DATACENTER REFERENCE GUIDE

White Paper
July 2007

Keith Black
Jason Carolan
Gary Combs
Brian Couling
Ron Graham
Laura Hartman
Eugene Ho
Patrick Hudelot
Vernon Le Sueur
Mikael Lofstrand
Steve Nelson
Tim Read
Jeff Savit
Eran Steiner

Contents

Chapter 1

Introduction

Virtually every business in operation today understands that its survival depends on the quality and capability of its IT organization. This understanding spans the spectrum from start-up companies whose whole business is conducted on the Internet, to established companies that use their datacenters to run their internal operations. Indeed, most organizations are both inwardly and outwardly focused, and must support both a strong and secure network presence and support a range of internal applications:

- Companies that conduct their business on the Internet depend on their information technology to support every single business transaction — whether they use the Internet to sell products directly to customers, conduct transactions with business partners, or provide a range of innovative, collaborative, and participatory services to customers. For some externally focused companies, the user is the programmer, and the growth of their customer base and their network strains their infrastructure.
- Companies whose IT organizations are more focused on internal activities depend on their IT organizations to provide services that boost efficiencies to the point where they can operate on razor-thin margins — whether they use their information technology to operate their supply chain with near zero inventory and just-in-time manufacturing, or run complex high-performance computing (HPC) applications to crack genetic codes, simulate automobile crashes, or model oil reservoirs.

Today's competitive marketplace demands more and more of their IT organizations. IT organizations must be able to implement and roll out new applications — internal and external — at an unprecedented rate. The very business model of delivering applications is also changing. Rather than building monolithic applications that are deployed with much fanfare, many organizations are adopting a model that uses what is essentially natural selection. They roll out a larger number of smaller applications, let the marketplace decide which of them are used by customers, and then polish, enhance, scale, and support the ones that prove themselves.

A Formidable Set of Challenges

The dynamic nature of the marketplace and the demand for IT organizations to be flexible, agile, and more productive than ever, poses a difficult set of challenges to IT organizations.

The Space, Power, and Cooling Crunch

IT organizations large and small are up against hard constraints in their datacenters. They are running out of space and are up against their power and cooling limits. They have realized that, in some electricity markets, it already costs more to power a server that it does to purchase it in the first place, adding economic weight to the argument for purchasing more energy-efficient, eco-friendly servers, storage, and network infrastructure components. In some areas, the cost of inefficiency is so high that IT organizations require vendors to include carbon footprint data along with their proposals.

Most existing IT organizations are saddled with legacy applications built with silo architectures. These applications are supported with server, storage, and network components purchased and assembled for a single specific purpose and cannot be shared with other applications. These systems were purchased to handle maximum expected peak workloads and are normally underutilized, resulting in costly inefficiency and compounding the space, power, and cooling crunch.

Focus on the -ilities

The growing importance of energy efficiency has for some companies obscured the importance of focusing on the underlying characteristics of their datacenters — the '-ilities," including scalability, reliability availability, serviceability, manageability, and security.

Security is one characteristic that is just as important as space, power, and cooling, because it affects a company's ability to business on the network, protect its business-critical applications and data, and conform to regulations and policies dictating who can see what personal and customer data. IT organizations understand that security isn't an add-on, it's a systemic feature that has to be integrated through all infrastructure and software layers.

Scalability is another characteristic that affects a company's ability to conduct business. When the success or failure of an application is left in the hands of customers, IT organizations must be ready to scale up an application's capacity at a moment's notice. Indeed, if a company's success depends on the serendipitous acceptance of a service offered to millions of Internet users, the ability to capitalize on that success depends on the application's scalability. A popular application whose performance lags generates only bad press, not increased revenues.

Accelerated Application Lifecycle

Rolling out new applications is more important than ever, and building applications using a flexible, re-usable set of hardware and software components is the only way that they can be successful. The traditional, silo-based application architectures that constrain many IT organizations cannot adapt to this new way of doing business.

Pressure to Cut Costs

Cost pressures are forcing IT organizations to find ways to increase administrator productivity in terms of servers managed per administrator, reduce the cost of rolling out new applications, and to squeeze more servers and storage into datacenters without undertaking expensive remodeling efforts.

Complying With Government Regulations

A new generation of government regulations dictate how customer and employee personal information is handled and retained. In some industries, including the financial industry, even application availability is regulated and disaster-recovery capabilities are no longer optional.

Sun Technology to Meet the Challenges

Information technology is no longer just a component of a business, in many cases it *is* the business — leading today's IT organizations to realize that their datacenters are critical to their strategy. Sun Microsystems has the long-held belief that "The Network is the Computer™," and has delivered every one of its server and workstation products with built-in networking. Over time, this belief has proven itself over and over, and nowhere is it more valid than in helping today's IT organizations architect, deploy, and manage their datacenters. Sun has maintained a systems-level view of information technology, providing solutions that encompass the complex interactions between a datacenter's key components, whether they are the servers, storage, and networks that provide the underlying power that drives the datacenter, or the network services, software, virtualization, management, and provisioning capabilities that can help turn a datacenter into a global pool of resources that can support new capabilities at a moment's notice. Sun understands the growing need for eco-responsibility, and it offers highly efficient server technologies that, when combined with Sun's approach to architecture, consolidation, and virtualization, can help IT organizations do more with less power.

Sun has a range of architectures, products, and services that can help IT organizations establish and maintain the flexibility and agility they need to adapt and excel in a world where business begins and ends in the datacenter, and some of the technologies available to Sun's datacenter customers include the following:

Dynamic Infrastructure and Service-Delivery Networks

Dynamic infrastructures and service-delivery networks allow IT organizations to re-define their datacenters as a flexible, scalable, adaptable pool of resources that can be leveraged to meet a rapidly-changing set of application and workload demands. Dynamic infrastructures allow, for example, new applications to be deployed using existing infrastructure, scaling their capacity as rapidly as workloads demand. These topics are discussed in the first three chapters:

- Chapter 2, “Datacenter Design Principles” on page 7
- Chapter 3, “Service Delivery Networks” on page 15
- Chapter 4, “Dynamic Infrastructure” on page 25

Virtualization and Consolidation

One technology that can help IT organizations create a global pool of resources, free themselves from the constraints of their silo-based architectures, and address their space, power, and cooling challenges is virtualization. Virtualization enables organizations to consolidate multiple workloads from existing systems onto a smaller number of more powerful servers, allowing workloads to be created and re-allocated across them in real time. Some virtualization techniques even allow legacy applications to be transitioned into a virtualized environment without re-installing operating system or application software, providing a pathway to retire obsolete, energy inefficient servers while still supporting the mission-critical software that runs on them. Virtualization presents new challenges, specifically managing servers, middleware, networks, and storage as a system. Sun’s management technologies can help IT organizations get the most flexibility out of their virtualized environments. Virtualization and Consolidation are discussed in:

- Chapter 5, “Server Virtualization Technology” on page 30

Matching Servers with Applications

Sun understands that every application is unique, and that supporting different applications with the required level of performance, scalability, availability, and reliability takes more than a server product line with one, two, or four processors per server. Sun offers a highly differentiated server product line with products that are designed to meet the requirements of different workloads and contribute to power efficiency in the datacenter. For applications such as database management systems that scale vertically, Sun offers servers that can scale up to 64 dual-core SPARC64® VI processors per server. For applications that scale horizontally, Sun offers UltraSPARC®, AMD Opteron™, and Intel® Xeon® processor-powered servers that give customers the choice of running Solaris™, Microsoft Windows, or Linux operating systems and, with virtualization technology, the ability to run all three at once. Sun offers high-density blade servers and systems designed for streaming video; it offers servers whose characteristics — 24 TB of storage in only four rack units — blur the distinction between server and storage. Sun’s server technology portfolio is discussed in:

- Chapter 6, “Sun Server Technology” on page 39

Storage — An Organization's Most Valuable Asset

IT organizations realize that their most precious asset is their data, and consolidating storage into a single, central pool allows them to manage each logical volume to best meet application and organizational requirements. Storage can be configured onto systems having the right capacity and performance characteristics for the job, replicated in remote locations for disaster recovery purposes, and snapshot and archived according to record-retention requirements. Sun StorageTek™ storage systems offer a capacity range of up to 336 TB per system, and up to 247 PB of additional storage using a crossbar switch architecture that offers market-leading performance of 3.5 million IOPS and a throughput of 106 GBps. StorageTek tape devices and libraries allow valuable data to be stored at rest, and moved offsite for the utmost in data security. Compared to disk storage, tape is inexpensive and requires no power when it is not being used. A range of Sun software products give IT organizations the tools they need in order to fully manage their data, the systems that store it, and plan for future storage requirements. Sun's storage products, and ways to use them to establish disaster-recovery strategies, are discussed in:

- Chapter 7, "Sun Storage Solutions" on page 48
- Chapter 8, "Disaster Recovery" on page 57

Sun's Systemic Security

Sun's systemic approach to security recognizes that security is not an add-on, it is a property that must be woven into the very fabric of the datacenter. Sun's systemic security approach uses architectural building blocks and patterns to build security into each step of the process of designing a datacenter or application. Sun's security practices uses a best-of-breed approach that combines technologies from Sun and its affiliates. Sun technologies range from Solaris OS security technologies to identity management software available as part of the Sun Java™ Enterprise System. Sun chooses its security software affiliates carefully, and depends on them for firewall, antivirus, and anti-spam technology. Sun's systemic security is discussed in:

- Chapter 9, "Systemic Security" on page 62

Management Solutions

IT organizations are under constant cost pressure, and Sun management solutions provide tools to help increase administrator efficiency and increase the number of systems that each administrator can provision and manage. Solutions such as the Sun™ Customer Ready HPC Cluster help organizations to deploy and manage servers a rack at a time rather than a server a a time. Management software from Sun's N1™ product line can help provision and manage systems from bare metal to application provisioning. The Sun℠ Connection service provides dynamic patching and "phone home" service tagging to help detect and repair potential failures before they result in down time. Sun's management solutions are discussed in:

- Chapter 10, "System and Server Management" on page 68

Getting Started with Sun Services

Sun offers the services to help you get started in the way most appropriate for your organization. Sun believes in tailoring its service offerings depending on an organization's needs and the capabilities of its staff, rather than forcing its customers into a one-size-fits-all model of how it does business. Sun's products, philosophies, architectures, and step-by-step instructions on how to best utilize its technologies are available in a range of solution briefs, white papers, and Sun BluePrints™ articles. Together, this white paper and these resources can help you to understand the depth of experience and range of products that are available for your datacenter when you do business with Sun. A brief overview of Sun datacenter services is provided in:

- Chapter 11, "Sun Datacenter Services" on page 76

Chapter 2

Datacenter Design Principles

Good architectures are critical to designing complex, network-oriented systems, including datacenters. Historically, architectures have provided the solid frameworks onto which applications can be deployed. Increasingly, however, they are also providing the flexible, scalable foundation that helps datacenters cope with change.

- The best practices and experience that Sun has developed over the last 25 years find themselves embodied in its products, services, solutions, and its datacenter design principles. Sun products provide an excellent set of building blocks, including the server technology that allows IT organizations to support a range of applications, and systemic qualities ranging from scalability and availability to security.
- Sun's architectural expertise can be brought to bear on datacenter design and re-design efforts, maintaining high quality-of-service levels while providing a solid foundation for growth and change.
- Sun understands that the best architectures are meaningless unless they can be deployed, powered, and cooled in the constraints of existing datacenters. Sun's focus on power efficiency and eco-responsibility helps customers keep their power and cooling budgets under control as well as help reduce their carbon footprint.
- Sun's managed service offerings can provide a foundation for managing change by integrating the people, processes, and technology that, combined, are critical to the success of any IT organization. Sun can provide a range of managed services, from the complete outsourcing of datacenter functions to working side-by-side with customers, educating them and helping to instill Sun's best practices in your organization.

The Challenge of Constant Change

There is probably no more constant in IT organizations today than constant change. YouTube, for example, grew from a startup to a billion-dollar enterprise in a mere 18 months. The several months it would take to determine exacting requirements and design specifications may be necessary for some parts of an IT organization, such as storing critical data, but other aspects of information technology must move at a lightning pace in order for a company to be successful. YouTube would not have been able to rise to the level that it has if it had spent 18 months developing design specifications. Although this is an extreme example, it does illustrate how architecture helps to strike the balance between requirements, implementation constraints, and time to market.

Companies still must spend time on critical infrastructure areas, including messaging-oriented middleware, database services, and other areas that need secure, scalable infrastructure, and which protect the data critical to a business.

But a major shift in application development is affecting how rapidly an IT organization must change, incorporate, and support new services on an ever-expanding set of devices ranging from desktops to mobile devices. Companies are adding value by interfacing with customers, partners, suppliers, and even competitors, as they develop new applications and services based on general sets of requirements. Development cycles have thus become shorter, more iterative, and more interactive with both business stakeholders and customers. All of these factors point to the need for datacenter design principles that can help organizations better adapt to a world that is characterized by constant change.

Architecting Dynamic Data Centers

In a world where applications must be developed practically overnight, with few formal requirements, and then scaled as customers demand the service, patterns can help provide guidance. Patterns encapsulate core problems and solutions that can be used as elements of specific business solutions. Architectural patterns may encapsulate part of a solution, for example a firewall sandwich configuration, or they may represent an entire solution, such as an enterprise mail subsystem. Patterns can also help demonstrate the *forces* or *systemic qualities* of a solution.

Patterns are useful tools for architects to use when exploring how various solution components can be combined and interconnected to create successful architectures. They raise the level of abstraction, allowing architects to focus on solving a specific business problem rather than becoming mired in an undifferentiated sea of details. Patterns also can be used to develop standards, since they often highlight how a problem is being solved and the techniques and products used.

Patterns may be arbitrarily complex, but they are ultimately based on a set of building blocks — individual elements with no interdependencies (Figure 1)

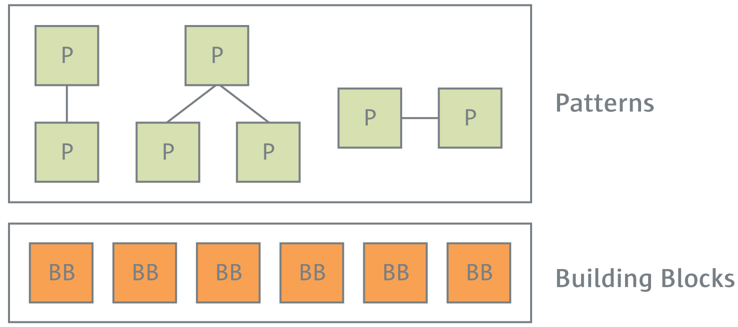

Figure 1. Patterns include interrelationships, and are based on independent building blocks.

Building Blocks

Building blocks represent the smallest, irreducible component or object within the context of a pattern or architecture. As a general rule, building blocks provide clear features, clean interfaces, and have limited interdependencies with other building blocks. An example building block may be compute nodes or servers that process specific types of workloads, or a content load balancing switch.

Patterns

The concept of patterns in IT system design has evolved over the years, originating from Christopher Alexander's seminal book in the architecture field, *The Timeless Way of Building* (Oxford University Press, 1979).

From Sun's perspective, a pattern is defined as a generalized, reusable design solution to a recurring network architecture design problem. Patterns derive from the experience of solving the same or similar design problem over and over. Once defined, the generalized design solution can be used heuristically, applied to problems with similar design challenges, and used to establish standards. Patterns can be used for the purpose of specifying a "best fit" solution that is shaped by the prevailing and relevant forces that are present in each implementation project.

Patterns are observed, not invented, so they are descriptive, not prescriptive. They are defined at a sufficiently high level that allows architects to avoid working through and getting bogged down with specific details. While patterns are adequate for high-level discussions and direction, the devil is always in the details, and patterns help to standardize solutions that have the details already ironed out.

The concept of patterns provides an excellent basis for documenting and describing the various attributes that make up a solution to a problem. In practice, there are patterns and there are also "bad," or anti-patterns that specify solutions to be avoided.

For example, a service requiring an implementation that incorporates redundancy through horizontal scaling, network traffic distribution for load balancing, and traffic redirection for high availability, would use a *load balancing/failover* pattern (Figure 2). Similar, but inappropriate pattern matches include: *clustering*, which provides horizontal redundancy and traffic redirection but not traffic distribution; and *load balancing*, which implements horizontal redundancy and traffic distribution but not redirection (failover).

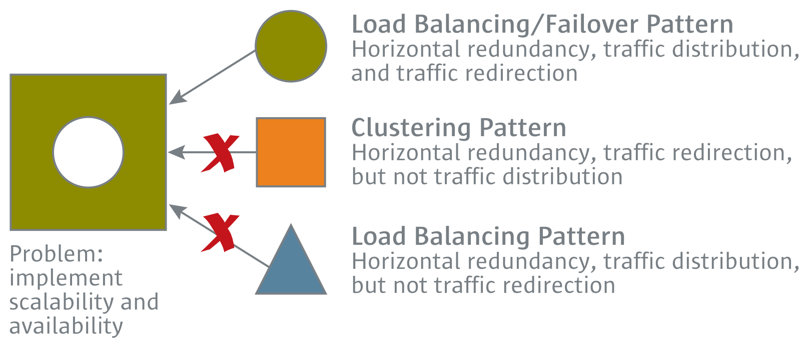

Figure 2. Sun's library of patterns helps its architects to quickly determine the right solution to a problem.

Standardization

Many large organizations rely on standards to help manage the challenge of creating and deploying good datacenter and application architectures. Good standards dictate both the products, configurations and reference architectures that should be used to solve a recurring problem. Unfortunately, for most companies, 'standards' only consist of a list of approved products. In addition, standards often do not keep up with rapid change in technology or with business requirements. Organizations that depend on standards alone may find that they cannot keep them up to date fast enough.

In contrast, Sun's approach is to use patterns to highlight common architectural-building blocks, and then use those patterns to derive standards. This approach leads to better standards that are based on real-world observations, not optimistic scenarios that are invented by committee and not always tried in the field.

Standards are everywhere in the mature data center. Whether formal or informal, they include architectures, product selections, configurations, and processes for managing the components. Making changes in datacenter environments requires considering many complex dependencies. Standards are the basis for moving up the operational maturity scale, and introducing automation and other means of increasing efficiencies.

Some important standards used by Sun include Points of Delivery (PODs), Standardized Operating Environments (SOEs), and use cases.

PODs and Workloads

Points of delivery provide environments that are optimized for specific workloads. They encapsulate storage, networking, management, and servers. One example of a POD is the Sun Customer Ready HPC Cluster, a platform that allows IT organizations to deploy a standard set of preintegrated servers, switches, and storage devices with rack-at-a-time granularity (see “Sun Customer Ready HPC Cluster and Scalable Storage Cluster” on page 46.)

Workloads help describe the resources an application needs to run. In order to meet target quality-of-service levels, they must run on an underlying architecture that optimizes their service levels. For example, front-end Web server workloads have significantly different characteristics than database workloads. Database workloads themselves differ based on the types of transactions they process. Because of these differences, different applications require different mixes of resources including disk I/O, network bandwidth, and CPU processing power.

Systemic qualities such as scalability, availability, and security also affect how workloads are deployed. If a workload requires high availability, then the application architecture and deployment architecture supporting it should meet the required service levels. Similarly, multiple applications with the same workload footprints may have different security requirements, and they may not be able to run on the same infrastructure.

The POD hardware platform layer consists of compute hardware, network, and storage. The availability and scalability requirements and the service tier that the hardware is intended to support often drive the specifications of the servers. Applications can scale independently. As applications need more resources than are available in a POD, additional PODs can be added, providing more capacity. Both horizontal and vertical scaling can be used as appropriate for each application.

Horizontal Scaling for Web-Centric and Compute-Centric Workloads

Applications that perform lightweight processing and are stateless can easily be scaled through a horizontal scaling model. This model uses multiple identical servers combined with load balancing and network traffic redirection mechanisms to create an infrastructure with horizontal redundancy and high availability. These applications often process network-centric workloads and reside at network edges. Examples include:

- A Web site’s presentation tier, with horizontally-scaled Web servers
- A Web site’s business tier, where state is pushed down to a database tier
- Infrastructure services, such as Lightweight Directory Access Protocol (LDAP) and Domain Name Service (DNS) services

Some compute-intensive workloads are also candidates for the horizontal-scaling model. High-Performance Computing (HPC) applications including climate modeling, automobile crash simulation, and seismic analysis are often built to run on compute clusters: large numbers of small servers each of which work in parallel on some small portion of a larger problem.

