AdvancedTCA® and Sun Microsystems: A Technical Overview of Sun's AdvancedTCA

White Paper
February 2006

Table of Contents

Chapter 1

Introduction

Overview

As the telecommunications field becomes increasingly competitive and a greater breadth of next-generation services is demanded—from smart phones to VoIP, and video to wireless broadband and other high-bandwidth services—both network equipment providers (NEPs) and telecom carriers look for faster, more flexible, and more cost-effective means of competing effectively and meeting market needs. As a result, in order to reduce capital expenditures (CAPEX) and operating expenses (OPEX), narrow proprietary systems are being abandoned in favor of commercial off-the-shelf (COTS) systems.

Advanced Telecom Computing Architecture is a major step forward for telecom in this direction. Integrating highly available digital data processing and data-transport infrastructures, ATCA is the first open industry specification for carrier-grade equipment incorporating high-speed switched fabric technology. By delivering an open, multivendor platform architecture for carrier-grade telecom applications, ATCA ushers in a new era of time and cost efficiencies for the telecommunications industry. Sun Microsystems, Inc.—the computing industry's thought leader—is in position to provide the telecom industry with a strong, broad ATCA product line, and plans to align with evolving SAF standards as they emerge. Supporting use of multiple, high-performance processors and high-speed interconnect technologies in a highly dense blade system, ATCA provides the headroom needed for next-generation data services by overcoming the limitations of previous PICMG standards.

This white paper provides an overview of the role of ATCA in the marketplace, and outlines Sun's unique strategy for delivering highly compelling, comprehensive ATCA solutions, software, and services.

Chapter 2

Advanced Telecommunications Computing Architecture

Advanced Telecom Computing Architecture, known also as AdvancedTCA or ATCA, is PICMG 3.0 and its derived specifications—the largest specification effort in the history of the PCI Industrial Computer Manufacturers' Group (PICMG). With over 100 companies participating, including Sun, Intel, and Motorola—among many—ATCA is gaining wide acceptance as the new industry standard. ATCA is built on the input and feedback of a community of developers and experts, rather than the assumptions of a single proprietary interest. ATCA takes the already pervasive CompactPCI standard to the next level, reducing space needs, enhancing performance, and improving energy efficiency—thereby improving businesses' return on investment.

Key advancements of ATCA include the following:

- A high-capacity (up to 2.5 terabits), packet-based backplane coupled with low latency enables ATCA systems to process workloads more quickly.
- A highly scalable, switched fabric architecture—initially based on Gigabit Ethernet and eventually accommodating InfiniBand, StarFabric, and PCI Express interconnects—overcomes the I/O bottlenecks created by conventional bus-based architectures.
- Support for up to 200 watts per board and as many as 16 boards per shelf significantly increases performance density, enabling each blade to have two or more high-end processors through improved heat dissipation.
- At least two power modules and two power rails to each slot provide hardware redundancy, reducing single points of failure.
- Support for carrier-grade features such as NEBS and ETSI drive availability up to 99.999 percent.

Figure 1: This table illustrates the advancements provided by the ATCA specifications over the previous PICMG 2.x specifications.

Feature	**ATCA (PICMG 3.x)**	**CompactPCI (PICMG 2.x)**
Board area	140 sq. in.	60 sq. in.
Pitch	1.2 in.	0.8 in.
Power	200 W	50 W
Cooling	200 W	< 50 W
Redundancy	-48 V power feeds/rails	Not specified
Bandwidth options	1 Gb/10 Gb/2.5 Tb	1 Gb/2 Gb/3 Gb
Architecture	Gigabit Ethernet, InfiniBand, StarFabric, and PCI Express	PCI bus
System management	Integrated	Limited

Until recently, NEPs have had no alternative but to purchase or vertically develop in-house proprietary hardware and software to meet the demands of telecom and its high service level requirements. As traffic grows in both traditional and wireless networks, carriers require higher capacities—in the face of rising costs. As a result, both NEPs and carriers are moving away from proprietary solutions toward open standards-based, commercial off-the-shelf (COTS) software and hardware. NEPs will increasingly implement transport, control, and service network elements on ATCA platforms, particularly in 3G and above wireless infrastructures that demand greater I/O and compute capacity. The AdvancedTCA standards and other affiliated guidelines are predicted to exceed $3 billion of the total merchant market, growing at a rate in excess of 140 percent per year for the near future—displacing proprietary solutions and accelerating time to market.

ATCA is intended to meet the demands of a wide range of communications applications, as it uses a common backplane architecture to support multiple switch fabric technologies and new blade and chassis form factors optimized for communications. This architecture enables systems to change with users' needs and with the growth of new technologies. It is a new series of industry-standard specifications aimed at carrier-grade communications equipment, providing improved reliability, manageability, and serviceability as it incorporates the latest processors and high-speed interconnect technologies.

With ATCA's reduced development costs and faster time to market—but with no loss in terms of performance or quality—companies are able to focus their attention on developing business advantages for their products and services rather than "reinventing the wheel." Components built to the specifications of ATCA are scalable, cost-effective, minimize the system footprint, and deliver top performance with streamlined system management.

AdvancedTCA specifies connectors, mechanics, power distribution, blade dimensions, and system management that are robust and independent of the switch fabric link technology used. The architecture provides wider vendor and product choice and flexibility, promotes interoperability as hardware vendors develop blades, chassis, and backplanes that are integrateable. As a result, ATCA protects existing IT investments, reduces acquisition costs, and assists the migration toward COTS components for the telecom IT infrastructure.

And because of wider choice and interoperability, ATCA enables enterprises to access an expanded support infrastructure and pool of expertise, rather than relying on locked-in vendors or in-house development for support. Under ATCA, NEPs are able to lower development costs and speed time to market by replacing proprietary solution components with modular, carrier-grade COTS products that adhere to the new specifications.

Designed specifically to address growth issues within the telecom industry, ATCA includes specifications for power and performance density. ATCA supports up to 200 watts per blade and a maximum of 16 blades per chassis, with superior cooling to allow each blade to use multiple high-end processors. The specifications also support a minimum of two power modules and two power rails in each slot, to reduce single points of failure and provide hardware redundancy. ATCA also includes specifications for electronic keying, comprehensive system management, and a packet-based backplane with a capacity of up to 2.5 terabits, as well as low latency, for faster workload processing. A dual redundant management infrastructure, designed to resist service failures, is also specified by ATCA.

AdvancedTCA provides standardized platform architecture for carrier-grade telecommunication applications, with support for carrier-grade features such as NEBS and ETSI to enable 99.999 percent availability. ATCA is a plug-and-play architecture—multiple vendors' components interoperate seamlessly. Improved choice in products and vendors, reduced acquisition costs, interoperability, broader support infrastructure and pool of expertise—ATCA leverages all of these factors, eliminating marginalization and vendor lock-in, to lower TCO.

ATCA platform applications include core telecom network elements such as Gateway GPRS Support Nodes (GGSNs) and Serving GPRS Support Nodes (SGSNs), Voice over IP (VoIP), Radio Network Controllers (RNCs), IMS, softswitch, data encryption, streaming data technologies, smartphones, cellphones, and 3G and above wireless infrastructures demanding increased compute capacity and I/O. ATCA is expected to serve as the chief standard of the telecommunications industry until 2012 at the earliest, guaranteeing long-term value and avoiding obsolescence.

Figure 2: Network equipment provider needs and the Sun products and services that answer them.

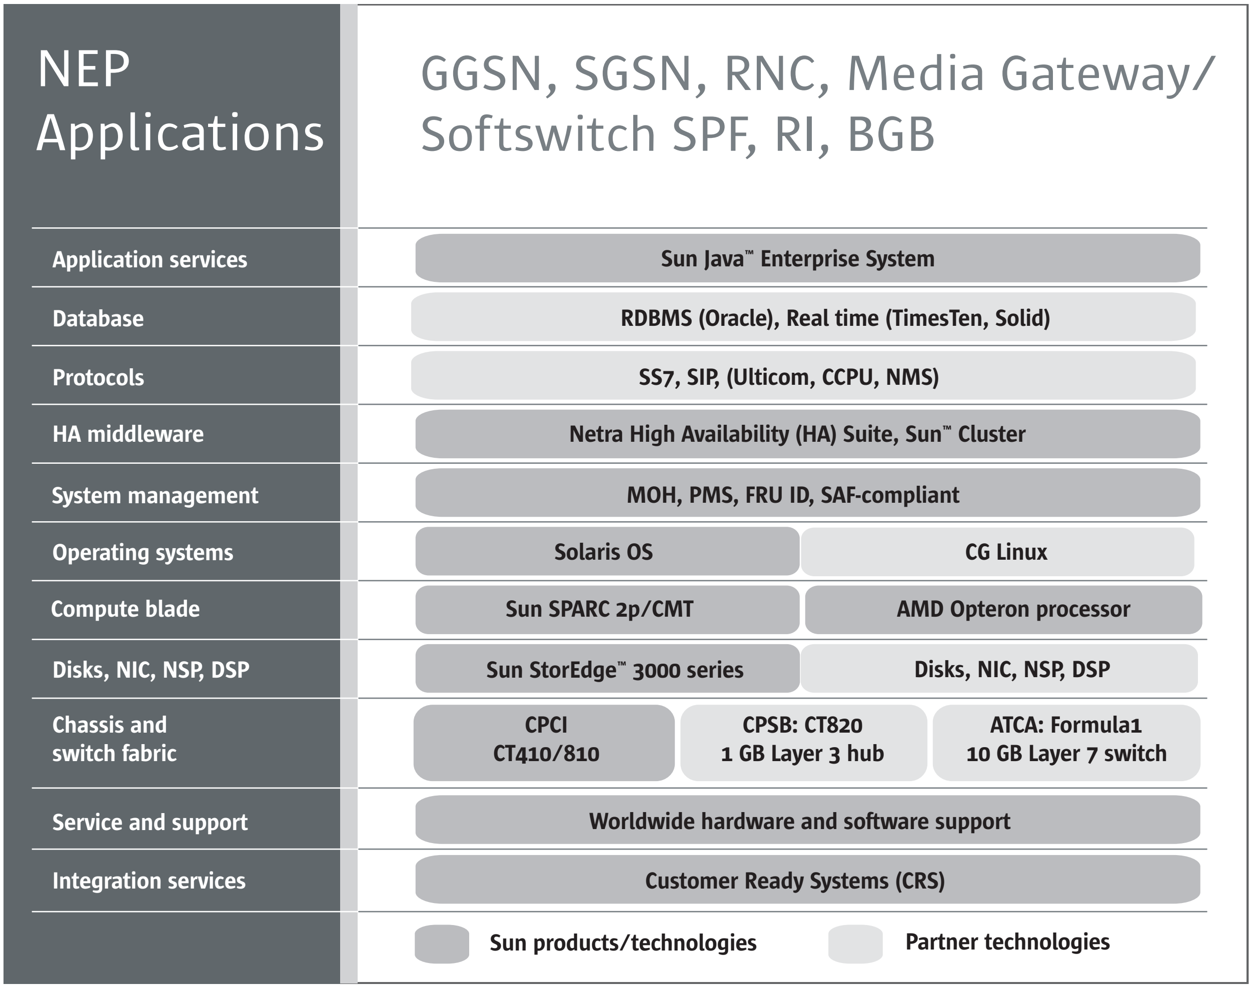

Finally, ATCA introduces the new, open hardware development framework referred to as a modular communications platform (MCP). MCP is a series of engineering specifications for the chassis and blades in a telecom system and is designed to ensure that components are physically compatible for purposes of integration.

