Computing at the Edge

White Paper

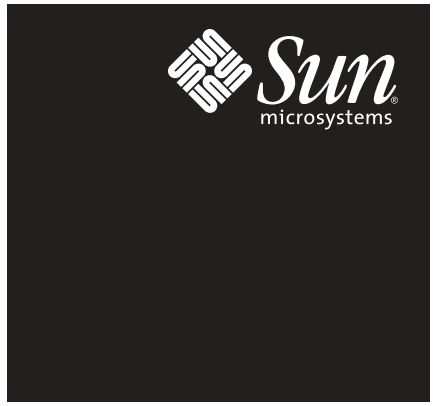

Table of Contents

Executive Summary

Edge Computing has emerged as an important strategy for proving scalable and highly available Web services for organizations around the globe. Build-out at the edge of the network is bringing about new ways for deploying and managing servers and services, and ultimately changing the ways that IT purchasing decisions are made. Beyond simply finding better ways to manage growth and bandwidth, Edge Computing technology is helping to enable new interactive, collaborative, and transactional application services that can promise to change some of the fundamental ways people make use of the Internet and other networks.

Edge Computing: Deploying computing resources at the appropriate point in the network to bring the greatest benefit to the customer and the business

Sun™ views Edge Computing as an essential area of focus for helping organizations get more from their existing computing resources as well as providing cost-effective ways to deploy new resources for improved productivity and business efficiency. Technologies like Grid Computing are providing improved utilization of existing heterogeneous resources while increasing productivity and collaboration for valuable technical professionals. Server appliances are combining key functionality like security and transactional capabilities with powerful computing platforms, allowing organizations to deploy scalable services without higher administrative costs. Increased levels of integration, standardization, and higher levels of abstraction are finally allowing management of services, rather than individual hardware, operating system, and application platforms.

Chapter 1

Edge Computing Essentials

As use of the Internet and demand for Web services continues to grow, organizations are adapting to new and constantly changing pressures that drive the ways that applications and computing infrastructure are purchased, and deployed. Millions of new devices are gaining access to the Web — especially mobile wireless devices — all requiring fast access to services from anywhere. Web services too are changing, increasingly moving from static replicated content to delivering customized, dynamic information. At the same time, the availability and scalability of Web-based services has never been more important. Services must be able to scale with unpredictable demands and remain available around the clock, and around the globe.

As organizations move to improve the customer experience, increase collaboration, and compete, they must also accommodate current economic realities. Curtailed IT infrastructure spending means that existing systems must be fully utilized and any new purchases must contribute strategically to the organization's goals. Return on investment (ROI) and lower total cost of ownership (TCO) are essential for survival. Without effective strategies, simply deploying and managing larger numbers of servers can quickly overwhelm scarce administrative resources, driving costs higher.

Edge Computing has emerged as an essential practice that turns these challenges into opportunities and focuses computing resources at the appropriate point in the network for maximum benefit. Sun understands that the promise of Edge Computing goes well beyond the benefits of mere inexpensive servers. Well designed Edge Computing infrastructure can offer cost-effective ways for organizations to solve problems on a new scale, increasing productivity and business efficiency by reducing the amount of visible complexity. Edge Computing also has serious implications for the ways that systems are designed, organized, and managed.

Deploying Resources for Greatest Effect

Edge Computing has its roots in infrastructure such as content delivery networks (CDNs) that move content, and increasingly business logic and computing resources out to the edge of the network — away from centralized data centers and closer to customers and end-users. Now Edge Computing is evolving to exploit higher levels of integration, pervasive use of standard protocols and application programming interfaces (APIs), and higher levels of abstraction. These trends are reducing management complexity, driving efficiency, and enabling innovative new Edge Computing applications.

Sizing and placing Edge Computing resources at the appropriate point in the network can bring considerable benefit to organizations along with their partners and customers:

- ***Better response time for users***
  By moving content out of centralized data center and closer to customers, CDNs and similar architectures have vastly improved response-time for clients and customers. More than mere convenience, faster response times mean new applications, new Web services, and competitive advantage.
- ***Better throughput for applications***
  As application components and select business logic move out from the data center, they need scalable computing resources to drive them. Horizontal scalability — through replicated smaller servers — is increasingly shaping Edge Computing infrastructure. But small needn't mean slow, and scalable multithreaded multiprocessor platforms, along with hardware acceleration for standard activities like encryption, are enabling new levels of service.
- ***Better utilization of resources***
  Most computational resources are sadly under-utilized. Technology such as Grid Computing is allowing technical and business organizations alike to get the most out of their existing heterogeneous computing resources while providing effective ways to deploy new resources on a large scale. Grid Computing links distributed computational resources and allows users to address them as a single inexhaustible computing entity. In addition to productivity gains for end-users, Grid Computing also provides an effective platform for collaboration and lets organizations access immense computational resources beyond their individual capabilities.
- ***More cost-effective infrastructure***
  TCO is driven by more than inexpensive systems. True cost savings only occur when computing resources are *deployed and managed* as commodities. Edge computing infrastructure such as dedicated server appliances offer integrated components with a higher-level of functionality — combining best-of-breed applications, fast deployment, and high levels of performance with cost-effective centralized management models.
- ***Better manageability and availability of services***
  Increasingly, integration of management capabilities in the form of service processors along with powerful new approaches to service management are combatting the complexity brought by large numbers of horizontally-scaled servers. This approach reduces the amount of visible complexity, allowing administrators to manage larger numbers of systems with better availability and reliability for the services they provide.
- ***Simplified operation***
  The high-level approach implied by Edge Computing goes directly to simplified operations — both for Edge Computing services, as well as back-end data center infrastructure. Simplified operations mean lower costs, and the opportunity to refocus valuable talent at solving business problems.

These and other strengths help enable Edge Computing infrastructure to deliver results that were simply not possible only a few years ago. Scalable and available Web services now support millions of clients. New kinds of electronic transactions are increasingly commonplace. New kinds of collaborations are allowing technical professionals to work together on problems that are too large for any one person or organization to solve.

Technology Innovation at the Application Layer

Any discussion of Edge Computing begins with software, for applications are now truly defining the ways that IT architectures are deployed. Business requirements, rather than legacy platforms, are determining the direction of application services and flexible software platforms are helping to enable new ways for people to interact and collaborate with less concern for underlying platforms or operating systems.

Innovative applications such as peer-to-peer, collaboration, presence, and instant messaging all presume distributed Edge Computing infrastructure. With IP networking assumed throughout the industry, application development is now focused at a platform-independent layer. The open source movement has brought a wealth of popular and portable applications, including Tomcat, Zope, PHP, PERL and others. Even legacy platform-specific, client-server applications are now frequently accessed using Java™ technology, XML, and Sun™ Open Net Environment (Sun ONE) tools resulting in portable and modular Web services.

