

Consolidation Through Virtualization

With Sun™ x64 servers

IT organizations must perform a nearly impossible juggling act. Deliver more services at less cost. Increase performance yet manage fewer systems. Buy the right server for the job yet adapt to change at a moment's notice. Sun Fire™ x64 servers and Sun Blade™ modular systems powered by AMD Opteron™ and Intel® Xeon® processors let businesses run the Solaris™ Operating System (OS), Linux or Windows today, and change direction tomorrow. But consolidate multiple applications and operating systems onto the same server? It's straightforward with virtualization technologies from Sun and its partners (Figure 1).

Doing more with less

Almost every enterprise depends on the IT organization to support the applications that make it run. Yet today, most businesses are cutting costs in order to deliver products and services more cost-effectively and efficiently than the competition. IT organizations are now pressured to do the same, and must cut cost and complexity while delivering highly available, business-critical applications that perform at lightning speed.

Highlights

- Consolidate multiple applications onto fewer servers and reduce cost and complexity, increase agility, and lower datacenter power and cooling costs
- Utilize Sun x64 servers and run the Solaris™ Operating System, Linux or Microsoft Windows, and gain the flexibility to re-deploy the moment needs change
- Consolidate multiple Solaris OS and Linux applications onto a single server using Solaris Containers technology
- Consolidate any combination of Solaris 10 OS and supported Linux and Microsoft Windows applications onto a single server using Sun xVM hypervisor or VMware Virtual Infrastructure technology

Virtualization and consolidation are two key tools that help IT organizations do more with less.

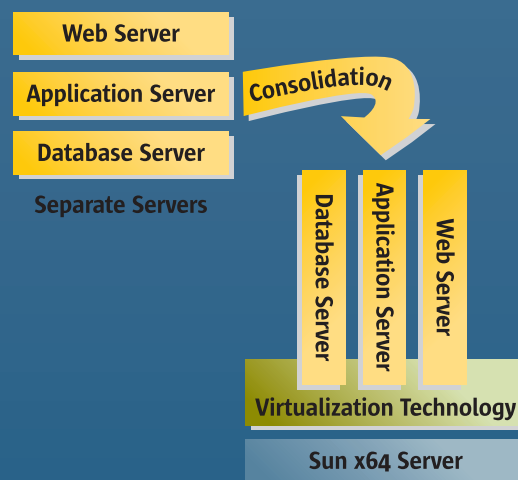


Figure 1: Consolidation through virtualization enables different applications to run in secure, isolated environments on a single server platform—such as Sun x64 servers.

These demands translate into a unique set of challenges:

- Reduce capital expenses by using a smaller number of larger, more powerful and cost-effective servers
- Provide for continued operations by distributing applications geographically
- Reduce operating expenses by managing a smaller number of servers and by supporting a small, core set of enterprise operating systems
- Increase server utilization and reduce power and cooling costs by deploying multiple applications onto a single server, sharing resources, and reducing the number of wasted CPU cycles
- Optimize flexibility and protect investments by choosing servers that can be used to support one application and operating system today and support a different combination tomorrow
- Raise availability levels with local replication and reliable servers equipped with redundant components

Virtualization and consolidation

In the 1990s, a key strategy in designing Internet architectures was to decompose applications into separate components. Each component was replicated for availability, hosted in its own security domain (usually a dedicated server), and tuned for optimal performance. This resulted in a large number of 1RU and larger servers deployed across datacenter environments, each contributing to inefficiency.

- Each server was sized to handle the maximum expected workload, leaving its CPU (and other resources) under-utilized most of the time
- Each 1U server needed infrastructure, including power supplies and cooling fans, resulting in higher initial capital costs and ongoing operational costs than with larger servers offering better economies of scale
- Each server contributed to the number of systems to maintain, licenses to track, and operating system instances to support

Virtualization and consolidation techniques have evolved since the 1990s to the point where today they can be used together to support Internet architectures that are logically decomposed but which physically share the same infrastructure. This helps IT organizations achieve the same security, availability, and performance benefits of decomposed architectures while realizing a whole new level of efficiency. Virtualization and consolidation are two key tools that help IT organizations do more with less.

