IDC
Analyze the Future

WHITE PAPER

Addressing Virtualization and High-Availability Needs with Sun Solaris Cluster

Sponsored by: Sun Microsystems

Jean S. Bozman
October 2009

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

EXECUTIVE SUMMARY

Business applications running on IT infrastructure require high levels of availability in order to minimize the amount of downtime experienced during any planned and unplanned outages. The business need for this high availability is apparent: Without access to the system, end users and end customers cannot continue their business activities — and this impacts revenue generation, profitability, and the way that users feel about their mission-critical systems and data stores.

Now, with the accelerating virtualization of all IT infrastructure in the datacenter, more enterprise applications must be protected, whether they are running on physical or virtual servers.

Virtualization technologies enable consolidation of multiple applications onto one physical server. If that physical server goes offline, this consolidation of workloads opens up much higher risks of failures for any one of the application services that are running in the same virtualized environment. Redundancy is needed to mitigate the potential risks, along with software protection that provides high availability for the applications. With virtualized resources, the security isolation, resource management, and fault isolation of a virtual instance are the given requirements. High-availability software should also be implemented for timely data and services recovery from failure, especially for enterprise applications and databases.

Today, the world of enterprise computing is changing, with software delivered as a service (SaaS); the advent of cloud computing to deliver information services via "private" and "public" cloud infrastructure; and the convergence of once independent elements of IT infrastructure — servers, storage, software, and networking devices.

All of this IT activity signals the arrival of a new inflection point in the industry, when business solutions will be delivered to end users — and the supporting IT technology will be refreshed, over time, and maintained in a behind-the-scenes way so that end customers will have little, or no, visibility into the workings of IT activities and internal IT processes. Instead, end users will look to quality of service (QoS) and to service-level agreements (SLAs) as the metrics that will measure how well IT is providing enterprise-class data services to the business organization.

SITUATION OVERVIEW: HIGH-AVAILABILITY SOFTWARE AND VIRTUALIZATION IN THE ENTERPRISE

The Business Value of High-Availability Software

Enterprise applications must continue running, even in the event of an unforeseen/unplanned outage. Outages can be caused by a number of events — hardware or software failures, network outages, power outages, or natural disasters. No matter the cause of the outage, enterprise applications need to be restarted on alternate resources, addressing unplanned downtime. During periods of maintenance/repair of systems, planned downtime must also take place.

In today's distributed IT infrastructure, the cause of downtime can occur anywhere in the enterprise — but, importantly, the effects of downtime can be felt everywhere in the enterprise. However, there are many approaches to ensuring high availability, ranging from deployment of highly reliable hardware to the presence of high-availability software that ensures that applications can run on alternate resources and, if needed, on fault-tolerant systems that have fully redundant hardware coupled with lockstep software. Most systems fit somewhere along this high-availability spectrum in terms of the way that they are protected from downtime, based on the types of workloads running on the machines.

The business value of high availability is clearly seen with few, if any, interruptions to business applications, productivity of office workers and IT staff, and smoothly running business processes. If systems are unreliable, or if availability is hampered by downtime, then business processes are disrupted, and IT is forced into a crisis mode to address system failures, data recovery, and the resumption of normal IT processing.

IDC defines a "spectrum" of availability, including simple backup, replication of data, protecting applications through high-availability failover and restart of applications on alternate resources, workload balancing within computing tiers (e.g., the Web tier, the application tier, the database tier), and fault-tolerant systems.

The broadest coverage is provided by Availability Levels 2 and 3 (AL2 and AL3), as defined by the IDC Availability Spectrum, which are the levels supported by failover clusters and workload balancing. These solutions require the ability to do the following:

- Support existing applications and databases, providing them with alternate resources for restart of applications and databases
- Enable a fast failover from one server resource to another, in the event of an outage
- Provide a single, coherent view of the systems on which workloads are being managed and restarted (including both physical servers and virtual servers)

IDC notes that the two other levels — AL4 (for fault-tolerant servers with no interruptions in processing) and AL1 (for servers with no ability to restart applications) — fill in the rest of the IDC Availability Spectrum.

