An Oracle White Paper
April 2010

Guide to Deploying Carrier-Grade Services Using Open Source Technologies

ORACLE

Executive Overview

Open source technologies and commodity carrier-grade hardware can combine to create a highly available, scalable, and responsive infrastructure for deploying next-generation communications services. This paper examines a solution stack based on integrating Oracle GlassFish Server, MySQL Cluster, and OpenSolaris, Oracle's Sun Netra ATCA Blade servers, and Sun open storage from Oracle into a carrier-grade solution for converged services—without the traditionally high carrier-grade price.

Introduction

The telecommunications industry is rapidly changing, consolidating communications networks, and moving toward service delivery on Internet Protocol (IP)-based networks. Communication carriers are turning to IP networking to support virtually all of their operations, from fixed line and mobile communications to rich media, content delivery, and IP Multimedia Subsystem (IMS) applications. In addition, the rapid adoption of Voice over Internet Protocol (VoIP), Internet Protocol Television (IPTV), and video on demand is causing IP-based workloads to skyrocket. The convergence of communications networks is driving extensive opportunities for growth, enabling communications service providers (CSPs) and network equipment providers (NEPs) to launch rich portfolios of next-generation services.

In markets that feature intense competition, CSPs and NEPs must differentiate offerings to attract customers, increase loyalty, and decrease churn. At the same time, such companies are striving to establish a dominant presence in emerging markets that are characterized by razor-thin margins and low average revenue per user (ARPU). To increase revenues (especially in tight economic conditions), these vendors must build and deploy a range of attractive and creative low-cost services that can increase and retain the subscriber base.

Innovation alone, however, is not sufficient to guarantee successful market adoption. New services must be brought to market quickly, at low price points and with continuous availability. Product acceptance calls for service continuity and consistent performance, especially as services gain in popularity and usage scales. If the supporting infrastructure fails to scale in sync with demand, service response times can become inconsistent and slow, which can alienate subscribers and derail adoption. Because steady innovation tends to shorten product life spans, the infrastructure must also be extremely agile—scaling up quickly to deploy new services, scaling down as services mature, and being easily reconfigured (at the end of product lifecycles) to support the next generation of services. A low-cost, highly reliable, scalable, and agile infrastructure is critical to revenue generation and favorable business outcomes.

In recent years, the computing industry has witnessed an evolution in open source, resulting in technologies that can help meet goals for business continuity, consistent performance, and ongoing growth. Open-source-based components—such as Oracle GlassFish Server, MySQL Cluster, OpenSolaris, Sun Netra ATCA Blade servers, and Sun open storage—offer enterprise-level features and carrier-grade reliability, along with the compelling advantage of avoiding up-front, costly licensing fees.

"Using Oracle GlassFish Server, Personeta is uniquely positioned to foster innovation and address the emerging trend of open application development of converged multimedia services. Combining [Oracle's] product expertise with Personeta's extensive solutions of next-generation intelligent network and fixed/mobile convergence for business and consumer markets helps our customers self-create new revenue streams and improve customer loyalty by providing a market-leading set of differentiated, convergent multimedia services."

—Avraham Rosenbach, President and CEO, Personeta

Launching Successful Converged Services

Telecom companies are increasingly receptive to the benefits of using open source. In addition to lowering the cost of service development and avoiding vendor lock-in, open source provides transparency into the code base, which helps foster innovation, improving core technologies through continuous enhancements and new features. Today's open source offerings can create solutions that deliver levels of reliability and performance that can meet and often exceed those of proprietary offerings. In addition, because enhancements are available more quickly than proprietary technology stacks, using open source can help decrease time to market and expedite the implementation of emerging industry standards. Combining open source software with commodity carrier-grade Sun Netra ATCA Blade servers creates a low-cost, scalable, reliable, and flexible infrastructure for delivering future communications services.

A Proof-of-Concept Application Based on Open Source Technologies

Mobile World Congress 2009, in Barcelona, included a demonstration of the successful integration of open source components and Sun Netra ATCA Blade servers to construct an effective solution stack. This integration enables delivery of a new generation of converged communication services. Using an infrastructure based on Oracle GlassFish Server and MySQL Cluster, an application was developed that provides subscriber registration, presence, and session management functions typical of converged services.

From a subscriber's perspective, the proof-of-concept application presents a Web page that first enables the subscriber to register and sign in to access the service. Once registered, the subscriber initiates a call to another registered subscriber from a Session Initiation Protocol (SIP) software phone (such as the X-Lite software phone shown in Figure 1). The application supports live audio communications, text-based chat, and video display over HTTP protocol—functions typical in Web-based conferencing and communications applications.