Consolidation

Consolidation is the strategy of moving multiple applications from separate servers onto fewer shared servers. This technique is used by IT organizations wishing to leverage the efficiency and cost-effectiveness of larger, more powerful servers.

Virtualization

Consolidation is the goal, but virtualization and partitioning are the means by which multiple applications or application instances can share a platform and resources without interfering with each other. When each application instance runs in its own isolated environment, each one can have its own unique configuration and data files, and not interfere with the others.

**Consolidation is the goal.
Virtualization is the means.
The end result is lower cost,
greater agility, and less
complexity.**

Virtualization enables applications to access the resources they are authorized to use — and not exceed the boundaries of their security domains. In addition, virtualization lets IT organizations manage resources such as CPU, memory, and network bandwidth for greater utilization, and dynamically adjust resource allocation. This helps give IT organizations the flexibility they need to respond quickly to rapidly changing workloads.

Virtualization begins with a single environment and creates the illusion of multiple ones. It can take place at several levels, but regardless of level the effect is that the application or the guest operating system itself has the illusion that it ‘owns’ its environment. Several approaches are used in products from Sun and its affiliates today, two of which are of particular importance in the x64 server market.

- Virtual machine monitors, or hypervisors, create the illusion that each operating system has its own dedicated hardware—despite the fact that each operating system only ‘owns’ a part of the hardware platform.

Virtual machine monitors include Sun xVM hypervisor, Xen, Microsoft Hyper-V, and VMware Infrastructure, in which ESX Server is the key component. Sun xVM Server and VMware ESX Server are highlighted in this brief.

- Containers partition a single operating system instance to give each application the illusion that it has its own environment and its own dedicated set of resources. Solaris Containers is the technology highlighted in this brief.

Unprecedented opportunity with Sun x64 servers

The Sun x64 server product line makes Sun a “one-stop shop” for IT organizations needing to support multiple operating systems on the same server, and also for those wishing to consolidate multiple applications onto a smaller number of servers.

Sun x64 servers

Sun x64 servers run existing 32- and 64-bit operating systems and applications with blinding speed. Built around AMD Opteron and Intel Xeon processors, Sun x64 servers and blade systems are built to deliver the performance and memory capacity that is necessary when consolidating many existing workloads on to a single server.

With servers built using processors from both AMD and Intel, customers can choose the architecture that best suits their needs. Whether it’s the processor they’re most familiar with, or whether it’s one that excels in running their specific application, Sun builds servers with virtualization and consolidation in mind.

The Sun Blade Modular System product line is a dense consolidation platform with a wealth of modular I/O options. Both the Sun Blade 6000 and 8000 modular systems support up to 10 Sun Blade server modules.

Server modules for the Sun Blade 6000 and Sun Blade 8000 modular systems are available with up to four sockets equipped with quad or 6-core processors and up to 128 GB of main memory.

One-stop shop

The ability to run the Solaris OS, Linux, or Microsoft Windows on Sun x64 servers gives IT organizations the flexibility to use one vendor to meet a wide variety of requirements. One set of servers and storage can be purchased and deployed for one purpose today, and redeployed with a different operating system the moment needs change.

Virtualization enables all three operating systems to run on the same platform at the same time. Applications can run at full processor speed, without the need for time-consuming processor set emulation. Three operating system choices give organizations tremendous flexibility and investment protection.

Lights out management

All Sun x64 servers include integrated lights out management through Ethernet interfaces that typically connect to a datacenter management network. Integrated Lights Out Management (ILOM) enables Sun x64 servers to be logically extended to provide access to peripherals such as CD, DVD, and floppy drives on remote systems, making it easy to load software into virtualized systems without entering the datacenter.

AMD Opteron™ processor-powered servers

Sun's AMD Opteron processor-based server product line includes both rack-mount servers and Sun Blade server modules that offer a range of options for memory capacity, I/O expansion, scalability, and server density. Server products range a single dual-core AMD Opteron processor and up to 8 GB of main memory to 8 sockets capable of hosting single, dual, or quad-core processors and up to 256 GB of main memory.