Chapter 3
Sun and ATCA

The Sun offering is a highly scalable, switched fabric architecture that prevents I/O bottlenecks. Sun's Netra ATCA products are fully compliant with industry-standard ATCA specifications and are fully interoperable with third-party offerings. Sun's decades of experience in the telecom industry and its solid relationship with PICMG have also ideally positioned it for its role in the development of the ATCA standard. As a result, Sun has designed hardware and software systems that are ideal to meet this new standard's requirements. Sun systems offer the benefits of long industry experience, unexcelled choice, a complete ATCA system offering, and superior software system management development with world-class services.

Sun's experience with system-level hardware and software is a vital asset to numerous standards bodies in the creation of their standards and the rules for applying system management specifications based on informational models, allowing for easier integration of third-party COTS components. Sun's hands-on involvement in advancing and defining standards for the industry give Sun the necessary knowledge to meet and exceed emerging specifications. Sun firmly supports industry standards and is a long-time leader in worldwide standards organizations and consortiums, actively helping define new specifications and sharing innovations with standards bodies. This level of commitment and activity is one aspect of how Sun is working with standards bodies to improve the value of COTS components in NEPs' next-generation solutions. Among these bodies are PICMG, SAF, the PICMG-SAF advisory group, the Network Processor Forum (NPF), the Distributed Management Task Force (DMTF), and the Open Source Development Lab (OSDL).

Sun helps define the evolving ATCA standard with PICMG and uses this expertise to deliver ATCA-compliant innovation. Sun has actively participated in SAF from its inception and leads the SAF-AIS working group in helping SAF develop system management and high-availability service standards, as well as actively participating in the PICMG-SAF advisory group. Sun helps speed and ease the development of next-generation networking and telecommunications products based on network processors with NPF, and is a founding member of and active participant in the DMTF—developing management standards and integration technology for Internet and enterprise environments. And Sun participates in the Carrier-Grade Linux and Data Center Linux working groups of OSDL to help advance the use of the Linux® OS and its functionality in the telecommunications and enterprise environments.

Because multiple customers must adopt hardware standards to drive down hardware costs in the immediate future, Sun has developed ATCA components that have sufficient market interest and adhere to these new standards and is driving and adopting middleware standards to help reduce the operational costs of managing multivendor hardware. Because emerging standards have not yet addressed major issues of availability and management—presenting substantial risk of hidden costs to NEPs—Sun has adopted these new middleware standards.

Sun's ATCA solution improves reliability, manageability, and serviceability. It is compatible with PICMG 3.0—a core specification defining architecture mechanicals, power, system management, fabric connectors, and Base interface (10/100/1000 Base-T)—as well as PICMG 3.1, a specification for Ethernet and Fibre Channel Fabric interfaces. Further paving the way for high-bandwidth services, in the future Sun will meet additional ATCA specifications for incorporating multiple high-speed interconnect fabrics into central office equipment.

Sun understands the needs and concerns of NEPs and telecom providers alike, having met those needs around the world for many years. Sun has been a standard bearer in the computerization of telecommunications for decades, and has the largest installed base of any manufacturer by far. Eight out of 10 of the world's largest telecom carriers use Sun solutions, supporting over 500 million subscribers, and over 90 percent of the world's NEPs use carrier-grade Sun Netra servers. Sun has extended the same integrated information models and robust instrumentation that enable superior fault management in its CompactPCI-based Netra servers to its next-generation, SAF-compliant ATCA systems.

Carriers and NEPs are able to strategize with Sun's clear roadmap for future product design—which includes ATCA blades based on UltraSPARC® IIIi, UltraSPARC® T1, and single-core and dual-core AMD Opteron processors—and Sun's binary compatibility protects existing IT investments. Sun also plans to incorporate chip multithreading (CMT) technology into its ATCA products—carriers will realize valuable gains in throughput with processors having up to eight cores, each one able to simultaneously process up to four threads. Sun will also upgrade to 10 Gigabit Ethernet bandwidth on the backplane upon approval of related standards.

Sun's ATCA solution meets or exceeds the demanding NEBS certification process, and is designed to provide cooling and power that ensures continuous operation in even the most challenging environments. Sun already provides a broad range of products for the telecom industry, including the industry-leading Solaris OS and a complete line of Netra servers. Sun's ATCA solution expands on this choice by accommodating UltraSPARC and AMD Opteron processors, as well as operating systems from the Solaris OS to Carrier-Grade Linux (CGL)—within the same chassis. This flexibility reduces cost and simplifies systems, as well as preventing solution lock-in.

Figure 3: A simplified system diagram showing an ATCA system for data and voice services.

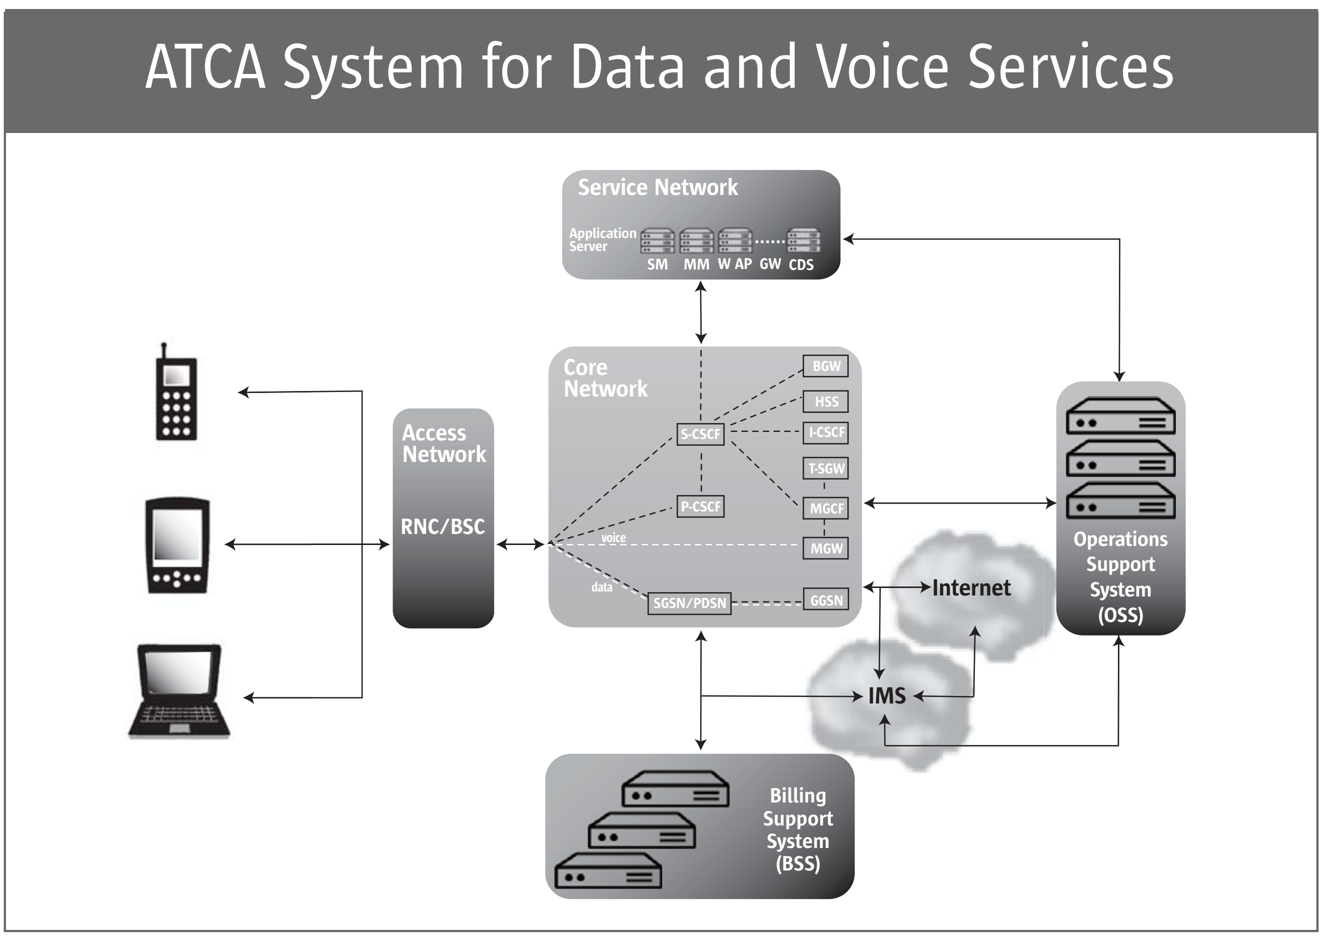

Figure 4: The Sun ATCA line offers unparalleled flexibility, with exceptional interoperability and a choice of processors, operating systems, and deployments.