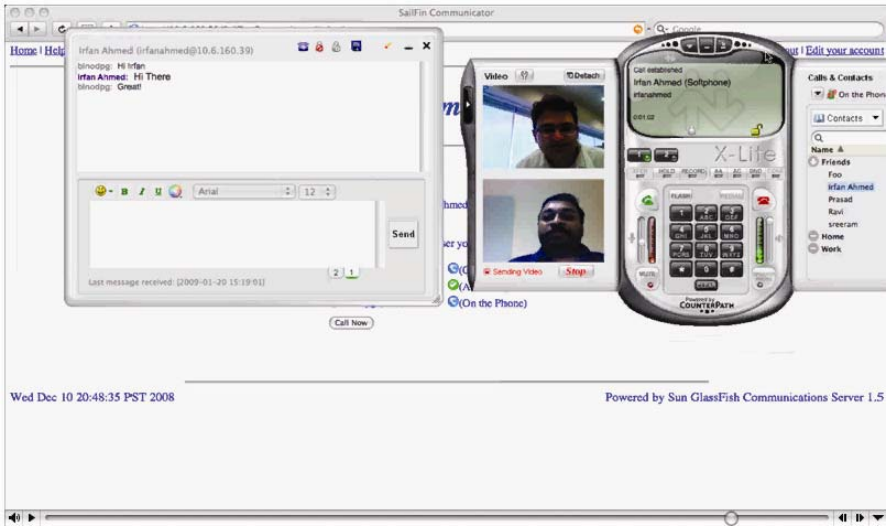

Figure 1. The proof-of-concept application presents a Web page that enables the subscriber to initiate a call to another registered subscriber from a SIP software phone.

The application shows the power of SIP servlets technology (part of Project SailFin), which adds session initiation, session media negotiation, load balancing, and high-availability features to GlassFish, a widely deployed open source implementation of the Java Platform, Enterprise Edition (Java EE) specification. Java EE is a mature technology that has been proven in more than 10 years of implementations, often as part of a platform for deploying communications services. Oracle offers a commercial edition of Project SailFin—Oracle GlassFish Server—for building Java EE applications that invoke SIP servlets. Project SailFin enables the full integration of Java-based services over telecommunications protocols such as SIP.

Perhaps most importantly, the proof-of-concept application highlights the integration of open source and commodity products to deliver carrier-grade communications services. The proof-of-concept application demonstrates the integration of the open source Oracle GlassFish Server with the highly available and reliable MySQL Cluster database.

In the proof of concept, the MySQL Cluster database stores subscriber preferences and entitlements and validates subscribers as they log in and initiate sessions. In telecom environments, it has become a leading choice for the network subscriber database and is deployed in support of a wide variety of applications, including offline/online charging gateways, VoIP, intelligent network (IN), and service delivery platforms. Developed in close cooperation with some of the largest CSPs and NEPs in the world, MySQL Cluster provides a 99.999 percent available, shared-nothing, distributed database with predictable real-time performance. In a clustered architecture, it can provide continuous data service availability in the event that a node fails due to a hardware or software malfunction.

"SailFin is an open source project supported by not only the community but also [Oracle]. We thus have all the advantages of an open source project with professional support behind the commercial product. With Oracle GlassFish Server, any Java EE developers can now create converged telecom applications such as VoIP, IM, multimedia applications—even an IMS-compliant IPTV middleware core."

— Ayse Erinc, COO, Argela

Together with Sun Netra ATCA Blade servers and OpenSolaris, the open source technologies of Oracle GlassFish Server and MySQL Cluster form a low-cost, highly available, and open-standards-based infrastructure for deploying new and innovative services.

An Open Source Solution Architecture

This paper presents a carrier-grade solution architecture based on open source software and commercial off-the-shelf (COTS) hardware technologies from Oracle. As shown in Figure 2, solution components include the following:

