


*A Developer's Perspective on Sun's
Solaris 8 Operating Environment
Reliability, Availability,
Serviceability (RAS)*



February 2000


Prepared for Sun by:

D.H. Brown Associates, Inc.
222 Grace Church Street
Port Chester, NY 10573
(914) 937-4302
Fax: (914) 937-2485



TABLE OF CONTENTS

EXECUTIVE SUMMARY	1
<i>UPTIME AND SECURITY</i>	<i>1</i>
<i>PROBLEM RESOLUTION AND CHANGE MANAGEMENT</i>	<i>2</i>
<i>SYSTEM-MANAGEMENT FEATURES.....</i>	<i>2</i>
<i>SUN CLUSTER</i>	<i>2</i>
SUN BUILDS ON ITS UNIX SUCCESS	3
UPTIME AND SECURITY	3
UNPLANNED DOWNTIME	3
PLANNED DOWNTIME	4
SECURITY	4
PROBLEM RESOLUTION AND CHANGE MANAGEMENT	5
SYSTEM-MANAGEMENT FEATURES.....	5
RESPONSE TIMES	6
<i>SOLARIS RESOURCE MANAGER</i>	<i>6</i>
<i>SOLARIS BANDWIDTH MANAGER.....</i>	<i>6</i>
ADMINISTRATIVE INTERFACE	7
<i>SOLARIS MANAGEMENT CONSOLE.....</i>	<i>7</i>
<i>SOFTWARE DEVELOPMENT KIT</i>	<i>7</i>
SUN CLUSTER	8
SUN CLUSTER 2.2.....	8
SUN CLUSTER 3.0.....	9
<i>GLOBAL DEVICE NAMING</i>	<i>9</i>
<i>CLUSTER FILE SYSTEM.....</i>	<i>9</i>
<i>GLOBAL NETWORKING.....</i>	<i>9</i>
DEVELOPING HIGHLY AVAILABLE, CLUSTER-AWARE APPLICATIONS.....	10
<i>APPLICATION REGISTRATION WITH CLUSTER SOFTWARE.....</i>	<i>10</i>
CONCLUSIONS AND RECOMMENDATIONS.....	12



A Developer's Perspective on Sun's Solaris 8 Operating Environment Reliability, Availability, Serviceability (RAS)

EXECUTIVE SUMMARY

Fueled by ongoing successes in enterprise markets, Sun has generated significant momentum and leading mindshare in Internet and other mission-critical installations. Independent software vendors (ISVs) find Sun's platform increasingly attractive because the enterprise markets in which it is a major player comprise a key battleground for their ongoing economic success.

Recent announcements from Sun further strengthen the company's position. With enhanced Sun Enterprise server hardware and the new Solaris 8 Operating Environment,¹ Sun now offers a host of new features, procedures, and technologies that directly target the needs of leading-edge users² and represent priorities for developers as well. New features that enhance Sun's rich suite of tools and its environment for independent software developers include:

- Support for higher availability with customizable security – critical factors in boosting application availability (which depends on server uptime)
- Control of resource allocation to achieve predictable response times (combined with security, response-time management represents a key differentiator among deployment platforms for e-commerce applications)
- Tools that help developers maximize ease of use
- Clustering software that enhances availability, scalability, and ease of use

UPTIME AND SECURITY

Sun continues to implement features that further minimize unplanned and planned downtime.

- Sun Enterprise server hardware includes reliability, availability, and serviceability (RAS) characteristics such as ECC memory, RAID, hot-swap support, and other features, as well as techniques for predictive-failure analysis and remote hardware diagnosis.
- Solaris 8 introduces identification and lockout of bad components upon startup or during operation, capabilities that supplement the company's existing leadership in support for online repair (Dynamic Reconfiguration and Alternate Pathing).
- Customizable security features incorporate state-of-the-art IPSec-based encryption and user authentication (including smart-card support) to protect against downtime caused by compromises of the system.

¹ For the purposes of this paper, the Solaris 8 Operating Environment will be referred to as Solaris 8.

² For more detail, see *Sun Server Hardware and Software and the New Solaris 8 Operating Environment Optimized for Mission-Critical and E-Business Environments*, D.H. Brown Associates, Inc., January 2000.

- Dynamic Reconfiguration, Alternate Pathing, and hot-plug support prove useful in repair, but also enable online expansion of systems to add capacity or enhance functionality.

PROBLEM RESOLUTION AND CHANGE MANAGEMENT

Sun provides problem-resolution capabilities to assist administrators in returning systems to service from either planned or unplanned outages.

