the workload of a specific job (for example, a database server). This instance can also be used for disaster recovery to quickly replace a failed server.

Image management consolidation. Storing boot images on a SAN enables consolidation of homogeneous boot images on the same LUN in the RAID unit, as well as consolidation of incremental changes on a snapshot LUN.

Choosing applications that are well suited to boot from SANs

Applications that require high data availability are good candidates for boot-from-SAN implementations. Examples of such applications include but are not limited to:

- Microsoft® Exchange servers as well as database, e-mail, and Web servers located in data centers or server farms
- Commerce servers that handle revenue-critical applications
- Enterprise servers that provide access to critical network resources and services
- Blade server installations, in which small servers built on cards plug into stackable racks that offer shared power supplies and networking capabilities, allowing administrators to add servers quickly and cost-effectively on an as-needed basis

Defining boot-from-SAN standards

Today, no standards are associated with booting from SANs—either programmatically or in best practices. However, the Storage Networking Industry Association (SNIA) has formed a Host Technical Working Group (TWG) to address the needs and technological challenges associated with booting from SANs. These standards groups are dedicated to help minimize costs and maximize interoperability of boot-from-SAN environments. Dell and Sun Microsystems are co-chairs of the Host TWG; other active participants include HP, Adaptec, Engenio, AppIQ, and VERITAS. The SNIA Host TWG focuses on host-side storage management and is responsible for generating Common Information Model (CIM) profiles and subprofiles that address storage resources associated with the host, which include boot-from-SAN implementations.

Currently, vendors must provide systems management functionality to manage disks, RAID, enclosures, and boot methodologies. This adds complexity to systems management, especially in heterogeneous environments—which can increase the likelihood of errors and the consequent financial burden on a data center that must safeguard against and recover from errors that may occur in a complex heterogeneous environment. Standardized boot-from-SAN practices can help streamline systems management and help administrators avoid common problems associated with such configurations.

Understanding how booting from SANs works

Administrators can enable a server to boot from a SAN by configuring the server with its own virtual boot device. The boot device allows the server to fetch boot-loading instructions from a storage array and execute them. The boot device is mapped as a LUN on the storage array. Depending on an organization's requirements, all LUNs may boot servers with the same profile. Alternatively, a storage array may have some LUNs mapped to servers running a Microsoft Windows® OS, other LUNs mapped to servers running a Linux® OS, and still other LUNs mapped to servers running a UNIX® OS. In this scenario, each server is configured to point to the appropriate LUN at the specific RAID controller from which the server is to boot.

Setting up the LUN to serve as a boot disk occurs at the Fibre Channel HBA. For example, an administrator can use the following procedure to configure a LUN for a server with a QLogic SANblade HBA:

1. Choose "Selectable Boot" from the host options provided by the HBA.
2. Select a Fibre Channel RAID port from which to boot.
3. Select the LUN associated with the selected RAID port.

Organizations with fully redundant and highly available SAN configurations may have two HBAs installed in each server attached to the SAN. In such installations, this procedure is repeated for the second HBA so that it points to the same LUN via the same RAID controller or a different RAID unit.

After performing the preceding procedure, the administrator should power up the server, which will cause the following sequence to occur:

1. During the boot process, the Extension ROM in the Fibre Channel HBA is initialized.
2. The Extension ROM in the Fibre Channel HBA finds the boot LUN at the Fibre Channel RAID unit and configures it as a boot device.
3. A booting protocol tells the server BIOS (or equivalent) to scan the Fibre Channel bus for the appropriate LUN for the OS image.
4. When the server locates the appropriate LUN, the server CPU executes the OS boot image.

Identifying OS dependencies

Various operating systems require that certain guidelines be followed to allow servers to boot from a SAN. Understanding these requirements is key to the successful deployment of a boot-from-SAN environment within a data center. This section describes the requirements that system administrators should consider.

Microsoft Windows dependencies

Before servers boot from a SAN in a Microsoft Windows environment, the following conditions must be met.

Segregated hosts. The SAN must be either configured in a switched environment or directly attached from each host to one of the storage subsystem's Fibre Channel ports. Fibre Channel Arbitrated Loop (FC-AL) is not supported, because it does not permit hosts attached to the SAN to be segregated from each other sufficiently.

LUN 0. The Microsoft Windows NT® and Windows 2000 operating systems both require the storage array LUN to appear as LUN 0 to the server. The Microsoft Windows Server™ 2003 OS does not have this restriction.

Exclusive access to the bootable disk. The host must have exclusive access to the logical disk from which it is booting. No other host on the SAN should be able to detect or have access to the same logical disk. This can be accomplished by using a LUN management facility such as LUN masking, zoning, or a combination of these methods.

Microsoft Windows Server 2003. Typical issues associated with boot-from-SAN environments range from not booting from the correct partition to page-file retry errors resulting in an OS crash. A commonly known problem with diskless servers within the Microsoft OS environment is that the OS may crash if an extended delay in writing to the page or swap file occurs. Windows Server 2003 is designed to extend the retry time-out values and perform up to three retry attempts. Note that previous versions do not have retry capabilities or the ability to extend the time-out values associated with the page and swap files. Failover between HBAs is not available for the boot LUN. Microsoft documentation can guide administrators through the setup and challenges associated with booting from a SAN.[1]

The default behavior for a Windows OS is to attach and mount all LUNs that it detects when an HBA driver loads. However, this process causes one of the most common problems that administrators encounter when configuring a SAN: multiple hosts have access to the same logical disk—a situation that will likely damage the file system. To help avoid this problem, administrators should follow proper planning and SAN management practices.

Latency issues associated with a page file can be difficult to diagnose. Problems can range from slow response times to system hangs to crashes. Administrators should begin troubleshooting by investigating the system event log (SEL). An event ID of 51 in the SEL may indicate a latency problem. Event 51 could indicate that the

[1] For Microsoft documentation about booting from SANs, visit download.microsoft.com/download/f/9/7/f9775acc-baa6-45cc-9dec-b82983705620/Boot%20from%20SAN%20in%20Windows.doc and support.microsoft.com/default.aspx?scid=kb;EN-US;q305547.

memory manager was attempting to copy data to or from memory and experienced a problem. For example, the following could appear in the Microsoft SEL:

```
Event ID: 51
Event Type: Warning
Event Source: Disk
Description: An error was detected on device \device\
    harddisk0\DR0 during a paging operation

Event ID: 11
Source: %HBA_Driver_Name%
Description: The device, \device\ScsiPort0, did not respond
    within the time-out period
```

Another indicator is a blue screen, which signifies an OS crash. If administrators see the following messages, then latency issues may be the culprit:

```
0x00000050 PAGE_FAULT_IN_NONPAGED_AREA
or
0x0000000A IRQL_NOT_LESS_OR_EQUAL
```

A quick resolution to the preceding types of problems is to place a disk within the server for page and swap files. This approach can help prevent access from being influenced by devices, fabrics, or hosts on the SAN.

Linux dependencies

Linux has many of the same dependencies as Microsoft Windows with regard to host segregation, LUN 0, and exclusive access to the bootable disk. However, Linux does not have the same latency sensitivities associated with page and swap files as Microsoft Windows because the Linux kernel is not a pageable OS—that is, it does not swap for kernel operations, but it will swap for applications. Linux pages applications in and out after the boot sequence. If latency occurs when accessing the LUN (more than 60 seconds without response), SCSI time-out messages similar to the following will appear on the console and possibly in the system log:

```
kernel: scsi : aborting command due to timeout : pid
    50212683, scsi0, channel 0, id 0, lun 0 Read (10) 00 00
    48 01 8f 00 00 80 00
kernel: scsi: device set offline - command error recover
    failed: host 0, channel 0, id 0, lun 0
kernel: SCSI disk error: host 0, channel 0, id 0, lun 0,
    return code = 6000000
kernel: I/O error: dev 08:11, sector 22903784
```

If the system is configured with multipathing HBA drivers, a failover should occur before a kernel panic is initiated. Linux does not automatically mount file systems on all discovered LUNs as Windows does. By default, the Linux OS accesses only the LUNs specified in the /etc/fstab file. Most Linux file systems are not cluster aware, so administrators should not mount the same LUN and file system on two servers simultaneously because file corruption may occur.

A properly designed SAN can help minimize the possibility of a kernel panic. Data center evaluations and SAN consulting can dramatically reduce the current risks associated with boot-from-SAN time-out issues.

Recovering from failure in a boot-from-SAN environment

Booting from a SAN is a powerful capability, and in a disaster recovery scenario, it is critical. Thus, administrators must know how to replace a mission-critical server in a boot-from-SAN environment. To perform this task, administrators should adhere to the following best practices:

- When replacing a server, administrators should remove the HBA from the failed server, install it in the replacement server, and configure the system BIOS in the replacement server to boot from the Fibre Channel HBA.
- If the replacement server is a different model from the one it is replacing, either the OS will reconfigure itself and prompt the administrator to insert the driver CD for the replacement hardware, or the administrator will have to manually install the driver for the replacement hardware. The OS will not require the reconfiguration of applications such as Microsoft Exchange Server.
- If an HBA is replaced by an HBA of the same model or from the same vendor, administrators should update the access rights on the RAID ports to reflect the change in the World Wide Name (WWN) of the replacement HBA. The WWN of the Fibre Channel HBA can be found by accessing the HBA's ROM extension utility during system power-up. In addition, administrators must configure the HBA to boot the server from the appropriate LUN.
- If the replacement HBA is from a different vendor, either the OS will reconfigure itself and prompt the administrator to insert the driver CD for the replacement hardware, or the administrator will have to manually install the driver for the replacement HBA.
- For x86-based systems, if a RAID controller fails during the reboot of a server containing two HBAs, the designated boot LUN may not be visible on the HBA's primary path. In this case and depending on the system BIOS design, the system may boot from the path configured

on the second HBA—or manual intervention may be required to change the system BIOS to boot from the secondary HBA.

Considering other troubleshooting issues

Boot time is the amount of time a storage system requires to boot. This time is a function of the number and types of LUNs in the storage system, and it is typically less than five minutes. If a server and a storage system are powered up (cold start) at the same time, the server's extended BIOS may scan the Fibre Channel bus before the storage system is ready to respond. When this occurs, the server will not be able to find the boot LUN. To avoid this problem, administrators must either power up the storage system and fabric before the server or, if the server is so equipped, set it to delay the scan until the storage system is ready. Note that the Microsoft Windows NT and Windows 2000 platforms require clustered servers to keep boot disks on data paths that are separate from shared-storage paths—thus requiring one dedicated HBA for booting and another HBA for shared disks.

In a boot-from-SAN environment, replacing a failed server can be time-consuming; this process is not as easy as simply swapping out one server for another. Administrators should transfer the failed server's HBA to the replacement server or, if the HBA itself has failed, use an identical HBA as the replacement. If the failed server was using an internal disk for the paging file system, administrators can move this disk to the replacement server. Alternatively, administrators can use a replacement disk, but they must format it and manually reset the paging to that disk.

Booting from SANs for enhanced systems management

Understanding the complexity behind any systems management architecture or technology can help administrators optimize data center operations and efficiency. By examining the advantages and potential issues that may arise when implementing boot-from-SAN configurations, administrators can maximize the benefits of this technology—including enhanced reliability, high availability, and streamlined systems management.

Matthew Brisse is a technology strategist in the office of the Dell CTO and is vice chair of the SNIA board of directors. At Dell, Matthew chairs the Standards Review Board and is a member of the Systems Management Architecture team.

Ahmad Tawil is a technology strategist in the Storage Architecture and Technology Group within the office of the Dell CTO, where he focuses on future networked I/O technologies.

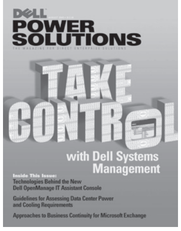

Link-Level Error Recovery with Tape Backup

Error recovery in a Fibre Channel environment is crucial for maintaining data integrity during periods of link errors, noisy lines, bit errors, lost or corrupted data frames, or even SCSI commands that time out. Depending on how data integrity problems occur, Fibre Channel architecture implements processes for error detection and recovery. According to the Fibre Channel specifications, a bad bit may appear every 1,000 seconds during 1 Gbps data transmission. The FC-TAPE technical report addresses how error detection and recovery processes apply to tape backup devices.

BY RICHARD GOLASKY

Related Categories:

Fibre Channel

Storage

Tape library

Troubleshooting

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

FC-TAPE is a technical report from Technical Committee T11 within the InterNational Committee for Information Technology Standards (formerly NCITS).[1] This report defines a profile of the features and options required to operate streaming devices and media changers in a Fibre Channel environment. For example, FC-TAPE defines the standards that are required to operate streaming devices, such as the Dell™ PowerVault™ 132T, PowerVault 136T, and PowerVault 160T tape libraries, in a Fibre Channel environment.

Key areas referenced by FC-TAPE include the Fibre Channel Protocol-2 (FCP-2) and Fibre Channel Framing and Signaling (FC-FS) specifications for link error recovery with tape backup devices. Unlike tape backup software applications, which conduct error recovery at the host application level, devices that comply with FC-TAPE conduct error detection and recovery at the link level, between the Fibre Channel host bus adapter (HBA) and the target device, such as a Fibre Channel–SCSI bridge or a native Fibre Channel tape backup device (see Figure 1).

Handling backup errors with tape backup software

Tape backup devices operating in a Fibre Channel environment do not recover from errors as well as disk storage systems do. Tape backup units read and write data in sequential order, one block at a time. Data on tape will be

[1] For more information about Technical Committee T11 and the FC-TAPE report, visit www.t11.org.

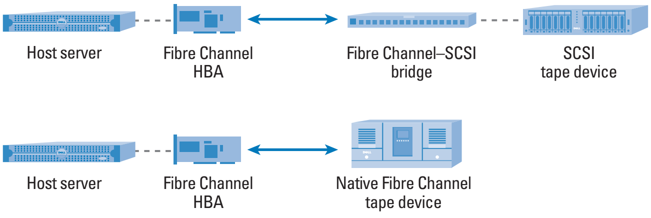

Figure 1. Link level at which error recovery occurs with FC-TAPE

corrupted if any of the data blocks are written out of sequence. Once data is written incorrectly to tape, it cannot be corrected—hence the need to prevent the errors from occurring. In most cases, tape backup operations are interrupted because of an external event such as a target reset, misbehaving hardware devices, a SCSI command time-out, or a faulty link. Mechanisms are in place to detect and recover from many types of errors, and one such non–Fibre Channel mechanism is the error recovery process built into tape backup software applications.

When a data transmission fails, a well-written tape backup software application can perform some error recovery with a tape backup device. However, such error recovery is often limited by the type of error and can take several minutes. Although error recovery techniques are proprietary to each software vendor, a general process is implemented to identify the problem and attempt to recover from it. If an error is encountered, the tape software application will rely on a SCSI check condition and sense data from the tape target. The tape software will then identify the current block position and compare it to the expected block position. As part of the error recovery process, tape backup software applications are designed to verify whether all data was successfully transferred to tape by reading back the last data transfer. Once this process is complete, the tape software will continue with the tape backup operations.

Although error recovery can be performed effectively by tape backup software, this approach is a time-consuming process that often takes up to 10 minutes. This may appear to be an acceptable amount of time to administrators, but it is an extremely long time compared to the error recovery time frame that an FC-TAPE device implements, which is typically less than five seconds. With FC-TAPE, error detection and recovery operations are conducted at the link level (between the HBA and the tape device) and do not involve OS components such as the port driver. This approach is designed to make FC-TAPE error recovery and data retransmission more efficient than possible with tape backup software.

Handling backup errors with FC-TAPE devices

The error recovery process for FC-TAPE devices occurs at the link level. During a Fibre Channel process login (PRLI) extended link service request, ports will exchange service parameters and declare the capabilities of each device. Inside the parameter field of the PRLI request are two entities related to recovery: the RETRY bit and the Confirm Completion Allowed bit. If the RETRY bit is set, this indicates that the request initiator supports retransmission of data frames, or that the target has the ability to request retransmission of data. Once the RETRY bit is set, two functions are enabled to assist in error recovery: Read Exchange Concise (REC) and Sequence Retransmission Request (SRR).