- **Oracle GlassFish Server.** Backed by Oracle's mission-critical support, this software adds SIP servlet support to Project GlassFish, the open source Java EE platform. Oracle GlassFish Server is a full-fledged commercial Java EE application server with clustering, high availability, and secure administration.
- **MySQL Cluster.** MySQL Cluster combines the popular open source MySQL database with a fault-tolerant clustering architecture to support mission-critical applications. Designed specifically for telecom use, this real-time, shared-nothing clustered database is a hardened version for carrier-grade applications that require continuous availability and dynamic online scalability.
- **OpenSolaris.** Backed by world-class support from Oracle, OpenSolaris provides an ideal open source platform for developing next-generation applications. With the same robust features as Oracle Solaris 10 OS—including DTrace for runtime observability and performance tuning, network virtualization, and the advanced Oracle Solaris ZFS file system technology—OpenSolaris offers performance and reliability for demanding transaction workloads.
- **Oracle's Sun Netra CP3250 and CP3260 ATCA Blade servers.** These commodity processor blades feature carrier-grade reliability and can be combined within a single Sun Netra CT 900 ATCA Blade server rack from Oracle to accelerate networking performance and deliver high throughput for converged services. The Sun Netra CP3250 ATCA Blade server features two sockets for low-power, quad-core Intel Xeon processors. It provides applications with leading-edge x64 performance and incorporates Oracle's advanced 10 Gigabit Ethernet (GbE) networking technology. The Sun Netra CP3260 ATCA Blade server is based on the UltraSPARC T2 processor, which includes built-in, high-speed packet processing capabilities

"Project SailFin offers much to Gintel and our customers. There is a clear need for a robust, carrier-grade SIP application server, and SailFin has already attracted considerable interest from our customer base. It's available on platforms that are familiar to us and our customers, which is a great proof point. Add to this the fact that it is a complete solution, offering both the middleware we need and database capabilities, backed by Oracle's global presence and reputation, and a compelling package emerges. It's definitely going to cause a splash."

—Tore Saeter, Gintel

and chip multithreading technology (CMT) and consumes less power and space than many competitive offerings. Notably, the UltraSPARC T2 processor design is based on OpenSPARC, the world's only open source reference implementation of a 64-bit, multithreaded, multicore design.

- **Sun open storage products.** Sun open storage reflects experience in open storage development, offering a choice of commodity offerings that blend best-of-breed technology, reliability, and affordability. In Oracle's Sun Storage 7000 Series Unified Storage Systems, the open systems approach combines general-purpose servers and storage components with innovative technologies and storage software. The Sun Storage 7000 Series provides a simple way to implement storage at a fraction of the cost of proprietary storage solutions. For ruggedized storage requirements, Oracle's StorageTek 6140 Array blends high availability, reliability, and performance in a single economical package. With redundancy that optimizes data protection for business-critical converged services, this array is a NEBS Level 3-compliant solution that easily scales as services increase in popularity and the need for storage capacity and performance grows.

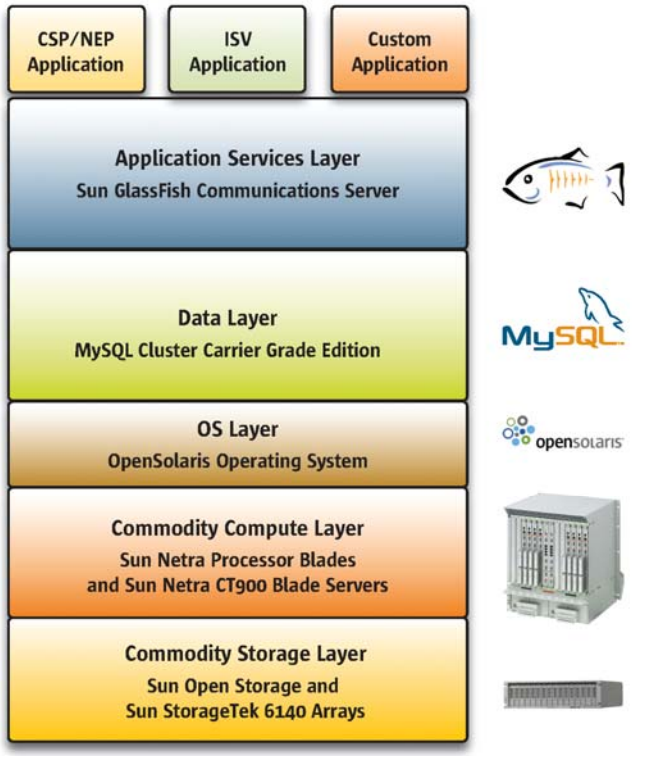

Figure 2. These software components constitute a carrier-grade solution architecture based on open source software with support from Oracle and commercial off-the-shelf (COTS) hardware technologies.

When integrated with each other, the components in Figure 2 form an open-technology-based solution stack that delivers carrier-grade capabilities without the traditional carrier-grade price. The solution is designed to meet typical requirements of converged services:

- **High availability.** The infrastructure is designed to withstand hardware or software failures without a loss in service. Immediate detection of failures causes services to be instantly failed over to other instances on operating cluster nodes (see Figure 3). Disaster recovery strategies can also replicate data and services across geographically distributed nodes. In addition, redundant server instances and cluster nodes enable software and hardware upgrades to occur without downtime, enhancing availability and dynamic scaling.

