Sun Server Hardware and Software and the New Solaris 8 Operating Environment Optimized for Mission-Critical and E-Business Environments

Prepared for Sun by:

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

January 2000

TABLE OF CONTENTS

EXECUTIVE SUMMARY: E-BUSINESS REQUIREMENTS RAISE THE STANDARD FOR SYSTEMS MANAGEMENT 1

Table 1: How Solaris 8 Increases Application Availability for Users 2

MINIMIZING UNPLANNED DOWNTIME 4

HARDWARE 4
ERROR CAPTURE AND CORRECTION 4
HOT SWAP AND DYNAMIC RECONFIGURATION 5
SUN CLUSTER 2.2 5
SUN CLUSTER 3.0 – SUN CLUSTER 2.2 AND MUCH MORE 6
GLOBAL DEVICE NAMING 6
CLUSTER FILE SYSTEM 6
GLOBAL NETWORKING 7

DISASTER TOLERANCE 8

PROTECTION AGAINST HACKERS 8

MANAGING SYSTEM CHANGES 8

REDUCE PLANNED DOWNTIME 9

MAINTAINING SERVICE LEVELS 10

WHAT'S MY SHARE? 11
MULTIPLE ALLOCATION FILES 11
SOLARIS RESOURCE MANAGER IN A CLUSTER 11
BANDWIDTH MANAGEMENT 12
ROLE-BASED ACCESS CONTROL 12

SOLARIS 8 MANAGEMENT CONSOLE 13

CONCLUSIONS AND RECOMMENDATIONS 14

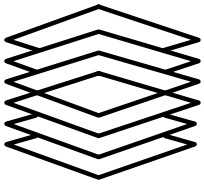

Sun Server Hardware and Software and the New Solaris 8 Operating Environment Optimized for Mission-Critical and E-Business Environments

EXECUTIVE SUMMARY: E-BUSINESS REQUIREMENTS RAISE THE STANDARD FOR SYSTEMS MANAGEMENT

Sun server hardware and software and Sun's new Solaris 8 Operating Environment[1] offer a host of new features, procedures, and technologies to support business-critical applications. These capabilities increase the availability of mission-critical systems by reducing both planned and unplanned downtime. They also support additional requirements to ensure service-level compliance in support of Internet commerce and other e-business applications.

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. This requirement encompasses three complementary goals:

- Minimizing unplanned downtime (where systems go down due to errors and failures)
- Minimizing planned downtime (where systems are taken offline for maintenance, upgrades, etc.)
- Ensuring predictable, consistent response times

Continuous uptime requires an operating environment that supplies features to minimize downtime. This paper provides IT managers with an overview of features that promote continuous operation and mitigate the impact of both planned and unplanned downtime. Sun Enterprise servers, the new Solaris 8 (Table 1), and related products and services all offer features to mitigate the effects of planned and unplanned downtime.

To maximize the level of application availability, organizations must also ensure that users can enjoy 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 computing and networking resources. This paper goes beyond uptime issues, examining features that support the ability to allocate or dedicate resources to users or applications, an important first step to managing system response time. Sun Enterprise servers and Solaris 8 provide several features that organizations can use to allocate and control system

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

resources. For enterprises with conspicuous dot-com presences, resource-control features can prove particularly important.

The paper also summarizes important new administrative functions in Solaris 8 that enhance its ease of use.

TABLE 1: How Solaris 8 Increases Application Availability for Users

Reduces Unplanned Downtime
• Improved error capture and faulty component lock-out • Additional hot-plug capability and application-aware hot-plug capability • Sun Cluster 3.0 failover software • IPSec security architecture • Patch-management software and hot-diagnostic capability
Reduces Planned Downtime
• Live Upgrades
Ensures Consistent Access and Service Level Support
• Support for Solaris Resource Manager and Solaris Bandwidth Manager
Provides Administrative Functions
• Role-based access control • Solaris Management Console

D.H. Brown Associates, Inc. (DHBA) believes Solaris 8 extends the capabilities of Sun Enterprise servers to enhance support of critical business activities and high availability of user applications. This environment addresses both unplanned and planned downtime issues. It also meets the need for a continuous computing-system presence with stringent response times and compliance with service-level agreements.