Read Exchange Concise function

The REC is an extended link service that requests additional information on an open exchange mainly based on the initiator's Source Identifier (S_ID) and the Originator Exchange Identifier (OX_ID), which are specified in the payload of the REC. A REC is transmitted by a device whenever the time-out value (REC_TOV) has expired. The REC_TOV is defined as one second more than the Error Detect time-out value (E_D_TOV). Usually, the E_D_TOV default is two seconds; therefore, the REC_TOV is set to three seconds. When the REC_TOV timer expires after three seconds, a REC is sent to the target device requesting information on the open exchange. For example, a SCSI command may take longer than three seconds to get a response and thus the REC_TOV timer would expire. If the target accepts the REC, it will respond with an accept (ACC) indicating the status of the exchange, via its payload. Otherwise, the target will reject the REC with a link service reject indicating why the REC was rejected.

Data on tape will be corrupted if any of the data blocks are written out of sequence. Once data is written incorrectly to tape, it cannot be corrected—hence the need to prevent the errors from occurring.

Note that a failed data transmission is not the only reason a REC can be transmitted. Tape backup operations that take longer than the REC_TOV timer also cause the initiator to transmit a REC. The target will respond with an ACC indicating that the sequence is still in progress. The initiator will then transmit a REC every three seconds (the REC_TOV) until the operation is complete. Examples of this scenario include time-consuming commands such as a long erase to tape drives and media movement within libraries.

The ACC payload, in response to a REC, contains two word values indicating the current status of the exchange: the data transfer count (DTC) and the exchange status block (E_STAT). The DTC is the number of bytes successfully received by the device during a write operation, or the number of bytes transmitted by the

> The extended link service Sequence Retransmission Request is sent by the initiator (the host) to the target device to request that a target retransmit data back to the initiator, or to inform the target that data will be retransmitted back to the target. The invocation and acceptance of an SRR indicate that a previous command or data frame was lost.

target during a read operation. The DTC is important in determining whether all data was successfully transmitted or received. The E_STAT word value has two bit values that identify the state of the exchange. The first value is the Completion bit, which specifies whether the exchange is still open; and the second value is the Sequence Initiative bit, which indicates which port holds the initiative (that is, which end of the transmission is supposed to make the next move).

Sequence Retransmission Request function

The next part of FC-TAPE error recovery involves the extended link service SRR. This service is sent by the initiator (the host) to the target device to request that a target retransmit data back to the initiator, or to inform the target that data will be retransmitted back to the target. The invocation and acceptance of an SRR indicate that a previous command or data frame was lost.

The SRR payload consists of an OX_ID, a Relative Offset value, and a Routing Control (R_CTL) value, all of which assist in determining which data needs to be retransmitted. The Relative Offset field identifies an offset value that the initiator requires for data retransmission, and the R_CTL field defines the frame function type that is to be retransmitted. The R_CTL value identifies whether the retransmission is a data descriptor (FCP_XFER_RDY), command status (FCP_RSP), or solicited data (FCP_DATA). If the target accepts the SRR, the target will respond with an ACC, which will then allow the data to be retransmitted.

Confirmed Completion (FCP_CONF) is used to acknowledge successful receipt of a response frame. In other words, once the target sends an FCP_RSP to the initiator, the initiator will send back confirmation informing the target that the response was successfully received. The confirmation is merely an added protection mechanism to help ensure successful delivery of the response—allowing the target to safely close the exchange. Confirmed Completion in a Fibre Channel environment is equivalent to the SCSI function REQ/ACK for command completion.

Example error recovery implementation with FC-TAPE

To demonstrate the error recovery process related to FC-TAPE, the Figure 2 example shows how a data frame (FCP_DATA), lost during a SCSI write command (FCP_CMD), is recovered using FC-TAPE as follows:

- **Section A:** An uninterrupted flow process writes data to a tape device. Once the data is successfully transmitted to the tape device, the tape device returns good status back to the initiator.
- **Section B:** An error causes the third data frame to be lost. In this instance, the initiator is waiting for a response, and the target is still waiting for more data. During this deadlock period, the three-second REC_TOV timer expires. Sensing a delay in the response from the target, the initiator then transmits a REC requesting information from the target regarding the status of the exchange. The target responds to the REC by transmitting an ACC back to the initiator. The ACC payload from the target contains the DTC value, indicating that the number of bytes successfully received by the target was incomplete. Next, the E_STAT bit in the ACC payload informs the initiator that the exchange is still open, and that the target does not own the sequence initiative. In other words, the target is telling the initiator to take the next action to continue the exchange.
- **Section C:** To correct the error condition, the initiator will transmit an SRR to the target requesting that the target resend an FCP_XFER_RDY (request to allow data transfer). The specific request for the FCP_XFER_RDY is defined in

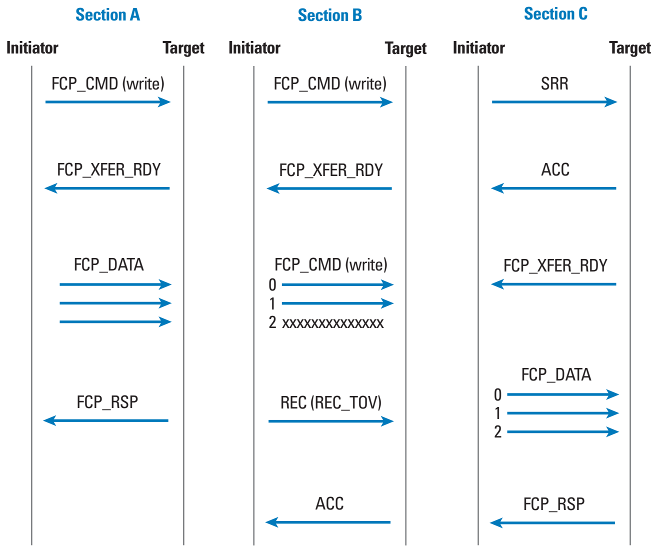

Figure 2. Example FC-TAPE error recovery process

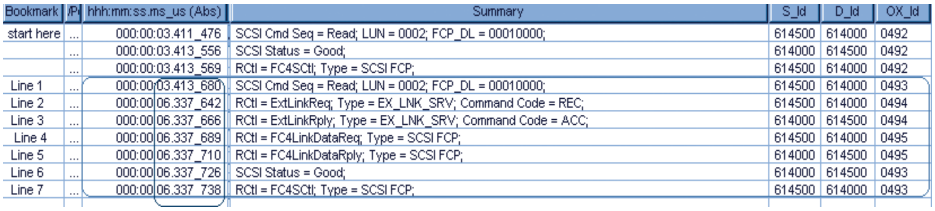

Bookmark	/P	hhh:mm:ss.ms_us (Abs)	Summary	S_Id	D_Id	OX_Id
start here	...	000:00:03.411_476	SCSI Cmd Seq = Read; LUN = 0002; FCP_DL = 00010000;	614500	614000	0492
	...	000:00:03.413_556	SCSI Status = Good;	614000	614500	0492
	...	000:00:03.413_569	RCtl = FC4SCtl; Type = SCSI FCP;	614500	614000	0492
Line 1	...	000:00:03.413_680	SCSI Cmd Seq = Read; LUN = 0002; FCP_DL = 00010000;	614500	614000	0493
Line 2	...	000:00:06.337_642	RCtl = ExtLinkReq; Type = EX_LNK_SRV; Command Code = REC;	614500	614000	0494
Line 3	...	000:00:06.337_666	RCtl = ExtLinkRply; Type = EX_LNK_SRV; Command Code = ACC;	614000	614500	0494
Line 4	...	000:00:06.337_689	RCtl = FC4LinkDataReq; Type = SCSI FCP;	614500	614000	0495
Line 5	...	000:00:06.337_710	RCtl = FC4LinkDataRply; Type = SCSI FCP;	614000	614500	0495
Line 6	...	000:00:06.337_726	SCSI Status = Good;	614000	614500	0493
Line 7	...	000:00:06.337_738	RCtl = FC4SCtl; Type = SCSI FCP;	614500	614000	0493

Figure 3. Example Fibre Channel trace showing FC-TAPE error recovery

the R_CTL payload in the SRR. The target will respond to the SRR with an ACC and then transmit an FCP_XFER_RDY back to the initiator, indicating that the target is ready to receive data transmission. Once the initiator receives the FCP_XFER_RDY, the initiator will immediately retransmit the FCP_DATA data frames.

The amount of time to recover from the original lost data frame is no more than three seconds (the REC_TOV), which is typically far less time than the error recovery mechanisms used by tape backup software. For this reason, FC-TAPE can be a valuable asset in link-level error recovery procedures for tape backup devices.

Fibre Channel trace showing FC-TAPE error recovery

Figure 3 displays a segment from a Fibre Channel trace in which FC-TAPE error recovery is invoked because of a lost FCP_RSP (good status). The debug trace data shown in Figure 3 corresponds to the following process:

- **Line 1:** The host tape backup server issues a SCSI read command.
- **Line 2:** No FCP_RSP is received within the REC_TOV timer (three seconds), and thus a REC is transmitted. (The SCSI read command was issued at the 3.4 second mark, and the REC was issued at the 6.3 second mark.)
- **Line 3:** The target accepts the REC by responding with an ACC.
- **Line 4:** The initiator issues an SRR.
- **Line 5:** The target replies to the SRR.
- **Line 6:** The target resends the "lost" FCP_RSP (good status).
- **Line 7:** The initiator sends an FCP_CONF to the target indicating that it received the FCP_RSP.

This example demonstrates the efficiency of FC-TAPE in performing error recovery without incurring overhead on the tape backup software. FC-TAPE error recovery is conducted at the Fibre Channel link level and completed within seconds between the initiator and the target tape device.

Recovering from tape backup errors quickly and efficiently

Best practices for tape backup operations recommend that administrators enable FC-TAPE error recovery on the Fibre Channel adapter as well as the data verification setting in the tape backup software. By combining FC-TAPE and tape backup software error recovery, administrators can quickly and easily enhance the capability to detect and correct errors on a Fibre Channel link, thereby helping to preserve data integrity.

Richard Golasky is a development engineer senior consultant in the Dell/EMC Enterprise Storage Group. His responsibilities include all areas of tape backup, including Fibre Channel storage, high-availability clusters, network attached storage, and performance analysis. He has a B.S. in Electrical Engineering from Florida Atlantic University.

FOR MORE INFORMATION

Dell tape backup devices:
www.dell.com/tapebackup

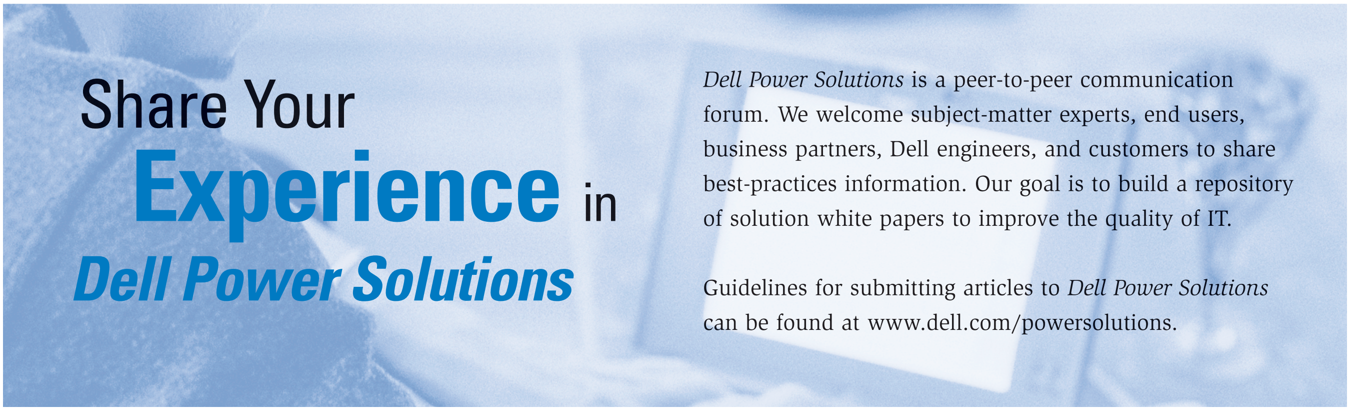

Migrating from Dell OpenManage Array Manager to

Dell OpenManage Server Administrator Storage Management

Dell™ OpenManage™ Server Administrator Storage Management is designed to streamline the administration of locally attached storage. This article explains the differences between Storage Management and Dell OpenManage Array Manager as well as the enhanced capabilities of Storage Management. It also describes how to install these storage management tools and migrate from Array Manager to Storage Management.

BY LISA FILEMYR, ERIN GEANEY, JAMEE CUSHMAN-RAMSEY, NADINE LATIEF, AND TERESA TAYLOR

Related Categories:

Dell OpenManage

Dell PowerEdge servers

Storage

Storage array

Storage software

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

As storage needs grow, organizations often require enhanced data availability and reliability. To help meet these requirements, system administrators can deploy RAID or other storage systems. Storage management software can enable administrators to create, configure, and monitor redundant virtual disks and manage physical disks.

Since 1999, Dell has offered Dell OpenManage Array Manager as its primary enterprise-level storage management software. In 2004, Dell introduced the Dell OpenManage Server Administrator Storage Management suite as a follow-on to Array Manager. This article examines the differences between these two products and explains the features and capabilities of Storage Management. In addition, this article discusses best practices for installing these two storage management applications and for migrating from Array Manager to Storage Management.

Identifying differences between Array Manager and Storage Management

Although Array Manager and Storage Management are both designed to help manage storage devices, several differences exist between the user interfaces and feature sets of these two tools. From a user perspective, there are two major differences. First, Array Manager can be installed with other Dell OpenManage applications or as a stand-alone tool, while the Storage Management architecture is integrated into the Dell OpenManage suite and thus is installed only as part of the Dell OpenManage suite.

The second major difference is the user interface. Because Storage Management plugs into the Dell OpenManage suite, it uses the Dell OpenManage Web server interface (see Figure 1). Array Manager has its own user interface that can be launched independent of other Dell OpenManage applications. Also, in environments running the Microsoft® Windows® 2000 OS, Array Manager is a Microsoft Management Console and can be used in place of Windows Logical Disk Manager.

Besides user interaction, Storage Management and Array Manager differ in the environments they support—including OS, storage controllers, platforms, and other configurations. Figure 2 compares the products and capabilities that these two tools support.

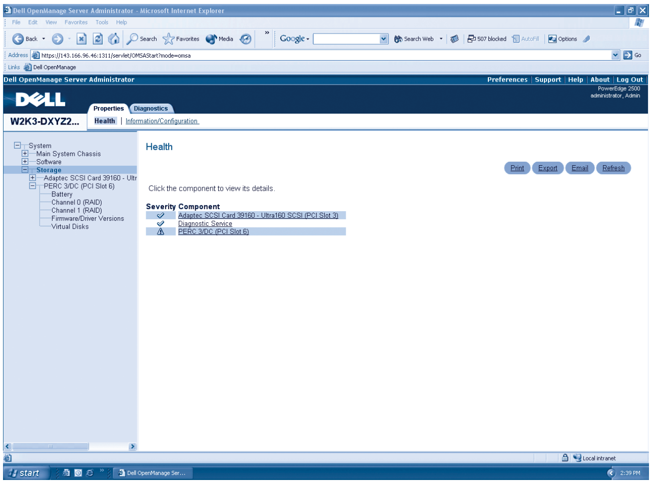
Dell OpenManage Server Administrator - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites Media
Google Search Web 507 blocked AutoFill Options
Address https://143.166.96.46:1311/servlet/OMSAStart?mode=omsa
Go
Links Dell OpenManage
Dell OpenManage Server Administrator
Preferences | Support | Help | About | Log Out
PowerEdge 2500
administrator, Admin
DELL
Properties
Diagnostics
W2K3-DXYZ2...
Health | Information/Configuration
System
Main System Chassis
Software
Storage
Adaptec SCSI Card 39160 - Ultr
PERC 3/DC (PCI Slot 6)
Battery
Channel 0 (RAID)
Channel 1 (RAID)
Firmware/Driver Versions
Virtual Disks
Print
Export
Email
Refresh
Local intranet
start
Dell OpenManage Ser...
2:39 PM

Health

Click the component to view its details.