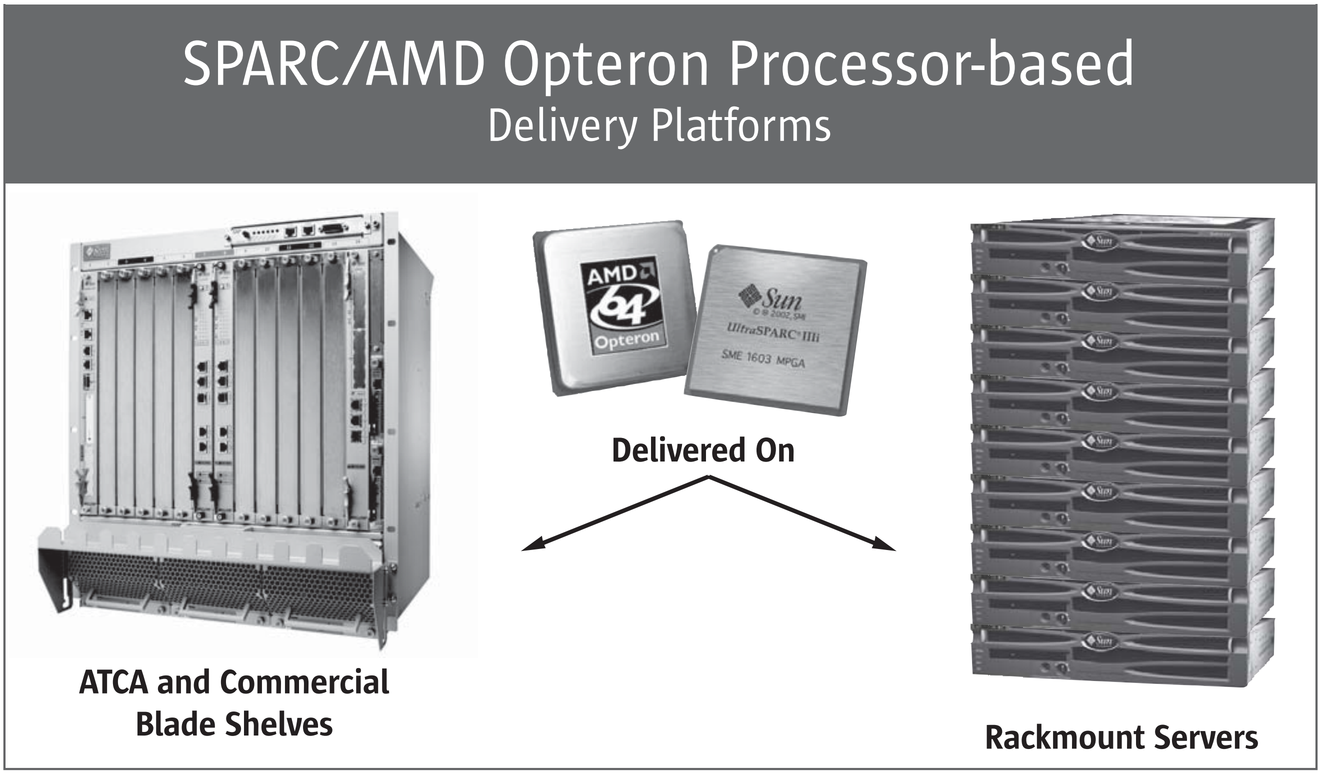

The Sun ATCA line also delivers powerful fault management and diagnostics capabilities, with immediate fault detection through hardware sensors, OS probes, and health monitors, as well as fast fail-over to hot standby nodes via the Netra High Availability Suite (NHAS) to prevent dropping of calls. (Additionally, NHAS will be SAF AIS-compliant.) Sun's service management facility disables and automatically restarts faulty components rather than shutting down an entire system, and Solaris Fault Manager in the Solaris 10 OS uses a sophisticated rules engine to identify the root cause of the first fault. Along with Sun™ Remote Services (SRS) to facilitate direct repairs, these features help ensure that NEPs are able to provide up to six 9s system availability.

Sun is a one-stop vendor for a complete, integrated system—Sun offers operating system, chassis, blades, switches, alarm cards, power supplies, and management software tools. The Sun ATCA solution is a comprehensive offering, providing a complete ATCA system from a single source—yet the Sun solution may be integrated with existing and diverse systems. Sun offers all the necessary hardware—blades, racks, switches, NEBS-certified storage arrays, and more—and recommended vendors for additional hardware that are reliable Sun partners. These partners offer approved, complementary components such as network processing blades and DSP blades, and every component is best in class and complies with both ATCA and Service Availability Forum (SAF) specifications.

A number of valuable services are also available from Sun—NEP Lifecycle Services with loaner systems, custom integration, go-to-market programs, end-to-end support, and development support; Sun™ Customer Ready Systems that ship customized, ready-to-run systems to the customer's location; OEM business practices including extended availability of products and advanced notification of product changes; and consulting, support, and training programs. Sun offers worldwide customer service and support that carriers and NEPs can rely on, with extensive design experience in delivering manageability and serviceability.

Further, Sun's ATCA solution offers total interoperability, so that all components can be mixed and matched across every ATCA-compliant vendor. Sun also provides the exclusive Solaris Application Binary Guarantee—an investment in Sun software continues to yield benefits for many years, because it is guaranteed to operate smoothly from one iteration of an OS to another, generation after generation. The industry-standard Solaris Operating System includes such valuable features as ZFS, Predictive Self-Healing, Solaris Containers, and DTrace to complement its seamless application porting on SPARC-based and x86-based architectures.

The variety of components purchased by NEPs requires a uniform system of management to keep heterogeneous solutions working together. This has been a continuing challenge for companies in the telecom industry. As new standards emerge, development of unified telecom management systems is becoming more streamlined. Sun is actively engaged in this process, uniquely overcoming the gaps in NEPs' development that result in high development costs and long development cycles.

Only Sun is in a position to carry out this role in advancing system management for telecom. Despite the standards set by the SAF, there remain no APIs defined among informational models of solution component layers—these would need to be developed in-house if not for Sun's unique, forward-thinking ATCA strategy. Sun has developed standards-based mapping layers for existing PICMG 2.x Netra products, and is extending these mapping layers to ATCA offerings that are SAF-compliant as well.

When these SAF APIs are overlaid on the Sun management stack, the standards gap is bridged. Sun management tools such as Managed Object Hierarchy (MOH) and Processor Management Services (PMS) are enhanced even further in this way, and NEPs' time to market and development costs are dramatically reduced. (Additionally, MOH will be compliant with SAF management specifications.) This system-level solution allows companies to manage groups of server racks rather than individual systems, providing powerful fault management and diagnostics capabilities by means of proven SAF compliance and API mapping at every layer.

Comparison: Sun ATCA vs. Intel, IBM, and Motorola

When measured against its competitors, Sun has a product line that leads the field. Essentially, Intel is an OEM hardware provider for third-party ATCA integrators. Third-party integrators must provide network switches, operating systems, certification and compliance testing for NEBS standards. Intel's ATCA solution combines components, testing, and certification from separate vendors—with all the complications inherent in such a mixed environment. BladeCenter T—IBM's attempt to provide a telecom-specific blade chassis solution—is, in fact, merely a product derived from their existing product for the commercial data center market. By comparison, Sun's telecom products are designed from the ground up to meet the exacting demands of telecom applications. IBM's chassis is not compatible with any other vendor's blades, and none of its products appear to be NEBS-certified. In contrast with Sun's comprehensive ATCA solution, Motorola has designed only part of an ATCA solution. Unlike Sun, Motorola only uses third-party operating systems such as Carrier-Grade Linux, requiring hardware and middle to be supported by Motorola and the OS by ODSL. Motorola's ATCA offering retains all of the complications inherent in environments with mixed-vendor hardware, software, testing, and certification—and none of the advantages of an integrated system. Sun is the only complete, fully integrated solution that meets every carrier and NEP need.

Figure 5a: A brief comparison of Sun and Intel ATCA offerings.

Features	**CP3020**	**Intel MPCPL0010**	**Intel MPCPL0001**
CPU type	AMD Opteron processor	Low-Voltage Xeon	Low-Voltage Xeon
Number of CPUs	One	One	Two
CPU speed	Min. 2.2 GHz (single-core) Min. 1.8 GHz (dual-core)	2.8 GHz	2.0 GHz
Memory	Up to eight GB	Up to four GB	Up to eight GB
Support for mezzanine cards	2x PMC 2x PIM at RTM	2x AMC	1x PMC
OS	Solaris OS, Linux, Windows	CG Linux 3.1 (validated)	CG Linux 3.1 RH AS 3.0 (validated)
Storage	2x SAS/SATA SFF 2x SAS Interfaces at RTM	SATA via mezzanine card	2.5 ATA (optional)
Solid-state storage	Compact Flash Socket (Up to two GB CF card)	Compact Flash Socket	None
Network connectivity	2x GbE (Base) and 2x GbE (SERDES) 1x async serial	2x GbE (Base Fabric) and 2x GbE (Extended Fabric) 10/100 Ethernet sync serial	2x GbE (Base Fabric) 2x FC fabric (optional)
NEBS Level 3	Certified	Capable	Capable
Standards compliance	PICMG 3.0 R1.0 PICMG 3.1 R1.0 PMC-IEEEP1386.1	PICMG 3.0, PICMG 3.1, AMC 0.3	PICMG 3.0
Availability			
Features	**CP3020**	**Intel MPCPL0010**	**Intel MPCPL0001**
Management	2 x IPMI with shelf manager	MPCMM001 Chassis Management Module (PICMG 3.x and IPMI compliant)	MPCMM001 Chassis Management Module (PICMG 3.x and IPMI compliant)
Availability	Five 9s	Unspecified	Unspecified
Hot-swap	All	All	All
Hot-plug	All	All	All

Figure 5b: A brief comparison of Sun and IBM.

Features	**CP3020**	**IBM HS20/HS40**	**IBM JS20**	**IBM LS20**
CPU type	AMD Opteron processor	Xeon	PowerPC 970	Opteron
Number of CPUs	One	Two/four	1 or 2	1 or 2
CPU speed	Min. 2.2 GHz (single-core) Min. 1.8 GHz (dual-core)	2.8 GHz - 3.8 GHz/ 2.2 GHz - 3.0 GHz	2.2 Ghz	2.6 GHz (single-core) or 2.2 GHz (dual-core)
Memory	Up to eight GB	Up to eight GB/ Up to 16 GB	Up to eight GB	Up to eight GB
Support for mezzanine cards	2x PMC 2x PIM at RTM	Proprietary I/O expansion card (not on board)	Proprietary I/O expansion card	Proprietary I/O expansion card
OS	Solaris OS, Linux, Windows	Windows, Linux	AIX, Linux	Windows, Linux
Storage	2x SAS/SATA SFF 2x SAS Interfaces at RTM	2x U320 SCSI (SFF)/ 2x IDE	2x IDE	2x U320 SCSI
Solid-state storage	Compact Flash Socket (Up to two GB CF card)	None	None	None
Network connectivity	2x GbE (Base) and 2x GbE (SERDES) 1x async serial	2x GbE 4x GbE	2x GbE	2x GbE
NEBS Level 3	Certified	Compliant	Compliant	
Standards compliance	PICMG 3.0 R1.0 PICMG 3.1 R1.0 PMC-IEEEP1386.1	Not ATCA-compliant	Not ATCA-compliant	Not ATCA-compliant
Availability				
Features	**CP3020**	**IBM HS20/HS40**	**IBM JS20**	**IBM LS20**
Management	2x IPMI with shelf manager	Alarm management software	Alarm management software	Alarm management software
Availability	Five nines	Not specified	Not specified	Not specified
Hot-swap	All	Optional PSU	Optional PSU	Optional PSU
Hot-plug	All	Blades, PSU, Ethernet switches, management module, fans	Blades, PSU, Ethernet switches, management module, fans	Blades, PSU, Ethernet switches, management module, fans

Figure 5c: A brief comparison of Sun and IBM.

