Comprehensive Data Management Using Solaris™ Volume Manager Software

A Technical White Paper

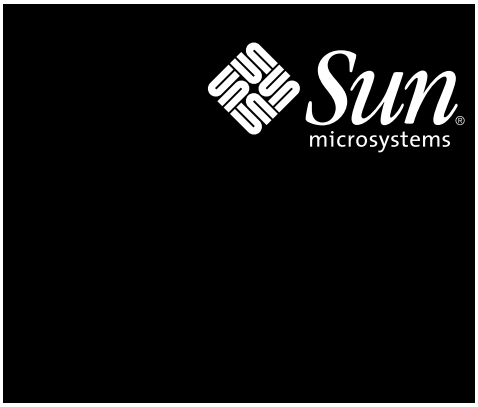

Table of Contents

Chapter 1

Comprehensive Data Management Using Solaris™ Volume Manager

To support mission-critical business applications, IT departments are faced with the challenging task of delivering continuous access to strategic corporate information assets. For many environments, access is required 24 hours a day, seven days a week, with periods of unplanned downtime often translating directly into lost revenues.

Today's data center managers must carefully administer the computing infrastructure — both systems and storage resources — to provide reliable, continuous data access. At the same time, they must operate within tight budget constraints, providing high service levels while containing personnel and equipment costs. In this regard IT managers face a daunting task — how to maintain high levels of service while simultaneously reducing the total cost of ownership.

Sun Microsystems, an industry leader in supplying solutions for mission-critical business computing, understands the challenges faced by data center managers. In recent years, Sun has focused on delivering products that can help reduce downtime, improve service levels, and lower the total cost of ownership. Today Sun offers a fully scalable product line with built-in availability features, and a reliable, mature platform — the Solaris™ Operating Environment — which is proven around the world in numerous mission-critical computing environments.

Integrated Storage Management — Solaris™ Volume Manager Software

With the introduction of the Solaris 9 Operating Environment, Sun is integrating key technologies that can help to enhance service levels even further, including a robust storage management solution known as Solaris™ Volume Manager. As described in this white paper, Solaris Volume Manager software can be used to configure multiple storage components into storage volumes, with redundancy and fail-over capabilities that can help provide continuous data access even in the event of a device failure. With an easy-to-use interface, the software greatly simplifies storage administration and allows many operations — such as recovering volumes or expanding the size of a file system — to occur online, minimizing the need for costly downtime.

Solaris Volume Manager software is based on the Solstice DiskSuite™ storage management tool, which Sun has offered for the past 10 years. Recognizing the importance of reliable storage management, Sun is incorporating the Solaris Volume Manager software directly into the Solaris 9 Operating Environment (OE), offering customers a comprehensive, highly integrated data management solution.

Solaris Volume Manager software is one of several key technologies that Sun is integrating into the Solaris 9 OE release (see the whitepaper, *"Better by Design — The Solaris 9 Operating Environment"* for a detailed discussion of what is included in this release). By integrating strategic technologies, Sun is providing a more complete, robust environment for deploying mission-critical, networked business applications.

Advantages of Integration

Incorporating the Solaris Volume Manager software in a single source base with the Solaris OE offers these advantages:

- Because key technologies are integrated, customers do not need to purchase additional software at additional cost. Solaris Volume Manager software gives customers a powerful solution for storage management, and its license is included with the operating environment.
- Integration means that Solaris Volume Manager software receives increased testing at all stages of development and maintenance in conjunction with the operating environment, which reduces deployment risk and enhances compatibility with future releases.
- Support issues are greatly simplified because customers can have one point of contact and require only a single maintenance contract.
- The upgrade process is clean and simple. For example, customers can migrate seamlessly from previous Solaris OE releases and earlier versions of Solstice DiskSuite software. Sun performs extensive testing for common upgrade scenarios, which helps to minimize deployment risks.

With the integration of the Solaris Volume Manager software into the Solaris 9 OE, Sun can help customers control costs, reduce risk, and continue delivering high levels of service.

