Workload Consolidation on Sun Fire™ UltraSPARC® Servers Using Solaris Containers

Technical Brief

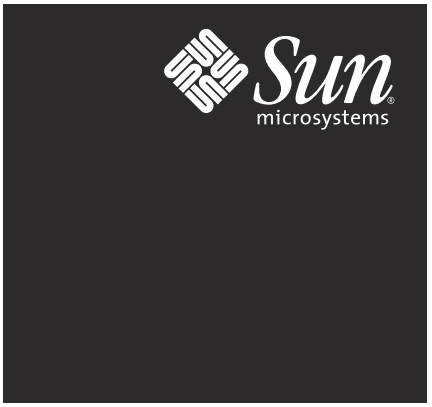

Table of Contents

Executive Overview

Consolidation of enterprise applications from multiple small servers onto larger servers has become a proven means to reduce costs and better manage system resource utilization.

The chief benefits offered by information technology (IT) consolidation include:

- Increased utilization by dynamically matching system resources to workloads
- Ability to deliver predictable service levels
- Reduced costs by consolidating to fewer servers with the associated savings in floor space, power and cooling, on-going maintenance, and software licensing
- Simplified administration from improved rationalization and reduced management

Managing multiple workloads on a single system requires enterprise-class technologies such as highly scalable hardware and operating system software, resource management controls and monitoring tools, and a set of goals and policies in which to maintain preestablished service levels.

Sun offers a range of products and services that enable customers to consolidate and/or migrate applications onto midrange to high-end platforms, enabling them to achieve the benefits of increased utilization, reduced costs, and simplified administration.

Sun Enterprise™ Servers have a proven record of nearly ten years of enabling multiple applications to be effectively managed using Dynamic System Domains and Resource Manager software. With the introduction of the Solaris™ 10 Operating System (OS), the benefits offered by consolidation can for the first time be realized on the entire range of Sun Fire™ server platforms.

When combined with breakthrough multi-threading UltraSPARC™ IV processor technology introduced in Sun's midrange server platforms, customers have a dynamic proposition of creating a secure, consolidated infrastructure on low-cost server platforms.

This Brief demonstrates how Solaris Containers deployed onto a midrange Sun Fire server platform enables an organization to effectively consolidate applications onto a single system while maintaining predictable service levels and security isolation.

An Introduction to Solaris Containers

With the escalating costs of managing vast networks of servers and software components, attention is shifting toward finding new ways to reduce IT infrastructure costs and better manage end-user service levels.

Server consolidation and virtualization techniques help by enabling systems within data centers to be visualized and managed as interconnected computing resources rather than as individual systems, but better ways must be found to provision applications and ensure shared resources are not compromised.

For these resource management techniques to be effective, companies must be able to manage their applications independently, control resource utilization according to business need, isolate faults, and ensure security between multiple applications.

As an integral part of the Solaris 10 OS, Solaris Containers isolate software applications and services using flexible, software-defined boundaries. A breakthrough approach to virtualization and software partitioning, Solaris Containers let many private execution environments be created within a single instance of the Solaris OS. Each environment has its own identity, separate from the underlying hardware, yet behaves as if it is running on its own system, making consolidation simple, safe, and secure.

Because Solaris Containers are location independent, applications can be moved as needed. Each application runs in its own private environment—without dedicating new systems—and many applications can be tested and deployed on a single server without fear that they will impact one another. System and network resources can be allocated and controlled on a fine-grained basis, helping simplify computing infrastructures and improving resource utilization. As a result, companies can better consolidate applications onto fewer servers without concern for resource constraints, fault propagation, or security, simplifying service provisioning.

Benefits of Solaris Containers

Increase Resource Utilization and Deliver Predictable Service Levels

No two organizations have the same type of workload or use system resources in the same way. Some utilize batch compute servers and application or database servers, while others share systems in more complex ways. Regardless of how the data center is arranged, a vast amount of potential computing capacity often remains untapped. Users are continually searching for more computing resources to help solve problems, resulting in systems that are alternately overloaded or underutilized.