Severity	Component
✓	Adaptec SCSI Card 39160 - Ultra160 SCSI (PCI Slot 3)
✓	Diagnostic Service
⚠	PERC 3/DC (PCI Slot 6)

Figure 1. Storage Management user interface

While both storage management tools can run on 32-bit Windows operating systems, only Storage Management offers support for the 64-bit Windows OS. Array Manager is available for the Novell® NetWare® OS, and Storage Management is available for the Red Hat® Enterprise Linux® OS.

Both products support Dell PowerEdge™ RAID Controller 2 (PERC 2), PERC 3, PERC 4, and Cost Effective RAID Controller (CERC) product families as well as Dell-supported SCSI controllers. Future releases of RAID controllers will be managed through Storage Management.

The support of storage management software also differs based on the server platform. Future Dell server platforms may support only Storage Management. Also, for organizations that have Dell PowerVault™ storage area network (SAN) systems or Fibre Channel devices, only Array Manager is available for storage management. For PowerVault network attached storage (NAS) systems, Array Manager is supported.

Another difference between Array Manager and Storage Management is in the Simple Network Management Protocol (SNMP) feature. The SNMP trap IDs of these two applications are different—so if an administrator has configured triggers on specific trap IDs for Array Manager, new event triggers will need to be created after migrating to Storage Management.

Exploring enhanced capabilities of Storage Management

Storage Management provides enhanced features for configuring and managing a system's locally attached disk storage. Using Storage Management, administrators can protect data by configuring data redundancy, assigning hot spares, or rebuilding failed drives. This tool enables administrators to perform controller and enclosure management functions for supported RAID and non-RAID controllers and enclosures from a single graphical user interface (GUI) or command-line interface (CLI). The Web-based GUI offers features for novice and advanced users along with detailed online help. The CLI is fully featured and scriptable. This section describes Storage Management features accessed through the GUI; however, the CLI can also be used to run such management tasks.

Monitoring storage status and tasks. When Storage Management is installed as part of the Dell OpenManage suite, the Storage object in the Dell OpenManage Server Administrator tree view expands to display the storage components attached to the system. A quick way to review the overall status of storage components is to select the Storage tree view object and view the Health tab (shown in Figure 1). The Health tab displays the current status of the storage components and their lower-level objects. Clicking the controller name on the Health tab displays additional health information about the controller, including the status of its array disks, virtual disks, and so on. Many storage components also have an Information/Configuration tab, which displays property information and may have drop-down menus and buttons for launching tasks and wizards.

Managing controllers. Controllers read data from disks and write data to disks. A RAID controller handles the logic needed to help protect data by making it redundant—while minimizing the additional CPU time needed for RAID calculations. When a controller is selected in the Server Administrator tree view, the drop-down menu on the Information/Configuration tab displays the controller's tasks (see Figure 3). Because different models of RAID controllers have different

Supported products and capabilities	Array Manager	Storage Management
Dell OpenManage suite	Partial integration with Dell OpenManage Server Administrator (status reporting only)	Full integration with Dell OpenManage Server Administrator
OS	Windows NetWare	Windows Linux
User interface	Windows console Limited CLI	Web HTML user interface Full CLI
Drive interface	SCSI ATA Serial ATA (SATA) Fibre Channel	SCSI ATA SATA
RAID levels	RAID-0 RAID-1 RAID-5 RAID-10 RAID-50	RAID-0 RAID-1 RAID-5 RAID-10 RAID-50
RAID types	Hardware RAID Software RAID through Windows disk and volume management (Windows 2000 only)	Hardware RAID
Languages	English	English French German Spanish Simplified Chinese Japanese

Figure 2. Comparing the products and capabilities supported by Array Manager and Storage Management

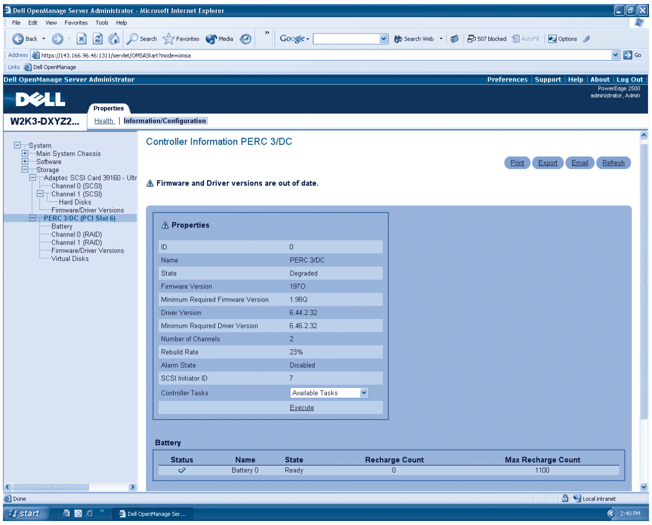

Controller Information PERC 3/DC

⚠ Firmware and Driver versions are out of date.

⚠ Properties

ID	0
Name	PERC 3/DC
State	Degraded
Firmware Version	197O
Minimum Required Firmware Version	1.98Q
Driver Version	6.44.2.32
Minimum Required Driver Version	6.46.2.32
Number of Channels	2
Rebuild Rate	23%
Alarm State	Disabled
SCSI Initiator ID	7
Controller Tasks	Available Tasks
	Execute

Battery

Status	Name	State	Recharge Count	Max Recharge Count
✓	Battery 0	Ready	0	1100

Figure 3. Managing RAID controllers with Storage Management

capabilities, the drop-down menu dynamically displays only those tasks appropriate for the selected controller.

For example, Set Rebuild Rate is a Storage Management controller task that can be used to manage system performance during the rebuilding of a virtual disk. Other tasks include creating virtual disks, reconditioning the controller battery, and exporting the controller log to a file. In addition, controllers perform storage management tasks such as rebuilding data, initializing disks, and so on. Many of these tasks are displayed in the GUI when the lower-level object (the array disk or the virtual disk) is selected in the tree view.

Managing array disks and hard disks. A controller's channels are typically attached to a server's backplane or an external disk enclosure. The backplane or enclosure contains physical disks, each labeled in Storage Management as either an array disk or a hard disk. In Storage Management terminology, an array disk is attached to a RAID controller and can be part of a virtual disk. A hard disk is attached to a non-RAID controller or to a non-RAID channel of a RAID controller.

Depending on the capabilities of the controller, Storage Management enables administrators to perform various functions on array and hard disks (see Figure 4). For example, the Prepare to Remove task spins down an array disk so that it can safely be removed from an enclosure. Administrators can also use the Blink task—which makes the LEDs on a specific disk blink—to assist in locating a disk within an enclosure.

Administrators can protect data that is stored on redundant virtual disks by assigning an array disk as a hot spare. A hot spare is an unused backup disk that is reserved to take the place of a failed array disk for redundant virtual disks. Hot spares remain in standby mode. When an array disk that is used in a redundant virtual disk fails, the hot spare is activated to replace the failed array disk.

Managing virtual disks. A virtual disk consists of one or more array disks. RAID controllers use RAID algorithms to store data on virtual disks. Different RAID levels provide varying degrees of redundancy and performance. Storage Management has an Express Wizard and an Advanced Wizard for creating virtual disks. These wizards are designed to assist administrators in creating virtual disks that effectively utilize disk space and provide sufficient data protection.

If data capacity or protection requirements change, administrators may need to either expand the virtual disk by adding array disks or change the virtual disk's RAID level. The Reconfigure Virtual Disk Wizard is designed to assist in these tasks.

Another example of a virtual disk task available through the Storage Management tool is Check Consistency. If the virtual disk is configured with a redundant RAID level, the Check Consistency task enables administrators to verify the accuracy of the redundant information and correct it if necessary. Administrators can also modify the virtual disk's performance characteristics by using the Change Policy task to change the virtual disk's read, write, and cache policies.

Installing Dell storage management software

Since March 2005, Array Manager and Storage Management have been available as components of the Dell OpenManage Systems Management CD and have been able to take advantage of the native installation strategy. Using the Dell OpenManage Systems Management CD, administrators can install the full version of Array Manager that includes the Node and Console components. The Array Manager Console can also be installed independent of the Array Manager Node through the Dell OpenManage Management Station CD. Installing the Array Manager Console can be useful in environments where both Array Manager– and Storage Management–managed systems exist. In subsequent references to installation and migration in this article, the term *Array Manager* refers to the Array Manager Node and Console package.

Beginning with the March 2005 Dell OpenManage release, Windows-based installations use Windows Installer technology, Linux-based installations use Red Hat Package Manager (RPM™) technology, and NetWare-based installations use IPS scripts. Previously, Dell OpenManage components were installed using the

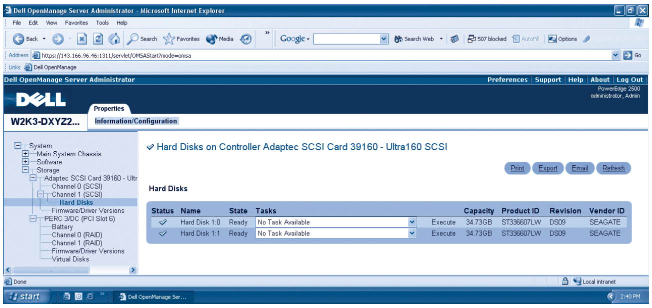

✓ Hard Disks on Controller Adaptec SCSI Card 39160 - Ultra160 SCSI

Hard Disks

Status	Name	State	Tasks		Capacity	Product ID	Revision	Vendor ID
✓	Hard Disk 1:0	Ready	No Task Available	Execute	34.73GB	ST336607LW	DS09	SEAGATE
✓	Hard Disk 1:1	Ready	No Task Available	Execute	34.73GB	ST336607LW	DS09	SEAGATE

Figure 4. Managing disks with Storage Management

Dell OpenManage installation framework that was consistent across operating systems.

In version 4 or earlier of the Dell OpenManage suite, Array Manager is the available option for storage management. Array Manager can be installed either as a stand-alone product or as part of the Dell OpenManage suite. In the express installation of Dell OpenManage, Array Manager is installed.

Dell OpenManage 4.1 introduces Dell OpenManage Server Administrator Storage Management, which can be installed only through the Dell OpenManage suite. Both Array Manager and Storage Management are available beginning in version 4.1, but only one of these applications can be installed on a system. In versions 4.1 and 4.2, the default installation is Array Manager—which means that, during an express installation, Array Manager will be installed as the storage management software; Storage Management must be installed through the custom installation option in versions 4.1 and 4.2.

Beginning with Dell OpenManage 4.3, Storage Management is the default installation for storage management software. If no storage management software is installed on the system, Storage Management will be installed in an express installation. It is also the default in a custom installation; administrators must use the custom installation option to install Array Manager in 4.3.

Migrating from Array Manager to Storage Management

To migrate systems from Array Manager to Storage Management, administrators can choose from several methods, depending on whether the native installation or Dell OpenManage installation framework is used.

For systems running a version of Array Manager earlier than version 3.1.1, a direct migration to Storage Management is not possible. To preserve metadata such as virtual disk names, administrators can upgrade to the latest Array Manager version and then migrate the system according to the procedure discussed later in this section. Alternatively, administrators can manually uninstall Array Manager before installing Storage Management, but metadata will not be preserved.

The next upgrade scenario comprises systems running Array Manager version 3.1.1 to 3.6 that are being upgraded to Storage Management through Dell OpenManage 4.1 or 4.2. In this case, administrators can migrate the system directly through the Dell OpenManage upgrade process. When performing the upgrade, administrators should use the custom installation option and select Storage Management as the storage management software.

For systems running Array Manager version 3.1.1 to 3.6 that are being upgraded through Dell OpenManage 4.3 (or later), administrators can run the native installation through the Dell OpenManage Systems Management CD. Because Array Manager is already installed on the system, the express (default) installation is Array Manager. To install Storage Management, administrators must perform a custom installation. In this case, Array Manager will automatically be removed from the system and metadata will be preserved.

In organizations that have previously upgraded or installed Array Manager through Dell OpenManage 4.3 (or later) and plan to migrate their systems to Storage Management, administrators can use the modify feature of Windows Installer. Administrators should modify the installation to include Storage Management and remove Array Manager.

Implementing integrated storage management software

A system's storage architecture comprises many components. To use storage resources effectively, administrators must choose a well-suited management tool. Dell OpenManage Server Administrator Storage Management is designed to ease storage management and enable organizations to optimize their storage resources. Using Dell OpenManage tools, administrators are enabled to easily and effortlessly migrate systems from Array Manager to Storage Management. In addition, Storage Management is seamlessly integrated into the Dell OpenManage Server Administrator suite, allowing administrators to use one consolidated approach for managing Dell enterprise platforms.

Lisa Filemyr is a software engineer on the Enterprise Storage development team at Dell. Lisa has a B.S. in Computer Science, with a minor in Applied Mathematics, from the University of Virginia.

Erin Geaney is a software engineer on the Enterprise Storage development team at Dell. Erin has a B.S. in Computer Engineering from the University of Virginia.

Jamee Cushman-Ramsey is a software engineer on the Enterprise Storage development team at Dell. Jamee has a B.S. in Computer Science from Michigan State University.

Nadine Latief is a software engineer on the Dell OpenManage Server Administrator Storage Management team. Nadine has a bachelor's degree in Computer Science, with a minor in Psychology, from Cornell University.

Teresa Taylor is an information developer on the Dell OpenManage Server Administrator Storage Management team. Teresa has a bachelor's degree from Michigan State University and is currently pursuing a master's degree in the School of Information at The University of Texas at Austin.

FOR MORE INFORMATION

***Dell OpenManage Server Administrator Storage Management User's Guide*:**
support.dell.com/support/edocs/software/OMSS10UG

Fast, Flexible Change Management

for Backup and Recovery Systems Using SQL LiteSpeed

SQL LiteSpeed™ backup and recovery software from Quest can help streamline the change-management process, particularly for enterprises that have invested heavily in Microsoft® SQL Server native command scripts. By translating native backup and restore commands using Native Command Substitution, SQL LiteSpeed helps IT organizations to save time and money—enabling quick business response while helping to reduce total cost of ownership. In addition, SQL LiteSpeed's innovative data compression technology is designed to dramatically reduce the size of backup files, which can minimize the time and storage capacity required for backups while freeing valuable network resources for business-critical applications.

BY TOM JOHANSMEYER

Related Categories:

Backup

Database

Disaster recovery

Microsoft SQL Server

Microsoft Windows

Storage

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

All too often, resource constraints prevent organizations from optimizing their data backup and recovery systems in Microsoft SQL Server environments. However, continually patching holes in backup and recovery software with SQL Server native command scripts is a time-consuming process that can result in code that is difficult and costly to maintain. Even when migrating to a new backup system is the preferred option, the associated expense and service disruption can be prohibitive, particularly for 24/7 enterprise operations.

One practical alternative is to build upon existing SQL code by creating an extensible framework for growth. SQL LiteSpeed enables this approach using Native Command Substitution (NCS) to convert SQL-native backup and restore commands to the SQL LiteSpeed equivalent. Using NCS, SQL LiteSpeed enables rapid migration from existing backup software and procedures instead of requiring administrators to rescript complex code evolved over years of changes in the IT environment.

In this way, SQL LiteSpeed helps free valuable database administrator (DBA) resources from the requirement to rewrite and test code. SQL LiteSpeed allows administrators to use familiar SQL Server commands and tools in an environment that is designed to be nearly plug and play with native SQL Server implementations. The capability to transform disparate scripts into a manageable backup and recovery environment enables rapid IT response to changing business requirements while helping to keep total cost of ownership (TCO) low.

Consider an example enterprise network configured with 100 Dell™ PowerEdge™ servers. To manage this environment, the IT organization relies on an increasingly complex set of scripts using native SQL Server commands to govern backup processes. Developed on an ad hoc basis, the backup scripts have evolved over years of in-house and outsourced software development. The resources required to maintain such a complex programming environment are considerable and certainly could be used more profitably. Time saved through more effective code development and management could be focused on activities that would help advance business goals, including upgrades, application development, and other initiatives that contribute to revenue-generating business projects.

