



An Oracle White Paper  
April 2014

# Consolidation Using the Fujitsu M10-4S Server

Executive Overview .....	1
Why Server and Application Consolidation? .....	2
Requirements for Consolidation.....	3
Consolidation on Large, Vertically Scalable SMP Servers .....	3
Vertically Scalable High-End SMP Servers .....	4
The Fujitsu M10-4S SMP Server .....	4
Scalability and Capacity of the Fujitsu M10-4S Server .....	5
Fujitsu M10-4S Virtualization and Resource Management.....	8
Oracle Solaris .....	9
Physical Partitions.....	9
Oracle VM Server for SPARC .....	9
Oracle Solaris Zones .....	10
Oracle Solaris Resource Manager .....	10
Fair Share Scheduler .....	11
Oracle Enterprise Manager Ops Center .....	11
Fujitsu M10-4S High-Availability Features for Consolidation .....	11
Case Study .....	12
Conclusion .....	13
Further Information .....	13

## Executive Overview

The benefits of enterprise consolidation are well understood. Consolidating workloads, applications, databases, operating system instances, and servers makes it possible to reduce the number of resources under management, resulting in improved system utilization rates and lower costs. Higher utilization rates reduce the need to make additional hardware purchases. At the same time, consolidation can contribute to strategic goals such as improving security, delivering more-predictable service levels, and increasing application deployment flexibility.

For successful consolidation deployments, it is necessary to select a server platform that has the scalability to support many application instances. Additionally, the server platform must have the high availability needed for mission-critical applications, the resource management and virtualization capabilities to simplify managing numerous applications, and the tools to manage the consolidated environment.

The Fujitsu M10-4S server delivers on all these requirements and is an ideal solution for server consolidation. With this server, IT managers can create pools of compute resources that can be rapidly and dynamically allocated to meet new and changing workloads.

## Why Server and Application Consolidation?

Traditionally, applications have been deployed on a single server for each application instance. With complex enterprise applications, this style of deployment means that data centers require many servers for a single application, with separate servers for the Web tier, application tier, and database tier.

Furthermore, many enterprise applications require test and development servers in addition to the production servers. Commonly the production servers, when initially deployed, have enough headroom to support spikes in the workload, but as the applications grow, the only way to add more capacity is to add more servers, thereby increasing complexity. As the number of servers increases, the number of OS instances that need to be managed grows, adding further layers of complexity and reducing IT flexibility.

In the one-application-per-server deployment model, server utilization is normally very low—between 10 percent and 30 percent—which is a very inefficient use of server resources. Each server needs to be large enough to handle spikes in workload but normally will need only a small part of the server capacity.

Figure 1, which shows many small servers running a single application instance, illustrates this point. Each of these servers needs to have enough headroom to meet peak capacity requirements and cannot “share” headroom with other servers that need more capacity or have excess capacity.

If these servers could share headroom, loaning it out or borrowing it as needed, they would have higher utilization rates. But with multiple applications consolidated on a single larger server, where resources shift dynamically from application to application, the workload peaks and troughs tend to even out and the total compute requirement is less variable. The more applications that are consolidated, the more even the server usage. Applications that are consolidated on a larger server benefit from shared headroom, so consolidating applications can lead to much higher server utilization as excess capacity is reduced significantly.

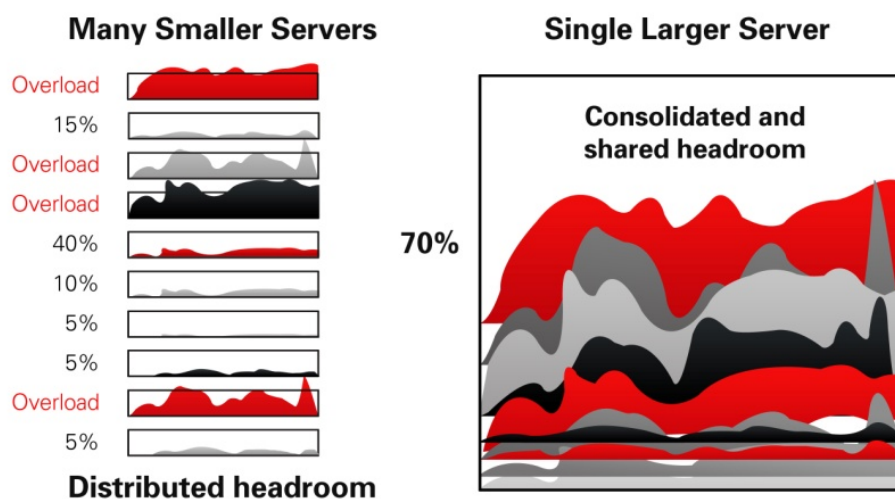


Figure 1. Large symmetric multiprocessing servers can consolidate and share headroom.