The Role of Virtualization

In the broad IT marketplace, much of the discussion about the virtualization revolution is centered on the x86 server tiers. Historically, these servers did not have the built-in reliability features or high-availability software capabilities of mainframe or scalable Unix servers, although improved reliability features have been added to some x86 models in recent years. IDC notes that both mainframes and Unix servers are already highly virtualized and offer a combination of reliability, availability, and serviceability (RAS) features in the hardware components and very granular control of all software components running on the systems. Given virtualization's ability to significantly reduce operational costs by consolidating workloads on fewer server "footprints," interest in virtualization as an agent of change has grown in recent years.

The business benefits of virtualization apply to all types of hardware platforms because the ability to virtualize hardware resources, and to move workloads between these resources, supports workload consolidation and improved workload management — a benefit unrealized with unvirtualized servers. Today, customers are leveraging virtualization to accomplish a number of goals in the datacenter:

- To improve resource utilization on x86 servers and reduce the power/cooling requirements by supporting the hosting of more workloads per server
- To increase workload consolidation onto a smaller number of server "footprints" within the datacenter, reducing operational costs
- To isolate applications so that they do not interfere with one another, improving uptime for the business
- To improve IT flexibility by increasing the mobility of workloads across the infrastructure by means of migrating virtual machines (VMs)
- To improve high availability by failover of applications from one virtual server to another, often tying in with a disaster recovery plan

The Rapid Growth of VMs

The rapid growth of VMs in recent years has been dramatic. Technical capabilities on the x86 server platform have changed, given the presence of multicore processors and virtualization along with improved energy efficiency. Only a few years ago, support for two to four VMs per physical server was typical; however, IDC's 2008 Server Virtualization Study shows that today, an average of eight or more VMs are running per virtualized server — and higher numbers of VMs are possible, depending on workloads and the specifications of the physical servers. This increase in VM density is raising awareness about the business and IT requirements for high availability because any downtime for these multi-VM servers will cause disruption to the business.

The growth in both physical servers and virtual servers within the worldwide installed base of servers demonstrates the increasing need to be able to manage either kind of server. Both must support a number of application restarting scenarios, including migration of workloads between virtual servers, as IT resources and application "services" are increasingly mapped to business processes rather than to specific physical machines.

High Availability in a Virtualized World

Virtualization technologies alone do not address thoroughly the high availability that is required for business continuity by mission-critical services. The combination of high-availability technologies and virtualized server technology extends and optimizes compute resource utilization across the IT infrastructure while maximizing the availability of data services being delivered to end users.

Using virtual servers in place of physical servers for failovers reduces the amount of capital expenditure for additional physical servers, any or all of which might otherwise be underutilized (prior to virtualization). High-throughput, high-performance enterprise-class servers are usually shipped with a substantial amount of storage capacity and network bandwidth, making recovery from any downtime even more challenging than it is for volume servers.

Importantly, the increasing density of VMs per physical server means that any outage at the physical level would cause all of the VMs hosted on that server to go offline. This idea of having such a high density of workloads running on a single server is equivalent to having "all the eggs in one basket." This realization on the part of businesses that many workloads would be disrupted in the event of any sort of outage affecting the host physical servers now raises the visibility and importance of high-availability software protection in a virtualized world.

Manually moving virtual machines across physical servers will allow only a limited level of high availability for purposes of planned downtime. IDC notes that availability for planned downtime can be increased through the use of automation and policy-based control of the environment, as servers are repaired and software is upgraded. However, typically, such live migrations do not address the issue of downtime in an unplanned downtime scenario, when applications go offline.

In contrast, unplanned outages in a virtualized architecture require a different approach that includes the following elements:

- Reliable monitoring of the hardware and software components, such as network, storage, and server
- Policy-based and application-specific recovery actions
- Resilient configuration of the overall infrastructure, with no single point of failure to ensure a complete and high-availability business continuity solution

A recent IDC study highlighted these trends, showing that datacenters are deploying high-availability technologies to ensure business continuity for a range of virtualized server platforms. Respondents said that the increasing number of VMs and the need to keep enterprise workloads available to end users were drivers for high-availability software deployment (see Figure 1). Further, the data indicated that the need for high availability would rise over the near term, as VM density grows in the x86 space and as enterprise applications increasingly run on virtualized x86 server platforms.