But the alternative—migrating to an entirely different enterprise backup system—could present significant challenges as well. Residual impacts arising from potentially steep learning curves and altered business processes could result in lost staff productivity. During that time, the enterprise's ability to respond to changing market conditions might be compromised. In either case—maintaining a complex legacy backup system or switching over to a completely different backup system—TCO is likely to remain high, and worse yet, business disruptions may lead to customer dissatisfaction and lost revenue.

Understanding the SQL LiteSpeed approach

SQL LiteSpeed provides a robust enterprise backup tool that can be implemented cost-effectively with minimal disruption to existing SQL Server environments. Built for SQL Server, SQL LiteSpeed uses native SQL Server application programming interfaces (APIs) that can help organizations streamline backup management and reduce backup times significantly compared to SQL Server native backup environments.

SQL LiteSpeed uses the functionality of Microsoft Management Console and standard Microsoft tools, which are designed to help simplify implementation and minimize the learning curve. Because SQL LiteSpeed scripts are developed using SQL Server Transact-SQL, screen elements and functionality in SQL LiteSpeed are consistent with what DBAs typically use in their day-to-day jobs. As a result, extensive (and expensive) training can be avoided because DBAs are working with familiar tools and on-screen elements that have clear, intuitive labels.

Minimizing total cost of ownership

Innovative data compression technology enables SQL LiteSpeed to minimize the enterprise backup and restore window by reducing file size and thus shortening backup time compared to uncompressed native SQL backups. As a result, data compression also enables SQL LiteSpeed to help keep enterprise network resources free for business-critical application traffic. In addition, SQL LiteSpeed's data compression technology can enable organizations to reduce the amount of disk storage capacity that must be provisioned, compared to the space that would be required to back up uncompressed files.

When the need to restore data arises, time delays can lead to significant losses. The longer data remains unavailable, the greater the potential that productivity, revenue, and intangible assets such as customer satisfaction will suffer.

The opportunity to reduce TCO can depend heavily on the frequency with which databases are unavailable. Minimizing the time required to restore an unavailable database can contribute to the residual benefit of low TCO. For example, consider an enterprise backup and restore system in which each database restoration costs approximately $50 per resource-hour. In an environment that requires just 100 hours of database restores per year—which is very low for an organization supporting development, test, and production environments—the cost of owning and managing the database environment could amount to $5,000. Thus, for every 100 resource-hours of database restores, the TCO could increase by approximately $5,000. Moreover, this example projection of straightforward backup and recovery systems management costs does not account for the considerable potential loss of soft assets such as customer revenue and employee productivity while a business-critical enterprise application is unavailable.

In such a scenario, if an organization requires the equivalent of one full-time administrator to focus on database restores, it will likely budget 2,000 hours of database restoration time per year at a total of $100,000. An additional administrative expense of $100,000 to maintain a backup and restore system could have a significant impact on TCO—and such additional costs could cause the IT budget to quickly spiral out of control if unexpected problems require an increase in the time spent restoring databases. The funds could be redeployed for growth initiatives.

Besides helping to minimize the systems management cost of enterprise backups and restores, SQL LiteSpeed can help accelerate the SQL script development process.

Streamlining the backup and recovery process

Through innovative data compression technology, SQL LiteSpeed is designed to significantly reduce the TCO for data backup and recovery systems in enterprises deploying Dell PowerEdge servers. Besides helping to minimize the systems management cost of enterprise backups and restores, SQL LiteSpeed can help accelerate the SQL script development process, facilitating improvements in the existing IT environment while enabling fast, flexible response to changing business requirements. Together with SQL LiteSpeed's data compression technology, Native Command Substitution functionality allows organizations to use their existing scripts without having to retrofit code—and to create an extensible framework for growth by streamlining the development of future scripts. As a result, SQL LiteSpeed can enable organizations to free IT administrators and resources from routine systems maintenance to pursue more profitable business initiatives.

Tom Johansmeyer writes extensively about software, compliance, and project management. He has a B.A. from Ripon College.

FOR MORE INFORMATION

SQL LiteSpeed:
www.quest.com

Scaling Out SQL Server

with Data-Dependent Routing

To scale out database applications efficiently and cost-effectively, enterprises can use a data-dependent routing (DDR) method to partition and access data across an array of industry-standard, symmetric multiprocessing nodes. This article explains how administrators can implement a DDR approach to scale out highly available Microsoft® SQL Server™ database–based enterprise applications, and then presents benchmark results from a pilot study in which Microsoft engineers simulated a real-world enterprise scenario—part of the Communication Services Platform for MSN (The Microsoft Network). The pilot configuration ran Microsoft SQL Server 2005 Beta 2 on Dell™ PowerEdge™ servers, Dell/EMC storage, and Emulex network adapters.

BY MAN XIONG, BRIAN GOLDSTEIN, AND CHRIS AUGER

Related Categories:

Pilot study

Database

Dell/EMC storage

Emulex

Microsoft SQL Server

Microsoft Windows

Performance

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

The past decade has witnessed tremendous growth in enterprise applications on the Internet, and along with that, an explosion of data storage requirements. Today, many organizations must contend with the possibility that millions of online users will be using the Internet for common transactions such as shopping, storing e-mail messages, and viewing financial information. Database systems are at the heart of most enterprise applications, and Microsoft SQL Server is a leading database platform in large data centers—particularly those supporting online transactions.

A scalable database platform allows application designers to start small and grow a system as large as necessary. Traditionally, administrators have achieved scalability with symmetric multiprocessing (SMP) servers—using a scale-up approach by adding processors, memory, disks, and network adapters to a single server when required. As the workload increases, however, a single database server—also referred to as a node in this article—will eventually hit a bottleneck and be unable to support further capacity. This situation typically becomes evident in terms of diminishing performance returns when additional system components are added to the chassis—or prohibitively expensive hardware upgrades to achieve the requisite performance.

To enable growth much beyond a factor of 10—which is generally considered to be the practical limit for scaling up a monolithic SMP server—application designers may adopt a scale-out architecture in which the workload and database can be partitioned among an array of industry-standard SMP nodes. A scale-out architecture enables administrators to grow systems incrementally by adding more nodes to the array. Such an array is often referred to as a federation of servers. Ideally, the partitioning is easy to manage and transparent to the application.

The scale-out approach based on a federation of servers is designed to provide several important advantages, including the following:

- **Cost-effectiveness:** SMP node arrays can be expanded in small increments of standards-based commodity components quickly and flexibly, in response to immediate business needs.
- **Fault tolerance:** Redundant system components can help prevent downtime so that a failure in one node does not necessarily cause the entire SMP array to become unavailable.
- **High availability:** The relative independence of nodes in the SMP array helps facilitate server failover, enabling high availability of enterprise applications.

Despite such attractive benefits, the scale-out approach can pose complex management challenges because a federation of servers typically involves many more components than a monolithic SMP server. Also, every application cannot be conveniently partitioned across nodes, so the scale-out approach may work well for some applications but be unsuitable for others.

Data-dependent routing

A key design decision for any scaled-out database platform is to determine the best way to partition data among the different nodes in an SMP array. Many applications can be partitioned by a particular value such as customer name, store location, time/date, register name, and so forth. The important consideration is how the data will be accessed. For example, the retail application for a large insurance company might be partitioned by branch office, allowing each branch office to maintain client records locally. In such a scenario, all data access would be local and a batch job could replicate new or modified records to the central SQL Server array at headquarters every night. In most cases, agents working at the branch offices would not access the central SQL Server array, and corporate analysts could run their reports against the central database without accessing every branch-office SQL Server node.

Data-dependent routing (DDR) requires intelligence in the client application, typically in the middle-tier layer, to route database requests to the appropriate node. However, the DDR approach does not provide views across nodes. With the exception of shared database schema, each node in a federation of servers is designed to operate independently. The middle tier contains the mappings to how data is partitioned and which node contains specific data.

In the preceding retail insurance scenario with partitioning, the application must be designed to track where records are located. This example assumes that a SQL Server database has been created on the middle-tier Web server. The database has been partitioned by customer ID across a number of federated servers, and a lookup table on the middle tier maps each customer ID to the node where the corresponding data partition resides.

Figure 1 shows a sample lookup table. When the customer service representative needs to display all transaction records for customer 10015, the application sends the request only to node 1. Because access is localized to a single node, no requests regarding customer 10015 are sent to nodes 2 and 3.

Customer ID	Node ID
10015	1
10016	2
10017	1
10018	3

Figure 1. Lookup table mapping customer IDs to nodes where the corresponding data partitions reside

Since the application was partitioned by customer ID, records for each product ID could be stored on every database node. That would require the application to query every node, bring back all records that match each product ID, and then combine and sort the results. When product ID access is not localized, this could be a time-consuming operation. However, product ID updates could be run as a background batch job to help avoid slowing the response time for user transactions.

Challenges to the scale-out approach

When scaling out an application, administrators must address several complexities involving management, data partitioning, application development and revision, and high-availability practices. Such challenges include the following considerations:

- **Management:** The larger number of nodes in a federation of servers compared to a monolithic SMP server can increase overhead for operations management. For example, administrators must plan maintenance across an array of SMP nodes rather than a single node, and find ways to add and remove nodes without affecting application availability.
- **Data partitioning:** As enterprise applications grow, business needs may change—which may in turn require administrators to alter the way data is partitioned across the federation of servers. In addition, load balancing across an array of SMP nodes can be difficult to achieve because "hot spots" can develop on some nodes while other nodes remain relatively idle.
- **Application development and revision:** As business requirements and data access needs change, administrators must modify database schema—and find ways to implement these modifications without affecting application availability.
- **High-availability practices:** Administrators must determine how to handle single-node failures without disrupting application availability, and find ways to minimize the time it takes to restore a database to a single node.

Such challenges will be addressed in the following sections, which describe a successful scale-out scenario for a large-scale enterprise application.

MSN Communication Services Platform

At the heart of the popular Microsoft MSN® Messenger and Microsoft Hotmail® services is the Communication Services Platform (CSP) application, which manages contact information and buddy lists for these and other MSN applications. Voluminous contact information is stored in a large SQL Server database that is partitioned on 100 four-processor back-end servers running SQL Server 2000 Enterprise Edition. This federation of servers is designed to support millions of user accounts.

The CSP application provides an example of how SQL Server is designed to scale out using a federation of servers with DDR, helping to provide the following benefits:

- **Processing power:** The workload volume that can be handled by an extensible array of industry-standard SMP nodes has the potential to far exceed the processing power of any single-server SMP system.
- **Query isolation:** The isolation of queries on a per-user basis can make the MSN CSP application an excellent fit for row-based partitioning and DDR.
- **Intelligent partitioning:** Because the SQL Server database is designed for flexible data partitioning, the CSP application can be run cost-effectively on commodity hardware such as industry-standard, four-processor servers.

Figure 2 presents an overview of the system architecture that supports the Microsoft MSN CSP application. System components are configured in four tiers, as follows:

- Web servers running Microsoft Internet Information Services (IIS) 6.0
- Lookup partition database servers running SQL Server 2000 Enterprise Edition
- Database back-end servers running SQL Server 2000 Enterprise Edition
- MSN scale-out management layer server

Records are ordered and partitioned by Passport User ID (PUID) across the back-end database servers. The scale-out management layer stores the mapping of data partition IDs to the back-end database servers in its own configuration database. The lookup partition server database stores the mapping of PUIDs to data partitions. Clients of the CSP post requests to the Web server, which queries the lookup partition server with the PUID to obtain the data partition ID where the records are located. Then the Web server queries the scale-out management layer to determine which back-end database server contains the information for that user. Information is typically returned to the client in a matter of seconds.

MSN scale-out management layer

The MSN scale-out management layer provides a platform for the MSN CSP application to deploy the partitioning, DDR, and failover topology for the back-end database servers. The MSN scale-out management layer defines a *fail-safe set* as a set of databases that contains one primary database and its replicas, called *secondary databases.* A fail-safe set is designed to be the high-availability unit for the MSN scale-out management layer. In practice, the primary and secondary databases of each fail-safe set are placed on different servers to help ensure high availability. Primary databases and secondary databases are synchronized by SQL Server transactional replication for the CSP application. A *partition* is defined as a set of partitioned data, which is the unit for data partitioning and DDR. A partition is stored in a fail-safe set with its master copy on the primary database and its replicas on the secondary databases. A *failover group* is defined as a group of servers that function as backups for one another. Since workload goes only to primary databases for the CSP application, primary databases for partitions are placed across servers carefully to distribute and balance the workload among servers of the failover group. Failover groups are made independent from one another by not allowing a fail-safe set to go across a failover group boundary.

The example shown in Figure 3 is a simple failover group consisting of two servers. This group hosts two fail-safe sets, highlighted in two different colors. In this particular example, each fail-safe set stores only one partition and has only one secondary database. The data on the primary database is replicated to the secondary database. The primary database of fail-safe set 1 is placed on server 1, and its secondary database is placed on server 2. The primary database of fail-safe set 2 is placed on server 2, and its secondary database is placed on server 1. Both arrangements are designed for high availability. The primary database for partition 1 is placed on server 1 and the primary database for partition 2 is placed on server 2 to help balance the workload.

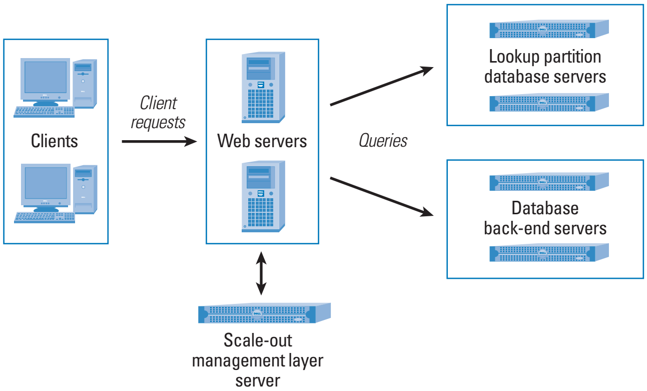
Clients
Client requests
Web servers
Queries
Lookup partition database servers
Database back-end servers
Scale-out management layer server

Figure 2. Overview of system architecture for the Microsoft MSN CSP application

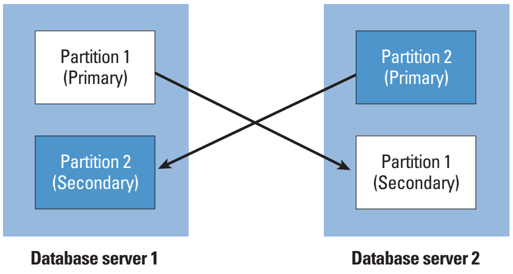

Figure 3. Replication-based fail-safe sets in a failover group

The scale-out management layer maintains a configuration database to store information about deployment and status, including:

- **Mapping:** The mapping of a partition to the SQL Server instance name and the database name of its primary database, by partition ID
- **Topology:** The topology of the fail-safe sets and failover groups
- **Status:** The current state of databases and SQL Server servers

The scale-out management layer reads configuration files for the partitioning and failover topology of the database servers and stores the information in its own SQL Server database, then configures the servers accordingly. It monitors the status of servers for failover operation and maintains the partition mapping for DDR. The scale-out management layer application also provides an administrative interface for common system maintenance operations on scale-out systems using replication to enable high availability, as shown in Figure 4.

Scalability to accommodate data and workload growth

The MSN CSP application architecture is designed to accommodate natural data growth and increased client requests. When address books are added to the system, the number of client requests increases and the CPU usage on the database servers increases. Best practices recommend a maximum operating usage of 50 percent across all CPUs to allow for the failover design shown in Figure 3. When the 50 percent CPU usage threshold is exceeded, administrators should add database servers to the system configuration. This is accomplished by adding a new failover group.

Additional failover groups

A new failover group can be added to the system by using the scale-out management layer interface. Administrators update the information in the scale-out management layer configuration database for DDR and failover topology. Lookup partition servers examine data distribution across the back-end database servers when new accounts are requested, so adding a new failover group results in new accounts being directed automatically to the new group. Figure 5 shows the DDR and high-availability architecture of a system, before and after adding a failover group. Details of actual execution will be explained in the "Pilot study of the scale-out approach" section of this article.

