

Managing More With Less — Maximizing System Administrator Resources

A Technical White Paper
November 2002



Table of Contents

Introduction	1
Solving Today's IT Challenges	2
Solaris System Management Solutions Overview	3
Simplifying Management	5
Sun Management Center 3.0	5
Solaris 9 Resource Manager	7
Sun Management Center Change Manager	7
Sun Cluster 3.0	8
Solaris 9 Operating Environment	9
Maximizing Existing Server Resources	11
Solaris 9 Resource Manager	12
Sun Management Center Change Manager	12
Sun Cluster 3.0	13
Solaris 9 Operating Environment	13
Managing Change	15
Sun Management Center Change Manager	15
Sun Cluster 3.0	17
Solaris 9 Operating Environment	17
Increasing Service Availability	18
Sun Management Center 3.0	18
Solaris 9 Resource Manager Software	20
Availability With Sun Management Center Change Manager	20
Sun Cluster 3.0	20
Solaris 9 Operating Environment	21
Conclusion	22
References	23

Chapter 1

Introduction

The Internet has dramatically changed the way companies provide services. The days of installing multiple copies of desktop PC or client-server software are yielding to a Services on Demand model, one where companies develop a single copy of their applications and deploy them to hundreds, thousands, and millions of users accessing the Internet from any IP-enabled device, anywhere, at any time. The shift toward a Services on Demand model moves the burden of administering the computing infrastructure from the end user and their PCs to the organizations providing the services, increasing the cost of management in environments with already limited IT budgets and staff.

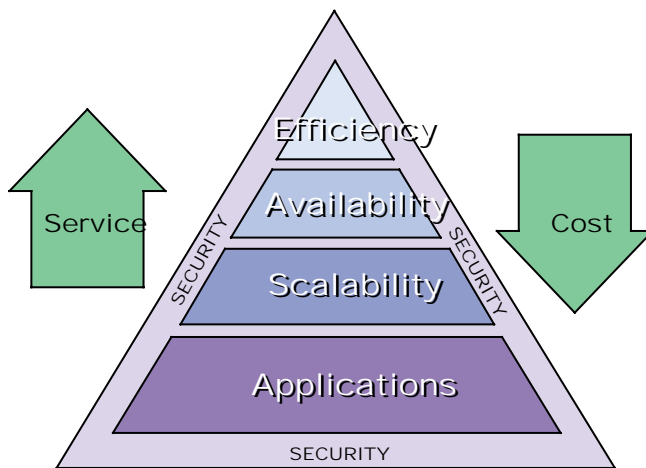
IT organizations are faced with competing goals to deliver these services. They need to consolidate multiple applications onto a single server to fully utilize system resources, and at the same time, supply each and every user and application with the resources they need. Server consolidation techniques and the use of service-level agreements (SLAs) are the latest trend for achieving these goals.

This challenging business environment, illustrated in Figure 1, leaves CTOs and IT departments faced with an increasingly difficult task. They must use a limited budget to maximize the return on existing IT investments — getting the most out of existing server infrastructure at minimum cost — in order to effectively support the Services on Demand model.

Doing more with less creates problems for IT managers and system administrators, including:

- Increasing service levels while simultaneously reducing service cost and risk
- Managing large numbers of complex servers, in a variety of different configurations
- Quickly adapting to changing workloads
- Providing access to critical applications around the clock — with service-level agreements and availability requirements that allow downtime measured in only minutes per year

Figure 1: Today's Computing Environment



Solving Today's IT Challenges

In order to solve these intricate challenges, IT managers and system administrators need:

- *Simplified Management:* IT organizations require an integrated set of systems management tools — a platform that offers common services for all enterprise management applications, with the ability to handle increasing numbers of users and more challenging workloads. Efficiency can be further increased with server monitoring and management tools that can automatically respond to a wide range of hardware and software faults, reducing the number of events requiring manual intervention by the system administrator.
- *Tools for Maximizing Existing Server Resources:* IT organizations need to consolidate servers and applications, improve resource utilization, and maximize the return on investment (ROI) on hardware that has already been purchased.
- *Change Management Tools:* Administrator efficiency can be greatly improved through tools for automating installation, deploying systems, and reducing the time required to execute routine tasks.
- *Service Availability:* IT groups need to ensure access to data, applications, and services. Availability is the cornerstone to delivering service-level agreements, minimizing both planned and unplanned downtime, and helping to ensure service levels for mission-critical applications.

Solaris™ System Management Solutions offer application and tools that excel at fulfilling these needs.

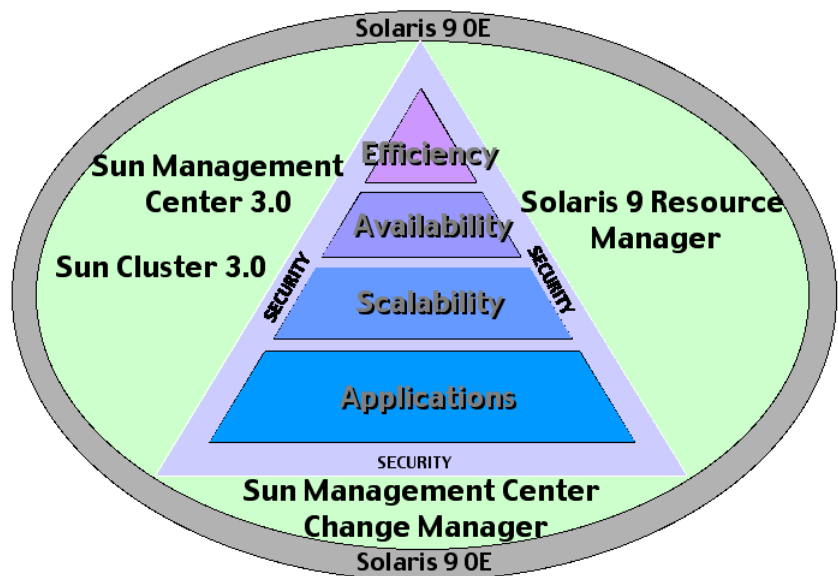
Solaris System Management Solutions Overview

