The Sun Infrastructure Solution for Grid Computing

Driving Innovation, Research, and Competitive Advantage

Technical White Paper

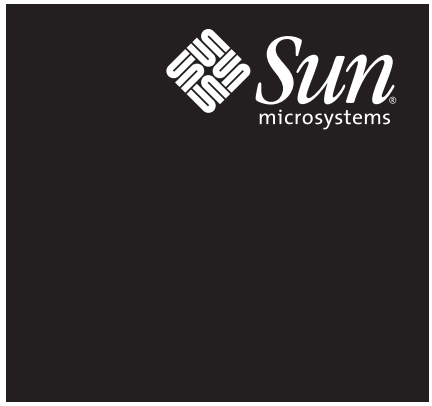

Table of Contents

Chapter 1

Executive Overview

For many technical and business organizations, computational power is literally of strategic importance, helping to enable new discoveries and new innovations for some while driving high-quality, on-time product delivery for others. Because of its ability to harness and focus computational resources, Grid Computing technology can make the crucial difference, literally letting organizations accomplish what they couldn't before. The results can be real technological, scientific, or competitive advantage along with increased productivity.

Computing infrastructure, and particularly Grid Computing infrastructure can be key to reducing costs, shortening time-to-market, and increasing product quality. Sun™ is committed to helping customers reap the benefits of this promising technology through the Sun Infrastructure Solution for Grid Computing — a comprehensive set of hardware, software, and services from Sun and its partners that help customers take better advantage of existing heterogeneous computational resources as well as deploy new resources quickly and successfully. Sun's Infrastructure Solution for Grid Computing includes:

- Industry-leading software components including award-winning Sun™ ONE Grid Engine, and Sun™ ONE Grid Engine, Enterprise Edition distributed resource management (DRM) software
- A wide range of powerful and economical Linux and UNIX® systems
- Pre-installed, configured and cabled Sun Fire™ Compute Grid rack systems that feature powerful multiprocessor Linux or Solaris™ Operating System based servers, network switching, along with powerful Grid Computing and management software.
- Proven, repeatable, Grid Computing reference architectures for key applications and industries
- Services to help architect, implement, and manage Grid Computing solutions that are tailored to customer's unique challenges and opportunities.

By taking an infrastructure solution approach, Sun can lower the risk and complexity of adopting Grid Computing strategies and greatly speed deployment of this important technology. The result is computing infrastructure that responds to the needs of the business and provides computational resources where they can make the greatest difference.

Chapter 2

Introduction

The compelling vision of *utility computing* has captured the imagination of the computing industry. Making resources available on-demand is a significant and powerful change that has equally powerful implications for both users and vendors of IT infrastructure. Grid Computing — named for the ubiquity of the familiar power grid — is a tangible early example of a utility computing model that holds significant promise for a wide range of industries and customers.

Harnessing and Deploying Compute Power for Scientific, Strategic, and Competitive Advantage

In today's research, manufacturing, and business environments, everyone is looking for an edge. Scientific, technical, and competitive pressures are creating new levels of urgency with little or no margin for error. Regrettably, these pressures are occurring at a time of stringent IT resource constraints. A challenging economy and a heightened focus on budget, total cost of ownership (TCO) and return on investment (ROI) means that every investment must predictably justify its acquisition.

For most organizations, time-to-market pressures or related issues are of critical importance. Being second with a drug discovery or a scientific insight or missing a critical market window can mean little or no return on considerable product-development investments. In this environment, getting an answer sooner, or getting to market first with a tested, quality product, can mean real competitive advantage.

With so many demands and opportunities, multiple concurrent projects with different schedules and milestones are also typical for most organizations. Product pipelines are literally stacked on top of each other and research and product teams must increasingly share the same finite set of computational resources. Datasets too are growing in size to accommodate larger and increasingly complex designs and to support cutting-edge scientific inquiries.

For many, these pressures culminate in the seemingly insatiable need for greater computational throughput — the demand for more and faster computational resources. Often there is simply not enough computational capacity available to get the necessary answers in the available time-frame. Table 2-1 lists several diverse industries with differing goals and specific computational requirements, but with the same appetite for computational power to address critical research and business needs. For these industries and others, the ability to provide appropriate levels of computational power at the right time and place can be directly related to success or failure.

Table 2-1: The need for available computational throughput spans a wide range of industries

Industry	Computationally Intensive Activities
Life Sciences	Genetic sequencing, database queries
Electronic Design	Simulations, verifications, regression testing
Financial Services	Risk and portfolio analysis, simulations
Automotive Manufacturing	Crash testing, stress testing, aerodynamics modeling, visualization
Scientific Research	Large computational problems, collaboration
Oil and Gas Exploration	Simulations, seismic analysis, visualization
Digital Content Creation	Frame rendering
Software Development	Code builds

The effect of applying appropriate resources to important applications can be stunning, directly resulting in innovation and enhanced competitive advantage. More computing resources can give organizations the critical ability to ask more questions, improve product quality, or more thoroughly explore a hypothesis before committing additional expensive resources.

In spite of this increasing demand for computing resources, much existing computing infrastructure remains greatly under-utilized. While a few individual servers may work at capacity, the vast majority of systems do not. As a result, many computing cycles are left unused — cycles that could ideally help organizations meet burgeoning computing needs.

The Sun Infrastructure Solution for Grid Computing enables organizations to harness their distributed computing resources more effectively, applying more usable computational power to their most important tasks. An effective Grid Computing infrastructure also makes it easier to transparently deploy massive amounts of new computational cycles quickly and efficiently, giving technical professionals seamless access to more computational power as they need it. With compute power removed as an impediment, organizations can literally accomplish things that they could not do before.

Infrastructure Solutions from Sun

During the economic boom of the late 1990s, companies expanded their infrastructures to meet the demands of increased network users and devices. Today, times have changed. Businesses must rethink how they create, manage, extend, and ultimately deliver products and services with greater functionality and reduced costs.

Pressure to deliver results to the bottom line is forcing organizations to find ways to drive cost out of their operations. Indeed, managers are seeking ways to do more with less. Service levels must be increased, or at least maintained. People and process issues — including maximizing uptime, ensuring security, increasing connectivity, and coping with complexity — must be addressed. Flexibility to react to new priorities resulting from mergers, changes in market focus, and integration with partners, is a must. In short, strategies that provide increasing services and applications to more people, and adapt to significant changes in corporate direction yet reduce costs, are essential.

Solving these challenging problems and getting IT assets under control requires a solid infrastructure that optimizes resources yet can evolve to meet changing business demands. Infrastructure — the network computing foundation on which business applications and services are delivered — is more important than ever. In fact, infrastructure can enable, or limit, the enterprise, and it demands competency if success is to be realized. Choosing the right

combination of architecture, hardware, software, storage, methodology, and services is paramount if IT infrastructures are to change from being reactionary to strategic and proactive. Complexity is the enemy of any IT infrastructure, and organizations must find ways to deploy the right technologies and methodologies with consistency and simplicity.

Understanding this, Sun has created a series of Infrastructure Solutions that organizations can use to meet these challenges head-on. Sun Infrastructure Solutions are the essential system elements of the enterprise and network, and provide simple solutions to complex problems. Sun Infrastructure Solutions are architected, implemented and managed to reduce deployment risks. Using the SunTone™ design methodology and tested in Sun iForce™ Centers and customer environments around the world, each Sun Infrastructure Solution includes a reference architecture, best-practices methodology, sizing guide, Sun Lifecycle services and support, a list of qualified iForce partners, and financing and technology refresh services.

The Sun Infrastructure Solution for Grid Computing

Grid Computing represents a flexible and scalable architecture that collects and focuses available computational resources to solve business and mission-critical computational challenges. By organizing a collection of interconnected computers as one unified computational engine, Grid Computing offers a cost-effective approach to solving compute-intensive problems. Users access the Compute Grid as a single computational resource, freeing them from the need to seek out appropriate systems to run their jobs and applications. The grid and its resources can also be managed as a singular entity, greatly reducing visible complexity and lowering administrative costs.

Conceptually, a grid is comprised of a collection of computing resources connected through a network along with grid middleware that aggregates these resources — servers, storage, databases, even scientific instruments — and provides transparent, remote, and secure access to computing power wherever and whenever it is needed. Like the power grid which provides electricity to every wall socket, grid computing aggregates resources and delivers computing power to every user in the network. Distinct types of grids include:

- *Compute grids* — distributed compute resources consisting of desktop, server, and High Performance Computing (HPC) systems
- *Data grids* — distributed storage devices including disk and tape devices along with the necessary software to migrate data as needed
- *Access grids* — Remote or local workstations, PCs, or thin clients that allow users access to grid computing functionality through a browser or a Unix command-line interface

This paper focuses primarily on compute grids, computing resources that are managed and collectively made available to meet an organization's specific computing needs.

Compute Grids in the Real World

In general, Grid Computing is best suited for organizations with compute-intensive applications, heavy throughput requirements, or the need to run multiple large batch jobs simultaneously. In particular, applications that benefit the most from grid computing include:

- Compute-intensive simulations, often with hundreds or thousands of similar jobs varying only in datasets or parameters.
- Standard daily workloads consisting of many medium-size, single-threaded jobs, each running a few minutes to a few hours, without the need to provide interactive results.

- Parallel applications with medium scalability (up to 16 processors) and intensive inter-processor communication; or, parallel applications with high scalability (greater than 64 processors) and low inter-processor communication.