Beyond platform independence, the modular nature of these modern development environments allows applications to be fundamentally disaggregated or delayered and distributed. Since the Java™ Virtual Machine[1] abstracts the service or application from the underlying server and operating environment, application components can run where it makes the most sense for the application.

Increasingly, underlying technologies such as Java™ servlets, the JavaServer Pages™ (JSP™) framework, portals, and directory services are being deployed on Edge Computing infrastructure. Along with better functionality for services, these trends allow developers to identify business logic, data, and application code that stays inside the firewall while placing other components closer to the customer or end user for low-latency access. As an example, portal technology gives customers fast access to information and a personalized view of their information and applications. Customers might use the portal to search a product catalog and possibly configure a product, all without taxing data center servers or networks. Transactions such as placing or checking the status of an order could still be directed to the data center, keeping sensitive data and mission-critical processing in-house.

Sun's award-winning, open-source Project JXTA for peer-to-peer computing is another key example of technology enabled by Edge Computing infrastructure. Utilizing Project JXTA technology, enterprises can implement a scalable peer-to-peer network for instant messaging or file sharing with little or no administration, accelerating the development of content sharing networks.

Different Edge Computing Perspectives

Edge Computing infrastructure connects people and organizations across networks, and when it is working at its best, it is literally invisible to those who use it. Just as every organization has different needs, appropriate Edge Computing infrastructure depends on perspective. In particular, Sun characterizes Edge Computing infrastructure based on its location in the network, dividing solutions broadly into data center, network, and client access categories (Figure 1-1)

1. The terms "Java virtual machine" and "JVM" mean a virtual machine for the Java platform

Figure 1-1: Data centers, network carriers, and clients all have related but distinct requirements for Edge Computing infrastructure. Individual organizations may deploy solutions from any or all of these areas.

Data Center Edge Computing

Data center Edge Computing infrastructure is typically focused on deploying a secure and dependable interface to the data center — providing fast low-latency access to appropriate information, without compromising security or operations. Data Center Edge Computing comprises what has typically been referred to as Tier-0 and Tier-1 of a traditional multi-tiered data center architecture.

For the data center there is a strong focus on horizontally-scaled computing solutions in the form of large numbers of small replicated servers with maximal computational density to get the most of expensive real estate. High availability along with reliability and serviceability are paramount considerations for deploying seamless Web services. Additional infrastructure components typically include load balancers, firewalls, SSL accelerators, caching servers, and VPN servers. Primary activities supported by data center Edge Computing infrastructure include:

- Web serving
- E-mail services
- Network management
- Intrusion detection
- Portal computing
- Streaming media
- Content-distribution management
- Local load balancing (between servers)

Network Edge Computing

Network Edge Computing infrastructure connects clients to the data center through public network carriers. Network Edge Computing infrastructure supports services that focus on improved throughput, and reduced latency for key services provided by the carrier network including soft switches, voice over IP (VOIP) gateways, and equipment for the delivery of IP traffic to client sites. Network Edge Computing solutions are specifically focused at providing high availability network services and serviceability is a key concern. “Lights out management” is a requirement since systems are frequently deployed remotely without local administration personnel.

Network Edge Computing deployments require systems that are powerful and fast to deploy, with tightly integrated and pre-installed services that can be configured quickly. Carrier-grade

technology, certification with the Network Equipment Building Standards (NEBS), as well as adherence to other standards are also important for network Edge Computing infrastructure. Standards such as CompactPCI and those specified by the PCI Industrial Computer Manufacturer's Group (PCIMG) are particularly relevant. Network Edge Computing services include:

- Content distribution networks (CDNs)
- DNS Services
- Global load balancing (between geographically separated sites)
- Content adaptation and acceleration
- Streaming media
- Content caching, network proxy
- Firewall, VPN/IPSec
- SSL/SSH encryption
- Qos/SLA enforcement
- IP Traffic management

Client Access Edge Computing

Edge Computing infrastructure deployed at the customer or end-user premises must be full-featured, flexible, and serve the needs of both internal and external customers without compromising manageability or security. Branch offices or small businesses need infrastructure that is easy to deploy, easy to manage, and easy to administer remotely. Server appliances in the form of small rack-mountable systems with high-level graphical user interfaces (GUIs) are ideal for client access solutions. Client access Edge Computing services typically include:

- Application staging servers
- Local content caching
- Network management
- Access points, firewall/VPN
- File and print
- Distance learning
- Calendaring and collaborative computing
- Application execution

New Approaches to Technology

Edge Computing technology is evolving rapidly, driven by both technical innovation and business priorities. These challenges are changing the very ways that computing systems are designed and built, challenging traditional assumptions about everything from processor design and system architecture through services deployment and management. Sun is driving technology in a variety of areas with significant implications for Edge Computing infrastructure:

- ***Throughput Computing***

  Service deployment on Edge Computing infrastructure is highly dependent on the computational throughput that can be provided by individual systems in dense, horizontally-scaled environments. Just as symmetric multiprocessing (SMP) has expanded the throughput of individual computer systems, Sun is actively involved in delivering these same benefits at a microprocessor level. New chip multithreading (CMT) processors are being designed with multiple processor cores on a single piece of silicon. Unlike today's typical processors that can

only process one computational thread at a time and spend a majority of time waiting for memory, CMT processors will process multiple threads simultaneously with drastically improved chip utilization and orders of magnitude improvements in system performance.

- ***Blade Servers***

  Especially applicable for high-density Edge Computing environments, Blade servers promises extremely high levels of integrated components and both improved physical footprint and improved environmental characteristics. With the Sun Fire™ Blade Platform, multiple Sun Fire B100s Blade Servers — each with a CPU, memory, and hard disk — are attached side-by-side with in a Sun Fire™ B1600 Intelligent Shelf. By integrating heterogeneous high-density blade servers, dual-redundant power and network switching, and built-in management capabilities, the Sun Fire Blade Platform provides significant horizontal scalability while greatly simplifying deployment and management of services.

- ***Just-in-time Computing***

  N1 is Sun's vision, architecture, and products for the next-generation data center. N1 addresses the many problems that have hampered organizations from delivering just-in-time computing resources, particularly valuable for Edge Computing environments. Virtualization will allow N1 to combine pools of resources that can be dynamically assigned to provide services as they are needed, greatly improving utilization and improving an organization's abilities to respond to rapidly changing conditions. The N1 Provisioning Server 3.0 Blades Edition is now available for the Sun Fire Blade Platform.

The chapters that follow describe Sun's product offerings for Edge Computing environments. Chapter 3 details Sun products and Services for deploying ready-to-run Edge Computing infrastructure. Chapter 4 describes innovative software environments that aid in the rapid deployment and management of Edge Computing infrastructure. Chapter 5 lists some of Sun's many customers that have been successful using Sun products for Edge Computing infrastructure.