Physical and Electrical Specifications

Features	**Sun ATCA**	**IBM BladeCenter T**
Chassis width	19 in.	17.4 in.
Chassis height	12-RU	8-RU
Chassis depth	15.2 in.	20 in.
Blades/chassis	14	Eight
Chassis power	2.8 kW/3.2 kW	2.6 kW
Power/blade	200 W	325 W (two slots)
Power supply redundancy	2N	2N
Inter-blade spacing	1.2 in.	2.3 in.
Compute blade vendors	All ATCA-compliant product vendors	Single source (IBM)
Industry-standard architecture	Yes	Yes

Figure 5d: A brief comparison of Sun and Mototola.

System Specifications

Features	**CP3020**	**Motorola ATCA 715/717**	**Motorola ATCA 7101**	**Motorola ATCA 7120/7221**	**Motorola ATCA 6101**
CPU type	AMD Opteron processor	Pentium M	Xeon LV	Xeon LV	PPC970
Number of CPUs	One	One	Two	Two	Two
CPU speed	Min. 2.2 GHz single-core) Min. 1.6 GHz (dual-core)	1.8 GHz	2.0 GHz	2.8 GHz	1.8 GHz
Memory	Up to eight GB	Four GB	Eight GB	16 GB	Eight GB
Support for mezzanine cards	2x PMC 2x PIM at RTM	4x (100 MHz PCI-X)	1x PMC (133-PCI-X)	1x PMC (133 PCI-X) 2x GbE interface	1x PMC (133 PCI-X)
OS	Solaris OS, Linux, Windows	ODSL CG Linux 3.1	ODSL GC Linux 3.1	ODSL GC Linux 3.1	ODSL GC Linux 3.1
Storage	2x SAS/SATA SFF 2x SAS Interfaces at RTM	Onboard 2.5-in. HDD	2x U320 SCSI On board 2.5 in. (optional)	Nothing on board (only via RTM)	Nothing on board (only via RTM)
Solid-state storage	Compact Flash Socket (Up to two GB CF card)	Compact Flash Socket One MB (BIOS/System)	2x 1.0 MB BIOS Flash 2x 64 MB flash	2x 1.0 MB BIOS Flash 2x 64 MB flash	2x 64 MB flash
Network connectivity	2x GbE (Base) and 2x GbE (SERDES) 2x async serial	1x GbE (Base Fabric) 1x 2 GbE (fabric) option	1x 2 GbE (Base Fabric) 1x GbE (fabric) option	1x 2 GbE (Base Fabric) 1x GbE (fabric) option	1x 2 GbE (Base Fabric) 1x GbE (fabric) option
NEBS Level 3	Certified	Unspecified	Unspecified	Unspecified	Unspecified
Standards compliance	PICMG 3.0 R1.0 PICMG 3.1 R1.0 PMC-IEEEP1386.1	PICMG 3.0 PICMG 3.1 (fabric only)	PICMG 3.0 PICMG 3.1 (fabric only)	PICMG 3.0 PICMG 3.1 (fabric only)	PICMG 3.0 PICMG 3.1 (fabric only)

Availability

Features	**CP3020**	**Motorola ATCA 715/717**	**Motorola ATCA 7101**	**Motorola ATCA 7120/7221**	**Motorola ATCA 6101**
Management	2x IPMI with shelf manager	PICMG 3.0 IPMI-compliant blade	PICMG 3.0 IPMI-compliant blade	PICMG 3.0 IPMI-compliant blade	PICMG 3.0 IPMI-compliant blade
Availability	Five nines	Unspecified	Unspecified	Unspecified	Unspecified
Hot-swap	All	All	All	All	All
Hot-plug	All	All	All	All	All

Figure 6: Sun meets and exceeds industry standards, offering innovation at every level.

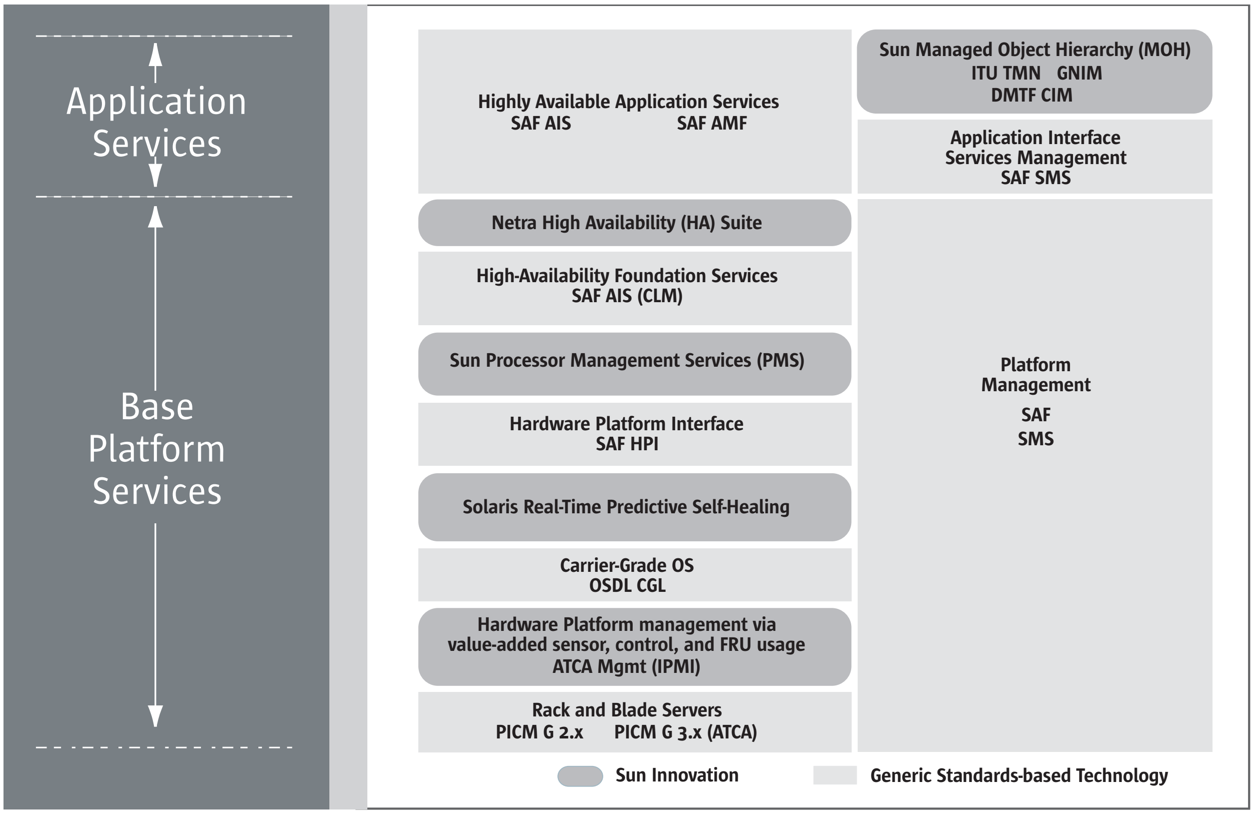

Sun is a systems company, and every Sun component is designed from the ground up with the system in mind. As a result, every Sun product is ready to be integrated into an overarching solution that fits the customer's needs and scales along with them in the future. Sun technology is built on principles of matchless interoperability, standardization, scalability, and flexibility. Sun's feature-rich ATCA solution is a compelling design that addresses multiple RFPs and set high standards in telecom.

Sun's ATCA strategy offers significant bottom-line benefits to telecom customers:

- Unparalleled choice in platform and an excellent feature-rich operating system to protect technology investments
- Unified system management of heterogeneous components lowers development and OAM costs
- Strong price/performance coupled with exceptional RAS help lower total cost of ownership
- Energy-efficient, reduced real estate needs, ROHS-compliant, and world-record benchmarks
- Proven technologies and guaranteed product lifecycles from a leading vendor reduce the risk of adopting new solutions
- High-density CMT computing maximizes throughput for increased computing efficiency
- Comprehensive ATCA offerings streamline solution development for NEPs by providing standards-based hardware, advanced software, and recommendations for third-party components
- A full complement of Sun Services, including the Sun Customer Ready Systems program, supports customers every step of the way

By turning to Sun for ATCA solutions, customers can leverage the strengths of a trusted systems company to reduce the time, cost, and risk involved in delivering next-generation telecom services.

Chapter 4

Sun ATCA Products and Services

Overview of the complete Sun ATCA solution

The complete base platform offering from Sun offers NEPs the greatest lifecycle cost savings. Sun and its best-of-breed partners provide NEPs with everything necessary to implement a standards-based IT infrastructure—working with a single vendor helps reduce integration and support costs and speed development cycles. The offering includes SAF HPI- and IPMI-compliant ATCA hardware, Predictive Self-Healing technology, processor management services, a carrier-grade OS, and NHAS—if combined with Sun's system management and high-availability stack, the gaps still left in the new standards will be addressed as well. Simply integrating multiple vendors' products in a loosely coupled fashion is not enough; the holes in the system must be addressed and will cost time and money to resolve. Sun's integrated COTS-based ATCA solution provides the necessary capabilities to users today.

Overview of the Sun ATCA product line

The uncompromising demands of the telecom industry and its customers grow every day, as carriers must meet federal regulations for service delivery and face high penalties for unplanned service outages—and NEPs, in turn, must meet increasingly high requirements, which must then be met by platform vendors through contractual obligations of systems availability. With extensive experience in the field, Sun offers numerous integrated availability and management technologies that support extremely high levels of service and help NEPs drive down the cost of delivering next-generation services. High-availability technologies such as clustering ensure service availability by providing immediate failover; fault management is built into every solution component to enable administrators to immediately detect any failure as it occurs, and advanced system management enables administrators to monitor availability of all network components. Sun technologies are proven, and offer significant business benefits.

The ATCA specifications enable hardware vendors to independently develop boards, backplanes, and chassis that will interoperate when integrated together, providing NEPS and carriers with greater choice and flexibility in designing central office infrastructures. Specification details include board dimensions, mechanics, connectors, power distribution, and a robust system management architecture that can be used independent of the switch fabric link technology. Sun has helped develop these standards, and has developed products that meet and exceed them.

Customers reduce the time, cost, and risk incurred in delivering cutting-edge telecom services when they turn to Sun's experience as a trusted systems company. Among the benefits of Sun's telecom offerings are flexibility and choice, advanced management and availability, and reduced risk. NEPs shave years off overall development cycles and reduce costs with Sun's middleware and hardware innovation, as Sun's proven availability and management technologies lower TCO and maximize uptime. Sun meets specific customer needs by offering a broad range of COTS solutions from servers to rack systems, and Sun's modular standards-based solutions provide implementation flexibility while Sun's Java™-based innovation delivers crucial portability

The ATCA system can be broken down into the following major components:

- Chassis (with its subcomponents)
- Node Boards
- Hub Board
- PMC Cards
- Shelf Manager
- AMC Carrier Board

Sun offers a wide range of products and services designed to meet the exacting requirements of ATCA and more specifically, the new paradigm of the modular communications platform (MCP). These include the ATCA CT 900 Server Chassis, the ATCA Hub Board, the Sun Netra ATCA AMD Opteron Server Blade, the Sun Netra CP3010 ATCA SPARC Server Blade, Sun Netra CP3030 AMC Carrier Blade, and Netra 210 server. This section will provide a brief discussion of these key components and their abilities, beginning with a description of the modular communications platform (MCP) chassis and blade specifications as they apply to Sun's systems.