An attractive aspect of Grid Computing is that it can be introduced into an existing computing environment by starting small and growing as needed. Grid Computing environments can be generally classified into three logical levels of deployment (Figure 2-1):

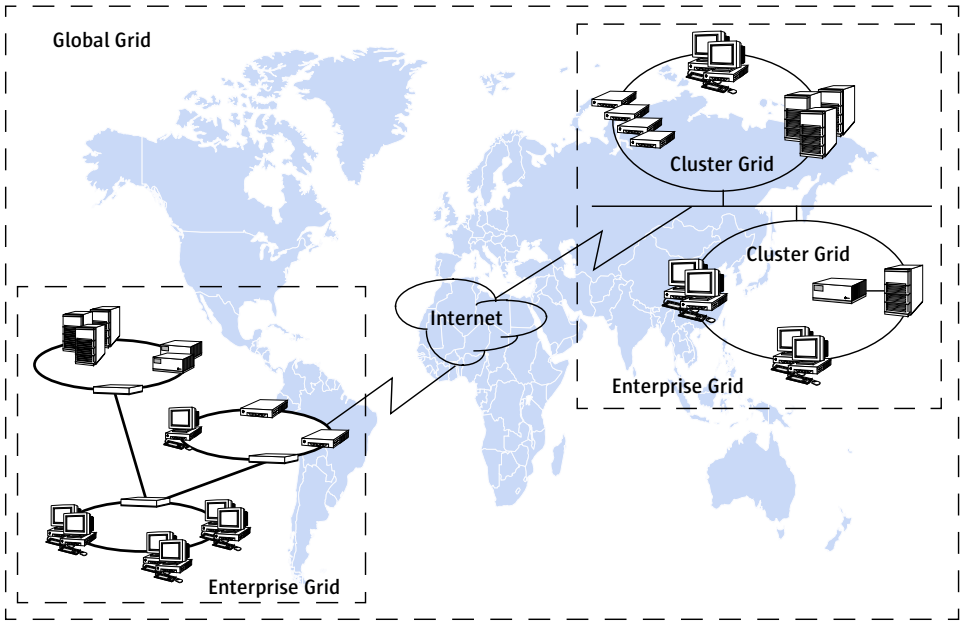

Figure 2-1: Cluster, Enterprise, and Global Grids

- *Cluster Grids (Departmental Computing)*
  The simplest and most common form of a grid, Cluster Grids consist of multiple systems interconnected through a network to provide maximized utilization and *priority-based* resource allocation. Cluster Grids may contain distributed workstations and servers, as well as centralized resources in a datacenter environment. Typically owned and used by a single project or department, Cluster Grids support both high throughput and high performance jobs. Common examples of the Cluster Grid architecture include compute farms, groups of multi-processor HPC systems, Beowulf clusters, and networks of workstations (NOW).

- *Enterprise Grids (Enterprise Computing)*
  As capacity needs increase, multiple Cluster Grids can be combined into an Enterprise Grid, enabling several projects or departments to share computing resources. Enterprise Grids typically contain resources from multiple administrative domains, but generally occupy the same geographic location. Enterprise Grids also differ from Cluster Grids in that they utilize *policies* in addition to priorities to ensure that multiple groups or departments get equitable, on-demand access to enterprise resources.

- *Global Grids (Internet Computing)*
  Global Grids represent a collection of multiple Enterprise Grids, all of which have agreed upon global usage policies and protocols, but don't necessarily share the same implementation. Computing resources may be geographically dispersed and are typically shared and accessed over the Internet. Effective data grids are an essential requirement for Global Grid implementations given the need to move and manage large amounts of data on a global scale.

Driving Innovation, Quality, and Competitive Advantage with Grid Computing

Because of its ability to focus resources where they are need most, Grid Computing can provide a host of benefits to both technical and business organizations:

- *Reduced time-to-market* — By applying more usable compute power to critical applications, Grid Computing gives researchers and development teams alike the tools they need to get their jobs done faster, often making the critical difference between success and irrelevance in the marketplace. This competitive advantage can also be expressed in terms of better innovation, more opportunity for discovery, and more time for testing to improve product quality.
- *Better utilization of resources* — By collecting and efficiently managing existing distributed resources, Grid Computing helps protect IT investments, containing costs while providing more capacity for important tasks. Grid Computing infrastructure also provides an architecture for growth, helping to ensure that future IT investments are also optimally utilized.
- *Increased productivity* — One of the most valuable resources of any organization are its people, particularly hard-to-find technical professionals. By providing transparent access to resources, Grid Computing infrastructure enables work to be completed more quickly. Users gain additional productivity since they can focus on design and development rather than wasting time hunting for resources and waiting for their compute jobs to provide results.
- *Scalability and flexibility* — Grid Computing infrastructure can grow seamlessly over time, allowing many thousands of processors to be integrated into a single Cluster Grid. Components can be updated independently and additional resources can be added as needed, reducing large one-time expenses. Grid Computing provides computing power where it is needed most, helping to better meet dynamically changing work loads. Grids can also contain heterogeneous compute nodes, with platforms selected based on the capabilities they deliver.

High-level Grid Solution Architecture

Depicted in Figure 2-2, Sun's Infrastructure Solution for Grid Computing takes a comprehensive view with the customer's goals as its focus. Because no two organizations are identical, Sun's approach considers the entire environment, employing Sun's products and solutions as well as those of key partners in the delivery of hardware, software, and services.

Figure 2-2: High-level Architecture of Sun's Infrastructure Solution for Grid Computing

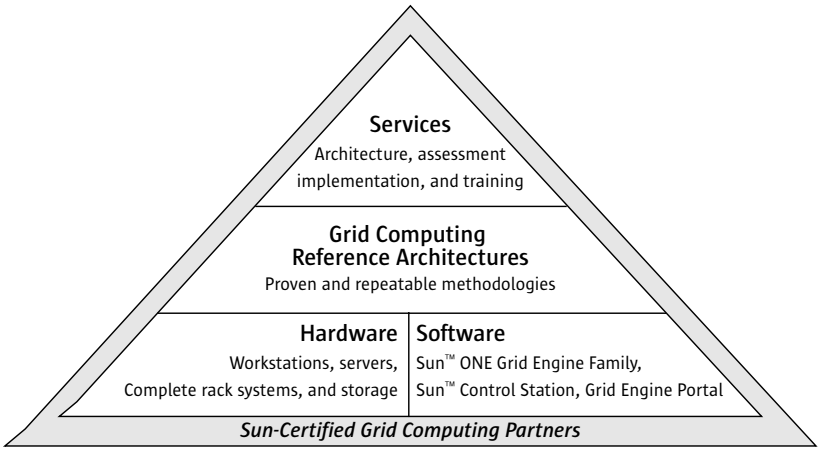

Sun works with each organization or group to understand what they need to accomplish, leveraging tried, tested, and proven Grid Computing reference architectures for key industries along with Sun's wealth of experience and expertise. Sun's vision for Grid computing goes beyond individual hardware, software, and storage products to deliver real business-critical results.

Powering Grid Computing into the Future

Grid computing is rapidly gaining momentum in many diverse industries and this promising technology will continue to evolve. Sun anticipates continued growth in Grid Computing infrastructure and is making key investments to foster innovation in this important area.

For now, most Grid Computing installations exist within individual organizations in the form of either Cluster or Enterprise Grids deployed on campus intranets. These grids are already paying dividends by providing organizations with the full advantage of their resources in order to take their technology to market faster and with greater value. Internet-based architectures will surely follow that combine geographically-dispersed resources within a company or between different companies or organizations.

These new Global Grids will have a phenomenal potential for increased power and capabilities over even today's largest Enterprise Grids. Grids on this scale will require new technology in the form of security and data bandwidth, both essential components that will allow organizations to send large quantities of their most valuable information and intellectual property over the network for processing. Open industry standards will also be needed in order for different grid implementations to interoperate seamlessly on a global scale.

With its considerable experience as both a user and provider of Grid Computing technology, Sun understands these key issues and is working both internally and with partners to provide effective solutions. As a leading manufacturer of computer systems and software, Sun makes extensive internal use of Grid Computing technologies to help design, test, verify, and build its products. Sun's own engineering and product development teams understand that effective deployment of Grid Computing infrastructure can make a critical strategic difference.

Sun has been instrumental in the early progress of Grid Computing, easing adoption by making the industry-leading Sun™ ONE Grid Engine software for distributed resource management (DRM) available at no cost, eliminating the price barrier for adoption. Sun also promotes Grid Computing through a collaborative, open source project. The goals of Sun's Grid Engine open source initiative include the advancement of Grid Computing technology and the establishment of community driven standards for distributed computing. In addition, the open source project seeks to enable developers to create innovative services and applications across distributed heterogeneous compute environments. For Grid Computing to evolve, it must be based on open standards. As a part of its commitment to open standards, Sun takes a leadership role in key Grid Computing standards organizations.

Sun's Commitment to Infrastructure Solutions

Sun Microsystems™ has been developing robust network computing solutions for over twenty years. In a world where technology advances at breakneck speeds, companies are looking to forge alliances that enable them to capitalize on each other's strengths. The combination of Sun systems, software, and services join forces in Sun Infrastructure Solutions to deliver simple and efficient end-to-end solutions that match application delivery needs from the workgroup to the data center. The latest contribution to a long-standing history of providing solutions that solve business challenges, Sun Infrastructure Solutions help businesses optimize resources while maintaining flexibility to adapt to changing business demands. With an unwavering commitment to solving complex problems, Sun continues to invest in new technology to create a complete line of Sun Infrastructure Solutions that promise a bright and interesting future.

Chapter 3

The Sun Infrastructure Solution for Grid Computing

Sun's Infrastructure Solution for Grid Computing offers a carefully-considered modular approach that provides significant advantages for users, system administrators, and software developers.

A Complete Grid Computing Solution

Though completeness is a key differentiator between Sun's Grid Computing solution and other implementations, one of its most compelling features is a modular and open design. Solution components are separate and have unique roles within the architecture. This approach is commonly referred to as a software stack, with each layer in the stack representing distinct functionality. The compute grid stack is illustrated in Figure 3-3. The products that provide the functionality required at each layer are shown in italics, and many of these key components are further described in this paper.