Solaris System Management Solutions are powerful applications and tools specifically designed to provide IT managers and system administrators with the ability to efficiently and cost-effectively manage Sun™ servers and application services.

- *Solaris 9 Operating Environment (OE):* The Solaris 9 OE delivers the effective support for high-level infrastructures that organizations need to provide increased services with fewer resources. Through its highly available, scalable, manageable, and flexible features, the Solaris OE provides the foundation for the wide range of services necessary to support both enterprise and Internet-based Services on Demand environments.

- *Sun Management Center 3.0*: Sun Management Center is the most advanced systems management tool from Sun to date. It is designed to provide a powerful, easy-to-use single management point for all Sun servers, the Solaris OE, applications, and services for data center and highly distributed computing environments.
- *Sun Management Center Change Manager*: Change Manager enables administrators to quickly and easily install, configure, upgrade, provision, and audit the Solaris OE and applications.
- *Sun Cluster 3.0 Software*: Sun Cluster delivers end-to-end availability for applications, data, and networking, as well as enables scalable services.
- *Solaris 9 Resource Manager Software*: Solaris 9 Resource Manager enables IT managers to control and allocate system and network resources on a fine-grained basis, including resource containment, accounting and controls, workload classification, and differentiated services.

Figure 2: Today's Computing Environment With Solaris Systems Management Solutions



Solaris System Management Solutions, as shown in Figure 2, can help increase the ratio of servers to administrators by:

- Increasing efficiency and manageability by reducing system administrator time and resources spent on daily tasks
- Increasing availability, thereby reducing costly system downtime and making applications more available to the user community
- Managing and scaling resources, and reprovisioning the network, servers, and services, providing the flexibility to meet challenging needs and SLAs
- Offering an integrated, single vendor solution, enabling administrators to learn new tools with little training, and manage systems from anywhere on the network

The technical benefits of using Solaris System Management Solutions are outstanding scalability, higher availability, excellent manageability, and unencumbered flexibility.

Scalability

Scalability is critical to companies needing to make the most cost-effective use of computing resources, gracefully handle peak workloads, and grow the computing environment as rapidly as the business grows. Companies offering Web-based services have lost control over how many customers might visit their sites — and successful companies are those that can grow to support millions of customers almost overnight.

Availability

With businesses operating around the clock and around the globe, organizations have lost control of when their customers might demand their services, making availability of paramount importance. Solaris System Management Solutions help IT departments deliver 24x7 availability that is so crucial in today's always-demanding, global computing environment.

Efficiency

With business models changing rapidly, and service-delivery platforms growing by orders of magnitude as a matter of course, deploying and reprovisioning servers are daily events. Facilities to manage growth are a must, but today's IT departments must also reduce risk through and complexity by using fewer hardware and software components as well as by using servers that can be repaired or upgraded with no downtime.

Monitoring and adjusting CPU utilization and network bandwidth allocation enables application response time and service-level requirements to be met. In addition, storing data for trend analysis enables effective capacity planning and resource management. When systems operate at peak efficiency, administrators are better equipped to work more productively and make better decisions — factors that can mean the difference between merely surviving and thriving in an increasingly competitive marketplace.

Applications

Constantly changing and growing business models require unrestrained flexibility. Solaris System Management solutions provide administrators with the flexibility to:

- Manage any system from any site
- Quickly reprovision systems to meet demands
- Freely allocate resources when and where they are needed
- Install and upgrade software during normal working hours without incurring system downtime

The ability to perform all of these functions effectively is the driving force behind delivering dependable mission-critical application services.

Chapter 2

Simplifying Management

Sun understands the difficult tasks that system administrators face today. IT managers look towards automating and integrating systems management in order to manage ever-increasing workloads with fewer administrative resources. That's why Sun offers a set of management products that simplify administration by streamlining procedures, enhancing existing skill sets with more powerful tools, and enabling the creation of processes that prevent failures and increase efficiency.

Sun Management Center 3.0

Sun Management Center software is designed to provide a powerful, easy-to-use single management point for all Sun servers and storage, independent of geographic location. System administrators can perform remote system configuration and performance monitoring as well as isolate hardware and software faults through a single interface at any time, from anywhere on the network, improving administrator efficiency. Sun Management Center easily integrates with enterprise management frameworks, and provides a central facility for managing events and alarms, automated responses, and diagnostics. This gives administrators the ability to manage systems and perform corrective actions without leaving the central office.

Easy-to-Use Management Console

All events and alarms are sent to a single, easy-to-use, Java™ technology-based console, enabling rapid response to problems as they occur from any managed system. Sun Management Center is the ideal complement to enterprise management systems. It has extremely fine-grained insight into the operation of Sun servers, enabling a detailed analysis of any situation and notifying the administrator and other higher-level enterprise management systems.