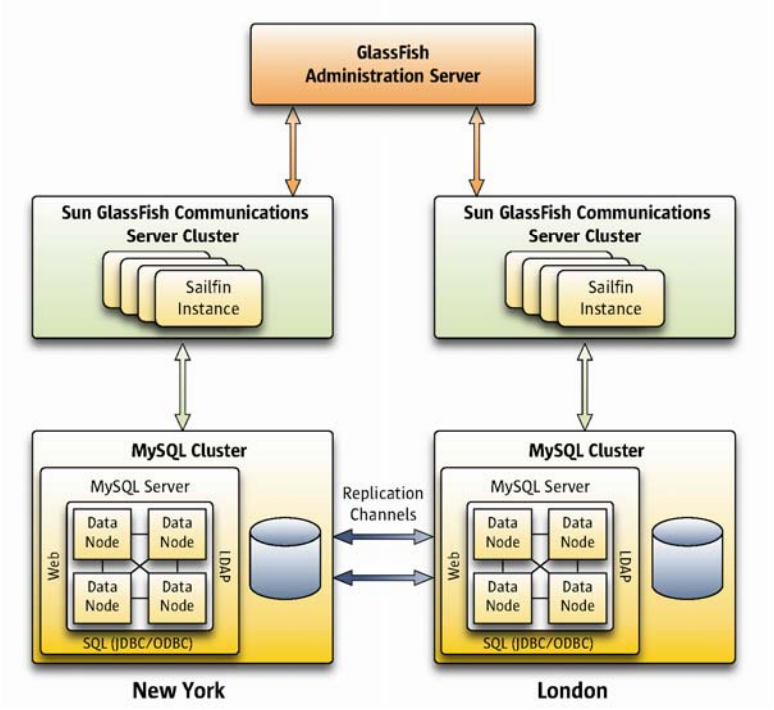

Figure 3. MySQL Cluster and Oracle GlassFish Server use redundant and distributed clusters to deliver continuously available services.

- **High throughput and optimal scalability.** To support large volumes of transactions and data, the infrastructure can expand readily, providing needed throughput as demand scales.
- **Low latency.** Response times must be short and consistent, especially as services gain acceptance and volumes grow. Generally, response times depend on subsystem processing times and are often affected by the speed of I/O operations and network connections. Carrier-grade databases such as MySQL Cluster help minimize latencies by using main memory storage of indexes and tables to limit disk I/O (with options for storing data on high-speed disks) and asynchronous operations to write log files to disk. Batching I/O operations also helps reduce network latencies.
- **Event-driven architecture.** Carrier-grade services are inherently asynchronous and event-driven. Typically, the core communications network generates events that need to be handled efficiently. A Java EE application server (such as the Project SailFin application server) that supports the Java Message Service (JMS) API and includes an enterprise service bus creates an event-driven, message-based solution.

The open source and commodity components in the solution stack help promote availability, scalability, and low latency, creating an architecture that can respond to service requests efficiently with minimal latencies. The remainder of this paper describes each component, highlighting features and capabilities that contribute to the solution's effectiveness.

Oracle GlassFish Server

Oracle GlassFish Server is built on the foundation of Java EE, which features industry-leading performance and availability. Ericsson, a leading provider of telecommunication and data communication systems, contributed scalable SIP technology to Project SailFin, extending the base Java EE application server functionality by adding support for SIP containers and servlets.

The use of a Java EE platform offers distinct advantages to carriers that are building next-generation communications services. There is a large pool of Java developer talent from which to draw, and the platform-independent nature of Java technologies prevents vendor lock-in. In addition, developers can mix SIP servlets with Java technologies such as Enterprise JavaBeans (EJB), the Java Persistence API, Java Database Connectivity (JDBC), and JMS.

The server's SIP stack is compliant with JSR 116 ("SIP Servlet API") and JSR 289 ("SIP Servlet v1.1"). JSR 289 updates the core SIP servlet API and defines a standard application programming model to mix SIP servlets and Java EE components. The ready availability of Java EE service technologies such as Web services, persistence, security, load balancing, and transactions helps accelerate development efforts—enabling vendors to bring products to market more quickly.

SIP servlets technology implements an application-level signaling protocol that creates, modifies, and terminates sessions between service endpoints. SIP sessions can be used to initiate a two-party call; a multiparty call; or even a multicast session for internet calls, multimedia calls, and multimedia distribution. Running within a SIP container, SIP servlets are similar to HTTP servlets but also support SIP. SIP technology is a foundation for many popular telecommunications services, including VoIP, instant messaging, handling chat presence, and managing buddy lists, as well as for Web conferencing.