Figure 3-3: Sun's Grid Computing Stack

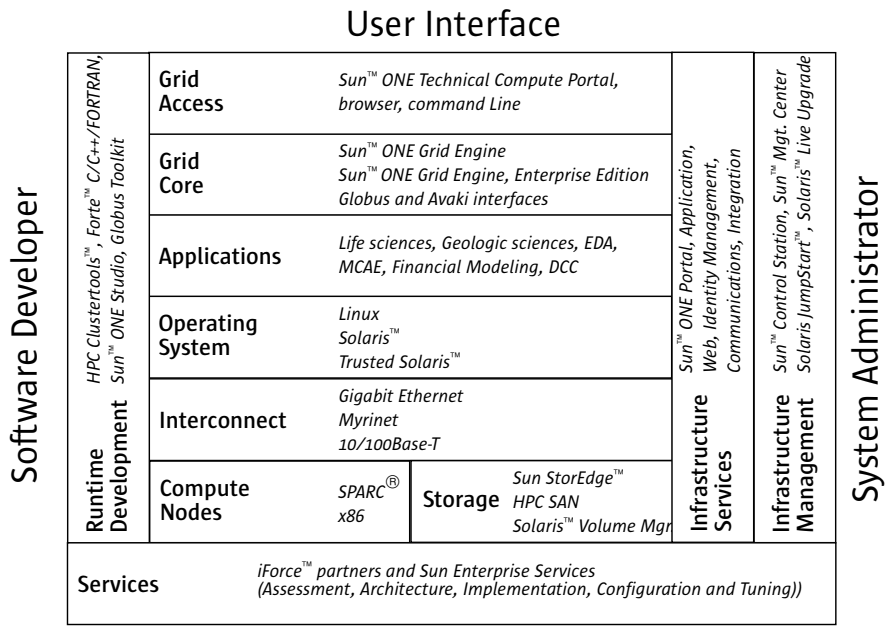

The software, hardware, and service components depicted in Figure 3-3 suggest the different functional requirements of a complete Grid Computing environment and the particular Sun

technology that can be employed. Software shown in a particular layer may not be present on all nodes, or a node might run only a small component of a software package. For each individual function, the Cluster Grid architect is free to substitute another vendors hardware, software, or services, while maintaining all other parts of the solution.

Offering a complete solution offers several significant advantages, including seamless integration of components, a single scalable platform for administrators to manage, and a single source for support. In addition, Sun offers powerful world-class platforms such as the Linux and Solaris Operating System, an extensive range of Sun servers, and Sun StorEdge™ products for a scalable and reliable foundation for building and deploying a successful Grid Computing solution.

Three-tier System Architecture

Grid Computing generally employs a standard three-tier system architecture, with front-end access nodes, mid-tier management nodes, and back-end server nodes (Figure 3-4). While the components of the Grid Computing stack map onto this logical architecture, there is not always a direct one-to-one correspondence. For example, Sun ONE Grid Engine software provides distributed resource management services at each tier in the architecture.

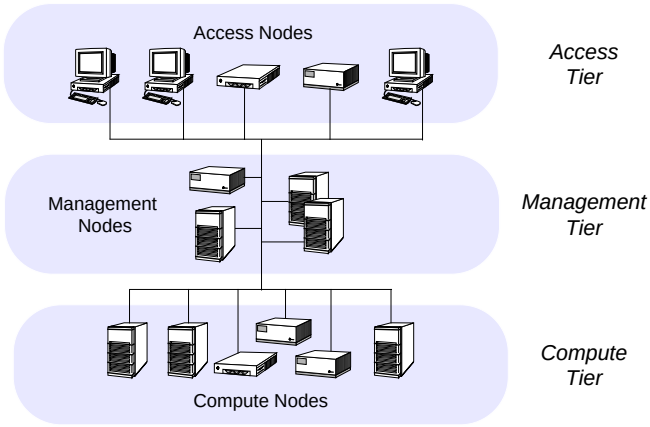

Figure 3-4: General three-tiered Grid Computing architecture

- *Access Tier*
  The access tier provides access and authentication services to Grid Computing users. Conventional command-line access methods, such as telnet, rlogin, or ssh along with Web-based portal services provide either open or tightly-controlled access to the facility.
- *Management Tier*
  The management tier includes one or more servers which run the server elements of client-server software such as Distributed Resource Management (DRM), hardware diagnosis software, and system performance monitors. Additional duties of servers in this tier may also include file servers, license key servers, or software provisioning servers.
- *Compute Tier*
  The compute tier supplies the computational power for the grid. Jobs submitted through upper tiers in the architecture are scheduled to run on one or more nodes in the compute tier. Nodes in this tier run the client-side of the DRM software, the daemons associated with message-passing environments, and any agents for system health monitoring. The compute tier communicates with the management tier, receiving jobs to run and reporting job completion status and accounting details.

Grid Access and Distributed Resource Management (DRM) Software

The ultimate goal of the Grid Computing paradigm is access to advanced computational resources from anywhere, anytime, and from any platform. The focus is no longer just the computing environment itself, but expanded application access, data access, and availability. As a part of its Sun ONE product line, Sun provides several products that are specifically designed for Grid Computing environments (Table 3-1). Together, these products provide the principal components that drive the grid and provide unified access to its resources.

Table 3-1: Features and uses of Sun's Grid Computing software

Product	Capabilities
Sun ONE Grid Engine	Cluster Grid implementation (free Internet download)
Sun ONE Grid Engine, Enterprise Edition	Enterprise Grid implementation with integral policy management module
Grid Engine Portal	Provides Web-based access to Grid Computing services

- *Sun ONE Grid Engine software*— Available free-of-cost from Sun's Website, Sun ONE Grid Engine software is ideal for departmental Cluster Grid implementations and provides traditional DRM functionality, including batch queueing, load balancing, job accounting statistics, user-specifiable resources, suspending and resuming jobs, and priority management of cluster-wide resources. Sun ONE Grid Engine software also includes a batch-aware shell, *Qtsch*, that allows interactive applications to be used with the Sun Grid Engine software. With *Qtsch*, resource-intensive applications can be automatically sent to appropriate servers without user intervention.
- *Sun ONE Grid Engine, Enterprise Edition software* — Sun ONE Grid Engine, Enterprise Edition includes the capabilities of Sun ONE Grid Engine but adds a critical policy management module that enables organizations to specify how resources are to be shared on a enterprise-wide basis. Individual groups are guaranteed access to their allotted resources and those resources are free for distribution to other groups when they are not being used.
- *Grid Engine Portal Software* — Open source Grid Engine Portal software helps integrate the capabilities of the Sun ONE Portal Server and Sun ONE Grid Engine software. High-performance technical computing users get secure anytime, anywhere access to Web-based services, content, and complex applications through a standard Internet browser and simple point-and-click interface.

  The Grid Portal Server software represents a services-centric, Web-based, shared-everything approach to technical computing. It offers an easy-to-use interface for job submission, job control, and access to results via the Sun ONE Portal Server (formerly iPlanet Portal Server) and the Sun ONE Grid Engine software. The Sun ONE Portal Server is a community-based server application that securely provides an aggregation of key content, application and services personalized based on user role/identity, user preferences, and system-determined relevancy.

Sun ONE Grid Engine Architecture

Distributed resource management (DRM) software and its batch queuing mechanism are the basic required element of a Grid Computing environment. As with traditional batch environments, in the normal operation of a compute grid, jobs are queued until required resources are available.

DRM further enhances traditional batch queuing by monitoring host computers in the cluster for proper balancing and load conditions.

Servers in a Sun ONE Grid Engine implementation are referred to as hosts. There are several types of hosts, as shown in Figure 3-5.

- *Master host* — A single host is selected to be the Sun ONE Grid Engine master host. This host handles all requests from users, makes job scheduling decisions, and dispatches jobs to execution hosts.
- *Execution hosts* — Systems in the cluster that are available to execute jobs are called execution hosts.
- *Submit hosts* — Submit hosts are machines configured to submit, monitor, and administer jobs, and to manage the entire cluster.
- *Administration hosts* — Grid managers use administration hosts to make changes to the cluster configuration, such as changing DRM parameters, adding new compute nodes, or adding or changing users.
- *Shadow master host* — While there is only one master host, other machines in the cluster can be designated as shadow master hosts to provide greater availability. A shadow master host continually monitors the master host, and automatically and transparently assumes control in the event that the master host fails. Jobs already in the cluster are not affected by a master host failure.

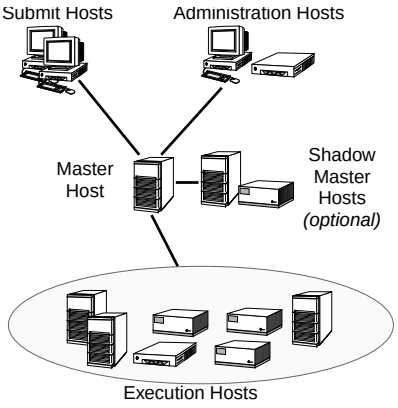

Figure 3-5: High-level Sun ONE Grid Engine system architecture

Software Job Flow

Jobs submitted to the master host are held in a spooling area until the scheduler determines that they are ready to run. Sun ONE Grid Engine software matches available resources to job requirements, such as available memory, CPU speed, and available software licenses. The requirements of different jobs may vary greatly, and only certain hosts may be able to provide the corresponding service. As soon as a resource becomes available for execution of a new job, Sun ONE Grid Engine software dispatches the job with the highest priority and matching requirements.

The Sun ONE Grid Engine software scheduler takes into account the order the job was submitted, what machines are available, and the priority of the job. Sun ONE Grid Engine, Enterprise Edition software weighs both policy and priority. Together these tools give administrators a great deal of flexibility, providing fine-grained control in managing an organization's resources. General software flow is illustrated in Figure 3-6.

Figure 3-6: Software flow in Sun ONE Grid Engine

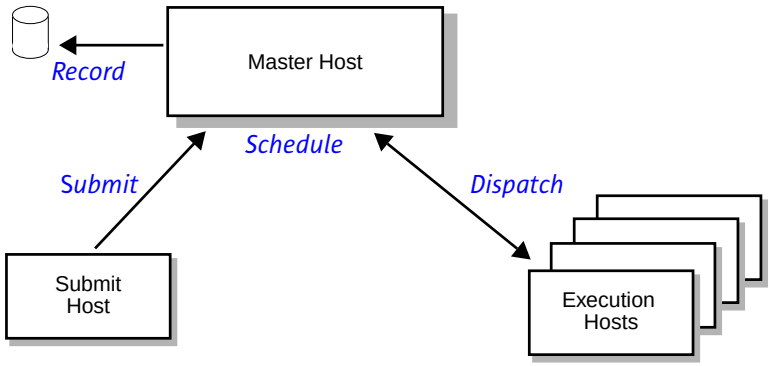

1. *Job submission* — When a user submits a job from a submit host, the job submission request is sent to the master host.
2. *Job scheduling* — The master host determines the execution host to where the job will be assigned. It assesses the load, checks for licenses, and evaluates any other job requirements.
3. *Job execution* — After obtaining scheduling information, the master host then sends the job to the selected execution host. The execution host saves the job in a job information database and initializes a shepherd process that starts the job and waits for completion.
4. *Accounting information* — When the job is complete, the shepherd process returns the job information, and the execution host then reports the job completion to the master host and removes the job from the job information database. The master host updates the job accounting database to reflect job completion.