Using Solaris Volume Manager Software for Data Management

The Solaris Volume Manager software incorporates the storage management functionality found in the 4.2.1 release of Solstice DiskSuite, along with several new features. These features combine to provide a comprehensive data management environment that can offer:

- Enhanced data availability
- Improved data reliability and integrity

- Better sustained I/O performance
- Greater configuration flexibility
- Simplified storage administration and management
- Lower deployment risk

This paper describes the features and benefits of the Solaris Volume Manager software, which are summarized in TABLE 1.

TABLE 1 Solaris Volume Manager software offers a comprehensive set of features.

Features	Benefits
Support for RAID 1 mirrored volumes and RAID 5 volumes (striping with parity)	Provides continuous data availability even when a disk device within the volume fails
Hot spare pools	Enables online system recovery
Alternate path support	Enhances data availability
State database replicas	Protects Solaris Volume Manager configuration information
RAID 0 striped volumes	Distributes the I/O workload over several disk devices to improve I/O performance
Soft partitioning	Provides greater configuration flexibility, enabling many partitions to be created on a single, high-capacity disk device
Disk concatenation and online expansion of volumes and file systems	Increases file system capacity without interruption or downtime
Device ID support	Preserves configuration information when disks or controllers are moved, providing greater flexibility
Support for disksets	Enables more effective namespace management
Graphical user interface (GUI) integrated with Solaris Management Console	Provides an easy-to-use, consistent GUI for storage administration
Command line interface (CLI)	Facilitates remote operations and scripting
CIM/WBEM API	Enables management of storage resources from any compliant tool
Storage monitoring	Simplifies administration and management of storage devices
Testing with Sun StorEdge™ storage products	Offers proven solutions that can reduce deployment risk
Upgrade support	Provides seamless upgrade process for Solstice DiskSuite software and earlier Solaris OE versions, minimizing downtime and risk

Enhanced Data Availability

Solaris Volume Manager software can enhance data availability through the use of RAID 1 mirrored volumes, RAID 5 volumes, hot spares, support for alternate pathing, and the use of state database replicas.

RAID 1 (Mirroring)

To protect valuable data assets, Solaris Volume Manager software allows multiple disk components to be configured into RAID 1 mirrored volumes (FIGURE 1-1:). In a mirrored configuration, data is duplicated (mirrored), over two or more physical disks. Data can be read from both parts of the mirror simultaneously to improve performance for read operations. If a physical disk fails, the surviving mirror provides continuous access to the data, and the mirrored copy can be regenerated on a hot spare component.

Figure 1-1: RAID 1 (Mirrored) Volume

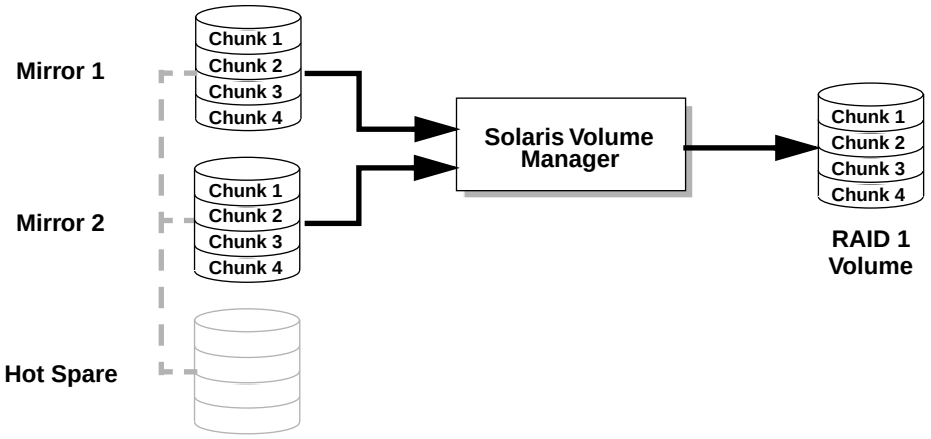

RAID 5 (Striping with Parity)

Solaris Volume Manager software also supports RAID 5 volumes, which use striping or segmentation to spread the data over multiple disks, and parity information to provide data redundancy. In a RAID 5 volume, every device has an area that contains a parity stripe or segment, while the other stripes contain data (FIGURE 1-2:). A RAID 5 volume can continue to provide data access when an underlying device fails, and a failed device can be automatically recreated on a hot spare without the need for downtime.