FIGURE 1

Why Virtualization Software and High-Availability Software Are Being Used to Protect Applications

Q. If you combine your virtualized servers with high-availability software, the reasons you do so include?

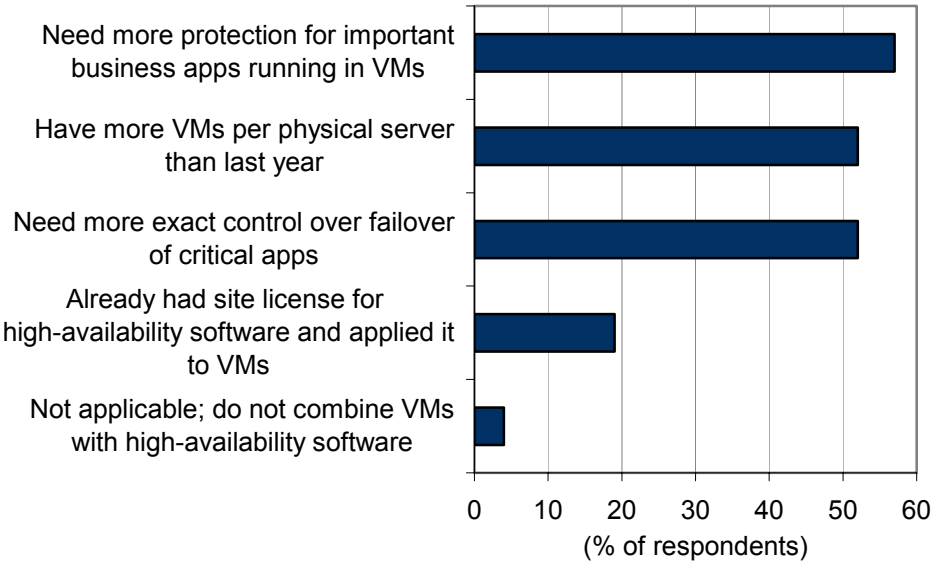

n = 369

Source: IDC, 2009

Overview of the Worldwide Market for Availability and Clustering Software

IDC supply-side data shows that the current market for availability and clustering software (ACS) grew from $1.75 billion in 2007 to $1.94 billion in 2008. The ACS market's continued growth, from 2008 through 2013 (see Figure 2), demonstrates business customers' need to minimize downtime by restarting applications on alternate resources (physical servers or virtual servers). Whether deployed on new servers or on existing servers, high-availability software improves the business value of investments in server systems and reduces operational costs due to maintenance, downtime, and IT staff costs associated with addressing that downtime.