Server consolidation provides the opportunity to lower costs by reducing the hardware and system administration required to run applications, but provisioning applications with the appropriate resources on a shared system can be tricky. Solaris Containers give businesses the ability to prioritize applications and have control over resource usage. Computing resources—CPUs, physical memory, network bandwidth, and more—can be dedicated to a single application one moment, or shared with others in an instant, without moving applications or rebooting the system. For example, a database, Web server, and batch application, each running on its own system, can be consolidated onto a single server configured to give each access to one-third of available system resources. That same server can be automatically reconfigured so that the Web server receives 75 percent of network bandwidth during peak load conditions. With the ability to dynamically move resources to where they are needed most, Solaris Containers help businesses increase resource utilization while ensuring that service-level agreements are met.

Ensure Application Isolation and Affect Greater Consolidation

Consolidation projects are frequently hampered by a lack of trust in how they will work, which discourages system sharing. With Solaris Containers, organizations can gain control and establish isolation mechanisms in order to improve resource utilization. Each application can be given a private environment in which to run, virtually eliminating error propagation, unauthorized access, and unintentional intrusions. Providing a fine granularity of control, Solaris Containers ensure all workloads have access to an appropriate amount of computing resources, and that no workload is able to consume the entire system unless authorized to do so.

Because resources are isolated and dedicated to a Solaris Container and its applications rather than a complete system, highly efficient application consolidation is now possible. For example, Web servers typically listen to network port 80, and in order to do that they require root privileges, which represent a high security risk. To reduce these risks and run multiple Web servers per system, each Web server can run in a Solaris Container and listen to its own unique port 80 and operate in an isolated and secure manner.

Rapidly Test and Deploy Applications

Developing new applications and services, and getting them operational quickly, is critical. Solaris Containers speed the deployment of new applications. Environments can be preconfigured on multiple systems and moved to where they are needed. Applications and services can be modified and tested in one Solaris Container and later moved to an online area without impacting other users. Multiple deployment scenarios can be tested with ease, and administrators can roll back to previous settings and configurations if needed. Because application testing no longer requires dedicated systems that can sit idle most of the time, companies can spend less to get services up and running.

Reduce Costs and Improve ROI

Maintaining flexibility and improving manageability are essential to effective cost-cutting strategies. Solaris Containers help businesses organize resources and understand how they are being used. With the ability to securely and dynamically manage and tune Solaris Containers, companies can host multiple applications on one system and use expensive resources to greater effect. In addition, Solaris Containers gather and use workload-based usage data rather than system data, making it easier to accurately charge for resources used. System-related administrative tasks are performed for the entire system, saving time and money.

Software Partitioning and Resource Management

In earlier Solaris OS versions, resource management required system administrators to add specific software packages to the server. With the release of the Solaris 9 Operating System, resource management facilities were built into the operating system, and resource controls could be implemented using *Projects, Resource Pools, and Fair Share Scheduler* facilities. This was the first implementation phase of Solaris Container technology.

The second phase is implemented with the introduction of the Solaris 10 OS. New to the Solaris 10 OS is a unique partitioning technology called Solaris Zones, which can be used to combine applications safely and with great manageability.

Another addition to the Solaris Containers technology is Dynamic Resource Pools, which provide the capability of separating workloads so that consumption of system resources does not overlap or conflict. They provide a persistent configuration mechanism for processor set and scheduling class assignment. In addition, the dynamic features of resource pools enable administrators to adjust system resources in response to changing workload demands as they happen, without any need to reboot the system.

The Solaris Zones partitioning technology can be used to virtualize operating system services and provide an isolated, secure environment for running applications. When a zone is created, an environment is produced in which processes are isolated from the rest of the system. This isolation prevents processes that are running in one zone from monitoring or interfering with processes running in other zones. Zones also provide means to isolate applications from physical devices on the system, such as network interfaces and storage devices.

Scalable Platform for Optimal Workload and Application Performance

The platform deployed to consolidate and manage multiple workloads and resources requires highly scalable, flexible enterprise-class technology. This platform must allow resources to be scaled, virtualized, managed, and deployed where and when the capacity is needed.

Based on UltraSPARC IV processor technology, Sun's line of midrange and high-end servers are powerful, scalable systems optimized for mid-tier, back-end, and mission-critical applications and deliver industry-leading price/performance and availability for enterprise-class computing.