Grouping Operations

Sun Management Center software helps increase productivity while decreasing errors by enabling administrators to implement logical element grouping. Grouping Sun systems by such categories as location, server role, and administrative responsibility helps administrators automate tasks to manage systems as a group rather than individually.

Automating Performance/Data Management and Application Monitoring

Sun Management Center software includes a variety of tools to help administrators proactively and automatically manage both common and complex and common administrative tasks, such as managing performance and data and monitoring applications. These tools increase efficiency, reduce the likelihood of costly errors, and help to ensure availability.

- *Health Monitoring*: Incorporates a large body of administrator knowledge into an intelligent rules-based monitor for the system and applications, giving suggested steps for problem resolution, resulting in simplified administration
- *Physical Viewer*: Displays photo-realistic images of hardware components, enabling junior administrators to quickly determine which components to replace
- *Logical Viewer*: Presents a tree hierarchy of managed hosts or domains, enabling administrators to quickly identify a component's exact location
- *Real-time Performance Analysis*: Enables administrators to isolate potential and existing bottlenecks
- *Event and Alarm Management*: Provides administrators with information when they need it
- *Automatic Discovery*: Enables fast inclusion of new systems into the Sun Management Center management framework

System Reliability Manager

System Reliability Manager is an optional add-on package for Sun Management Center 3.0 that helps proactively manage and monitor Sun servers and the Solaris OE.

- The *Script Launcher* feature in System Reliability Manger software enables administrators to conveniently manage scripts from any Sun Management Center console on the network. It makes it easier to manage a variety of scripts, increases security, and helps ensure successful script execution, reducing the risk of operator error when managing system scripts.
- The *File Watch Module* monitors an enterprise-wide list of system or configuration files for record deletion, addition, or modification, and detects syntax errors. This results in improved system reliability and availability. It proactively alerts administrators about missing system and configuration files, automatically monitors default system files, and enables administrators to specify additional files to be monitored.

Automating tedious, time-consuming, and error-prone tasks with Script Launcher and the File Watch Module frees administrators to do other tasks and manage more systems.

Performance Reporting Manager

Performance Reporting Manager is an add-on to Sun Management Center 3.0 that enables administrators to create reports detailing the status of their systems. Performance data is gathered by the history logging capability included with Sun Management Center software and used by Performance Reporting Manager to generate performance reports. The reports can be generated at schedule times and viewed via a Web browser. Examples of how Performance Reporting Manager can be used include:

- IT service managers can use reports on the availability of services and performance levels from the previous day to ensure that they are maintaining their SLAs.
- Performance Reporting Manager reports can be used for diagnostic and trend analysis, reporting trends leading to the point in time of a problem or outage. A unique feature of Performance Reporting Manager enables administrators to click on a spike in the report to see alarm data and/or processes that were running at that time, providing further information to help troubleshoot problems. The reports can also provide baseline information to help administrators tune thresholds for alarms or action.
- Performance Reporting Manager reports can be used as input for other capacity planning tools, and to do inventory analysis. The reports are designed to provide a very accurate, detailed inventory of hardware and software installed on Sun servers, which can be exported to third-party tracking systems.
- Performance Reporting Manager can generate a very detailed level of information regarding the demand and workload endured on an application service and the resource utilization of memory, CPU, and I/O by applications and service components. This information can be analyzed and employed to make decisions about application performance tuning and dynamic resource allocation.

Solaris 9 Resource Manager

Application consolidation provides the opportunity to lower costs by reducing the number of hardware components and system administration personnel required to run a set of applications. However, the problem of provisioning applications with the appropriate resources on a shared system can be a difficult task. Solaris 9 Resource Manager software simplifies this task by providing administrators with a powerful tool to meter, monitor, and control resource utilization, enabling them to deliver predictable service levels.

The Solaris 9 Resource Manager enables administrators to efficiently control system resources by creating resource pools that are reserved to be used exclusively by an application or set of applications. Using resource pools, the administrator can create isolated computing environments to help ensure that applications have access to a consistent set of resources — regardless of the resource usage of the remainder of the system. The administrator can also place caps on resource usage, preventing workloads from over-consuming resources.

Sun Management Center Change Manager

Installing and changing the Solaris OE and other applications are tasks that can consume a great deal of a system administrator's time. The Sun Management Center Change Manager console can drastically diminish installation time and complexity.

Easy-to-Use Change Manager Console

The Change Manager console can be accessed anywhere on the network via an easy-to-use Web browser interface or a command line interface (CLI). The key features of the Change Manager are:

- *Flash Archive Manager:* The Flash Archive manager enables administrators to import Solaris Flash archives (images of entire systems), group them into folders, delete and add functionality, search for an archive, and edit archive properties. This tool enables administrators to easily manage hundreds of different system profiles.
- *Target System Manager:* The Target System Manager enables administrators to perform single server or group operations for upgrades, software auditing, and reprovisioning services on target systems. These groups can be groups that were previously created using Sun Management Center.
- *Configuration Manager:* The Configuration Manager provides administrators with the ability to create and manage server configuration profiles and apply them to individual servers or groups of servers, reducing the complexity that currently exists with Solaris JumpStart™ technology.
- *Software Auditor:* The Software Auditor eases administration tasks by quickly managing changes across a farm of servers and helping to manage change reliably for upgrades and maintenance. It gives the administrators the ability to create a software stack manifest of one or more deployed servers. The manifest can be used in rule-based comparisons to check software configurations of a server or group of servers between two points in time — between the current state and the original state or between servers. It also includes the ability to write auditing rules such as detecting file content changes.
- *Task Scheduler:* Change Manager includes robust and flexible scheduling features that enables administrators to schedule service provisioning during off-peak hours. Operations such as upgrades and software auditing can be scheduled on a single server or group of servers, offering yet another level of time-saving automation.