Improved server utilization means more-efficient use of server resources, which improves ROI and reduces the total server hardware required to meet workload requirements.

Consolidating many older and smaller servers onto fewer larger and newer servers provides many benefits beyond improved utilization. The newer servers will have more capacity, better performance, better energy and space efficiency, and improved availability features and will be easier to manage.

## Requirements for Consolidation

Servers used for consolidation must provide scalability and high capacity, high availability, and simple upgrade paths. They must also enable reuse of existing applications and have effective virtualization and resource management tools. Because applications are combined on consolidated servers, these servers need the capacity to handle dozens of workloads of all types. The performance of each application when consolidated with other applications must match or exceed the performance it achieves when it is deployed by itself on its own server.

Consolidation, by definition, means putting “more eggs in one basket,” so a system failure will have a greater effect on application availability than if each application were deployed on its own server. Servers used for consolidation must have high-availability features, in both hardware and software, to reduce planned as well as unplanned downtime. Consolidation servers must be extremely reliable, so that they rarely go down. They also need to have advanced serviceability features so they can be reconfigured, upgraded, and repaired with minimal or no downtime.

Consolidation servers are used mainly to run older applications in a newer environment, so they must be able to run legacy applications as well as new applications.

A consolidation environment will have many workloads of different types, and these various workloads will all have specific patch, resource, security, and performance requirements. In many cases, the operating system will have enough tools to manage multiple applications, but in other cases, applications will require separate environments to run effectively. Virtualization and resource management tools are required so that the pool of resources in a consolidation server can be partitioned and deployed as needed for multiple applications. Virtualization enforces application separation, and resource management guarantees that the performance requirements of each application are met.

## Consolidation on Large, Vertically Scalable SMP Servers

Large symmetric multiprocessing (SMP) servers, such as Fujitsu’s M10 servers, have terabytes of RAM and dozens of processors and I/O slots, all housed in a single rack that can be deployed in a single OS instance.

In essence, vertically scalable servers are large pools of resources that can support dozens of workloads of various sizes and types to simplify consolidation and application deployment. New applications can be deployed on a large SMP server, eliminating the need to install a server for each new application. Existing applications can grow by taking advantage of the extra headroom available.

## Vertically Scalable High-End SMP Servers

All servers consist of the same essential components, but different server architectures combine, connect, and utilize these components in different ways.

Vertically scalable servers—generally larger SMP servers hosting eight or more processors—have a single instance of the OS for managing multiple processors, memory subsystems, and I/O components, which are contained within a single chassis. Most vertically scalable servers, such as the Fujitsu M10-4S server, can also be partitioned with virtualization tools to create multiple instances of the OS that use subsets of the server’s resources. Virtualization tools are used to share or separate resources as required, based on the workload and the security and availability requirements.

In a vertically scalable design, the system interconnect is a high-speed interconnect that provides both low latency and high bandwidth. In vertical or SMP systems, memory is shared and appears to the user as a single entity. All processors and I/O connections have equal access to all memory, eliminating data placement concerns.

The cache-coherent interconnect maintains information on the location of all data, regardless of its cache or memory location. There are no cluster managers or network interconnects in SMP servers, because the internal interconnect handles all data movement automatically and transparently. Adding resources to the rack involves inserting “building blocks” with additional processors, memory, and I/O subassemblies. Vertical architectures can also include clusters of large SMP servers that can be used for a single large application.

High-end SMP servers greatly simplify application deployment and consolidation. Large SMP servers have a huge pool of easily partitioned processor, memory, and I/O resources. This pool of resources can be assigned dynamically to applications with system management tools.

With these such servers, workloads of any size are easy to deploy, because resource partitions can be as small as a single processor or as large as a Fujitsu M10-4S server with 64 processors, 1,024 cores, 32 TB of RAM, and 128 PCI Express (PCIe) I/O slots. Access to all resources is equal. A large SMP server functions as a cloud of resources that is tightly integrated and extremely flexible.

Because a large SMP server can handle virtually all sizes and types of workloads, a data center can standardize on a single system to simplify management, service, and deployment. Data centers can standardize on one type and size of SMP server and deploy applications on the servers until they are full, only then deploying a new server of the same type to provide extra capacity for new workloads.