Sun Fire™ X4600 M2 server

One of Sun's AMD Opteron processor-powered servers that offers particular value to virtualization customers is the Sun Fire X4600 M2 server. Sun builds this server for organizations that need more raw compute power than typical four-socket servers. Combined with the scalability of VMware Infrastructure 3 software, it demonstrates near-linear scalability on VMware's virtualization benchmark (see the technical brief "Proven Virtualization Scalability with the Sun Fire X4600 Server.")




Figure 2: The Sun Fire X4600 (top) and the Sun Fire X4450 server (bottom)

In a mere four rack units, the Sun Fire X4600 M2 server can pack up to eight dual- or quad-core processors for a total of 32 cores and up to 256 GB of memory (Figure 2).

In addition to its processing capability, the server can support up to four hot-pluggable SAS disk drives with on-board RAID 0 and RAID 1, and four Gigabit Ethernet ports. The server has four low-profile 8-lane PCI Express card slots, two low-profile 4-lane PCI Express slots, and two low-profile PCI-X slots.

Intel® Xeon® processor-powered servers

Sun's server product line based on Intel Xeon processors delivers 4 and 6 processing cores per socket in both rack-mount servers and server modules. Rack-mount servers range from a 1U, 2-socket to a 2U, 4-socket server with up to 24 CPU cores and 128 GB of main memory.

Sun Fire X4450 server

With up to twice the CPU and memory density of competitive platforms, the Sun Fire X4450 server supports four CPU sockets that can be populated with dual-quad or 6-core Intel Xeon processors. With up to 24 cores in only two rack units, the server can support up to 128 GB of main memory, making it an excellent consolidation platform.

In addition to its processor and memory strengths, the server supports SAS disk drives with on-board RAID. The server can support up to eight hot-pluggable SAS drives with RAID levels 0,1, 1E, 5, 5EE, 6, 10, 50, and 60 based on customer choice of optional controllers. The server has six 8-lane PCI Express card slots.

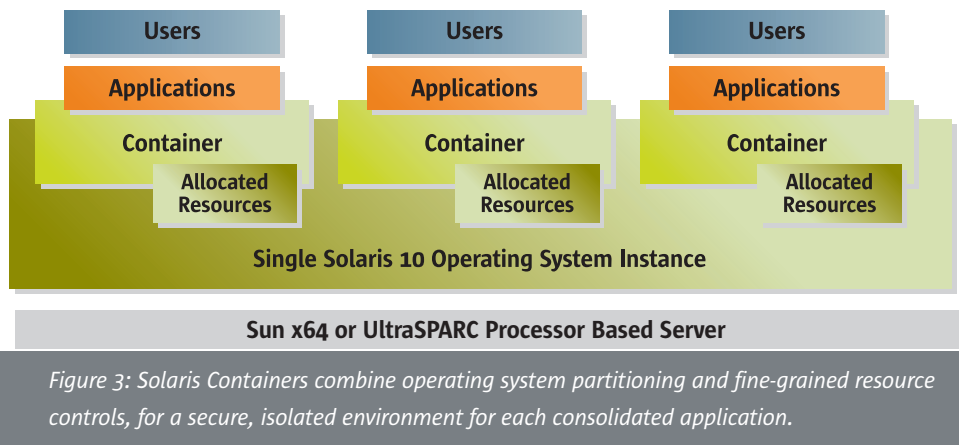


Figure 3: Solaris Containers combine operating system partitioning and fine-grained resource controls, for a secure, isolated environment for each consolidated application.

Resource partitioning with Solaris™ Containers

Solaris Containers combine operating system partitioning with fine-grained resource controls to enable servers to be partitioned at sub-CPU granularity without replicating the operating system image (Figure 3).

Solaris Containers provide a virtualized Solaris 10 OS image, including a unique root file system, a shared, read-only set of system executables and libraries, and whatever resources the root administrator assigns to the container at creation time. Branded Zones (BrandZ) allows different versions of the Solaris OS and Linux to run within a container, supporting applications written for the OS defined by the zone. Solaris Containers can be booted and shut down just like any instance of the Solaris 10 OS, and rebooted in seconds if the need arises.