Sun offers a range of server products that support the horizontal-scaling model, including the Sun Fire™ T1000 server powered by the UltraSPARC T1 processor with CoolThreads™ technology. This unique processor technology supports the processing of up to 32 parallel threads per processor with minimal power consumption, ideal for applications such as front-end Web servers. For horizontally-scaled applications that are compute-intensive, performance per thread is most important, and Sun's AMD Opteron processor-powered servers including the 1U Sun Fire X2100, X2200, and X4100 severs, and the 2U Sun Fire X4200 serves are frequently-specified systems.

Servers used for a horizontally scaled service or application typically have 1-4 CPUs per server. Differences around standards for this class of server are usually related to the amount of compute capacity needed and reliability requirements. For example, a set of horizontally-scaled Web servers can use servers with a single power supply because any single server can fail without affecting the overall service. Long-running HPC applications, where the failure of a single server may require re-calculating a result that takes a week to calculate, often require redundant components (including power supplies) to minimize the impact of a server failure.

Sun has made deploying large numbers of horizontally-scaled servers even easier with the Sun Customer Ready HPC Cluster. Rather than purchasing, receiving, installing peripherals, racking, and cabling on a server-by-server basis, customers can now order servers a rack at a time, with options preintegrated by Sun Customer Ready Services. This helps IT organizations to rapidly scale their infrastructure and minimize the time from receiving a rack of systems to putting them to use.

Vertical Scaling for Data Centric Workloads

Sun offers high degrees of scalability in both its UltraSPARC and AMD Opteron processor-powered server product lines. Sun's UltraSPARC processor-powered servers scale up to 64 dual-core UltraSPARC IV processors. Sun's x64 server product line scales up to eight processor sockets populated with dual-core AMD Opteron processors. The Sun Fire X4600 M2 server can be upgraded to use up to eight quad-core AMD Opteron processors as they become available.

Stateful applications such as database management systems and other resource-intensive services share state between an application's multiple threads, requiring vertical resource scaling through the addition of CPUs and memory. Vertically scaled applications can require scaling to the point where the server has no more headroom for additional CPUs and memory. IT organizations can choose between two different strategies. They can choose a large server that can scale the application to its maximum capacity, adding CPUs and memory as workload demands dictate. Or they can choose a server product line that offers servers with increasing levels of capacity, replacing the server each time that the maximum number of CPUs and memory has been reached.

Standardized Operating Environments

Standardized Operating Environments (SOEs) describe the runtime stack of applications and the services they require on the operating platform or server. They represent a contract between the hardware and software and the business services that run on them.

An SOE may be very specific for certain types of applications, or more generic to support hundreds of applications. Virtualization technology has an impact on SOEs as well. Today, a server can be configured with a generic Solaris operating system global zone, and specific containers can be configured later and “hardened” to support each application.

SOEs are versioned, as are all parts of the infrastructure model. They include common agents, such as backup and management software, a specific operating system version and patch level, and the tools necessary for support (Figure 3). SOEs include not only the binary representation of the operating system and patches, but also the application and application configuration that runs on the OS. To an increasing degree, the Sun N1 product family can provision and manage an entire SOE from the bare metal to the application configuration.

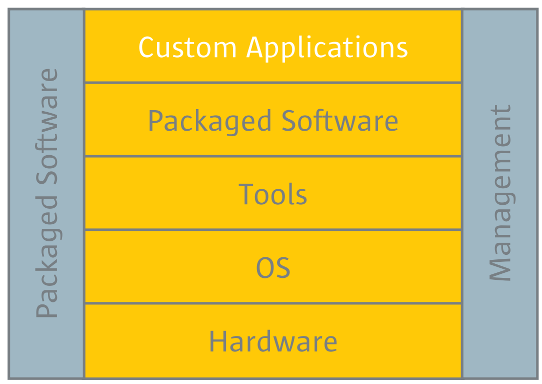

Figure 3. Standardized operating environments include all of the components necessary to run and support a service.

Use Cases

Use cases provide a tool by which common, repeatable operations can be well defined and then automated, helping to leverage an IT organization’s staff and reduce errors. Everyday steps in the operation of a datacenter include:

- Installing a Server
- Requesting an DNS Name and IP Address
- Installing an Application
- Moving an Application
- Failing over to another data center

In the dynamic datacenter that Sun promotes, common activities can be modeled, and automated, providing a basic level of capability that can be used across environments and applications.

Service Delivery Networks

Service Delivery Networks (SDNs), the topic of the next chapter, provides a standard way to design standard, scalable, fast, and secure datacenter networks. Service delivery networks can be implemented using PODs, and managed using Sun’s Dynamic

Infrastructure (DI) solution. PODs can be plugged in to a properly designed SDN to provide scalable services, and it provides the foundation for communication and growth in the network and network services.

People, Process, and Technology

It is easy to focus on the technologies that provide services in a modern datacenter, and not account for the people and processes involved in day to day operations. Sun's datacenter solutions take all parts of the equation into account: the checks and balances of ensuring proper IT lifecycle management, understanding business drivers that may affect the technology infrastructure and architecture, and the elements of solutions that provide management tools and functionality.

Sun’s approach of using use cases to drive pattern-based solutions helps to design dynamic datacenters that are ready to adapt to rapidly-changing business requirements, and are manageable today and into the future.

Chapter 3
Service Delivery Networks

The Sun datacenter design principles discussed in the previous chapter stress the use of patterns in datacenter designs. Through the use of patterns, network architects can recognize common scenarios for which there are well-understood solutions. Sun Service Delivery Networks (SDNs) is a modular, pattern-based approach to creating logical architectures that can then be deployed onto physical architectures.

The service delivery network methodology results in a logical design that is specifically created to meet the needs of a particular workload. It may, for example, specify the service components that must be assembled to provide a particular service to an end user, along with all of the required supporting infrastructure services such name, authentication, and messaging services.

IT organizations can take the logical architecture defined by the SDN methodology and then deploy it onto its physical architecture. As workloads change, the architecture can be replicated or scaled easily, allowing IT organizations to quickly meet the changing requirements. Service delivery networks take IT organizations one step closer to having a dynamic infrastructure, one that can rapidly adapt to changing business conditions, helping them to more quickly align with their company's business objectives.

SDN Building Blocks

The Service Delivery Network methodology uses the following set of building block components to construct service-optimized network design solutions:

- service instances
- service domains
- service and distribution modules
- optional specialty modules, including security and caching modules
- network components including switches, routers, and load balancers

These building blocks can be grouped into three entities that can be combined in various ways to deploy services (Figure 4). They are essentially classes that are instantiated in the process of deploying a service delivery network.

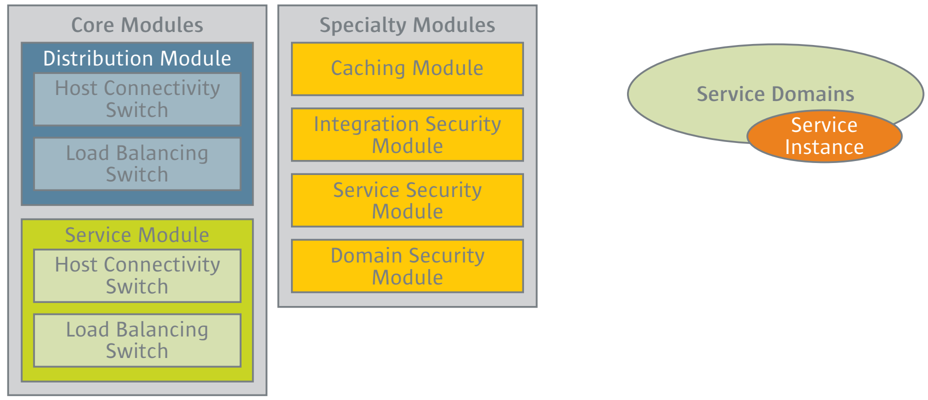

Figure 4. SDN building blocks include core modules, specialty modules, and service domains and instances.

Service Instances

A service instance is a running application that provides part of a service. Multiple service instances are typically deployed, each one providing the same service component but on different physical hardware or infrastructure.

Figure 5 illustrates two service instances, each of which is contained within a virtualized environment. The lower service instance is instantiated within a Solaris Container. The Solaris Container requires various services, including storage, execution, and network resources which are provided by the container and it's operating system.

The upper service instance illustrates how they can be deployed within physical or logical domains, including Dynamic System Domains that are physically isolated at the electrical level on Sun SPARC® Enterprise servers, and Logical Domains (LDoms) that provide software-level isolation on Sun's UltraSPARC T1 processor-powered servers.

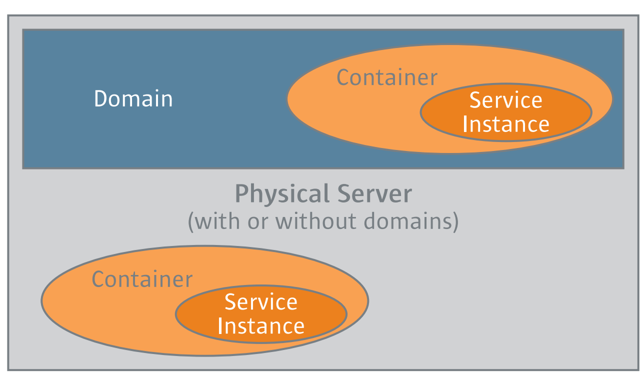

Figure 5. Two service instances, one deployed within a Solaris Container, and one deployed within a physical or logical domain.

Service Domains

Service domains are patterns that organize similar services into logical groupings that can then be managed at a higher level. A service domain consists of multiple service instances and the containers running on physical servers that provide the same, or similar, logically grouped services.

Separating services into service domains provides an additional scalability mechanism that allows services to take advantage of other service domains, and allow them to use all of the domain's features and capabilities.

This approach offers significant security benefits when the principles of compartmentalization and least privilege are applied to the offering of policy-based services. Elements in a service domain can be isolated from each other where inter-element communication is not required (for example, between Web server instances). This approach also mitigates the effects of undiscovered flaws, because the execution containers are not available on the network except through the specific service they provide for the service domain. Any other services running on the host are effectively out of reach. Finally, the approach of applying policies to network interfaces (in, out, and management interfaces) allows organizations to specify precisely how services can communicate. This significantly simplifies intrusion-detection and firewall policies.

Service domains provide:

- Network services to operating system instances, containers within them, and their dependent application services.
- Physical access layer and communication paths to each service instance. Each service domain provides network services for a single application. Applications and their service domains aggregate to provide a business service or services. Each service domain is a logical collection of similar service containers, each providing the same application service.
- Physical network access to containers and underlying servers. Each server has at least two physical connections to the service domain host connectivity switch to help ensure that the service instance will remain available in the event that a failure occurs within a network switch or cable.
- Security features, such as Access Control Lists (ACLs) and physical port isolation.

Service Modules

The next higher SDN building block is the service module, which consists of service domains and the hardware required to support them. A service module consists of physical network hardware, server connections for the service networks, and the software applications that make up the services to be delivered. Using these components, a service module provides the services, physical access, routing, distribution within the service module, availability features, integration to other networks (for a single service module), and network services. A service module provides a service through its service delivery interface, as illustrated in Figure 6.

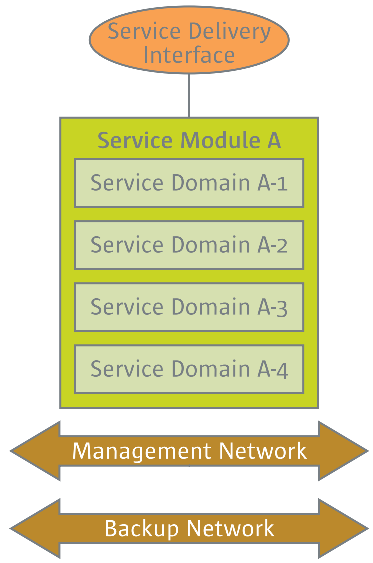

Figure 6. Service module provides a service through its service delivery interface.

A service module consists of one or more service domains to provide a specific business service, or a set of services that support a particular architectural layer, such as a presentation layer. For example, one service module may deliver messaging services, and another might provide identity services.

Service modules consist of more or fewer service domains depending on a number of factors, including the size of the deployment. In small deployments, a service module might contain a mixture of services, while in larger deployments might contain a more homogeneous mix of services. A service module can exist as a standalone component, or it can be linked together and scaled as needed.

Service Delivery Interface

The service delivery interface is the primary service interface. It provides the integration point to upstream access providers, LAN or WAN access, and is the primary connection point for clients and end users. Typically, the service delivery interface connects to a router, pair of routers, or other distribution network that connects to the data center. This allows the network design to be self-contained behind this interface, and limits the amount of configuration necessary above this point for service routing.

Distribution Modules

Distribution modules allow the integration of several service modules. They provide intelligent service routing, presentation services, and other features in a highly scalable, flexible design. For example, after capacity has been reached within a single service module, a distribution module and additional service modules can be added. Distribution modules allow a better logical layout of services and help reduce complexity while maintaining required service relationships. They are required when multiple service modules are used in a single environment. Figure 7 illustrates how a distribution module can be used to integrate multiple service modules.

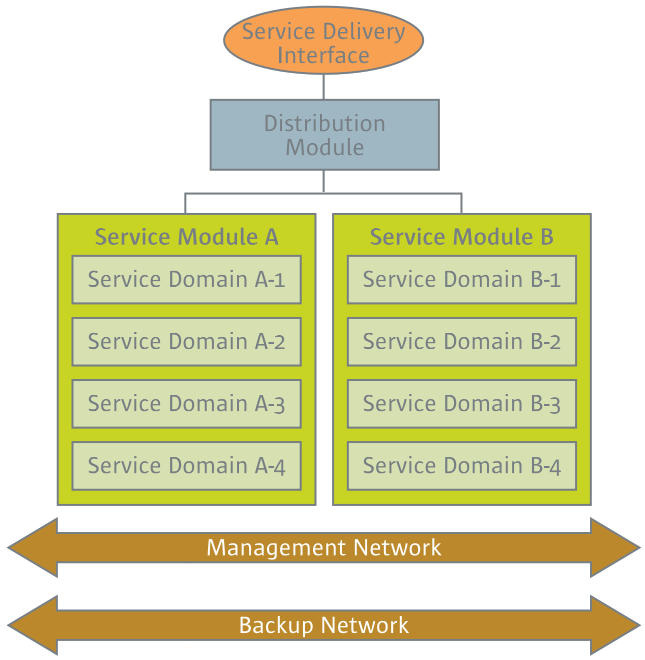

Figure 7. A distribution module allows multiple service modules to be integrated.

Distribution modules consist of network components that are similar to service modules. They enable several service modules to work together and aggregate services to a single service delivery interface that clients and users can access. Because distribution modules tie several service modules together, service modules may be isolated from one another, use different implementation hardware, or be managed by different organizations. The distribution module provides central access to the delivery of those services which can be managed by a centralized network team. They may also manage common specialty modules, such as proxies or caches, that provide these services to the entire network.

Specialty Modules

Specialty modules provide specific services in the SDN methodology. Depending on the deployment location within the network, a specialty module may provide services to the entire network architecture, to individual service modules, and to individual service domains. Specialty modules may be optional or required based on business and technical requirements.

Examples of specialty modules include the following:

- Security
  - Firewalls
  - Intrusion Detection Systems
  - VPN Gateways/SSL Gateways
- Performance
  - Caching
  - Application Acceleration
  - Network Shaping/QOS

Security Modules

Service delivery networks support pluggable security modules that can be used to augment the native security of individual platforms and services, as well as security controls implemented in the network infrastructure itself. This allows for the use of security appliances and devices such as firewalls, intrusion detection and prevention systems, and virtual private network concentrators to be integrated into the architecture. The decision to use these modules must be made based on an analysis of organizational policies and requirements, in conjunction with the threat profile for the network architecture.

SDN utilizes three types of specialty security modules: integration, service, and domain security modules. The architecture's in-depth security goal distributes security–related functions across the systems within the environment. The Solaris OS provides two security features that are particularly relevant to service delivery networks: the Solaris Security Toolkit, and support for IP-based filtering on each host.

Security and Management

A high security architecture must incorporate multiple, independent, mutually reinforcing, and different security technologies. A decentralized approach to security in the SDN architecture allows all of the components to protect themselves. This approach enables the ability to provide a high level of security with few, if any, inspection bottlenecks.

Differing, independent, and rationally configured technologies can be incorporated at every level of the network and software stack to reinforce the security of individual components. When designing a network architecture, architects must integrate the required security framework while considering any potential impact on network latency. Network architects must pay particular attention to the high-bandwidth segments of any proposed network architecture and minimize the number of complex security checks that could impact latency.

Integrated Security Framework

The SDN methodology uses an integrated security model that considers the security features in applications, containers, servers, network, security modules, and other security features in all of the different interfaces in the service delivery. Figure 8 illustrates this integrated security framework and how it is integrated with the various levels of the network and software stack.

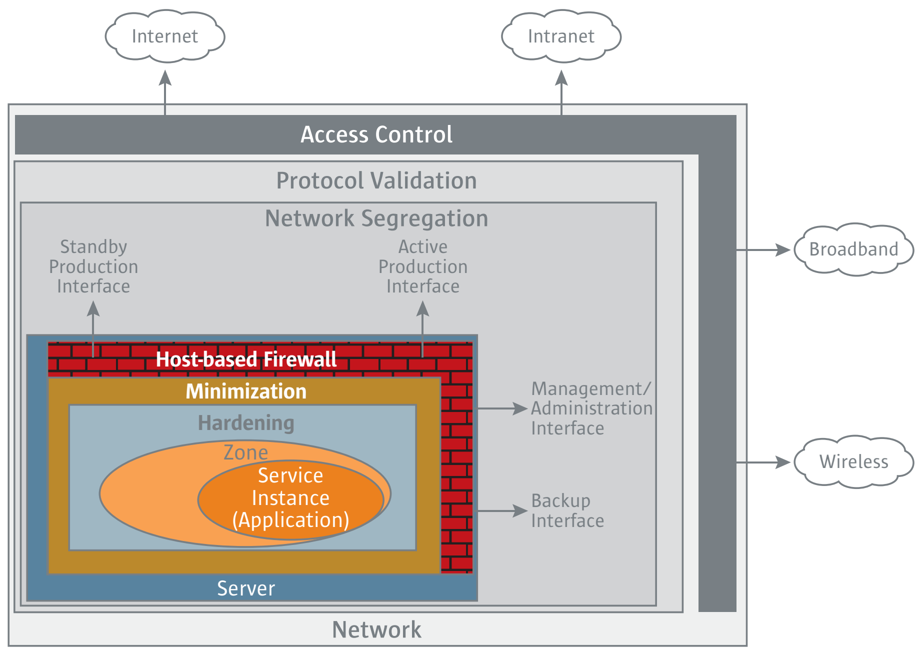

Figure 8. The SDN integrated security framework considers security at the various levels in the network and software stack.

High security depends on securing each aspect of the path by which traffic flows through the network, its components, and applications, all the while maximizing performance and reducing latency.

Security Zones and Compartmentalization

Security zones enhance security and flexibility by enforcing compartmentalization. Security zones can be implemented using virtual switching technology within a single load balancing switch. Each security zone in a single unit involves a virtual switch and includes a group of service domains that have similar security requirements. In addition, each security zone can be administered individually, which means that the administrator for one security zone cannot administer another security zone without explicit authorization.

Secure Management Networks

Network administrators must be able to manage and monitor services and applications in a network, as well as any related logical and physical components. The SDN methodology incorporates a management network into a network architecture in a secure and reliable way, minimizing risk to the production side of the network.