## The Fujitsu M10-4S SMP Server

Now that this paper has introduced the fundamental concept of vertical scaling, it is important to take a closer look at the underlying hardware and software technologies that deliver the benefits. This section describes the Fujitsu M10-4S server and the features that enable organizations to take advantage of the vertical scaling model to consolidate and achieve higher utilization rates, lower TCO,

and more-predictable service levels and achieve more-efficient use of resources. Figure 2 shows the specifications of the Fujitsu M10-4S server.

- **64 SPARC64 X+ CPUs**
  - 3.7 Ghz
  - 16 cores/cpu and 2 threads/core
  - 24 MB of Level 2 cache
- **32 TB of main memory**
- **128 disks**
- **128 I/O slots**
- **928 slots with PCI Expansion Unit**

Fujitsu M10-4S




Figure 2. The Fujitsu M10-4S's specifications enable the server to support vertical scaling.

### Scalability and Capacity of the Fujitsu M10-4S Server

Achieving great application performance requires a balanced system with fast processors; a fast interconnect; fast I/O; a scalable operating system; optimized applications; and a high degree of reliability, availability, and scalability. Oracle's approach to system design has always been focused on achieving a balance among these core elements.

#### Fujitsu M10-4S Modular Architecture

The Fujitsu M10-4S server is a two- or four-socket server powered by the latest SPARC64 X+ processor in a four-rack-unit (4U) form factor. Each Fujitsu M10-4S 4U server is a building block that can be connected with other Fujitsu M10-4S building blocks to create single-instance SMP servers with as many as 64 processors. Each Fujitsu M10-4S server supports as many as 64 DDR3 memory DIMM slots, eight PCIe Gen3 slots, and eight 2.5-inch hard disk drives. Figure 3 shows the modular architecture of the Fujitsu M10-4S. In this architecture, 1 to 4 Fujitsu M10-4S building blocks can be connected through just copper cables and can be mounted in any standard rack. For 5 to 16 building blocks, one or two Fujitsu expansion racks and two or four Crossbar Box (XB Box) interconnect switches are required. No matter how many building blocks are connected, the result can be a single instance of the Oracle Solaris OS or many Oracle Solaris OS instances in partitions.

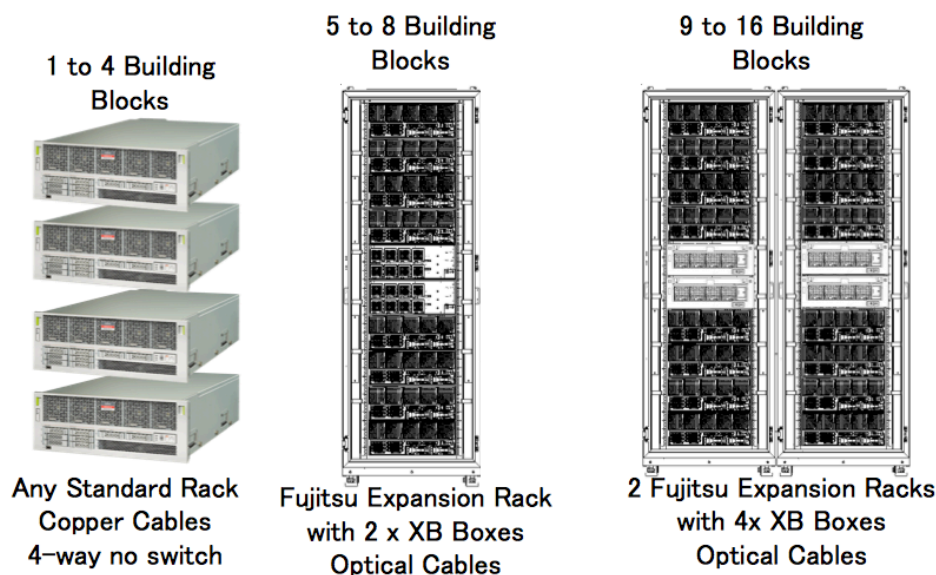


Figure 3. The Fujitsu M10-4S boasts a modular architecture.

#### Processors and System Interconnect

The Fujitsu M10-4S uses the new SPARC64 X+ 16-core processor running at 3.7 GHz. The SPARC64 X processor has 24 MB of Level 2 (L2) cache for excellent single-thread performance, making it ideal for just about any workload. The SPARC64 X+ processor has 2 threads per core and 16 cores per processor, for a total of 32 threads per processor. It supports 512 GB of memory per processor. Therefore, the Fujitsu M10-4S scales up to 64 SPARC64 X+ processors, 1,024 processor cores, 2,048 threads, and 32 TB of main memory, all of which can be used by a single instance of Oracle Solaris or by many instances with Fujitsu M10-4S partitioning technologies.