Figure 1-2: RAID 5 (Striped with Parity) Volume

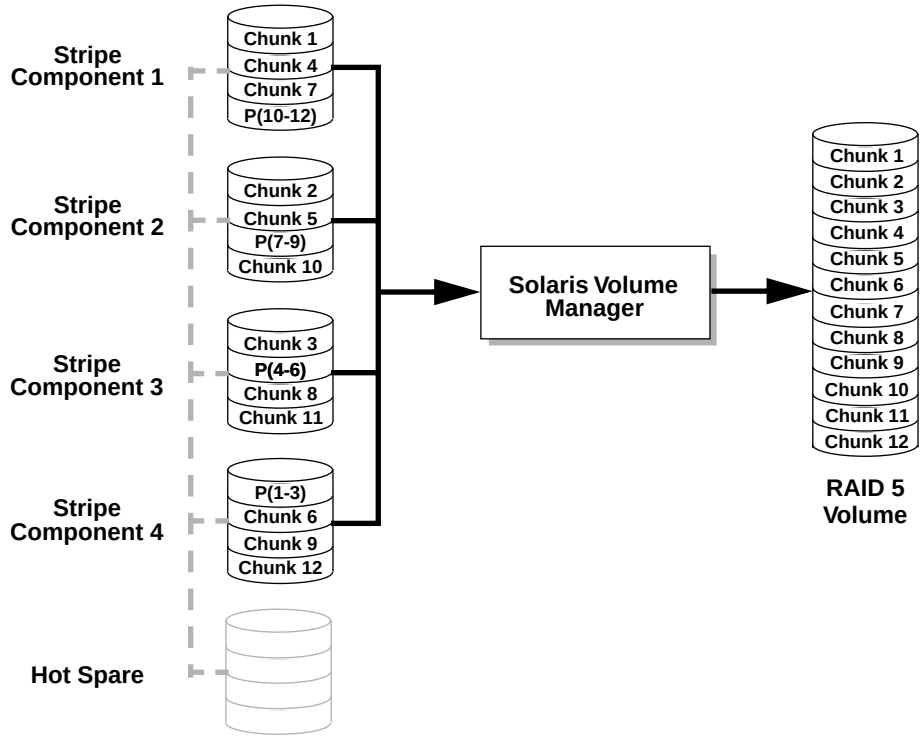

Hot Spare Pools

Hot spare pools enable online system recovery — without an interruption in service — when a mirrored or RAID-5 component fails. As shown in FIGURE 1-1: AND FIGURE 1-2:, a component configured as a hot spare can automatically replace a failed mirror or RAID 5 component. New slices are automatically created on the spare to replace the failed ones. Users can continue to access the surviving data (for example, the surviving mirror) while the replacement is generated on the hot spare component.

Alternate Path Support

Sun StorEdge Traffic Manager software is also integrated in the Solaris 9 OE, delivering multipath I/O capability that works in conjunction with Solaris Volume Manager software. Sun StorEdge Traffic Manager software offers the following features:

- *Single-Instance Multipath Device*. The device tree is restructured to represent all physical paths to a device as a single-instance, multipath device.
- *Load Balancing*. I/Os are balanced across all online paths in a round-robin manner, enabling increased I/O bandwidth to storage devices.
- *Automatic Failover.* I/Os are automatically failed over to the standby path(s) when all online paths have failed.

The Sun StorEdge Traffic Manager software uses the concept of a Virtual Host Controller Interface that implements multipath I/O services for SCSI-3 Fibre Channel devices (FIGURE 1-3:). It supports certain Sun StorEdge array configurations, including Sun StorEdge A5x00 arrays and Sun StorEdge T3 arrays. Solaris Volume Manager software works seamlessly with theSun StorEdge

Traffic Manager software Virtual Host Controller Interface, providing an alternate path to the data when the primary path to the data fails.