Security

To control access to the cluster, the Sun ONE Grid Engine master host maintains information about eligible submit and administration hosts. Only systems which have been explicitly listed as eligible submit hosts are able to submit jobs to the cluster. Similarly, only systems which have been added to the list of eligible administration hosts can be used to modify the cluster configuration. For example, to provide increased security for the cluster, administration hosts can be limited to only those hosts in a physically-secured computer room.

High Availability

The cluster can be configured with one or more shadow master hosts, eliminating the master host as a single point of failure and providing increased availability to users. If the master goes down, the shadow master host automatically and transparently takes over as the master. Jobs already in the cluster are not affected by a master failure, and users gain increased reliability and availability. This functionality is a fully-integrated part of the Sun ONE Grid Engine software. The only prerequisite for its use is a highly-available file system on which to install the software and configuration files.

In addition, Sun Cluster agents are available for Sun ONE Grid Engine software, enabling highly-available configurations to be built from multiple clustered servers. Sun Cluster configurations provide automatic failover of Sun ONE Grid Engine services.

Grid Management and Control

The scalability of the Cluster Grid architecture can result in hundreds or even thousands of distributed systems that must be installed, patched, provisioned, and eventually patched and

upgraded. Capable management tools that simplify and automate routine administrative tasks are an essential requirement for these environments. The right tools allow management costs to be reduced significantly by raising the number of systems that individual administrators can manage, saving their efforts for other important tasks.

Tools such as Sun™ Management Center provide comprehensive administrative and management operations which help decrease the costs of system management and provide improved service levels in a compute grid or other networked environment. Automated installation and deployment technologies such as Sun Management Center Change Manager speed the deployment and provisioning of new or re-purposed servers. Sun™ Control Station software combines grid management with large-scale distributed system management. The N1 Provisioning Server, Blades Edition provides an effective GUI for aggregated management of the Sun Fire™ B1600 Blade Platform. Other tools include Sun Validation Test Suite (SunVTS™) to test and verify hardware functionality across a network.

Sun Management Center

Sun Management Center software is an advanced system management tool designed to support Sun systems. It offers a single point of management for Sun systems, the Solaris Operating Environment, applications, and services for datacenter and highly distributed computing environments. A powerful tool for managing the enterprise network, Sun Management Center software enables system administrators to perform remote system management, monitor performance, and isolate hardware and software faults for hundreds of Sun systems, all through an easy-to-use Web interface. Enhanced, proactive event and alarm management provide early notification of potential service problems.

Sun Management Center is based on an intelligent agent-based architecture. With this approach, a manager monitors and controls managed entities by sending requests to agents residing on the managed node. Agents, in turn, collect management data on behalf of the manager. Sun Management Center employs autonomous agent technology so that agents are not dependent on other software components, and all data collection and processing is by the agent. Based on SNMP technology, these agents collect and process data locally, and can act on data to send SNMP traps and run processes, even if the connection to the manager is severed. These intelligent agents can also initiate alarms, administrative notification, or specific actions based on collected data or messages through customizable rules and thresholds.

Use of this intelligent agent-based architecture provides several benefits which are applicable to system management of a Cluster Grid environment:

- *Scalability* — Distributing responsibility to the agents greatly improves the Sun Management Center software's ability to scale as the number of managed nodes increases.
- *Increased reliability and availability* — Because Sun Management Center agents process data locally, they are not dependent on other software components.
- *Flexibility and extensibility* — Additional modules can be dynamically loaded to Sun Management Center agents, providing the ability to adapt to changing needs.
- *Decreased bandwidth requirements* — Intelligent agents offer a savings in network bandwidth, only reporting status and significant events when necessary.
- *Security* — All users are authenticated, limiting administrators access to and management of only the systems within their control.

Solaris based Installation and Deployment Technologies

The Solaris Operating System and Sun Management Center product lines provide several interactive and automated installation and deployment technologies that can ease the administrative burden of installing and managing large numbers of distributed systems. The Solaris Operating System and other applications can be installed interactively with a browser-based interface. Solaris JumpStart software provides automated installation and setup of multiple systems over the network. Furthermore, with Solaris Live Upgrade technology, systems can be upgraded while they run, reducing system downtime.

Sun™ Management Center Change Manager, Solaris JumpStart™, and Solaris™ Live Upgrade technologies are particularly relevant to the Grid Computing environment where large numbers of similarly configured systems must be managed.

- *Sun™ Management Center Change Manager*

  Sun Management Center Change Manager enables administrators to capture a complete system image of the Solaris Operating System, application stack, and system configuration and replicate that image onto multiple servers. Particularly applicable to Grid Computing environments, this technology provides numerous advantages:

  - *Rapid deployment* — Because entire systems can be patched, tested, secured, and replicated, servers can be rapidly deployed and re-deployed as needed, reducing considerable visible complexity and saving valuable time.
  - *Layered Flash deployment* — Sun Management Center Change Manager technology provides the ability to layer Flash Archives, increasing the flexibility of the installation while also reducing the disk space required to store Flash Archives.
  - *Server Snapshot* — Sun Management Center Change Manager can also be used to capture existing server configurations, making them quickly replaceable. In the event of a complete system failure, reinstallation can usually be done in a few minutes with little downtime.

- *Solaris JumpStart™ software*

  Solaris JumpStart software is an automated system that can install and set up a Solaris system anywhere on the network without any user interaction. With Solaris JumpStart software, the Solaris Operating System and application software can be placed on centralized servers, and the install process can be customized by senior system administrators. This capability helps reduce administrative costs by enabling less-experienced administrators or even end users to quickly and easily install the software they need.

  Solaris JumpStart technology is highly customizable. Administrators can set rules which automatically match the characteristics of systems to an installation method. Custom scripts can be used to implement complex software installations and configuration in addition to the base Solaris Operating System installation. With these capabilities, Solaris JumpStart technology is well-suited to heterogeneous Grid Computing implementations where the need for customization is of key importance.

- *Solaris™ Live Upgrade software*

  Solaris Live Upgrade promotes greater availability by providing a mechanism to upgrade and manage multiple on-disk instances of the Solaris Operating System, allowing operating system upgrades to take place while the system continues to run. Solaris Live Upgrade reduces the time needed for an operating system upgrade to that of a simple reboot, shortening system downtime for upgrades. Solaris Live Upgrade can be used for patch testing and roll-out, and can also provide a safe fall-back environment in the event of upgrade problems or failures.

Sun™ Control Station (SCS) and Integral Grid Management Module

In addition to allowing administrators to deliver applications to large numbers of appliances, Sun Control Station (SCS) software provides a scalable, centralized management solution for large numbers of remotely-deployed general purpose Linux and Solaris x86 servers. Sun Control Station 2.0 software provides a simplified task-based graphical user interface (GUI) that groups managed systems based on the tasks that can be performed on each. SCS provides the aggregated management capabilities required to manage hundreds of servers in hierarchical groups.

SCS also provides in-depth visibility into the hardware characteristics of systems using lights-out-management capabilities along with performance and health monitoring. Systems administrators can monitor critical system components such as fans, temperature, system status, CPU performance, and memory capacity. All transactions between managed systems and SCS software can be encrypted for security.

Push-, Pull-, and Auto-provisioning

SCS 2.0 can perform volume provisioning by *pushing* software (patches, updates, and software version control) out to servers for higher availability and server consistency. Servers can also *pull* desired software for installation that has been published as available by SCS. SCS 2.0 can also be used as a software repository hierarchy. Upstream SCS 2.0 servers publish software and packages, making them visible for downstream SCS 2.0 servers to see and distribute. Both Linux and Solaris Operating System images can be deployed using Kickstart and Solaris JumpStart respectively.

Flexible and Extensible Platform and Framework

Most system administrators have developed their own custom build-scripts and code for automating repeated tasks that are unique to their organization and existing infrastructure. One of the most attractive capabilities of SCS 2.0 is its modular and extensible architecture, specifically designed for the needs of system administrators. With this feature, administrators can add their custom scripts as SCS 2.0 modules, letting them manage their entire environment with a simple GUI interface. SCS 2.0 provides a simple software developer's kit (SDK) to help administrators develop custom modules and add them to the SCS framework.

Integrated Grid Management Module

SCS 2.0 integrates grid management capabilities directly by providing a Grid Management module that works with Sun ONE Grid Engine software. Developed using the SCS 2.0 SDK, this module provides a complete management solution that is accessible from a single Sun Control Station GUI. Administrators get an aggregated grid management capability along with basic management of volume systems and software deployment.

N1™ Deployment Server, Blades Edition

Consistent with Sun's N1™ utility computing vision, the N1 Provisioning Server 3.0 Blades Edition provides a management environment that helps enable the rapid design, configuration, provisioning, and scaling of server farms based on the Sun Fire Blade Platform. With the N1 Provisioning Server, administrators access resources within the datacenter via a Web-based interface known as the Control Center that provides drag-and-drop design, configuration, and modification of logical server farms.

Through the N1 Provisioning Server, administrators perform the steps required to activate a logical server farm, including:

- *Allocating* — When this request is received, the N1 Provisioning Server performs resource allocation. Resources are allocated from the resource pool and tracked within the N1 Provisioning Server.
- *Image provisioning* — Disk volumes are populated from user-designated software images.
- *Virtual wiring* — Following the physical allocation of resources, network connections are configured for the Sun Fire Blade Platform. This process includes allocating virtual resources such as IP subnets and virtual network (VLAN) definitions for the blade platform switches.
- *Dispatching* — After virtual wiring is completed, the N1 Provisioning Server initiates automated DHCP and DNS services. Once these services are available, the devices within the farm are powered on.
- *Activating* — The final step in the activation of a logical server farm is the configuration of load balancer and firewall devices as defined with the network design.