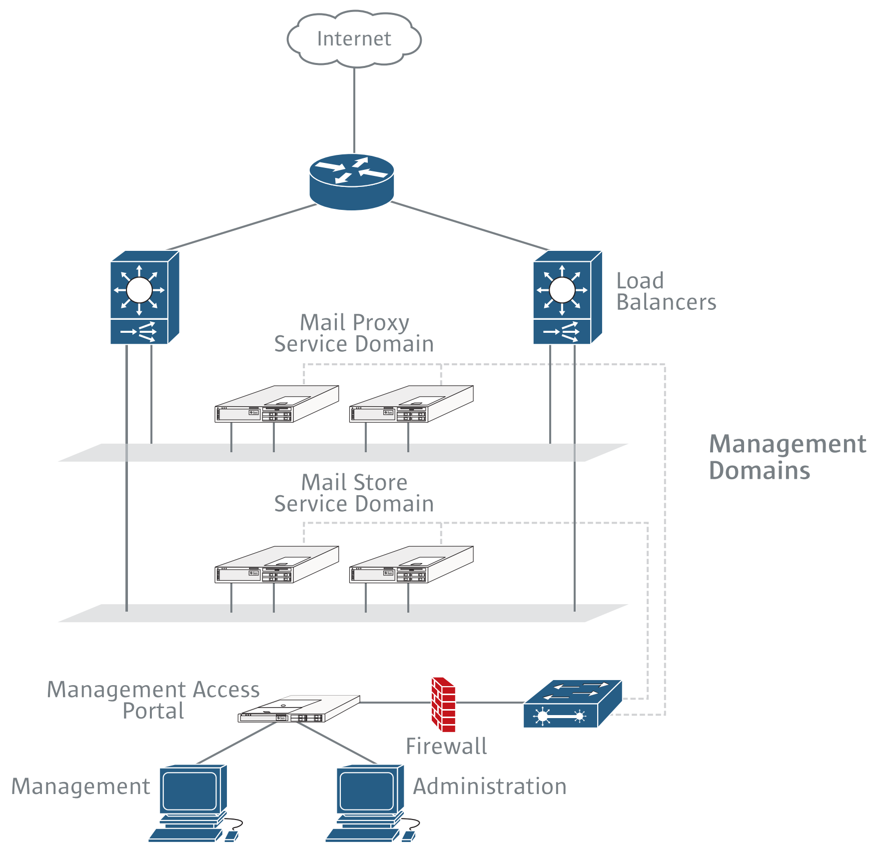

Figure 9. The service delivery network architecture uses segregated management and production networks.

Figure 9 illustrates how the operational, client-delivery side of a network can be segregated from the management network. Inbound client requests are distributed by a pair of redundant load balancers. The management network connects to each server and to a switch front of a firewall that secures the organization's management systems. The figure illustrates an access portal with various administration workstations attached. The management network infrastructure follows the security requirements of each service, and is separated into management domains.

Management Domains

Management domains is another design pattern that can used to satisfy network management requirements. Incorporating management domains into a network architecture design provides access to common management features including those described in Table 1.

Table 1. Common management features

Feature	Description
Network Management System Integration	Includes network and system monitoring, usually performed using SNMP v2 usec, v3, or v3 usec, Sun Management Center or Customer Network Systems (CNS) agents, dynamic infrastructure management's tier, or third-party products. Includes support for fault monitoring, performance monitoring, and security
Access to Common Management Services	Services including DNS, NTP, SYSLOG, RADIUS
Automated Provisioning	Includes the ability to automatically manage network configuration, especially for common functions such as VLAN creation, modification, and load balancer updates. This coordination can be initiated by a provisioning product or a workflow management product, and integrated with third-party network management products.
Administration Access	Usually provided through a command line interface, such as `ssh`, which is recommended, or (worse-case) `telnet`. Access through the NMS can provide another channel (usually via SNMP) for additional management options, including serial connections for out-of-band management and access.
Network Equipment	Provides vendor-specific management capabilities, such as SNMP, HTTPS, or XML.
Hosts and Servers	A management network can provide access to agents such as SNMP, Sun Management Center, Teamquest, or BMC Patrol.
Staging Network Access	Enables support for operating system loads, initial setup, and provisioning. Using Sun N1 Service Provisioning System, systems can be automatically configured with an OS, then later configured for specific applications or containers.

NMS Server Network and Network Operations Center Access

Management networks can optionally use a firewall to provide network access s to the various management domains, the network operation center segments, and network management system server networks. For example, network management system software could be configured to allow both polling and push management traffic. SNMP traps could be enabled and could be configured to forward the traps to specific event servers located on an internal management network segment.

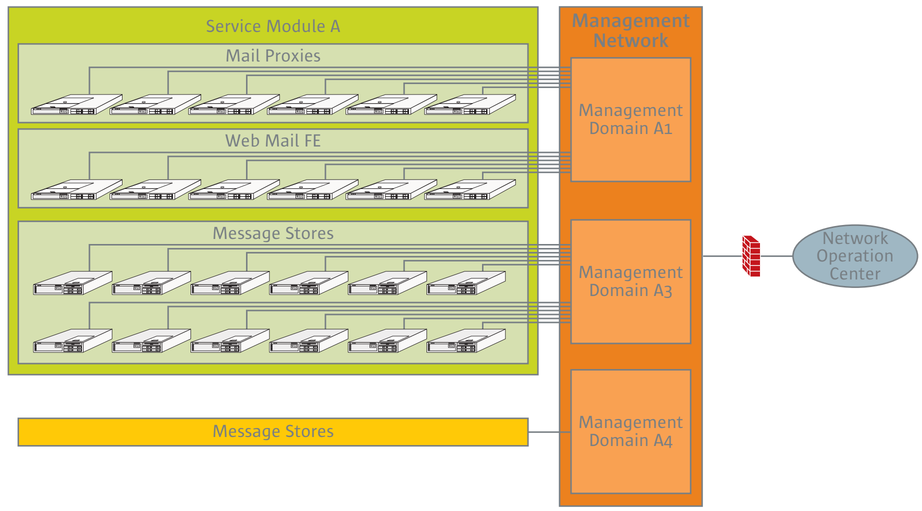

Figure 10. Management domains can be segregated securely while still allowing common access from the network operations center.

A logical representation of management domains segregated in this way is illustrated in Figure 10. In an actual implementation, each management domain could use separate switches or chassis-based switches using VLANs depending on security requirements and cost constraints. A network management system can store data polled from the network on servers located in an internal network management system server network segment. Because the management network is a standard IP network, it can be used to integrate a variety of commonly used network management system platforms, which can in turn be integrated into an network operation center LAN or other networks as needed — including customer networks. The network operation center can access management servers to retrieve and view network and system events and other management data through a secure channel to the management server segment. This design can be further enhanced via user profile-based VPN technology.

Chapter 4

Dynamic Infrastructure

Sun's datacenter design principles illustrate Sun's approach to using patterns and standards in state-of-the-art datacenter design. Service delivery networks illustrate a patterns-based mechanism for deploying datacenter architectures. Sun's dynamic infrastructure suite of services bring these architectures together along with Points Of Delivery (PODs) and management. The Sun Dynamic Infrastructure for Web Services service is an example of how Sun's approach works to deploy dynamic datacenters.

The Sun Dynamic Infrastructure for Web Services service helps address technology adoption cycles, manage virtualization, and architect for flexibility and agility. Sun can help customers build an environment for network applications that can be deployed easily and then managed using both Sun and third-party products. The Sun Dynamic Infrastructure for Web Services service provides:

- Architecture and design of a standardized physical platform in terms of a POD
- Integrated service delivery network
- A Dynamic Infrastructure (DI) management tier
- Reference use cases that have been implemented within the dynamic infrastructure, such as "add a resource" and "add a service."

The DI management tier provides functional capabilities to realize common process automation and control of dynamic infrastructure POD resources, applications and the services that they provide. The DI control framework is implemented based upon a set of reference use cases. Examples of reference use cases include the following:

- Add A Resource (such as servers, racks, switches)
- Manage A Service (define a composite service to the dynamic infrastructure)
- Deploy A Service (to servers, PODs, and datacenters)
- Control a Service (start, stop, register, etc.)
- Migrate A Service (move a service instance)
- Patch a Platform or Service (update a service)
- User Defined (extensible to meet future requirements)

The DI management tier is logically comprised of the following components:

- Bare-metal provisioning (using Sun N1 System Manager and the JumpStart™ Enterprise Toolkit
- Service provisioning (via Sun N1 Service Provisioning System)
- Network Provisioning (Provided by Sun partners)
- Managed access portal (Provided by Sun partners)
- Monitoring/instrumentation (Sun Management Center)

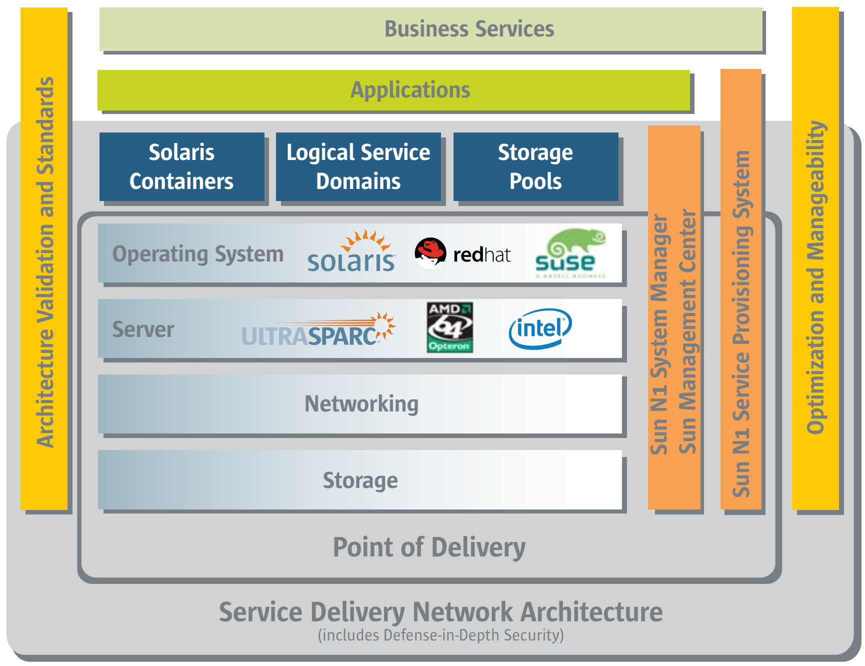

Figure 11. Sun's Dynamic Infrastructure for Web Services is an example of how Sun puts the concepts from the previous chapters into practice.

The dynamic infrastructure management tier can integrate with a Configuration Management Database (CMDB) for storing important system information including IP address and application configurations.

DI Management Platform

The DI management platform is a production-ready management suite that provides IT organizations with a foundation for managing highly available, business-critical applications in a virtualized datacenter. This can improve the utilization of existing servers, storage, and network devices, helping IT organizations to provision applications rapidly on existing or new hardware.

Management Platform Requirements

The management platform is a key requirement for delivering and managing services within a virtualized datacenter. Virtual servers or domains need to be created and destroyed, applications need to be provisioned, updated and removed. Services need to be monitored to ensure they remain within their service levels. Services and their supporting infrastructure need to be monitored and operators alerted when components fail. Ideally the failure of any component does not affect the service depending upon it, however a failed component must be flagged as having failed and eventually repaired.

Information on the state of the services and their supporting infrastructure must be available to the operations team. A centralized management portal is one way to provide this consolidated view of the datacenter. It can consist of multiple servers, multiple services and multiple service views, providing information to different management tiers in the organization.

A simple management platform can consist of a single server providing basic management services. A highly available management platform may consist of two or more clustered servers. A more sophisticated management platform can consist of multiple tiers, providing a help desk, asset management, service provisioning, IT service management, information portals, centralized system alerts and a CMDB.

Dynamic Infrastructure Management Platform Overview

Sun Dynamic Infrastructure for Web Services service includes a management platform to manage servers and network switches, deploy standard operating environments, create containers, deploy software applications to provide services, patch applications and monitor the servers and services thus deployed. The platform is built on either a single Sun server running the Solaris OS or two networked Sun servers in a highly available configuration using Sun Cluster software.

Services deployed on the DI management platform include:

- Management and monitoring
  - Platform management (power on/off)
  - Platform monitoring
  - Operating system provisioning
  - Container management
  - Application provisioning
- Operational procedure documentation (provided through the Sun i-Runbook service)
- NFS file system service
- Centralized system messages log

Wherever possible, service instances are deployed in Solaris Containers, taking advantage of their built-in security and resource-management capabilities of the Solaris OS. The management platform provides a provisioning server for the application and computation servers within the datacenter or POD, enabling the operating environment and business applications to be quickly, accurately and consistently deployed.

As Figure 12 illustrates, the DI management platform uses Sun Management Center software, Sun N1 Service Provisioning System, and Sun N1 System Manager software to manage every aspect of service delivery from bare metal power management and OS provisioning to monitoring the servers and security zones configured on them. The platform can also leverage third-party products not shown in the platform, including CMDBs, workflow management, trouble ticketing, and source code control systems.

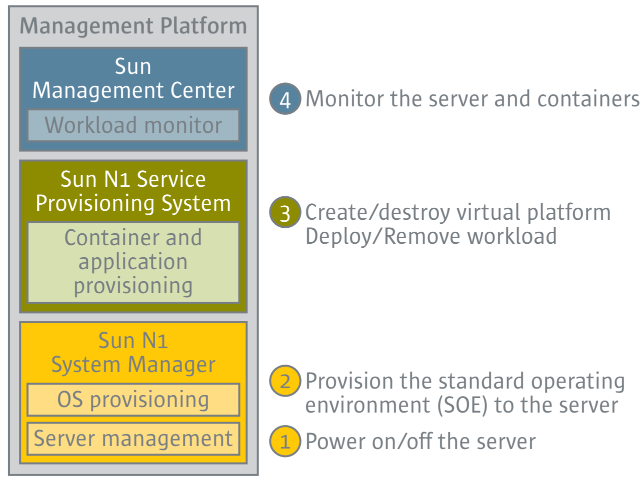

Figure 12. The DI management platform uses Sun Management Center and Sun N1 products to handle every aspect of POD operation, from bare metal configuration to monitoring servers and workloads in operation.

Management Platform Approach

Consolidating multiple services onto a single, reliable, highly available, Management Platform illustrates Sun's virtualization capabilities. Virtualization helps to reduce total cost of ownership over the use of dedicated servers, storage, and network devices. Through the DI management platform, a datacenter can benefit from rapid deployment of management and infrastructure services and of the virtual servers it provisions. The automation and integration testing done by Sun helps to deliver higher levels of quality and consistency in datacenter system deployments, helping to lower project risk and shorten deployment time scales.

Management Platform Capabilities

Typical capabilities from the operational point of view include:

- Rapid provisioning of standard operating environments and business applications
- Support for Solaris OS and Linux environments
- Centralized workload monitoring and management tools for servers, storage, and business applications
- Centralized server operational consoles and logging
- On-line build specifications, configuration information and technical procedures through the Sun i-Runbook service

Management Platform Customer Benefits

Sun’s dynamic infrastructure management platforms allows Sun customers to leverage their best practices and expertise to help achieve the following:

- *Simplicity*. Datacenters can benefit from a consolidated dynamic infrastructure management platform that will help reduce the complexity of application lifecycle management. Building and initial testing of a management platform configuration can be done at the Sun Customer Ready facilities, allowing a pre-integrated and pre-tested solution to be delivered with minimal disruption.
- *Speed*. Sun can help customers achieve their goals faster by leveraging proven configurations. Once implemented, the management platform is ready to provision and run datacenter applications.
- *Cost savings*. By building efficiencies based on best practices, Sun can provide a management platform in the most cost effective manner.
- *Smooth transition*. Customers receive run books and a transfer-of-information session with Sun experts to facilitate a smooth and efficient transition to production.
- *Scalability*. A management platform and Sun's proven configurations are flexible, helping meet customer requirements whether it's a large implementation or a small addition.

The Sun Dynamic Infrastructure for Web Services service is an example of how Sun can utilize a management platform to provide a highly scalable solution for rapid provisioning capability and ongoing management. Sun’s dynamic infrastructure solutions can scale quickly using servers deployed a rack at a time using the Sun Customer Ready HPC Cluster as a pre-integrated, pre-tested point of delivery.

Chapter 5

Server Virtualization Technology

IT organizations everywhere are being asked to do more with less, and nowhere do they face more hard limits than in the space, power, and cooling capacities of their datacenters. While the previous chapters have discussed principles and techniques for creating a dynamic datacenter that can help with rapid application deployment and scaling, services ultimately must run on server, storage, and network infrastructure, and this is where IT organizations face real limits:

- Legacy, silo-architected applications have contributed to server sprawl because each server is sized for maximum workload requirements of a single application, and resources generally cannot be shared. The result is a large number of aging, energy-inefficient servers that run with low utilization levels most of the time.
- State-of-the-art servers, such as those from Sun, put so much computing power in the hands of IT organizations that they often can be tasked to support multiple applications at the same time. Consolidation efforts, however, are hampered by the need to align operating system levels, administrative domains, and security contexts on single operating system instances.

Virtualization is a key technology that addresses these issues. It helps consolidate legacy applications from multiple obsolete hardware platforms onto a smaller number of up-to-date, more powerful, and more energy-efficient servers. It supports moving today's applications from a large set of under-utilized servers to a smaller set of more powerful servers, helping to reduce the number of servers to house, power, cool, and maintain. Increasing utilization levels helps to reduce inefficiency, helping with the space, power, and cooling crunch.

Nearly every server virtualization technique provides a mechanism for putting an execution environment into a container that can be moved between servers, giving IT organizations the freedom to change the allocation of services to servers. This helps maintain service levels by allowing IT organizations to vacate a server for maintenance or upgrade without a large effort to re-host the services running on it. Rather than dedicating resources to an application, and wasting excess dedicated capacity, IT organizations can deploy them in virtual environments that can be scaled up or down, depending on their requirements.

Virtualization Defined

Some form of virtualization has been in commercial use since the 1960s, and is experiencing a resurgence because of its ability to help solve so many issues in today's datacenters. There are many kinds of virtualization. Storage virtualization presents mass storage capacity and service levels as a virtual resource without exposing the details of disk subsystems. Network virtualization provides virtual networks with security isolation and different service levels using the same physical infrastructure. Finally, server virtualization, the focus of this chapter, provides virtual servers on physical ones.

Server virtualization creates the illusion of multiple computing environments on a single computer. This allows IT organizations to run multiple, different application environments, even multiple operating systems, on the same server at the same time. Each behaves as if it had its own dedicated computer system.

Sun supports several different virtualization technologies, each of which provide different degrees of isolation, resource granularity, and flexibility. Sun supports virtualization technologies that allow multiple OS (and application) instances to run on the same server, while each instance has the illusion of owning its own hardware resources. Sun Dynamic System Domains provide hardware partitioning capabilities on selected Sun SPARC Enterprise servers. Virtual machine technology including Xen, VMware Infrastructure 3, and Microsoft Virtual Server software provide software emulation of underlying hardware that supports multiple environments even on systems with only a single processor. Sun also has innovative virtualization technology that overcomes the limitations of both Dynamic System Domains and virtual machine environments: Solaris Containers provide security and resource isolation that allows multiple virtual Solaris OS environments to share the same OS instance. Logical Domains offers a hybrid of partitioning and virtual machine technology, allowing multiple OS instances to share the same hardware without the high overhead typical of virtual machine environments.

Virtualization Technologies

There are two different types of hardware virtualization technologies that have been in use for years, domains (or hard partitions) and virtual machines. (Figure 13):

- Domains are a hardware or firmware capability that divides a computer system's assets so they can host multiple OS instances. This was introduced to the mainframe world in the 1980s by Amdahl Corporation's Multiple Domain Facility, and introduced in open systems by Sun’s Dynamic System Domains beginning in the mid-1990s.
- Virtual machines date to the 1960s with IBM's VM operating system. Operating system software, called a *hypervisor* or *Virtual Machine Monitor* (VMM), creates virtual machines, each of which has the illusion of owning its own hardware and is capable of running its own separate operating system instance.