Sun Cluster 3.0

The Sun Cluster 3.0 software architecture enables a potentially large pool of hardware and software resources to be managed as a single server, resulting in increased flexibility in choice of server, storage, interconnect, file system, and volume management technologies.

Cluster management can be accomplished via SunPlex™ Manager, a Java browser-based interface that allows complex tasks to be performed with ease. This enables administrators to manage any resource on a cluster from anywhere on the network where Java technology is available — from a different building, city, or even country. SunPlex Manager integrates easily with Sun Management Center, reducing the need to learn new tools and further reducing costs.

Resource groups can be defined to integrate applications with the cluster resource necessary for their execution. Clusters use the resource group specifications to monitor the health and performance of applications, allowing automatic failover in the event of a failure. Resource groups are also used to schedule and guarantee resource availability to applications while simultaneously maintaining the relative priority of each resource group. The result is a system that automatically manages itself to meet service level agreements.

Additionally, developers can use the SunPlex Agent Builder to develop Scalable or HA Agents. The SunPlex Agent Builder generates agent code with two simple clicks, reducing the time and cost of deploying application services by automating agent development.

Solaris 9 Operating Environment

The Solaris 9 Operating Environment offers many features designed to simplify management and administrative tasks.

Binary Compatibility

Sun's binary compatibility means that applications can move up the product line from desktop workstation to the most powerful servers that Sun sells — all without touching the application software or changing operating environments. Binary compatibility across the product line reduces administration and training costs because only one operating environment needs to be supported.

Integrated Directory Services

For user management, the Sun ONE Directory Server is integrated with the Solaris 9 OE, creating a Lightweight Directory Access Protocol (LDAP)-based environment that can scale to millions of users. The Sun ONE Directory Server is designed to provide comprehensive information asset management, including user names and attributes (such as passwords) as well as group, project, and task definitions. LDAP can link a complex, fragmented, and heterogeneous infrastructure by managing all of its network resources from a central, interoperable server. For those developing Web-based services, the Sun ONE Directory Server central repository provides the basis for storing user preferences, context, and attributes necessary for delivering services on demand.

Role-Based Access Control

Studies have found that human error causes more system downtime than physical failures. Unplanned downtime is an administrative time-sink that should be eliminated whenever possible. One way to minimize human error is to limit the scope of changes that different administrators can make. Using role-based access control (RBAC) in the Solaris 9 OE, administrative tasks can be delegated to different classes of administrators through a simple, centralized, easy-to-use graphical user interface (GUI). For example, a database administrator is allowed access only to necessary functions without having superuser access to the operating system. RBAC profiles only need to be defined once and are stored in the naming service (files, NIS, NIS+, and LDAP), which helps lower the cost of administration and allows delegated administration to be used anywhere on the network.

Modular Debugger

Sun's Modular Debugger (MDB) is an extensible utility for low-level debugging and editing of the live operating system, crash dumps, user processes, core files, and object files. The framework integrates a set of debugger modules, each of which is written to have specific knowledge of a kernel component or device. The result is a framework that speeds problem resolution by effectively managing information.

Solaris Patch Manager Tool

Solaris Patch Manager Tool is a patch tool for the Solaris OE that does automated intelligent analysis, dependency resolution, and secure patch transfer to reduce the time required to keep systems optimally configured for maximum availability. Solaris Patch Manager Tool increases the efficiency and accuracy of patch management on systems running the Solaris 2.6, 7, 8, and 9 Operating Environments.

Key features of Solaris Patch Manager Tool include:

- Automatic patch inventory and configuration-based analysis of target systems
- Automatic resolution of patch dependencies and patch order
- Simplified analysis, download, installation, and removal of patches on the Solaris platform
- Automatic security verification of digitally signed patches
- Integrated into the Solaris 9 Operating Environment
- Also available via download for the Solaris 2.6, 7, and 8 Operating Environments

Solaris Volume Manager

Solaris Volume Manager software, as well as software RAID, is bundled into Solaris 9 OE. It helps administrators maintain data integrity by enabling them to aggregate multiple disk volumes — which might also be large disk arrays — into single, large, logical volumes. The task of managing disk volumes is made even easier with a new GUI that is accessible through Sun Management Center software.

Soft Disk Partitions

Using built-in Solaris Volume Manager software, up to 8,192 soft disk partitions can be created per disk volume, enabling each user in a shared environment to have a personal, private virtual disk. In addition to providing privacy, soft partitions also create a firewall between users, and make it easy to establish different backup policies for each partition. Soft disk partitions can be expanded as long as free disk space exists in the underlying partitions.

Snapshot Backups

The Solaris 9 OE eliminates the need to take file systems offline for backups. Through UNIX® file system (UFS) snapshots, a temporary copy of a file system can be created so that a consistent point-in-time image of the file system can be backed up to tape.

MKFS Performance

Growth in database tables means growth in disk capacity, and for database management systems storing tables in UNIX file systems, creating a file system structure on a large disk array can be a time-consuming operation. With improvements to the Solaris 9 OE *mkfs* command, the time required to create a UNIX file system is reduced by up to a factor of ten.