The Sun Fire V490–E25K servers deliver up to double the application throughput in the same footprint with UltraSPARC IV chip multi-threading (CMT) technology and the Solaris OS, helping ensure optimized application delivery and lowering total cost of ownership.

Sun™ systems have a proven record of nearly ten years of enabling multiple applications to be effectively managed using Dynamic System Domains and Resource Manager software. In addition, with Sun's legacy application support and Solaris binary compatibility guarantee, there is no need to migrate the OS or port applications to take advantage of new hardware features—offering unprecedented levels of investment protection for owners of Sun's previous-generation systems.

Consolidating applications using Solaris 10 and Solaris Container technology on the Sun Fire Enterprise platform helps customers reduce costs, increase utilization, and simplify administration.

Sun Fire V890 Server Architecture

At sub-$50K price points, the Sun Fire V890 server is a midrange system that enables customers to run demanding mission-critical workloads in a low-cost, secure, consolidated environment.

With the computational throughput of multi-threaded processors, the Sun Fire V890 server gives enterprises the capabilities of large systems with the economies of footprint often found only on smaller platforms. Multiprocessing and multi-threading technology delivers up to twice the productivity of Sun's previous generation of servers by running tasks in parallel—speeding database queries, providing remote file service, and accelerating computationally intensive applications. Running the Solaris 10 OS can further improve throughput performance through improved threading libraries and the Dynamic Tracing (dTrace) capability.

This system also features a number of enhancements to improve the availability of systems running mission-critical or data-center applications, including redundant hot-swappable components and optional redundant paths to internal storage. Multiprocessing, high-speed interconnects, high-performance and capacity storage systems, and the Solaris OS all help provide the highest level of scalability, reliability, security, and performance needed for many classes of applications.

Benchmark Test Environment

The resources and processes to consolidate and manage three disparate enterprise-class applications on the Sun Fire V890 server demonstrated that the system could handle mixed and diverse workloads in a single system while maintaining predictable service levels.

The testing activity described in this section made use of the Solaris Container technology, which is available as standard in the Solaris 10 Operating System. In this benchmark, the Solaris Containers managed a combination of diverse workloads on a Sun Fire V890 midrange server, such as:

- **A Secure Web Serving Application**—Multiple interactive users accessing a Web Server with a mix of secure HTTP requests over a wide range of file sizes, using industry-standard Secure Sockets Layer (SSL) encryption. The Sun Java™ System Web Server 6.1 was used for this workload.
- **An Enterprise JavaBeans™ Business-to-Business Application**—Multiple interactive users accessing a J2EE™ Application Server with a mix of transactions that emulate both an order/inventory and a business-to-business supply-chain management environment. The Sun Java System Application Server 8.0 was used to host the Enterprise JavaBeans application.
- **An RDBMS-based Decision Support Application**—Simulating a data mart, with users submitting complex SQL queries on a large database. This is a batch-oriented workload with high system resource demands, which are both I/O and compute-intensive. The Oracle 10g Database was used as the RDBMS for both the J2EE application server and DSS workloads, using separate database instances.

To obtain a baseline of system utilization measurements and throughput capacity, each of the workloads were run concurrently without resource controls. Then, for comparison, all three workloads were run simultaneously using Solaris Containers to isolate applications and to control consumption of system resources.

Baseline Performance of Mixed Workloads without Resource Management

The following diagrams demonstrate the cumulative CPU utilization and transaction throughput of running all three workloads together on the Sun Fire V890 server without the benefits of Solaris Containers.