On the hardware side, Sun Enterprise servers have consistently led the pack in single-system reliability, availability, and serviceability features that are fully supported in Solaris. Sun Enterprise servers already have unique features unmatched by the competition. With Solaris 8, Sun boosts its lead, for example extending hot-plug support to even more devices.

In business environments that require high availability, Sun provides a competitive clustering capability with Sun Cluster 2.2. Sun will expand and enhance clustering capability with Sun Cluster 3.0, which will be introduced after the release of Solaris 8. As well as increasing the number of cluster nodes, Sun Cluster 3.0 introduces new ease-of-use features.

For environments that require guaranteed resource allocation to ensure response times and to meet prescribed service-level requirements for management and maintenance, Sun provides Solaris Resource Manager. This product guarantees

Copyright 2000 © D.H. Brown Associates, Inc.

that critical applications and users get the resource allocations they need. Solaris Bandwidth Manager provides similar capabilities for network resources. Both products stand at the forefront and provide granular system-resource control to ensure application availability.

New e-business applications that require predictable response times, combined with the need for high availability, have dramatically increased the complexity of system management and monitoring. Solaris 8 has responded with several new initiatives to simplify system management and make the best use of corporate administrative expertise.

DHBA further believes that when Sun completes its rollout of Solaris 8, it will likely have a product offering that will meet or exceed the competition for ease in delivering Internet and e-business applications. Sun's offerings will need to be on any user's short list for mission-critical purchase consideration.

Copyright 2000 © D.H. Brown Associates, Inc.

MINIMIZING UNPLANNED DOWNTIME

HARDWARE

For some time, Sun has led in developing innovative techniques to improve the availability of its servers, as illustrated by the following examples[2]:

- Sun uses ECC in all memories and caches,[3] enabling correction of the vast majority of data-corruption errors. Sun's Alternate Pathing (AP) capability allows the operating system to make use of redundant network and storage hardware connections to find new paths to an I/O device if there is a failure. This rerouting is totally transparent to users and applications, completely masking the fact that there has been a failure.
- Sun supplies hardware RAID controllers and highly available storage subsystems that allow transparent recovery from disk device failures.
- Sun Enterprise platforms support redundant environmental elements such as power supplies and fans for cooling. If one of these devices fail, there is sufficient capacity in the remaining components to keep the server operating within proper environmental limits.

ERROR CAPTURE AND CORRECTION

Sun goes beyond hardware-based measures for reducing unplanned downtime, starting with improved error capture in Solaris 8 compared to previous Solaris releases. Examples include:

- Solaris 8 has better methods for collecting and managing core dumps.
- It has basic error-capture capabilities available now, with more to come.
- There will be constant improvements in predictive failure analysis, allowing the system to identify components on the verge of failure.
- Sun Enterprise servers can dial out and notify Sun support personnel of failures or potential failures ("call home"), assuming the appropriate support program is in place. This allows remote diagnosis of problems because the servers have a remote console and diagnostic capability.
- A new feature allows the operating system to lock out bad components from the usable configuration, either dynamically on a running system or when the server is rebooted. The server will continue to operate until the failed component is replaced.

[2] *Competitive Analysis of Reliability, Availability, Serviceability and Cluster Features and Functions*, D.H. Brown Associates Inc., September 1998.
[3] ECC is not used in the L2 cache of the Sun Enterprise 10000 server.

 Copyright 2000 © D.H. Brown Associates, Inc.

HOT SWAP AND DYNAMIC RECONFIGURATION

To ensure continued resilience of a server, systems must support repair or replacement of failed or marginal components. Online replacement (hot swap) of environmental elements that are largely beyond the control of the operating system – such as power supplies and fans – have become commonplace on even inexpensive servers. Sun Enterprise servers fully support this capability.

Sun has pioneered the support of online hardware replacement and repair of logic modules. A highlight of this effort is Dynamic Reconfiguration (DR), which was first introduced with Solaris 2.5.1 and the Sun Enterprise 10000 and allows an administrator to prepare a server for repair and subsequent return to production.

Systems with DR allow an administrator to command the operating system to bring a device in need of repair into a quiescent state and move activity to a redundant component. The extensive use of hot-plug components in Sun Enterprise servers enables I/O boards and even processor/memory boards to be replaced online. When repairs are complete, the administrator can reintroduce the component into the configuration. The latter capability allows server capacity to be expanded online without requiring the server to be rebooted.