Chapter 3

Maximizing Existing Server Resources

In order to do more with less, system administrators need to maximize existing server resources, improving resource utilization and increasing the return on investment (ROI) on hardware that has already been purchased. Consolidating systems and applications is the most efficient way to attain this goal. Maximizing server resources improves the ratio of managed servers to system administrators by allowing them to:

- *Allocate Resources With Greater Flexibility:* Consolidating applications onto fewer servers enables administrators to allocate existing resources on the fly between those applications.
- *Easily Provide Scalable Services:* Scalable services simply need more of the same hardware, requiring little to no additional administrative support.
- *Offer Predictable Service Levels:* Predictability means no unplanned downtime and meeting SLAs, making the administrator's life that much easier.
- *Effectively Manage Conflicting Demand for Limited Resources:* The ability to automate resource management can put resources where they are needed, when they are needed.

Solaris 9 Resource Manager

With Solaris 9 Resource Manager software, administrators can configure a single server to run multiple applications, yet maintain isolation between each application. Furthermore, the ability to partition a system into resource pools enables administrators to share parts of the system with different departments. Workloads within departments, such as database instances, can be labeled and managed, enabling administrators to modify CPU resource allocation policies to increase resource utilization and eliminate resource contention. These features enable administrators to offer more predictable service levels.

Another feature within Solaris 9 Resource Manager is the Fair Share Scheduler. When there is contention for a resource, the Fair Share Scheduler helps ensure that each user or application receives its fair share of CPU allocation. If users and applications are consuming less than the CPU allotments designated to them, the Fair Share Scheduler allows other active users and applications to take advantage of the unused processing power. This feature enables administrators to protect key applications and fully utilize resources.

Sun Management Center Change Manager

Using Sun Management Center Change Manager, administrators can instantly reprovise systems, giving them the ability to immediately respond to changes in workload in order to comply with SLAs. For example, an Internet service provider (ISP) can install servers for primarily e-mail services during one part of the business day, while a small number of servers provide Web hosting. Later in the day, as user demand for e-mail drops and Web hosting demand increases, the server farm can be reprovise so that more servers are reconfigured from e-mail servers to Web servers (using Solaris Flash archives and Solaris Live Upgrade software), as shown in Figure 3. Using Solaris Live Upgrade software to reprovise the servers in a cascading fashion, both the e-mail and Web hosting services are never offline.




Figure 3: Reprovising Servers to Meet Changing Business Needs and SLAs

Sun Cluster 3.0

Sun Cluster 3.0 software is a powerful service-level management platform. It allows a single application, or series of applications from a service, to run across multiple domains (see Solaris 9 OE below) and systems, providing scalable services. Scalable services are designed to allow access to the same data and network address, regardless of the physical node the application is running on. By adding more domains or systems to the cluster, the capacity of the application can be dynamically increased.

With a scalable service, instances of an application can be started on additional nodes at any time, making extra capacity automatically available to clients accessing the cluster. Similarly, instances of the application can be stopped at any time, enabling those system resources to be reallocated to different tasks. Additionally, the Prioritized Service Management (PSM) feature, tightly coupled with the pooling of resources, helps ensure better resource utilization for the entire cluster.

The global networking facilities designed into Sun Cluster software can perform internal load balancing and distribution of scalable service client requests, resulting in higher performance because load balancing within a cluster is more efficient and accurate than external load balancing.

Because a scalable service has multiple instances of an application running simultaneously and accessing the same data, service levels can be maintained in the event of any number of potential outages — either planned or unplanned.

Solaris 9 Operating Environment

The Solaris 9 Operating Environment is excellent in all aspects of scalability, providing rock-solid reliability as workloads increase, and delivering the most processing power available for applications for Sun systems. Because Sun maintains a constant focus on a single operating environment, companies can scale their services horizontally (by adding more servers) or vertically (by adding more computing resources to a single server) without worrying about whether their application software will run.

Solaris 9 OE has an unprecedented number of resource management facilities that enable administrators to manage service levels for services and networks. Depending on the environment, Dynamic System Domains can assign an entire electrically isolated operating environment to a single service, with multiple domains possible on a single server. Within a single server domain, multiple services can run in separate *containers*, each of which has fine-grained control over the service's resource consumption and performance. This resource management capability increases administrator productivity because they can consolidate many services onto a single server that is easier to maintain.

Dynamic System Domains

Dynamic System Domains are designed to provide high-level control over allocation of processors, memory, and I/O devices to separate electrically isolated domains on supported servers. The Solaris platform supports dynamic system domains, where system resources can be reallocated on the fly or on a scheduled basis to rapidly adapt to changing workloads. For example, Dynamic System Domains might be used to support online transaction processing during the day and long data warehouse queries at night.

Dynamic Reconfiguration

As workloads change rapidly over time, server capacities inevitably must be adjusted by adding and subtracting processors, memory, and I/O devices. Likewise, when tools like Sun Management Center suggest proactive replacement of marginally performing components, server configurations must also be modified. With Dynamic Reconfiguration, processors, memory, and I/O interfaces can be hot-plugged into the server and dynamically configured into or out of the operating environment's domain, enabling such critical components as CPUs to be replaced without incurring down time.

Note – Dynamic System Domains and Dynamic Reconfiguration are currently only available on Sun Fire™ V880, 3800-6800, 12K, and 15K servers.