SIP servlets enable

- The discovery of communications endpoints
- The authentication of subscribers
- The negotiation of session media capabilities
- Session creation and teardown
- Multiparty session state

Key features in the Oracle GlassFish Server that enhance availability, improve scalability, and reduce latency include the following:

- **Java EE and SIP replication.** Oracle GlassFish Server defines a redundant service layer that consists of clustered Java EE and SIP application server instances, helping provide continuous service availability. If an instance fails, the load balancer will redirect requests to other instances.

"MySQL Cluster is a product of high quality, is extremely robust, and meets our demands in terms of performance and high availability. We evaluated shared-disk clustered databases, but the cost would have been at least 10 times as much."

— François Leygues, Alcatel-Lucent

- **Overload protection.** When a server falls behind in handling SIP traffic, packets are often retransmitted. This can lead to a snowball effect that causes the amount of traffic to increase rapidly; degrading the quality of service; and, in some cases, even causing total service failure. Oracle GlassFish Server employs overload protection features that monitor system performance metrics and trigger logic to restrict traffic.
- **Converged load balancing.** In converged applications, one or more subscribers may interact with the same service and use multiple protocols (such as SIP and HTTP) concurrently. Converged load balancing refers to a unique feature in Oracle GlassFish Server that recognizes if incoming SIP and HTTP messages belong to the same session instance and, if so, routes them to the same target server.
- **Management and monitoring.** Oracle GlassFish Server supplies a management framework for configuring and monitoring instances and clusters, securely and remotely, from a Web-based central administration console. In this way, an administrator can be alerted to a potential problem and take preemptive action. Management and monitoring capabilities can often help prevent problems before they occur, which helps improve availability.

Oracle GlassFish Server promotes continuous service availability at the application tier, whereas MySQL Cluster protects data availability in the database tier. The current release of Oracle GlassFish Server has hooks to work with Service Availability Forum (SAF) specifications, enabling application services to integrate smoothly with high-availability middleware. Work continues within Project GlassFish to implement availability enhancements for full SAF integration under JSR 319 ("Availability Management for Java").

MySQL Cluster

MySQL Cluster is a real-time, fault-tolerant database that offers robust relational database functionality in conjunction with the compelling cost benefits of open source. It is designed specifically for telecom providers that must meet the stringent availability requirements of SLAs.

MySQL Cluster features a "shared nothing" distributed architecture designed for reliability with no single point of failure. Because the architecture is a hybrid model that supports both in-memory and disk-based data, it enables predictable, millisecond response times and can scale to serve tens of thousands of transactions per second.

MySQL Cluster is an open source, high-availability database that enables carriers to start building innovative next-generation services—at a fraction of the cost and time to market often associated with proprietary solutions. The database solution includes the following key capabilities that help increase availability, enhance throughput, and minimize latencies:

- **Node distribution, self-healing, and scalability.** The database system consists of multiple nodes, which are distributed to ensure continuous availability. In the event of failure, cluster nodes automatically restart, recover, and reconfigure themselves in a self-healing manner. The database system uses three types of nodes (data, application, and management nodes) in an architecture that enables flexibility and scalability over time.
- **Synchronous replication.** Data is replicated across data nodes, with a two-phase-commit protocol—this eliminates the time needed to re-create and replay log files as in shared-disk architectures (which can also add cost to a solution).
- **Geographical replication.** By replicating data to multiple data center sites, carriers can eliminate planned and unplanned downtimes. When a failure occurs, services fail over and requests are redirected to available nodes and instances, enabling continuous service delivery. Geographic replication is especially critical in facilitating disaster recovery and data reconstruction after a systemic failure or a catastrophic event.
- **Data partitioning.** The MySQL Cluster architecture scales in a linear fashion through the use of automatic and user-defined data partitioning. Partitioning enables more-efficient data access, reducing cluster node intercommunication.
- **Asynchronous writes to transaction logs.** By asynchronously writing transaction logs to disk, MySQL Cluster reduces I/O bottlenecks, which helps make response times more consistent, even for processing large transaction volumes.

OpenSolaris Operating System

The OpenSolaris OS delivers performance, security, scalability, and reliability while providing the total cost of ownership (TCO) advantages of open source. OpenSolaris is built on the foundation of Oracle Solaris 10, with ongoing six-month release cycles to furnish new features.

Underlying OS technologies (such as a high-performance networking stack, advanced file system technology, and a modern memory management model) combine to optimize throughput and performance. OpenSolaris supports near-linear scalability from 1 to 72 CPUs and addressability of up to 264 bytes of memory—well beyond the physical memory limits of even Oracle's largest server. In addition, with the ability to automatically recover from hardware faults, OpenSolaris optimizes data and service availability.