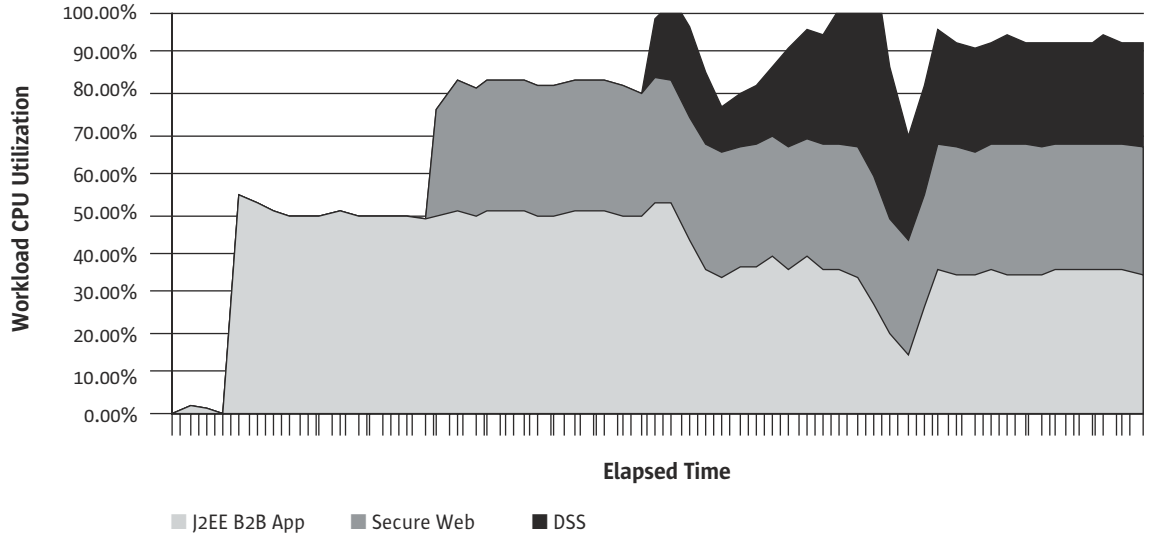

Figure 1. Mixed Workload CPU Utilization, Without Containers

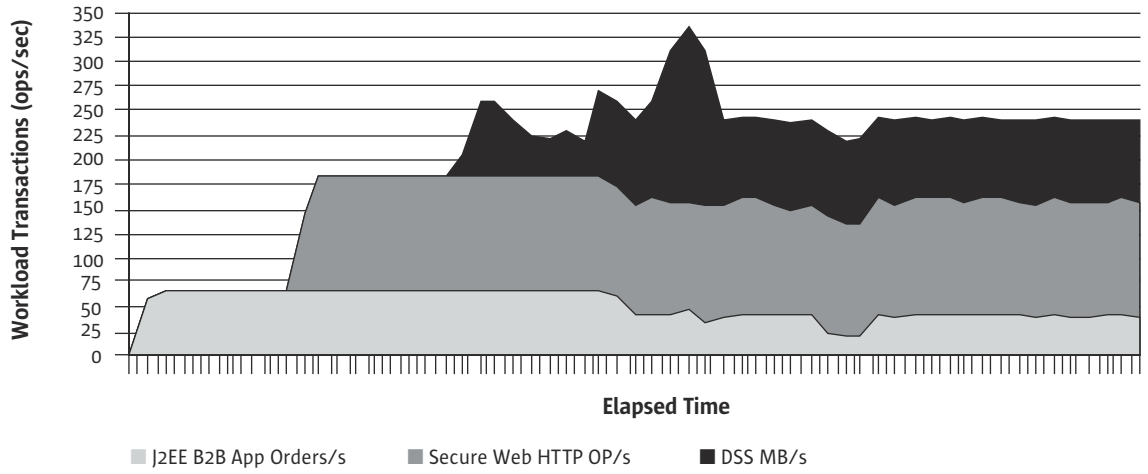

Figure 2. Mixed Workload Transaction Throughput, Without Containers

As demonstrated above, when all three workloads run simultaneously, the J2EE B2B application throughput suffers significantly and the DSS workload shows inconsistent throughput.

Individually, the J2EE B2B Application initially consumes approximately 50 percent of system resources but is eventually reduced to only 33 percent of resources when the other workloads are introduced. This results in a reduction of transaction throughput, which can also negatively affect user response time.

For a production environment in which the transaction response times are a critical component of the business, running multiple workloads without resource controls can result in interruptions or reduced service levels.

In addition, the mixture of secure and non-secure applications on the same system require the implementation of controls to ensure that users or processes in one application space cannot affect data or resources in another.

Mixed Workload Performance Using Solaris Zones and Resource Pools

To ensure predictable service levels and workload isolation, each application was configured in a separate Container and enabled with a unique resource pool in order to meet the following business criteria:

Application	% CPU Demand	% CPU Allocated
J2EE B2B App	50	50
Secure Web Server	30	25
RDBMS	25	25

Although the CPU demand of the Secure Web Server transactions consumed as much as 30 percent of the total system processor resource, in this demonstration the resource pool was configured to allow only 25 percent of system resources.