Workload balancing

When a new account is created, lookup partition servers estimate how busy each database server is and add the new account to the least stressed server. This task is accomplished by determining a heuristic indicator of each database server's load, which is calculated from the sum of existing accounts, weighted by the number of contacts in each account. The load of a database server is approximately proportional to the value of this heuristic indicator. Newly added database servers are configured to process new user accounts and the corresponding workload until their load reaches a level similar to that of the preexisting database servers. This approach helps ensure a smooth scale-out process that can enable the total throughput to scale linearly with the number of nodes.

Enhanced availability with transactional replication

The database uptime requirement for the MSN CSP application reads is 100 percent. Given the current Internet usage environment, 10 minutes of downtime per year is allowed for write access. The CSP application has enabled Microsoft to achieve these service levels in the two years it has been in operation.

Scale-out management layer administrative operation	Description
Promoting a database	Converts a secondary database into a primary database by redirecting workload and establishing replication from the primary to the secondary database.
Demoting a database	Converts a primary database into a secondary database. This results in draining the replication queue and dropping replication for that database. If this is the primary database in a fail-safe set, the appropriate secondary database will be promoted.
Marking a database as "offline"	Prevents client applications from querying a database and pauses all replication processes. If this is the only primary database, the appropriate secondary database will be promoted.
Marking a database as "online"	Resumes replication processes to and from a database.
Marking a database as "needs repair"	Results in draining the replication queue and dropping replication for a database. If this is a primary database, the appropriate secondary database will be promoted.
Repairing a database	Re-creates a database that is marked for repair through a backup/restore process, after which the database is left in an offline state.
Marking a server as "offline"	Causes all databases on a server to be marked as "offline."
Marking a server as "online"	Causes all databases on a server to be marked as "online."

Figure 4. Common scale-out management layer administrative operations

The MSN CSP application uses SQL Server transactional replication to help ensure high availability because transactional replication enables a combination of low latency and transactional consistency guarantees. The CSP application is not designed to read or write to both the primary and the secondary databases concurrently for two reasons:

- Application design would be much more complex. It would be necessary to implement bidirectional replication between primary and secondary databases.
- The secondary copies coexist with primary copies of different data set partitions on the same nodes. Reading from the secondary database would take resources from those primary databases.

Transactional replication uses the transaction log to capture incremental changes that were made to data in a published table. Microsoft SQL Server monitors INSERT, UPDATE, and DELETE statements, or other modifications made to the data, and stores those changes in the distribution database, which is designed to act as a reliable queue. Changes are then propagated to subscribers and applied in the same order as they occurred by opening a connection to the subscriber database and issuing SQL commands to the subscriber database. In applications where there are high write-to-read transaction ratios, replication may lag behind the transaction processing. The common limiting factors are network latency, index overhead on the subscriber database, and the number of connections to the subscriber that is executing the commands. When the source system's transactions per second exceed the replication capacity of the subscriber's system, replication latency will keep climbing until the transaction load is reduced. Losing the primary system causes transactions to be lost in the replication queue. The tolerable level of transaction loss depends on business needs and the desired end-user experience.

There are several ways to avoid the replication bottleneck:

- The data and workload should be spread out across more physical servers, which reduces the workload per distributor—although this approach can lead to underutilized hardware. This is the option elected for the CSP application.
- Since each server running SQL Server has only one distributor, having multiple SQL Server instances provides multiple distributors per server. By spreading the original workload across these instances, the replication load can be processed by more distributors per server. In this way, additional servers are not needed but additional management is required for allocating hardware resources among instances, such as CPU and memory.
- Microsoft SQL Server 2005 is designed to support parallel replication and to help ensure that transactions are processed to the subscriber in the same order as to the publisher. At press time, SQL Server 2005 had not been released so this feature was not yet implemented on the CSP production site.

On an ongoing basis, the MSN CSP application development team runs a stress test lab at Microsoft to determine the stress-level threshold where replication queue build-up exceeds specified design levels. In live production, the MSN CSP operation team monitors the stress level of client requests per node. As the number of accounts and size of the accounts increase, so does the number of queries per node. When the stress level reaches the threshold, the CSP team adds another failover group to the system.

System failure detection and failover

The MSN scale-out management layer monitors the status of all nodes. It detects the failure of a server or a database and promotes the secondary database for the failed partition by redirecting the workload traffic to the secondary database.

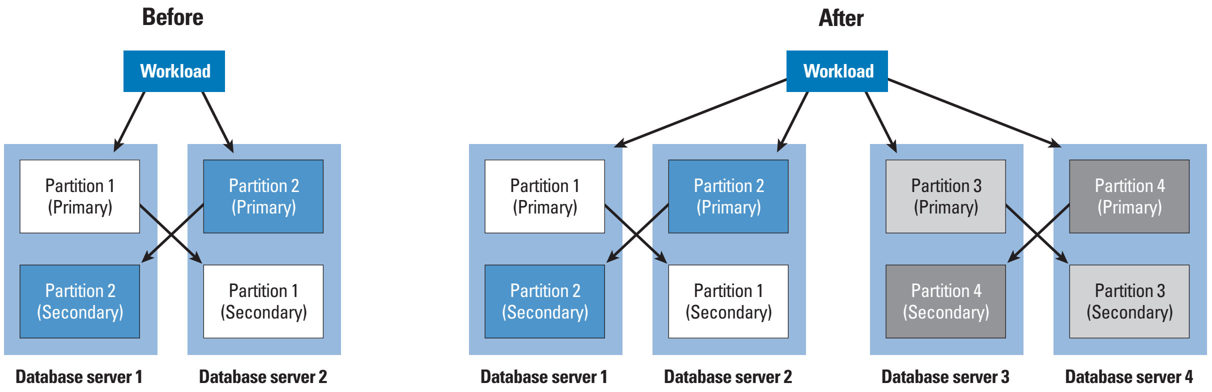

Figure 5. Adding a failover group

System role (number of systems)	Model	CPU	Physical memory	Storage	OS version	Application
Database server (4)	Dell PowerEdge 6650	4 Intel® Xeon™ processors at 2 GHz	8 GB	Dell/EMC CX600 array–based SAN	Windows Server 2003, Enterprise Edition	SQL Server 2005 Beta 2
Lookup partition server (2)	Dell PowerEdge 6650	4 Intel Xeon processors at 2 GHz	8 GB	Direct attach SCSI disk array with five 146 GB drives	Windows Server 2003, Enterprise Edition	SQL Server 2005 Beta 2
Web server (3)	Dell PowerEdge 2650	2 Intel Xeon processors at 2.4 GHz	4 GB	Local disk	Windows Server 2003, Enterprise Edition	IIS 6.0
Scale-out management layer server	Dell PowerEdge 2650	2 Intel Xeon processors at 2.4 GHz	4 GB	Local disk	Windows Server 2003, Enterprise Edition	
Web client (12)	Dell PowerEdge 1650	2 Intel Pentium® III processors at 1.4 GHz	2 GB	Local disk	Windows Server 2003, Standard Edition	

Figure 6. Configuration for pilot study deployment of Microsoft CSP running on SQL Server 2005 Beta 2

The Web server establishes connections to the back-end databases according to the information in the scale-out management layer configuration database. The Web server makes requests against the correct physical database instance during processing. Return codes from SQL Server indicate connection problems. Connection time-out is also treated as a failure. The Web server runs a client of the MSN scale-out management layer, which communicates the failure to the management layer. The management layer "blacklists" the failed databases and is designed to redirect the Web server to its backups in a matter of seconds.

Systems maintenance

Maintenance presents a particular challenge for scale-out systems. For example, an online transaction processing (OLTP) application like the MSN CSP application requires routine systems maintenance tasks including OS and application patching, node replacements and additions, database backups, and index defragmentation. Such common tasks must be accomplished in a way that incurs no appreciable impact on data availability and workload performance. Some maintenance tasks can be performed without taking databases offline; others require taking a particular database offline, while still others require bringing down an entire server. Administrators can use the administrative interface of the scale-out management layer for common system maintenance tasks.

Pilot study of the scale-out approach

In January 2005, Microsoft engineers configured an example enterprise deployment of the large-scale MSN CSP application in the Microsoft SQL Server Product Group Scalability Lab. This pilot study was designed to demonstrate how a scaled-out, end-to-end OLTP database application that enables high availability and enhanced manageability can be built using the Microsoft SQL Server 2005 platform.[1]

The test configuration included 12 Web clients, three Web servers, two lookup partition servers, four database servers, and one scale-out management layer server (see Figure 6). The three Web servers were connected to the network through a switch that distributes queries from clients to the Web servers in a round-robin pattern to enable load balancing. Data and log files for the database servers were stored across the same group of disks on a storage area network (SAN), which consisted of a Dell/EMC CX600 storage array with 146 GB 10,000 rpm drives. Each database server was connected to the 2 Gbps switched Fibre Channel SAN with two Emulex LP9802 Peripheral Component Interconnect Extended (PCI-X)–based host bus adapters to enhance I/O capacity and I/O failover.

The test team configured the enterprise scenario described in this article using SAN storage instead of direct attach storage (DAS) because the SAN approach is designed to enable the following benefits compared to DAS:

- **Centralized management:** Helps IT organizations streamline administration.
- **Flexible scale-out:** Enables administrators to increase capacity and/or the number of spindles without adding servers.
- **High availability:** Allows administrators to upgrade and deploy storage resources without disrupting business operations running on individual servers.

Although the SAN approach can be significantly more expensive to deploy and maintain than DAS, SANs can be used to enable sophisticated high-availability data deployments. However, that discussion is beyond the scope of this article. *Note:* Best practices recommend careful planning when using a SAN to support high data availability, to help avoid configurations in which a SAN may become the single point of failure for an entire enterprise system.

Increasing throughput with additional nodes

Figure 7 shows that the total number of transactions per second processed by the four database servers increased proportionally to

[1] The Microsoft SQL Server 2005 Beta 2 release was used in the pilot study described in this article. Actual features of Microsoft SQL Server 2005 as it is released to market may differ slightly; actual performance depends on various factors, including but not limited to the specific hardware and software configurations on which Microsoft SQL Server 2005 is deployed.

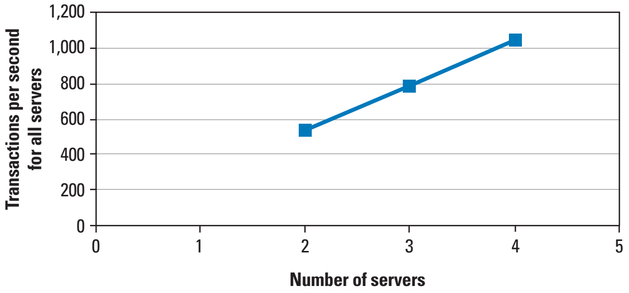

Figure 7. Workload performance scaling versus number of database servers

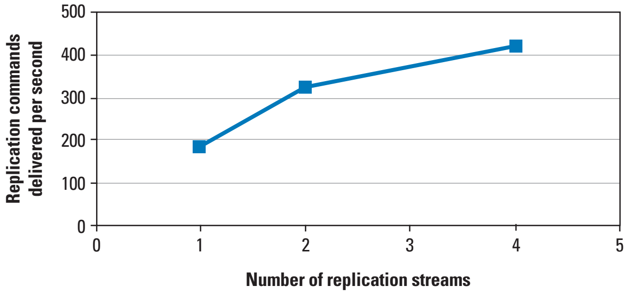

Figure 8. Replication throughput versus number of replication streams

the number of database servers in this pilot study. These results demonstrate that true linear scaling is possible with proper application architecture. The MSN CSP application is designed to be scaled out for any number of users; currently, millions of users are supported by the online MSN CSP application.

Combining the scale-out management layer and replication for enhanced availability

Together with the scale-out management layer application, replication can help provide excellent data availability. Single-threaded transactional replication in SQL Server 2000 prevented the MSN CSP application from fully utilizing available server resources. To help overcome that limitation, SQL Server 2005 is designed to support parallel streaming of replication commands—enabling multiple streams of 1 to 64 commands to be processed simultaneously. The optimal number of streams and how well the performance will scale typically depend on a variety of factors. Best practices include:

- **Number of CPUs:** Do not use more replication streams than the number of CPUs processing those streams to help minimize context-switching overhead.
- **Blocking:** Avoid overlapping transactions on tables being replicated to help prevent replication streams from blocking one another.
- **CPU cycles:** Be aware that adding a replication stream creates a task, thus requiring more work and potentially increasing CPU usage.

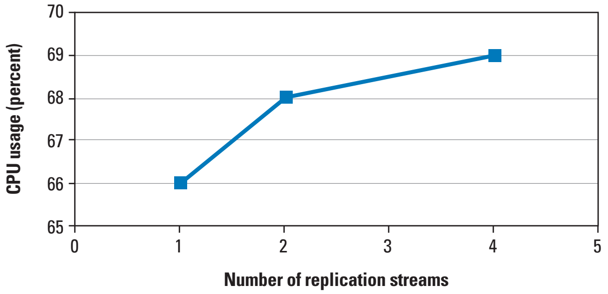

Figure 9. CPU usage versus number of replication streams

- **Number of replication commands:** Run only the necessary number of streams; avoid adding streams if the replication queue can be drained in a timely way.

Figures 8 and 9 show that, in the pilot study, the replication throughput indicator (number of replication commands delivered per second) increased significantly when the number of replication streams was increased up to four streams. Each additional stream added roughly 1 to 2 percent of CPU usage.

A highly manageable SQL Server scale-out system

In the example enterprise deployment of the Microsoft MSN CSP application, engineers observed linear performance scaling for up to four database servers. These results were achieved by partitioning and accessing data across an array of industry-standard SMP nodes that were configured as a federation of Dell PowerEdge servers. In addition, the pilot study explored ways to enable high availability using the scale-out management layer application in conjunction with SQL Server transactional replication. The findings discussed in this pilot study demonstrate that Microsoft SQL Server 2005 can be used to effectively scale out an enterprise database application using DDR to help accommodate data and workload growth.

Man Xiong is a software design engineer in test on the Microsoft SQL Server product team. Man has a master's degree in Computer Science and Engineering from the University of Washington.

Brian Goldstein is a lead program manager on the Microsoft SQL Server product team. Brian has a master's degree in Engineering from the Massachusetts Institute of Technology and is a Microsoft Certified Systems Engineer.

Chris Auger is an enterprise technologist in the Dell Advanced Systems Group.

FOR MORE INFORMATION

Microsoft SQL Server 2005:
www.microsoft.com/sql/2005/default.asp

Migrating an Oracle10*g* RAC Database

from Oracle Cluster File System to Oracle Automatic Storage Management on Dell/EMC Storage Arrays

Oracle® Real Application Clusters (RAC) database file storage options have further matured with the release of Oracle 10*g*™ software and its Oracle Automatic Storage Management (ASM) feature. Before ASM, the options available for Oracle RAC database file storage were limited to raw devices or the open source Oracle Cluster File System (OCFS). This article compares raw devices and OCFS with the ASM approach, and provides basic steps and best practices for migrating from OCFS to ASM.

BY ZAFAR MAHMOOD, UDAY DATTA SHET, AND BHARAT SAJNANI

Related Categories:

Database

Dell/EMC storage

File systems

Oracle

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

Oracle Real Application Clusters (RAC) databases require direct access to a physical storage system. Before the availability of Oracle Cluster File System (OCFS) and Oracle Automated Storage Management (ASM), RAC databases required that data files, control files, and redo log files be placed on raw devices so that they could be shared among cluster nodes. Database administrators were required to create the raw devices, using OS-specific commands, before building the database. Administrators could bind a raw device to an existing block device, such as a disk, and use the raw device to perform I/O operations with that block device. Such I/O bypasses the Linux® OS buffer cache that is normally associated with block devices and can help eliminate file system overhead. Figure 1 depicts the association of physical block devices to Oracle database files for the creation of a RAC seed database.

Although raw devices fulfill the requirements for a RAC database, they can be inflexible and difficult for system administrators to manage. The unformatted disk partitions cannot use several key Oracle9*i*™ and Oracle 10*g* functions, including the automatic database file extension and Oracle Managed Files (OMF) features. Each raw device is bound to only one Oracle database file and has a fixed size. If the database runs out of space on that partition, the database administrator must create another partition and add another database file to the Oracle tablespace. Also, the maximum number of raw devices—and consequently the maximum number of Oracle database files—that administrators can create on Red Hat® Enterprise Linux is limited to 255.