Other features of the Solaris 9 OE that are intended to provide increased performance and scalability include:

- *Shared Libraries*: Allowing multiple applications to share large amounts of system software, helping to consolidate multiple functions on the same server.
- *Enterprise DHCP*: Sun's Enterprise Dynamic Host Configuration Protocol (DHCP) server enables host systems to receive IP addresses and network configuration at boot time, and is essential for organizations wishing to deploy networks on a massive scale.
- *Sun ONE Grid Engine Software*: Sophisticated software that aggregates available computing resources and delivers compute power as a network service. With Sun ONE Grid Engine software, administrators can treat a collection of distributed computing systems as a single, large computational resource and balance the workload across the systems.
- *Processor Sets*: Enables the administrator to partition a server's processors so that only a process or group of processes can use them.
- *Device Hot Plug Management*: Supports hot-pluggable disks, CPUs, memory, and I/O boards that can be connected to a variety of high-capacity storage devices and cluster interfaces. Sun understands that when administrators dynamically reconfigure servers, they must have the ability to hot-unplug as well as hot-plug devices.

Chapter 4

Managing Change

Running a business in a networked world and delivering Services on Demand means managing systems with the predictability of a data center and the agility of the Internet. Providing high levels of availability while quickly responding to change are now prerequisites for increasing profitability and retaining a competitive edge. Sun understands these needs, and is continually developing applications and tools to help IT managers and administrators to do more with the resources at hand.

Installing and maintaining software — operating systems, applications, and patches — are some of the most time-consuming and error-prone tasks facing system administrators today. Installing large numbers of similar servers can be complex, inconsistent, and inefficient. Upgrading and patching software can typically require hours of planned downtime and unplanned downtime if the upgrade fails.

With the Solaris System Management Solutions discussed below, administrators can move towards optimizing systems provisioning and change management by automating the practice of building, testing, deploying, and maintaining systems.

Sun Management Center Change Manager

Many of today's enterprises rely on horizontally scaled server farms to provide software services. It is common to see tens or even hundreds of replicated servers — each running an identical software configuration — that provide a set of services such as Web, e-mail, directory, caching, and other applications. The task of administering and provisioning a consistent software stack across this vast array of systems has been complex, time-consuming, labor-intensive, and error-prone. Until now.

Many organizations have large numbers of servers all delivering the same services, but each requiring their own similar yet unique software configuration. For example, a customer might provide Internet search services using 200 Sun servers, each one similarly configured and loaded with similar software so that the entire server complex can support millions of customers.

Managing this number of servers can be costly in terms of administrator time without automated tools. Sun Management Center Change Manager makes it easier to manage large numbers of servers in any type of environment.

Change Manager is designed to provide a fast and easy way to deploy software stacks securely to a single server or groups of servers using Solaris Flash archives, enabling formal IT procedures in a build-test-deploy-maintain model. The concept is simple: The administrator builds and tests a prototype server, captures the entire software stack (Solaris OE, middleware, applications, and configuration) as a single, integrated software image in a Solaris Flash archive, then uses the archive to install or upgrade a single server or a group of servers. In the case of the customer example above, the administrator would configure one system and simply copy the changes to all 200 of the systems, as opposed to manually installing and applying the changes to each of the 200 servers individually, one at a time.

Change Manager helps to drastically increase the number of servers an administrator can manage by:

- Reducing configuration complexity and installation time
- Improving deployment scalability
- Significantly saving time and manpower resources for server provisioning

Upgrading Existing Servers

Using Solaris Live Upgrade software, Change Manager enables automated software updates while servers continue to run, avoiding interruptions to services. Integrated to work with Solaris Flash technology, Solaris Live Upgrade software significantly reduces the downtime associated with an operating system upgrade or maintenance. This functionality provides a mechanism to install a Solaris Flash archive on an inactive boot environment while the active boot environment is fully functional and unaffected by the installation. When the process is finished, a simple reboot will migrate the system to the updated environment. Another reboot can roll the server back to the previous system software configuration in the event that the upgrade has errors or is unsatisfactory. Again, when administering hundreds of similar systems, this functionality can save a lot of time and system downtime.

Reprovisioning Existing Servers

Using Change Manager, administrators can easily reconfigure a server from one type to another in response to service-level demands and to meet SLAs. For example, at peak traffic periods on a Web site, the administrator might reconfigure a cache server into a Web server to increase the level of the Web server pool, sustaining a peak load.

Providing Revision Control

Using Change Manager, administrators can implement a high level of revision control for applications deployment. For example, revision 1.0 of a Web server — consisting of the OS, patches, Web server, and applications — can be tested and deployed to n servers. As new patches and revisions of software arise, the administrator can test modifications to the Web server stack and label it revision 1.01, deploying the update to n servers as time permits.

If a problem is discovered in production with v1.01, the administrator can roll back to v1.0. Treating the complete application stack as a single, manageable object is a key step to providing a build-test-deploy-maintain model and increasing efficiency.

Sun Cluster 3.0

Sun Cluster scalable services can help administrators deploy multiple instances of an application. For example, more nodes can be added as the need for higher service levels grows.

The Global File Service (GFS) can be used for managing application changes. The administrator can install multiple versions of an application on the GFS and then start and/or stop them when necessary. If a new version of an application is unsatisfactory for some reason, the administrator simply stops it and restarts the previous version. This functionality enables the administrator to upgrade application software without having to reboot the system.

Sun Cluster software also makes updating applications easier. For example, it is far easier to update SAP application server versions if they are deployed in a Sun Cluster environment because only one set of SAP scripts needs to be updated. In a nonclustered environment, each server would have to be updated individually.