Other features in Solaris 8 further enhance online repair capabilities. Support for SCSI and generic PCI hot-plugging has been added to an already lengthy list of supported hot-plug devices.

SUN CLUSTER 2.2

A single Sun Enterprise server running Solaris 8 can often continue to operate if errors or failures occur. Even when a failure does cause an outage, recovery occurs rapidly, perhaps with a simple reboot. In an environment that cannot afford any outages, however, the system must eliminate the limited downtime required for such recover-and-reboot procedures.

Environments that require continuous service may consider using Sun Cluster 2.2, an add-on option to Solaris 8.[4] Sun's current cluster offering is a multi-computer solution that can adjust to the loss of a server (cluster node) and keep the cluster operational. This capability allows Sun Cluster 2.2 to prevent software errors from causing outages and to minimize outages caused by hardware failures or repairs.

Sun Cluster 2.2 supports up to four nodes. The nodes in a Sun Cluster environment connect to each other using private, high-speed networking connections. To ensure uninterrupted access to disk storage, disks are connected to at least two nodes in the cluster. Sun Cluster 2.2 also supports Fibre Channel (FC-AL)-connected Sun StorEdge disk arrays, which provide greater connectivity and can boost I/O throughput.

[4] Sun Cluster 2.2 will also work with Solaris 2.6 and Solaris 7. Sun's next version of this product, Sun Cluster 3.0 (discussed in a subsequent section), is designed to work only with Solaris 8.

Copyright 2000 © D.H. Brown Associates, Inc.

Sun Cluster 2.2 increases application availability by providing fast recovery from events that can cause an application to become unavailable. Locally, an application can be restarted or a network adapter can be failed over. If a node fails, a node's workload or part of its workload can be failed over and restarted on another cluster node. The cluster application programming interface (API) included in the cluster software enables developers to make most applications highly available.

SUN CLUSTER 3.0 – SUN CLUSTER 2.2 AND MUCH MORE

Sun plans to ship Sun Cluster 3.0, the successor to Sun Cluster 2.2, with an update to Solaris 8 in the second half of 2000. The new version of Sun Cluster will increase the number of cluster nodes the system can handle to eight. It will also introduce important new features, such as:

- Cluster-wide device access through global device naming
- A cluster file system
- Global networking

Solaris 8 provides kernel support for these new clustering features, which simplify the process of creating and managing highly available applications within a cluster.

GLOBAL DEVICE NAMING

For any cluster to support continuous access to critical data, the disk devices containing that data must connect to more than one node. Highly available disk devices that support this requirement include dual-ported disk arrays and Fibre Channel-connected storage subsystems. Sun Cluster 3.0 gives these highly available devices a global, logical device name that is maintained on the global file system.[5]

Global device naming makes disks accessible from any cluster node using the same name. This device-naming scheme makes it possible to easily migrate an application to another cluster node, since applications and volume-management software typically use logical device names. Another important benefit of this approach is the fact that raw reads and writes to disk – a technique extensively employed by database-management systems – will always succeed because the clustering software can automatically discover the secondary path to the disk.

CLUSTER FILE SYSTEM

A cluster file system allows any node to access a file within the cluster file system and does not require specifying the actual physical file location. The file can be accessed identically from any cluster node because the cluster file system – which serves as a proxy file system between the kernel and the underlying file system –

[5] The system maintains the logical device name in addition to the device name in the local device name space.

 Copyright 2000 © D.H. Brown Associates, Inc.

is transparent to existing applications. All existing on-disk data formats can still be used.

The cluster file system provides continuous access to data by creating mini-transactions that are atomically retried if there are failures. As long as the system has a physical path to the disk where the file resides, the application will not experience errors in data access.

The combination of global device names and a cluster file system in Sun Cluster 3.0 vastly simplifies the process of specifying how applications fail over to other cluster nodes in the event of a node failure. With Sun Cluster 2.2, the failover script for a mission-critical application needed to know which node an application would run on in order to properly define paths to all the resources it required. If an application ran on a different node than specified, the paths to its required resources had to change. The administrator had to supply alternative device-routing information in the script to enable application failover.

With Sun Cluster 3.0, by contrast, an application accesses cluster-wide resources the same way regardless of where that application resides or where on the cluster the resources reside. Sun Cluster 3.0 automates the tedious and exacting work of determining the new path names to required resources when failover occurs.