Unlike virtual machines, which must intercept every single interrupt and allocate it to the right OS instance, Solaris Containers support mainframe-level partitioning capabilities with almost zero overhead.

Operating system partitioning

Solaris Containers provide a set of up to 8,191 virtualized environments per Solaris 10 OS instance, and each container appears to users, administrators and applications as an isolated, independent system. Global administrators can create, allocate resources to, and boot containers as if they were an operating system instance.

Once booted, Solaris Containers provide a secure environment that includes:

- A virtual platform containing a unique root, shared user, and other administrator-configured file systems —

plus network interfaces, inter-process communication objects, console devices, and sub-container resource management facilities.

- System identity settings, including host name, time zone, RPC domain, and locale
- Secure isolation from other containers enforced by the kernel and capable of preventing a process within a container, even if compromised, from escalating privileges to compromise another container
- Fault isolation that restricts the propagation of software faults to a single container. If an error causes a container to fail, it can reboot in only a few seconds because the single operating system instance running on the server remains intact.

A Solaris Container hosting a Web server might be allocated an IP address with rights to bind to port 80, and a disk device containing a file system for the Web site content. The Web server cannot see or access any resources not allocated to its container. If the Web server fails, or its security is penetrated by an intruder, it cannot affect other containers or the applications running in them.

Fine-grained resource control

Solaris Resource Manager software gives administrators almost unlimited flexibility to assign and isolate resources to specific containers. Resource Manager can be used to allocate resources to multiple applications within a single container. Across multiple containers, Dynamic Resource Pooling enables administrators to allocate discrete pools of resources such as CPUs to specific containers. Administrators can dynamically change the content of resource pools manually or automatically on a rule basis. For example, an additional CPU can be added automatically to a container when its utilization exceeds 80 percent—without rebooting.

The Fair Share Scheduler supports dynamic resource allocation, enabling proportions of resources—such as fractions of a CPU—to be allocated to containers. When resources like CPUs and memory are dynamically allocated, resource-capping controls can be used to set limits on the amount of resources consumed by specific containers. Finally, Solaris IP Quality-of-Service (QoS) can be used to manage network bandwidth used by multiple containers, helping administrators to maintain specified QoS levels in a consolidated environment.

Virtualizing the hardware with Sun xVM hypervisor

In contrast to Solaris Containers, which partition the operating system, Sun xVM

hypervisor is a virtual machine monitor that supports running multiple operating system instances on the same server.

This approach helps Sun xVM hypervisor to support a different class of consolidation requirements than Solaris Containers, namely support for applications requiring kernel-level isolation and running multiple OS instances on the same server.

Sun xVM

By virtualizing the hardware, a single system running Sun xVM hypervisor can support multiple guest operating systems, including the Solaris OS, Linux, and Microsoft Windows XP, 2003, and Vista operating systems.

Based on the Solaris OS

Sun xVM hypervisor is based on the Solaris OS, and is currently available through the OpenSolaris™ program. The thin hypervisor software layer supports multiple guest operating systems on Sun x64 servers, each securely isolated from one another.

One instance of the Solaris OS runs as the control domain. It creates new domains, allocates resources to them, and manages CPU and memory allocation. In operation, it mediates I/O requests between guest operating systems and the hardware. This allows the multiple OS instances to securely share server resources with minimal overhead. Sun xVM hypervisor combines two technologies, virtualization and paravirtualization.

Broad guest OS support with virtualization

Sun xVM hypervisor provides a layer of abstraction between the server hardware and unmodified operating system instances that run on it. With this model, each guest operating system is given the illusion that it owns the server itself, down to the ability to execute privileged instructions. This requires the hypervisor to intervene when the guest OS issues privileged instructions.

Virtualization helps Sun xVM hypervisor to support a variety of operating systems, including Microsoft Windows XP, 2003, and Vista, in such a way that the operating systems can run unmodified.