For both the Secure Web Server and J2EE B2B application Containers, private network interfaces were configured using the Container administration utilities. This was done to ensure each Container could send and receive data from its respective client systems.

In addition, each of the three Containers had separate file systems mounted with the data required to support the application software. These were not visible to the other Containers.

The following diagrams demonstrate processor utilization and transaction throughput when running all three workloads together, each within their own Container.

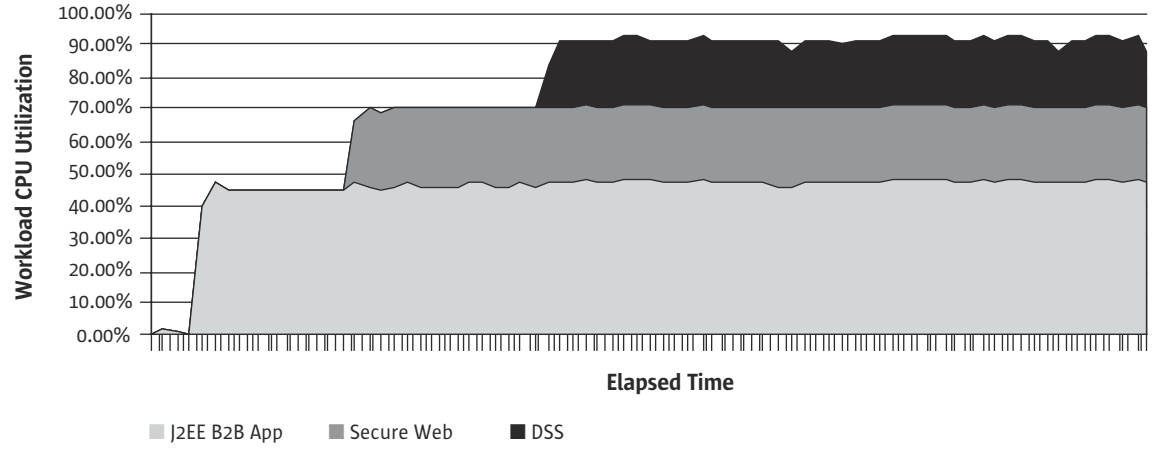

Figure 3. Mixed Workload CPU Utilization, With Containers

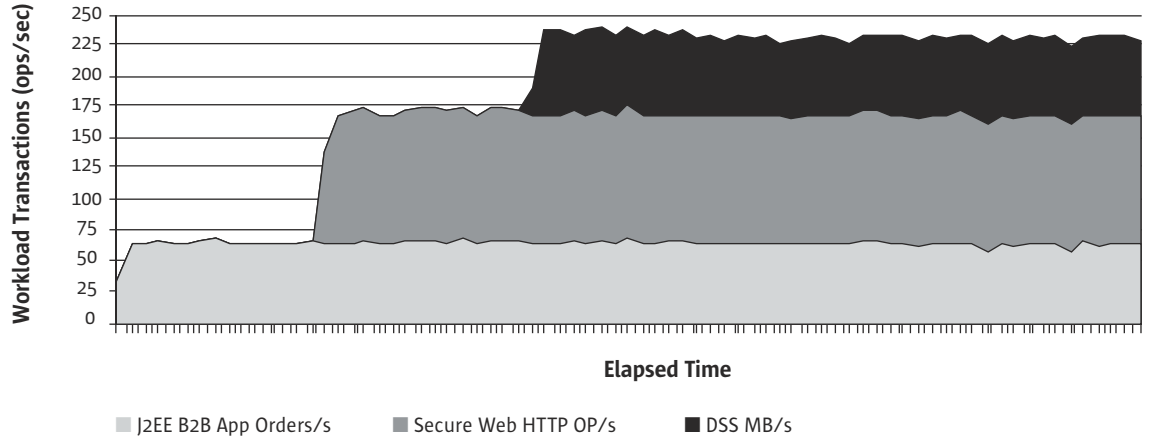

Figure 4. Mixed Workload Transaction Throughput, With Containers

After configuring each of the workloads to use the preselected amount of system resources for the specified Container, the throughput of each of the workloads reaches its desired level of performance. This remains constant through the duration of the workload execution to consume 93 percent of available processor resources.