FIGURE 2

Worldwide Availability and Clustering Software Revenue, 2008–2013

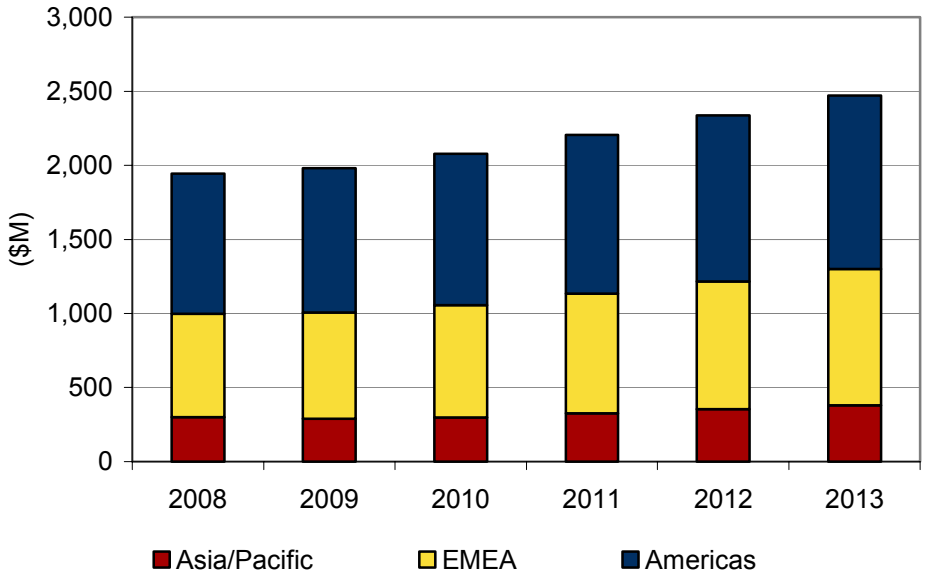

Source: IDC, September 2009

IDC notes that the ACS market has been growing faster than the worldwide server market, demonstrating the need for high availability in IT infrastructure containing scale-out configurations of small and midrange enterprise systems and also scale-up midrange or high-end enterprise servers within the datacenter.

Increasingly, ACS products are being combined with virtualization software to protect workloads, whether they are running on physical servers or virtual machines, for the purposes of moving applications to alternate resources, if needed. Unplanned downtime, which is caused by power outages, network outages, or man-made or natural disasters — or by the failure of hardware or software components —- is the primary reason for moving workloads to alternate computing resources. However, ACS software can also be used for purposes of planned downtime. In these planned downtime scenarios, ACS software is used to support the planned maintenance for purposes of hardware repairs, software updates, or replacements of physical servers.

In disaster recovery scenarios, lengthy mean time to recovery (RTO) and recovery to the point when the failure occurred (RPO) can be time-consuming, and the downtime itself is a problem due to lost employee productivity and downtime that causes lost revenue. Importantly, IT staff need to have the right skill sets to write custom scripts for failover of important applications, unless that failover is automated — something that can be accomplished if the failover software already supports a number of widely used enterprise software products.

With respect to virtualization, the increasing number of VMs, when combined with the presence of enterprise applications, means that any downtime will be more disruptive to a business with a heavily virtualized computing infrastructure than one that is more lightly virtualized. That is why strong support for virtualization (support of virtual servers) is so very important when acquiring software to protect against unplanned downtime.

SUN'S SOLARIS CLUSTER SOFTWARE

Sun Microsystems provides a high-availability software solution for enterprise workloads called Solaris Cluster. Originally known as Sun Cluster, the Solaris Cluster product has evolved over time and now supports both SPARC servers and x86 servers, along with the latest features of the Sun Solaris operating system and the OpenSolaris open source distribution of Solaris.

Solaris Cluster Addresses Customer Needs

Sun Solaris Cluster offers functionality and features that are being demanded by ACS customers, and as such, it addresses many of the pain points experienced by IT operations looking to improve availability for applications and databases.

These customer pain points include:

- Lengthy or complicated failover processes when applications have to be restarted on alternative computer systems
- Lack of sufficient onsite IT staff to write, test, and maintain custom scripts to orchestrate the failover of important applications to another server platform in the event of an outage
- Increasing numbers of applications running in virtualized computing environments that need to be protected from the consequences of any failure in their physical host platforms

Solaris Cluster Features

Sun Solaris Cluster supports both SPARC servers and x86 servers (servers based on x86 processors from Intel or AMD). This continued support across RISC and x86 architectures gives Solaris Cluster wide applicability within the Sun installed base of servers and for a wider group of customers on their non-Sun x86 servers. Solaris Cluster addresses customers' high-availability needs in three major areas: virtualization, high availability, and flexibility.

Virtualization

- Support for Solaris Container Clusters. Solaris Container Clusters allow multiple, logical clusters to be instantiated on top of a single, physical Solaris Cluster. With this, distributed applications such as Oracle Database products and open source software, such as Apache and MySQL, can now reside inside Solaris Container virtual nodes; importantly, they can run in fully isolated virtual clusters, with operational advantages in delegated management, application isolation, and resource capping.

- Support for Solaris Containers running Solaris 8, Solaris 9, Solaris 10, or Linux applications. Within the Containers, the applications cannot interfere with one another. In case of outages, the full Container, together with the embedded application, can be failed over to another physical server using the Solaris Container High-Availability Agent.

- Support for logical domains (LDoms), Sun's virtualization technology for Sun SPARC CMT-based, or CoolThreads, servers, allowing one or more images of the Solaris 10 operating system to run on the same underlying hardware. Solaris Cluster adds high availability to the consolidated environment of LDoms.