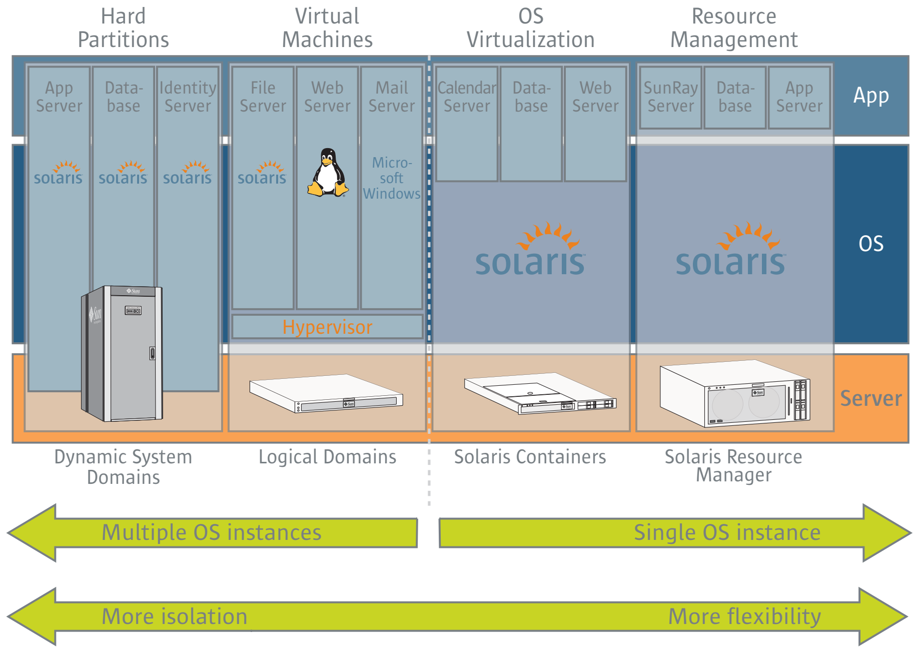

Figure 13. Sun offers a range of hardware and software-based virtualization technologies.

Sun supports Dynamic System Domains on supported UltraSPARC processor-powered Sun Fire servers, and supports VMware on its x64 server product line, which in turn supports the Solaris, Microsoft Windows, and Linux operating systems.

Domains vs. Virtual Machines

Although there is some overlap between these categories, the simplest way of distinguishing between them is that domains are a hardware or firmware function that subdivides a computer's physical CPU and RAM. Domains, or partitions, typically have low overhead and high reliability with strong isolation between different domains. They usually have a fixed, architectural limit on the number of OS instances or amount of resource granularity or sharing.

Virtual machines are typically supported by a hypervisor, which is itself an operating system implemented in software. The hypervisor time slices between guest operating systems just as a conventional multiprogramming OS time slices between application processes. Hypervisors usually don't have fixed limits on the number of OS instances, and often permit fine-grained resource control. Many virtual machines can run on a single CPU, but have substantial overhead that can affect performance and limit the number of OS instances that can be run effectively.

Hypervisors give virtual machines the illusion that they can run with the processor in both privileged (OS level) and unprivileged (application program level) states. When operating systems execute privileged instructions, for tasks like changing memory mapping for virtual memory, performing I/O, these state-changing instructions and events in one virtual machine must not affect others. To enforce this, hypervisors run virtual machines without full machine access privileges.

This causes a program exception, or trap, which causes a context switch to the hypervisor every time the guest virtual machine's operating system executes a privileged instruction. The hypervisor then determines what privileged instruction the virtual machine was attempting to execute (start I/O, enable or disable interrupt masks, flush cache, switch address spaces), and then the hypervisor emulates that instruction on behalf of the guest operating system. Similarly, the hypervisor must reflect an event in the guest OS such as a timer or I/O interrupt.

Hypervisors give the guest OS the illusion that it has its own machine, but at a price in complexity and substantial overhead that sometimes exceeds 50 percent. The memory footprint of multiple complete operating systems - kernels, libraries and applications, and the housekeeping tasks each OS runs also adds to the overhead in running multiple virtual machines.

Solaris Containers

Solaris Containers is a virtualization technology from Sun that address the limitations of domains and virtual machines. Solaris containers combine Solaris Zones, a technology that provides isolation, with Solaris Resource Manager, which can be used to allocate resources to zones. Solaris Containers, introduced with the Solaris 10 Operating System, provides OS-level virtualization because it gives the appearance of multiple OS instances rather than multiple physical machines. Instead of virtualizing the hardware architecture of the underlying machine as described above for virtual machines, Solaris Containers provides multiple instances of the OS abstractions that Solaris provides to applications. This can be accomplished without the high overhead typical of virtual machines because there is no need to do processor instruction-level emulation. Another advantage of Solaris Containers is that because they run on the Solaris OS, they run on any UltraSPARC, AMD Opteron, or Intel processor-powered computer that runs the Solaris OS.

Solaris Container Implementation

The Solaris OS installs with a single global zone that has system-wide access to the entire Solaris instance. The global zone can be used to configure up to 8191 non-global zones. Each zone operates in a private, isolated context, with its own network address and identity, security context, and potentially different applications. Figure 14 illustrates how a single Solaris OS instance can support multiple customers on a single hardware platform. Each customer can run their own applications and has access to their own resources.

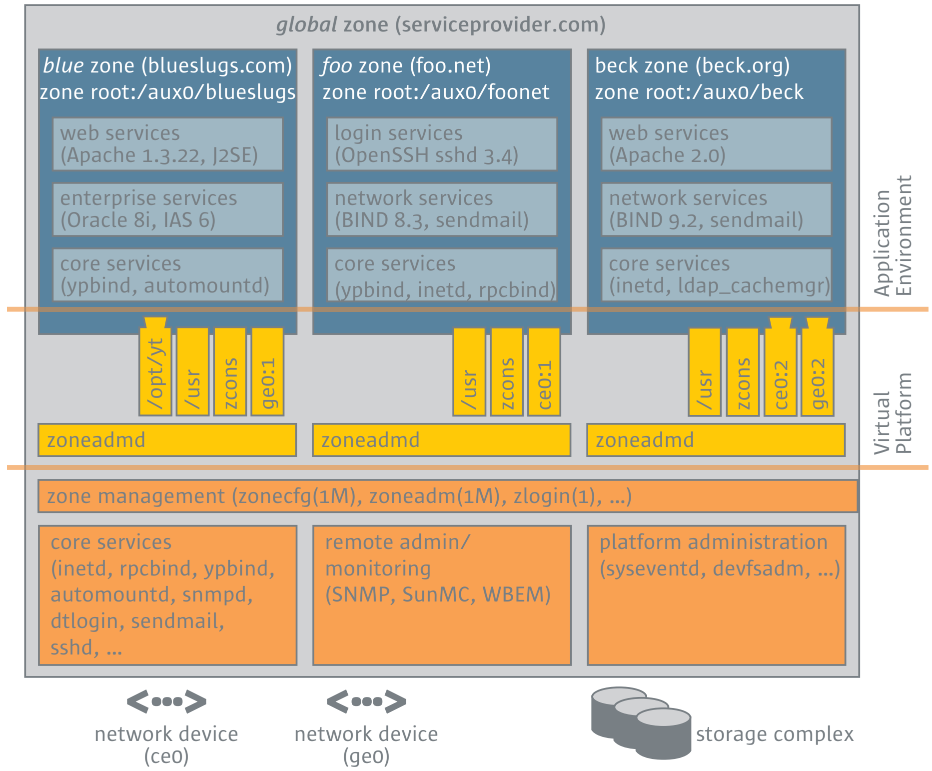

Figure 14. An example service provider server supports three different zones for three different customers, each of which run their own applications and have access to their own resources.

When a zone is created, it is provided with the following:

- A private startup and shutdown sequence and service management facility: each zone can be separately booted or rebooted without affecting other zones, and can run its own set of services.
- Private IP addresses and port ranges: these allow each container to be separately addressed on the network so that each container can independently host its own network applications. For example, an organization can create multiple zones, each of which hosts its own Web site using its own unique HTTP and HTTPS ports. This can be used to isolate different customers or to enable developers to create and test applications in private environments without being able to harm or conflict with one another. Applications in different zones communicate with one another via TCP/IP, just as if they were on separate physical machines. Inter-zone network traffic is handled by memory-to-memory transfer rather than the network interface. This allows extremely high bandwidth and message rates, and it allows network traffic to pass between zones securely, without the possibility of it being intercepted.
- Private authentication and name services: each zone can independently use its own password, DNS, LDAP or NIS naming services for users, passwords, services, and host

names. A given userid can be defined in one zone and not exist in other zones. Having root access in one non-global zone has no value in the other zones.

- Private process lists: each zone has its own processes and cannot see processes in the global zone or other non-global zones. This provides privacy and safety for the different zoned environments, just as if they were on separate physical computers.
- Separate security, resource management, and failure scopes: each zone can have its own rules for authentication and authorization, allocation of computer resources, and containment of application or even system faults.
- Private file systems: these appear to the zone as the root of the file system. They can be set up in two ways: *whole root* and *sparse root*. Sparse-root zones optimize physical memory and disk space usage by sharing some directories, such as `/usr` and `/lib`, in read-only mode. Sparse-root zones have their own private file areas for read-write directories, including `/home`, `/etc`, and `/var`. Whole-root zones increase flexibility by providing complete read-write file structures for the zone. This lets the zone administrator install products or patches into what would be read-only directories in a sparse-root zone, but at the cost of increased disk space. Additionally, the global zone administrator can assign zones devices, UFS mount points, ZFS pools, loopback filesystems and disk assets. Filesystems that are private to zones typically are part of the global zone's filesystem hierarchy, and are mounted through loopback filesystem mounts (Figure 15).

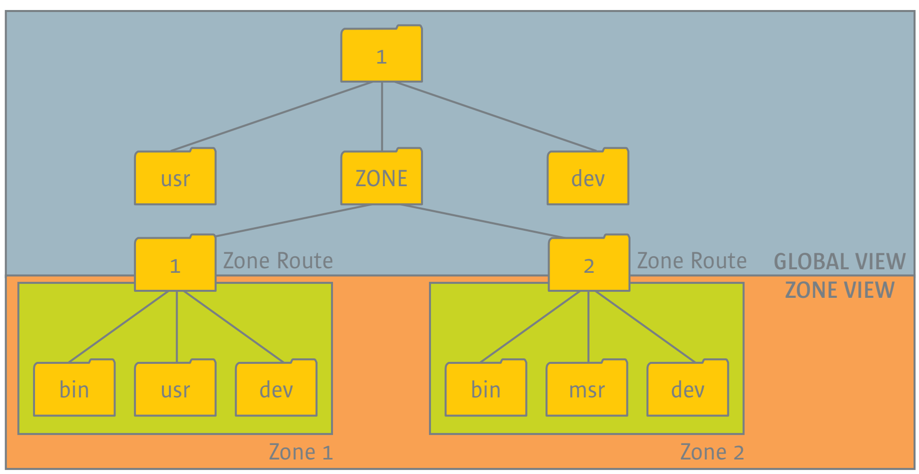

Figure 15. Filesystems that are local to a zone are typically part of the global zone's file hierarchy.

Low Overhead

Unlike virtual machines, Solaris Containers have minimal overhead. Isolation between zones is accomplished by restricting the scope of system calls to the zone from which they are issued, rather than the CPU-intensive task of emulating hardware architectures and instruction sets in software. This makes it possible to create hundreds, even thousands, of Solaris Containers on a single processor. Because of this negligible overhead, and unlike partitioning or virtual machines, Solaris Containers can be created in large numbers, for example, to give individual developers safe, isolated test environments.

Solaris Containers are available on any platform that runs the Solaris 10 OS, whether it is based on UltraSPARC, AMD Opteron, or Intel processors. Zones can be created by the global zone administrator in a few minutes. Solaris Containers can be cloned from a golden master, or even detached from one Solaris system and attached to another. Solaris Containers are included with the Solaris 10 OS, making it possible to have high-performance, low-overhead, low-cost virtualization on any Solaris OS platform.

All Solaris containers on the same platform run under the same OS kernel, so it is not possible to mix and match software that requires conflicting OS versions. The use of a single OS instance has the benefit that a patch or update applied to the global zone is applied to all of the non-global zones at once, instead of requiring individually maintained, separate OS instances within domains or virtual machines.

Logical Domains

Sun has recently introduced Logical Domains (LDoms), a virtualization technology for systems based on the UltraSPARC T1 processor, including the Sun Fire T1000 and T2000 servers. LDoms support separate OS instances without the overhead of virtual machines, and with more flexibility than hardware partitioning features such as Dynamic System Domains.

Minimal Context-Switch Overhead

Traditional virtual machines time slice physical CPUs among multiple virtual machines, intercepting and emulating privileged instructions. This is complex and CPU intensive — even simple context switching between virtual machines can require hundreds of clock cycles.

Time slicing is based on the historical design of computer systems where physical CPUs were relatively rare and expensive, and hence must be shared and time sliced. The UltraSPARC T1 processor turns the assumption of expensive and rare CPUs upside down. The processor's Chip Multi-Threading (CMT) design provides up to 32 logical CPUs (or *strands*) on a single processor. Using this multi-threading model, logical CPUs become plentiful, and they can be allocated to logical domains and dedicated for their use. This eliminates the cost of context switching a single processor between virtual machines.

Although context switching still occurs between processes in the same logical domain, the UltraSPARC T1 processor significantly reduces context switch overhead. Context switching typically occurs when a domain references memory that is not currently in cache. Accessing main memory to fulfill the request requires many clock cycles on any system architecture. While memory is fetched, a logical CPU stalls while executing the single instruction causing the cache miss. By switching to another CPU strand on the same physical CPU core, the UltraSPARC T1 processor lets another logical CPU continue instruction processing, during time that is unused memory wait time on most processors. On most CPUs, cache misses result in processor idle time. The UltraSPARC T1 processor can use what is dead time on other processors to continue doing useful work.

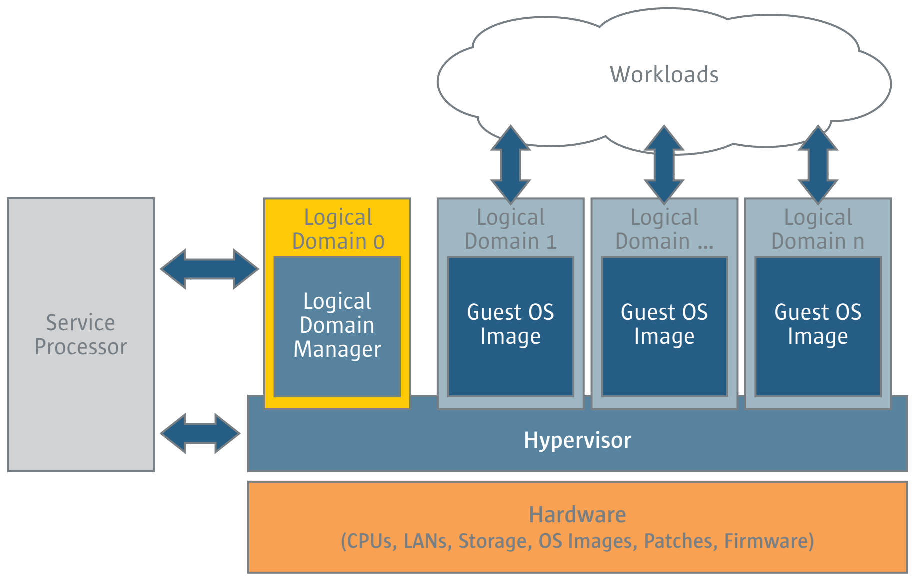

Figure 16. Logical domains allocate UltraSPARC T1 processor logical CPUs and memory to logical domains, allowing multiple guest operating system instances to be supported on the same CPU, but without the high cost of context switches between OS instances.

LDom Implementation

Logical domains are implemented using a very small hypervisor that keeps track of which logical CPUs, RAM locations, and I/O devices are assigned to which domain, and it provides communication channels to service domains (Figure 16).

Logical domains are designed for simplicity. Since each domain has its own dedicated logical CPUs, a domain can change its state (such as enable or disable interrupts) without having to cause a trap and emulation. That can save thousands of context switches per second, especially with workloads high in network or disk I/O activity. Since each logical CPU has its own private context in hardware, the UltraSPARC T1 processor can switch between domains in a single clock cycle, not the several hundred needed for most virtual machines.

Using logical domains, logical CPUs can be added to or removed from a Solaris instance individually and dynamically. A CPU can be added to or removed from a logical domain without requiring a reboot. The Solaris OS can make use of a dynamically added CPU for additional capacity immediately. Each domain has its complement of disk, network, and cryptographic resources, which are assigned by commands issued by the control domain, which also runs the Solaris OS (Figure 17).

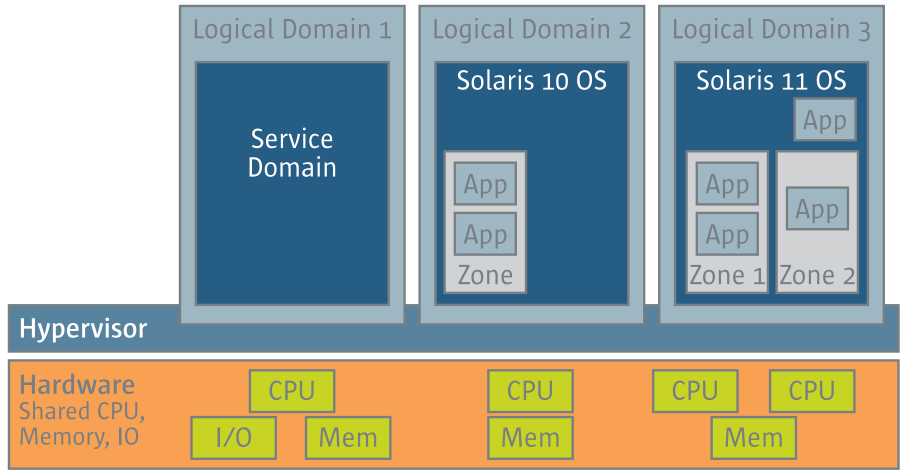

Figure 17. A service domain allocates CPU and memory resources to different logical domains, and can dynamically change the allocation without requiring a reboot.

Virtual network and disk I/O is provided to logical domains by service domains. A service domain is a separate instance of the Solaris OS that runs in a domain, but which has direct connections to network and disk devices. Regular domains, used for application processing, have network and device drivers that communicate with the service domains through communication channels provided by the UltraSPARC T1 processor firmware.

This allows service domains to export virtual LAN and virtual SAN services to other domains. Network resources can be set up for securely isolated VLANs, and disk resources can reside in virtual SANs, with actual disk images residing in physical disks, disk slices, ZFS pools, or even standard UFS files. The Solaris OS instances can use these disk and network resources exactly as if they were dedicated physical resources.

Service domains are designed for high availability: they can be set up in pairs so that system operation can continue if one is taken down for service or due to a fault without causing an interruption to guest domain operation.

Logical domains make it possible to provide low-cost, high-efficiency domains on UltraSPARC T1 processors with highly granular and dynamic resource allocation. They can be used in parallel with Solaris Containers for the highest degree of flexible virtualization. Separate OS instances can be configure d when different OS kernel levels required, with each OS instance hosting many lightweight, virtualized Solaris Container.

Chapter 6

Sun Server Technology

Patterns, service delivery networks, dynamic infrastructure, virtualization — all of the topics to this point have described the various ways in which Sun's expertise can help IT organizations design and deploy the most flexible and agile datacenters imaginable. When it comes to integrating the server, storage, and networking devices to make a datacenter perform, Sun's server technology forms the cornerstone.

Customers have a wide range of server vendors to choose from today. Many vendors believe that server technology is a commodity, and that every datacenter requirement can be met with a 1, 2, or 4-socket x86-architecture server. The most fundamental difference between Sun and the commodity server vendors is that Sun understands what its customers know: every application is different, and every application has different requirements that drive server choice.

Matching the right server to the application impacts everything related to service levels, including performance, scalabililty, reliability, availability, manageability, and serviceability. Even an IT organization's own inherent characteristics affect its server requirements: existing companies may be saddled with space, power, and cooling constraints that dictate the use of low-power, eco-friendly servers. A startup company, in contrast, often wants the most performance per dollar, with scalability the most important characteristic.

Matching Servers With Applications

Sun offers customers a range of server choices that allow them to closely match their application characteristics with a server product that offers optimal support, with some of the most eco-responsible choices on the market (Figure 18):

- *Back Office* applications require a high degree of scalability, and in some cases, large amounts of storage to match. Customers typically choose highly scalable symmetric multiprocessing servers that can grow with the speed of business. These servers deliver performance that allows back-office applications to run the most mission-critical applications in an enterprise.
- *High-Performance Computing* applications require the highest single-threaded performance possible, and many of them are highly floating-point intensive applications. Customers typically choose to run these applications on compute clusters populated with the most powerful processors available.

- *Web Infrastructure* applications are highly threaded applications that run many activities in parallel, where floating-point performance isn't the issue — it's the ability to context switch between many parallel activities with minimal overhead. Streaming-media applications need massive amounts of storage combined with the processing capacity required to push large number of packets onto the network.