Modular communications platform (MCP) specifications

Chassis

The ATCA chassis is 12 RU high x 493mm (19.41 inches) wide x 490mm (19.29 inches) deep, rackmountable, and provides for dual redundant input feeds for –48 V input power. Input power to the chassis connects to two isolated -48 V power sources. The chassis provides two Hub Board slots and 12 node board slots (with rear transition modules) that are each capable of supporting cooling for up to 250 W dissipation in power for each Hub Board and node board. This cooling ability is well above the 200 W per slot required by the PICMG 3.x specification; as a result, it allows higher performance blades to be installed into the ATCA system. The system supports rear access and some boards may provide both front and rear I/O connectors. In turn, this enables Sun to provide blades with performance, memory, and I/O capabilities far superior to third-party boards.

Blades

Because boards are the central building block of a system, they are compatible with multiple vendors' various ATCA boards and have the necessary redundancy required to meet the high availability demands of telecom carriers—five nines or better availability, functioning normally without any backup modules. This also enables users to use the boards wherever cost is a higher priority than system availability. These blades continue to function without interruption despite the loss of any single component, including FRU cooling units, and even in the event that a shelf manager is unavailable (however, shelf management functions are not functional in that case).

Figure 7: Sun ATCA solutions encompass the latest standards for airflow, power, and backplane support in the IT center, while maximizing the efficient use of space.

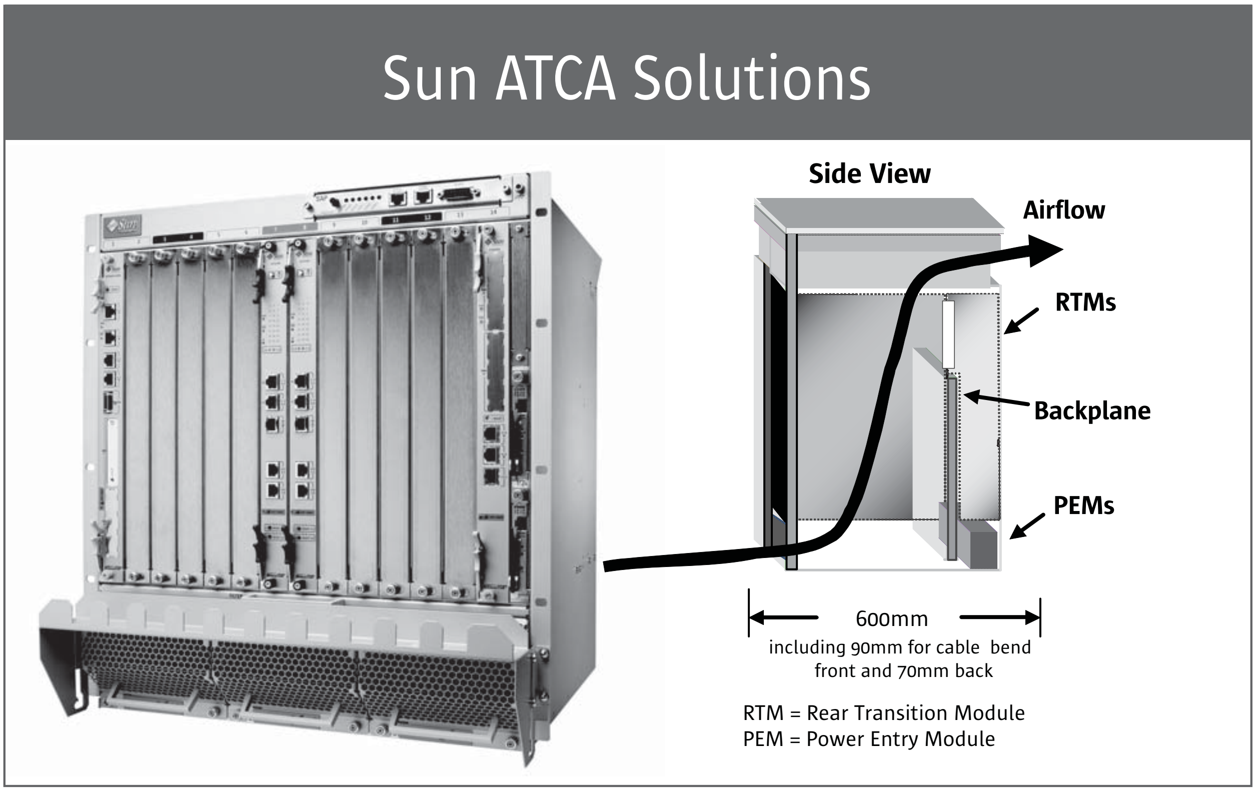

Connectors

All backplane connectors are press-fit connectors, while modules such as power entry modules (PEMs), fan tray, and filter tray use blind-mate connectors. Wherever shelf design cable connectors are used, they are polarized connectors with a latching mechanism to prevent vibration from loosening connectors.

Accessibility

Accessibility is enhanced and side interference avoided by the elimination of side-facing tool movements—serviceable components are not placed on the side of the board. Because rack designs and placements may make it time-consuming, difficult, and therefore expensive to service a shelf after it is placed on a rack, tool movements are designed to either be accessible from the front or from the rear of the board.

Rackmounting

The boards support mounting for English and metric systems, and both front and center mounting within a rack.

Backplane support

Boards with a backplane in Zone 3 support boards that do not have Zone 3 connectivity, and have shelf design connectors that are identical to those in Zone 2, for increased interoperability. The Zone 2/Zone 3 standardization enables OEMs to implement existing proprietary buses from Zone 2 on the Zone 3 area to ease the transition toward full ATCA compliance. With the backplane in the Zone 3 area, rear transition modules (RTMs) for rear I/O expansion/connectivity will not be possible in the corresponding RTM slots. The shelf design is able to support backplanes that are 8mm thick.

Figure 8: A blade configuration for the Sun ATCA solution.

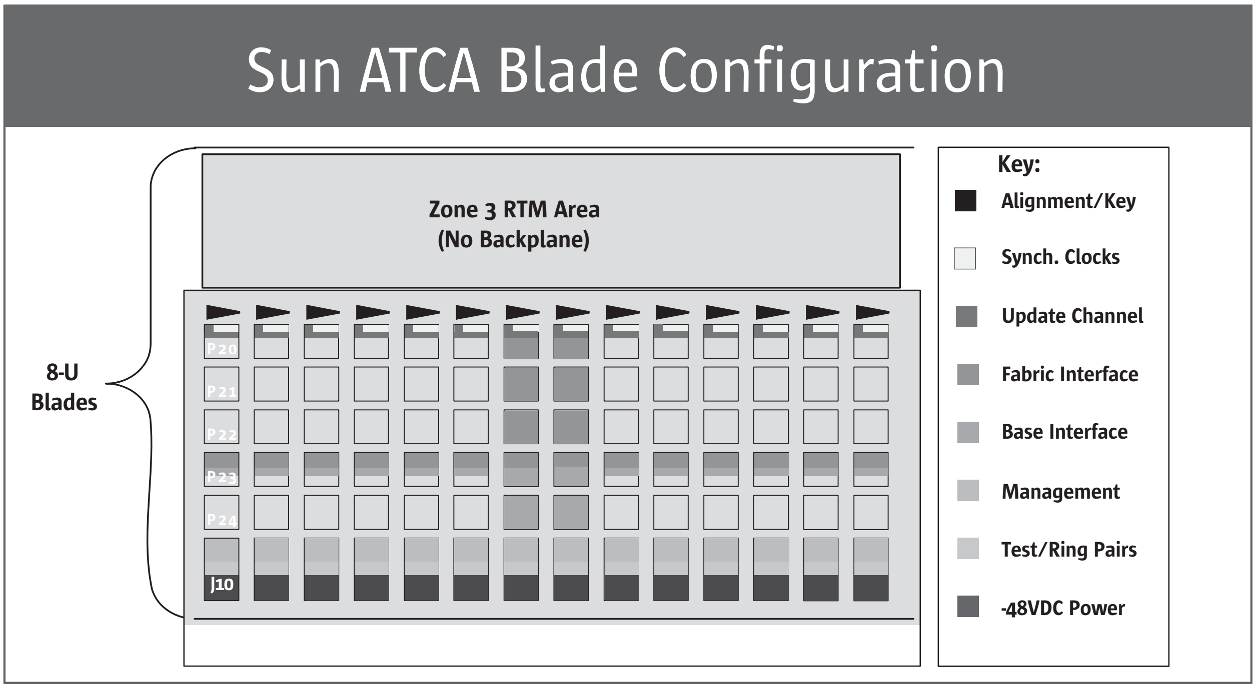

Fasteners

Metric fasteners are used to mount shelves, as well as modules within the shelves, further enabling standardization.

Handling

The shelf is able to rest stably on its front or rear surfaces without risk of falling over.

Power connections and power requirements

The ATCA architecture does not provide for power supplies within the system as in earlier cPCI or CPSB based systems (all based upon the PICMG 2.x family of specifications). Rather, each node board (and Hub Board/switch card) utilizes –48 V power from the midplane Zone 1 connector, and provides onboard DC-DC converters in order to generate the voltages required locally by the board in question. Power is rear access only, via PEMs.

Torque requirements are placed on the shelf beside the power connection and dual-hole ground lugs, and the entire power connection mechanism is designed to eliminate cable rotation under vibrations and high current. The shelf has strain relief mechanisms capable of withstanding a force of 90 lbs. The voltage drop across the power feed cables between board and power plant is 1 V or less. Because cable conductor size is often determined by the allowable voltage drop, the conductor size in the shelf is often larger than that required by UL. The cable size to be used for input in the Central Office/Install Base depends on the cable type, the number of cables, and how the cables are bundled and routed to the shelf. Cable size definition is therefore beyond the scope of this document.

The ATCA system supports node and hub boards that can dissipate up to 300 watts in the front slots and up to 15 watts for the RTMs. The system is capable of operating over a wide range of DC voltages—as input voltage falls, the input current requirements increase. As such, the maximum power supported on a per-slot basis—and therefore for the system overall—varies in order to remain within the supported systems power budget.