High Availability

- Support for the flexible fencing feature, which ensures data integrity in case of server outages. Solaris Cluster allows customers to choose and tune the fencing protocol according to the desired storage configuration.

- Support for multiple quorum solutions, including disk-based quorum, software quorum, as well as quorum server. Customers can tailor their quorum solution to satisfy a wide range of high-availability requirements. In doing so, they can reduce the cost of the IT infrastructure that is needed to construct an effective campus or metro cluster that is able to automatically survive a single sitewide failure.

- Support for disk-path monitoring. This means that multiple paths to data are continually monitored, allowing an automatic reboot to take place in the event that any of the paths become unavailable. This speeds failover of applications.

Flexibility

- Support for campus and metro cluster, where the physical servers within the cluster are widely separated. This separation is limited purely by site-to-site network latency rather than by specific distances. Data is mirrored by volume managers at the host level or through the replication products at the storage level.

- Support for business continuity through the layered Solaris Cluster Geographic Edition product, which integrates with a number of replication technologies, including EMC SRDF, HDS TrueCopy, Sun StorageTek Availability Suite, and Oracle Data Guard.

- Support for Oracle 9i, 10g, and 11g RAC (Real Application Clusters), running on top of Solaris Cluster, which links the server nodes supporting the RAC deployment. The integration of Oracle Clusterware with Solaris Cluster facilitates faster failure detection, more robust fencing, and a broader choice of data storage options.

- Direct support for more applications, without special scripting. Sun is providing pretested, precertified software agents that support more off-the-shelf applications, including Oracle Application Server, Oracle E-Business Suite, Oracle Databases, IBM Informix and IBM WebSphere, SAP, Sybase, MySQL, and PostgreSQL, among others.

USE CASES

The following sections describe "use case" scenarios for Solaris Cluster. That is, these are ways in which customers have deployed Solaris Cluster in conjunction with software products from Sun and other vendors. All of these use cases leverage virtualization technologies to support workload consolidation and improved levels of availability for the workloads being deployed.

Virtual Clusters Support Consolidation

Consolidation of workloads is helping many IT sites reduce the "footprints" of servers in the datacenter — a large number of which have been installed since 2001.

This kind of consolidation often brings rapid results in the form of reduced space requirements, power/cooling requirements, and management requirements within the datacenter. At the same time as it reduces "footprints," it leads to more efficient computing on the servers that remain in the configuration. This addresses the problem of maintaining deployed servers that are underutilized.

Solaris Cluster provides a way to consolidate multiple applications, and even multiple clusters, onto a single clustered system. This Solaris Container Cluster comprises multiple Solaris Containers, each of which represents a virtual "node." These Solaris Container Clusters work in the same way a standard Solaris Cluster works.

However, the business benefits are that the failure of a node simply requires a Solaris Container restart rather than a full node reboot, which significantly reduces the time for recovery. It also minimizes any interruptions to application services that are running in the same instance. Furthermore, the sharing of common networking and storage resources helps to maximize the utilization of expensive IT assets. This approach to virtualization is particularly beneficial to distributed database applications, such as Oracle databases deployments, maintaining isolation of the individual workloads while offering the capability to allocate compute resources in a flexible way.

Solaris Containers and Solaris Cluster

A single instance of Solaris 10 can be virtualized into multiple isolated environments using Solaris Containers. A Solaris Container is implemented as a collection of additional Solaris processes that share a common OS kernel and thus has a very low overhead, minimizing the impact of running an application inside such a construct. This is a major differentiation from other VMs, which are also a collection of processes but do not have low overhead.

Solaris Cluster supports both the failover of an entire Solaris Container together with the application services it contains (a "black box" solution) using the Solaris Cluster Container agent and the failover of services between Solaris Containers that are already instantiated on the individual servers within the cluster. Solaris Containers can also be used inside of a Solaris image running in an LDom.

Even so, IDC notes that for purposes of disaster recovery planning, it is always advisable to avoid configurations that have a single point of failure. Therefore, dual power feeds, access to more than one network path, and multiple paths to I/O and storage devices are elements of improving high availability in datacenter deployments.