Chapter 2

Sun Edge Computing Components, Platforms, and Services

Providing available and reliable services in the face of explosive growth and expanding complexity depends directly on the ability to manage services rather than individual servers. The increasing need to move quickly in response to business, competitive, or scientific imperatives means that most organizations can no longer afford to be saddled with the complexities of chip architectures, competing operating systems, and the management of individual systems and applications.

This transition implies a move from integrated discrete systems to an integrate-able component-based service delivery architecture based on high-level building blocks and high-level management concepts. The promise of this approach is services that are fast to deploy, easy to manage, easy to scale, and provide better reliability and availability. Sun is fully embracing this move toward higher level concepts by delivering software tools and integrated appliances in key application areas, all backed up by a scalable end-to-end server product line, and customized, pre-installed, and pre-configured systems that are ready to deploy.

High-level Components for Edge Computing Infrastructure

By deploying high-level infrastructure components, Sun's customers to save considerable time, effort, and cost. Though the offerings below represent different areas, together they help organizations build scalable Edge Computing infrastructure quickly while saving deployment and management costs.

Grid Computing

For many industries, available computational power has a direct connection to innovation and success in the marketplace. These industries span business and technical disciplines including:

- Manufacturing
- Life sciences
- Oil and gas exploration
- Financial services
- Educational and research institutions

Named for the ubiquity of the electrical grid, Grid Computing is helping many organizations accelerate and complete complex, compute-intensive projects that once would have been unthinkable. Grid Computing provides a number of significant benefits not available with traditional computing models, including:

- *Better utilization of resources* — Grid Computing manages heterogeneous distributed resources more efficiently, delivering more usable computing power and better utilization of expensive software licenses. This approach can decrease time-to-market, allow for innovation, or help enable additional testing and simulation for improved product quality. By employing existing heterogeneous computing resources, Grid Computing helps protect IT investments, containing costs while providing additional capacity. Grid Computing also helps streamline the deployment of new computing resources.
- *Increased user productivity* — By providing transparent access to resources, work can be completed more quickly. Users gain additional productivity because they can focus on design and development rather than wasting valuable time hunting for resources and manually scheduling and managing large numbers of jobs. The result is faster innovation and higher quality product development.
- *Enhanced Collaboration* — Collaboration, either for scientific pursuit, or commercial enterprise is increasingly essential to making future discoveries and design breakthroughs. Grid Computing can serve as an effective platform for collaboration, enabling individuals, research entities, or corporations to share large, and consistent execution environments.
- *Scalability* — Grids allow capacity to grow virtually seamlessly, with up to thousands of processors integrated into one cluster over time. Individual grid components can be updated incrementally and additional resources can be added as needed, reducing large one-time expenses.
- *Flexibility* — Grid Computing provides computing power where it is needed most, helping to better meet dynamically changing work loads. Grids can contain heterogeneous compute nodes, allowing resources to be added and removed as needs dictate.

Grid computing can be divided into three logical levels of deployment: Cluster Grids, Enterprise Grids, and Global Grids. The simplest form of a grid, a departmental or *Cluster Grid,* consists of multiple systems interconnected through a network. Cluster Grids may contain distributed workstations and servers, as well as centralized resources in a data center environment. Typically owned and used by a single project or department, Cluster Grids support both high throughput and high-performance jobs. Common examples include compute farms, groups of multi-processor HPC systems, Beowulf clusters, and networks of workstations (NOW).

As capacity needs increase, multiple Cluster Grids can be combined into an *Enterprise Grid*. Enterprise Grids let multiple projects or departments share computing resources in a cooperative way. Enterprise Grids typically contain resources from multiple administrative domains, but are located in the same geographic location.

Global Grids are a collection of Enterprise Grids, all of which have agreed upon global usage policies and protocols, but not necessarily the same implementation. Computing resources may be geographically dispersed, connecting sites around the globe. Designed to support and address the needs of multiple sites and organizations sharing resources, Global Grids provide the power of distributed resources to users anywhere in the world.

Sun is a world leader in Grid Computing technologies and an active participant in standards setting bodies. Sun offers a robust offering of leading Grid Computing products, including:

- ***Sun™ ONE Grid Engine***
  Ideal for departmental or cluster grids, Sun™ ONE Grid Engine software provides a leading Grid Computing implementation available for download from Sun's Website. Sun ONE Grid Engine is open-source and provides heterogeneous support for key platforms — allowing existing systems to be combined into effective Grid Computing configurations.
- ***Sun™ ONE Grid Engine Enterprise Edition***
  Sun™ ONE Grid Engine, Enterprise Edition software is the new flagship product in Sun's portfolio for grid computing software. In addition to all the features and benefits of Sun ONE Grid Engine 5.3 software, the Enterprise Edition features a policy module that allows computing resources to be assigned over specified periods of time (weeks, months, quarters) to each product, project, design team, and user. With this tool, management can plan compute resource usage over specific time intervals to achieve an organization's business and strategic goals.
- ***Sun Grid Engine Portal***
  Based on the Sun™ ONE Portal Server the Sun Grid Engine Portal (formerly known as the Sun Technical Computing Portal) is an enabling technology that gives organizations a single, unified interface to complex technical applications. Designed to support high-performance technical computing (HPTC) environments, the Sun Grid Engine Portal provides a distributed architecture that can deliver compute and data resources over the Web, making them available where and when they are needed most.

In addition to these products, Sun offers Grid Computing reference architectures for key industries along with an extensive list of partners that are certified to deliver custom solutions based on the Sun ONE Grid Engine and Sun ONE Grid Engine, Enterprise Edition (see Appendix).

Perimeter Security

As more and more valuable assets have been placed on-line, business continuity has come to rely heavily on effective security policies and procedures. Increasing levels of electronic transactions along with activities such as collaboration assume secure, uninterrupted communications. Unfortunately, the rate of attacks and viruses has steadily increased and systems and networks connected to the Internet are probed and attacked in many different ways on a daily basis.

The only way to successfully defend against these attacks is to deploy a comprehensive and integrated set of security tools. These tools must address complementary types of functionality and provide prevention, detection, and response to security threats. Sun understands that security is a process, not a product, and is focused on delivering best-in-class products that deliver real security. An integral part of Sun's iForce™ Perimeter Security Solution, the Sun™ LX50 VPN/Firewall appliance offers a single integrated product that can be rapidly deployed to deliver scalable Firewall and VPN services throughout the enterprise.

The Sun VPN/Firewall appliance "Secured by Check Point" software integrates the industry's leading Internet security software, Check Point VPN-1/Firewall-1 Next Generation (NG) on the Linux-based Sun™ LX50 server to create an enterprise-class appliance for maximum security and minimal total cost of ownership. Each Sun VPN/Firewall appliance is integrated and built-to-order pre-loaded with Check Point VPN-1/FireWall-1 NG software and available interfaces and options as specified by the customer.