Figure 1-3: Alternate path failover

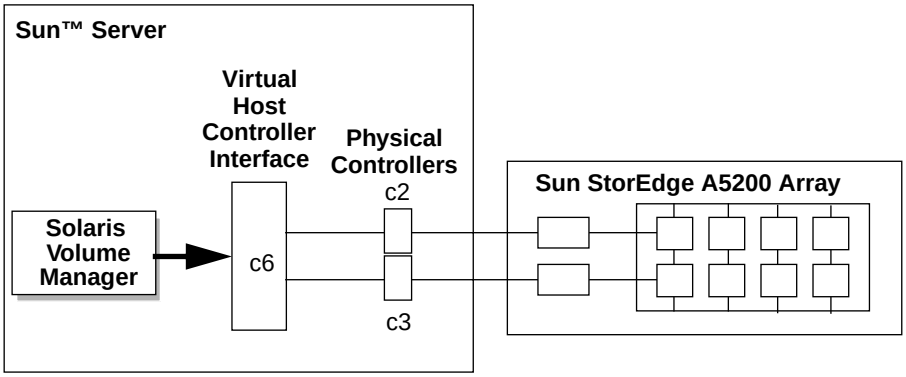

State Database Replicas

To enhance the availability of volume configuration information, Solaris Volume Manager software maintains a state database using a collection of multiple, replicated database copies. Multiple copies of the state database protect against data loss from a single point of failure. Each copy, called a state database replica, stores information about the state of the Solaris Volume Manager software configuration, and is updated constantly as the configuration changes. With the built-in redundancy of state database replicas, Solaris Volume Manager software protects important configuration information, which helps to improve data availability.

Performance

Solaris Volume Manager software supports disk striping and concatenated (RAID 0) volumes, which distribute the I/O workload across multiple devices to help improve performance. With striping, data is spread across relatively small, equal-sized fragments that are allocated alternately and evenly across multiple physical disks. A RAID 0 configuration alone does not provide redundancy — any single drive failure can cause data loss — but it can enable parallel I/O and load balancing for enhanced I/O throughput.

Solaris Volume Manager software also supports a combination of mirroring and striping (RAID 0+1 or RAID 1+0), depending on the underlying devices. (RAID 1+0 is automatically used when the underlying devices in the configuration can support it; for example, the devices are identical and not already configured as a virtual LUN.) Striping and mirroring are often used together to yield the combined benefits of better performance and greater data reliability.

Configuration Flexibility

Solaris Volume Manager software allows storage configurations to be easily modified and changed. These features provide greater configuration flexibility:

- Soft partitioning
- File system expansion
- Device ID support
- Support for disksets

Soft Partitioning

Soft partitioning enables an almost unlimited number of partitions and file systems on a single drive or volume, providing greater partitioning flexibility. In the past, the Solaris OE limited the number of partitions per drive by default to eight, which was generally adequate for the small-capacity drives that were available at that time. Today, 72-gigabyte drives are commonplace. As drive capacities increase, customers need to subdivide single, large-capacity disks into larger numbers of partitions.

With the soft partitioning feature in Solaris Volume Manager software, a disk can be subdivided into many slices that are controlled and maintained by software (hence the term "soft" partitioning). Soft partitioning allows up to 8192 partitions on a single volume, overcoming previous partitioning limitations. An administrator can create soft partitions either on top of individual disks, or on existing RAID 1, RAID 5, or RAID 0 volumes, simplifying the process of creating many file systems with the required data availability and performance characteristics.

Included with Solaris Volume Manager software and the Solaris 9 OE release, soft partitioning is also available through patches to the previous Solstice DiskSuite 4.2 and 4.2.1 versions. Consult a Sun service representative or the Sun support site, sunsolve.sun.com.

With soft partitioning capabilities, Solaris Volume Manager software is an extremely competitive storage management solution — one that gives customers tremendous flexibility and ease in configuring storage resources.