Introducing Oracle Cluster File System

The next evolution of RAC database file storage was the introduction of Oracle Cluster File System. OCFS is an open source cluster file system that provides a convenient alternative to raw devices for storing Oracle RAC database files on shared storage accessible by all cluster nodes. In environments using OCFS, all cluster nodes have the same view of the Oracle database files and can read and write to the shared storage concurrently.

OCFS is designed specifically for RAC databases. It eliminates the requirement for Oracle database files to be linked to unformatted partitions on logical drives. Administrators can create OCFS volumes on RAID disks so that the volumes span multiple shared disks for redundancy and performance enhancements. Figure 2 depicts the association between physical shared storage devices and Oracle RAC database files in an OCFS environment.

Advantages of OCFS over raw devices

On Linux platforms, OCFS behaves much like a generic file system, and many of the native Linux OS utilities—such as mkdir, rmdir, mv, cp, ln (softlinks only), tar, and cpio—work the same way on OCFS as they do on ext2 and ext3 file systems. Administrators can use an updated coreutils Red Hat Package Manager (RPM™), which supports the Direct IO feature to enhance the performance of these operations. Furthermore, the similarity of OCFS to the ext2 and ext3 file systems in terms of its operations, look-and-feel, and behavior can help simplify Oracle RAC database administration.

When compared to raw devices, OCFS provides several advantages. First, it eliminates the need to manage and set up raw devices, which can be a time-consuming, management-intensive process. Second, raw devices limit the number of database files that can be created to 255. In comparison, administrators do not have such a limitation when using OCFS.

Moreover, with a shared file system such as OCFS, the nodes within the database cluster can share archive logs, which helps streamline the media recovery process because every node has access to archived log files. Finally, OCFS can enable lower total cost of ownership than raw devices because it is designed to ease management tasks and minimize backup and recovery windows. Also, OCFS allows administrators to use the OMF and auto-extend features unavailable with raw devices, which enhances functionality.

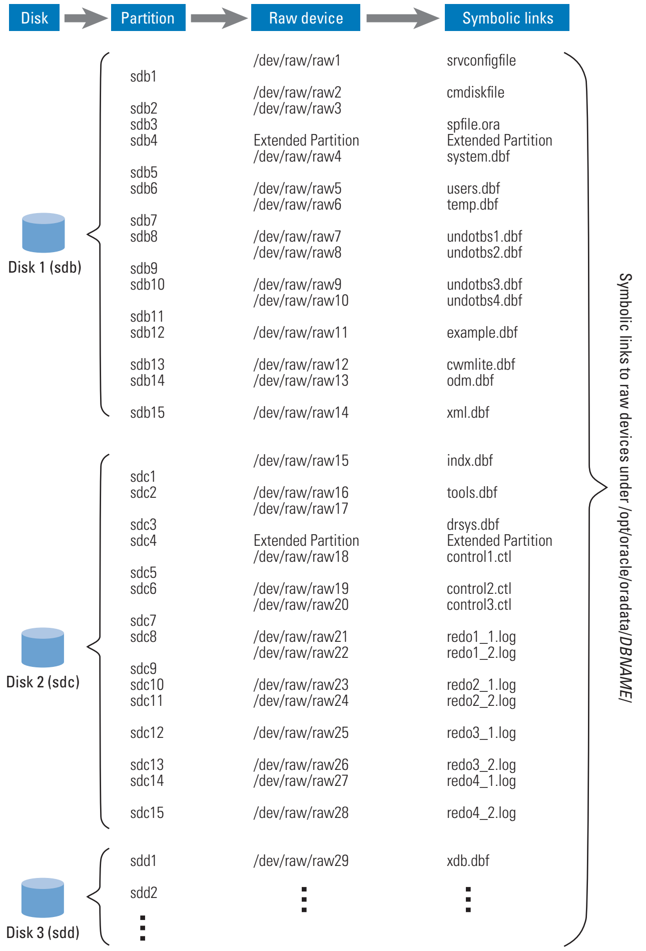

Figure 1. Using raw devices to associate physical block devices with Oracle RAC database files

Limitations of OCFS

OCFS version 1.0 supports only Oracle database–related files that include redo log files, archive log files, control files, and database data files. OCFS also supports the shared quorum disk file for the Oracle Cluster Manager and the shared configuration files. However, using OCFS for files that are not accessed by the Oracle relational database management system (RDBMS) server is not recommended. Also, Linux-based systems could not take advantage of the Async IO feature for OCFS releases before version 1.0.14 and Red Hat Enterprise Linux 3 Quarterly Update 4 or Red Hat Enterprise Linux 2.1 Quarterly Update 6 with the e57 kernel release. Oracle plans to release OCFS version 2 in summer 2005, and this version is expected to support the shared Oracle home capability. It is expected that the further planned releases of OCFS will be designed to support any type of files and to be fully compliant with POSIX (Portable Operating System Interface for UNIX®).

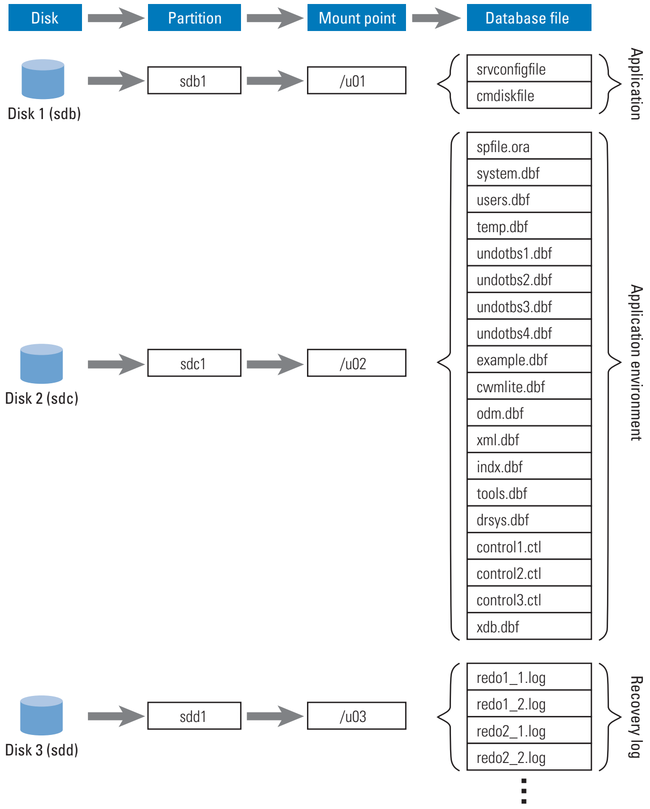

Figure 2. Using OCFS to associate physical storage partitions and mount points with Oracle RAC database files

Introducing Oracle Automatic Storage Management

ASM is an Oracle database file system and volume manager that is built into the Oracle Database 10*g* kernel. It provides both a graphical user interface (GUI) through the Oracle Enterprise Manager 10*g* Grid Control management application and a simple Oracle SQL*Plus® interface that allows database administrators to use familiar SQL Data Description Language (DDL) commands such as CREATE, ALTER, and DROP to manage storage volumes or disk groups for Oracle databases. Figure 3 depicts the association of Oracle database files with ASM disk groups.

ASM is designed to include benefits such as the elimination of manual I/O tuning, the ability to dynamically change storage resources, and automatic rebalancing. ASM also has the built-in capability to automatically tune, rebalance, and rebuild storage devices on demand.[1] Database administrators can save a significant amount of time by using ASM because ASM enables automatic management of these database tasks—freeing administrators to address other responsibilities.

The next section provides best-practices recommendations for migrating a RAC database from OCFS or raw devices to ASM. For step-by-step instructions on migrating to ASM, visit *Dell Power Solutions* online at www.dell.com/powersolutions.

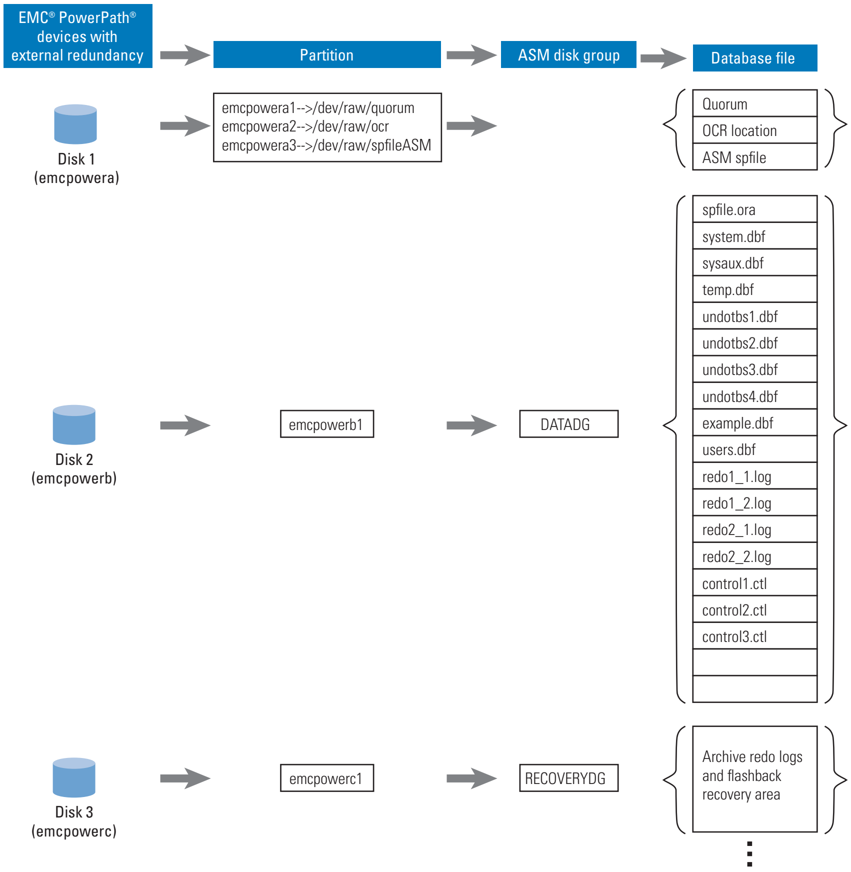

Figure 3. Using ASM disk groups on Dell/EMC Fibre Channel storage for Oracle 10*g* RAC database files

Preparing for migration from OCFS or raw devices to ASM

Administrators should take into account the following considerations before performing a database or storage migration:

- Administrators should always create a full backup of the database files using the Oracle Recovery Manager (RMAN) utility in case of data loss or system outage.
- ASM is available only with Oracle 10*g* Release 1 or later. If using OCFS or raw devices with an Oracle9*i* RAC database, administrators must first migrate the database to Oracle 10*g* and then migrate to ASM.
- For ASM migration, administrators need extra storage equal to or greater than the existing database files and the archive log files.
- Administrators can perform either a full database migration to ASM or a rolling migration, depending on the downtime window. In a rolling migration, each tablespace and data file is migrated to ASM individually. During this process, some data files and tablespaces will be using OCFS while others are using ASM. Typically, this does not cause problems for the RAC database functionality or performance.

Best practices for migrating to ASM

For migration to ASM, administrators should verify that at least two OCFS volumes are available to store certain cluster configuration and parameter files. Administrators should use an OCFS volume instead of an ASM disk group to store Oracle server parameter files. In addition, the cluster voting disk and the Oracle Cluster Repository (OCR) files should be stored on an OCFS volume.

ASM should be used to facilitate management of the Oracle database files, but it does not provide a file system–like view from the OS of the stored files. Administrators must use either Oracle Enterprise Manager Database Control or the SQL*Plus interface to query the file storage and its properties. Best practices also recommend that the Oracle ASM library driver be used to present logical volumes to the

[1] For detailed information about the architecture, benefits, and implementation of Oracle ASM, see "Enabling a Highly Scalable and Available Storage Environment Using Oracle Automatic Storage Management" by Zafar Mahmood, Joel Borellis, Mahmoud Ahmadian, and Paul Rad in *Dell Power Solutions,* June 2004; www.dell.com/downloads/global/power/ps2q04-008.pdf.

ASM instance instead of the raw device interface, because the ASM library interface offers enhanced user-friendliness.

Depending on the organization's business requirements and available downtime window, administrators should use the rolling migration method in which each tablespace and data file is migrated one at a time. When installing Oracle 10*g*, administrators should install ASM into a separate Oracle home, rather than the database home. This way, they can upgrade only the ASM binaries when updates become available, without affecting the database binaries.

Best practices also recommend that the OMF feature be used with ASM to minimize user file management errors, to enforce Oracle File Architecture (OFA) standards, and to take advantage of the automatic deletion of ASM files when database files are dropped. In addition, administrators should use Oracle Cluster Ready Services (CRS) to create dependency between the database instance and the ASM instance to help ensure that the ASM instance always starts before the database instance and that the ASM instance cleanly shuts down the database instance before it shuts down itself. Administrators can create this dependency by using the following command:

```
srvctl modify asm -d database name -i database
    instance name -s ASM instance name
```

Configuring storage to support the ASM database. When creating ASM disk groups on Dell/EMC storage arrays, administrators should configure the storage volumes as RAID-10 and then use the external redundancy option. To gain additional striping and to increase the number of spindles from the maximum of 16 per logical unit (LUN), administrators must create the disk group across multiple RAID-10 logical volumes. Administrators should separate the database area from the flashback area and make sure that both areas do not share the same physical spindles. They should use as many disk spindles of similar size and characteristics as possible in a single disk group.

If adding external redundancy by using storage arrays, administrators should make sure that the LUN stripe size is as close to 1 MB as possible to match the ASM stripe size. The disk partition boundary should start at 1 MB to help ensure proper I/O alignment because ASM writes to storage volumes in 1 MB stripes. In addition, administrators can use multiple initiators or host bus adapters and multipath software to enable high availability and load balancing for I/O.

Finally, administrators should set parameters according to the preceding best practices. The default value of the `SHARED_POOL_SIZE` parameter should be increased using the following formula:

Extra shared pool size required = (1 MB per 100 GB of ASM storage using external redundancy) + 2 MB

The value of the `PROCESSES` parameter should be increased by a value of 16 from the original value to take into account the ASM background processes after migration to ASM. The value of the `LARGE_POOL_SIZE` parameter should be increased by 600 KB after migration to ASM.

Easing management of Oracle 10*g* databases

In Oracle 10*g* environments, Oracle Automatic Storage Management is well suited for managing database storage. ASM is designed not only to provide ease of management and flexibility, but also to enhance the I/O performance of the Oracle database. In addition, migration from OCFS or raw devices to ASM need not be a complex process, thanks to flexible tools such as Oracle Recovery Manager. When migrating Oracle database systems to ASM, an organization should consider the best-practices recommendations discussed in this article and choose the migration method—whether the full database or one file at time—that best suits its business requirements.

Acknowledgments

The authors would like to thank their colleagues Ramesh Rajagopalan and CS Prasanna Nanda from the Dell Database Solutions Group for their valuable comments and help with this article.

Zafar Mahmood is a senior consultant in the Database and Applications team of the Dell Product Group. Zafar has an M.S. and a B.S. in Electrical Engineering, with specialization in Computer Communications, from the City University of New York.

Uday Datta 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).

Bharat Sajnani is a systems engineer on the Database and Applications team of the Dell Product Group. Bharat has a B.S. in Computer Engineering and a master's degree in Computer Engineering from The University of Texas at Austin.

FOR MORE INFORMATION

Oracle 10*g* ASM technology:
www.oracle.com/technology/products/database/asm/index.html

Oracle Real Application Clusters:
www.oracle.com/database/rac_home.html

Dell/Oracle certified and validated configurations:
www.dell.com/oracle

Dell/Oracle 10*g* RAC resources:
www/dell.com/10g

OCFS open source project:
oss.oracle.com/projects/ocfs

An Overview of Xen Virtualization

Xen virtualization technology—available for the Linux® kernel—is designed to consolidate multiple operating systems to run on a single server, normalize hardware accessed by the operating systems, isolate misbehaving applications, and migrate running OS instances from one physical server to another. This article provides an overview of Xen 3.0 architecture and how it is expected to utilize Intel® Virtualization Technology to enhance manageability and optimize resource utilization in Linux environments.