Monitoring and Automated Failover

The N1 Provisioning Server includes system-level monitoring for automated failover capabilities as well as automatically detecting and replacing failed servers and load balancers. The failed device is replaced with a device of the same type from the available resource pool. The replacement automatically takes on the network configuration and software stack of the failed device. For example, when the N1 Provisioning Server detects that a blade server has become unresponsive, it performs an automatic replacement of the failed blade server with another from the free resource pool. During the failover process, the software does the following:

- Allocates a new blade server from the free pool
- Loads the last software image snapshot of the failed blade server onto the new blade server
- Configures the switch to place the new blade server in the proper VLAN and subnet
- Boots the new blade server and confirms its availability and health
- Updates the load balancer and firewalls to work with the new blade server
- Places the failed blade server into an unusable pool for future repair

In addition to availability monitoring, the N1 Provisioning Server provides users with the ability to set monitoring and alarm thresholds for a basic set of performance monitors, including CPU and disk utilization along with physical and virtual memory. Administrators can specify notification preferences and receive e-mail messages when a threshold is exceeded.

Sun Validation Test Suite (SunVTS™)

SunVTS is a comprehensive diagnostic tool that tests and validates Sun systems by verifying the connectivity and functionality of most system hardware. SunVTS can be tailored to run on various types of machines ranging from desktops to servers, and supports testing in both 32-bit and 64-bit Solaris operating environments. Tests examine subsystems such as processors, peripherals, storage, network, memory, graphics and video, audio, and communication.

The primary goal of the SunVTS software is to create an environment where Sun systems can be thoroughly tested to help enable their proper operation or to find elusive problems. SunVTS can be used to validate a system during development or production, as well as for troubleshooting, periodic maintenance, and system or subsystem stressing.

Development Tools and Run-time Libraries

Both Sun HPC ClusterTools™ and Forte™ for High Performance Computing (HPC) software are commonly used to develop and run applications on compute grids. Sun HPC ClusterTools software provides an integrated software environment for developing and deploying parallel distributed applications. Forte HPC provides support for developing high-performance (non-parallel) applications in the FORTRAN, C, and C++ programming languages.

Sun HPC ClusterTools™ software

Sun HPC ClusterTools software is a complete integrated environment for parallel application development. It delivers an end-to-end software development environment for parallel distributed applications and provides middleware to manage a workload of highly resource-intensive applications. Sun HPC ClusterTools Software enables users to develop and deploy distributed parallel applications on Sun systems with continuous scalability from one to 2048 processes within a single well-integrated parallel development environment.

Parallel Application Development

Two primary programming models are supported for high-performance, parallel applications. The *single-process model* includes all types of multi-threaded applications. These may be automatically parallelized by Sun's high performance compilers using parallelization directives (e.g., OpenMP) or explicitly parallelized with user-inserted Solaris or POSIX threads. The *multi-process model* supports the Message Passing Interface (MPI) standard for parallel applications that run both on single SMPs and on clusters of SMPs or thin nodes.

Sun HPC ClusterTools software includes a high-performance, multi-protocol implementation of the industry standard Message Passing Interface (MPI), a full implementation of the MPI I/O protocol, and tools for executing, debugging, performance analysis, and tuning of technical computing applications. Sun HPC ClusterTools software is thread-safe, facilitating a third, hybrid parallel application: the mixing of threads and MPI parallelism to create applications that use MPI for communication between cooperating processes and threads within each process. Such codes may make most efficient use of the capabilities of individual Symmetric Multi-Processing (SMP) nodes in a high-performance cluster environment.

Sun HPC ClusterTools Software

Sun HPC ClusterTools software provides the features to effectively develop, deploy, and manage a workload of highly resource-intensive, MPI-parallel applications on Sun systems, and is integrated to work with Sun ONE Grid Engine software for use in compute grid environments. Sun HPC ClusterTools software supports standard programming paradigms like MPI message passing and includes a parallel file system that delivers high-performance, scalable I/O. These features, combined with a comprehensive suite of development and administration tools, make Sun HPC ClusterTools software one of the industry's most integrated, flexible, and affordable parallel development environments.

Integration with Sun ONE Grid Engine software

In a production computing environment, the Cluster Runtime Environment (CRE) component of Sun HPC ClusterTools software can be integrated within the Sun ONE Grid Engine framework to handle the details of launching and controlling Sun MPI jobs. Sun CRE provides Sun ONE Grid Engine software with the relevant information about parallel applications in which multiple

resources are reserved for a single job. The Sun ONE Grid Engine software uses the Sun CRE component to handle the details of launching Sun MPI jobs, while still presenting the same familiar interface to the user.

Integration of Sun HPC ClusterTools with the Sun ONE Grid Engine framework provides a distinct advantage to users of a Sun compute grid. By running their parallel jobs with Sun CRE under the distributed resource management of Sun ONE Grid Engine, users achieve both efficient resource utilization and effective control over parallel applications.

An integration package is available as an advanced starting point for a site specific customization. The package creates a standard integration between Sun HPC ClusterTools and Sun ONE Grid Engine software, helping users to set up the integration scheme with minimal effort.

Forte™ for High Performance Computing (HPC)

Forte Developer, formerly known as Sun WorkShop, is a suite of products that supplies an integrated programming environment with a full set of graphical tools for software development. Forte for High Performance Computing (HPC) provides a comprehensive, productive environment for developing reliable, scalable, high-performance applications in FORTRAN, C and C++ on the Solaris Operating System.

Key features of Forte HPC that help developers build high-performance applications for deployment in a compute grid environment include:

- *64-bit application development* — 64-bit technology offers many benefits, including an address space to handle large problems, 64-bit integer arithmetic to increase the calculation speed for mathematical operations, and support for files greater than 4 GB in size.
- *Sun Performance Library compatibility* — Compatibility with the Sun Performance Library helps provide optimized performance for matrix algebra and signal processing tasks on single-processor and multiprocessor systems.
- *Integrated programming environment* — Forte HPC includes an integrated programming environment that enables programmers to browse, edit, compile, debug, and tune their applications efficiently.
- *Software configuration management tools* — Forte HPC provides software configuration management tools to enable development teams to work together effectively and efficiently.
- *Multi-threading technology* — Forte HPC software enables users to develop and tune multi-threaded/multi-processing applications using capabilities such as OpenMP API support for C and FORTRAN programs.
- *Performance analysis tools* — Performance analysis tools enable developers to evaluate code performance, spot potential performance issues, and locate problems quickly.

Chapter 4

Building Effective Grid Computing Infrastructure

The advent and growing availability of small, inexpensive servers is driving Grid Computing implementations for many industries. Sun is leading these trends with powerful and compact systems such as the 1U (1 rack unit), dual-processor Sun Fire™ V60x and Sun Fire™ V120 servers.

Nonetheless, effective Grid Computing infrastructure solutions require more than just racks of inexpensive servers. Sun's Infrastructure Solution for Grid Computing provides the expertise to assess, design, architect, and support ideal Grid Computing infrastructure that matches specific industry and customer requirements. Sun's end-to-end server and storage product lines help assure the best performance for key applications.

Sun Reference Architectures

In an effort to speed deployment and lower total cost of ownership (TCO), organizations are seeking new ways to architect IT systems. Sun's Reference Architectures help enterprises address this challenge by delivering pre-tested, pre-integrated implementations that streamline IT decision-making and lower the risk, cost, and complexity involved in architecting new IT infrastructures. By simplifying the entire process of evaluation, selection, design, and deployment, Sun's Reference Architectures enable a less complex, lower risk path to success.

Sun's Reference Architectures combine technology products from Sun and best-in-class hardware and software providers to deliver an infrastructure that is tested and proven to meet specific business needs. Designed, tested, sized, and documented, Sun Reference Architectures serve as a starting point for a customized solution, helping businesses speed time-to-deployment and significantly reduce TCO. Enterprises can tailor the final solution to meet their specific requirements. This methodology reduces costly integration services that are often required by other technology vendors.

Businesses that utilize a Reference Architecture can rest assured that Sun has completed the up front work required to architect a complete implementation that is proven to work and deliver results. In addition, Sun's Reference Architectures are built to the industry developed standards of the SunTone Initiative to deliver the highest quality of service to customers. Using the Reference Architecture methodology, companies reap tremendous benefits — without sacrificing scalability, flexibility, and choice.

Each Reference Architecture includes a series of complementary elements which work together to simplify the entire process of designing and deploying new infrastructure implementations:

- A multi-tiered architecture
- Best-in-class hardware and software components
- Architecture, sizing, and implementation guides
- Support from Sun Services and iForce partners
- Proof-of-concept testing at a Sun iForce Solution Center
- Option of deployment through the Sun Customer Ready Systems program

Reference Architectures for Grid Computing

Though Grid Computing has broad appeal, individual industries and customers have their own unique requirements that dictate the ways that Grid Computing infrastructure is selected and deployed to yield the best results. Ideal Grid Computing solutions must minimize total cost of ownership for each organization while providing the best application performance for business-critical tasks. Ultimately, applications and data drive the selection of hardware and software components as well as for storage and network infrastructure.

- *Selecting the right hardware and operating system*
  Application performance and system utilization are key factors considered by Sun's Grid Computing reference architectures, requiring a careful match between applications and the systems configured to run them. While rated processor performance is generally important, it is only part of the overall equation. For example, a fast processor running at 50 percent of its capacity may actually provide less delivered performance than a slower processor running at 80 percent capacity.

- *Deploying the right network infrastructure*
  With loosely-coupled Grid Computing environments, properly-designed network infrastructure is critical to maintaining the manageability and productivity of the Grid Computing environment. Networking infrastructure must help enable seamless operation of the grid, removing bottlenecks that could compromise performance in applications, management, or movement of data throughout the grid. Networking infrastructure must also allow for scalability so that new resources can be added to the Grid Computing environment as they are required

- *Building the right storage solution*
  As with applications, different industries and applications process varying types and amounts of data. Some data is critical and must be retained while other data represents log files and other information that may be expendable. Storage solutions must meet the demands of the industry and application, providing access from each Grid Computing node as well as meeting the organization's needs for capacity, availability, bandwidth, and manageability.