For optimum application performance, it is generally more important that processors run at their maximum capacity, with high rates of utilization, than it is to have the fastest-available processor. A fast processor running at 50 percent of its capacity may actually be slower in terms of delivered performance than a slower processor running at 80 percent capacity. In addition, as the number of processors in a system increases, the bandwidth and latency of the system interconnect become more important than the speeds of individual processors in delivering performance.

The system interconnect moves data from disk, memory, and network interfaces to the processors. It is also used to move data locations, a critical task for maintaining cache coherence. If the system interconnect has slow address bandwidth, the processor will often become idle, awaiting data. The Fujitsu M10-4S uses a cable interconnect to connect 1 to 4 Fujitsu M10-4S building blocks and uses an optical interconnect and dual XB Box switches to connect 5 to 16 Fujitsu M10-4S building blocks. The Fujitsu high-performance interconnect has the high bandwidth and low latency necessary for scaling up to 64 processors.

The key benefit of large SMP servers with tightly coupled processors, memory, I/O, and interconnect is the fast internal data transfer rates. Compute capacity and capability are determined by how fast data



can be moved in a system. The bandwidth of the Fujitsu M10-4S server's interconnect is 6,553 GB/sec with 16 building blocks connected.

The Fujitsu M10-4S server is a large pool of resources that can handle workloads of any size, with no concerns about data layout. Because all processors access all memory and all I/O slots equally, a large workload is as easy to deploy and support as a smaller workload. The only difference is the number of resources used by each workload. Figure 4 shows a 64-processor Fujitsu M10-4S and how it can support varying workload sizes.

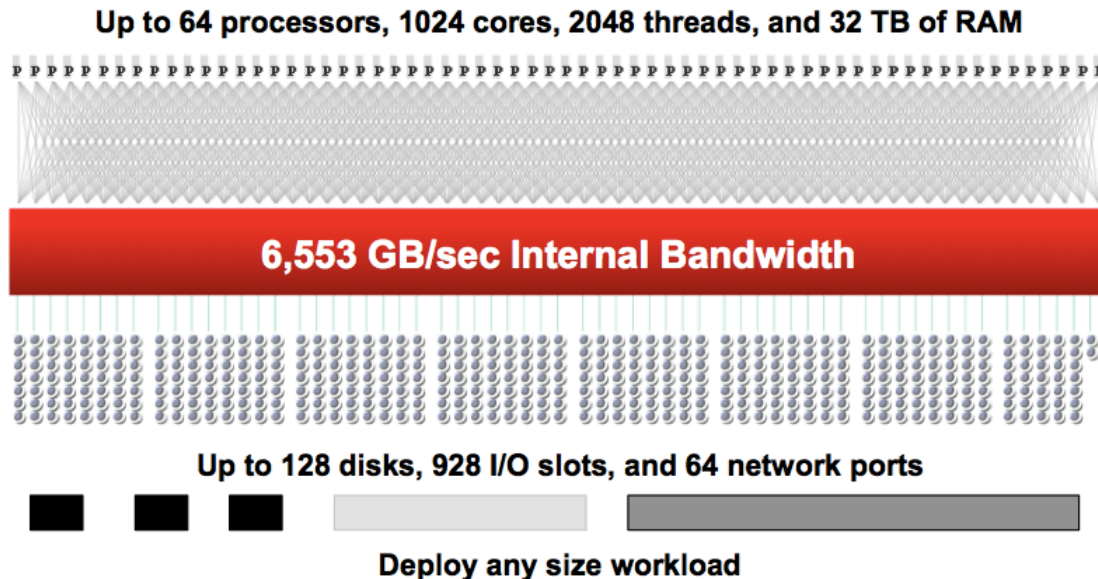


Figure 4. This schematic of a Fujitsu M10-4S server shows memory, processors, and I/O slots. All processors have equal access to all memory and I/O. Data is moved at extremely high rates between resources by a high-bandwidth, low-latency interconnect.

### Input and Output

For many applications, fast I/O is essential for getting data from the disks and the network to the system interconnect and on to the processors. An I/O bottleneck can adversely affect even the fastest interconnect and processors. I/O bandwidth is a function of the throughput of the individual I/O subsystems and the number of I/O slots. The number of disk spindles can greatly affect I/O capability, so the more disks deployed for an application, the better it will run. The Fujitsu M10-4S has as many as 128 PCIe Gen3 I/O slots internally and up to 928 PCIe Gen3 I/O slots with the optional PCI Expansion Unit.