BY TIM ABELS, PUNEET DHAWAN, AND BALASUBRAMANIAN CHANDRASEKARAN

Related Categories:

Dell PowerEdge servers

Linux

Virtualization

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

Recent advances in virtualization technologies—enabling data centers to consolidate servers, normalize hardware resources, and isolate applications on the same physical server—are driving rapid adoption of server virtualization in Linux environments. This article provides an overview of Xen 3.0 architecture, which is near-native-speed virtualization software for Intel x86 architectures.

Virtualization software abstracts the underlying hardware by creating an interface to *virtual machines* (VMs), which represent virtualized resources such as CPUs, physical memory, network connections, and block devices. Software stacks including the OS and applications are executed on top of the VMs. Several VMs can run simultaneously on a single physical server. Multiplexing of physical resources between the VMs is enforced by a VM monitor, which is also designed to provide the required translations of operations from the VMs to the physical hardware.

Full virtualization versus para-virtualization

There are several ways to implement virtualization. Two leading approaches are full virtualization and para-virtualization. *Full virtualization* is designed to provide total abstraction of the underlying physical system and creates a complete virtual system in which the guest operating systems can execute. No modification is required in the guest OS or application; the guest OS or application is not aware of the virtualized environment so they have the capability to execute on the VM just as they would on a physical system. This approach can be advantageous because it enables complete decoupling of the software from the hardware. As a result, full virtualization can streamline the migration of applications and workloads between different physical systems. Full virtualization also helps provide complete isolation of different applications, which helps make this approach highly secure. Microsoft® Virtual

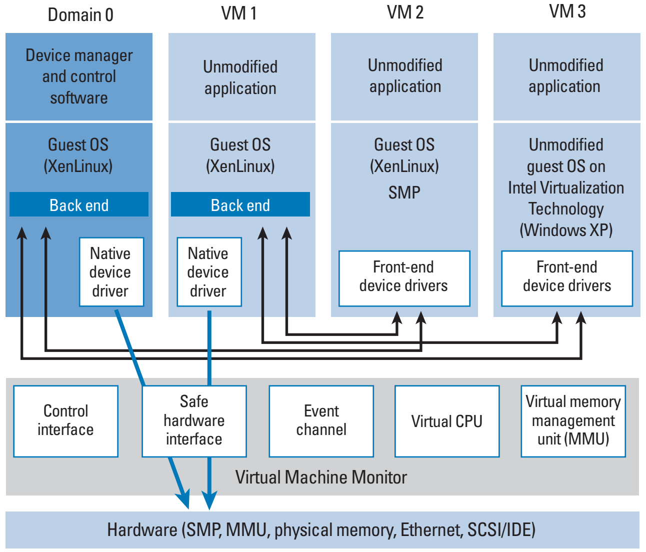

Figure 1. Xen 3.0 architecture: Hosting four VMs

Server and VMware® ESX Server™ software are examples of full virtualization.

However, full virtualization may incur a performance penalty. The VM monitor must provide the VM with an image of an entire system, including virtual BIOS, virtual memory space, and virtual devices. The VM monitor also must create and maintain data structures for the virtual components, such as a shadow memory page table. These data structures must be updated for every corresponding access by the VMs.

In contrast, *para-virtualization* presents each VM with an abstraction of the hardware that is similar but not identical to the underlying physical hardware. Para-virtualization techniques require modifications to the guest operating systems that are running on the VMs. As a result, the guest operating systems are aware that they are executing on a VM—allowing for near-native performance. Para-virtualization methods are still being developed and thus have limitations, including several insecurities such as the guest OS cache data, unauthenticated connections, and so forth.

Xen 3.0 architecture

Xen is an open source virtualization software based on para-virtualization technology. This section provides an overview of the Xen 3.0 architecture.[1]

Figure 1 shows the architecture of Xen 3.0 hosting four VMs (Domain 0, VM 1, VM 2, and VM 3). This architecture includes the Xen Virtual Machine Monitor (VMM), which abstracts the underlying physical hardware and provides hardware access for the different virtual machines. Figure 1 shows the special role of the VM called Domain 0. Only Domain 0 can access the control interface of the VMM, through which other VMs can be created, destroyed, and managed. Management and control software runs in Domain 0. Administrators can create virtual machines with special privileges—such as VM 1—that can directly access the hardware through secure interfaces provided by Xen. Administrators can create other virtual machines that can access the physical resources provided by Domain 0's control and management interface in Xen.

In this example, the guest operating systems in VM 1 and in VM 2 are modified to run above Xen and also have Xen-aware drivers to enable high performance. Near-native performance can be achieved through this approach. Unmodified guest operating systems are also supported, as discussed in the "Xen and Intel Virtualization Technology" section in this article. In addition, the developers of Xen 3.0 plan to include support for virtual machines with symmetric multiprocessing (SMP) capabilities, 64-bit guest operating systems, Accelerated Graphics Port (AGP), and Advanced Configuration and Power Interface (ACPI).

In a virtual data center framework, CPU, memory, and I/O components need to be virtualized. Xen 3.0 is designed to enable para-virtualization of all three hardware components.

Para-virtualization techniques require modifications to the guest operating systems that are running on the VMs. As a result, the guest operating systems are aware that they are executing on a VM—allowing for near-native performance.

CPU operations. The Intel x86 architecture provides four levels of privilege modes. These modes, or *rings*, are numbered 0 to 3, with 0 being the most privileged. In a non-virtualized system, the OS executes at ring 0 and the applications at ring 3. Rings 1 and 2 are typically not used. In Xen para-virtualization, the VMM executes at ring 0, the guest OS at ring 1, and the applications at ring 3. This approach helps to ensure that the VMM possesses the highest privilege, while the guest OS executes in a higher privileged mode than the applications and is isolated from the applications. Privileged instructions issued by the guest OS are verified and executed by the VMM.

Memory operations. In a non-virtualized environment, the OS expects contiguous memory. Guest operating systems in Xen para-virtualization are modified to access memory in a non-contiguous

[1] Note that Xen 3.0 is still under development by a large community of open source developers, and the actual features and architectural details may vary from what is discussed in this article. For more information, visit wiki.xensource.com.

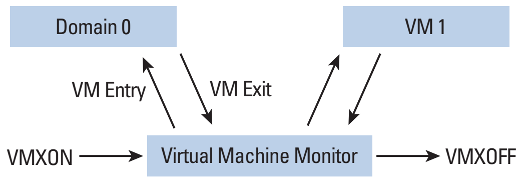

Figure 2. VMM support in Intel Virtualization Technology

manner. Guest operating systems are responsible for allocating and managing page tables. However, direct writes are intercepted and validated by the Xen VMM.

I/O operations. In a fully virtualized environment, hardware devices are emulated. Xen para-virtualization exposes a set of clean and simple device abstractions. For example, I/O data to and from guest operating systems is transferred using shared-memory ring architecture (memory is shared between Domain 0 and the guest domain) through which incoming and outgoing messages are sent.

Modifying the guest OS is not feasible for non–open source platforms such as Microsoft Windows® 2000 or Windows Server™ 2003 operating systems. As a result, such operating systems are not supported in a para-virtualization environment. The following section explains how Xen works with Intel Virtualization Technology to support unmodified operating systems.

Xen and Intel Virtualization Technology

Intel Virtualization Technology (IVT), which is expected to be available in future Intel processors, is designed to support virtualization. IVT-enabled processors will have additional instructions that can be used by the VMM to create and support VMs. In terms of ring architecture, IVT places the VMM one level below ring 0, allowing the VMs to execute in ring 0.

Figure 2 depicts a simple processor state of IVT. The system starts the VMM by executing the VMXON instruction. The VMM can then schedule any VM by executing the VM Entry instruction. The processor state of the VMM is automatically stored and the saved processor state of the VM is loaded. This level of hardware support enables faster context switching, which can enhance overall system performance. Under certain conditions, the VM may relinquish control to the VMM. Such conditions are called VM Exits and can be caused by several factors, including external interrupts, nonmaskable interrupts, page faults, and other high-privileged instructions executed by the VM.

Processor-level support for virtualization allows Xen 3.0 to support an OS whose source code cannot be modified. In such cases, the Xen implementation would have the capability to enable full virtualization with support from IVT. That is, Xen is designed to provide a fully abstracted VM—including virtual BIOS and virtual devices—to allow full support for unmodified guest operating systems.

Xen 3.0 management tools

Two important management tools in Xen 3.0 are xend and xm. The Xen daemon, xend, is a Python program that forms a central point of control for starting and managing VMs. Major functions include invoking setup of virtual networking for VMs, providing a console server, and maintaining the Xen events log.

The Xen command-line interface, xm, provides functionality to create, destroy, save, restore, shut down, migrate domains, and so forth. It also enables administrators to configure the CPU scheduler, list active domains, adjust the memory footprint using ballooning, and call an xend HTTP application programming interface (API) directly.

Xen looking forward

Xen enables administrators to implement two types of virtualization environments: para-virtualized VMs, which can enhance performance but require guest OS modifications, and fully virtualized VMs, which are highly portable and do not require guest OS modifications. Developers are planning to include Xen in the Linux 2.6.12 kernel so that organizations can take advantage of this powerful virtualization tool in their Linux environments.

Tim Abels is a senior software architect currently developing scalable enterprise computing systems at Dell. Tim has an M.S. in Computer Science from Purdue University.

Puneet Dhawan is a systems engineer on the Scalable Enterprise Team at Dell. Puneet has a bachelor's degree in Electrical Engineering from Punjab Engineering College (PEC) in Chandigarh, India, and a master's degree in Computer Engineering from Texas A&M University.

Balasubramanian Chandrasekaran is a systems engineer in the Scalable Enterprise Computing Lab at Dell. His research interests include virtualization of data centers, high-speed interconnects, and high-performance computing. Balasubramanian has an M.S. in Computer Science from The Ohio State University.

FOR MORE INFORMATION

Xen virtual machine monitor:
xen.sf.net

XenSource:
www.xensource.com

Improving Fault Tolerance

Using Memory Redundancy and Hot-Plug Actions in Dell PowerEdge Servers

Features that enable redundancy across physical memory can enhance server reliability and help keep critical business applications available 24/7—particularly when combined with hot-plug capabilities designed into the Dell™ PowerEdge™ 6850 server. Allowing administrators to replace failing memory devices and add incremental memory upgrades while a server is running can help reduce system downtime and enhance scalability considerably. This article discusses memory redundancy and hot-plug features of PowerEdge 6850 servers.

BY CHANDRA S. MUGUNDA, VIJAY NIJHAWAN, DENNIS STULTZ, SAURABH GUPTA, AND HARISH JAYAKUMAR

Related Categories:

Dell PowerEdge servers

Memory

Multiprocessor (MP)

System architecture

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

RAID technology helps provide redundancy, fault tolerance, and high availability in enterprise disk drive subsystems. Different RAID levels, such as RAID-0, RAID-1, RAID-5, and RAID-10, can be configured to enable secondary storage benefits that suit specific long-term organizational requirements. Currently, the low cost of physical memory allows servers to support 32 GB to 64 GB of server RAM cost-effectively in dual in-line memory modules (DIMMs).

Unfortunately, the requirement for large memory capacities can increase the chance of memory errors simply because physical memory devices have the potential for failure. Thus, the more memory that resides in a system, the greater the potential for memory failure in that system. Safeguarding against potential memory failures and helping to ensure uninterrupted application availability, fault tolerance, redundancy, and hot-plug capabilities can be crucial for IT environments.

Hard and soft memory errors

Memory is an electronic storage device, and electronic storage devices have the potential to incorrectly return information—that is, data read from memory can be different from what was originally written to memory. Dynamic RAM (DRAM) stores 1s and 0s as charges on small capacitors residing on the DIMM, which must be continually refreshed to help ensure that the data is not lost. But even a small electrical disturbance near the memory cell can alter the charge in a capacitor, thus causing a memory error.

Typically, two types of error occur in a memory system. The first is called a repeatable, or *hard,* error.

Hard errors consistently return incorrect results and usually indicate that a piece of hardware is broken. For example, a bit may be stuck such that it always returns a 0 regardless of whether a 1 or a 0 is written to it. Hard errors are relatively easy to diagnose and fix because they are consistent and repeatable.

The second kind of error is called a transient, or *soft,* error. Soft errors occur when a bit reads back the wrong value once, but subsequently functions correctly. Soft errors are more common than hard errors—and are also more difficult to diagnose. They are not caused by circuit problems and, once corrected, do not reoccur.

Memory error detection and correction

Reliability in memory starts at the DIMM level. To help ensure a reliable memory system in servers, it is essential to provide protection from both hard and soft memory errors. Memory detection or correction protocols such as parity checking or error-correcting code (ECC) are designed to provide true protection from hard and soft memory errors.

Parity checking is one of the oldest and most basic forms of memory checking. It is a simple method of detecting single-bit errors in a memory system. Along with the eight bits of data stored in memory, parity checking uses one additional bit to determine the parity of the byte—odd or even. However, parity checking detects only odd-numbered single-bit errors, and does not enable administrators to locate and correct these errors. When a parity error is detected, the parity circuit generates what is called a nonmaskable interrupt (NMI), which is generally used to instruct the processor to halt immediately. The processor is halted to help ensure that the incorrect memory does not corrupt other data on the system.

ECC, on the other hand, not only detects both single-bit and multibit errors, but it also corrects single-bit errors. Each time data is stored in memory, ECC memory uses an algorithm to add a block of bits known as check bits. When this data is retrieved, the sum of the check bits (the checksum) is recomputed. The checksum of the written data is then compared with the checksum of the read data to determine whether any of the data bits are corrupted. If the checksums are identical, this indicates to the ECC memory that there is no error. If they are different, the data contains one or more errors. The ECC memory logic then isolates the errors and reports them to the system. For a correctable single-bit error, the ECC memory logic corrects the error and outputs the corrected data without halting the system. For multibit errors—that is, errors involving two or more bits—ECC memory is capable of detecting but not correcting the errors. When multibit errors occur, ECC handles them by generating an NMI that instructs the system to halt.

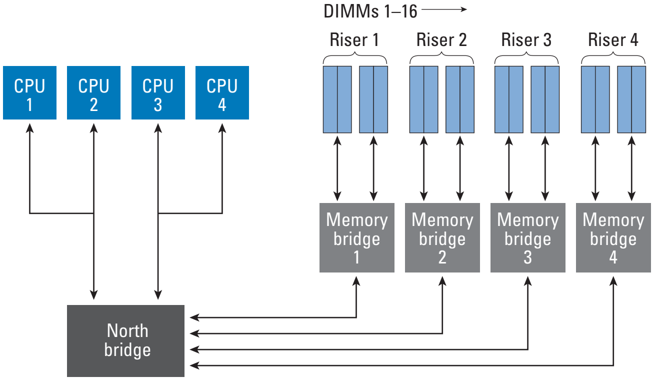

Figure 1. Dell PowerEdge 6850 planar block diagram

The PowerEdge 6850 offers three memory redundancy options, which are set in the BIOS: spare-bank memory, memory mirroring, and memory RAID.

As memory capacity increases, the number of soft errors will rise. Typically, a percentage of soft errors are multibit errors, which ECC cannot correct. As a result, administrators should expect the potential for failure in ECC systems to increase as memory is increased.

Memory redundancy options in PowerEdge 6850 servers

To help provide server fault tolerance, maximize memory capacity, and enhance reliability, Dell server hardware is designed with memory redundancy options that can help improve the performance and uptime of servers in memory error situations. The Dell PowerEdge 6850 memory subsystem resides on up to four memory riser boards, or *cards,* supporting a system maximum of 64 GB when 4 GB DIMMs are installed. Each riser has four double data rate 2 (DDR2) slots arranged logically as two banks and one memory bridge (see Figure 1).

The PowerEdge 6850 offers three memory redundancy options, which are set in the BIOS: spare-bank memory, memory mirroring, and memory RAID. Organizations select these options when ordering the PowerEdge 6850. They may also opt not to use these three options. To disable these options, organizations should select the Redundancy Disabled option. Figure 2 shows BIOS options for memory redundancy modes.