Sun's Grid Computing reference architectures weigh these considerations and others providing proven and repeatable strategies for a wide range of industries and situations. Reference architectures can accommodate a variety of heterogeneous computational resources, including workstations, smaller rack-mounted 1U and 2U servers as well as larger SMP servers and even high-end clustered environments.

Sun's End-to-end Server Product Line

Grid Computing environments often must support a variety of job types, including serial and multi-process applications. Multi-process applications implement one of several possible message passing methods for either distributed memory or shared memory platforms. The mix of different job types ultimately affects the choice of hardware for compute nodes.

For some industries and customers, Grid Computing may simply facilitate the purchase computing cycles in their least expensive, or most compact form. Small form-factor 32-bit 1U servers or blade servers may appeal to customers where expensive space is at a premium and small-footprint, high-density servers are desirable. At the same time, particular applications and grid support services may require larger symmetric multiprocessor (SMP) servers to obtain optimum performance and overall grid resource utilization.

For example, customers in electronic design (EDA) often require applications with very expensive software licenses that can quickly dwarf hardware costs. These key applications must run on the most capable platforms available and Grid Computing software must provide for optimal utilization. Industries such as Oil and Gas have stringent storage requirements to support large datasets and important floating point calculations often require the increased accuracy and capacity of powerful 64-bit platforms and operating systems.

In reality, Grid Computing environments often mix clusters of blades or 1U servers — appropriate for running distributed memory applications — with larger symmetric multiprocessing (SMP) servers to cope with important applications as well as shared-memory and large-memory jobs. Sun's end-to-end product line helps ensure that the right servers are available to provide the best possible application performance.

The Sun Fire™ B1600 Blade Platform

The Sun Fire B1600 Blade Platform integrates multiple heterogeneous blade servers into the Sun Fire B1600 Intelligent Shelf, providing shared dual-redundant power supplies and network switching. The Intelligent shelf replaces the individual chassis, power supplies, and external networking of traditional servers. This innovative approach helps customers manage the continued growth of the datacenter while still controlling operating expenses, dealing with space constraints, and leveraging existing infrastructure.

The Sun Fire Blade Platform also provides integrated management capabilities along with features that enable simplified installation and serviceability for lower total cost of ownership and increased service availability. The Sun Fire Blade Platform supports the following blades:

- Sun Fire B100s Blade Servers based on the UltraSPARC® processor running the robust and proven Solaris Operating System
- Sun Fire™ x86 Blade Servers running the Solaris Operating System or Linux
- Sun Fire™ Content Load Balancing Blades
- Sun Fire™ SSL Proxy Blades
- Validated third-party specialty blade servers

Occupying just three rack units (3U, 5.25 inches high), the Sun Fire B1600 Intelligent Shelf contains dual-redundant power supples, Layer 2 switches capable of supporting up to 16 blades, and system controllers for full remote access and administration. Heterogeneous blades executing different applications can exist in one Intelligent Shelf. In addition, all components in the Sun Fire Blade Platform are hot swappable with built-in Lights Out Management (LOM), gigabit Ethernet, and front-to-rear cooling.

Sun x86-Based Servers

Sun provides a growing product line of servers and server appliances based on the 32-bit x86 architecture.

- The Sun Fire™ V60x and Sun Fire™ V65x servers are Sun's next-generation, x86-based, entry-level servers. Capable of running Standard Linux Distributions or the Solaris Operating System (x86 Platform Edition), these 1U and 2U servers are ideal for performance-minded Grid Computing deployments.
  - With up to two Intel Xeon processors running at 2.8 GHz, up to 6 GB of memory, and up to three Ultra320 SCSI hard drives, the Sun Fire V60x server is a powerful, yet reliable, general-purpose thin server for building out Grid computing environments, deploying compute farms, or serving as small workgroup servers.
  - Featuring up to two 2.8-GHz or 3.06-GHz processors, up to 12 GB of memory, up to six 36-GB or 73-GB hard drives, and up to six PCI-X slots, the Sun Fire V65x server offers exceptional value, innovation, and choice.
- The Sun™ LX50 server is a high-performance, dual-processor, general-purpose server in a compact 1U rack-mountable form factor. Equipped with either Sun Linux or the Solaris™ Operating System (x86 Platform Edition), each Sun LX50 server provides up to two Intel Pentium III processors, dual SCSI disks, up to 6 GB of dual interleaved SDRAM, dual 10/100MHz Ethernet interfaces, and many other features.

Sun Fire™ V Series Edge Servers

Sun also provides a line of 64-bit UltraSPARC based Sun Fire™ V Series servers running the robust Solaris Operating System. Sun Fire V Series servers bring the thousands of applications that run in the Solaris Operating System to Grid Computing environments.

- The Sun Fire™ V210 server provides high-performance multiprocessing in a small 1U rack-ready enclosure featuring up to two UltraSPARC IIIi processors and up to 8 GB of memory for significant application performance. In addition to support for two internal Ultra 160 disk drives, the Sun Fire V120 server provides a PCI expansion slot along with external support for USB and Ultra 160 SCSI devices. In addition to the gigabit Ethernet interface, both Ethernet and serial management ports are available to access the system's Advance Lights Out Management (ALOM) capabilities.
- The 2U Sun Fire™ V240 server shares the capabilities of the Sun Fire V210 server but provides three internal PCI slots. For enhanced reliability and availability, the Sun Fire V240 server features hot-swappable dual-redundant power supply units, each with a separate power cord. Either power supply unit alone is capable of running the Sun Fire V240 server.
- The Sun Fire™ V100 server provides a low-cost, full-featured package with a 1U form factor, helping customers leverage their SPARC/Solaris experience to Edge Computing environments.
- The 1U Sun Fire™ V120 server offers capacity and performance at a low cost along with expandability through USB, PCI, and SCSI-2 for environments with AC electrical systems. The ruggedized 1U Netra™ 120 server is one of several servers from Sun that have received full NEBS Level 3 certification, complying with the stringent physical and electrical requirements of the telecommunications industry. The Netra 120 server features a DC power supply.

Other members of the Sun Fire V series family include the two-processor Sun Fire™ 280R server, the four-processor Sun Fire™ V480 server, the eight-processor Sun Fire™ V880 server, and the twelve-processor Sun Fire V1280 server. Sun also offers a line of 64-bit mid-range and high-end servers for applications that are best suited to larger SMP environments. These systems include the Sun Fire™ 4800 and 6800 server as well as the Sun Fire™ 12K and 15K servers that offer up to 52 and 106 processors respectively.

Custom Grid Computing Solutions from the Sun℠ Customer Ready Systems Program

To help customers provide faster Grid Computing deployments, Sun now offers factory-integration of Sun and complementary third-party hardware and software products. The Sun℠ Customer Ready Systems (Sun CRS) program delivers integrated systems that area built in Sun factories and based on customer specifications. Additionally, the Sun CRS program offers a broad range of services to deliver complete life-cycle management for infrastructure.

Helping to take the time and complexity out of deployment, all Sun CRS systems are pre-installed, pre-tested, pre-configured, and interoperability tested in Sun's ISO 9002 certified factories. Customers avoid on-site assembly and integration problems and reap the rewards of faster deployment of services. In fact, with Sun CRS, Sun has seen 90- to 95-percent reductions in deployment times coupled with up to 80-percent reductions in early-life system issues.

All Sun hardware and software products can be integrated by Sun CRS. Customers can order pre-installed and preconfigured appliances from Sun CRS such as those discussed above or they can specify their own configurations. For example, customers can work with Sun Professional Services or other partners to design custom racks of servers for Grid Computing configurations that are ready-to-run upon delivery. Customers can also incorporate selected third-party hardware and software as well as their own customized software images.

Sun Fire™ V60x Compute Grid Rack Systems

With the Sun Fire™ V60x Compute Grid rack system, Sun offers a complete Grid Computing environment integrated and cabled into a single rack that can accommodate as many as 32 low-profile 1U servers along with essential switching and management capabilities. By providing a comprehensive, tested, single-rack solution, this approach allows rapid deployment of initial or additional Grid Computing resources.

Each Sun Fire V60x Compute Grid rack system includes:

- Up to 32 dual-processor Sun Fire V60x compute nodes running either standard Linux distributions or the Solaris for x86 Operating System
- Sun Cluster Grid Manager (Control Station Software with the Grid Management module) running on a Sun Fire V60x server under either Linux or the Solaris for x86 Operating System
- One or two Gigabit Ethernet switches for connection to each compute node
- A 24-port gigabit Ethernet switch (two-switch configurations are also available)
- A 48-port terminal server to enable console and management access to each server
- A 1U keyboard/video/mouse (KVM) unit to provide local GUI access to the Cluster Grid Manager

Shown schematically in Figure 4-7, this factory-integrated, custom-built solution is delivered to a customer's site ready to run in a single rack enclosure so that deployment is quick and easy.

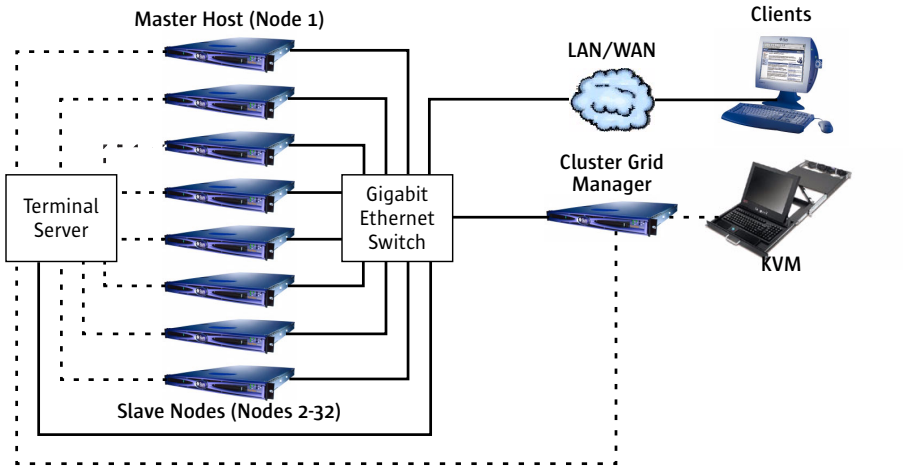

Figure 4-7: Schematic for basic Sun Compute Grid Rack

Each of the Sun Fire V60x server compute nodes is connected to at least one gigabit Ethernet switch. One compute-node runs the Sun ONE Grid Engine master software. The Sun Fire V60x based Sun Cluster Grid Manager runs Sun Control Station software along with its integral Grid Management Module. Up to four individual Sun Fire V60x Compute Grid rack systems can be connected together into a single logical farm using a 32 Gbps high-speed stacking interconnect.