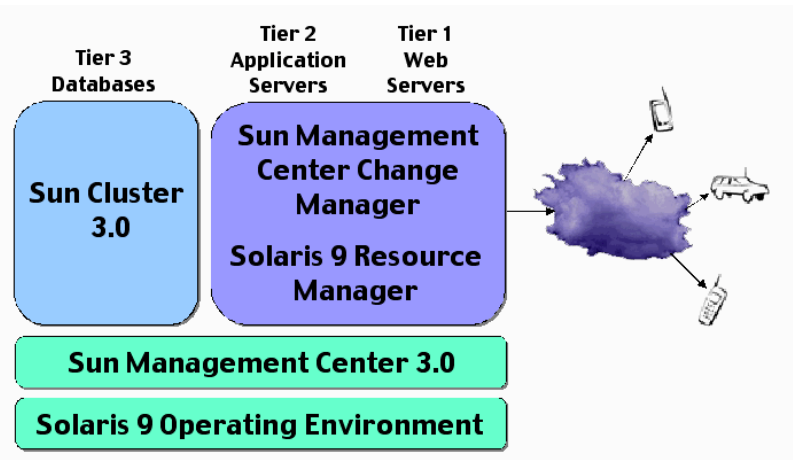
Solaris 9 Operating Environment

Solaris Patch Manager Tool

Patch management features in the Solaris 9 OE are the most comprehensive ever offered for Solaris software administrators — and they are designed to help reduce costs through increased efficiency. With these patch management functions, administrators can easily analyze the patch state of a system and automatically download the recommended patches. The administrator is provided with the install order necessary to accommodate patch dependencies. Administrators can use the patch management tools on local and remote systems. All patches delivered through Solaris Patch Manager Tool are digitally signed to help ensure that the patches came from Sun and have not been altered in transmission.

Figure 4 illustrates how Solaris System Management Solutions can be applied across the tiers of a service architecture to manage change.

Figure 4: Using Solaris System Management Solutions to Manage Change



Chapter 5

Increasing Service Availability

The advent of globally accessible Internet applications has given rise to the need for continuously available, 24x7 services. Companies can no longer predict when clients will demand their services, making it difficult to impossible to schedule planned downtime, and making unplanned downtime unacceptable and very costly in terms of lost revenue and reputation.

Solaris System Management Solutions help provide continuous service availability by:

- *Enabling System and Application Failover:* Increasing profitability by offering reliable services
- *Minimizing Planned and Unplanned Downtime:* Increasing availability and administrator efficiency and decreasing overtime costs
- *Increasing Service Predictability and Resource Availability:* Allocating resources when and where they are needed

Sun Management Center 3.0

Sun Management Center increases availability by establishing an agent on each managed server that allows it to evaluate both operating environment parameters (such as CPU utilization) and hardware parameters (such as power supply temperatures). This enables the tool to do predictive fault analysis, monitor performance, and isolate hardware and software faults — all increasing server availability.

Sun Management Center can direct the execution of server diagnostics and take pro-active steps to prevent problems. It can, for example, notify an administrator when diagnostics indicate an imminent failure of a CPU or disk drive, or when power supply operational parameters go out of range. Identifying problems before they result in a server shutdown helps increase uptime. It can also perform predictive failure analysis, enabling administrators to predict potential memory or disk hardware failures on a statistical basis, enhancing decision making and increasing availability.

System Reliability Manager

The OS Crash Dump Analyzer included in the System Reliability Manager helps administrators monitor and analyze OS crash dump data to find the root cause of the problems and take corrective action. This feature is crucial in preventing future crashes and helping to ensure higher system reliability. It increases system availability by proactively notifying the administrator (via pager, e-mail, etc.) when a crash dump is detected.

Service Availability Manager

The Service Availability Manager, another add-on application for Sun Management Center, monitors the availability of the following Internet services:

- Web Servers utilizing HTTP protocol
- Directory services using LDAP/DNS/NIS protocols
- Mail services using SMTP/IMAP4/POP3 protocols
- FTP services
- Telnet protocol
- Solaris Calendar Services

To optimize the availability and accessibility of the services running on the Solaris OE, Service Availability Manager enables system administrators to:

- Constantly monitor and confirm service availability at the end user's site
- Continuously measure and validate the delivered quality of service against the user's predefined performance thresholds
- Proactively manage and receive notification of service availability problems
- Set alarm thresholds on the monitored properties

To help administrators provide consistent service and save time tracking and troubleshooting potential problems, Service Availability Manager measures and displays:

- Specific configurations for a service request
- Service availability and uptime
- Total service processing time, network time, response time, resolution time, and connect time
- Data transfer rate

Solaris 9 Resource Manager Software

A server that is operating normally but whose applications are resource-starved may as well be down. Solaris 9 Resource Manager software greatly enhances application availability by helping to ensure that available resources — including network bandwidth, CPU utilization, and physical memory consumption — are divided so that one application cannot impair another's ability to provide services. This increases availability and service predictability by allocating of resources, reducing costs, and giving IT managers the ability to offer consistent resource availability.

Note – Support for managing network bandwidth within Solaris 9 Resource Manager software is expected to be available in Solaris 9 OE Update 1.