Volume and File System Expansion

Solaris Volume Manager software allows multiple physical components to be concatenated, enabling large volumes that can span across multiple drives. An administrator can use the graphical user interface (or the equivalent Solaris Volume Manager software commands) to create a new volume that consists of multiple components, or expand an existing volume, without the need for downtime or an interruption in service. The Solaris OE also provides a grow file system command (`growfs`), which may be used to expand a file system across multiple components while the file system is online.

Device ID Support

The configuration database for Solstice DiskSuite software — the precursor of Solaris Volume Manager software — relies on major and minor device numbers to describe the physical devices used in each volume. This method works well for storage systems that are static, but does not easily accommodate more dynamic systems and frequent configuration changes.

To address the needs of these dynamic environments, the Solaris 9 OE and Solaris Volume Manager software use unique device identifiers (device IDs) to define the physical devices used in each volume. By using device IDs in the configuration database instead of major and minor device numbers, Solaris Volume Manager software can automatically adapt to configuration changes, allowing disks to be more easily moved or rearranged. For example, when a controller is moved to a different slot or a disk is moved to a new location, Solaris Volume Manager software can still find and access these resources without the need to modify the configuration database. This support for device IDs helps to provide greater configuration flexibility and protects configuration and data resources.

Support for Disksets

Solaris Volume Manager software supports the concept of disksets. A diskset is a set of disk drives containing volumes and hot spares that can be shared (but not at the same time) between multiple hosts. An administrator can manage the volumes in a diskset as a separate namespace, allowing volumes to be easily configured and attached or detached from a node in a cluster.

Simplified Storage Administration and Management

System administrators are a valuable resource in any IT organization — one that must be conserved to contain the total cost of ownership. Data center managers must manage highly complex IT environments while controlling costs and staffing levels. To make it easier for administrators to accomplish routine tasks, and to simplify the configuration and management of storage resources, Solaris Volume Manager software supports both a command line interface (CLI) and a new GUI.

Consistent, Intuitive GUI

The new GUI for Solaris Volume Manager software (FIGURE 1-4:) is based on the Solaris Management Console — a Java™ technology-based interface that administrators use to access other administrative tools for the Solaris OE. The look and feel of this GUI is customizable, providing a consistent, intuitive interface for all Solaris OE management tools. Wizards guide administrators through easy-to-follow, step-by-step instructions, automating many storage management operations. The new Solaris Volume Manager software GUI allows system administrators to quickly review storage configurations and perform routine storage management tasks, freeing up their time for more proactive work.