Storage Solutions for Grid Computing Environments

In Grid Computing environments, each compute node must have access to essential file systems, making storage architecture a critical challenge. Advanced storage systems and file system technology are required to provide the capacity, data rates, availability, and manageability demanded by today's growing datasets.

Sun offers a wide range of flexible and scalable storage solutions that are easy to deploy, manage, and expand — ideal for environments such as Grid Computing that demand storage systems with high capacity and performance. The Sun StorEdge product family delivers complete storage solutions with availability, exceptional performance, cost-effective scalability, and streamlined management to help address increasing storage demands.

Network attached storage (NAS), high-availability NFS clustering, and storage area networks (SAN) are all highly applicable to Grid Computing environments. Sun's StorEdge™ 3000 family offers multiple attachment options to help enable a broad range of solutions. Available in both SCSI-attached and NAS configurations, the Sun StorEdge™ 3310 array provides a compact, 2U form factor for up to 876 gigabytes of 10,000 RPM disk drives. Support for up to two expansion units allows scalability up to 2.6 terabytes. For either direct-attached FiberChannel or SAN configurations, the 2U Sun StorEdge™ 3510 FC array offers up to 1.72 terabytes.

Sun's HPC SAN is an advanced, high-performance Storage Area Network (SAN) iForce solution for optimizing data movement of vast datasets with extreme data-handling requirements for many-terabyte datasets. HPC SAN provides a dedicated high-speed subnetwork for moving data to and from applications at speeds exceeding 2 GB/second. Built on Sun servers, Sun StorEdge™ T3 arrays, Sun StorEdge™ QFS file system and Sun StorEdge™ SAM-FS hierarchical storage system, and Tivoli SANergy partner software, this scalable solution consolidates fragmented storage systems, eliminating need to physically distribute storage across heterogeneous servers and establishes a single point of management and administration for storage on all HPC computing platforms.

Chapter 5

Sun^{SM} Services — Transforming Technology into Business Results

Sun provides a host of programs and services that help organizations select, prototype, deploy, and manage Infrastructure Solutions and ensure success.

Sun Services

Sun^{SM} Services, and its partners, provide a continuum of expertise, technology, and global coverage to assess business needs, and implement and manage Sun Infrastructure Solutions to help organizations realize the full value of IT investments. For each Infrastructure Solution, Sun Services offers specific consulting, training, and pre-emptive support, providing the service offerings that meet business needs. Sun Services gets solutions up and running quickly and efficiently, and provides the operational support and management capabilities that help maximize service levels while minimizing costs.

Lifecycle Services that Meet Business Needs

Every business has unique needs — needs that cannot be solved by technology alone. Companies need great people to help architect, implement, and manage complete solutions. Through professional consulting, training, and support services, Sun Services can help businesses meet challenges head-on, and offer a total solution that can help achieve business needs. Deep technology expertise, broad service offerings, and global experience serving enterprises make Sun Services the best choice for companies looking to reduce the time, cost, and risk of transforming businesses with technology. Key aspects of Lifecycle Services include:

Expert consultants

From business concept through architectural design and deployment, Sun's expert consultants can help create the innovative services that give companies sustained business advantages — maximizing return on investment while minimizing total cost of ownership.

Specific consulting services for Grid Computing include:

- *Architecture Workshop for Grid Computing* — A working session with key customer staff helps customers understand the benefits of Grid and clarify the means of achieving their vision. Customers and senior Sun consultants work collaboratively, specifically focusing on the high-level business and technology planning issues and implications of Grid Computing.
- *Architecture Assessment for Grid Computing* — This service performs a detailed assessment of an existing or proposed Grid Computing architecture. Grid architects analyze the customer's current and anticipated service level requirements and business requirements, evaluate the ability of the existing or proposed design to meet these requirements, and provide recommendations for a specified application or application area.
- *Architecture Services for Grid Computing* — The Architecture Service for Grid Computing takes as its input, the customer's architectural requirements as identified in the Architectural Assessment for Grid Computing service. This service provides a design and specification for the Grid Computing production environment addressing the complete hardware and software components as well as management and operational requirements.
- *Implementation Services for Grid Computing* — The Grid Implementation service is designed to provide the essential implementation and project management services necessary to smoothly implement and configure a Grid Computing environment. This service delivers a system that is ready to accept the customer's application.

The SunTone Architecture Methodology

The SunTone Architecture Methodology enables the development and delivery of services based on best practices. Sun consultants work through the entire architectural stack using a service-driven approach to establish Quality of Service (QoS) standards across the enterprise. The SunTone Architecture Methodology offers proven results in driving operational and management efficiencies and improving service quality and availability.

Spectrum of learning solutions

Sun and its authorized training centers educate hundreds of thousands of students annually in over 250 training facilities in more than 60 countries around the world. Training experts develop and deliver global learning solutions for the enterprise, IT organizations, and the IT professional. Curriculum is offered online, through instructors at training facilities, or at customer sites. Sun also works with organizations to understand their unique learning needs, and utilizes its network of worldwide consultants to deliver educational solutions that blend services, curriculum and certifications. With learning needs as the foundation, Sun assesses individual and team skills and identifies gaps that may pose risks to business, project, or career objectives. Sun experts then conduct job, work, and role analyses to help ensure the right people, with the right skills, are working in the right jobs. Sun offers a range of opportunities for training in Grid Computing technology including:

- *Administering and Supporting Sun ONE Grid Engine, online training*
  This interactive web-based course equips system administrators and technical support professionals with the knowledge and skills required for successful implementation and support of Sun ONE Grid Engine technology. The course explains the process of installing and configuring the product, and includes case studies that address which configuration settings are optimal for a variety of customer needs and network scenarios.

- *Sun ONE Grid Engine Advanced Administration*
  The Sun ONE Grid Engine Advanced Administration course provides students with the necessary skills to install, configure, use, and troubleshoot the Sun ONE Grid Engine software and the Sun ONE Grid Engine, Enterprise Edition software. This instructor-led course describes performing a full installation of the Sun ONE Grid Engine, Enterprise Edition software on a Network File System (NFS) server.

  Students learn how to use the command-line interface to submit interactive and batch jobs. In addition, students perform basic configuration tasks, such as displaying, adding, modifying, and deleting queue, host, and global configurations and complexes. The course also describes resource management and troubleshooting of failures in the grid. Finally, students learn about the special scheduling features available in Sun ONE Grid Engine, Enterprise Edition, including the share tree, functional, deadline, and overrides scheduling policies.

- *Grid Computing Certification Program*
  Sun helps enable Sun's resellers, System Engineers, and other partners to become more proficient in the emerging Grid Computing market. Sun Certified Grid Computing Partners have an understanding of Grid Computing, the ability to add value and differentiate their business, be involved in co-marketing activities, and have the ability to get back-up support from Sun.

  The course consists of a Web-based training module and a two-day instructor-led training class. The intent of the certification program is to educate companies so they can provide Grid computing and Sun ONE Grid Engine services and support. More information on the services provided by these partners can be found at *http://sun.com/software/gridware/partners/*.

Global support

Teaming with its extensive partner base, Sun extends its system support capabilities to span the globe, offering comprehensive proactive support solutions that best meet business needs. From traditional hardware and software support, to 24/7 access to online resources, to network-delivered system management services, Sun Services offers comprehensive support programs to meet unique business requirements — from proactive, mission-critical services to basic self-maintenance support.

Trusted Support Delivered Around the World

Sun has nearly 1,500 Sun employees located worldwide in 36 Customer Care Centers that strive to keep customer computing environments running, and to arm in-house technical experts with the right processes, tools, and technologies to help avoid issues and minimize business disruption. Sun experts are skilled and highly trained in their specialities, escalate critical situations and answer customer questions, as well as provide coordinated support with other IT vendors. With a qualifying Sun support contract, companies have 24/7 global access to support and information by phone or online.

Resolving Vendor Interoperability Issues

The Sun Vendor Integration Program (SunVIP™) simplifies and consolidates the management of multiple vendors, leading to rapid problem management and resolution. The SunVIP program ensures qualifying customers have a single Sun contact for their services and implementation needs, and streamlines problem resolution. Predefined mutual call handling procedures and problem management processes between Sun and participating SunVIP vendors are in place and ready to implement when qualifying customers call Sun.

Anywhere, Anytime Access to Information and Services

The Sun Online Support Center (OSC) offers support features in many countries and districts. Using the resources of the Sun OSC, companies can research problems, gain access to patches and manuals, and much more. Feature availability varies by country, and additional Web-based solutions are added continuously. The OSC gives companies the power to search for information, guides, and documentation, enables service contract customers to download patches and speak to the Sun community in a forum, and keeps them up-to-date on the latest resources that Sun provides. Expanded access includes service request tracking, support contract maintenance, software downloads, and other Web-based services.

Better Controlled IT Environments Through Web-Delivered Services

Sun Remote Services (SRS) Net Connect offers an innovative, self-driven way for companies to monitor IT infrastructures via the Web and to identify problems before they occur, enabling preemptive remediation. This approach provides a pipeline for services that tap into Sun's extensive knowledge base, driving availability and Quality of Service to new levels through proactive monitoring and reporting. These services are applicable to all Sun servers, from Sun Blade servers to large-scale SMP systems. SRS Net Connect is available as a free download, with additional fee-based services.