Performance through paravirtualization

Paravirtualization incorporates knowledge of the hypervisor in the guest operating system so that it can make requests of the virtualization layer directly rather than relying on traps and instruction re-writing as is the case in with full virtualization. Operating systems that can run in the more efficient paravirtualized mode include the Solaris OS, and Linux.

As Figure 4 illustrates, Sun xVM hypervisor software is layered on top of a Sun x64 server. The hypervisor supports both unmodified and paravirtualized operating system instances. An instance of the Solaris OS, known as the control domain, or ‘dom0,’ handles I/O for virtualized environments. Sun xVM hypervisor software is currently available with the OpenSolaris Operating System.

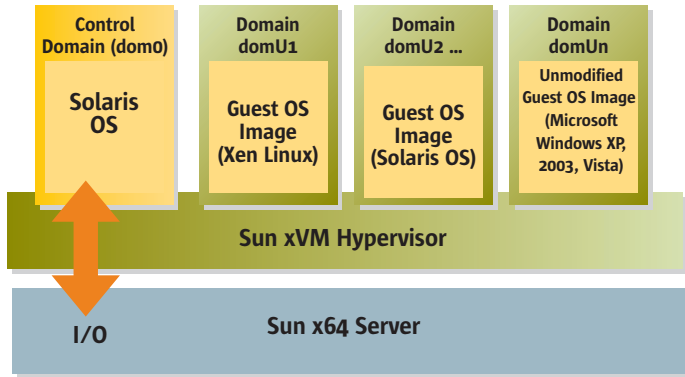


Figure 4: Sun xVM hypervisor supports both unmodified and paravirtualized operating system instances, providing an optimal mix of flexibility and performance.

Managing virtualized environments with Sun xVM Ops Center

Sun xVM Ops Center is designed to help customers mitigate the complexity inherent in virtualization solutions by addressing both the virtualization and management of heterogeneous assets.

The first product in the Sun xVM family, Sun xVM Ops Center, is a comprehensive tool for managing thousands of Linux and Solaris OS-based multi-vendor x64/86 and SPARC systems.

Sun xVM Ops Center simplifies the discovery, provisioning, updating, monitoring, and reporting of both the virtual and physical assets through one unified interface. It consolidates a number of key management capabilities:

- Advanced software update management
- Compliance reporting
- Hardware and software discovery

- Provisioning of server firmware, lights-out bare metal operating systems, and applications
- System and operating system monitoring
- Threshold settings and alerts
- Scalable for global IT environments with thousands of assets to manage

With Sun xVM hypervisor and Sun xVM Ops Center, customers have a comprehensive virtualization and management tool to support their operations across the datacenter.

Virtualizing the hardware with VMware Infrastructure

No two companies provide more complete, mature, and proven virtualization solutions than Sun and its strategic partner, VMware.

Sun and VMware

Thousands of enterprises worldwide use VMware Infrastructure software to maximize the manageability, flexibility, and efficiency of their datacenter infrastructures. And Sun's vision of a virtual datacenter is redefining how IT services are designed, deployed, and managed in mission-critical environments. Together, Sun and VMware offer a virtual datacenter solution that works in the real world.

Sun is an OEM reseller of key VMware products and technologies, making it easier than ever for customers to consolidate through virtualization. VMware on Sun has also yielded outstanding consolidation results for customers.

VMware Infrastructure, comprised of VMware ESX Server and VirtualCenter, is another virtual machine technology. VMware ESX Server provides a virtual machine platform with resource management capabilities, all managed through VMware VirtualCenter.

As with Sun xVM hypervisor, VMware ESX Server creates a layer of abstraction between the server and the guest OS (Figure 5).




Figure 5: Each virtual machine supported by VMware Infrastructure software supports an idealized hardware environment, including CPUs, memory, disk, and even BIOS.

Idealized hardware platform

VMware ESX Server runs directly on supported Sun x64 servers to provide a secure, uniform platform for deploying, managing, and remotely controlling multiple operating system instances.

Hardware interfaces, such as device drivers, enable hardware-specific service delivery while hiding hardware differences from other parts of the system. The virtualization layer provides an idealized hardware environment and virtualizes underlying physical resources.