- A Live Upgrades capability improves management of changes to server configuration, allowing backout without loss of information.
- Hot-diagnostic capabilities implemented in the Solaris 8 kernel assist in troubleshooting.
- Solaris 8 patch-management software enhances tracking and backout of system changes.

SYSTEM-MANAGEMENT FEATURES

Maximizing application availability also entails being able to control a system and guarantee predictable response times. Sun offers features that organizations can use to manage system and network resources.

- Solaris Resource Manager software enables allocation and prioritization of system resources to applications, users, or groups of users, complementing hardware features such as Dynamic Domains.
- Solaris Bandwidth Manager extends Solaris Resource Manager functions to the network.
- A consistent, easy-to-use interface in Solaris Management Console minimizes the chance of administrative mistakes.
- A software development kit (SDK) enables ISVs to develop interfaces that integrate with Solaris Management Console to leverage its infrastructure services. The SDK includes an application GUI toolkit that enables ISVs to match the look and feel of Solaris.

SUN CLUSTER

Sun's current and upcoming clustering products feature additional availability and management capabilities, some specifically targeted at helping ISVs leverage cluster features.

- Sun Cluster 2.2 not only provides high availability, but also supports application scalability.
- The upcoming Sun Cluster 3.0 will greatly simplify implementation and management of cluster failover with cluster-wide devices, a cluster file system, and global networking.
- Introduction and support of logical host capabilities enables ISVs to take full advantage of clustering capability for applications. The Sun Cluster Data Service APIs register applications with the cluster to support failover.

SUN BUILDS ON ITS UNIX SUCCESS

Sun continues to lead as a platform of choice among enterprises deploying Internet and other mission-critical applications. The company won this strong position by maintaining a focus on UNIX, which paid off as the company built on its lead in the UNIX server market both in terms of units sold and revenue. In 2Q99, the company had growth rates of 25% in revenue and 34% in shipments.³ Sun also holds a 35% market share in the midrange server market, which has boomed in recent years fueled by web hosting and e-commerce growth. Sun's presence in these server and midrange markets represents an attractive target for ISVs, especially with the explosion of e-business. In fact, these markets will likely prove an essential factor in the future economic well-being of ISVs.

Continuing developments in Sun Enterprise server hardware and the new Solaris 8 Operating Environment further increase the company's momentum. New features for developers allow easy exploitation of the capabilities of Sun Enterprise servers, the Solaris Operating Environment, and complementary products. The result: highly available applications that are easy to use, can seamlessly integrate into the administrative environment, and can be easily supported.

UPTIME AND SECURITY

Users now place a significant value on making sure their systems are both available and secure. The growth of global trade and the Internet, as well as the need for rapid response to customers' requests, have made 24x7 operation a nearly universal concern. Today, the Internet and other networks expose critical corporate data and information from suppliers, partners, and customers to an unprecedented extent. Industries that a few years ago paid no attention to these issues now make them a priority. Industries that have traditionally focused on availability and security – such as banking, finance, and telephony – now define them as critical requirements.

For most application developers, increasing the availability and security of applications to meet user needs is critically dependent on the capabilities of the platform on which the application runs. Sun continues to develop existing features and to introduce new capabilities that improve uptime and security.

UNPLANNED DOWNTIME

To mitigate hardware failures, Sun Enterprise servers feature the full complement of RAS features, including ECC memory; RAID controllers; redundant, hot-swap environmental elements (fans, power supplies, etc.); and hot-swap disk

³ *Q2CY99 Server Tracker*, International Data Corporation, October 1999.

drives. Features such as Dynamic Reconfiguration and Alternate Pathing (discussed in the section on planned downtime, below) even support online repair of systems.

With Solaris 8, Sun introduces the capability of having the operating system lock out components – such as certain memory and disk – from the server configuration if the system identifies them as bad upon startup or during operation. The system will not use components that it has locked out. Administrators also have the ability to lock out components on a running server.

Sun support agreements can offer warranty replacement of components prior to failure, if identified by predictive failure-analysis techniques. When the system sees a component it thinks may fail or is subject to imminent failure, the server will dial out and notify Sun support personnel. Sun servers also support remote hardware-diagnostic capabilities.

PLANNED DOWNTIME

Addressing growing needs for 24x7 operations, Sun offers many features to minimize planned downtime for system upgrades or expansion. The company has pioneered support for online hardware replacement and repair of logic modules. Dynamic Reconfiguration (DR) – first introduced with Solaris 2.5.1 and the Sun Enterprise 10000 – allows an administrator to prepare a server to be expanded or repaired online. While not new or unique to Solaris 8, DR and Alternate Pathing – combined with the extensive use of hot-plug components (including logic modules) – help significantly to minimize planned downtime.