Using an appliance approach to VPN/Firewall infrastructure provides numerous benefits:

- ***Easy to use, deploy, and manage***
  The appliance offers rapid and simple deployment, featuring integrated VPN and Firewall functionality, a user-friendly interface, and the ability to manage multiple enforcement points (either Linux or Solaris Operating System) from a single console.
- ***Secure right out of the box***
  In addition to obtaining SunTone certification for Check Point VPN-1/FireWall-1 from Sun, the VPN/Firewall appliance has also undergone a rigorous testing process to earn "Secured by Check Point" certification as well as Open Platform for Security (OPSEC) certification, guaranteeing interoperability and tight integration with Check Point infrastructure products. The appliance also features a hardened Sun Linux operating system.
- ***High performance***
  The Sun VPN/Firewall appliance offers a truly powerful combination that delivers high-throughput at a low price-point. By default, each appliance features a gigabit Ethernet interface and additional interfaces can be added up to a total of eight. In addition, an optional VPN accelerator card can be added to perform VPN encryption and decryption in hardware, off-loading the appliance CPUs for other functions.

Sun FIX Appliances

Straight through processing (STP) of electronic securities transactions has long been a goal in the financial services industry. Straight through processing can help traders increase productivity in the face of shrinking profit margins and industry consolidations that have left them struggling to consolidate legacy trading systems with multiple proprietary communication protocols. In addition, there is a regulatory push to shorten settlement cycles from three days to one — a task that would be impossible without straight through processing.

The Financial Information eXchange protocol or FIX has emerged as a common messaging standard that facilitates real-time life-cycle management for electronic financial transactions and quotations and provides a way to bridge the gap between legacy systems. Unfortunately, much custom-designed FIX software has remained expensive and out-of-reach for smaller buy-side firms. Though low-cost solutions do exist, they often lack the support, warranties, and value-added extras of custom solutions.

By integrating best-of-breed FIX implementations with scalable and affordable Sun computing platforms, Sun's FIX appliances offer considerable benefit to financial service organizations, including:

- Error reduction through elimination of manual transactions and full life-cycle electronic processing of securities transactions
- Greater productivity through the ability to handle increased transaction volume
- Platform independence and interoperability from the FIX protocol and multiple industry-leading FIX implementations.
- Affordability for cost-sensitive buy-side traders along with lower processing costs
- Scalability and flexibility to handle varying processing needs

A variety of Sun FIX appliances are available providing from two to eight processors, SPARC® and x86 computing platforms, Linux and Solaris operating systems, and a choice of leading FIX software implementations.

Sun LX50 FIX Appliance Powered by B2BITS FIXEdge software

The Sun LX50 Fix Appliance Powered by B2BITS FIXEdge software provides a high-performance multi-threaded engine that can handle multiple processing queues for over 10,000 messages per second and up to 200 simultaneous sessions. High-availability configurations are also available with hardware and software failover. The FIX appliance features:

- B2BITS' FIX Antenna V12
- A dual-processor Sun LX50 server running the Sun Linux operating system
- Access to the B2BITS FIX testing gateway
- An affordable three-year lease program that includes 24/7 support

Sun FIX Appliance Powered by Financial Fusion Global FIX software

The Sun FIX Appliance powered by Financial Fusion Global FIX software includes appliances built on a range of multiprocessor Sun systems. The Sun LX50 FIX appliance features:

- Lightweight limited edition of Global FIX 5.0
- A dual-processor Sun LX50 server running the Sun Linux operating system

Sun Fire FIX appliances are also available that provide:

- A full-featured edition of Global FIX 5.0
- Sun Fire™ 280R, Sun Fire™ V480, and Sun Fire™ V880 servers
- The Solaris Operating System
- Integrated hardware, software, and support for purchase

Global FIX 5.0 software provides an open architecture based on the Java™ 2 Platform, Enterprise Edition (J2EE™ platform). Financial Fusion also provides value-added products for integration with other trading systems including adapters for OMGEO, SWIFT, and Message broker.

Heterogeneous Edge-Computing Platforms

Given the varying requirements for Edge Computing deployments, it is clear that customers need a broad range of platforms choices. With increasingly modular and platform-independent software, choosing an effective Edge Computing platform is often a question of deployment needs rather than one of software development constraints. Edge Computing infrastructure will likely continue to contain heterogeneous platforms with servers selected for their capabilities or the ability to run applications the best. Powerful but compact Linux and UNIX® servers have proven themselves in edge environments by coupling strong performance with popular applications, security, and remote-administration capabilities.

Linux has become extremely popular for Edge Computing applications given its open-source culture and availability of key edge applications such as the Apache Web server, Samba file and print server, Squid Web proxy cache server, and sendmail mail program. Linux is particularly effective for those customers who want operating-system source code in order to provide low-level optimization or customization.

Also popular for edge deployments, the industry-leading Solaris Operating System provides a powerful and robust 64-bit high-performance environment that offers scalability, security, and manageability. Available in versions for both x86 and SPARC hardware architectures, the Solaris Operating System provides compatibility with Sun's extensive line of SPARC servers — including the 106 processor Sun Fire™ 15K server.

The Sun Fire™ Blade Platform

The Sun Fire™ 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 data center 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 the availability of special-purpose blades, Sun integrates more functionality into smaller packaging. The Sun Fire Content Load Balancing Blade and the Sun Fire SSL Proxy Blade help enable the deployment of secure Web servers that can scale horizontally. The SSL Proxy Blade accelerates SSL and RSA transactions while off-loading Sun Fire Blade server CPU processing — allowing the blade servers to focus on other applications or on higher-level transactions. With built-in Lights Out Management (LOM), built-in gigabit Ethernet, and front-to-rear cooling, the rack-ready Sun Fire Blade Platform is designed to integrate into existing environments.

Sun x86-Based Edge Servers

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

- The Sun Fire(TM) V60x server and Sun Fire(TM) V65x server are Sun's next-generation, x86-based, entry-level servers. Capable of running Standard Linux Distributions or the Solaris(TM) x86 Operating System, the Sun Fire V60x server and Sun Fire V65x server are designed for Tier 0, 1 and 2 applications. These 1U and 2U servers are powerful additions to the portfolio of Sun products designed for horizontally scaled compute environments.

  With up to two Pentium 4 Xeon processors running at 2.8 or 3.06 GHz, the Sun Fire V60x server and Sun Fire V65x server utilize the latest x86 technology to run and drive Solaris-x86-based and Linux-based applications. Drivers for Standard Linux Distributions will be introduced over the next few quarters, enabling customers to choose the OS that best meets their requirements.