By presenting a platform-independent x86 architecture hardware environment that is appropriate for each guest operating system, ESX Server makes it easy to move virtual environments from machine to machine without having to exactly match CPUs, disk drivers, and network interfaces. Platform-independent virtual environments also make it possible to support operating systems and applications no longer supported by hardware vendors.

For example, ESX Server can be configured to support Microsoft Windows NT guest operating systems, enabling enterprises to migrate existing applications from obsolete hardware platforms to state-of-the-art Sun x64 servers.

Granular resource management

The resource manager in ESX Server uses a proportional share mechanism to allocate CPU, memory, and disk resources across multiple virtual machines. Network bandwidth is controlled with network traffic shaping. Minimum and maximum percentages of a single physical CPU's processing power can be specified for each virtual machine.

Hyper-V

Sun's Hyper-V collaboration continues Sun's commitment to help maximize flexibility and deliver value for its SMB and enterprise customers. Sun x64 customers can now run Hyper-V with Windows Server 2008. This next-generation hypervisor-based server virtualization technology, allows customers to make the

best use of server hardware investments by consolidating multiple server roles as separate virtual machines running on a single physical machine.

With Hyper-V, customers can also efficiently run multiple different operating systems—Windows, Linux, and others—in parallel, on a single server, and fully leverage the power of x64 computing.

Consolidation through virtualization

Virtualization with Sun xVM hypervisor, VMware Infrastructure, Microsoft Hyper-V, and partitioning with Solaris Containers give Sun customers a flexible set of choices for consolidating multiple applications onto a single server. Putting these technologies to work in an IT environment opens up a whole new realm of possibilities.

Complimentary workloads

Organizations can deploy a set of applications onto servers based on complimentary resource requirements. On-Line Transaction Processing (OLTP) systems, for example, might occupy most resources in a server during the daytime while customers are placing orders. At night, long-running data warehousing queries can absorb unused CPU cycles, helping to increase overall IT resource utilization and efficiency.

Complementary Web applications can be deployed in similar manner. For example, a news site can be deployed in secure virtualized environments across a number of servers dedicated to different purposes.

When a major news story results in an astronomical increase in activity, the standby Web servers can be pressed into service simply by adding them to a load-balanced service group. IT organizations can respond to spikes in workload without the delay of wheeling in new servers and loading operating system and application software.

Flattening architectural layers

A high-availability, three-tier Web application requires at least six servers, two at each tier. Using virtualization technologies to flatten the architecture, only two servers are needed for availability. Both servers host Web, application and database server software, each in a secure, virtualized environment. Network traffic can be secured outside containers by allocating a dedicated network interface to each environment, or by sharing interfaces and implementing separation through VLANs.

IT organizations can flatten layers in a number of contexts where related applications require multiple tiers today. For example, a Web application firewall and a Web server can be deployed in separate environments on the same server. A mail server, virus scanner, and spam filtering suite can be installed into separate environments so that if an intruder is successful in breaking a component's security, the intruder no longer has access to all three.

Development, staging, and production

Virtualization and partitioning technologies let IT organizations deploy development, test, staging, and production versions of an application onto the same server.

Developers can work in a personal environment. Once an application has been tested and readied for deployment, it can be installed in a virtualized environment for staging. When deploying the new application, simply change IP address and resource allocation—and the staging environment becomes the production one.

This approach reduces the number of resources that an IT organization must support, and it provides an automatic fallback mechanism if a new version must be rolled back to the previous one.

Legacy application consolidation

Sometimes a set of applications running on multiple operating system versions needs a performance boost that can be accomplished by re-hosting onto a faster server. At other times, applications run on no-longer-supported server or operating system versions. Using virtual machine technology, these applications can be consolidated onto high-performance x64 servers and receive the needed platform upgrade and performance boost. Host each operating system and application instance in its own separate virtual machine, assign an x64 processor to each VM, and watch performance soar.

Geographic replication, disaster recovery

Deploying applications into containers or virtual machines makes it easy to package and redeploy them anywhere around the globe for both geographic replication and disaster recovery purposes.