SECURITY

To guard against system failures caused by security breaches, Solaris 8 provides a flexible set of facilities for strong user authentication. These facilities can be used out of the box, or developers can integrate them into applications as specific security features. For example, Pluggable Authentication Modules (PAMs) can be constructed to support authentication requirements (e.g., biometric scanning devices or smart cards) for applications or even for access to Solaris 8 itself.

Solaris 8 implements the Internet Engineering Task Force (IETF) specifications for the Internet Protocol Security Architecture (IPSec). IPSec allows administrators to establish secure encrypted networks and provides mechanisms for restricting access and authenticating users. The Solaris IPSec implementation includes an extensive framework for supporting the features of smart cards. Solaris 8 provides out-of-the-box support for smart-card authentication at login time. The smart-card support framework provides a set of tools that enable custom applications to access smart-card features such as the ability to install, remove, and manage Java applets on smart cards.

Other security features such as certificate-management systems are under development as part of Solaris 8 and are available from third parties.

PROBLEM RESOLUTION AND CHANGE MANAGEMENT

Many of the most difficult and insidious support problems occur as a result of changes to the server configuration: administrators add new hardware, install a new version of operating system or database software, or load a new set of patches that cause a problem. The bad effects impact administrators and users, and are especially significant to an ISV when the problem relates to their application.

Solaris 8 provides tools to help resolve these difficulties. For example, it introduces Live Upgrades. In a deployment scenario, the Live Upgrades capability dramatically reduces the amount of downtime normally associated with operating-system upgrades.

Live Upgrades allow installation of Solaris 8 software and other applications on a partition separate from that of the currently running environment. When installation is complete, a simple reboot puts the new system software in control. In a support environment, Live Upgrades then enables easy fallback to a previous version of the software with just a simple reboot. This fallback does not lose administrative information.

Solaris 8 also has patch-management software that offers new tools for determining the details of a user's operating environment. As patches are applied to the system, the patch-management software records the software revision and other information. Administrators can use this information to easily back out patches if necessary.

For the most difficult problems that result from system changes, Sun provides a hot-diagnostic capability that allows code to be dynamically added to the Solaris 8 kernel. This code gathers additional data to assist specialists in pinpointing interaction problems.

SYSTEM-MANAGEMENT FEATURES

For many organizations, particularly those that have made the Internet a critical part of their business infrastructure (e.g., organizations pursuing e-business strategies), high availability means more than minimizing downtime. What matters to such organizations is making sure users can have optimal interactions with the applications that run on their systems.

To maximize the level of application availability, organizations must guarantee predictable response times. This requirement often forms a central part of service-level agreements, but can prove difficult for organizations to satisfy while also making efficient use of limited computing and networking resources. Sun provides several features that organizations can use to allocate and control system resources. For enterprises with conspicuous dot-com presences, resource-control features can prove particularly important.

RESPONSE TIMES

The expansion of the Internet throughout business processes – from supplier to customer – poses new challenges to applications exposed to its effects. Unlike data center applications that traditionally experienced gradual increases in utilization, applications affected by the Internet are notorious for their wildly fluctuating workloads. When a system approaches capacity utilization, it must give critical applications the resources they need, while making other, less-critical applications share the remainder.

SOLARIS RESOURCE MANAGER

When an application fails to meet service expectations, a lack of processor cycles is the most likely culprit. Software-based and policy-driven functions for resource management, such as those found in Solaris Resource Manager, provide predictable response times for critical applications and users under varying workloads.⁴

Solaris Resource Manager allows a system administrator to allocate portions of the total processing resources on a server to specific applications or users. Once administrators set a policy, the system controls the scheduling of work to meet minimum allocations when the system is fully utilized. Administrators can control the allocations given to applications, users, or groups of users in various combinations.

For example, administrators might use Solaris Resource Manager to set priorities on a department-wide basis, giving marketing 20 shares and finance 15 shares. All department users would subsequently inherit capability from this specification. When a share allocation is changed, administrators make the change only once, and all department users then inherit the new allocation.

Solaris Resource Manager tools also work with more hardware-oriented forms of resource management, such as processor sets, which bind specific applications to specific processors in an SMP system, and the Dynamic Domains features available on the Sun Enterprise 10000.