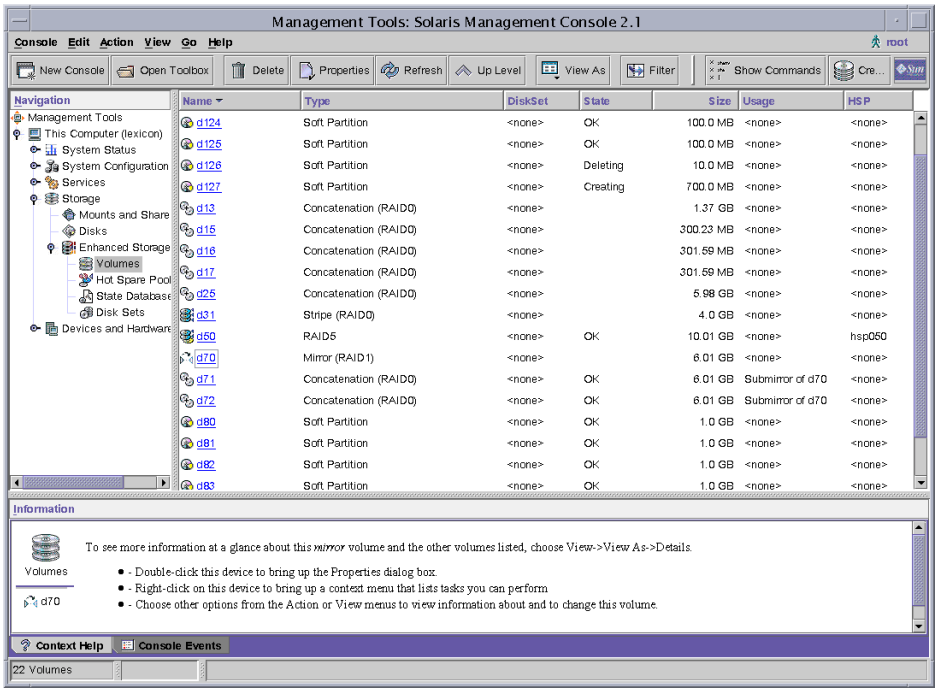

Figure 1-4: GUI for Solaris Volume Manager software

CLI and Command Logging

More experienced administrators sometimes prefer a CLI over a GUI. These administrators use a CLI to manage storage resources remotely or generate scripts that can replicate storage configurations. Using a command logging feature in the GUI, administrators can automatically capture CLI commands for a particular operation, which allows them to generate scripts quickly in order to automate repetitive tasks.

CIM/WBEM Application Programming Interface (API)

Solaris Volume Manager software includes support for an API that allows standards-based management of storage resources. This API adheres to the Web-Based Enterprise Management (WBEM) infrastructure and uses the Common Information Model (CIM) object model — standards which are specified by the Distributed Management Task Force (DMTF).

WBEM is an industry-wide initiative that includes vendor-independent standards for managing systems, networks, users, and applications using Web-based technology. It provides a way for management applications to share management data, and it is completely compatible with existing management protocols such as Simple Network Management Protocol (SNMP), Desktop Management Interface (DMI), and Common Management Information Protocol (CMIP). CIM defines the language and object model that compliant network management applications can access (see www.dmtf.org/standards).

For Solaris Volume Manager software, the CIM/WBEM API provides a industry-standard, programmatic interface to monitor and configure storage resources. Any CIM/WBEM-compliant tool can access and manipulate storage volumes that have been created using Solaris Volume Manager software.

Storage Monitoring

Active monitoring can remotely detect and report device failures, which provides convenience and better manageability for system managers. Solaris Volume Manager includes a software monitoring daemon (`mdmonitord`), which proactively monitors and probes storage volumes to identify device errors and failures. In addition, SNMP traps can signal events to network management tools such as Sun™ Management Center or third-party systems management applications.

Fewer Deployment Risks

In addition to comprehensive testing as part of the Solaris 9 OE, Solaris Volume Manager software is extensively tested with Sun storage products and in upgrade scenarios.

Compatibility with Sun Storage Products

Solaris Volume Manager software is tested and compatible with all Sun storage products, including:

- Sun StorEdge 9900 series
- Sun StorEdge 6900 series
- Sun StorEdge 3900 series
- Sun StorEdge T3 array, private loop and fabric modes
- Sun StorEdge A5200 array
- Sun StorEdge A1000/D1000 array

- Sun StorEdge S1 array
- Sun StorEdge D2 array
- Netra™ st A1000 array
- Netra st D1000 array
- Netra st D130 storage array
- Sun StorEdge D240 media enclosure
- Sun StorEdge Multipack and Unipack devices

In addition, Solaris Volume Manager software can function with other vendors' storage solutions running the Solaris OE. Contact the storage vendor for more information.

Upgrade Support

Solaris Volume Manager software is also tested to facilitate smooth transitions in operating environment and storage management releases. For customers migrating from Solstice DiskSuite 4.1, Solaris Volume Manager software is designed to enable seamless upgrades, including the automatic conversion of the configuration database replicas into the appropriate format. Even currently mirrored root disks can be upgraded without the need for modifying the configuration database. Since Solaris Volume Manager software is part of the Solaris 9 OE release, it will be constantly maintained and tested as future operating environment releases evolve.

Chapter 2

Conclusion

Sun has integrated Solaris Volume Manager software in the Solaris 9 OE to provide a powerful, comprehensive data management solution — one that is extremely competitive with other unbundled storage management products.