- 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.

Linux-equipped Sun LX50 servers run Sun Linux 5.0, an optimized and hardened implementation that is highly compatible with Red Hat Linux 7.2. These servers are pre-configured with a wealth of applications (Web, file/print, caching, grid engine, and streaming), sophisticated Sun ONE development tools, and the mySQL database. The Sun LX50 server is also available with the Solaris Operating System making it ideal for customers with Solaris Operating System and Sun ONE expertise that wish to protect their investments while leveraging the low-cost Sun LX50 server.

- The Sun Cobalt RAQ 550 server appliances is a ready-to-deploy Linux-based server with a Web-based administration model that offers reduced administrative costs and rapid return on investment. The appliance features a 1U rack-mount form factor that integrates the hardware, software, database and development tools needed to deploy applications extremely quickly without any prior server experience. Running Sun Linux, Sun Cobalt RaQ servers are ideal for those deploying Edge Computing infrastructure for their business or for small-to-medium size businesses who are bringing their Web-hosting infrastructure in-house.

Sun Fire™ V Series Edge Servers

Sun also provides a line of UltraSPARC based Sun Fire™ V Series servers running the robust Solaris Operating System. Sun Fire V Series servers bring the very large number of applications that run in the Solaris Operating System to Edge Computing and other rack-based environments.

- 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 provides a DC power supply.
- 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, an optional Sun™ Crypto Accelerator 500 daughter card can be installed, leaving the single PCI expansion slot free for other options. A gigabit Ethernet interface is included along with external support for USB and Ultra 160 SCSI devices. 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.

Sun™ ONE Web Server software, Sun™ ONE Active Server Pages, and the Apache HTTP server software are provided and pre-installed on these systems, making them ideal for fast Edge Computing deployments. Other members of the Sun Fire V series family include the two-processor Sun Fire 280R server, the four-processor Sun Fire V480 server, and the eight-processor Sun Fire V880 server.

Ready-to-deploy Systems from Sun

To help customers provide faster Edge 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 Services for Effective Deployment

Sun Services, teaming with partners, provides a continuum of expertise, technology, and global coverage to assess customer needs, then implement and manage solutions so that organizations can realize the full value of their IT investments.

For each solution, Sun provides consulting, training and preemptive support, providing the service offerings to fit customer needs. Sun can provide assistance throughout every phase in the lifecycle of the project, from architecting and implementing, to managing the solution. Sun offers expert consulting, education, and support services that can help achieve business goals, lower total cost of ownership, simplify administration and manageability, optimize staff utilization, and improve availability, efficiency and profitability. Available services include:

- Assessment services to gain a deep understanding of existing systems and identify any potential risks before beginning the design process
- Proven methodologies and recommended best practices in the design and implementation of architectures that can achieve optimal performance in day-to-day operations as well as quickly respond to new business challenges, opportunities, and emerging technologies
- Education consulting services to help plan, deploy, and manage learning solutions
- A comprehensive portfolio of learning products that help reduce complexity and optimize IT investments by ensuring that staff are properly skilled and certified
- SunSpectrum^{SM} and Sun^{SM} Software Support Services to help get the most out of solution investments and proactively prevent problems before they happen
- Sun Vendor Integration Program (SunVIP^{SM}) cooperative support programs to help identify the source of a technical problems in computing environments
- Network-enabled services to allow self-management of Sun systems via Web tools --- at virtually no cost to the customer

Sun experts get solutions up and running quickly and efficiently and provide the back-end support and management capabilities to maximize service levels while helping to minimize operating costs.

In addition to these products and services, Sun has an extensive list of partners that are certified to deliver solutions based on the Sun ONE Grid Engine, and Sun ONE Grid Engine, Enterprise Edition (see Appendix).

Chapter 3

Accelerating Service Deployment Through Management

In most organizations, modern multi-tiered data-center environments have served the demands for availability and scalability of application services. Unfortunately, these environments are still difficult and expensive to manage — especially with horizontally-scalable Edge Computing deployments. The sheer number and variety of different things to manage in the data center is rapidly dwarfing most organizations ability to cope. Rising numbers of smaller transactional and Web-service-related servers have accelerated the level of physical and logical complexity making it difficult to track problems and locate failed physical systems.

Regrettably, most management tools today are focused at the individual server, and principally monitor for faults with minimal configuration management. As a result, change management is a complex and slow manual process. Service and policy management often reside in the heads of administrators, making those with long-term institutional knowledge extremely valuable and hard to replace.

Despite improvements, service deployment is still largely a serial process with individual servers built one at a time. The evolving nature of software means that a final, tested, secure server undergoes a predictable set of steps on the way to deployment, including:

- Hardware installation
- OS installation and configuration
- Security configuration
- Patch installation
- Application installation and configuration
- Testing
- Deployment

Since this process scales quite linearly with the number of server applications, it grows more cumbersome and complex with larger numbers of shared-application servers. Though most organizations have developed at least some in-house methods to accelerate or automate various parts of the deployment process, these techniques may provide only small incremental improvements in TCO.

Automating Service Deployment

The considerable horizontal scalability implied in Edge Computing environments requires a new approach to system management. Sun provides sophisticated management tools that add considerable leverage to the service deployment lifecycle across very large numbers of servers. Tools such as the Sun™ Management Center Change Manager and the N1 Provisioning Server 3.0 Blades Edition offer considerable labor and time savings through the capture and deployment of tested software stacks (Figure 3-1) that include a configured operating system along with pre-configured applications.

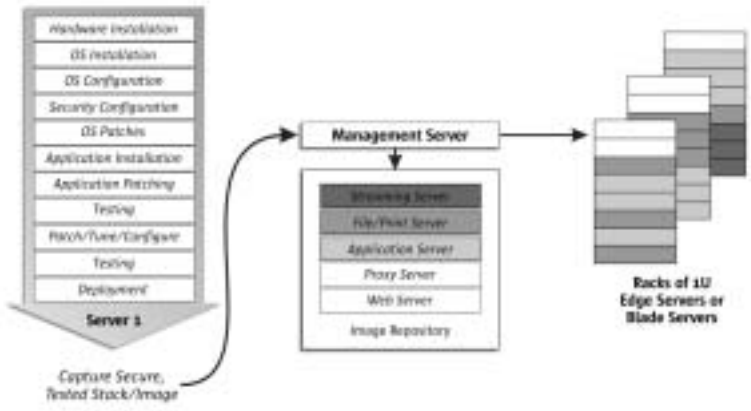

Figure 3-1: Software such as the N1 Provisioning Server 3.0 Blades Edition and Sun™ Management Center Change Manager allow rapid deployment and re-deployment of tested and secure OS and application stacks onto stateless Edge Computing servers.