OpenSolaris includes features designed to enhance performance, throughput, and reliability, including the following:

"Telecom is mission-critical. In other words, it requires rapid, successful call completion. Think about how frustrating it is as an end user when you make a call and it takes a second or more to connect or never does."

—Mario Aluf-Medina, founder and CEO, Marteleron

- **Dynamic Tracing.** DTrace is a powerful tool that provides a true system-level view of application and kernel activities. System administrators and developers can use this dynamic instrumentation tool to identify performance bottlenecks and diagnose problems, greatly reducing the time it takes to optimize performance and resolve problems.
- **Containers.** Containers provide a breakthrough approach to virtualization and software partitioning, supporting the creation of many private execution environments within a single operating system instance. Using this technology, organizations can improve resource utilization, reduce downtime, and lower costs by consolidating services on a single system.
- **Predictive Self-Healing.** Predictive self-healing supplies built-in capabilities for diagnosing and automatically recovering from many hardware or application faults. As a result, mission-critical applications and essential services can often continue uninterrupted in the event of software failures, hardware component breakdowns, or software misconfiguration problems.
- **Resource management.** Built-in resource management facilities enable computing resources to be allocated among individual tasks in a structured, policy-driven fashion. By using these facilities to proactively allocate, control, and monitor system resources (such as CPU time, processes, virtual memory, connect time, or logins) on a fine-grained basis, developers can often reach and maintain more-predictable service levels.
- **Advanced networking features.** OpenSolaris offers sophisticated networking features that enhance capabilities and performance, including the following:
  - **Performance and latency improvements.** These include dynamic polling, hardware and software fanouts to multiple cores, and stack parallelization.
  - **Virtualization enhancements.** These include virtual IP instances within zones, software and hardware-based virtual NICs, and "virtual wire" capabilities (that is, the ability to simulate a complete network, including elements such as switches, routers) within a system.
  - **Service virtualization or "flows."** This is based on IP addresses, IP subnets, transport, and ports.
  - **Resource partitioning.** This includes segmenting bandwidth for NICs, for virtual NICs, or for flows, as well as CPU resource and priority assignment per datalink.
  - **Analytics/observability.** Real-time usage and history are available for flows and datalinks; fine-grained statistics per link are also available.

"Mobile WiMAX deployments must be driven by a flexible approach that can start small but grow to almost infinite scale. The combined solution of Aricent WiMAX software products and Oracle's unified network platform offers a powerful, scalable growth path, while slashing CAPEX and OPEX. In emerging economies, Aricent ASN Gateway Software and Sun Netra systems will speed WiMAX deployments, and we will jointly develop proofs of concept that demonstrate the benefits of reduced time to market, risk, and cost."

—Sajal Gupta, Aricent

OpenSolaris Deploys Mission-Critical Telecom Applications

In searching for a platform for developing and delivering IPJunction, a new service for the telecom industry, the company Marteleron required continuity of service, performance, and reliability. IPJunction enables direct connectivity between gateway carriers. Specifically, it bypasses the public switched telephone network (PSTN) and routes VoIP calls, text, and video directly to destination gateways over IP. A centralized gateway registry service such as IPJunction is necessary for streamlined connectivity between all carriers.

To help Marteleron get started, servers running OpenSolaris for use in development and testing were provided. OpenSolaris provides the performance, reliability, and scalability needed for next-generation telecom services.

Sun Netra ATCA Blade Servers

Sun Netra ATCA Blade servers are high-performance, COTS processor blades with high-speed packet processing capabilities that can help accelerate IP-based services and transaction workloads. The Sun Netra CP3250 Blade server is an ATCA board with dual sockets for low-power, quad-core Intel Xeon ATCA processors. The Sun Netra CP3260 ATCA Blade is based on the UltraSPARC T2 processor—the industry's first massively threaded system on a chip. NEBS Level 3-certified and compliant with PICMG 3.0 R2 and PICMG 3.1 Option 9, both processor blades are compatible with Oracle's Advanced RTM options for rear-accessible 10 GbE hard disk drives and interfaces as well as with Fibre Channel interfaces. Sun Netra ATCA processor blades provide high throughput and energy efficiency to support the challenges of carrier-grade applications.

These processor blades can be plugged in to any of the 12 slots of a Sun Netra CT 900 ATCA Blade server to create a commodity platform that can run OpenSolaris, Oracle Solaris, Microsoft Windows, or Linux, all in the same enclosure. With dual redundant, hot-swappable GbE switches, optional 10 GbE switches, shelf managers, and power and fan modules, the Sun Netra CT 900 ATCA Blade server is designed for "six nines" reliability.