The power entry to the system consists of dual redundant power entry modules (PEMs). Each PEM is hot-swappable, with hot-swap and status LEDs. One PEM provides power feed A, the other power feed B. Removal of either PEM will not affect system operation, and a single PEM will support full system operation at full load. In addition, a user push button is provided to initiate a hot-swap event prior to the unit's removal. The PEMs also provide fusing, power filtering, and I2C-based FRUID with a minimum 8K SEEPROM. The units support operation of the system at 250 W per slot down to −36 V with a single PEM.

Figure 9: System power and current requirements for ATCA.

System Power and Current requirements

Slot	Input Voltage (V)	Input Power per slot (W)	Total Input Power (W)	Total Input current (Amps) (at -48V)	Total Input current (Amps) (at -36V)	Comments
Fan tray	48	N/A	360	7.5	10	Assumed power. Two fan trays
Shelf managers	48	15	30	0.63	0.83	Assumed power Two shelf managers per system
Front hub slots (x2)	48	300	600	12.5	16.67	Very conservative—for future upgrade. ATCA spec limits to 200 W
Rear hub slots (x2)	48	15	30	0.63	0.83	
Front node slot (x12)	48	15	180	3.75	5	
Rear node slots (x12)	48	15	180	3.75	5	
Totals			**4,800**	**100**	**133.33**	

Internal wiring and cable management

Printed circuit boards (PCBs) are used wherever possible to interconnect components, rather than cables—cables are prone to failure due to pinching, scraping, and cutting when components are inserted or removed. Wherever cables are used, they are fixed to the shelf to avoid obstructing field-replaceable unit (FRU) insertion and removal. The system is designed so that cables do not drop down and impede front or rear airflow, and so that service is kept easy by avoiding cables blocking door operation and component insertion and removal. The cables and fibers are also designed for 45-degree draping to reduce mechanical stress.

Subrack

The subrack for the system meets all requirements of Section 2.6 of PICMG 3.0, helping guarantee interoperability of boards in the shelf.

Field-replaceable units (FRUs)

All FRU devices are hot-swappable modules, enhancing availability and flexibility.

PMC card overview

The ATCA system provides for PMC slots on the node boards to provide flexibility in adding additional I/O functionality as dictated by the system functional requirements. This will include PMC cards for storage, Fibre Channel, and more.

PMC Hard Disk

The PMC Disk Card is used to add storage and provide a boot server when local boot disk is needed.

Key features:

- PMC hard disk with CMC Form factor (compliant to IEEE1386), single wide 32-bit/33M Hz PCI interface to carrier
- Support 24x7 enterprise-class hard disk drive, 2?-inch, 9mm hard drives
- Minimum drive capacity of 40 GB
- Solaris OS driver support
- RoHS-compliant to the new European Directive
- NEBS Level 3-certified at system level
- Reliability of 500,000 hours under continuous operation with 50 percent duty cycle when operating in temperature range of 35-40º C ambient
- Service life of at least five years under continuous operation with 50 percent duty cycle and temperature range of 35-40º C ambient
- Throughput of at least 20 MB/s in I/O mode

Telecom alarm modules (TAMs)

Wherever the system requires a dedicated TAM, it resides on IPM controller and is connected to IPMB-0.

Backplane

The backplane is a vital nexus of operations for the system. It incorporates the ring generator bus, Metallic Test Bus, power distribution, and low-level shelf management signals through Zone 1 connectors; it also incorporates Base Interface, Fabric Interface, Update Channel Interface, and Synchronization Clock Interface signals distributed though Zone 2 connectors. With so many crucial functions hinging on the reliability of the backplane, it is very important that it never fails. One method of helping ensure the reliability of the backplane is to ensure that it is entirely passive. The ATCA backplane has no active components of any sort, not even simple integrated circuits, and contains only traces, connectors, and a minimal number of resistors and capacitors. It is also compliant with the mechanical specifications described in Section 2.5 of PICMG 3.0. The mean time to replace a failed ATCA backplane is no more than three hours, although repair typically requires power, front boards, and RTMs to be removed from the shelf. All PICMG 3.0 backplanes support Update Channels (P20 connectors), which are private links between boards in any two slots. These backplanes also include an Update Channel routed between nonadjacent slots to better support wide boards, and an Update Channel interconnect diagram visible with all boards, on the shelf.

Midplane

The midplane provides dual redundant radial IPMI buses (IPMI0 in the ATCA specification) for the use of the shelf manager for system management. The midplane supports the Dual Star Base Fabric as well as the Dual Star Extended Fabric topology, and conforms to the mechanical requirements as defined in PICMG 3.0. The midplane provides 14 blade slots in total—two slots for the Hub Boards and 12 slots for node boards—with Zone 1 Connector(s) providing the system management interface between the midplane slots and the power distribution interface for telecom –48 V power. The Zone 2 connector is comprised of five ZD connectors and provides the Data Transport Interface support for four separate interfaces, providing connectivity along the backplane among up to 16 slots in:

- Base Interface
- Fabric Interface
- Update Channel Interface
- Synchronization Clock Interface

The midplane does not provide Zone 3 connectors, as each blade is provided with its own Zone 3 connector to interface with its RTM for rear access. In addition, RTM connector definition is outside the scope of the ATCA specification—in order to remain an open platform environment, Zone 3 connectors are not provided on the midplane.

Logical and physical slot mapping

Note also that the mapping between logical and physical slots in an AdvancedTCA backplane is implementation-dependent, but is optimized to reduce trace length over the backplane—thereby reducing cost for a backplane supporting only star topology. All shelves have easily legible slot numbering on the front and rear subracks, and by locating the Fabric Slots in the center slots of the backplane, the trace length is reduced by half. This makes it possible to implement topology supporting up to 3.125 Gbaud signaling rates on a low-cost PCM.

Figure 10: Example of logical and physical slot mapping in an ATCA backplane.

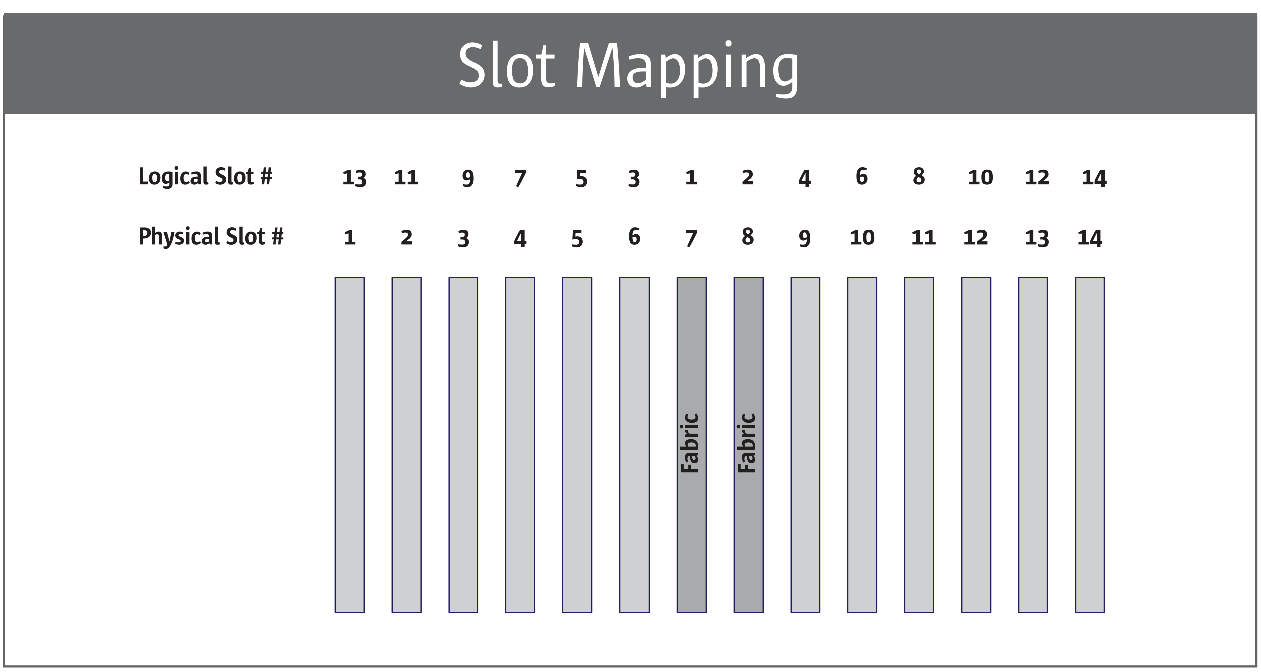

Synchronization clocks

The backplane supports three redundant synchronization clocks, per PICMG 3.0 Section 6.2.2.

Shelf Manager Card

The Shelf management system allows monitoring, control and proper operation of the ATCA node boards and other shelf components. The Shelf Management Card is one entity in the complete shelf management system (which includes the SMC controllers on the node boards and Hub Boards, as well as connectivity between these via the Midplane). One key aspect of the PICMG 3.0 architecture that exceeds the management features for previous PICMG systems is Electronic Keying. The shelf manager determines compatibility between the interfaces supported by the system and Boards for the Base Fabric, Extended Fabric, Update Channel Interface, Synchronization Clock Interface, and Metallic Test Bus Interfaces, and enables compatible interfaces only in order to prevent malfunction or board damage. Dual Redundant Shelf Managers is provided for greater system availability. Each shelf manager connects to each Hub card for redundancy. The Shelf Manager provides I/O (including telecom alarm I/O) via the front panel.

Zone 2 Routing

The backplane complies with PICMG 3.0 Sections 6.2.3.1, 6.2.3.2, and 6.2.3.3, and its Fabric Update Channel Interface is capable of supporting a 3.125 Gbaud signaling rate.

Platform services

A high availability (HA) application running on the ATCA platform leverages the management interface (IPMI) and operating system services to manage the platform.

Intelligent Platform Management Interface (IPMI)

The core components of the Extended Platform Management Daemon is based on IPMI protocol, PICMIC 2.9. This interface specification enables system monitoring and control between hardware devices and system management software. IPMI is standard message based protocol with a set of commands for operations such as setting threshold.

Possible implementations of the IPMI interface include:

- Keyboard controller style (KCS): Utilizing the legacy KCS interface to transfer data
- System management interface chip (SMIC): Alternative to KCS interface
- Block Transfer (BT): For higher performance option
- IPMI-over-LAN (RMCP): For IPMI messages sent and encapsulated in remote management control protocol over UDP datagrams
- Serial/modem: Transferring IPMI messages over serial or modem connections
- Other: IPMI messages can be careered over other interfaces such as ICMP, IPMB, and PCI buses
- An IPMI controller provides IPMI services to the platform management software

IPMI services include:

- Sensor
- Environmental monitoring
- Sensor data record (SDR): Describes sensors and device entities, such as the type, events, and threshold information
- Repository: Central storage repository of data records
- FRU inventory: FRU information

- System event log (SEL)
- Historical log of events generated by platform
- Watchdog timers
- Selectable time-out actions and automatic logging of time-out events
- System power management
- Control of system shutdown or power on the system

Distributed resource management

Management software or management agent (MOH) runs on each computing blade and/or Shelf Management Controller in a Sun ATCA system. The MOH agent can provide local resources management, shelf resources management, or rack resources management, depending on the role it plays.