Solaris Container Cluster supports a variety of use cases such as database consolidation, test and development consolidation, and multiple application consolidation — all of which result in cost containment. In cases where software licenses are priced upon the processing power of the system, administrators can configure a Container to use a specific number of processors (CPUs) for the applications that run in it. The limits set on the use of CPUs will determine the number of application licenses needed and the corresponding fee required, which can result in significant cost reductions.

Sun LDoms and Solaris Cluster

Sun SPARC CMT-based servers (e.g., chip multithreaded servers, which Sun calls CoolThreads servers) support virtualization through LDoms. Solaris Cluster is supported on two types of LDom: I/O domains and guest domains. A physical server may have only a few I/O domains but can support multiple guest domains. Each domain runs a single instance of Solaris, regardless of its type.

By virtualizing two or more separate, physical Sun SPARC CMT-based servers, multiple Solaris Clusters can be constructed. Each Solaris Cluster runs on Solaris instances that are completely isolated from one other on the same server. Furthermore, the version of Solaris used to support each individual cluster can differ, allowing a single set of servers to fulfill a range of roles such as production, development, regression, or user acceptance testing. This enables consolidation decisions to be made based on compatibility of resource demands rather than on operating environment, patch levels, or application software coexistence limitations.

Solaris Cluster supports both virtualization technologies, LDoms and Solaris Containers, independently or combined. When the technologies are used together, a Solaris Container runs inside of a Solaris image in an LDom and then fails over to another LDom that resides in a different physical server when outages occur.

CUSTOMER SNAPSHOTS

Following are two brief accounts of customer experiences with Solaris Cluster software running on SPARC and x86 server platforms. Both customers were interviewed by IDC, and the content was based on those interviews. These "snapshots" of customer deployments show how the Solaris Cluster and Sun server products were used in production to address business requirements for 24 x 7 x 365 availability of information services supported by the organizations' IT departments.

Liverpool Direct

Liverpool Direct Limited (LDL) is a joint venture between British Telecom (BT) and the Liverpool City Council; BT retains 80% ownership of the joint venture. The company is the largest public/private partnership of its kind in the United Kingdom, providing contact centers, ICT, human resources, payroll, welfare benefits, and revenue services. The success of LDL has caught the attention of the public sector bodies and private sector organizations across the United Kingdom seeking to provide similar services to end customers.

For many of its key application services, LDL is leveraging clustered Sun servers that enable the high levels of availability required by its customer base. The company, which has been using Sun SPARC-based server platforms since its founding in 2001, is running Oracle Database Server Enterprise Edition, Oracle R12 eBusiness Suite, and the Civica Comino DMS for document management on SPARC CMT T2000 servers. All of the applications on the six production servers run inside Solaris 10 Containers for purposes of workload isolation and management control on an application-by-application basis. This deployment applies the virtualization technology in the clustered Solaris 10 Containers to maximize availability across the campus cluster between twin datacenters.

The business benefits of this approach are clear, according to Enterprise Architect Frank Oakley. "Solaris Cluster is an end-to-end solution," Oakley said, "while usually, you have to find bits and pieces from different suppliers and build them into a cohesive, unified solution. The Solaris Containers also fit well with the Solaris Cluster technology by supporting the consolidation of workloads. Overall, business continuity is supported, and applications can continue to run, even if other workloads must be restarted, without affecting the ongoing business processes."

Eurecom

Eurecom is a privately held education and research center, based in southern France, near Nice, that maintains a well-established international partnership between academia and industry focusing on the work of engineers from industry leaders and academic institutions in the telecommunications industry. With a small IT group of just 10 staffers, Eurecom supports hundreds of end users, including researchers and engineers, and it operates 150 servers (100 of which are Unix servers) and three clusters of x86 servers running the Solaris 10 operating system and the Solaris Cluster high-availability software. In all, the IT organization supports 100 researchers and more than 200 students accessing the campus server systems.

The Solaris Cluster configurations support two primary workloads — email for the entire organization and end-user access (through hundreds of Windows PC desktop systems) to Eurecom-based research. The technical information is stored on files within the Solaris ZFS (zetabyte file system) files on the organization's central-site storage area network (SAN). A third workload will soon be included — support for dynamic Web-enabled applications that will be used by researchers and Eurecom partners. This Web computing environment, as delivered on Eurecom servers,

supports the development of future telecommunications algorithms supporting converged technologies (e.g., WiFi 3G, GSM) on next-generation cell phones.