Availability With Sun Management Center Change Manager

Sun Management Center Change Manager enables administrator to update software while servers continue to run, avoiding costly interruptions to services. Using Solaris Live Upgrade software, Change Manager significantly enhances service levels by:

- Allowing software upgrades to be performed on an alternate system disk partition while the system continues to run
- Enabling the administrator to roll back to a previous configuration if an upgrade is unsatisfactory or fails
- Managing multiple boot environments for system maintenance or software upgrades

Additionally, Change Manager enables administrators to audit the software stack of any single server or group of servers to compare them to a reference configuration. This functionality improves reliability and security by providing the administrator with the ability to track changes, who made the changes, and when the changes occurred. It also enables the administrator to reset the system back to a known state.

Field Replaceable Unit (FRU) Servers

Replacing a failed system is often much faster and efficient than trying to troubleshoot it. Change Manager enables administrators to apply an FRU approach to a complete server by installing a Solaris Flash archive from the Flash Archive Manager onto a spare server, providing a replacement server in a relatively short period of time.

Sun Cluster 3.0

Sun Cluster software enables two or more servers to work together as a single entity to provide high-availability applications, system resources, and data to users. The Sun Cluster architecture is designed so that no single point of failure can bring the cluster down, with many multiple points of failure protected. Redundant high-speed interconnects, storage system connections, and public networks help ensure that cluster connectivity is immune from single failures. The software failover feature automatically and transparently migrates an application to another server in the cluster in the event of a failure so that clients continue operation uninterrupted.

Solaris 9 Operating Environment

Sun designed the Solaris Operating Environment with a small, compact kernel that limits the exposure to errors that can bring a system down. Its ability to quickly recover from errors is due in part to powerful features such as, memory scrubbing, file system journaling, NFS failover, IP network multipathing, and I/O multiplexing.

Memory Scrubber

Sun servers are built with ECC memory throughout, enabling them to correct any single-bit error and detect double-bit errors. Detecting and correcting single-bit errors early reduces the possibility of a single-bit error later becoming a double-bit error. The cache and memory scrubbers built into the Solaris OE scan for undetected single-bit errors, writing back correct data so that subsequent program references succeed without error. The memory scrubbers note when memory errors are persistent, enabling administrators to replace defective memory modules before they can cause a system outage.

File System Journaling

File System Journaling is bundled with the Solaris 9 OE, increasing recoverability by reducing file system check times when the operating environment is rebooted. Journaling writes every file system transaction to a journal before writing it to the actual file system structure, protecting against file system corruption and making recovery from unexpected shutdowns easier.

NFS Failover

NFS failover enables client systems to retry network file accesses on an alternate server in the event that the primary server fails. This feature enables the Solaris OE to continue executing applications — even when its NFS server goes down.

Network Multipathing

Solaris IP Multipathing (IPMP), Sun's Network Multipathing implementation for the Solaris 9 Operating Environment supports multiple paths for IP networks. Under normal operations, network traffic is balanced across multiple devices. If a path becomes unresponsive, due to a failed network interface card or an upstream switch or router, IPMP redirects network activity to the devices that remain operational. When the problem is repaired, IPMP software restores traffic on the interface. This feature enables the Solaris OE to automatically recover from single point of failures in a network infrastructure.

Sun StorEdge™ Traffic Manager

Sun StorEdge™ Traffic Manager software in the Solaris OE supports multiple paths for I/O devices such as SCSI and Fibre Channel-accessible storage. Not only does I/O multiplexing balance the workload across multiple devices, it increases reliability by redirecting requests from a failed interface card or storage device to one that is operational.

Chapter 6

Conclusion

Today, IT departments are facing tighter budgets and increased expectations, forcing them to accomplish even more with thinly stretched resources. They need solutions for simplifying management processes, maximizing existing resources, managing change, and providing higher degrees of service availability. When these solutions are implemented the right way, they can contribute significantly to reducing the time and money spent by IT departments on managing system infrastructures.

Sun servers capitalize on Sun's strength in scalability, availability, manageability and flexibility. Enterprise management tools, however, must be designed to understand the systems they intend to manage — and no one knows how to manage Sun hardware and operating systems better than Sun. Because Sun develops both the processors and the platforms that run Services on Demand, it is free to make processor- and architecture-specific software optimizations necessary to help provide an automated, integrated approach to system management. As a result, Solaris System Management Solutions provide administrators with the most comprehensive tools for managing Sun systems, offering:

- Server consolidation, resulting in fewer, more cost-effective servers
- More efficient use of administration staff
- The most scalability, availability, manageability, and flexibility available anywhere

For organizations taking and maintaining a competitive edge by delivering Services on Demand, Solaris System Management Solutions are ahead of the competition, providing services companies need today to cost-effectively manage systems into the future.

Chapter 7

References

Sun Microsystems posts complete information on Sun's hardware and software products and service offerings in the form of datasheets, specifications, and white papers at sun.com.

The following white papers are available at sun.com/software/whitepapers, and address a variety of Solaris System Management Solutions:

- Provisioning and Change Management Tools
- Sun Management Center — Managing the Integrated Enterprise
- Better By Design — The Solaris 9 Operating Environment
- Solaris 9 Resource Manager

SUN™ Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, California 95054 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, Solaris, Solaris JumpStart, Sun Fire, SunPlex, and Sun StorEdge are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

SUN™ Copyright 2002 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95054 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, Java, Solaris, Solaris JumpStart, Sun Fire, SunPlex, et Sun StorEdge sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Please
Recycle



Adobe PostScript

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA Phone 800 786-7638 or +1 512 434-1577 Web sun.com



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 +421-2-4342-9485, 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 +44-1-276-20444, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800 11/02 FE1913-0