Solaris Volume Manager software allows storage resources to be configured for continuous data availability and performance, with features such as RAID level volumes, hot spares, and alternate pathing. New features, such as soft partitioning and the ability to move disks and maintain the configuration using device IDs, provide greater configuration flexibility than ever before. A new GUI offers a consistent, more intuitive user experience, simplifying the tasks involved in day-to-day storage administration. Wizards make it easy to configure and maintain storage volumes, and command logging allows scripts to be generated rapidly to automate repetitive tasks.

The integration of Solaris Volume Manager software into the Solaris 9 OE eliminates the cost of add-on storage management software and lowers deployment risks. Sun has qualified Solaris Volume Manager software with other elements of the operating environment, including upgrade scenarios and the family of Sun storage products, to facilitate smooth transitions in IT deployments. Solaris Volume Manager software is designed to ease the burden of configuring and maintaining data center storage, which can help to lower the total cost of ownership.

Sun is an industry leader that provides systems to support mission-critical applications, and offers a mature, proven operating environment known for its continuous availability and scalability features. With the Solaris 9 OE, Sun adds robust data management capabilities to help further enhance service levels. Solaris Volume Manager software can help IT managers deliver the continuous data access required for mission-critical applications, while reducing cost and deployment risk.

Chapter 3

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers at www.sun.com. Look for these and other white papers:

- *Better by Design — The Solaris 9 Operating Environment,* March 2002.
- *The Solaris Volume Manager Administration Guide* (Sun P/N 806-6111-10) describes how to manage storage resources using Solaris Volume Manager software. This document is available on docs.sun.com™, Sun's product documentation site.

For more details on the WBEM initiative and CIM specifications, see www.dmtf.org/standards/.

Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, CA 94303-4900 USA **Phone** 800 786-7638 or +1 512 434-1577 **Web** sun.com

We make the net work.

Sun Worldwide Sales Offices: Africa (North, West and Central) +33-13-067-4680, Argentina +5411-4317-5600, Australia +61-2-9844-5000, Austria +43-1-60563-0, Belgium +32-2-704-8000, Brazil +55-11-5187-2100, Canada +905-477-6745, Chile +56-2-3724500, Colombia +571-629-2323, Commonwealth of Independent States +7-502-935-8411, Czech Republic +420-2-3300-9311, Denmark +45 4556 5000, Egypt +202-570-9442, Estonia +372-6-308-900, Finland +358-9-525-561, France +33-134-03-00-00, Germany +49-89-46008-0, Greece +30-1-618-8111, Hungary +36-1-489-8900, Iceland +354-563-3010, India–Bangalore +91-80-2298989/2295454; New Delhi +91-11-6106000; Mumbai +91-22-697-8111, Ireland +353-1-8055-666, Israel +972-9-9710500, Italy +39-02-641511, Japan +81-3-5717-5000, Kazakhstan +7-3272-466774, Korea +822-2193-5114, Latvia +371-750-3700, Lithuania +370-729-8468, Luxembourg +352-49 11 33 1, Malaysia +603-21161888, Mexico +52-5-258-6100, The Netherlands +00-31-33-45-15-000, New Zealand–Auckland +64-9-976-6800; Wellington +64-4-462-0780, Norway +47 23 36 96 00, People's Republic of China–Beijing +86-10-6803-5588; Chengdu +86-28-619-9333; Guangzhou +86-20-8755-5900; Shanghai +86-21-6466-1228; Hong Kong +852-2202-6688, Poland +48-22-8747800, Portugal +351-21-4134000, Russia +7-502-935-8411, Singapore +65-6438-1888, Slovak Republic +421-2-4342-94-85, South Africa +27 11 256-6300, Spain +34-91-596-9900, Sweden +46-8-631-10-00, Switzerland–German 41-1-908-90-00; French 41-22-999-0444, Taiwan +886-2-8732-9933, Thailand +662-344-6888, Turkey +90-212-335-22-00, United Arab Emirates +9714-3366333, United Kingdom +44 0 1252 420000, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800

FE1818-0