Eurecom leverages Solaris Containers, which isolate applications within a software-defined "zone," to protect workloads from interfering with one another. Multiple applications can be supported within a single Solaris 10 operating system domain — with one application running in each Solaris Container. Combining Solaris Cluster with Solaris Containers, in addition to check-pointing and firewalls and data encryption, enhances system-level security and prevents any outside hackers from compromising the Eurecom software and applications while providing high availability to the services offered to the users.

Pascal Gros, a senior IT manager at Eurecom, said that the organization has been using Sun servers and software technology for nearly 20 years and that the relationship with Sun for services, including direct support for servers and software, is very close. He added that Sun has been able to provide custom solutions, where needed, to meet Eurecom's technical requirements.

CHALLENGES AND OPPORTUNITIES

As in other segments of the worldwide IT market, the market for ACS is highly competitive. This is especially true in the x86 server space, with software products shipping from multiple vendors for Microsoft Windows, Unix, and Linux deployments on x86 servers. There are literally dozens of high-availability software solutions, including failover software, workload-balancing software, and software for grid computing. In the broader market space, Solaris Cluster is offered on both SPARC-based and x86 server systems.

Given its product set and history of supporting high-availability requirements for enterprise workloads in the datacenter, Sun is well-positioned to compete in the ACS market space. By leveraging its 15 years of experience in building and supporting highly available server systems in both the RISC segment and the x86 segment, Sun is bringing well-tested technology to bear on a new set of IT challenges associated with bringing improved availability to the world of virtualized, scale-out infrastructure.

Sun also understands the synergistic opportunity of combining its operating system virtualization technology and its high-availability technology. Here's how it works: Solaris virtualization opens up an instance of an operating system for numerous isolated application services. This basic level of availability can be brought to the next level when running with Solaris Cluster, which allows application services to be clustered under full protection of a high-availability software product. Resources can be shared, both vertically and horizontally, reducing IT costs while protecting more application and database workloads running on the system. To address a broader audience of potential customers, Sun must communicate to customers and channel partners the business benefits of running Solaris Cluster in concert with Solaris Containers and virtualized servers (SPARC based and x86 based) running Solaris.

CONCLUSION

Achieving high availability for production applications is a lofty goal, indeed, for many IT organizations. The current economic environment, coupled with uncertainty about IT requirements and the future shape of the enterprise, make it very difficult to predict exactly how the next round of business strategies will be supported from an IT perspective. And yet business owners continue to demand reliable and available application services from the IT organizations that acquire, deploy, and maintain IT infrastructure. Providing a stable foundation for end-to-end computing is an important driver for evaluating current IT infrastructure with respect to its support for business continuity (for business processes) and high availability (for workloads).

The ability to minimize system downtime and to prevent lengthy downtime is now a top business requirement — even for Web-enabled workloads that once were considered expendable, compared with the workloads running in the application and database tiers of the datacenter.

Solaris Cluster addresses many of the current pain points within IT organizations, with respect to operational costs associated with high-availability requirements and IT staff costs for technology refresh cycles and maintenance. The 2009 enhancements to Solaris Cluster Version 3.2 have built-in support for consolidation projects designed to improve datacenter efficiency, for virtualization technologies, and for ongoing maintenance that can be accomplished during daytime work shifts rather than catch-up maintenance that is performed during overnight or weekend shifts. More enhancements to the Solaris Cluster product are also expected to ship in 2009 and 2010.

Today's enterprise needs to know that although unplanned outages occur rarely, applications and databases can be quickly restored to production. Moving from an occasional "window" of opportunity for planned downtime and repair, the modern enterprise needs to account for high availability and disaster recovery as elements of everyday, ongoing IT operations, protecting applications on a 24 x 7 x 365 basis. Doing less than that would undermine the progress that many business organizations are already achieving via virtualization and consolidation projects, all of which is saving real costs and operating more efficiently during these challenging economic times.