GLOBAL NETWORKING

The global networking in Sun Cluster 3.0 allows each cluster node to retain its own IP address. In addition, it is now possible to connect to "the cluster" in the abstract, rather than a specific cluster node. A specific service that runs on multiple cluster nodes can now be given a unique, global IP address. New TCP connections for this address get forwarded transparently to one of the cluster nodes that runs the service.

Global networking automatically load-balances incoming connections to provide better network utilization. Furthermore, with the use of global networking, single node outages can be completely masked from users. The user's incoming service request will be forwarded to one of the remaining cluster nodes to provide the service.

Copyright 2000 © D.H. Brown Associates, Inc.

DISASTER TOLERANCE

Users who are anxious about continuous uptime must plan for situations that go beyond the failures of software or hardware. They also have to deal with certain extraordinary events – including fires, storms, or even earthquakes – that can disrupt computers and networks and create lengthy outages.

Sun's Campus Clusters product allows the nodes within a Sun cluster to be separated by 10 km. This limits the risk of downtime due to loss of a computer facility at one site. In addition, Sun's Fibre Channel Sun StorEdge disk arrays allow mirroring of data at remote locations up to a distance of 10 km. Having a copy of current data will expedite the process of application recovery.

PROTECTION AGAINST HACKERS

In conventional as well as e-business environments, computer security breaches may result in loss of proprietary information to a competitor, theft of service, or theft of revenue. Security breaches also can impact a system's availability by disabling systems and networks. These can result in users experiencing outages of a system's services.

Solaris 8 introduces several new security features to address this problem. For example, it provides an implementation of the Internet Engineering Task Force (IETF) specifications for the Internet Protocol Security Architecture (IPSec). IPSec enables the establishment of secure encrypted networks and provides mechanisms for restricting access and authenticating users. The Solaris IPSec implementation offers additional mechanisms that enable security policies to be customized, including a framework for incorporating support for smart cards.

MANAGING SYSTEM CHANGES

Many of the most difficult, insidious problems faced by administrators occur when system components – devices, operating systems, middleware, and applications – suddenly begin behaving unpredictably. The system gets new hardware, a new version of an operating system or database, or a new patch, and suddenly the system does not work correctly. Such situations will frequently cause unplanned downtime.

Solaris 8 provides tools to help resolve such difficulties, including the following:

- A Live Upgrades feature that enables the system to revert to a previous Solaris release with just a simple reboot, a process that preserves administrative information.

Copyright 2000 © D.H. Brown Associates, Inc.

- Patch-management software that enables easy patch installation. It also records all patches that have been applied to the system, giving administrators access to software revision control information. Administrators can use this revision control information to easily back out patches if necessary.
- For the most recalcitrant of system problems, Solaris 8 has a hot-diagnostic capability that allows code to be dynamically added to its kernel. This code gathers additional data to assist a specialist in pinpointing a problem.

REDUCE PLANNED DOWNTIME

To reduce planned downtime, associated with software upgrades, Solaris 8 introduces the concept of "Live Upgrades." While it is still not possible to completely eliminate outages when upgrading to a new Solaris version, Live Upgrades dramatically reduces the amount of planned downtime normally associated with an upgrade.

With Live Upgrades, Solaris 8 software is installed on a partition separate from that of the currently running environment. When the installation is complete, a simple reboot allows Solaris 8 to take control. As a result, the only downtime associated with an operating system upgrade is now the amount of time it takes the operating system to reboot.

Live Upgrades represent a major difference between Solaris 8 and previous Solaris versions in the management of planned downtime. The Live Upgrades capability should prove a particularly strong differentiator for Solaris 8 in the e-commerce world, where 24x7 operation and minimal planned downtime are the norm. As noted earlier, DR and AP, while not new or unique to Solaris 8, also contribute to minimizing planned downtime for system expansion and certain repairs.

Future developments in Live Upgrades capabilities will likely include support for rolling upgrades of the operating system and of the cluster software. Eventually, Sun could even support rolling upgrades of the applications within a cluster (of course, applications must also be recast in many cases). For example, if an administrator could dynamically change the Oracle database (while it continues to run), that would increase the availability of a database system. The Oracle database could then take advantage of additional or changed computing resources on the fly. The latter capability may be available in the third update of Solaris 8.

Copyright 2000 © D.H. Brown Associates, Inc.