Spare-bank memory

Within a riser board, memory can be set as a spare bank. Once sparing is enabled, when the error rate—that is, the rate of correctable errors—of a failing DIMM reaches a preset threshold (set by the administrator in the BIOS), the DIMM's contents are copied to the spare bank. When the copy is in progress, all reads from the

BIOS options	Sparing	Mirroring	RAID	Hot addition	Hot replacement
Spare-bank memory	Support depends on memory card	Not supported	Not supported	Not supported	Not supported
Memory mirroring	Not supported	Supported if riser 1 and riser 2 have equal memory and/or riser 3 and riser 4 have equal memory (only memory mirroring is enabled)	Not supported	Not supported	Supported
Memory RAID	Not supported	Not supported	Supported if all four risers have equal memory (only memory RAID is enabled)	Not supported	Supported
Redundancy Disabled	Not supported	Not supported	Not supported	Hot addition in previously empty slots is supported	Not supported

Figure 2. Memory redundancy modes

memory space mapped to the failing DIMM will be serviced by the spare bank, while all writes will be made to both the failing DIMM and the spare bank. Once the copy is complete, the failing DIMM is disabled and the spare bank services requests to that memory space. The BIOS logs an event in the system event log (SEL) indicating that a spare-bank switch occurred.

The sparing process is transparent to the OS. When sparing is enabled, only three-quarters of the total installed memory is available to the OS. Uncorrectable errors will not generate sparing failover events; instead, such errors will cause the system to halt.

Memory mirroring

Memory on one riser board can be mirrored to memory on another riser board if total memory is identical on both. Mirrored memory consists of redundant copies of the system memory. If a multibit error is detected when memory is accessed, the system will not crash because uncorrectable errors are considered to be correctable when memory mirroring is enabled. In that case, the chipset will automatically try to regenerate the data from the redundant mirrored copy of the memory, which resides on the redundant system memory. (If the retry limit exceeds a certain threshold, redundancy will be lost.) The BIOS logs an event in the SEL indicating whether memory mirroring is enabled and another event if redundancy is lost.

Memory mirroring can provide a high level of fault tolerance. The minimum configuration required for mirroring is two riser boards and two identical DIMMs on each board. Mirroring is transparent to the OS. When mirroring is enabled, only half of the total installed memory is available to the OS.

Memory RAID

The Dell PowerEdge 6850 disk drive subsystem can be configured as RAID-5 in server memory, in much the same way that RAID-5 is implemented in storage devices. Implementing RAID-5 requires that all four risers be present in the system and populated with equal amounts of memory. RAID is transparent to the OS, and when enabled, less than half of the total installed memory is available to the OS. Usable memory will be three times the total effective memory present in each riser. Behavior of correctable and uncorrectable errors for RAID is the same as it is for memory mirroring. In addition to correctable and uncorrectable memory SEL events, the BIOS logs an event in the SEL when memory RAID is disabled and logs another event in the SEL when memory RAID is enabled.

The Dell OpenManage™ systems management suite presents administrators with a graphical view of the system components. Figure 3 shows the memory redundancy options on a PowerEdge 6850.

Hot-plug functionality in PowerEdge 6850 servers

Dell PowerEdge 6850 servers offer two hot-plug functions designed to improve the uptime of servers: hot replace and hot add.

Hot-replace functionality

Hot-replace functionality refers to the capability to replace a failed DIMM on a removed riser board while the system continues to operate. PowerEdge 6850 servers provide hot-replace capability without requiring OS support. In fact, the hot-replace process is transparent to the OS.

A riser board can be hot replaced—which entails removing the board from the system, replacing failed DIMMs, and reinserting the board—only if a memory mirroring or memory RAID configuration is enabled as described in the preceding "Memory mirroring" and

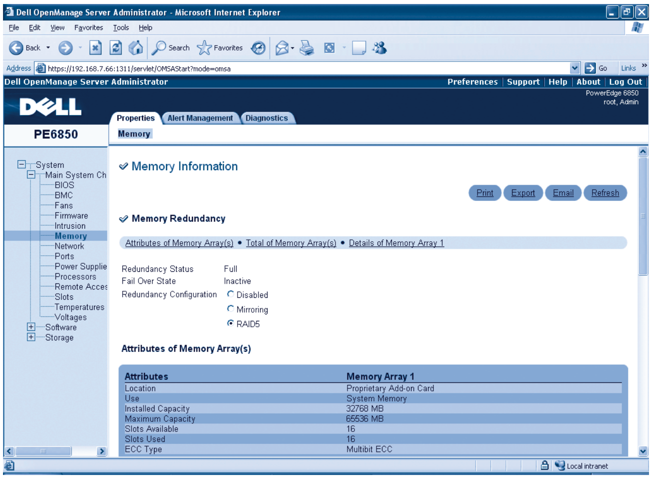
Dell OpenManage Server Administrator - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address https://192.168.7.66:1311/servlet/OMSAStart?mode=omsa
Go Links
Dell OpenManage Server Administrator
Preferences | Support | Help | About | Log Out
PowerEdge 6850
root, Admin
DELL
Properties
Alert Management
Diagnostics
PE6850
Memory
System
Main System Ch
BIOS
BMC
Fans
Firmware
Intrusion
Memory
Network
Ports
Power Supplie
Processors
Remote Acces
Slots
Temperatures
Voltages
Software
Storage
Print
Export
Email
Refresh
Attributes of Memory Array(s) • Total of Memory Array(s) • Details of Memory Array 1
Local intranet

✓ Memory Information

✓ Memory Redundancy

Redundancy Status Full
Fail Over State Inactive
Redundancy Configuration ○ Disabled ○ Mirroring ● RAID5

Attributes of Memory Array(s)

Attributes	Memory Array 1
Location	Proprietary Add-on Card
Use	System Memory
Installed Capacity	32768 MB
Maximum Capacity	65536 MB
Slots Available	16
Slots Used	16
ECC Type	Multibit ECC

Figure 3. Memory redundancy options shown on a system using Dell OpenManage

"Memory RAID" sections in this article. *Note:* It is not necessary for DIMMs to fail before the riser board can be hot replaced; administrators can hot replace riser boards as needed.

When the chipset detects a DIMM failure on memory riser 1, all memory accesses are routed to the healthy memory riser 2 of the mirrored set. The system automatically turns off the green power LED and activates the amber attention LED on riser 1, and then initiates the hot-removal event. Prompted by the amber attention LED, an administrator pushes the attention switch, which switches on the blinking green power LED and switches off the amber attention LED. After BIOS completes setting status registers, the system turns off power to riser 1 (indicated by the green power LED switching off). At this point, the administrator can safely remove riser 1.

Hot insertion to a failed mirror or RAID set is initiated when an administrator inserts a riser board into the slot from which the board was previously removed. After seating the riser board, the administrator pushes the attention switch. The green power LED starts blinking and the chipset begins its recovery initialization. BIOS then performs memory initialization. Upon successful completion, re-silvering begins. Once this process is complete, risers 1 and 2 are once again redundant. The solid green power LED remains on and the mirror/RAID LED turns on.

Reliability in memory starts at the DIMM level. To help ensure a reliable memory system in servers, it is essential to provide protection from both hard and soft memory errors.

The BIOS logs a memory removal event in the SEL when an administrator successfully removes a riser board, accompanied by a mirror or RAID redundancy lost SEL event (memory mirror redundancy requires two risers; memory RAID requires four risers). When an administrator successfully replaces a riser board, BIOS generates a SEL event indicating that memory mirroring or memory RAID redundancy has been regained as well as a memory-add event (see Figure 4). BIOS also generates a memory configuration error event if the current memory configuration does not support either adding or removing the riser board.

Hot-add functionality

Hot-add functionality allows administrators to increase memory size dynamically by adding a riser board to open slots. *Note:* A memory riser board should not be hot removed to add more memory. Only a previously empty slot may be used, and configuration rules must be followed.

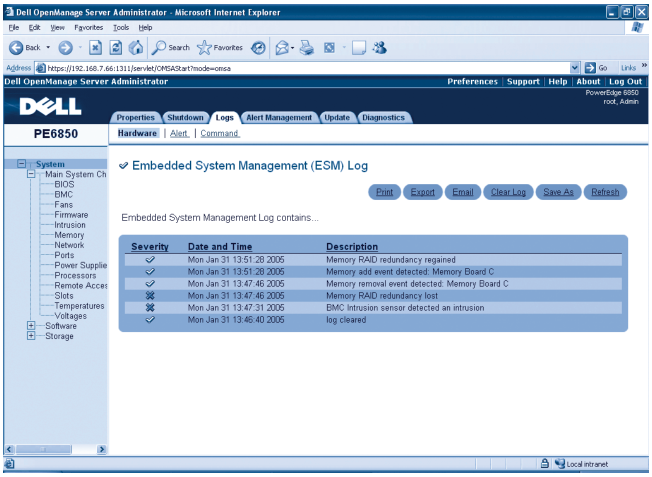
Dell OpenManage Server Administrator - Microsoft Internet Explorer
File Edit View Favorites Tools Help
Back Search Favorites
Address https://192.168.7.66:1311/servlet/OMSAStart?mode=omsa
Go Links
Dell OpenManage Server Administrator
Preferences | Support | Help | About | Log Out
DELL
PowerEdge 6850
root, Admin
Properties Shutdown Logs Alert Management Update Diagnostics
PE6850
Hardware | Alert | Command
System
Main System Ch
BIOS
BMC
Fans
Firmware
Intrusion
Memory
Network
Ports
Power Supplie
Processors
Remote Acces
Slots
Temperatures
Voltages
Software
Storage
Print Export Email Clear Log Save As Refresh
Local intranet

✓ Embedded System Management (ESM) Log

Embedded System Management Log contains...

Severity	Date and Time	Description
✓	Mon Jan 31 13:51:28 2005	Memory RAID redundancy regained
✓	Mon Jan 31 13:51:28 2005	Memory add event detected: Memory Board C
✓	Mon Jan 31 13:47:46 2005	Memory removal event detected: Memory Board C
✗	Mon Jan 31 13:47:46 2005	Memory RAID redundancy lost
✗	Mon Jan 31 13:47:31 2005	BMC Intrusion sensor detected an intrusion
✓	Mon Jan 31 13:46:40 2005	log cleared

Figure 4. ESM logs shown on a system using Dell OpenManage

A hot-add event is initiated when an administrator inserts a riser board into a previously empty slot. After seating the board, the administrator pushes the attention switch, the green power LED starts blinking, and the chipset starts the initialization process. BIOS then performs memory initialization, and the solid green power LED turns on.

Unlike hot-replace functionality, hot-add functionality requires the support of the OS. Upon successful completion of the hot add, control is transferred to the OS, which allocates memory and makes required changes to system properties. If an initialization step fails, power to the slot is turned off and all LEDs are turned off. Currently, only Microsoft® Windows Server™ 2003, Enterprise Edition, and Windows Server 2003, Datacenter Edition, support hot addition of system memory.

BIOS generates a memory-add event when an administrator successfully adds a memory card.

Notifying the OS about hot addition of memory

To notify the OS that the system is capable of supporting the hot addition of memory,[1] Dell BIOS provides a static resource affinity table (SRAT).[2] The SRAT also indicates to the OS the memory range that can potentially be hot added in the system.

To report added memory to the OS when an administrator adds a riser board, the BIOS instructs the chipset to generate a system control interrupt (SCI) after the added memory is

[1] To support hot addition of memory, Dell BIOS adheres to the best practices described in "Hot-Add Memory Support in Windows Server 2003" by Microsoft Corporation, www.microsoft.com/whdc/system/pnppwr/hotadd/hotaddmem.mspx.

[2] For more information about the SRAT, see the ACPI 2.0 specification at www.acpi.info.

initialized. The SCI causes control to be transferred to the OS Advanced Configuration and Power Interface (ACPI) driver. In Windows Server 2003, the ACPI driver calls the _Lxx control method of system BIOS. This control method notifies the OS that an administrator has added main memory to the system. Following this notification, the OS calls the _CRS control method of the memory device. The _CRS method returns the memory range added in the system.

SMBIOS provides an interface for motherboard and system vendors to present management information about their products in a standard format on Intel architecture systems.

SMBIOS update. Server management BIOS (SMBIOS) provides an interface for motherboard and system vendors to present management information about their products in a standard format on Intel® architecture systems.[3] SMBIOS includes information about hardware in the system such as CPU, memory, BIOS, and Peripheral Component Interconnect (PCI) devices. This information is stored in the SMBIOS table during the power-on self-test (POST). This table is designed to provide generic hardware instrumentation to deliver this information to management applications.

The SMBIOS specification is intended to provide enough information for developers of management instrumentation to create generic routines that translate from SMBIOS format to the format used by their chosen management method.

SMBIOS contains information about total memory in the system and the DIMMs present in the system. When a hot-plug operation is initiated, fields in the SMBIOS table need to be updated at the end of the operation—otherwise, the management software using SMBIOS information will contain outdated information. BIOS POST code updates the SMBIOS table to ensure that the management software using SMBIOS has current information.

Synchronizing MSRs and MTRRs

Memory Type Range Register (MTRR) is a processor feature that allows the processor to optimize memory operations for different types of memory, such as RAM, ROM, and memory-mapped I/O. MTRRs configure an internal map of how physical-address ranges are mapped to various types of memory.

MTRRs and Model Specific Registers (MSRs) should be synchronized among all the processors in a system. Anytime MTRRs or MSRs are updated for any processor, all the processors in the system should be in sync; otherwise, the system may hang. Because the amount of memory in the system increases during hot-add operations, MTRRs should be updated so that the processor is aware that memory was added. At the end of a hot-plug operation, the BIOS updates the MTRRs and all of the processors' MSRs and MTRRs will be synchronized.

The drive for IT reliability

To enhance reliability and enable 24/7 operations for servers supporting business-critical applications, administrators can combine server redundancy features—such as spare-bank memory, memory mirroring, and memory RAID—with hot-plug capabilities such as those available in the Dell PowerEdge 6850 server. As IT organizations are called upon to provide the memory fault tolerance necessary to minimize downtime, features that enhance redundancy and hot-plug memory capabilities are expected to become increasingly important tools in the administrative arsenal.

Chandra S. Mugunda is a senior development engineer in the Dell Instrumentation Software Group. Chandra has an M.S. in Computer Science from the Indian Institute of Technology, Roorkee, and a B.S. in Electrical Engineering from Andhra University in India.

Vijay Nijhawan is senior consulting engineer in the Dell Server BIOS Group. He is currently pursuing his M.S. in Engineering Management from The University of Texas at Austin. He has a bachelor's degree in Computer Science from MD University in India.

Dennis Stultz is the lead BIOS engineer in the Dell Server BIOS Group for enterprise systems. He has a B.S. in Electrical Engineering from Mississippi State University.

Saurabh Gupta is a BIOS engineer in the Dell Server BIOS Group. He has an M.S. in Computer Engineering from Texas A&M University and a B.S. in Electrical Engineering from Gujarat University in India.

Harish Jayakumar is a test engineer in the Dell OpenManage software development and test organization. He has an M.S. in Computer Science from Arizona State University and a B.S. in Computer Science from the University of Madras in India.

FOR MORE INFORMATION

ACPI:
www.acpi.info

SMBIOS:
www.dmtf.org/standards/smbios

[3] For more information about SMBIOS, visit www.dmtf.org/standards/smbios.

MEGABYTE:

What not having a
Linux strategy can take out
of your bottom line.

If you're paying unreasonable licensing fees for software that constantly needs security patches, you're getting eaten alive. But there's a solution. With SUSE® LINUX, Novell® can help you unleash the cost-saving power of a flexible, end-to-end open source strategy. Only Novell supports Linux from desktop to server, across multiple platforms. We'll integrate our industry-leading security, management and collaboration tools seamlessly into your environment. We'll provide award-winning technical support 24/7/365, and train your IT staff to deploy Linux-based solutions. And we'll make sure your open source strategy actually meets your number-one business objective – making money. Call 1-800-215-2600 to put some teeth back into your tech strategy, or visit www.novell.com/linux ➔ **WE SPEAK YOUR LANGUAGE.**

suse
A NOVELL BUSINESS

Novell®