"Sun Netra CMT-based ATCA Blade and rack mount servers, in addition to Sun Netra Data Plane Software Suite, enable SDC Labs to provide our telco customers with innovative services at a fraction of the previous cost. Our digital media processing and security solutions demand the highest levels of performance, reliability, and scalability. [Oracle's] technologies give our customers higher session capacity, faster throughput, and improved quality of service with lower power consumption and TCO."

—Steve DeLaney, President, SDC Labs

Sun Netra CP3250 Intel Xeon ATCA Blade Servers

Sun Netra CP3250 Intel Xeon ATCA Blade server is based on a nonthrottling embedded class of quad-core Intel Xeon processors, with two processor sockets per blade. In this processor blade, Intel Xeon processors incorporate a low-power design and power management features that minimize power consumption and reduce the need for cooling.

Features of the Sun Netra CP3260 ATCA Blade processor include the following:

- Dual-socket, quad-core Intel Xeon L5408-LV processors
- Six DDR-2 memory slots for up to 24 GB of main memory
- One x8 PCI Express (PCIe) AMC I/O slot that also supports SAS signaling between the processor blade and a Netra CP3200 ARTM-HDD Advanced RTM
- Two GbE channels connected to the base fabric and two 10 GbE channels for the extended fabric
- Two front-panel-based USB ports
- One GbE management port routed to the optional Advanced RTM through the Zone 3 connector
- CompactFlash Type II
- Support for Oracle Solaris 10, carrier-grade Linux, and Microsoft Windows 2003 Advanced Server

Sun Netra CP3260 ATCA Blade Servers with UltraSPARC T2 Processors

In the Sun Netra CP3260 ATCA Blade server, the UltraSPARC T2 processor incorporates CMT, which provides as many as eight cores and 64 threads. The chip architecture supplies increased computational density while staying within highly constrained power and cooling envelopes. High levels of integration also help reduce latency, lower costs, and provide carrier-grade reliability.

Features of the Sun Netra CP3260 ATCA Blade processor include the following:

- Single UltraSPARC T2 processor socket with processor options for six or eight cores, eight threads per core
- Eight FB-DIMM memory sockets for up to 32 GB of main memory
- Eight GbE channels:
  - Two channels for base fabric (1 Gb/sec—PICMG 3.0)
  - Two channels for extended fabric (10 Gb/sec—PICMG 3.1, Option 9)
  - Three channels for Zone 3 (1 Gb/sec—ARTM implementation-specific)
  - One channel for front panel (1G)
- One RS-232 serial craft port and two USB ports (front panel)
- 8 GB to 16 GB CompactFlash socket
- Oracle Solaris 10 and Sun Netra Data Plane Software Suite support

Sun Open Storage Solutions for Carrier-Grade Applications

In many data centers, a shift in computing technology has united industry-standard, commodity hardware platforms with open source software technologies. This trend has resulted in greater flexibility and has helped reduce costs. Now the same change is taking place with respect to storage, enabling customers to achieve greater agility in responding to changing storage needs, while lowering acquisition and operational costs for storage.

Sun open storage brings together general-purpose storage and open source software. The Sun Storage 7000 Series Unified Storage Systems—which differ from highly proprietary, special-purpose devices that often make up traditional storage solutions—are part of this initiative.

The Sun Storage 7000 Series Unified Storage Systems appliances redefine storage economics, integrating a hybrid storage pool architecture, industry-standard components, and open storage interfaces. Advanced technologies—such as low-cost JBOD arrays, solid-state devices (SSDs), OpenSolaris, and Oracle Solaris ZFS—fundamentally change the price/performance storage equation, exploiting the power of enterprise Flash and the inherent value of hybrid storage pools. Sophisticated features such as DTrace and ZFS promote real-time performance and debugging analysis, easy computational and storage volume scaling, and end-to-end data integrity. To support carrier-grade applications that require maximum availability, Sun Storage 7410 Cluster System features an architecture with no single point of failure, enabling high performance and protection against downtime at commodity price points.

Certified for NEBS Level-3 compliance, StorageTek 6140 Array is also an ideal choice for commodity storage in carrier-grade applications. This array provides an affordable entry point

that expands up to 112 TB, supporting storage capacity and performance that scales according to service volume and data storage needs. Advanced data protection features—such as redundant and hot-swappable components and automatic failover—enable continuous uptime. The array architecture supports change in a nondisruptive manner, including drive and volume additions, RAID and segment size migration, and system upgrades, providing uninterrupted data access.