The Local MOH provides the following functionalities:

- Provide local view of resource management
- Cooperates with the shelf MOH, by making available its local resources for shelf view via object proxy mechanism

The Shelf MOH provides the following functionalities:

- Constructs chassis view, based on resources under the Shelf Management Controller and resources from Local MOH agents it discovers
- Updates the shelf topology tree based on events (device insertion or removal, node card hot-swap, and so forth)
- Cooperates with the Rack MOH by making available its shelf resources for rack view via object proxy mechanism

The Rack MOH provides the following functionalities:

- Constructs the rack view based on resources from shelf MOH agents it discovers
- Updates the rack topology tree based on events from Shelf MOH agents (device insertion or removal, node card hot-swap, and so forth)

It is assumed that the system management bus is already configured properly when starting up the Local MOH, Shelf MOH, or the Rack MOH and the MOH agent will make connection to the system management bus on start up. View and Rack agents use this bus to discover other MOH agents and to proxy their managed objects

Operational and regulatory specifications

The ATCA system and components are designed to comply with Sun Specification 990-1151-04 REV B, and are tested to the enterprise levels specified in Section 4 and defined in Table 1 of 990-1151-04 REV A; the acceptance criteria for the product are those given for Work Group products. The system and components also meet the specifications of NEBS GR-63 (Level 3), ETS 300-019-2-1 May 1994 Class T1.2, ETS 300-019-2-2 September 1999 Class T2.3, ETS 300-019-2-3 September 1999 Class T3.1E, and ETS 300-753 October 1997.

Advanced fault management and availability

Sun designs its hardware with numerous intelligently placed sensors, OS probes, and health monitors for rapid detection of system problems. These tools continually monitor a broad range of parameters—including active states, voltages, and temperatures—to ensure that components function as required and expected. If a tool detects an event falling outside set thresholds, it executes a previously defined action or triggers an alarm to accelerate recovery. Each layer in the stack—software, I/O, OS, and hardware—has a standard information model applied to it, and the instrumentation built into the hardware has been mapped into usable models. As a result of this comprehensive monitoring capability, NEPs are enabled to have in-depth knowledge at each layer within vertically integrated solutions. Administrators are no longer required to pull boards to check for pin damage, send components out for replacement only to discover no problem found, or to change power feeds—with faults immediately located and isolated, efficiency increases and operational costs are reduced.

Key features of Sun ATCA fault management and availability:

- Immediate fault detection through hardware sensors, OS probes, and health monitors
- Fast failover to hot standby nodes via the Netra High Availability Suite (NHAS) to prevent dropping of calls
- Speedy recovery through a service management facility that only disables and automatically restarts faulty components instead of taking down the entire system
- Real-time diagnosis using the Solaris Fault Manager in the Solaris 10 Operating System (OS) to identify the root cause of the first fault through use of a sophisticated rules engine
- Straightforward repair by leveraging Sun Remote Services (SRS) Net Connect to automatically send service calls to a given NEP or Sun Support and to provide actionable messages that clearly identify repair part numbers and repair procedures

Because Sun's innovation ensures that monitoring information reaches both high-availability middleware for service failover and the service management interface for availability monitoring, if a failure occurs in a given power domain, the power domain sensor immediately notifies both the system management software and NHAS. NHAS immediately recognizes the loss of the components maintained by that power domain and moves services to other boards, speeding recovery. Sun's built-in innovation also saves NEPs considerable time by eliminating the need to retest a solution whenever components are moved or changed.

Chapter 5

Sun Hardware

Key ATCA hardware features:

- PICMG 3.0 and 3.1—telecom industry standards, 3.1 defines Extended Fabric
- Redundant Star Base and Extended fabrics
- 14-slot 12-U (height) chassis
- 12 node board slots—8-U x 1.2 in.
- SPARC: 2x UltraSPARC IIIi processors, up to eight GB, 2x PMC mezzanine, 2x serial
- Single-socket 939-pin AMD Opteron processor (single-/dual-core-capable)
- AMC Carrier Blade—Up to four AMC-connected mezzanine cards
- Two Hub Boards: Dual Star 24-port Gigabit Ethernet Switch Fabrics, L2 or L3/4 VLANs, trunking, IEEE 802.3z, 802.1Q, 802.1P, 802.1ad
- Two system management cards: Redundant Radial IPMI Management Controllers
- Redundant -48 V DC Power Entry Modules (PEMs)
- Midplane-enabled full Extended Fabric routing for future 10 Gigabit XAUI
- Chassis mechanicals: 12 RU high x 493mm (19.41 in.) wide x 490mm (19.29 in.) deep
- Rack compatibility: Traditional 19-inch, universal 19-inch, 600mm ETSI, 23-inch and 24-inch racks
- Front-to-back forced air flow
- Rear-access cabling (with exception for HUB board egress ports)
- RoHS compliance
- NEBS Level 3-certified
- High-density CMT computing maximizes throughput for increased computing density

ATCA CT 900 Server Chassis

The new ATCA CT 900 Server Chassis is the only chassis capable of accommodating three different kinds of blade (AMD Opteron processor-based, SPARC processor-based, and Advanced Mezzanine Card (AMC)-based) if necessary, providing unparalleled flexibility and choice. The chassis is designed to enable the creation of high-performance systems with maximum availability and low cost. The chassis is designed to mount in traditional 19-inch, universal 19-inch, 600mm ETSI, 23-inch, and 24-inch racks. The midplane is engineered so that the system provides 240 watts for all node boards and both hub boards, as well as 15 watts per RTM slot, with one PEM failure with input voltage to –37 volts DC. The ATCA CT 900 Server Chassis comes with an optional rear-entry module, 14 PICMG 3.x-compliant node board (blade) slots, two PICMG 3.x-compliant Hub Board slots, a nominal operating voltage of –48 volts DC, node and hub boards that can dissipate up to 300 watts in the front slots and up to 15 watts for the RTMs. With front-to-back airflow cooling and enhanced rack compatibility, Sun's ATCA CT 900 Server Chassis is a solid match for the telecom industry's needs.

ATCA Hub Board

The Hub Board in the ATCA system is analogous to the switch card used in PICMG 2.16 cPSB systems; it provides switch functionality and interconnectivity across the Midplane rather than requiring cables to connect node boards. The Hub Board is hot-swappable, providing high availability and extended functionality for ATCA systems. Dual redundant Hub Boards are used to provide a Dual Star configuration for the ATCA system.

Key features of the Hub Board include:
- Base Fabric interface to connect to node boards, between Hub Boards, and to the shelf manager
- Extended Fabric interfaces to connect to node boards
- Out-of-band management ports
- Layer 2 functionality, upgradeable to Layer 3 and beyond
- Front egress ports
- IPMI management functions in compliance to PICMG 2.9 and the PICMG 3.0 specifications

Due to its built-in redundancy, the ATCA Hub Board is a highly available solution with a managed Layer 2 switch with VLAN tagging support that may be upgraded to Layer 3 and beyond. It has superior networking, I/O, and throughput with 16 base Fabric Channels supporting Dual Star 10/100/1000 Base-T and 15 Extended Fabric Channels supporting dual star 1000 Base-BX, as well as FRUID support and dual IPMI controllers supporting dual redundant IPMB buses. The Hub Board also features enhanced load balancing for further availability.

Netra servers

Sun's Netra rack and blade servers are not merely compliant with NEBS Level 3 but are certified to meet that specification, having passed tests showing they can operate under demanding conditions. Netra servers are ideal for rugged environments demanding continuous availability, easy management, and optimum horizontal scalability. ATCA blades are the equivalent of rackmounted servers for form factor, and these systems can utilize rackmount (for small to medium applications) or blade deployment (for large applications) or both.

Netra CP3020 ATCA AMD Opteron Server Blade

The high-performance, single-board Netra CP3020 ATCA AMD Opteron Server Blade has single-core and dual-core processor capabilities, utilizing AMD Opteron 100 series processors (1.8 GHz for dual-core and 2.2 GHz for single-core). It supports two PMC mezzanine cards and two PIM cards at the RTM, with up to eight GB of main memory. The Netra CP3020 ATCA AMD Opteron Server Blade is hot-swappable and hot-pluggable, with five 9s availability and NEBS Level 3 certification. The server also includes two SAS/SATA SFF and two SAS interfaces at the RTM. The blade is designed to interoperate with Sun and third-party ATCA systems. The blade is designed for high availability and Switched Network Computing environments, with longevity and breadth of application in mind, and is and compliant with ATCA specifications PICMG3.0 and PICMG3.1. The board provides one 939-pin AMD Opteron Processor socket to support single-core or dual-core processors, and provides four standard DDR-1 unbuffered DIMMs installed as pairs. It supports ECC as well, and achieves far greater performance levels than prior “cPCI” standards-based products targeted for telecom markets.

Figures 11: The Netra CP3020 ATCA AMD Opteron Server Blade.