### CPU Activation

CPU Activation is a “capacity on demand” (COD) acquisition model whereby a customer buys processors but cannot use all of the processor cores until they purchase CPU Activation licenses. Buying new licenses is quick and allows for almost instantaneous capacity increases. This reduces the initial cost of the hardware and increases flexibility. CPU Activation has a granularity of a single processor core which makes it easy to deploy just the right amount of processor resources for each

application. CPU Activation licenses can be moved from one Fujitsu M10 server to any other Fujitsu M10 server or Physical Partition which simplifies meeting changing workload requirements.

### Operating System

A poor operating system can bottleneck the best hardware and have a negative impact on performance, capacity, and availability. Oracle optimizes the Oracle Solaris OS for the scalability, availability, and performance requirements of large SMP servers. Moreover, because Oracle engineers the entire stack—application, database, processor, system interconnect, and OS—it is able to innovate at each level and optimize the whole system for specific types of workloads.

The result is a balanced system and robust performance tailored to the needs of the application. Oracle Solaris has been deployed on large SMP servers since the mid-90s, and it has the proven scalability to enable all the processor, memory, and I/O resources of the Fujitsu M10-4S server to be used for a single large application or dozens of consolidated applications.

## Fujitsu M10-4S Virtualization and Resource Management

The following sections examine the virtualization and resource management technologies that enable the deployment of many applications together to improve system utilization, optimize the use of computing resources, and deliver greater ROI from IT investments. Figure 5 shows the various levels of virtualization technologies available, at no cost, on the Fujitsu M10-4S server. At the bottom of the virtualization stack is the SPARC64 X+ processor. The first level of virtualization, physical partitions (PPARs), are hardware partitions. In each PPAR are hypervisor-based Oracle VM Server for SPARC partitions. The next level of virtualization is Oracle Solaris Zones, a feature of Oracle Solaris. Each instance of Oracle Solaris has various resource management tools that are very useful for managing many applications on a single server. The next few sections describe these virtualization and resource management technologies in detail.

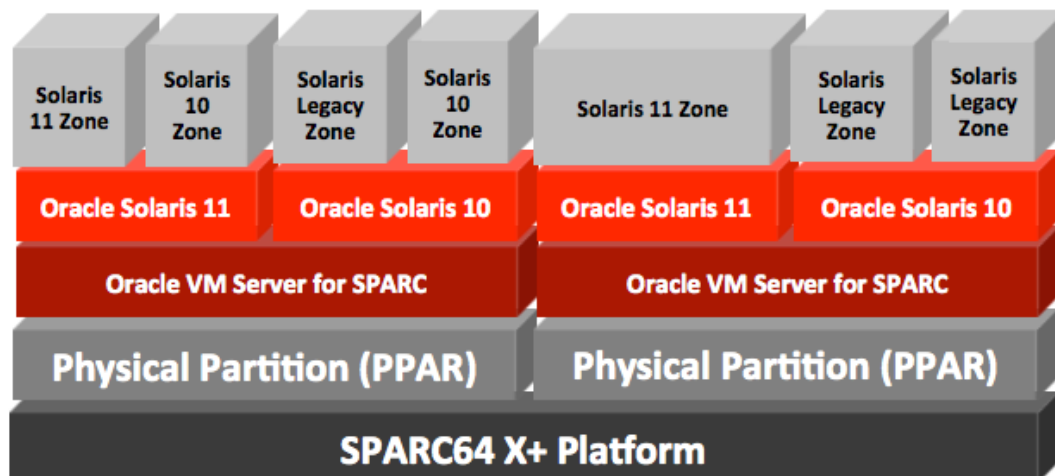


Figure 5. This figure illustrates the virtualization technology stack of the Fujitsu M10-4S server.

## Oracle Solaris

The Oracle Solaris OS is very efficient at scheduling large numbers of processes among all the processors in a given server or domain, dynamically migrating processes from one processor to the next, depending on the workload. Many enterprises run more than 100 Oracle Database instances on single SPARC servers, using no virtualization tools. Oracle Solaris is able to effectively manage and schedule the database processes.

With this approach, a large vertically scalable server can assign resources as needed to the many users and application instances residing on the server. Using Oracle Solaris to balance workloads can reduce the processing resource requirements, resulting in fewer processors, smaller memory, and lower acquisition costs.

## Physical Partitions

Physical partitions are electronically isolated partitions. They make it possible to run multiple applications and multiple copies of Oracle Solaris on a single server.