SOLARIS BANDWIDTH MANAGER

A scarcity of other resources can also impact application performance significantly. This is particularly true of network or disk bandwidth. Sun's Solaris Bandwidth Manager, which comes bundled with Solaris Resource Manager, allows administrators to address bandwidth problems.

Solaris Bandwidth Manager provides capabilities similar to Solaris Resource Manager for network traffic. Administrators can prioritize network traffic, for example, so that applications requiring rapid responses get higher priority on the network. Solaris Bandwidth Manager can also guarantee a percentage of the available network bandwidth to different classes of network traffic. Administrators can base guarantees on several criteria, including the application, the source of the traffic, the destination IP address, the URL group, or a combination of these.

⁴ Sun does not bundle Solaris Resource Manager with Solaris 8.

ADMINISTRATIVE INTERFACE

UNIX has historically had relatively complex system-management procedures that required administrators to hand-edit a large and dispersed set of cryptic configuration files – a crude and error-prone process. This resulted from UNIX's ad hoc development by those largely unconcerned with the problems of large-scale production computing environments. As networks have expanded across organizations and the Internet has become a growing part of day-to-day activities, lowly administrative chores such as restarting a printer queue or handling a backup procedure often fall to less-experienced (and less expensive) system administrators.

Some UNIX vendors have responded by developing more user-friendly, GUI-based system management tools. Administrative tools have begun the transition from “remember and type” to “recognize and point.” Sun's Solaris Management Console addresses the management needs of its systems. The company makes its well-developed management interface accessible to ISV developers by supporting application integration with the Solaris Management Console and by providing an SDK.

SOLARIS MANAGEMENT CONSOLE

Solaris Management Console represents another key differentiator for Solaris 8-based ISV applications. Solaris Management Console software provides a consistent, easy-to-use interface for managing and administering a Solaris server, the clients of the server, and the applications running on the server. The console can be used for monitoring and tuning applications as well as being the administrative interface for user-written or third-party applications.

With Solaris Management Console, a Solaris server can be managed from almost anywhere – a browser, any Solaris workstation or server, or an NT server or standalone. Solaris Management Console can also plug into other consoles.

Solaris Management Console can consolidate the management functions of Solaris-branded optional products, as well as base components of Solaris that require an administrative interface. This gives the Solaris administrative interface a common look and feel across the entire range of Solaris products. For administrators, the interface provides a familiar Windows look and feel, with Windows-type GUIs, point and click, tab panels, and wizards.

SOFTWARE DEVELOPMENT KIT

Solaris Management Console also includes an SDK. The SDK allows developers to integrate administrative interfaces for their applications into Solaris Management Console. This ties them into a full set of runtime infrastructure services, including authorization, authentication, messaging, and logging. When developers integrate an application's administrative interface into Solaris Management Console, it gains the same look and feel as other Solaris

SUN CLUSTER 3.0

Sun plans to ship Sun Cluster 3.0 in an update to Solaris 8 during the second half of 2000. Sun Cluster 3.0 increases the number of nodes in a cluster to eight and introduces new ease-of-use features, including new capabilities for global device naming, a cluster file system, and global networking that simplify the process of creating and managing application failover among cluster nodes.

GLOBAL DEVICE NAMING

In Sun Cluster 3.0, disk devices used by more than one cluster node have a global, logical device name maintained by a global file system. The system maintains this name in addition to the device's name in the local device name space. Global device naming makes disks always accessible by the same global name from any cluster node. This scheme should prevent failures of any raw reads/writes to a disk (a technique extensively employed by database management systems) because clustering software can automatically discover another path to the disk.

CLUSTER FILE SYSTEM

The cluster file system in Sun Cluster 3.0 allows file access from any cluster node to any file within the cluster file system. The request does not need to specify the actual physical file location, and any cluster node can access the file identically.

The cluster file system is a proxy file system that supports all existing on-disk data formats. This file system provides continuous data access by creating mini-transactions that are atomically retried if there are failures. As with global device names, applications will not experience any data-access errors so long as a physical path exists to the disk where the file resides.

The combination of global device names and a cluster file system simplifies the process of specifying how applications fail over to other cluster nodes. When a node failover occurs, Sun Cluster 3.0 automates the tedious and exacting work of determining the new path names to required resources. Sun Cluster 3.0 allows applications to access cluster-wide resources the same way, regardless of where on the cluster they reside.

GLOBAL NETWORKING

Sun Cluster 3.0 makes client access to a cluster far more powerful through global networking. This feature gives the entire cluster a global IP address in addition to the IP address still retained by each cluster node, which allows users to connect to the cluster as a whole, rather than to a specific cluster node.