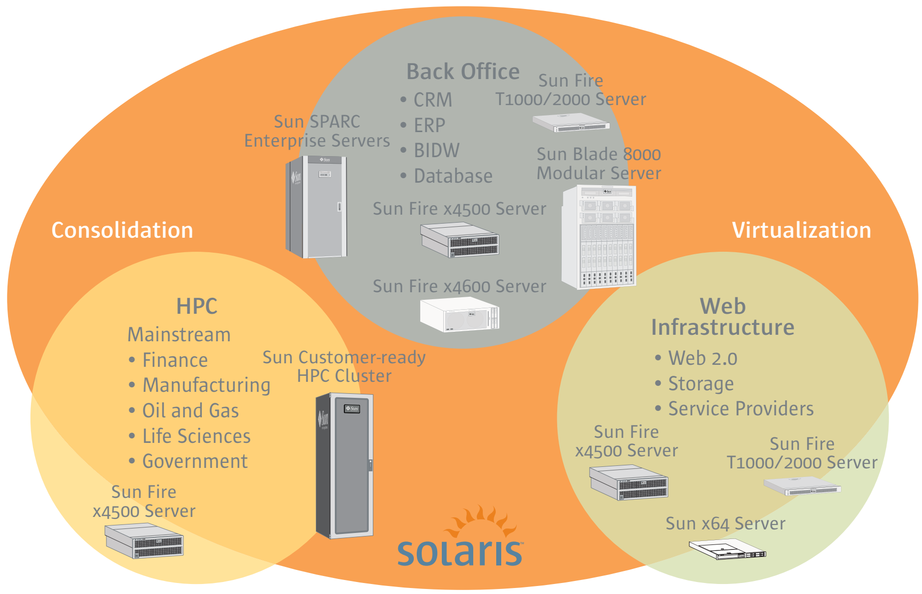

Figure 18. Applications tend to cluster when described in a two-dimensional space.

Matching Servers with Organizations

Just as different applications impose different requirements on their servers, so does the nature of the different companies hosting them:

- *Small and startup companies* tend to have meteoric growth. These customers push Moore's law and desire server platforms that can provide them with the highest performance at the lowest cost. Where these companies see value worth paying for is in platforms that can scale to support their growth.
- *Traditional, established companies* tend to have slower, more measured growth, and their pain points revolve around the real limitations they face: space, power, and cooling in their datacenters. These organizations often must focus on consolidating their legacy applications to make room for growth, which means choosing the most power-efficient and eco-friendly server technology available.

Over time, each type of company tends to assume some of the characteristics of the other. Established companies may need to deploy technology and applications more quickly, leveraging the benefits of communities, standards, and the open-source movement. Startup companies, in order to maintain their growth, need to do what established companies already know how to do: operate with procedures, broaden markets, and leverage the benefits of commercial software. Sun has designed its server product line so that it can accommodate the needs of both kinds of companies and also continue to support them as they grow and mature.

Industry-Leading Processor Support

Sun's server product line begins with industry-leading processors from Sun, AMD, and Intel. Sun's choice of processors allows it to build servers that can support a wide range of application requirements, ranging from data intensive to network intensive; from single threaded to highly multi-threaded. The processor selection that forms the foundation of Sun servers includes the following:

- *SPARC64 VI processor*
  The new SPARC64 VI processor consists of two SPARC V9 cores with two vertical threads per core. Sun's UltraSPARC processors are designed for the utmost scalability; Sun's server product line features the ability to have up to 64 dual-core SPARC64 VI processors in a single server.
- *UltraSPARC IV+ processor*
  The UltraSPARC IV+ processor represents the fifth generation 64-bit UltraSPARC processor. With significantly enhanced cores, dual-core UltraSPARC IV+ processors are used in Sun servers hosting up to 72 processors and 1 TB of main memory. This processor provides a seamless upgrade path for previous UltraSPARC III and IV processor-powered servers.
- *UltraSPARC T1 processor*
  Sun's UltraSPARC T1 processor with CoolThreads technology is designed for highly threaded applications where the ability to run a large number of threads simultaneously is extremely important. The processor is available with up to 8 cores, with each core capable of processing 4 threads simultaneously for a total of 32 threads. The processor is one of the most eco-responsible processors ever developed, providing the best performance per watt available.
- *AMD Opteron processors*
  The AMD Opteron processors that Sun has selected for use in its x64 server products represent a highly efficient balance between performance and power consumption. With high single-threaded and floating-point performance, dual-core AMD Opteron processors power some of the most compute-intensive workloads. Their Intel Architecture compatibility allows them to run all three enterprise operating systems, the Solaris OS, Microsoft Windows, and Linux.
- *Intel Xeon processors*
  Sun's choice to bring Intel Xeon processors into its server product line offers customers the ability to choose the most appropriate processor for their application when they choose Sun as their server vendor. For example, Intel's Core

Microarchitecture-based processors have a different cache structure than AMD Opteron processors, allowing them to excel on applications that make most effective use of the cache. The shared, 4 MB L2 cache configuration in the quad-core Intel 5300 "Clovertown" processor can provide some HPC applications with just the performance they need. As an addition to Sun's x64 server product line, Intel Xeon processors bring compatibility with the Solaris OS, Microsoft Windows, and Linux.

Operating System Support

Sun's choice of processors for its server platforms results in an ability to support a broad range of operating systems. Sun's Solaris OS runs on all of the processors that Sun supports, from AMD Opteron and Intel Xeon to UltraSPARC processors. Customers needing to run the Solaris OS, Microsoft Windows, and Linux, can choose servers based on AMD Opteron and Intel Xeon processors for the greatest flexibility and investment protection. And for those customers working on consolidating their datacenters, they can run all three operating systems on the same server using VMware Infrastructure 3 software on AMD Opteron processor-powered servers.

Compatibility

Processors and operating systems determine binary compatibility, and despite the wide range of processor options that customers have, the compatibility story is quite simple. Customers using the Solaris OS can enjoy binary compatibility up and down each of Sun's server product lines: applications compiled for 64-bit SPARC processors will run on any of today's SPARC processor-powered servers, and applications compiled for any of Sun's x64 servers will run on all of them. Customers running Microsoft Windows or Linux have binary compatibility across Sun's x64 server product line.

Customers using the Solaris OS have another dimension of compatibility that enhances both their choice and investment protection. Customers wishing to switch between Sun's SPARC and x64 server product lines need only re-compile their 64-bit Solaris OS applications in order to run them on the other product line, so long as their software is built following Sun's source code guidelines, including the ability to operate in both big-endian and little-endian environments.

Sun's Server Product Line

Sun has used its SPARC, AMD Opteron, and Intel processor choices to build one of the most comprehensive server product lines available anywhere:

- *Sun SPARC Enterprise Servers*
  Built around the SPARC64 VI processor, Sun SPARC Enterprise servers deliver world-class performance and scalability, with mainframe-heritage features to improve reliability, availability, and serviceability. Sun SPARC Enterprise servers range from the Sun SPARC Enterprise M4000 server (up to 4 dual-core SPARC64 processors and 128GB of memory) to the Sun SPARC Enterprise 9000 server (up to 64 processors and 2 TB of memory. These server platforms are ideal for organizations running large online transaction processing systems, data warehousing, CRM, and ERP systems.
- *Sun Fire Servers*
  Sun's Sun Fire server product line encompasses a range of servers based on the UltraSPARC IIi, IIIi, III, IV, IV+ and T1 processors, including from one to 72 processors per system, and a range of choices including compact 1U servers and mainframe-class systems. Sun offers a range of systems from rack-optimized, entry-level systems designed for throughput computing to large Symmetric Multiprocessing (SMP) servers offering massive computing power with high levels of I/O and networking capacity. For example, the Sun Fire T1000 server delivers the unique multi-threaded capabilities of the UltraSPARC T1 processor in a compact, 1U form factor, while the Sun Fire T2000 server offers the same processing power with additional I/O expansion capability.
- *Sun x64 Servers*
  Sun's x64 server product line includes options ranging from 1 and 2-socket, 1U servers, to its unique Sun Fire X4600 M2 server. This server is a modular, scalable system that can pack up to eight CPU sockets and up to 256 GB of memory in a mere four rack units, putting roughly twice the resources in the same space as other 4 RU servers on the market today. Where the Sun Fire X4600 M2 server packs compute density, the Sun Fire X4500 server packs up to 24 TB of disk storage and two processor sockets in the same 4U. Sun's x64 server product line currently uses energy-efficient AMD Opteron processors, and the M2 designation on many servers indicates that they can be upgraded from dual-core AMD Opteron processors to quad-core versions as they become available.
- *Sun Blade™ Modular Systems*
  Sun has taken the best ideas that blade systems have to offer and has incorporated them into a modular system product line that beats even rack-mount servers on price, compute density, processing power, energy consumption, and cooling. The use of single, shared power supplies, cooling, and I/O infrastructure contribute to the efficiency of these servers, while delivering mainframe-level RAS capabilities through the use of redundant, hot-swap components.

The Sun Blade 8000 Modular System offers the ability to host up to 80 AMD Opteron processor cores and 64 GB of memory in only 19 rack units. It provides the ability to configure on-board disk drives and a myriad of I/O expansion options. The Sun Blade 8000P server increases density even further, packing the same 10 Sun Blade x8400 server modules in only 10 rack units. The new Sun Blade 6000 Modular System provides even more flexibility, with the ability for customers to mix and match up to 10 UltraSPARC T1, AMD Opteron, and Intel Xeon processor-powered server modules in only 10 rack units.

- *Sun Customer Ready HPC Cluster, Sun Customer Ready Scalable Storage Cluster*
  In many datacenter environments, servers sit on the loading dock for days before they can be opened, configured, racked, powered, cabled, networked, and put to use to solve the problems they were purchased to solve. The Sun Customer Ready HPC Cluster and Scalable Storage Cluster discussed beginning on page 46, raises the granularity of deploying servers in datacenters from a system at a time to a rack at a time. With the ability to support custom configurations from 1U servers to modular systems, what was once only within reach of specialized compute farms can is now available for general-purpose use in enterprise datacenters.
- *Sun Netra™ Server Products*
  For the demanding environmental and DC power requirements of telco datacenters and central offices, Sun has designed its Netra server product line with a range of NEBS-compliant rack-mount servers and blade systems that incorporate UltraSPARC and AMD Opteron processors.

Datacenter Benefits

Sun's focus on developing a server product line to meet the needs of datacenter applications results in a compelling set of tangible benefits to customers choosing Sun:

- *Operating System Choice*
  On Sun x64 servers, customers have the choice of running all three enterprise operating systems: the Solaris OS, Microsoft Windows, and Linux, and with the right virtualization technology, they can run all three on the same server at once. For organizations needing to scale a single system beyond the eight processors sockets (16 cores) of the Sun Fire X4600 server, they can use the Solaris OS on Sun's Sun Fire and Sun SPARC Enterprise server product line on servers ranging from 1 to 72 processors.
- *Investment Protection*
  The ability to invest in a single pool of server infrastructure not only allows organizations to have a uniform platform to support a dynamic datacenter; it allows IT organizations to buy servers from a single vendor without knowing what OS they need to support. Customers can run Microsoft Windows or Linux on their servers one day, the Solaris OS the next.

- *Processor Choice*
  Sun's support for SPARC, AMD Opteron, and Intel Xeon processors allows IT organizations to choose the processor architecture that best matches their requirements. They can choose the low-power UltraSPARC T1 processor for highly threaded front-end Web servers and application servers, while they can choose the quad-core Intel Xeon 5300 processor to meet the most demanding HPC application requirements. Customers with many single-threaded applications, or large decision support systems or data warehouse environments most often choose highly scalable systems based on the UltraSPARC IV+ or SPARC VI processors.
- *Dialing In The -ilities*
  The datacenter space, power, and cooling crunch has taken some of the industry focus away from the traditional requirements of scalability, availability, reliability, manageability, and security. Sun, however, has not lost its focus, and its server product line allows customers to choose the balance of these characteristics they need, paying for what's required, but not paying for features they don't need.

  For example, consider a set of horizontally-scaled front-end Web servers. These servers provide performance and availability through their numbers, making the reliability of individual servers less important. For these applications, customers can choose servers with minimal internal storage and single power supplies. For mission-critical systems that must stay running no matter what, customers can choose high-end servers that allow on-line servicing while the system continues to operate. These servers support physical hot-swapping of not just power and cooling components, but also the addition and removal of CPU/memory and I/O boards.
- *Power, Cooling, and Density*
  Sun understands the power, space, and cooling crunch that today's IT organizations are facing, and has been addressing these issues through its server product line. To help alleviate the space crunch, Sun has created innovative blade servers, and high-density rack-mount servers such as the 4U, 8-socket Sun Fire X4600 server.

  Learn about how much power Sun has saved in its own datacenter through the Sun BluePrints article *Consolidating the Sun Store onto Sun FIre T2000 Servers*, available at http://*www.sun.com/blueprints*.

  Density alone, however, offers no benefit if it is accompanied by ever increasing power demands. In its UltraSPARC server product line, Sun has developed the low-power UltraSPARC T1 processor, which is available in the rack-mount Sun Fire T1000 and T2000 servers and on the Sun Blade T6300 server module that can be configured into the new Sun Blade 6000 Modular System. Sun's x64 server product line has been powered for years with power-efficient AMD Opteron processors that excel in providing high-performance with low power consumption.
- *Virtualization Support*
  Virtualization discussed in detail in Chapter 5, "Server Virtualization Technology" on page 30, is a key technology that enables consolidation, helping datacenters to further address the space, power, and cooling crunch. Virtualization also enables a dynamic datacenter where applications can be deployed onto a shared, virtual pool of server resources. As the discussion in Chapter 5 illustrated, Sun's support of

virtualization technologies range from hardware partitioning through Dynamic System Domains, to the highly-efficient, sub-CPU virtualization that can be accomplished with Logical Domains. One of the more powerful virtualization technologies, Solaris Containers, is available on any server running the Solaris Operating System, and Sun supports virtualization using Xen and VMware Infrastructure 3 software as well.

Sun Customer Ready HPC Cluster and Scalable Storage Cluster

More information on the Sun Customer Ready HPC Cluster is available at *http://www.sun.com/hpccluster.* More information on the Sun Customer Ready Scalable Storage Cluster is available at *http://www.sun.com/servers/cr/scalablestorage.*

The Sun Customer Ready HPC Cluster and Scalable Storage Cluster raise the level at which datacenters can deploy compute clusters from a server-at-a-time to a rack-at-a-time model, helping to speed the time from delivery to deployment. These products provide the server, storage, and network resources needed to deploy self-contained points of delivery that can help dynamic infrastructures grow at the speed of business.

The Sun Customer Ready HPC Cluster and Scalable Storage Clusters include servers, storage, networking, interconnects, and software racked, powered, and networked in a Sun Rack cabinet. The Sun Customer Ready HPC Cluster with Sun Fire servers is optimized for applications such as scientific research, Mechanical Computer-Aided Engineering (MCAE), Electronic Design Automation (EDA), financial analysis, and any other compute-intensive application. These configurations focus on delivering computing power with high space density, in a cluster or grid architecture.

The Sun Customer Ready Scalable Storage Cluster provides the data storage capacity and throughput that high-performance applications require. Key components of this data storage solution include high-performance Sun Fire servers, the Lustre scalable cluster file system, and high-speed InfiniBand cluster interconnects integrated into one system.

For more information on the Sun Customer Ready program, please visit *http://www.sun.com/customerready.*

All Sun Customer Ready HPC Clusters are configured and tested by the Sun Customer Ready program, which offers factory-integrated solutions that incorporate both Sun and third-party products. These factory -integrated systems help enable organizations to deploy IT solutions in a simpler, safer, and swifter manner. As a result, customers gain on productivity and save on deployment costs and risk.

The Sun Customer Ready program can help customers build out their data center by streamlining and accelerating deployment of customized configurations that can be used for applications ranging from cluster computing to Web services. By integrating and testing the configuration in Sun's factories, the Sun Customer Ready program can reduce on-site deployment times by up to 90-98 percent. Depending on the configuration, the Sun Customer Ready program can simplify the overhead-intensive receiving process by reducing the amount of physical materials shipped to the Customer's site by 95-99 percent, helping to eliminate unnecessary packaging waste and helping to minimize an organization’s environmental impact. The Sun Customer Ready program also can simplify the vendor management process by integrating both Sun products and third-party products, acting as the one-stop-shop for all configuration requirements.

To further accelerate datacenter build-out, the Sun Customer Ready program offers a co-planning service that can reduce the total time experienced from purchase order issuance to materials receipt in some cases to less than a week. By collaborating on shared forecasts and establishing a catalog of configurations, Sun can respond very quickly to specific customer requirements, resulting in greater datacenter efficiencies.

Chapter 7
Sun Storage Solutions

While most IT organizations consider their server, storage, and network infrastructure components to be their greatest assets, those assets depreciate and must be refreshed continually with newer technologies. The only asset that appreciates in a datacenter is the organization's data. For most companies, their value is represented by their data, for it includes everything from engineering drawings and software source code to the information on which their ERP and CRM systems run.

The loss of data, even the temporary loss of access to data, is a risk whose consequences can be quantified by the business lost while e-commerce capabilities are unavailable, while production ceases, while the supply chain is disrupted, and while employees are unproductive.

To mitigate this threat, IT organizations consider characteristics such as storage system reliability, availability, and scalability as they do with their server infrastructure. Indeed, many IT organizations protect their data with remote replication backed up with disaster-recovery plans, while they make frequent storage snapshot-based backups knowing that there is no substitute for a stable, offline, offsite copy of their business-critical data.

Sun's storage solutions offer IT organizations a range of storage system features that can be used to complement its server technology. Every application is different, and every application's data needs are different. Sun offers a range of disk storage and tape products that allow IT organizations to best match their storage policies, best practices, and application requirements with storage products that best meet their needs. Sun's enterprise storage systems support centralized storage with resource management, virtualization, and business continuity support. Sun's modular storage components offer pay-as-you-grow scalability for data that does not require the central storage and management of enterprise systems. And network-attached storage supports multiple protocols and data characteristics with strategies that support mission-critical, business-critical, and task-critical availability levels. Sun's disk storage products are complemented with tape backup systems that range from enterprise-level tape libraries to server and desktop-hosted tape drives.

Solving Business Problems

IT organizations must face increasingly difficult decisions around their data. The *Participation Age* is here, and increasing numbers of people and devices are connecting to the Internet for the first time. By some estimates, a million new subscribers join the Internet daily; all of them creating data and contributing to the high Compound Annual

Growth Rate (CAGR) of data and the storage hardware that must be purchased to support it.

Compounding the problem is the number of regulations that require organizations to retain data for longer periods of time and to ensure it is disposed of in a secure fashion with complete audit trials. These regulations include Sarbanes-Oxley (SOX), HiPPA, Securities and Exchange Commission (SEC 17a), Basel II and many others. All this data must be maintained according to governmental regulations and business requirements. And all of this is in an environment where IT organizations are facing space, power, and cooling constraints, with the inevitability of mandated limits on energy footprint and carbon dioxide emissions.

The key storage-related problems facing today's IT organizations can be summarized as follows:

- CAGR of data and storage assets
- Information lifecycle management and archiving
- Regulations
- Environment including real estate, footprint, power consumption and cooling

Sun Storage Solutions

Any storage solution that an IT organization deploys must implement solutions to the above issues. Fortunately, most problems can be addressed by implementing architectures and solutions that are complimentary. Sun uses a systems approach that combines servers, software, services and storage to address customers business issues.

Storage Virtualization

Just as with their servers, IT organizations can help reduce their power consumption through storage consolidation, reducing their space, power, and cooling requirement as well as their carbon footprint. Deploying an enterprise class storage system such as the Sun StorageTek 9990V System enables customers to pool and virtualize heterogeneous storage behind the same subsytem. For example, the Sun StorageTek 9990V System can be used as a front-end system to virtualize storage from Sun and third parties. The Sun StorageTek 6540 array, with high-capacity SATA drives, can be used in such a configuration as a lower cost tier in a tiered storage architecture.

Integrating Heterogeneous Servers and Storage Systems

Once virtualized, the storage can be accessed by a wide variety of operating systems. The internal and external storage is managed by the same software, offering performance and tuning tools that reduce the complexity of controlling many islands of storage. The Sun StorageTek 9900 System also offers the highest level of performance and availability required for mission-critical environments. To help achieve maximum utilization, and access to the data on both a block and filesystem level, a Network Attached Storage (NAS) gateway can be added to the architecture.