PPARs enable administrators to isolate hardware or security faults and constrain their scope. The result is a superior level of system availability and security. PPARs, now in their third generation, are a very successful and established partitioning option for the UNIX server market.

With PPARs, software and hardware errors and failures do not propagate themselves beyond the domain in which the fault occurred. Complete fault isolation between PPARs limits the effect on applications of any hardware or software errors. This helps maintain a high level of availability in Fujitsu M10-4S servers, which is necessary when consolidating many applications. PPARs keep the administration of each domain separate, so a security breach in one domain does not affect any other domain.

PPARs are combinations of Fujitsu M10-4S building blocks, so the granularity is the resources in a single building block. Each building block can have two to four SPARC64 X+ processors, up to 2 TB of RAM, eight hard drives, eight network ports, and 16 PCIe Gen3 I/O cards. Building blocks can be moved dynamically from one PPAR to another with no downtime of the Oracle Solaris instance or the applications inside the PPARs. Building blocks can be added or removed for service or reconfiguration or to add more capacity. The ability to dynamically resize PPARs makes it much easier to meet changing workload requirements, which is important in server consolidation.

## Oracle VM Server for SPARC

Supported in all servers from Oracle using Oracle's multicore/multithreaded technology, Oracle VM Server for SPARC provides full virtual machines that run independent operating system instances. Each operating system instance contains virtualized CPU, memory, storage, console, and cryptographic devices. Within the Oracle VM Server for SPARC architecture, operating systems such as Oracle Solaris 10 or Oracle Solaris 11 are written to the hypervisor, which provides a stable, idealized, and virtualizable representation of the underlying server hardware to the operating system in each domain. Each domain is completely isolated, and the maximum number of virtual machines created on a single platform depends on the capabilities of the hypervisor rather than the number of physical hardware

devices installed in the system. For example, a Fujitsu M10-4S server with four Fujitsu M10-4S processors in a single PPAR supports as many as 256 virtual domains, and each individual virtual domain can run a unique OS instance.

Oracle VM Server for SPARC 3.x has the ability to perform a live migration from one domain to another. As the term *live migration* implies, the source domain and application no longer need to be halted or stopped. A logical domain on a Fujitsu M10-4S can be live-migrated to another PPAR on the same server or to another Fujitsu M10-4S server.

By taking advantage of domains, organizations gain the flexibility to deploy multiple operating systems simultaneously on a single platform. In addition, administrators can leverage virtual device capabilities to transport an entire software stack hosted on a domain from one physical machine to another. Domains can also host Oracle Solaris Zones to capture the isolation, flexibility, and manageability features of both technologies. Deeply integrating Oracle VM Server for SPARC with the SPARC64 X+ processor, Oracle Solaris increases flexibility, isolates workload processing, and improves the potential for maximum server utilization.

## Oracle Solaris Zones

In a consolidated environment, it is sometimes necessary to maintain the ability to manage each application independently. Some applications may have strict security requirements or may not coexist well with other applications, so organizations need the ability to control IT resource utilization, isolate applications from each other, and efficiently manage multiple applications on the same server.

Oracle Solaris Zones technology (formerly called Oracle Solaris Containers), available on all servers running Oracle Solaris, is a software-based approach that provides virtualization of compute resources by enabling the creation of multiple secure, fault-isolated partitions (or zones) within a single Oracle Solaris OS instance. Running multiple zones makes it possible for many different applications to coexist in a single OS instance.

Oracle Solaris Zones was first introduced in Oracle Solaris 9 to provide full resource containment and control for more-predictable service levels. The Oracle Solaris Zones environment also includes enhanced resource usage accounting. This highly granular and extensive resource tracking capability can support the advanced client billing models required in some consolidation environments.

## Oracle Solaris Resource Manager

Oracle Solaris Resource Manager uses resource pools to control system resources. Each resource pool can contain a collection of resources, known as resource sets, which can include processors, physical memory, or swap space.

A processor set generally is assigned a subset of the total number of processors available on the system or the partition. For example, on a 16-processor system, an administrator could define two resource pools, each containing a processor set, with pool A's processor set containing 11 processors and pool B's processor set containing the remaining 5 processors. Resources can be dynamically moved between resource pools as needed.

## Fair Share Scheduler

Oracle Solaris Resource Manager incorporates an enhanced fair share scheduler that can be used within a resource pool. When using the fair share scheduler, an administrator assigns processor shares to a workload that may comprise one or more processes.