The SunTone™ Initiative — Helping Achieve Service Quality Goals

Sun Infrastructure Solutions provide companies with a solid solution design, architecture and implementation. The final step in the quality of any solution is the operational management practices to meet service level agreements. Disciplined operations management practices, such as change management controls, are critical to service quality and availability. The SunTone Initiative provides best practices for IT service management that help customers streamline operational practices and drive efficiencies up front, while addressing people and process requirements to meet on-going service delivery and management quality goals, such as availability, reliability, response times, and security. SunTone helps companies:

- Meet service level agreements (SLAs)
- Utilize existing IT resources more effectively through the adoption of streamlined processes
- Build and maintain staff competencies to preempt people and process errors that result in unplanned downtime
- Enhance service security, performance, manageability, and disaster recovery capabilities

SunTone Initiative Key Components

The SunTone Initiative consists of several key components, including a SunTone Service Delivery Specification, SunTone quality enablers and reference architectures, and certification audits.

The SunTone Service Delivery Specification

The SunTone Service Delivery Specification provides a benchmark to measure quality service delivery and management. The SunTone specification builds on and complements the frameworks of other industry standards such as ITIL, SysTrust, COBIT and ISO17799. It provides a comprehensive technical guide covering 12 assessment areas under the categories of service delivery and service support that includes: service level management, service architecture,

availability management, capacity management, security management, service continuity, data center management, configuration management, service desk, incident and problem management, change management, and release management.

SunTone Quality Enablers and Reference Architectures
SunTone Quality Enablers are pre-qualified tools and offerings for ensuring SunTone quality standards throughout the service delivery lifecycle. These pre-qualified building blocks — reference architectures, consulting and managed services, as well as SunTone certified applications, allow for easy adoption and implementation while helping to reduce complexity, costs, and risks of new solutions.

The SunTone Certification Audit
Through an assessment process and on-site audit provided by Sun Services or an authorized iForce Partner, SunTone certification validates the quality of ongoing service delivery and management practices. It also verifies the solution's ability to meet customer SLAs, whether operated and managed in-house or outsourced.

iForce™ Solution Centers

Sun's iForce Solution Centers enable customers to simulate real-word production environments in a risk-free setting. These centers — located in Menlo Park, California; Tokyo, Japan; and Paris, France — enable customers to work jointly with experts from Sun and appropriate iForce partners to build and test proof-of-concept demonstrations and run pilot programs. Sun's iForce Ready Centers help organizations validate their architectures and accelerate the deployment of high quality IT Infrastructure Solutions. Sun’s iForce Ready Centers provide:

- *Assessments*, providing consultation and evaluation of business processes to identify areas for improvement.
- *Solution Demonstrations*, providing access to a complete array of hardware and software products that are ready for custom demonstrations.
- *Proof-of-Concepts (PoC)*, prototyping business solutions that match business needs. Proof-of-Concepts determine the proper combination and specific product versions of systems, software, and services needed for a solution, and ensure solution viability.
- *Sizing and tuning*, specifying the optimal configuration of hardware and software to serve business needs.

iForce™ Partners

Sun’s Infrastructure Solutions can be extended through the addition of components, applications, and services from Sun’s market-leading, best-of-breed iForce Partners — independent software vendors (ISVs), original equipment manufacturers (OEMs), service providers, integrators, systems providers, solution resellers, and value-added resellers of Sun systems and Sun platform-based solutions. Sun understands the value of partners, and works with them to co-create, market, and deliver network computing solutions that reduce costs and increase customer satisfaction. All Sun partner programs are unified under the iForce Initiative, making it fast, easy, and safe for partners to deliver business solutions based on Sun’s robust Infrastructure Solutions.

Grid Partner Service Offerings

A significant number of partners offer services that complement the Sun Infrastructure for Grid Computing. Available in geographies around the world, these service providers have completed Sun Microsystems' Grid Certification Course and offer a variety of specializations in different industries. More information on the services provided by these partners can be found at *http://sun.com/software/gridware/partners/*.

Chapter 6

Putting the Sun Infrastructure Solution for Grid Computing to Work

While every situation is unique, the following scenario illustrates the common growth of Grid Computing within an organization — and the challenges that accompany this growth.

- *Setting up grid infrastructure* — Often, a grid computing environment is initially implemented to meet the needs of a small set of users, perhaps on a single project. A basic cluster grid is deployed to better utilize existing computing resources and improve productivity for that set of users. Sun's consultants can be especially helpful at this phase to design and implement a Grid Computing architecture that both fully utilizes existing resources to expand compute capacity as well as providing an environment for seamless growth and access.
- *Enhancing the grid* — Once a basic grid has been implemented and tested, it can be rolled out and put into production. As the organization ramps up its workload, the first thing to grow is typically the back-end environment of HPC servers, rack mounted blades, workstations, large SMP servers, storage units and interconnect technology. In addition to providing the computing horsepower needed, these systems also require system management tools and technologies that are integrated into the support environment. Sun can help customers meet the ever-increasing demand for compute resources, providing an environment that allows customer to continue to innovate, meet deadlines, enhance products, and increase product quality.
- *Ongoing configuration and tuning* — As the organization grows, it is not unusual to find that everyone's needs are not exactly alike, leading to a greater number of applications and additional compute node configurations. As the number of applications and users continues to grow, the size and complexity of the computing environment also increases. Additional resources may be needed to process, schedule, launch, and track the results of jobs executing in the grid. Middleware servers may be needed to handle job routing, licensing, message services, network security, back-end system provisioning, and service management.

Throughout the entire process, Sun can act as a trusted advisor, assisting with everything from architecture to implementation, to ongoing tuning of the Grid Computing environment to help ensure that it continues to meet business needs.

Chapter 7

Conclusion

Whether competing in today's challenging marketplace or pursuing complex technical or scientific inquiries, computing infrastructure needs to serve the needs of the organization. Grid Computing infrastructure can provide the critical difference, allowing organizations to get the most from resources they already own, and providing a model that helps assure that new resources are deployed in an efficient and effective manner.

By providing a complete and comprehensive solution, Sun's Infrastructure Solution for Grid Computing helps take the risk and cost out of deploying, growing, and managing Grid Computing infrastructure. Because Sun solves the technology and infrastructure issues, customers are free to concentrate on what they know best, their business or their science. This approach can literally provide results and competitive advantage that couldn't be obtained before.

Sun is committed to providing Infrastructure Solutions that solve real-world problems through products like the Sun Fire V60x Compute Grid rack system that enables streamlined deployment of Grid Computing technology. Sun's extensive end-to-end product line, award-winning Sun ONE Grid Engine software, and powerful system- and grid-management capabilities make Grid Computing easier than ever to deploy.

For More Information

To learn how to get the most out of Sun Infrastructure Solutions, sign up for an Infrastructure Solution Workshop that can help customize a specific solution and provide direction on tailoring an Infrastructure Solution to business environments. Table 7-1 identifies other sources of information related to Sun products and service offerings. Organizations can also contact a local Sun sales representative to learn how Sun can help build competitive advantage with Sun Infrastructure Solutions that match application delivery needs from the workgroup to the data center.

Table 7-1: Web links for additional information on Sun products and service offerings.

Web Site URL	Description
http://www.sun.com/solutions/infrastructure/	Sun Infrastructure Solutions
http://www.sun.com/solutions/infrastructure/grid	Sun Infrastructure Solution for Grid Computing
https://www.sun.com/solutions/infrastructure/12_0.html	Sun Infrastructure Solution Workshop
http://www.sun.com/service	Sun Service Offerings
http://www.sun.com/suntone	SunTone Certification and Branding Program
http://www.sun.com/iforce	Sun's iForce Initiative
http://www.sun.com/servers	Sun Hardware Servers
http://www.sun.com/solaris	Solaris Operating Environment
http://www.sun.com/storage	Sun Network Storage Solutions
http://www.sun.com/security	Sun Security Solutions
http://www.sun.com/sunone	Sun Open Net Environment (Sun ONE)

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

Sun Worldwide Sales Offices: Africa (North, West and Central) +33-13-067-4680, Argentina +5411-4317-5600, Australia +61-2-9844-5000, Austria +43-1-60563-0, Belgium +32-2-704-8000, Brazil +55-11-5187-2100, Canada +905-477-6745, Chile +56-2-3724500, Colombia +571-629-2323, Commonwealth of Independent States +7-502-935-8411, Czech Republic +420-2-3300-9311, Denmark +45 4556 5000, Egypt +202-570-9442, Estonia +372-6-308-900, Finland +358-9-525-561, France +33-134-03-00-00, Germany +49-89-46008-0, Greece +30-1-618-8111, Hungary +36-1-489-8900, Iceland +354-563-3010, India–Bangalore +91-80-2298989/2295454; New Delhi +91-11-6106000; Mumbai +91-22-697-8111, Ireland +353-1-8055-666, Israel +972-9-9710500, Italy +39-02-641511, Japan +81-3-5717-5000, Kazakhstan +7-3272-466774, Korea +822-2193-5114, Latvia +371-750-3700, Lithuania +370-729-8468, Luxembourg +352-49 11 33 1, Malaysia +603-21161888, Mexico +52-5-258-6100, The Netherlands +00-31-33-45-15-000, New Zealand–Auckland +64-9-976-6800; Wellington +64-4-462-0780, Norway +47 23 36 96 00, People's Republic of China–Beijing +86-10-6803-5588; Chengdu +86-28-619-9333; Guangzhou +86-20-8755-5900; Shanghai +86-21-6466-1228; Hong Kong +852-2202-6688, Poland +48-22-8747800, Portugal +351-21-4134000, Russia +7-502-935-8411, Singapore +65-6438-1888, Slovak Republic +421-2-4342-94-85, South Africa +27 11 256-6300, Spain +34-91-596-9900, Sweden +46-8-631-10-00, Switzerland–German 41-1-908-90-00; French 41-22-999-0444, Taiwan +886-2-8732-9933, Thailand +662-344-6888, Turkey +90-212-335-22-00, United Arab Emirates +9714-3366333, United Kingdom +44-1-276-20444, United States +1-800-555-9SUN or +1-650-960-1300, Venezuela +58-2-905-3800