The Storage Value Platform (SVP) is tool and methodology that helps organizations rapidly identify their storage assets and highlights investment opportunities to optimize the storage environment. This helps organizations to better understand the Return On Investment (ROI) for any solution being proposed, and it helps to reduce total cost of ownership (TCO). When coupled with a storage assessment, the SVP also can provide architectural recommendations and consolidation possibilities.

Business Continuity and Disaster Recovery

Once an IT organization consolidates its storage, it needs to establish a business continuity and disaster recovery infrastructure to match the criticality of the data. There is a distinct difference between business continuity and disaster recovery. A disaster recovery plan is a reactive process that helps an organization return to operational service, while a business continuity plan is a proactive set of architectures and services designed to keep the business in an operational state when a disaster strikes. Various governmental regulations mandate effective business continuity planning to protect stakeholders, investors, customers and employers. In a business continuity plan, critical data would require an active/active solution while non-critical data might require a backup and restore procedure. Sun can help organizations evaluate their data to determine its importance to the business and help match the protection architecture to the importance of the data.

The Sun Business Impact Analysis Service helps IT organizations identify the operational and financial impacts of a business disruption. In addition, the Sun Disaster Recovery Plan Audit Service analyzes the customers business continuance plans and procedures and suggest improvements.

Information Lifecycle Management

The constant growth in data that an organization is required to retain has resulted in an explosive growth in storage budgets and the space, power, and cooling required to maintain it. Organizations need to have retention policies that are in line with their business requirements and governmental regulations. For data that must be retained beyond the time when it is in active use, offline storage should be considered.

The solution to this is to implement an Information Lifecycle Management (ILM) solution. An ILM solution puts the right data on the right storage tier to offer the best performance and economy for an organizations entire data and its lifecycle. Data has a value attached to it when it is created, and over time this value changes. It is therefore not always necessary to keep legacy data (which is very rarely accessed) on tier 1 storage. Instead, this data can move to the next tier, helping to reduce cost while still meeting performance requirements.

If an organization must manage its data according to governmental regulations, then the appropriate third-party software can help maintain the correct retention policies associated with the data. The Sun StorageTek 9990V System offers this functionality in the form of its Tiered Storage Manager that migrate volumes from one tier to the next. An alternative to tiered storage is Sun StorageTek SAM-FS software which offers the ability to migrate data on a per-file basis. This approach supports application-transparent file migration.

Sun's Information Management Maturity Model (IM3) helps customers discover where they are in their overall ILM strategy, compares them with their industry peers and identifies areas for improvement. The value to IT organizations is that it focuses in on areas of over and under-investment to help drive down cost, reduce risk, and provide an infrastructure that is both integrated and supports the companies business objectives. This enables IT assets to be used as an integral part of generating new revenue streams.

Sun’s storage-related services advise customers resolve the business issues that affect them on a daily basis. Sun’s systems approach to storage brings together its four practices of Servers, Storage, Software, and Services to deliver on its recommendations and provide solutions that effectively interoperate in heterogeneous environments.

Disk Storage

Sun splits its disk storage portfolio into three categories based on a range of market requirements: mission critical, business critical and task critical. The mission-critical category defines environments where there are numerous servers with mixed workloads that require continuous availability. Sun’s enterprise storage products meet the requirements of this category. The business-critical category is for environments where there are multiple servers running multiple workloads needing high availability and data services. The task-critical category is for environments with several servers running single workloads requiring high availability. Sun’s modular storage portfolio covers the business-critical and task-critical categories.

Enterprise Storage

The enterprise class product in Sun’s portfolio that best manages storage requirements for large IT organizations today is the Sun StorageTek 9990V System. The system is built with a third-generation crossbar switch that provides high performance and scalability. The crossbar architecture support performance of up to3.5 million I/O Operations Per Second (IOPS) with a throughput of up to 106 GBps. In addition, the Sun StorageTek 9990V System can scale up to 336 TB per subsystem with the ability to virtualize and manage up to 247 PB of external storage behind it. It has availability and reliability that can be supported with an optional data availability guarantee.

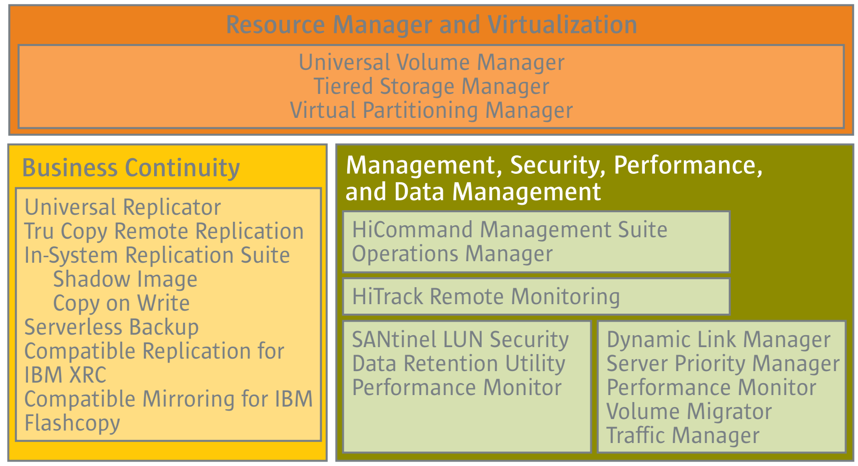

Figure 19. The hardware features of the Sun StorageTek 9990V System are complemented with three categories of software products that extend its capabilities even further.

In addition to its hardware capabilities, the Sun StorageTek 9990V System is complemented by a comprehensive portfolio of optional software applications, which are grouped into three categories (Figure 19). These software solutions operate with internal storage as well as virtualized external storage, giving IT organizations the choice of a single suite of software for all of their storage types.

Resource Manager and Virtualization

Virtualization helps IT organizations increase their storage system utilization while helping to protect their investment. It also drives down the cost by having a common management for mixed systems and masks the complexity.

- *Universal Volume Manager* software helps organizations to virtualize their existing storage behind the Sun StorageTek 9990V System.
- *Tiered Storage Manager* supports policy-driven volume migration that helps to put data on the right storage tier, both internal or external, to match an organization's access, performance and cost requirements.
- *Virtual Partitioning Manager* allows organizations to dedicate specific resources to an application, department or even a mainframe platform to meet service level agreements requiring specific response times and quality of service.

Business Continuity

Sun's Business Continuity software suite allows organizations to replicate data within the storage system using In-System Replication software suite, and externally using either Universal Replicator or Sun StorEdge™ TrueCopy software-based remote replication:

- *In-System Replication* uses high-speed, non-disruptive local mirroring to rapidly create multiple copies of mission-critical data in both mainframe and open-systems environments.
- *Copy on Write* can be used in connection with both Sun StorEdge ShadowImage and TrueCopy software. By copying only changed data, replicas can be created faster and use significantly less space.
- *Sun StorEdge 9900 TrueCopy* software performs remote replication, automating the data recovery process so that in case of an interruption, business operations can resume in a matter of seconds.
- *Universal Replicator* provides enterprise-class heterogeneous data replication for significant improvement to business continuity requirements.

System Management, Performance and Security

System Management, Performance and Security tools are critical when managing heterogeneous storage in a multi-vendor host environment. The management tools support any storage product that complies with the Storage Management Initiative Specification (SMI-S) standard set by the Storage Networking Industry Association (SNIA).

- *Device Manager* software allows organizations to manage their storage including any other storage products that comply with the Storage Management Initiative Specification (SMI-S) from a single, central platform.
- *The Performance Enhancement Suite* consists of a number of software packages that help to increase performance service levels for business-critical on-line transaction processing and decision support system applications.

Modular storage

Not all data requires enterprise-class storage, and modular storage offers an excellent alternative in many cases. Modular storage offers excellent performance for transactional applications, it is extremely flexible with pay-as-you-grow scalability, it has a small footprint, and can be managed easily with an intuitive interface.

Sun's modular storage portfolio includes the Sun StorageTek 6540 and Sun StorageTek 6140 arrays. These arrays use a switched-loop architecture enabling detailed diagnostics and analysis with multiple point-to-point communications, allowing the full compliment of disk drives to be utilized concurrently with excellent performance.

Network Attach Storage (NAS)

Network attached storage can also be divided into the three categories of mission-critical, business-critical and task-critical functionality. File-based sharing of data is another cost-effective way to place the right data, on the right storage platform, at the right performance and cost by utilizing NFS, CIFS or FTP protocols. Sun StorageTek 5300 NAS appliances are based on the fast, reliable, and energy-efficient Sun x64 server product line and can be configured either as dedicated NAS storage with its own captive storage, or as a gateway which utilizes Storage Area Network (SAN) based storage. The two models available are the Sun StorageTek 5220 NAS appliance for task-critical environments and the Sun StorageTek 5320 NAS appliance for business-critical environments. Mission-critical environments are serviced by clustering the ST5320 or utilizing a NAS gateway.

Tape

There is no substitute for having stable, offline backups in the form of tape. Tape not only protects data due to its stability and transportability; it also helps reduce datacenter energy costs by storing legacy data that is not frequently accessed. Once data is stored to tape, power costs are no longer a factor as they are with disk. Tape is also a key part of an ILM solution, providing the right price, access and performance for most legacy data. The archived data is stored in compliance with governmental regulations and is easily retrievable. Tape provides enormous capacity at a very low cost, which is why the majority of the world's data is stored on tape. Tape is not only used for backup, there are a multitude of industries and functions that utilize tape. Archiving, hierarchical storage management, document management, common repositories, and backup and disaster recovery are just some of its many uses.

Today's Intelligent Tape technology is one of the primary reasons that tape remains such a highly-viable storage mechanism. Intelligent Tape supports new formats on existing cartridges, high density and transfer rates, shared robotics, physical and logical WORM and encryption in the drive. Some of Sun's tape products allow organizations to encrypt their data in order to prevent data loss due to the loss of a tape cartridge.

As with disk drives, organizations can choose the trade-offs between performance or capacity that best suits each use, and they can install them into tape libraries that offer investment protection through the ability to support multiple drive and cartridge types available today and into the future.

- *Enterprise tape drives* and their associated cartridges are designed to have read/write activity of up to 17 hours per day, and handle the load/unload activity of a busy tape library.
- *Mid-Range tape drives* are most useful for tape libraries that are primarily used for backup operations that are not continuous throughput the day.

Tape Libraries

Sun provides a range of choices for automated tape backup in the form of tape libraries that can support task-critical, business-critical, and mission-critical operation. Two particularly important innovations are the Sun StorageTek StreamLine™ SL8500 modular library system and the Sun StorageTek L1400 tape library.

These tape library systems have been designed to accommodate both heterogeneous tape drives and media. The ability to store any cartridge in any slot enables these libraries to store a LTO, DLT, T9840, T9940 or a Sun StorageTek T10000 tape drive cartridge in any slot. Library investments are protected because they can handle new tape technologies as they come to market. The Sun StorageTek SL500 library provides the same heterogeneous drive technology support for Super DLT and LTO form factor drives, however it does not accommodate the T-Series tape drives.

Sun StorageTek tape libraries have a specific design that separate the data paths (tape drives) and the command path (automated tape library commands). This architecture provides some of the best flexibility in the market to handle multiple tape drive technologies and to share the automated tape libraries between different applications using the Sun StorageTek ACSLS Manager software

Virtual Tape Libraries

The Sun StorageTek Virtual Storage Manager (VSM) demonstrates Sun's leadership in virtual tape solutions. Sun StorageTek VSM provides features required by IT organizations that need to efficiently manage tape cartridges in mainframe environments. VSM supports multiple tape drive technologies including all T-Series tape drives, providing a fast access path for critical data with Sun StorageTek T9840C and T10000 tape drives to store data long term, and optionally provide application and operating system-transparent encryption. VSM helps the customer to manage the media and tape drives more efficiently and is a key component to simplify migration technology.

The Sun StorageTek Virtual Tape Library Plus (VTL Plus) Storage Appliance is designed for open system environments needing an effective disk-to-disk-to-tape backup architecture. VTL Plus provides a high-speed transparent cache ahead of the physical tape, helping to speed backups and restores with throughput of up to 700 MBps. The system allows open systems to perform backups at disk speed with RAID reliability.

Innovative Storage Products

Sun continuously innovates in the storage market bringing unique storage solutions to common business problems.

Object-Oriented, Programmable Storage

The Sun StorageTek 5800 System provides application-aware, object-oriented, programmable storage. The system uses a unique design that combines low-cost hardware and a symmetrical cluster architecture to make it affordable yet reliable as an enterprise class system. It can withstand multiple disk and compute node failures with automatic failback. Organizations managing fixed and unstructured data such as medical imaging and broadcast material, and need it to be archived, are ideal candidates for this solution.

The object storage system has an extensible metadata layer that can be queried via an open API or scripting language, making application integration straightforward. For example, when storing photographs, each object (the photograph JPEG file) is stored and a unique object ID returned to the application. Software on the storage system then strips the photographs metadata (including resolution, date, time, lens, camera) and associates the metadata with the object. The data can then be retrieved either by the unique object ID or by browsing or searching the metadata. The value to organizations is moving the intelligence out of the application and putting it into a programmable storage system where it can be accessed by a number of different applications. The Sun StorageTek 5800 System is designed to scale virtually seamlessly, allowing an administrator to manage more than a petabyte of data behind a single management interface. The system can start small and scale over time into the petabyte range, while meeting the organization's throughput, utilization, and density needs. When the system is expanded, it automatically spreads existing data across newly added nodes, so capacity and workload stay evenly distributed across all nodes and disks.

Merging Servers and Storage

The Sun Fire x4500 server is an innovative product that bridges the gap between storage and server technology. Instead of separating processing from storage and connecting them through a relatively low-speed SAN connection, the Sun Fire x4500 server provides 24 TB of storage with up to two dual-core AMD Opteron processors and 16 GB of main memory in a compact, 4U package. Key applications for this system include digital video surveillance, digital media streaming, business intelligence and data analytics, high performance computing clusters, and mainframe batch processing. Used as a storage device, the Sun Fire x4500 server can be configured to take on different storage personalities including NAS, object storage, iSCSI target, Fibre Channel target, or VTL.

Chapter 8
Disaster Recovery

Virtually every business depends on its information technology, and disaster-recovery plans are no longer optional, they're required. Regardless of an organization's size, the ability to recover from a physical or even an IT disaster is a business requirement whose value can be measured by the amount of business lost in the event of a service outage. Indeed, the impact of losing a securities-trading system that manages the transfer of millions of dollars per hour can be quite high. Disaster recovery is not only a business requirement, it's also a legal requirement. Some organizations must adhere to governmental regulations that dictate service levels to which they must adhere. At minimum, every business must be able to protect its data for which it has record-retention requirements.

In order to maintain specified service levels in the event of a disaster, a secondary IT infrastructure must be available at a site sufficiently remote from the primary site so that it is not affected by the same disaster. The time to recover and restart the service is then the sum of the time it takes to decide to enact the disaster recovery plan and switch to remote site operations plus the time it takes for the service to be brought up on the system at the secondary site. The time that an organization targets for recovery from a disaster is known as a *recovery time objective*. The point in time from which operational data was frozen and from which operations can resume again (and beyond which any transactions or updates are lost) is known as a *recovery point objective*.

For less critical services where data is not continuously replicated to the secondary site, the time to bring up the service includes the additional time component of restoring data from tape or a point-in-time disk copy. Synchronous replication minimizes the recovery point objective because all transactions are continuously updated to a remote site. Tape backup as a sole disaster-recovery strategy can result in a recovery point objective of hours or days, as services are restarted from the last point-in-time backup.

Regardless of the technique employed and the acceptable recovery time and recovery point objectives, it is important to monitor and address any changes to the service and, in particular, to address any changes in data volume. If, for example, the primary site allocates more storage, but this change is not propagated to the secondary site, this mismatch can cause the service to fail to start on the secondary site in a disaster situation.

Backup And Recovery From Tape

Although RAID technology is designed to provide high levels of data availability through redundant components, it cannot protect against data loss if data is intentionally or accidentally deleted or corrupted. Regular backups to tape or virtual tape, such as a

Sun StorageTek Virtual Tape Library, helps to ensure that data can be recovered as quickly as possible to the latest possible point in time.

Although backup and recovery from tape is a robust approach to creating usable, point-in-time copy of the data, it is only applicable as the primary disaster recovery mechanism for non-critical services or data.

Data Replication Infrastructure

A key to disaster-recovery strategies that minimize both recovery time and recovery point objectives is data replication. A wide range of techniques are available for creating copies of application data at one or more remote sites. The appropriate mechanism for a particular application or service depends on the service objectives weighed against the cost of meeting these objectives. Most of these mechanisms require some level of permanent communication infrastructure between the primary and secondary site(s).

Secondary sites must be located a sufficient distance away so that they are unaffected by disasters that might strike the primary site. The distance can vary from a few kilometers to several thousand kilometers, depending on business needs, budgetary constraints, regulatory requirements, and availability of suitable facilities. Most organizations have a basic IP network infrastructure connecting the facilities and some may also have a an extended Fibre Channel SAN that connects the two sites. The first objective is to create an infrastructure that is resilient, highly available, and secure. This requires the core LAN or SAN to have redundant components on each site connected through links that follow independent paths, avoiding single points of failure. As Figure 20 illustrates, data can be replicated using host-based and storage-based replication technology.

Host-Based Replication

Host-based replication uses host-based software to replicate data from one site to another site. Remote mirroring software, such as Sun StorageTek Availability Suite software, operates at the Solaris OS kernel level to intercept writes to underlying logical devices as well as to physical devices, such as disk slices and hardware RAID protected LUNs. It then forwards these writes on to one or more remote Solaris OS-based nodes connected through an IP-based network.

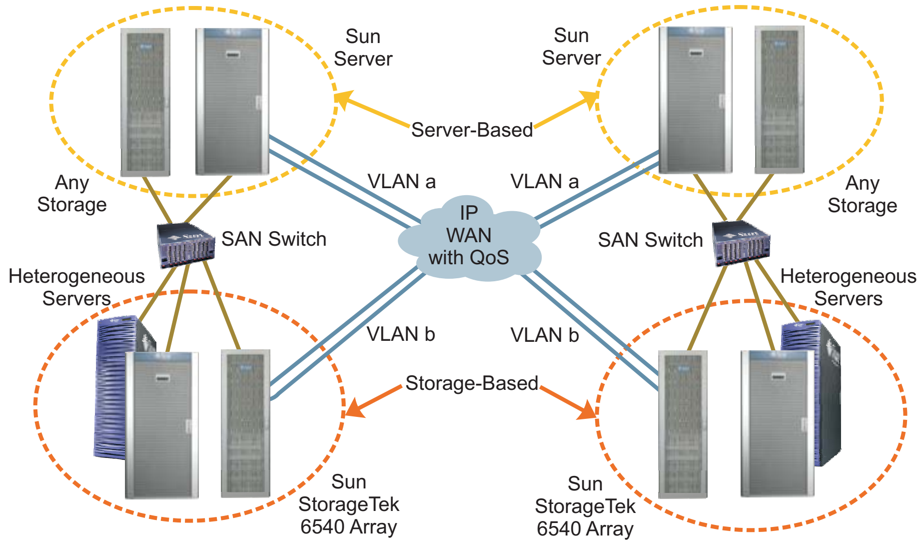

Figure 20. Remote replication requires highly available, secure interconnects between primary and secondary sites, and it may use server or storage-based technology.

Storage-Based Data Replication

Storage-based replication uses features of the storage systems themselves to replicate data to remote sites. The Sun StorEdge 6000 product family and the Sun StorageTek 9900V System perform data replication using the CPUs or controllers resident in the storage systems themselves. The functionality of both Sun StorageTek Data Replicator software and Sun StorageTek 9900 TrueCopy software is similar to that of Sun StorageTek Availability Suite software, but the software operates at a much lower level (Figure 21). The software features are also independent of the operating systems on the connected servers, allowing a single replication technique to be deployed in a heterogeneous environment and thus helping to lower management and administrative overhead.