Advantages of Using Open Technologies

With open source software and commodity ATCA blade servers and storage products from Oracle, NEPs and CSPs can cost-effectively build a reliable and scalable infrastructure for delivering converged services. The Oracle GlassFish Server and MySQL Cluster provide continuous service availability through redundant service layers: a layer of clustered SIP application server instances as well as a fault-tolerant data storage layer of redundant MySQL Cluster nodes.

In addition, the combination of Oracle GlassFish Server and MySQL Cluster offers synergy, in that these technologies are both data-centric in design. The converged load balancing and data-centric rules features in Oracle GlassFish Server—together with the data partitioning feature in MySQL Cluster—enable service requests to be efficiently targeted at a specific server instance and data partition. This approach increases throughput and decreases service response times, because requests are frequently directed to the server where the subscriber's data resides. As shown in Table 1, the open solution stack creates an infrastructure that meets the critical characteristics of carrier-grade service deployments, but at a price point well below that of most traditional carrier-grade solutions.

TABLE 1. CHARACTERISTICS OF AN OPEN SOURCE SOLUTION FOR CARRIER-GRADE SERVICES

ORACLE GLASSFISH SERVER	MYSQL CLUSTER	OPENSOLARIS	SUN NETRA CP3250 AND CP3260 ATCA BLADE SERVERS	SUN OPEN STORAGE
HIGH AVAILABILITY				
• Instance clustering • Replication of HTTP/SIP requests and session data	• Shared-nothing design; no SPOF • Subsecond failover • Synchronous replication	• Proven, mission-critical OS • Oracle Solaris fault management architecture • ZFS for integrity	• ATCA commodity hardware with redundant enclosure components • NEBS Level 3-certified	• NEBS-compliant StorageTek 6140 Array • Sun Storage 7410 Cluster System • Redundant/ hot-swappable components

HIGH THROUGHPUT AND SCALABILITY				
• Dynamic clustering • Scale-out by addition of instances • Data-centric rules	• Linear scalability via addition of nodes online • Multithreading for multiple cores • Transaction/ operations batching • Data partitioning	• Thread scheduler • Proven scalability • Scale out, scale up, scale within	CP3250: • Two low-power quad-core Intel Xeon processors • X64 compute density CP3260: • UltraSPARC T2 processor with as many as 64 threads • High-density server	• Scalable capacity and performance using commodity drives and open source
LOW LATENCY				
• Tuning of JVM, thread pools, and socket buffers • Load balancing to minimize latency	• Native APIs for low latency • Real-time extensions • In-memory tables to reduce disk I/O	• Real-time scheduling classes and priorities • DTrace	CP3250: • High-performance Intel Xeon processors CP3260: • Dual on-chip 10 GbE interfaces • High-performance UltraSPARC T2 processor	• SSDs that minimize I/O latency in Sun Storage 7000 Series • Hybrid storage pools to balance performance/cost
ONLINE UPGRADES AND MANAGEMENT				
• GUI/CLI admin within domains • JMX API for management • Rolling upgrade	• CLI console • Online backup, schema change, node addition	• Oracle Solaris Live Upgrade	• Built-in Netra Blade Management Suite for blade servers	• Common Array Manager software • Nondisruptive upgrades
FAILURES AND SELF-HEALING				
• Converged load balancing and failover • Self-healing • Overload protection	• Automatic load balancing • Automatic failover and node recovery	• Predictive self-healing in Oracle Solaris	• Hot-swappable components • Redundant GbE switches, optional 10 GbE switches, shelf managers, power supplies, and fan modules	• Hot sparing • Automatic failover configurations for continuous data availability

Conclusion

To facilitate the rapid adoption of next-generation services, telecom providers must take into account critical factors in designing the supporting infrastructure. Mass market adoption is successful only if services are highly available, if they scale in sync with demand, if performance is predictable and consistent, and if the infrastructure demonstrates agility with respect to ongoing change.

Recently, NEPS and CSPs are embracing the compelling cost advantages of open source software technologies and commodity hardware. By starting small and making a relatively low initial investment, they are quickly able to create solutions that provide carrier-grade reliability and real-time performance. As services grow in popularity, an open-technology-based infrastructure can scale seamlessly by adding more server instances and cluster nodes, increasing throughput, and minimizing latencies. When open source software and commodity hardware are integrated with each other, these technologies create an agile solution for delivering new generations of highly available converged services—at price points far below those of traditional carrier-grade solutions.

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

Guide to Deploying Carrier-Grade Services
Using Open Source Technologies
April 2010

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

Worldwide Inquiries:
Phone: +1.650.506.7000
Fax: +1.650.506.7200
oracle.com

SOFTWARE. HARDWARE. COMPLETE.