These shares enable the administrator to specify the relative importance of two workloads, and the fair share scheduler translates that into the ratio of processor resources reserved for each workload. If the workload does not request processor resources, those resources can be used by other workloads. The assignment of shares to a workload effectively establishes a minimum reservation of processor resources, guaranteeing that critical applications will get their required server resources.

## Oracle Enterprise Manager Ops Center

One of the key goals of server consolidation is to simplify server management by reducing the number of servers and OS instances that need to be managed. Oracle Enterprise Manager Ops Center 12c achieves this by merging the management of systems infrastructure assets into a unified management console.

Through its advanced server lifecycle management capabilities, Oracle Enterprise Manager Ops Center 12c provides a converged hardware management approach that integrates the management of servers, storage, and network fabrics, including firmware, operating systems, and virtual machines. Oracle Enterprise Manager Ops Center 12c provides asset discovery, asset provisioning, monitoring, patching, and automated workflows. It can also discover and manage virtual servers as well as physical servers, simplifying the management of high-end servers such as the Fujitsu M10-4S as well as all other SPARC servers in a data center. Oracle Enterprise Manager Ops Center 12c is available free of charge to all Oracle server customers with Oracle Premier Support contracts.

## Fujitsu M10-4S High-Availability Features for Consolidation

Servers used for consolidation must have very high availability, given that a consolidation platform failure can affect many applications. These servers must be highly reliable to minimize unplanned downtime, and they also need a high level of serviceability to minimize planned downtime.

The Fujitsu M10-4S server has many sophisticated reliability, availability, and serviceability (RAS) features and technologies that deliver extremely high levels of availability. It is designed for “five nines” availability, which means that for many enterprises, a single Fujitsu M10-4S server has sufficient availability to make clustering unnecessary, greatly reducing costs. Here are some examples of the RAS features available in Oracle’s high-end SPARC Enterprise M-Series servers:

- **Full hardware redundancy.** The Fujitsu M10-4S has two or more of each component, so the servers can automatically recover from any hardware failure without requiring service or a component replacement. Hardware redundancy includes redundancy of I/O paths and internal interconnects. I/O cards, disks, power supplies, fans, and building blocks with processors and

memory are all hot-pluggable, meaning that servicing, adding, or deleting any of those components doesn't require any OS or application downtime.

- **PPARs.** Hardware and software errors are limited to only the PPAR in which they occur. Downtime due to repair or reconfiguration is limited to the PPAR being modified or repaired.
- **Predictive self-healing.** Oracle Solaris predictive self-healing further enhances the reliability of SPARC servers. The implementation of Oracle Solaris predictive self-healing for Fujitsu M10 servers provides constant monitoring of CPUs and memory. Depending on the nature of the error, resolving persistent CPU soft errors is a matter of automatically off-lining either a thread, a core, or an entire CPU. In addition, the memory page retirement function supports the ability to take memory pages offline proactively in response to multiple corrections to data access for a specific memory DIMM.

## Case Study

The following customer example highlights how the Fujitsu M10-4S “dynamic scaling” model is ideal for consolidating existing workloads. A recent customer had several older SPARC servers it wanted to consolidate onto newer hardware to take advantage of faster processors, faster I/O, and improved availability and manageability features. The goal was to reduce management costs and improve resource utilization as well as increase applications capacity to meet the growth in the business. The customer took advantage of the modular architecture of the Fujitsu M10-4S, its CPU Activation feature, Oracle VM Server for SPARC, and Oracle Solaris Legacy Containers.

The customer deployed a Fujitsu M10-4S with four processors and 64 cores, but most of the cores were not licensed. The plan was to license processor cores as applications were added. By using CPU Activation, the customer did not have to pay for processors until its older applications were migrated to the newer system. It started out with five applications on five different servers. Some of the applications could run on Oracle Solaris 11 and Oracle Solaris 10, but a few required Oracle Solaris 8 or Oracle Solaris 9 environments. The customer created five Oracle VM Server for SPARC partitions that were to be the destinations for the respective applications. As each application was deployed onto the new Fujitsu M10-4S server, processor cores were activated. CPU Activation's core granularity makes it easy to deploy just the right amount of processor resources as applications are migrated to the new server.

Oracle Solaris 8 Legacy Containers and Oracle Solaris 9 Legacy Containers were created to deploy the applications that required Oracle Solaris 8 and Oracle Solaris 9. As applications need more capacity, more processor cores can be activated. If an application needs less capacity, activated cores can be reassigned to other applications or new applications as more legacy servers are replaced. If the current Fujitsu M10-4S server runs out of capacity, another Fujitsu M10-4S server can be added to increase total system capacity. The new Fujitsu M10-4S server does not require a new OS instance, but the resources can be applied to any of the existing OS instances, because the Fujitsu M10-4S is a shared memory architecture and not a distributed memory architecture. Figure 6 shows how the customer has migrated its legacy servers to a new Fujitsu M10-4S server.