With these tools, reference server configurations can be built, patched and tested for functionality, performance and security, and then captured and deployed to any number of servers simultaneously. The benefits of this approach include:

- Rapid server provisioning and application deployment
- Better administrator to server ratios
- Better utilization of computing resources through rapidly re-provisioning
- Leveraged deployment of patches, upgrades, and updates
- Pre-secured, tested, and locked-down software stacks

Sun™ Management Center Change Manager

For non-blade servers platforms running the Solaris Operating System, Sun offers a variety of tools to both manage servers as well as rapidly deploy predefined software and operating system stacks. Together with Solaris™ Flash software and Solaris™ Live Upgrade software, Sun Management Center Change Manager software helps enable servers to be rapidly redeployed for different uses in as little time as it takes to reboot. This approach provides significant savings in time and effort over traditional serial deployment techniques.

Sun’s N1 Vision

A new vision of the data center is emerging where administrators will finally be able to manage services directly with less concern for deployment details — Sun calls this vision N1. Sometimes known as fabric-based computing, N1 works by *virtualizing* the traditional data network into an anonymous and essentially stateless pool of resources including computing resources, network elements, and storage. N1 promises to leverage the disaggregation of function in the data center and dynamically recombine resources into manageable services that customers can use.

Through virtualization, N1 will allow customers to define service-level objectives while deploying just the right amount of resources need to meet those requirements. N1 will manage network, storage and computing resources — from small servers and appliances to domains on massive SMP servers.

N1 Provisioning Server Blades Edition

Consistent with Sun's N1 vision, the N1 Provisioning Server 3.0 Blades Edition provides a management environment that helps enable users to rapidly design, configure, provision, and scale server farms based on the Sun Fire Blade Platform. With the N1 Provisioning Server, administrators access resources within the data center via a Web-based interface known as the Control Center (Figure 3-2) that provides drag-and-drop design, configuration, and modification of logical server farms.

Figure 3-2: The N1 Provisioning Server 3.0 Blades Edition organizes and configures Sun Fire Blade Server Platforms into logical server farms.

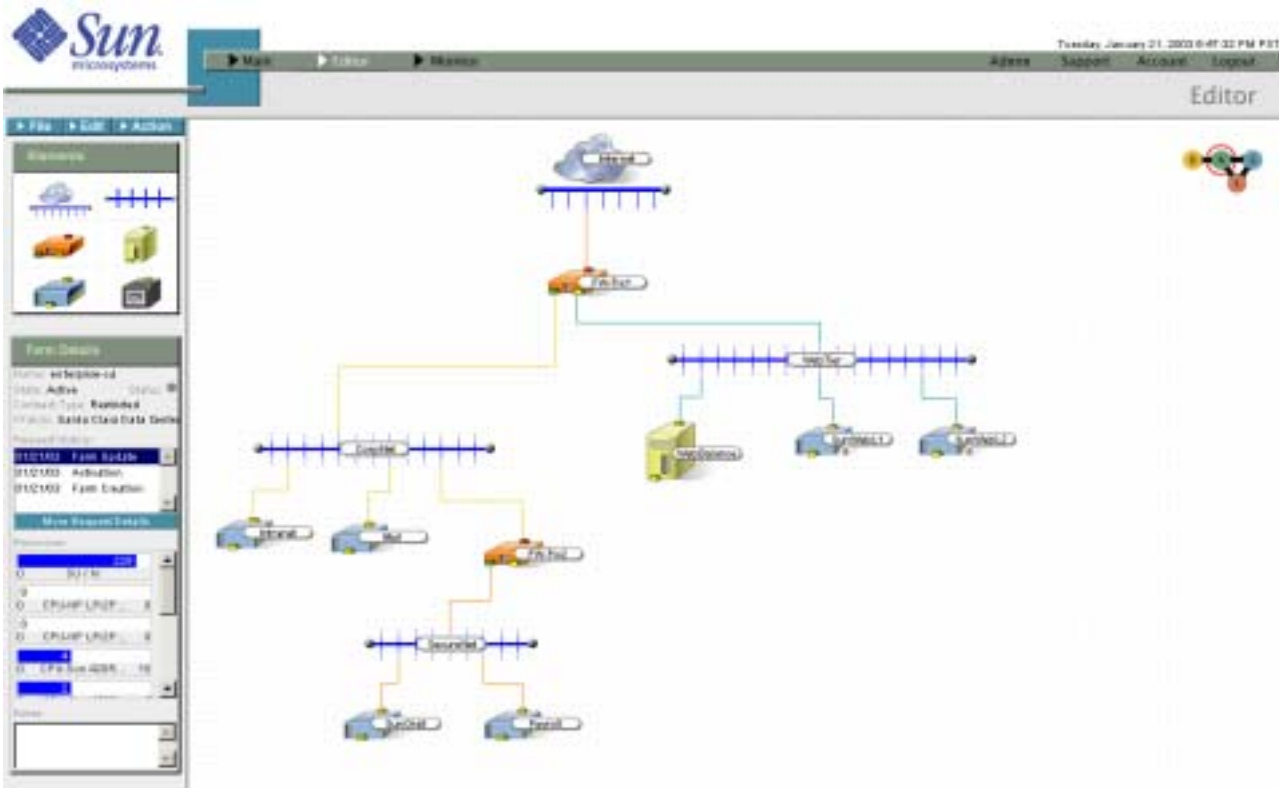

Administrators can drag icons representing resources such as servers, load balancers, firewalls, subnets, and network connectivity into the workspace and are free to design network topology and connectivity between devices by graphically connecting the network ports to a shared subnet within the logical server farm.

Server Farm Activation

The N1 Provisioning Server manages and automates the ongoing evolution of logical server farms as well as their initial activation. As the components of a server farm change, the N1 Provisioning Server automatically adds and removes resources as needed.

The N1 Provisioning Server helps enable administrators to 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 (from the image library)
- 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 Cobalt™ Control Station

In addition to allowing administrators to deliver applications to large numbers of appliances, the Sun Cobalt™ Control Station provides a scalable, centralized management solution for Sun Cobalt RaQ appliances and Linux-equipped Sun LX50 servers. Beyond basic management functionality, the Sun Cobalt Control Station includes factory-installed software modules for inventory management, performance monitoring, software management, and health monitoring.

Sun Cobalt Control Station along with the Sun Cobalt BlueLinQ service gives owners of Linux based Sun Cobalt servers instant access to product updates and new services as they become available. Through Sun Cobalt BlueLinQ technology, server administrators are notified when a new package is available, allowing fast and easy downloading and installation of software and updates. As illustrated in Figure 3-3, the Sun Cobalt Control Station *pulls* software updates from a Sun Cobalt BlueLinQ server and allows updates and software payloads to be *pushed* to Sun Cobalt appliances and Sun x86-based servers running Linux.