MAINTAINING SERVICE LEVELS

Many organizations must move beyond the basic availability functions discussed thus far to ensure that the system maintains prescribed service levels through resource management. As noted in the introduction to this paper, user application availability depends not only on minimizing unplanned and planned downtime, but also on predictable system-response times (which are part of service-level agreements). Typically, organizations configure their computer hardware so it can just handle peak loads, which means it will have capacity to spare for handling typical loads.

When no system component or network is overloaded, there is little need for resource allocation. Users and applications get the resources they require, and there are resources to spare. If there are several processes waiting to be run, the operating system uses a priority-based scheme to determine which process gets the next available processor resources.

When systems approach capacity utilization, however, the situation changes dramatically. To ensure that the system can meet the resource needs of the most critical applications, the system must guarantee resource allocation. Critical applications must get the resources they need, while other, less-critical applications (or users) share the remaining resources.

Software-based and policy-driven resource-management functions are effective tools to provide predictable response times and maintain and manage service levels under varying conditions. Sun's resource-management tool is Solaris Resource Manager, which is not bundled with Solaris 8. Solaris Resource Manager software works in conjunction with more hardware-oriented forms of resource management such as processor sets (which bind specific applications to specific processors in an SMP system). Note that the Sun Enterprise 10000 has Dynamic Domains, which allow partitioning of the server into separate instances of the Solaris Operating Environment. Thus, a critical application can be run in its own domain without having to share resources with other applications. Organizations can use Solaris Resource Manager within such domains.

Solaris Resource Manager enables system administrators to allocate portions of resources to specific applications or users. Once an administrator has input a policy, the system controls the scheduling of work to meet minimum allocations when the system is fully utilized. Administrators can make resource allocations to applications, users, or groups of users in various combinations. Solaris Resource Manager can also ensure the maximum amount of resources – such as virtual memory, number of processes, number of logins, or connect time – does not get exceeded.

 Copyright 2000 © D.H. Brown Associates, Inc.

WHAT'S MY SHARE?

Solaris Resource Manager uses an allocation scheme based on the concept of "shares" to allocate processor resources. An application or user's percentage of processor cycles varies in proportion to the number of shares it has compared to the total number of shares in the server's current workload. An application's importance and entitlement to resources will vary depending upon the other users and applications running on the system.

When new applications start up or new users log into the system, the system will automatically adjust resource usage based on the new workload. In contrast, resource-allocation schemes based on percentages may require the administrator to redefine all allocations in order to change a single allocation (resource allocations cannot exceed 100%).

Using Solaris Resource Manager, an administrator can dynamically change how the system schedules and allocates resources. Solaris Resource Manager provides several ways to adjust both resources and scheduling priorities.

MULTIPLE ALLOCATION FILES

An administrator can create different allocation files, for example, to provide different share assignments to users and applications. By using scripts executed at predetermined times, an administrator can shift from one resource-allocation plan to another plan. In addition, an administrator can use Solaris Resource Manager to change several of the internal parameters that drive processor scheduling. These parameters include scheduling frequency and how past usage is used to recalculate the scheduling priority of a particular process. Such changes will automatically be reflected in job priorities and new usage profiles without the need for a system reboot. Put simply, Solaris Resource Manager enables administrators to control processor resource scheduling at a very granular level if they so desire.

SOLARIS RESOURCE MANAGER IN A CLUSTER

Use of Solaris Resource Manager in the nodes of a Sun cluster provides a powerful tool to ensure that critical applications get the services they require when they have failed over. Failure of a cluster node reduces the amount of processor resources in the cluster as a whole. Thus, it is especially important to guarantee that the most critical applications receive the proper share of the processor resources that remain.

As an example, consider a two-node cluster where Node1 is used to run applications A, A1, A2, and A3, with A being the most critical application. Node2 is used to run applications B, B1, and B2, with B being the most critical. When both nodes are operational, the cluster has ample resources for all applications to run well.

Copyright 2000 © D.H. Brown Associates, Inc.

If either node fails, however, the remaining node will be running all six applications. When that occurs, applications A and B should get most of the processing time, but the system must dedicate a small amount of time to applications A1, A2, A3, B1, and B2 – it cannot deprive them completely of resources.