Conclusion

The results of these benchmarks show that consolidating applications to the Sun Fire V890 server running the Solaris 10 Operating System provides an ideal solution to handle mixed and diverse workloads in a single system while maintaining predictable service levels.

Solaris 10 Containers demonstrate the value of operating system virtualization and partitioning to isolate application execution environments with the flexibility to effectively manage consolidated commercial workloads.

As a result, companies can increase utilization, reduce costs, and simplify administration of their IT infrastructure.

Specific benefits of deploying Solaris Containers on Sun midrange servers are demonstrated as:

Increased Utilization

- Improve usage of system resources by dynamically configuring, creating, deleting, or resizing containers according to workload size, type, or schedule.
- Meet changing resource demands by shifting resources from one application to another with no downtime.
- Increase utilization by deploying up to 8000 secure, fault-isolated software partitions—to run multiple applications on the same server with one instance of the Solaris OS.

Reduce Costs

- Reduce costs by consolidating to fewer servers and using partitioning technologies to dynamically allocate and share existing resources.
- Reduce floor space and power costs by consolidating to fewer servers and creating multiple “systems within a system” with Solaris Containers.
- Reduce cost of ownership with a shared server among departments or projects, using Sun’s partitioning technology to dynamically control system resources without impacting the entire system.
- Reduce the costs of maintaining multiple operating systems by using Solaris Containers supporting multiple independent execution environments within the same instance of the Solaris OS.

Simplify Administration

- Reduce the complexity of managing hundreds, even thousands of applications, multiple servers, and multiple operating systems by consolidating and utilizing partitioning to run multiple applications on fewer servers.
- Rapidly deploy new applications in minutes by using Containers that can be instantly created or modified with no impact to other applications on the system.
- Reduce the complexity of applying upgrades and enhancements to multiple application environments with Solaris Containers that share a single Solaris OS instance.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN **Web** sun.com

The Network is the Computer™

Sun Worldwide Sales Offices: Argentina +5411-4317-5600, Australia +61-2-9844-5000, Austria +43-1-60563-0, Belgium +32-2-704-8000, Brazil +55-11-5187-2100, Canada +905-477-6745, Chile +56-2-3724500, Colombia +571-629-2323 Commonwealth of Independent States +7-502-935-8411, Czech Republic +420-2-3300-9311, Denmark +45 4556 5000, Egypt +202-570-9442, Estonia +372-6-308-900, Finland +358-9-525-561, France +33-134-03-00-00 Germany +49-89-46008-0, Greece +30-1-618-8111, Hungary +36-1-489-8900, Iceland +354-563-3010, India–Bangalore +91-80-2298989/2295454; New Delhi +91-11-6106000; Mumbai +91-22-697-8111, Ireland +353-1-8055-666 Israel +972-9-9710500, Italy +39-02-641511, Japan +81-3-5717-5000, Kazakhstan +7-3272-466774, Korea +822-2193-5114, Latvia +371-750-3700, Lithuania +370-729-8468, Luxembourg +352-49 11 33 1, Malaysia +603-21161888 Mexico +52-5-258-6100, The Netherlands +00-31-33-45-15-000, New Zealand–Auckland +64-9-976-6800; Wellington +64-4-462-0780, Norway +47 23 36 96 00, People's Republic of China–Beijing +86-10-6803-5588 Chengdu +86-28-619-9333; Guangzhou +86-20-8755-5900; Shanghai +86-21-6466-1228; Hong Kong +852-2202-6688, Poland +48-22-8747800, Portugal +351-21-4134000, Russia +7-502-935-8411, Saudi Arabia +9661 273 4567 Singapore +65-6438-1888, Slovak Republic +421-2-4342-94-85, South Africa +27 11 256-6300, Spain +34-91-767-6000, Sweden +46-8-631-10-00, Switzerland–German 41-1-908-90-00; French 41-22-999-0444, Taiwan +886-2-8732-9933 Thailand +662-344-6888, Turkey +90-212-335-22-00, United Arab Emirates +9714-3366333, United Kingdom +44-1-276-20444, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800, or online at sun.com/store

1/05