Global networking automatically load-balances incoming connections for better server and network utilization. It also allows for complete masking of single node outages from users.

DEVELOPING HIGHLY AVAILABLE, CLUSTER-AWARE APPLICATIONS

To help ISVs develop highly available, cluster-aware applications, Sun Cluster software introduces the concept of a logical host that has processing resources, disk storage, and a network IP address. A logical host is mapped to a physical server node in the cluster, and the physical node impersonates the node name and IP address of the logical host. A logical host can reside on two or more cluster nodes that have physical access to the logical host's disk storage.⁵ A Sun cluster can support any number of logical hosts, and a physical server or cluster node can impersonate more than one logical host.

For access to an application over the network, the application advertises the logical host name as the server name clients should use. When clients access the application, they connect to the server node impersonating that logical host name. If a server node fails over, a client connecting to a service will experience a brief disconnect at most.

Sun Cluster software continuously monitors the health of the cluster's physical servers. If a node fails, Sun Cluster software moves the logical host, including its disk storage and logical network IP address, to a surviving node.

Applications exploit the logical host concept to become highly available. Sun offers a number of Sun Cluster Data Service APIs that enable developers to ensure that applications they develop take full advantage of Sun Cluster capabilities.

APPLICATION REGISTRATION WITH CLUSTER SOFTWARE

To become cluster-aware, an application first registers with Sun Cluster software so that it will be informed of key events that occur within the cluster. The application can register up to three callback programs (or methods) that the cluster software will activate depending upon the event:

- Start – when an application needs to be restarted on a new node following a failover
- Stop – when an application needs to perform an orderly shut down because the node on which it is executing will be removed from the cluster by operator command
- Abort – when an application needs to execute fast clean-up code prior to a shutdown initiated by the operating system due to an imminent node failure

⁵ Sun Cluster supports the use of UFS, VxFS, and raw partitions on disk storage.

Sun Cluster Data Service APIs use command-line utilities and a C-callable library. The two methods provide identical services and their use depends on developer preference. The API:

- Registers applications with Sun Cluster software
- Stores information about the current cluster configuration
- Can run a child program while holding a file lock
- Can run a child program under a time-out
- Can determine whether an application is already running
- Can determine the logical hosts being impersonated by a physical host
- Can move a logical host from one physical host to another
- Can move all logical hosts that a physical host currently impersonates to other Hosts and request a cluster reconfiguration

Sun supplies a high-availability data service module for NFS software operating with Sun Cluster 2.2. Additional (unbundled) high-availability data service modules are available for commonly used application software and middleware, including:

- Oracle, Sybase, and Informix databases
- Oracle and Informix parallel databases
- Netscape servers including the LDAP server
- Tivoli
- Lotus Notes
- SAP R/3

Highly available, cluster-aware applications developed using the Sun Cluster Data Service APIs can be used on all future cluster releases starting with Sun Cluster 3.0.

CONCLUSIONS AND RECOMMENDATIONS

Solaris 8 provides a rich set of tools that allows in-house and third-party application developers to make full use of the capabilities of Sun Enterprise servers, the powerful availability and management features of the Solaris 8 Operating Environment, and the leading features of Sun Cluster 2.2 and Sun Cluster 3.0.

- Sun Enterprise servers represent a leading hardware solution that has consistently led the pack in single-system RAS features that are fully supported in Solaris 7. Sun's servers have unique features unmatched by the competition.⁶ Sun continues to boost this differentiation, for example, by extending its hot-plug support to even more devices.
- The Solaris 8 Operating Environment provides impressive support for Internet and mission-critical applications that require a continuous presence, addressing unplanned and planned downtime, as well as security concerns.
- Solaris Resource Manager and Solaris Bandwidth Manager enhance application availability by allowing administrators to manage system and network resources to guarantee stringent response times and compliance with service-level agreements.
- Sun's Solaris Management Console and SDK bring a common Solaris look and feel to administrative interfaces for applications integrated with it.
- For business environments that require high availability and scalable performance, Sun provides competitive clustering capability with Sun Cluster 2.2, which it will expand and enhance with Sun Cluster 3.0. Sun Cluster 3.0 increases the number of nodes in a cluster to eight and introduces new ease-of-use features.

Sun's leading mindshare, combined with its strong capabilities in hardware and software, present ISVs with an excellent opportunity to produce industry-leading code that meets the application needs of leading-edge users.

⁶ *Competitive Analysis of Reliability, Availability, Serviceability and Cluster Features and Functions*, D.H. Brown Associates, Inc., September 1998.