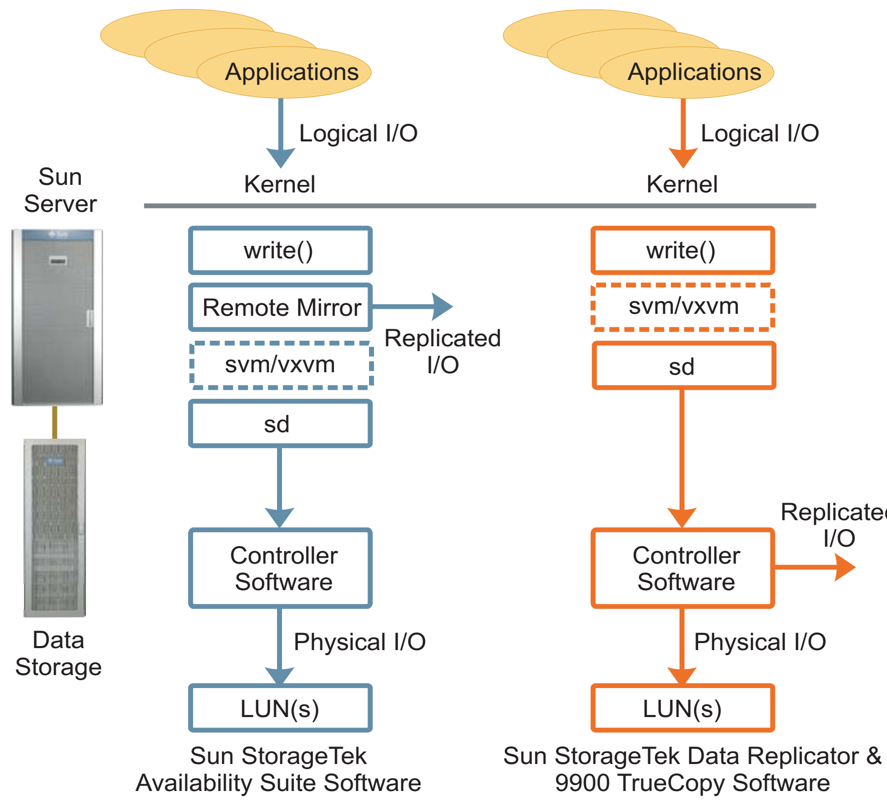

Figure 21. Data can be replicated in real time using host-based and storage-system based software.

Network Considerations

It is very important for IP networks and SANs to be properly implemented and configured for any solution that supports data replication.

In a LAN environment, traffic is typically separated using Virtual LANS (VLANs). One or more dedicated VLANs should be configured to carry IP data replication traffic. Each VLAN should be prioritized with a quality-of-service guarantee.

For host-based data replication, resilient network connections can be configured with a Solaris IP Multi- Path (IPMP) group that connects to separate network switches in each datacenter. For storage-based data replication, such as that provided by the Sun StorageTek 6920 System, the storage system is connected to the LAN by two connections from separate storage controllers (SIO) cards.

When the storage-based replication uses Fibre Channel connectivity, rather than an IP network, the outbound Fibre Channel connection should connect to an independent SAN switch to help ensure resilience and be configured into zone for security (Figure 22). If the SAN is extended more than 10 kilometers, sufficient buffer credits must be allocated to the links carrying the replication traffic.

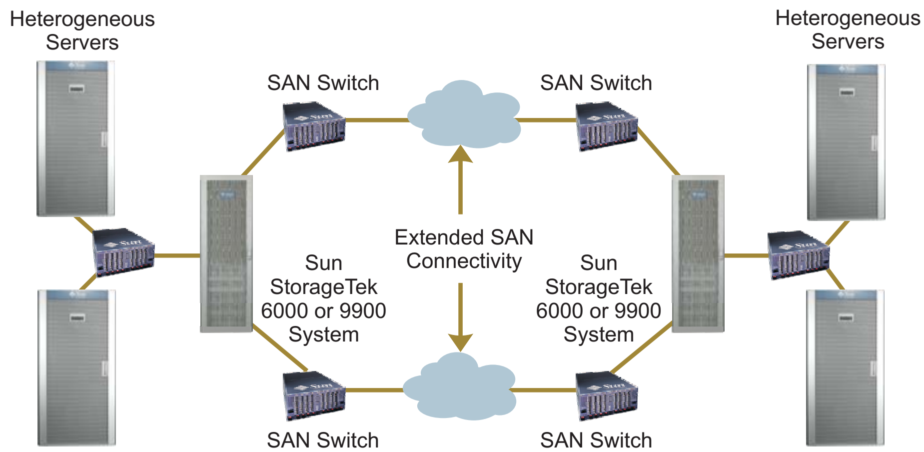

Figure 22. Storage-based replication can be accomplished using extended SAN connectivity.

Disaster recovery is a complex problem. Sun provides key software components that contribute to a solution, along with expertise that can integrate these components to solve individual customers' unique needs.

Chapter 9
Systemic Security

In the rush to deploy ever more applications, grapple with the space, power, and cooling crunch, and manage the rapid growth of business-critical data, the focus on the intrinsic properties of datacenter infrastructure has taken a back seat in many organizations — but not at Sun. Sun has maintained a constant focus on improving the intrinsic properties of its products, sometimes referred as the "*-ilities*," producing products and methodologies to deliver characteristics including scalability, availability, manageability, and security in all of its products and services.

Security is one such intrinsic property of a datacenter architecture. Security involves everything ranging from physical datacenter access to authentication and auditing. This chapter addresses the issues of protecting servers within the datacenter environment.

Managing Risk With Systemic Security

Sun uses a systemic model to address security. Systemic security manages risk using sound, preventive mechanisms. This approach uses architectural building blocks and patterns to build security into each step of the process. These building blocks leverage time-tested security principles that are sometimes applied in unconventional ways to reap greater levels of security than otherwise thought possible.

Systemic security helps to capture, reuse, and refine an organization's knowledge of IT infrastructure, processes and applications, as well as to better understand the interrelationships between the various building blocks and their underlying components. In the end, systemic security allows organizations to determine which patterns may be appropriate for a particular situation given a set of requirements, dependencies, and constraints. A good security architecture is required, is not sufficient a security strategy. Instead, a continuous improvement methodology, based on a security model, helps organizations to realize greater levels of integration, efficiency, and alignment with business goals.

Secure By Default

Datacenter architectures must guard information privacy, confidentiality, and data integrity by protecting against unauthorized access to, or tampering with, an organization's data. The Solaris 10 Operating System has been designed to provide a high level of security out of the box. When directed at install time, the Solaris OS activates only the minimal set of services that allow the system to run, eliminating a wide range of potential attack vectors before they can be exploited.

Security and Availability

Security and availability are closely related. Datacenter architectures are designed to provide high levels of availability. Security remediation, however, can significantly affect availability. Examining, fixing, testing, and patching systems can require systems to be taken offline periodically. To preserve high levels of availability, an architecture needs to minimize the impact of an individual system being unavailable due to security-related activities.

Best practices include security as a pervasive quality throughout a design. They leverage multiple, independent, different, mutually reinforcing security mechanisms to prevent unauthorized access to networks, systems and data, and to limit denial-of-service attacks. They are designed to provide multiple layers of security to satisfy most stringent security policy requirements of today's enterprises.

Security Model Components

A datacenter security model incorporates a set of components into the datacenter architecture as part of a comprehensive security framework. For example, an infrastructure stack of components addresses the datacenter architecture as a whole, while an operating system stack addresses issues that must be managed on a per-OS instance level. Some of the more commonly encountered components are illustrated in Figure 23 and are discussed below. These examples only represent a portion of a complete security solution. Layering, or defense-in-depth, integrates these stack components.

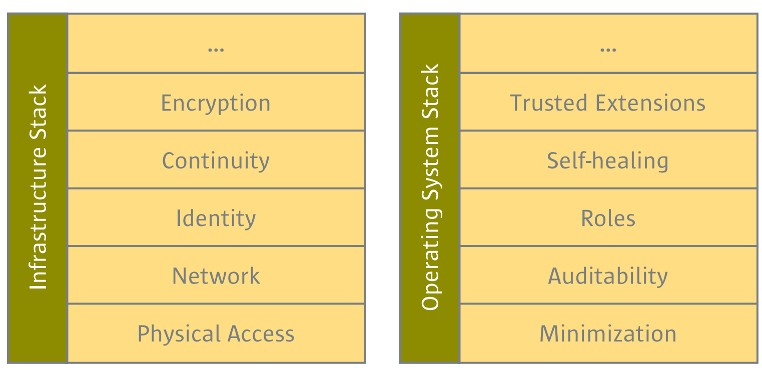

Figure 23. A defense-in-depth strategy integrates components within infrastructure and operating system stacks.

Infrastructure Stack

Integrating the infrastructure stack into datacenter design requires addressing issues including physical access, network infrastructure, identity management, business continuity, and encryption:

- *Physical security*
  Physical access to the systems must be controlled. Untrained or malicious personnel can power off, corrupt, or even remove systems or disk drives if they are given access. Equipment racks should be clearly labeled and secured within the datacenter as an extra precaution against accidental service disruptions.
- *Identity Management*
  Knowing the identity of all system users, from employees to customers, is mandatory. A single user may exist in multiple domains, such as development, testing and production, and have different levels of access in each domain. Trying to maintain consistent instances of usernames and passwords manually across domains is difficult. A unified identity management schema ensures that an individual retains the same identity across domains.
- *Network*
  A secure datacenter architecture uses a network design that isolates services to manage security and scalability requirements (Figure 24). Only public, externally accessible interfaces are on routable networks. All other networks are configured as private and are non-routable outside the datacenter.

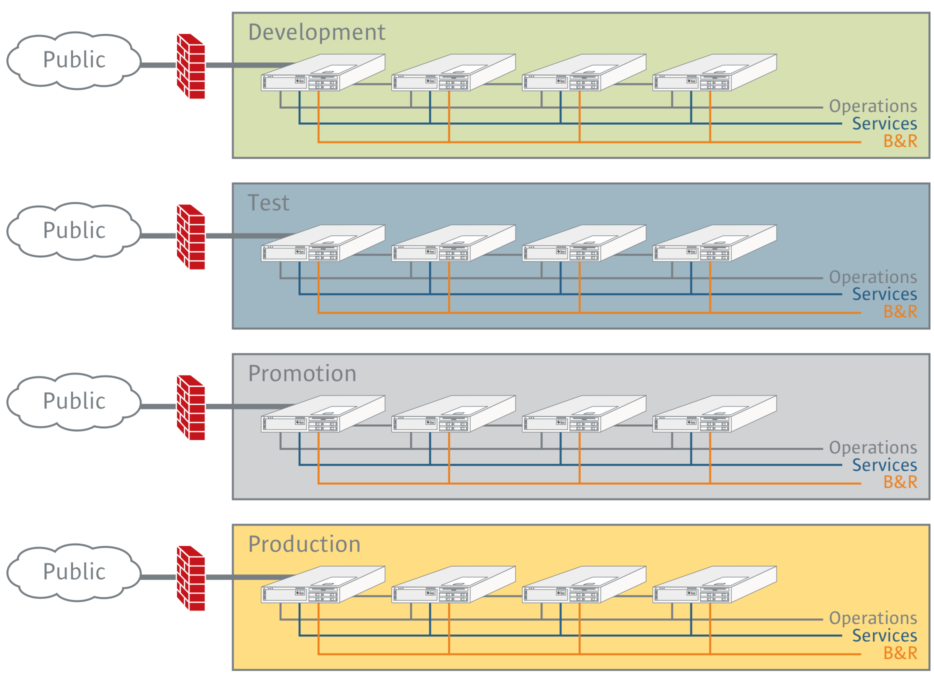

Figure 24. Good datacenter network design isolates networks that are exposed to public networks.

- *Data*
  Storage Area Networks can utilize protection techniques such as SAN zones and LUN masking. This assures that data is not available outside designated network or server parameters
- *Management services*
  Management access to each server console and other similar devices is limited to connections from a designated management network.
- *Service Continuity*
  As discussed in the previous chapter, system availability, disaster recovery, and business continuity may be mandated by government regulations, and must be addressed as part of an organization's security policy.
- *Encryption*
  Encryption is a mechanism that helps keep confidential data secure, as well as helping to protect against man-in-the-middle attacks. Data my be encrypted by servers, disks, and network devices. The Solaris Cryptographic Framework performs cryptographic functions in the OS itself, making the use of optional encryption-acceleration hardware completely transparent to applications.
- *Anti-Virus*
  A systemic approach to security includes the use of proxies to scan and remove malicious content from network traffic including file transfers, Web browsing, and e-mail. Using anti-virus technology is important to protect internal desktop systems from attack, especially those running operating systems less secure than the Solaris OS or Linux.

Operating System Stack

Security throughout the entire operating system stack is another example of systemic security in the datacenter environment.

- *Minimization*
  Standardized Solaris OS implementations start with a secure-by-default configuration that protects the OS beginning with the first boot. The OS configurations must meet functional requirements as well, and some customers may require additional OS services to be enabled. IT organizations must weigh the usability of the system against the cost of the additional workload and risk.
- *Solaris Security Toolkit*
  The Solaris Security Toolkit helps to automatically harden Solaris OS installations according to business requirements and make them easy to reproduce. In conjunction with the Service Management Facility (SMF), a minimized set of services can be enabled, which prevents enabling unnecessary services during installation.

  Using automated processes to build systems removes the possibility of operator errors and reduces the ability to make malicious changes to a system. A "golden master" installation instance, whether on physical media or available over a network, can provide an exact starting point for each system being configured. The Solaris Security Toolkit also supports validating network resources to limit an administrator's ability to install an incorrect image.

- *Audit*
  The Solaris OS provides an extensive suite of tools to audit system functions. By default, these are not enabled, but may be configured as part of a security policy that requires the ability to audit enabled OS services.
- *Role-Based Access Control*
  Role-Based Access Control (RBAC) creates specific roles that can be assumed by users to perform tasks. It allows administrators to define roles, the tasks they can perform, and the users that can assume each role. This improves upon the previous all-or-nothing model of root access. In addition, log entries are created when users assume roles, making auditing simpler and more secure.
- *BART*
  The Solaris OS includes the Basic Auditing and Reporting Toolkit (BART). BART can be configured to monitor critical files for changes. Using standardized installation procedures from the Solaris Security Toolkit allows a common BART manifest to be used to monitor many systems at once. Changes in any file attribute will create an alert that can be responded to by administrators. It is commonly used in conjunction with the Solaris fingerprint database to verify file integrity
- *Trusted Extensions*
  With the introduction of the Solaris 10 OS, features that were formerly part of the Trusted Solaris OS can now be added to any Solaris OS installation. The set of Trusted Extensions (TX) help to establish multiple levels of security. The flexibility of the Trusted Extensions allows for more than Unclassified, Secret, and Top Secret levels. Businesses can create their own levels such as Public, Internal Only, and Finance to segregate information according to users, their roles, and sensitivity.
- *Predictive Self-Healing*
  The Solaris 10 OS includes software to monitor the health of system components. The software can detect impending failures and direct an administrator to take certain actions to prevent or reduce unplanned system outages.

Additional Resources

Sun provides a wealth of solution briefs, whitepapers, Sun BluePrints articles, and online resources that provide more information on the components described above:

- *Systemic Security* describes time-tested security principles, architectural patterns, and iterative refinement processes. Please refer to: *http://www.sun.com/blueprints/0206/819-5605.pdf.*

- *Privilege Bracketing* describes the Solaris OS features that support least privilege access to programs, Please refer to: *http://www.sun.com/blueprints/0406/819-6320.pdf.*
- The *Solaris Fingerprint Database* describes a security tool that verifies the integrity of files in the Solaris Operating System. The second document includes the tools to configure BART. Please refer to: *http://www.sun.com/blueprints/0306/816-1148.pdf* and *http://www.sun.com/blueprints/0405/819-2260.pdf.*
- The *Solaris Security Toolkit* provides documents and scripts to harden and audit Solaris Operating Environments. Please refer to: *http://www.sun.com/software/security/jass/.*

Additional collections are available at:

- *http://www.sun.com/software/security/blueprints/index.xml*
- *http://www.sun.com/bigadmin/collections/security.html*
- *http://www.sun.com/security/whitepapers.jsp*
- *http://www.sun.com/security*
- *http://www.opensolaris.org/os/community/security/*

Chapter 10
System and Server Management

The push towards server consolidation through virtualization is likely to decrease the number of physical servers that an IT organization must manage while increasing the number of operating system and application instances they must support. This trend will make system and server management even more important than it is today. Large server platforms that serve as virtualization platforms will contribute more components necessary for service delivery, making their uptime even more important. The number of operating system and application instances will become even more of a challenge to manage, making the use of automated tools a necessity.

System management can be characterized by two distinct types of activities:

- Reactive management that includes system monitoring and response to problems
- Active management, including provisioning, change control, and process linkage.

Sun addresses both of these system and server management requirements through its Sun N1 software products, Sun Connection services, and its strict adherence to open system management interfaces.

Indeed, Sun products can be integrated with third-party hardware and software products to provide additional management capabilities. Sun products can interoperate with third-party products including HP Open View and Insight Manager, IBM Tivoli, BMC Patrol, and CA Unicenter. Sun's server and system management products provide best-of-breed management for Solaris OS and Linux-based environments.

System Management Products

Sun's system management products can help IT organizations automate every step of the system lifecycle, from bare-metal provisioning to automated patch management.

N1 System Manager

Sun N1 System Manager software provides comprehensive infrastructure lifecycle management capabilities for discovering, provisioning, monitoring, updating and managing SPARC processor-powered servers running the Solaris OS and x64 servers running the Solaris, Linux, and Microsoft Windows operating systems. Sun N1 System Manager software's unique and innovative hybrid GUI and CLI user interface gives IT organizations the choice and flexibility to efficiently manage systems with maximum ease of use from a single central console. Sun N1 System Manager software provides capabilities including the following:

- *Discovery*. The software automatically discovers bare-metal, single-domain systems based on subnet or IP address.

- *Grouping*. IT organizations can logically group systems together and perform management functions across a group of systems as easily as they can a single system.
- *Provisioning*. Using a centralized Web console, operators can remotely install operating systems, firmware, software packages, and patches onto selected systems to help reduce deployment time.
- *Monitoring*. Sun N1 System Manager software monitors hardware components (fans, temperature, disk, and voltage) and operating system attributes (CPU, memory, swap space, and filesystem) based on operating thresholds defined by the IT organization.
- *Event Logging and Notification*. The software records all events for the monitored components and alerts operators via pager, e-mail, and SNMP traps when the defined thresholds are exceeded.
- *Lights-Out Management*. Administrators can remotely power servers on and off, execute commands on remote systems via `ssh`, and access the serial console on remote systems.

Sun N1 Service Provisioning System

Sun N1 Service Provisioning System automates application lifecycle management across heterogeneous OS and server environments. IT organizations can provision the entire software stack, from operating system to application service, from a single Web-based user interface. The software eliminates many time-consuming and often error-prone manual administrative tasks, helping to increase efficiency and accelerate business service delivery. The use of Sun N1 Service Provisioning System can help raise application availability and consistency, increase compliance, reduce application deployment time, and lower operating costs.

Sun N1 Service Provisioning System provides the following key features:

- *Bare Metal OS Provisioning*. N1 Service Provisioning system can remotely install and configure operating systems onto specified systems or it can do so using Sun N1 System Manager software.
- *Application Provisioning*. The software provisions applications running on Solaris OS, Microsoft Windows, Red Hat Linux, IBM AIX, HP-UX and SuSe LINUX platforms.
- C*onfiguration Comparison*. The software checks for differences between two installed application instances on an on-demand or scheduled basis.
- *Deployment Simulation*. This feature stimulates the deployment process prior to implementing changes, helping to make sure that key deployment criteria, such as connectivity, permission, disk space, and other dependencies are met.
- *Role-Based Access Control*. RBAC grants specific management permissions, including write, execute, and check-in, to specific users.
- *Java Software API*. This interface enables N1 Service Provisioning System to be integrated into new or existing applications through a programmatic interface.

- *Version Control and Rollback*. The software records every action taken by system administrators across all applications and managed servers, enabling applications to revert back to previous states when needed.
- *Solaris Containers*. Operators can use N1 Service Provisioning System to provision applications onto Solaris Containers.
- *Broad Application Portfolio*. The software includes pre-built models for deploying industry-leading Java Platform, Enterprise Edition (Java EE) applications, Web servers, application servers, and databases.
- *OS Package and Patch Support*. Sun N1 Service Provisioning System includes pre-built models for installing Solaris OS packages and patches, and models for deploying RPM files in Red Hat Linux environments.
- *Software Development Kit*. The SDK provides APIs that developers can use to build custom models for provisioning specific applications

Sun Connection Services

To learn more about Sun Connection services, please visit the BigAdmin Sun Connection hub at *http://www.sun.com/bigadmin/hubs/connection/*.