Both Node1 and Node2 need to have the following allocation rule in place: A has 20 shares; B has 20 shares; A1, A2, A3, B1, and B2 each have two shares. When all the applications are executing on a single node, applications A and B will each receive 40% of the available processor cycles 20/(20+20+2+2+2+2+2), and applications A1, A2, A3, B1, and B2 each will receive 4% of the processor cycles.

BANDWIDTH MANAGEMENT

Solaris Bandwidth Manager – which is bundled with Solaris Resource Manager – provides capabilities similar to those of Solaris Resource Manager for managing network traffic. Solaris Bandwidth Manager allows prioritization of network traffic so that applications that require rapid responses get a higher priority.

In addition, Solaris Bandwidth Manager can guarantee a percentage of the available network bandwidth to different classes of network traffic. These percentages may depend on the application, source and destination IP addresses, URL group, or an administrator-chosen combination of these factors.

ROLE-BASED ACCESS CONTROL

Recent trends "conspire" to dramatically increase the complexity of managing corporate computing environments. Administrators must support networks of desktop devices, UNIX servers, and Internet access. Further, IT must meet many goals: fully utilizing computing resources, ensuring continuous application uptime, and maintaining application service-level time requirements as well as more complicated service-level agreements (particularly in e-business environments).

To meet their obligations, system administrators require a vast array of tools to monitor and control the computing environment. Unfortunately, with complexity comes an increased chance of inadvertent actions by users or administrators that cause serious unintended consequences. With Solaris 8, Sun has taken steps to minimize and mask complexity whenever possible to reduce the occurrence of such problems.

Copyright 2000 © D.H. Brown Associates, Inc.

Among the new system-management features of Solaris 8 is role-based access control. Role-based access control allows a system administrator to delegate limited sets of administrative functions to other users. This allows organizations to have less-experienced administrators perform useful work without risking serious consequences if mistakes are made; and even end users can be allowed to "kill" their own print jobs. As a result, system availability may increase as certain administrators can solve simpler problems quickly without waiting for more experienced administrator intervention.

SOLARIS 8 MANAGEMENT CONSOLE

The Solaris 8 Management Console provides a consistent user interface for all Solaris administrative, monitoring, and tuning applications. A consistent interface makes it easier for an administrator to learn and become proficient at managing new applications. In addition, Solaris Management Console software provides an extensible system-management framework for ISVs or in-house developers. The Solaris Management Console software development kit (SDK) provides a set of routines and services that allow easy development of new administrative and management applications with a Solaris "look and feel." Use of the management-console interface allows a Solaris server to be administered from almost anywhere – a browser, any Solaris workstation or server, or a Windows NT server or stand-alone.

Copyright 2000 © D.H. Brown Associates, Inc.

CONCLUSIONS AND RECOMMENDATIONS

Solaris 8 extends the capabilities of Sun Enterprise servers to enhance support of critical business activities and high availability of user applications. This environment addresses both unplanned and planned downtime issues. It also meets the need for a continuous computing-system presence with stringent "not to exceed" response times and compliance with service-level agreements.

On the hardware side, Sun Enterprise servers have consistently led the pack in single-system reliability, serviceability, and availability features that are fully supported in Solaris.[6] Sun Enterprise servers already have unique features unmatched by the competition. With Solaris 8, Sun boosts its lead, for example, extending hot-plug support to even more devices.

In business environments that require high availability, Sun provides a competitive clustering capability with Sun Cluster 2.2. Sun will expand and enhance clustering capability with Sun Cluster 3.0, which will be introduced after the release of Solaris 8. As well as increasing the number of cluster nodes, Sun Cluster 3.0 introduces new ease-of-use features.

For environments that require guaranteed resource allocation to ensure "not to exceed" response times and to meet prescribed service-level requirements for management and maintenance, Sun provides Solaris Resource Manager, which guarantees that critical applications and users get the resource allocations they need. Solaris Bandwidth Manager provides similar capabilities for network resources. Both products stand at the forefront and provide granular system-resource control to ensure application availability.

New e-business applications that require predictable response times, combined with the need for high availability, have dramatically increased the complexity of system management and monitoring. Solaris 8 has responded with several new initiatives to simplify system management and make the best use of corporate administrative expertise.

When Sun completes its rollout of Solaris 8, it will likely have a product offering that will meet or exceed the competition for ease in delivering Internet and e-business applications. Sun's offerings will need to be on any user's short list for mission-critical purchase consideration.

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

Copyright 2000 © D.H. Brown Associates, Inc.