Solaris Containers, Sun xVM hypervisor, and VMware Infrastructure facilitate backing up and restoring virtual environments.

In each case, two sets of data are stored: an archive of the environment's data, and configuration settings for the container or virtual machine environment.

Shared hosting environments

Internet Service Providers can offer subscribers a unique, complete Web hosting environment, giving them control over Web server software and even administrator passwords. Using Solaris Containers, Sun xVM hypervisor, or VMware Infrastructure, ISPs can offer virtual dedicated hosting with the efficiency of a shared hosting infrastructure.

Consolidation guidelines

At a high level, the choices for an IT organization wishing to consolidate multiple applications onto a single powerful Sun x64 server are straightforward:

- To consolidate multiple applications running on the Solaris 10 OS, use Solaris Containers.
- To consolidate multiple applications running on other versions of the Solaris OS and/or Linux, use Solaris Containers and Branded Zones
- To consolidate a mix of applications running on the Solaris 10 OS, Linux or Microsoft Windows, use Sun xVM hypervisor or VMware Infrastructure

Consolidating using Solaris Containers

For IT organizations wishing to consolidate multiple applications running on the Solaris OS, multiple open source applications, Linux applications, or any combination of the above, Solaris Containers is the consolidation technology of choice.

Solaris Containers are a low overhead partitioning approach, included with the Solaris 10 OS at no extra cost. With feature parity across platforms, Solaris Containers can be used on Sun servers regardless of the underlying processor architecture.

Consolidating using Sun xVM hypervisor and VMware Infrastructure

When an IT organization wishes to consolidate applications running on more than just the Solaris OS and Linux, a virtual machine monitor is the consolidation option of choice. Sun xVM hypervisor is available at no additional cost through the OpenSolaris program, and VMware Infrastructure is available from VMware.

Both of these options can support multiple versions of the Solaris OS, Linux, and Microsoft Windows. Solaris xVM hypervisor may be the best choice for some applications because paravirtualization should provide greater performance than true virtualization. On the other hand, if legacy systems that run on obsolete hardware need to be

consolidated, VMware Infrastructure can emulate old processors and peripherals to support them.

IT organizations consolidating onto Sun x64 servers running VMware Infrastructure enjoy the benefit of migration software that helps to package up an entire environment so that it can be installed in its own virtual machine. The Sun BluePrints™ article Consolidating Microsoft Windows NT Applications onto Sun x64 Servers using VMware ESX Server includes instructions for using VMware Physical-to-Virtual (P2V) software to perform just such a migration.

Sun — the ideal consolidation partner

For IT organizations attempting a nearly impossible juggle of competing priorities, consolidating multiple applications onto a smaller number of more powerful servers is one technique that helps to reduce capital and operational costs, increase utilization, raise availability levels, and provide for continued operation through geographic replication.

When IT organizations look for the right platform to support operations, there is no better partner than Sun. With a long history of pushing mainframe quality features into its mid-range and entry-level servers, Solaris Containers and Sun xVM hypervisor are only two of many technologies that Sun shares with all of its customers, regardless of how many or how large a server is purchased.

And because Sun develops new versions of the Solaris OS at the same time it develops new industry-leading server platforms, customers have early access to the newest server technology when they choose the Solaris OS.

When it comes to platform choice, Sun Fire servers, Sun Blade server modules, and Sun Blade modular systems offer the power of single to 6-core processors with flexibility and investment protection. Enterprises wishing to deploy a dedicated operating system can choose between the Solaris OS, Linux, and Microsoft Windows today, and re-deploy the same server with a different operating choice the moment needs change.

Those needing to consolidate Solaris OS and Linux applications have the flexibility to use any of three technologies from Sun and its affiliates: Solaris Containers, Sun xVM hypervisor, or VMware Infrastructure. When the range of operating systems to be supported includes Microsoft Windows operating systems, or Linux versions not supported by Branded Zones, customers can choose between Sun xVM hypervisor and VMware Infrastructure 3 software. Whichever choice an IT organization makes for consolidation through virtualization, the best choice is using Sun servers.



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