The Sun Connection service is a multi-platform, IT infrastructure management service. With Sun Connection services, organizations can register and track their IT inventory, update Solaris OS and Linux systems, and provision systems remotely.

Sun Connection Service Channels

The Sun Connection service supports the Solaris 8, 9, and 10 Operating Systems on SPARC, x64, and x86 platforms, as well as a number of Red Hat, SuSE, and Linux distributions from a single network service.

The Sun Connection service is divided into several channels:

The Inventory Channel is available, at no cost, at *http://sunconnection.sun.com/inventory*.

- *Inventory Channel*. This channel manages and reports on registered inventory. Organizations can create groups to organize the inventory, and later export the data into a CSV, PDF, or XML format. IT organizations also can receive RSS feeds that might be of interest.
- S*olaris 10 Update Channel*. This channel patches Solaris 10 systems automatically when organizations have an active Solaris or SunSpectrum[SM] service plan.
- *Update Channel*. This channel patches Solaris 8, 9, or 10 OS or Linux systems with a centralized patch management solution. Administrators can create patch rules and policies for patch deployment. The patch deployment channel can be run in simulation mode to identify patch dependencies and requirements before deploying patches on live systems. The Update Channel includes automated rollback and recovery options, and it can patch virtualized environments (including Solaris Containers.)
- *Provisioning Channel*. This channel can provision Solaris, Red Hat, and SuSe Linux systems running on Sun x64 servers from bare metal. The channel can help to replicate installations onto new systems.

The Update and Provisioning Channels are subscription-based services that support managing systems from a remote console.

Sun Connection Architecture

The Sun Connection service utilizes a number of components hosted by Sun in enterprise datacenters. (Figure 25). A Sun Connection server resides at Sun and provides Solaris OS updates to the System Dependency Server hosted in the customer datacenter. Sun maintains a software component repository for the Solaris OS. The Management server at the customer site downloads RPMs directly from RedHat and SuSe for those channels.

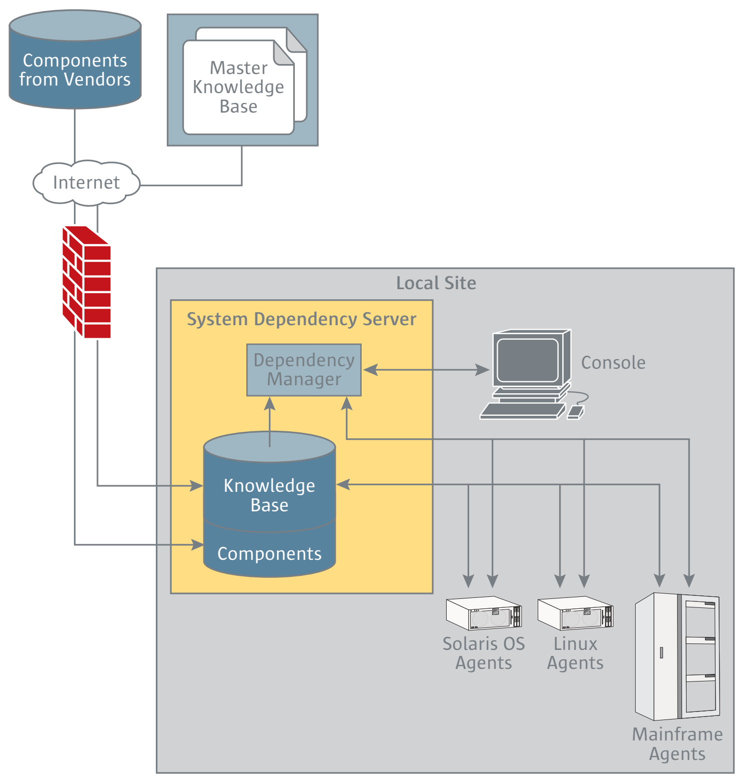

Figure 25. The local site architecture incorporates a number of components installed into enterprise datacenters.

The *System Dependency Server* (SDS) resides at the local site, acting as a local proxy to the Master Knowledge Base at Sun. The SDS includes the following:

- An *embedded Web server* that pulls certified updates from the Master Knowledge Base. and pushes local packages and scripts to the local knowledge base.
- A *Dependency Manager* (DM) that manages communications between applications, and performs job queuing and data storage.

- A *Knowledge Base* that caches both certified and generated deployment rules, as well as certified and local components.
- *Agents* installed on each managed host. Agents run the dependency resolver (DR) to find the best job solutions for the individual managed host.
- A *Console* that provides a graphical user interface (GUI) from which customers initiate tasks for the agents on the managed hosts. Organizations also can use a command-line interface (CLI) or application programming interface (API) to create site-specific user interfaces.

Sun Connection Service Benefits

The benefits to organizations using the Sun Connection service include the following:

- *Manage IT Inventory*. The Inventory Channel can help to organize and manage an organization's registered inventory.
- *Improve SLAs*. Sun Connection services can help to reduce managed and unmanaged downtime, helping to improve system availability and security.
- *Reduce Operational Costs*. Sun Connection services help to reduce operational costs of Solaris OS and Linux environments by helping organizations improve their server-to-administrator ratios.
- *Optimize Employee Effort*. Sun Connection services help to free highly skilled employees to focus on higher value tasks by automating critical, but repetitive, tasks.
- *Optimize Maintenance Windows*. Sun Connection services can deploy patches in simulate mode to analyze their impact, schedule deployment, and stage content for deployment.
- *Automate Applications and Patch Deployment*. Organizations can use the customizable profile, pre-deployment options, and post-deployment options to automate the deployment during non-critical hours.
- *Improve Efficiency*. Sun Connection services help organizations to avoid unnecessary patching, failed jobs, and downtime with complete and accurate vendor patch data with knowledge stream validation and augmentation.
- *Create Robust Audit and Compliance Reports*. Sun Connection services can help create and maintain reports that are critical for compliance with Sarbanes-Oxley and HIPAA requirements.

Sun Management Center Software

Sun Management Center is a core component of the Sun N1 software family. To help increase service levels and decrease administrative costs in enterprise computing environments, Sun Management Center provides in-depth monitoring and advanced management capabilities for Sun enterprise servers to enhance system performance, reliability, security, and utilization. Sun Management Center software:

- Can support hundreds of servers from centralized Web browser
- Manages advanced Solaris 10 OS features, including Solaris Containers and DTrace

- Monitors all Sun servers from single console with Sun N1 System Manager software integration
- Provides in-depth diagnostics for system hardware and operating systems
- Provides photo-quality images of hardware components including SPARC and AMD Opteron processor-powered systems
- Suggests corrective actions using knowledge-based information
- Integrates with Sun Cluster software to help improve management server availability
- Integrates with enterprise management solutions using standard protocols
- Java software API enables clients to access server information

Server Management Interfaces

The proliferation of servers and operating systems running in virtualized environments makes managing and monitoring them all a daunting task. Simple Network Management Protocol (SNMP) was created to provide a standard for managing network devices — ranging from switches to servers and storage systems — with a simple set of commands that allow them to be managed remotely.

Most network devices provide an SNMP agent that manages all of the SNMP communications to and from the device. Although SNMP was originally developed to manage network routers, it can also be used to manage hardware and software including operating systems, servers, printers, and power supplies. Agents monitor the devices constantly and can help to alert operators to failures before they cause outages. For example, an agent may notice whether the number of ECC memory errors is gradually increasing and suggest that memory is about to fail. Agents can report a wealth of information on just about any attribute from operating system characteristics such as swapping and CPU utilization to hardware status such as fan speeds, voltages, and temperatures.

While SNMP interfaces and protocols do not dictate specific tools or methodologies for managing platforms, its functions break down into two fundamental areas: interfaces that allow tools to autonomously manage devices, and interfaces that help third-party tools integrate with enterprise management systems.

Standalone Management Tool Interfaces

Autonomous, or standalone tools allow organizations to monitor and manage server platforms. These tools may be used on their own outside of the context of enterprise network management systems. Some examples of tools and the protocols they use for managing servers are illustrated in Table 2.

Table 2. Example autonomous management tools and the protocols they use.

Tool	Management Protocol
Sun N1 System Manager Sun Management Center	SNMP
BUI	HTTPS
SMASH	SSH

Third-Party Management Tools

One of the features that IT organizations require of enterprise network-management tools is the ability to integrate new types of systems into their network with minimal disruption of their management systems. These management platforms must support products and interfaces that allow third-party tools to interoperate easily. Some examples of tools capable of integrating third-party components are illustrated in Table 3.

Table 3. Third-party management tools and the protocols they use.

Tool	Management Protocol
IBM Tivoli CA Unicenter BMC Patrol Sun Management Center (can integrate with the above	SNMP
IPMI	IPMITool

Intelligent Platform Management Protocol Interface (IPMI) is an open-standard management interface specification used to query the state of a system. IPMI can be used to query system event logs, monitor system sensors, and system hardware status. IPMI can operate either in band or out of band to manage the system. Out-of-band management operates independently of the operating system and monitors the system regardless of the state of the operating system. With in-band management, operating system kernel drivers communicate with the service processor only when the operating system is running.

Wide Range of Options

Because SNMP is a standard used by virtually every device and operating system manufacturer, IT organizations can choose from a wide range of management tools that range from full-featured enterprise class network-management systems to open-source tools:

- Full-featured Network-Management Systems (NMS) include Sun N1 System Manager, IBM Tivoli Framework, Microsoft Systems Management Server (SMS), HP Open
- View, LANDesk, Novel Zenworks, BMC Patrol, and CA Unicenter.
- Basic server-monitoring tools include IPMITool, SMASH and BUI.

Total Network View

Network-management systems allow administrators to have a total network view of network devices such as servers, routers, and switches. These systems are easily configured and work in most network environments but some may not integrate well with devices from other hardware vendors. Some vendors, such as Sun, provide a Management Interface Base (MIB) for an NMS to use in decoding the SNMP traps provided by a server. These larger packages are designed for large, complex networks and require extensive training. It is important to make a careful assessment of potential NMS products to be sure that they can operate in your environment.

Chapter 11

Sun Datacenter Services

Designing, implementing, and managing a state-of-the-art datacenter is a significant challenge. Whether re-designing an existing environment or building it for the first time, IT organizations need an infrastructure that is sufficiently robust to meet quality-of-service requirements yet flexible enough to support deploying all of the new products and services that the company's line-of-business organizations demand.

Sun has accumulated decades of experience designing, deploying, and supporting enterprise datacenters, and through its professional services, support services, education programs, and managed services, these years of experience can be leveraged in your datacenter. Sun's IT service portfolio addresses the entire datacenter lifecycle, including design and implementation, system configuration, data migration programs, and managed services that allows organizations to outsource the operational aspects of their datacenters appropriate for meeting their business goals.

Design and Integration

Sun professionals can provide a high-level design of the technical and operational architecture required to support and optimize a production environment so that it meets business objectives. Sun can help design and deploy effective solutions that help IT organizations align themselves closely with their companys' business goals.

- *Service Delivery Network Architecture Design and Implementation Services* help IT organizations design scalable and secure datacenter network architectures. The services use Sun and third-party products to provide a network design, implementation, and best practices to support reliable and scalable growth, especially for virtualized infrastructures.
- *Service Oriented Architecture services* make Sun's SOA architects, best practices, and methodologies available for simplifying development, deployment, and management of applications and services in a service-oriented architecture.
- *Sun Dynamic Infrastructure Architecture and Design Service* can help to automate configuration and provisioning for both servers and networks. A dynamic infrastructure allows IT organizations to create, grow, and consolidate workloads dynamically without re-cabling.
- The *Sun Storage Assessment Services Suite* begins with an objective evaluation of current storage implementations and results in specific recommendations that can help increase service levels and positively impact an organization's bottom line. After a thorough, impartial review of a company's present storage environment, Sun applies best practices in tools and methodologies developed through extensive customer interactions at datacenters worldwide. Sun's best practices can help organizations maximize storage performance, increase security, minimize the risks

associated with data retention and regulatory issues, and improve an organization's ability to compete.
- *Security assessment services* deliver a comprehensive security review and assessment of a customer's existing security environment. Security exposures and risks are identified within a customer's policies, processes, procedures, networks, and systems. This gives customers the benefit of an outside security review of their environment which analyzes and measures their level of security versus industry standards and best practices.
- *Sun Enterprise Consolidation Architecture Service* can help IT organizations assess the cost and benefit of moving to a virtualized environment, resulting in an architecture designed to best match an organization's application and system characteristics with business and IT objectives.

Implementation Services

Sun offers a range of system implementation-related services that help to optimize system availability and achieve higher levels of serviceability.

- S*un Application Readiness Service* is designed to allow customer applications to be installed into a production-ready state. The service collects information about the requirements of the application and then develops an agreed -upon build specification and test plan to optimize application manageability, availability, performance and security. The completed system is then tested according to the plan.
- *Enterprise Installation Services* are designed to rapidly deploy new servers to a consistent state. This helps IT staff to stay focused on strategic business issues and helps facilitate the introduction of new products.
- *Sun Cluster Implementation Services* help IT organizations make the most of their Sun Cluster software, designing, implementing, and testing the environment to help ensure that it meets the high-availability requirements that it is designed to satisfy.
- *Sun Customer Ready Program* can design, validate, integrate, and test Sun and third-party components, helping to raise the granularity with which IT organizations can deploy systems. This helps IT organizations to maintain their focus on business issues, and it helps reduce the time needed to put new systems to work.
- The *Sun Reference Implementation Services for Data Centers* can help turn plans for state-of-the-art datacenters into reality. This service can help IT organizations consolidate, migrate, or refresh existing technology, or deploy a new datacenter following a time-proven solution. Sun can build, test, and implement customer-specific configurations using its best practices, helping to deploy a cost-effective, superior solution.
- *Data migration services* help organizations move data between storage systems, and assess, plan, and execute a migration project. Sun can provide the data management

experience and expertise to help overcome the obstacles typically associated with data migrations, such as time, cost, application downtime, and loss of productivity.

Operational Support

Sun provides a range of services to support a datacenter's day-to-day operations, from helping to establish and maintain critical infrastructure software, to providing complete managed services:

- The *Sun Suite of Identity Management Services* can help to implement identity-management solutions that allow IT organizations to know who their employees, customers, and business partners are.
- Sun's *IT services* allow customers to purchase computing and storage power on demand, augmenting their datacenter capabilities.
- The *Sun Enterprise Migration Suite* of services helps to make data and server migration a safe and virtually seamless exercise. Sun has incorporated its years of experience, skills, and engineering in tools and methods that support this suite of services.
- *Monitoring and management services* use proactive tools and recommendations from Sun to monitor systems for operational anomalies. These services help improve the reliability of datacenter environments through engineer-led system reviews and implementations.
- *Sun Managed Services* is a broad portfolio of dynamic, heterogeneous IT infrastructure and applications management capabilities that enhance the business value of IT investments through improved operational efficiency.
- *Sun Performance Analysis and Capacity Planning Services* provide information about system utilization and help customers to ensure their systems are upgraded or replaced before capacity thresholds are breached impacting service levels.
- *Storage implementation services* go beyond basic installation to configure and integrate storage hardware and software into multi vendor, heterogeneous storage technology environments. Sun implements quickly and smoothly, helping to avoid disruption that can cost organizations time, money, and business.
- *Business continuity and disaster recovery services* start by reviewing business continuity and disaster recovery procedures, including the people, processes, systems, network, data, and facilities that are critical for business continuity. Following a thorough analysis, Sun can design a detailed solution and process plan to help recover operations in the event of an interruption.

Sun has both the expertise and the services to help customers design, deploy, implement, and manage their data centers providing flexibility and agility to meet increasing service demands placed on them by internal business users and the external Internet market.

Chapter 12
Conclusion

Sun technologies touch virtually every aspect of a state-of-the-art datacenter's operations. Sun innovates at every level, including the processors that power its systems, the servers that deliver the power of SPARC, AMD, and Intel processors to operating systems and applications, and the architectural approaches and methodologies that help IT organizations manage the complexity of today's datacenter environments.

Sun has taken a leadership position in delivering eco-responsible servers to the point where, in the state of California, its Sun Fire T1000 server qualifies for Pacific Gas & Electric's rebate program. All of Sun's servers, from its single-processor UltraSPARC T1 processor-powered systems to its large symmetric multiprocessing systems, have built-in support for virtualization technologies that can help IT organizations consolidate workloads and save even more power while reducing their carbon footprint.

What differentiates Sun from the competition is its ability to deliver working solutions that integrate its architectures, best practices, and methodologies with its core software, systems, storage, and service products. Sun's ability to deliver working solutions helps IT organizations reduce risk and increase the effectiveness and agility of their datacenters. Through services provided by the Sun Customer Ready program, customers can deploy servers on a large scale while leaving it to Sun to work out any integration bugs. At the core of Sun's ability to deliver solutions is the excellence with which it delivers software, systems, storage, and services.

Software

Sun's industry-leading Solaris Operating System provides a rock-solid foundation for applications ranging from Web services to enterprise databases, with the same scalability, availability, manageability, and security available on platforms ranging from Sun's own SPARC processors to its AMD Opteron and Intel Xeon processor-powered x64 servers. Sun delivers enterprise infrastructure software in the form of its Sun Java Enterprise System that allows IT organizations to leverage one set of middleware for all of its applications, with support for basic functions ranging from identity management to messaging services. And of course, Sun's Java technology is just one proof point of Sun's excellence and leadership in software.

Systems

Sun delivers broad platform choice, allowing customers to choose the right processor for the job, from the high-performance, energy-efficient chip-multithreaded UltraSPARC T1 processor to multi-core AMD Opteron and Intel Xeon processors. Sun technologies change the way that IT organizations deploy server infrastructure in their datacenters. The Sun Blade 6000 Modular System provides a high-density, power and cooling-efficient platform that supports all three processor architectures on server modules supporting on-board disk drives— and up to 64 GB of memory on its AMD Opteron and Intel Xeon processor-powered modules. Sun raises the granularity with which IT organizations can deploy systems from a server at a time to a rack at a time with the Sun Customer Ready HPC Cluster, helping servers to spend less time on the loading dock and more time running the software they were purchased to operate. Finally, Sun supports a wide range of operating system technologies, from the Solaris OS, Linux, and Microsoft Windows to virtualization technologies including Logical Domains and VMware Infrastructure 3 software.

Storage

Sun's StorageTek storage product line offers IT organizations the choice of storage systems appropriate for the applications they run, including direct-attached storage, network-attached storage, and storage area networks. Sun understands that most IT organizations run datacenters populated with products from a variety of vendors, and innovative products such as the Sun StorageTek 9990V System embrace the competition with a storage solution that can incorporate storage from third-party vendors — including EMC and IBM — into a single, flexible, virtualized environment. There is no substitute for stable, offline copies of business-critical data, and Sun StorageTek tape libraries are built to support multiple cartridge types, embracing current technologies and helping IT organizations prepare for the future, with the capability of incorporating a wide range of current and future modular tape drives.

Services

Sun's years of experience designing, deploying, and operating state-of-the-art datacenters is reflected in the pages of this white paper. Sun's experience has led it to approach the challenges of datacenter design with a perspective that helps IT organizations deploy dynamic datacenters that can help raise resource utilization levels, conserve space, power, and cooling resources, all while increasing the level of agility with which a business can deploy new applications. In this participatory age, where applications are deployed first, and refined based on customer demand, IT organizations need to be prepared to deploy rapidly and scale at an unprecedented rate. With a pattern-based approach based on service-oriented architectures, Sun can design datacenter architectures ready to provide new, higher levels of service. When it comes to managing and monitoring those datacenters, Sun can provide managed services that provide IT organizations with custom support tailored to meet their specific needs, without the all-or-nothing approach imposed by some services organizations.

There is no other company like Sun. With its comprehensive Software, Systems, Storage, and Services portfolio, combined with the intellectual capital of its datacenter architectures and methodologies, IT organizations choosing Sun today will be ahead today, and ahead well into the future.

Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, CA 95054 USA **Phone** 1-650-960-1300 or 1-800-555-9SUN (9786) **Web** sun.com