The customer's successful migration of several legacy servers to a new Fujitsu M10-4S server has reduced its management complexity and increased server utilization. It can also gain a lot of room for future growth by activating new processor cores and/or adding new Fujitsu M10-4S servers to its existing Fujitsu M10-4S server.

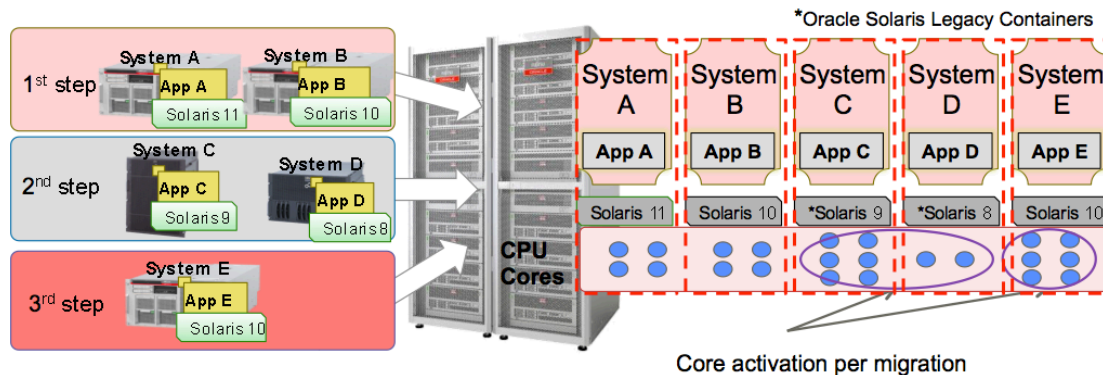


Figure 6. Legacy servers with different OS versions can be migrated to a single Fujitsu M10-4S server.

## Conclusion

The Fujitsu M10-4S is an ideal consolidation platform. It is the next generation of large, vertically scalable Fujitsu SMP servers that for more than 10 years have delivered scalability, capacity, high availability, virtualization, and system management capabilities for supporting dozens of diverse workloads of varying sizes.

With all resources connected by a high-speed interconnect, application deployment is far simpler and faster than on other server and deployment architectures. Virtualization technologies such as PPARs, Oracle VM Server for SPARC, and Oracle Solaris Zones make it easier to consolidate many applications on the Fujitsu M10-4S, with the goal of increasing server utilization and business flexibility.

Advanced RAS features, including hardware redundancy and hot-swappable components, dramatically reduce both planned and unplanned downtime. IT managers looking for a platform for consolidating mission-critical applications will find that the Fujitsu M10-4S server is an ideal solution.

## Further Information

For more information on the Fujitsu M10-4S server, please visit [oracle.com/goto/fujitsu-m10-4s](http://oracle.com/goto/fujitsu-m10-4s)



Consolidation Using the Fujitsu M10-4S Server  
April 2014

Oracle Corporation  
World Headquarters  
500 Oracle Parkway  
Redwood Shores, CA 94065  
U.S.A.

Worldwide Inquiries:  
Phone: +1.650.506.7000  
Fax: +1.650.506.7200

[oracle.com](http://oracle.com)



Oracle is committed to developing practices and products that help protect the environment

Copyright © 2014, Oracle and/or its affiliates. All rights reserved.

This document is provided for information purposes only, and the contents hereof are subject to change without notice. This document is not warranted to be error-free, nor subject to any other warranties or conditions, whether expressed orally or implied in law, including implied warranties and conditions of merchantability or fitness for a particular purpose. We specifically disclaim any liability with respect to this document, and no contractual obligations are formed either directly or indirectly by this document. This document may not be reproduced or transmitted in any form or by any means, electronic or mechanical, for any purpose, without our prior written permission.

Oracle, , Oracle Solaris, Oracle VM Server for SPARC, Oracle Enterprise Manager 12c, and Oracle Solaris Zones Oracle Solaris, and Oracle Legacy Containers are registered trademarks of Oracle and/or its affiliates. Fujitsu and Fujitsu M10-4S are registered trademarks of Fujitsu and/or its affiliates . Other names may be trademarks of their respective owners.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. UNIX is a registered trademark of The Open Group. 0413

**Hardware and Software, Engineered to Work Together**