Figure 3-3: The Sun Cobalt Control Station manages mixed deployments of linux-based Sun Cobalt RaQ 550, and other Sun x86-based servers such as the Sun LX50 server.

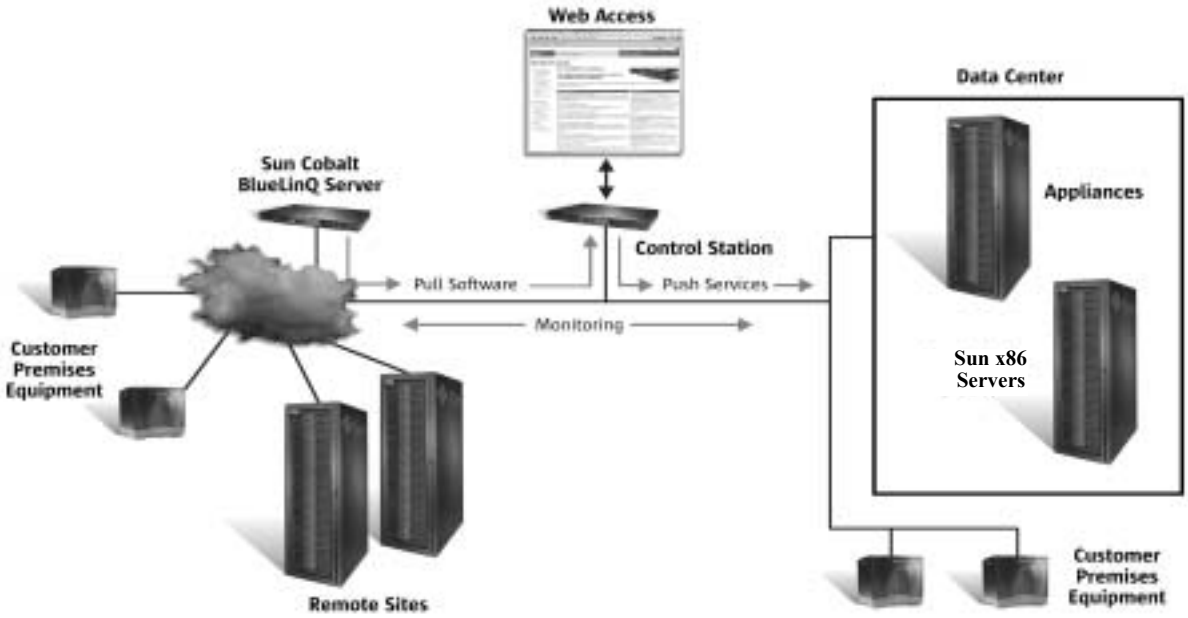

Chapter 4

Successful Sun Edge Computing Infrastructure Deployments

Sun has considerable experience with horizontally scalable computing and Sun servers are extremely popular for Internet data centers and edge-computing environments. The suitability of Sun hardware and software solutions for Edge Computing infrastructure is borne out by many successful installations in demanding and diverse deployments.

Data Center Edge Computing Deployments

Oxford GlycoSciences (OGS), PLC

Oxford Glycosciences (OGS), PLC, is a biopharmaceutical company applying integrated proteomics and genomics for the discovery, development, and commercialization of therapeutic products. Like many organizations in the life sciences and other industries, OGS depends directly on large-scale computational capabilities for development of its products and its competitive position in the marketplace.

Software such as the BLAST (Basic Local Alignment Search Tool) search tools are used to compare selected proteins against available sequencing databases. By comparing samples from both healthy and diseased individuals, proteomics can pinpoint disease-associated proteins. Some of these proteins represent potential novel targets for drug discovery and some may also provide biomarkers that can be used to monitor treatment response, select patients who might benefit from a particular drug, or improve the ability to diagnose disease.

To harness its considerable existing computational capabilities, OGS selected Grid Computing technology in the form of Sun ONE Grid Engine, Enterprise Edition. Before adopting Grid Computing, OGS experienced classic “hot-spots” in its computing infrastructure, individual systems were alternately either over- or under-utilized while project turnaround ran into numbers of months. In addition, users had to locate their own computing resources, taking considerable time from their primary activities.

Sun ONE Grid Engine, Enterprise Edition has provided OGS with an effective mechanism to harness significant compute power and considerable investment in its existing computing resources. A broad range of heterogeneous Linux and Solaris systems — from small rack-mounted 1U servers to large multiprocessor systems — is now addressed as a single unified computational

resource. As new systems are added, OGS is free to select platforms that provide the greatest computational performance for the lowest cost. Sun ONE Grid Engine, Enterprise Edition software makes the addition of these resources entirely transparent to the user.

The most important result for OGS has been massive improvements in turnaround times for projects, allowing peak demand to be addressed instantaneously. Projects that used to take as much as three months to complete are now routinely done in a week or two. OGS can now also use many tools and algorithms that simply couldn't be run before the addition of Grid Computing infrastructure. A part of OGS' "high-throughput Proteomics", Sun ONE Grid Engine, Enterprise Edition has helped provide essential competitive advantage.

Though turnaround time is the primary concern for OGS, it has also reaped other benefits. Computing resources are now considerably better utilized and the users of the system enjoy better scalability and productivity for their applications. The management of software licensing is also centralized and simplified with Grid Computing, helping to assure that licenses are better utilized while providing a convenient audit mechanism at the same time.

Looking forward, OGS is currently evaluating Sun Fire Blade Server technology for the improved management model it offers for administering large numbers of horizontally scaled servers. OGS is also evaluating the Sun Grid Engine Portal to enable collaboration with key strategic partners.

Garratt Callahan

A national water treatment company with a $50 million business and a large staff, Garret Callahan was increasingly limited by its Windows NT 4.0 network. The company was spending as much as $50,000 per year to maintain its NT servers and licenses with little prospect for improvement. Garratt Callahan reasoned that it could install a new NT 5 network at a cost of $150,000 with the same security problems or it could opt for a new Windows NT Pro network and repurchase the system licenses at full price every year, increasing annual expense while taking on a new set of security problems.

Instead, Garret Callahan added four Sun Cobalt RaQ server appliances to its network of seven NT servers and an IBM AS-400. NT services were turned off one at a time as replacement services were brought up on the Sun Cobalt RaQ servers. As a result, the one-week transition to the Sun Cobalt server appliances was completely transparent to users.

Garratt Callahan estimates that the company saved $100,000 compared with a new NT or XP network. Equally important, because there are no licenses and no suites of software to replace, the company estimates it will save another $20,000 to $70,000 per year — with vastly improved security over an NT solution.

Network Edge Computing Deployments

BlackSun

Internet service providers require affordable yet powerful servers that are easy to administer in large numbers. In Saskatoon, Canada, Internet service provider BlackSun was literally founded on Sun Cobalt RaQ server appliances and they remain the only servers offered to the company's over 5000 customers. BlackSun provides its membership with guaranteed performance, reliability, service and uptime. This strong record is possible because BlackSun's infrastructure is built *entirely* with Sun Cobalt servers and Cisco routers and switches.

With the increase in e-commerce and other transaction-based services, BlackSun began looking for a new platform that could offer increased processing power and capacity. The company chose the Sun Cobalt RaQ 550 server appliance, pleasing clients that needed greater speed due to increased traffic and deployment requirements. With a 1.26-GHz Pentium III CPU, up to two 80-GB 7200-rpm disk and up to 2 GB of RAM, the Sun Cobalt RaQ 550 server delivers superior performance and reliability along with Sun Cobalt ease-of-use features.

In a competitive market, Sun Cobalt RaQ servers have helped enable BlackSun to succeed by offering its customers high performance along with low entry-level pricing. Beyond the initial affordability of the Sun Cobalt RaQ 550 server, BlackSun recognizes considerable savings through reduced repair costs as well as reduced need for expensive IT talent to maintain and update the server appliances. The server appliances are extremely reliable, often running non-stop for more than 300 days at a time.

Verio

As one of the world's largest operators of Websites for businesses, Verio has established itself by offering a highly reliable and scalable Internet infrastructure to help enterprises create and manage their on-line presence. With more than 400,000 Web sites hosted globally, and client companies in more than 170 countries, Verio needs dependable, flexible and scalable servers to provide content quickly and without fail.

To meet demanding uptime requirements from its business customers, Verio places special emphasis on using server technology with a proven track record. Reliability is everything in Web hosting and Verio has deployed over 200 Sun Netra SPARC/Solaris servers throughout its global Web-hosting infrastructure.

The increasing demand for Web-hosting services puts pressure on industry leaders such as Verio to scale capacity and processing power in real time. By combining racks of Sun Netra servers with load-balancing software, additional servers can be added quickly. Provisioning is simple and fast so that new servers can be added as needs dictate. The low power consumption of Sun's Netra servers helps keep energy costs low in spite of large server deployments.

In addition to using Sun servers for Web-hosting infrastructure, Verio also provides Sun Netra servers as a part of its standard product offering to customers — a testament to the performance and reliability these servers deliver.

Client Access Edge Computing Deployments

Digex

Digex, Inc. is a leading provider of high-end managed Web- and application-hosting services, providing secure, scalable, high-performance business solutions to some of the world's leading companies. To meet stringent scalability and performance requirements, Digex has standardized its UNIX platform-based solutions on Sun hardware and the Solaris Operating Environment. A key part of the Digex hardware infrastructure is Sun's Netra servers.

By using the Lights-Out Management (LOM) function of Sun's Netra servers, Digex is able to manage most of its Web servers from a single, central location, providing increased system availability with fewer trained people needed on-site to manage each server. Servers are monitored, administered, patched, and rebooted remotely, all without dispatching a technician to work on each server.

One of the key factors that led Digex to select Sun's Netra servers is the ease with which they can be configured and deployed. This ability allows Digex to provide on-demand scalability for its customers by reducing the time it takes to build and bring the servers online.

As Digex continues to expand operations in the United States, Europe, and the Pacific Rim, it plans to continue using Sun Netra servers. In providing its customers some of the best-managed hosting solutions available, Digex will continue to rely on Sun for the core systems that power its customers' mission-critical Web sites.

Chapter 5

Conclusion

The natural evolution of Web services will continue to drive interesting and challenging technology changes for Edge Computing infrastructure. These new approaches will allow virtually unprecedented growth for Web services and other applications as they move away from the data center and closer to customers and end-users. Understanding these changes and moving quickly to anticipate and respond to new challenges will continue to be essential for successful organizations as they compete and collaborate in changing times.

Sun views Edge Computing as a key area of focus and will continue to drive technology that makes Edge Computing infrastructure a reality. Customers demand a range of effective high-level Edge Computing infrastructure and platforms that improve performance for Web services and reduce TCO without compromising performance, manageability, or security. Increased levels of integration for both hardware and software components will be key to this effort along with corresponding higher level concepts for managing and deploying resources.

With more than 20 years of dedication to open IP networking and an end-to-end product line featuring UNIX and Linux Edge Computing platforms, Sun is uniquely positioned to help companies define and deploy effective strategies for infrastructure. As Edge Computing evolves and matures, Sun will continue to innovate and listen to customers, providing practical, cost-effective solutions as well as fostering innovative initiatives such as N1, Grid Computing, Throughput Computing, and the Sun Fire Blade Server Platform.

References

Sun Microsystems posts product information in the form of data sheets, specifications, and white papers on its Internet World Wide Web Home page at: *http://www.sun.com*. These articles can be referenced for more information on Edge Computing:

Dynamic Data Center Computing: The Emerging Service-Centric Topology, IDC, Bulletin, February 2002

Internet Infrastructure 2002: Living at the Edge, Gartner, December 2001

Why the Network Will be the Computer, ZDNET, May 2002

Emerging Core Computing Technologies, Gartner, October 2001

Server Wars; Forget about PCs, Time, January 2002

Edge Computing, tele.com, June 2001

Edge Computing: Worth the Bother, Gartner, June 2001

N1: Branding the Data Center Cloud, Illuminata, April 2002

Appendix A

Sun Certified Grid Computing Partners

The service providers listed in Table A-1 have completed Sun Microsystems' Grid Certification Course as of this writing. 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/.

Table A-1: Sun Grid certified partners (as of this writing).

Partner	Geography	Website
Access Computing Limited	EMEA (Europe, Middle East, Africa), United Kingdom	www.accesscomputing.co.uk
Applied Computer Solutions	Northern California, United States	www.acsacs.com
Avcom	United States	www.avcom.com
Avinci AG	Europe	www.avinci.biz
The BioTeam, Inc.	United States, Europe, Asia	www.bioteam.net
Cards Engineering GmbH & CO. KG	Germany, United Kingdom, USA, Belgium, Spain, The Netherlands	www.cardse.com
CGI, Inc.	Montreal, Canada	www.cgi.com
Clarity Technology	United Kingdom	www.clarity-group.co.uk
Distributor Systems International Limited	United Kingdom	www.computefarm.co.uk
Eakins Open Systems	Northern California, United States	www.eos.com
EsEda Limited	Southern United Kingdom	www.eseda.com
Esteem Systems Plc	United Kingdom	www.esteem.co.uk
GNS	Germany	www.gns-systems.com
Grid Queuing Systems	United States	www.gridqueuing.com
GridFrastructure, Inc.	United States	www.gridfrastructure.com
IBEX AG	Germany	www.ibex-ag.de
Inserve Technology AB	Scandinavia	www.inserve.se

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