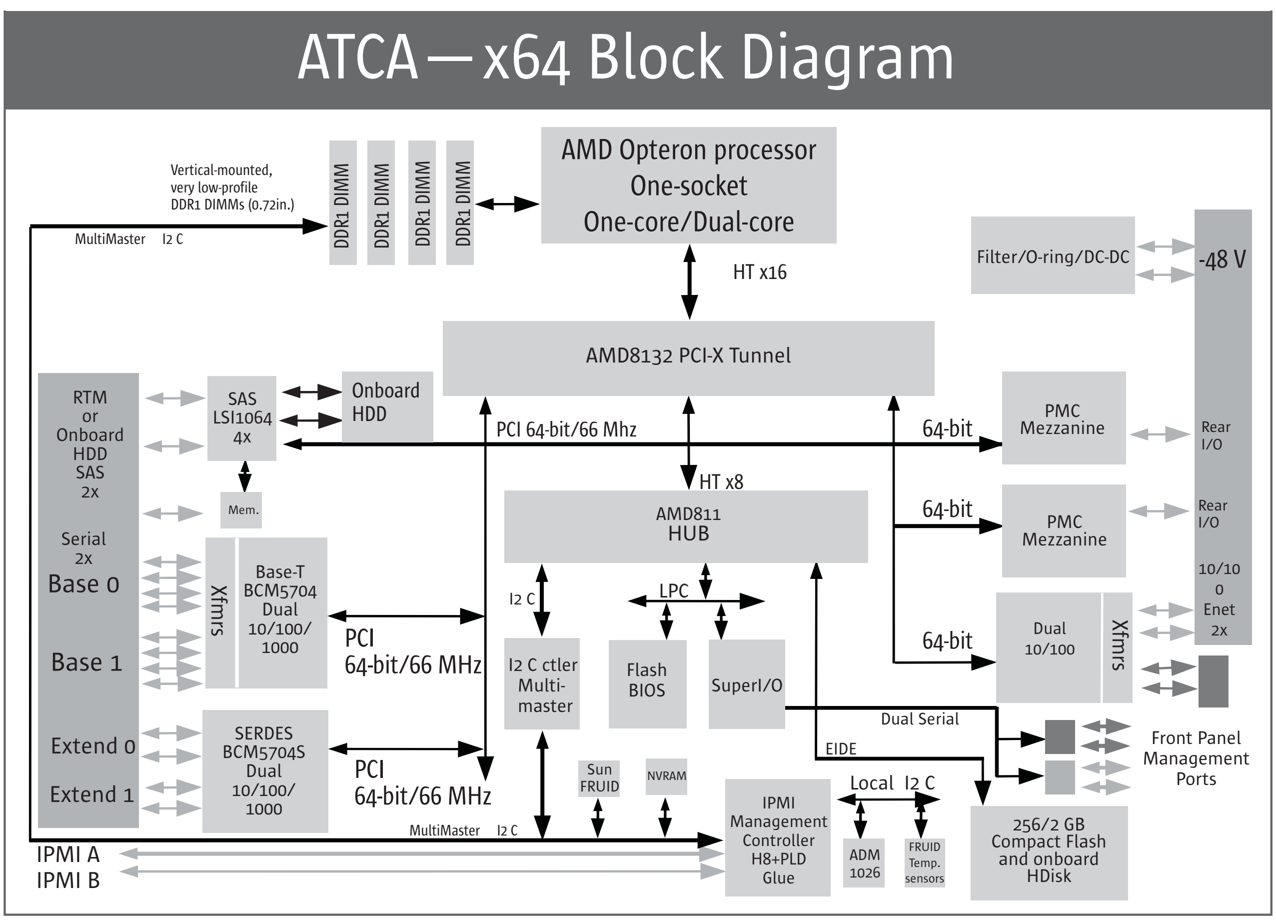

Netra CP3010 ATCA SPARC Server Blade

Enhanced cooling, increased bandwidth, and multiple high-speed interconnect fabrics are major features of the Netra CP3010 ATCA SPARC Server Blade, powered by the high performance and low latency of UltraSPARC IIIi processors. With an onboard power supply deriving local power from redundant –48 V from the midplane and flexible I/O with PMC I/O expansion, the blade has up to eight GB of main memory and supports the industry-proven Solaris OS. Two Tomatillos "Northbridge" chips enhance the system's available bandwidth and PCI/IO subsystems for even greater efficiency. The blade is designed to interoperate with Sun and third-party ATCA systems. The memory configuration of the board is also determined by the install base of the board: In Sun systems, the board can provide maximum memory and maximum CPU speed, since the maximum power supported in a Sun ATCA chassis is up to 300 W. Within the Sun ATCA chassis, the board operates with maximum memory and CPU speed configurations to provide the highest performance within a single slot. Under these conditions, the board will dissipate up to 150 W of power—hence will only operate at these rates within the SUN ATCA chassis, since only this chassis will allow for cooling of up to 150 W per slot.

Figure 12: The Netra ATCA SPARC Server Blade.

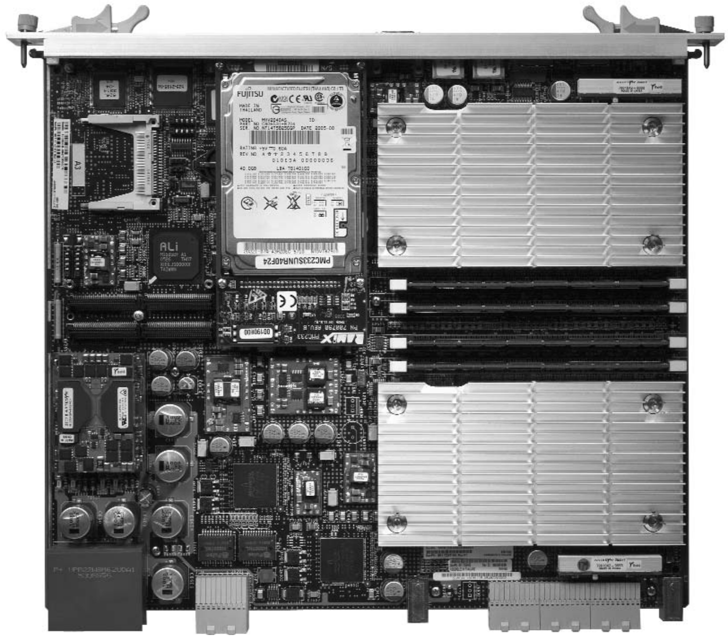

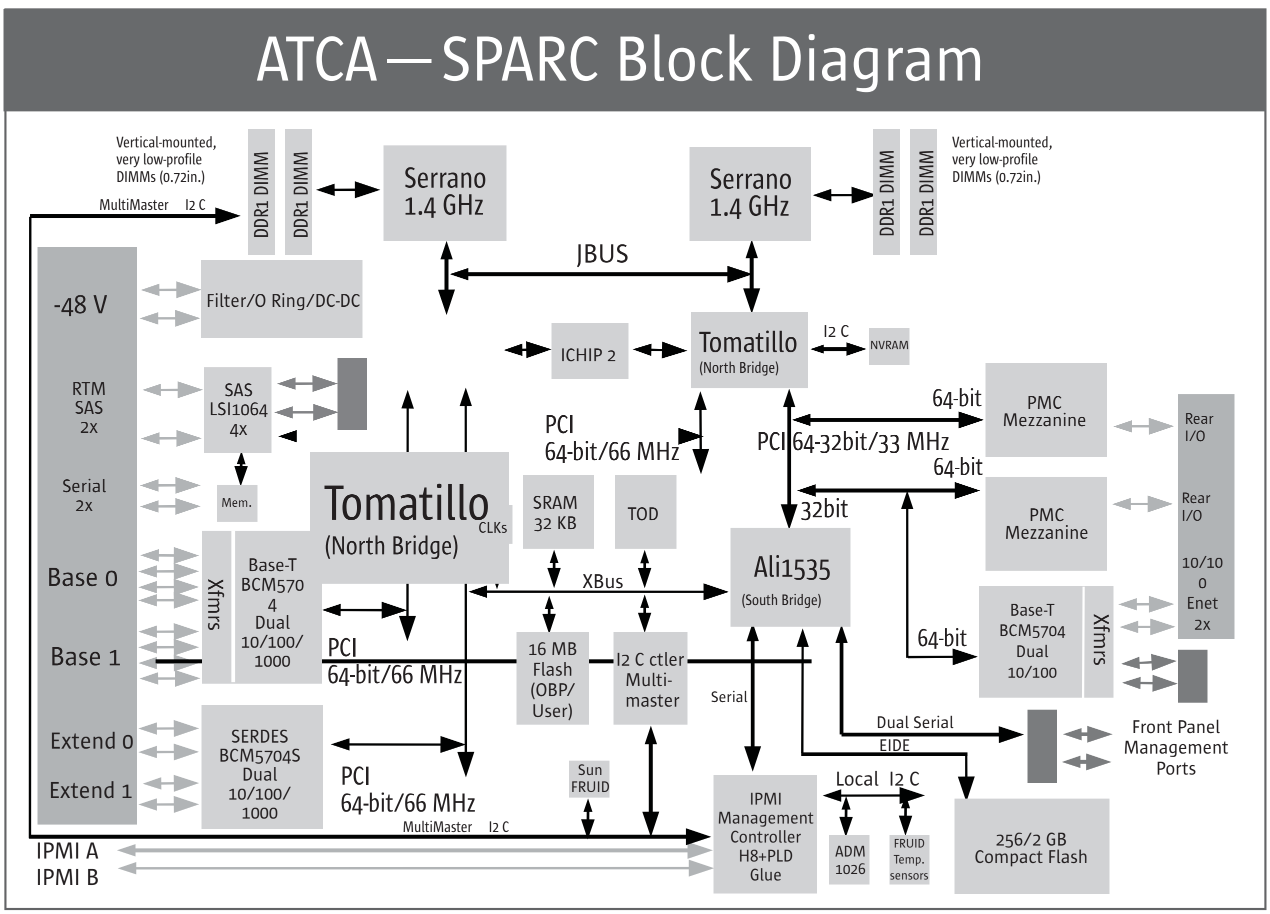

Netra CP3030 AMC Carrier Blade

The Sun Netra CP3030 AMC Carrier Blade includes hot-swappable AMC module support and a switch fabric-based, fully managed architecture. Featuring standard integrated interfaces, Dual Star 10/100/1000 Base-T, and 300-500 MHz 32-bit or 64-bit processors, the Sun Netra CP3030 AMC Carrier Blade offers superior power and I/O. The blade is designed to interoperate with Sun and third-party ATCA systems. The new generation of Advanced Mezzanine Cards, while not compatible with earlier CMC/PMC mezzanine cards, allows high-speed serial interface support and AdvancedTCA optimization.

Chapter 6

Sun Software and System Management

In order to control the costs of operations, administration and management (OA&M) and meet the high availability demands of the telecom industry, NEPs must aggregate COTS components into a unified system that can be managed as an integrated service. To support NEPs' drive toward COTS, Sun has developed an advanced system management stack for its CompactPCI-based Netra servers—innovative middleware that integrates with Sun's hardware to bridge the standards gap. Sun's system management stack provides integrated information models enabling system management at the hardware, I/O, operating system, and application levels. Once integrated with a NEP's element management system, it enables system management at the solution level as well. And Sun overlays next-generation APIs onto its existing technologies as new standards emerge, making COTS components more efficient and attractive to NEPs.

Key ATCA software features:

- Carrier-grade operating system (Solaris OS on SPARC/x86, OSDL CGL)
- Platform (system) management (PICMG 3.x system management, SAF)
  - Base systems management framework (JCP JMX)
  - System management model (DMTF CIM)
  - View models (TMN GNIM, SAF AIS/SYSMGT)
  - Management protocols (IETF SNMP, JMX RMI, DMTF CLI, CORBA IDL)
- Systems management bus
  - Fabric connectivity (redundant)
  - VLAN (private/redundant)
  - CGTP (private, reliable IP subnet)
  - Flow-based QoS support at fabric switch and OS end-points
  - DHCP/DNS (private-internal system level access only)
- Centralized systems management
  - JMX framework proxy mechanism
  - Rack-level system management bus (auto-configured)
- Integrated Services
  - High availability (SAF AIS)
  - NFS, DHCP, DNS, NTP, FTP
- Management application
  - WBEM (DMTF)
  - NMS/EMS (TMF OSS/J)

Figure 13: A simple diagram illustrating a Sun ATCA system seamlessly integrating multivendor companies and operating systems.

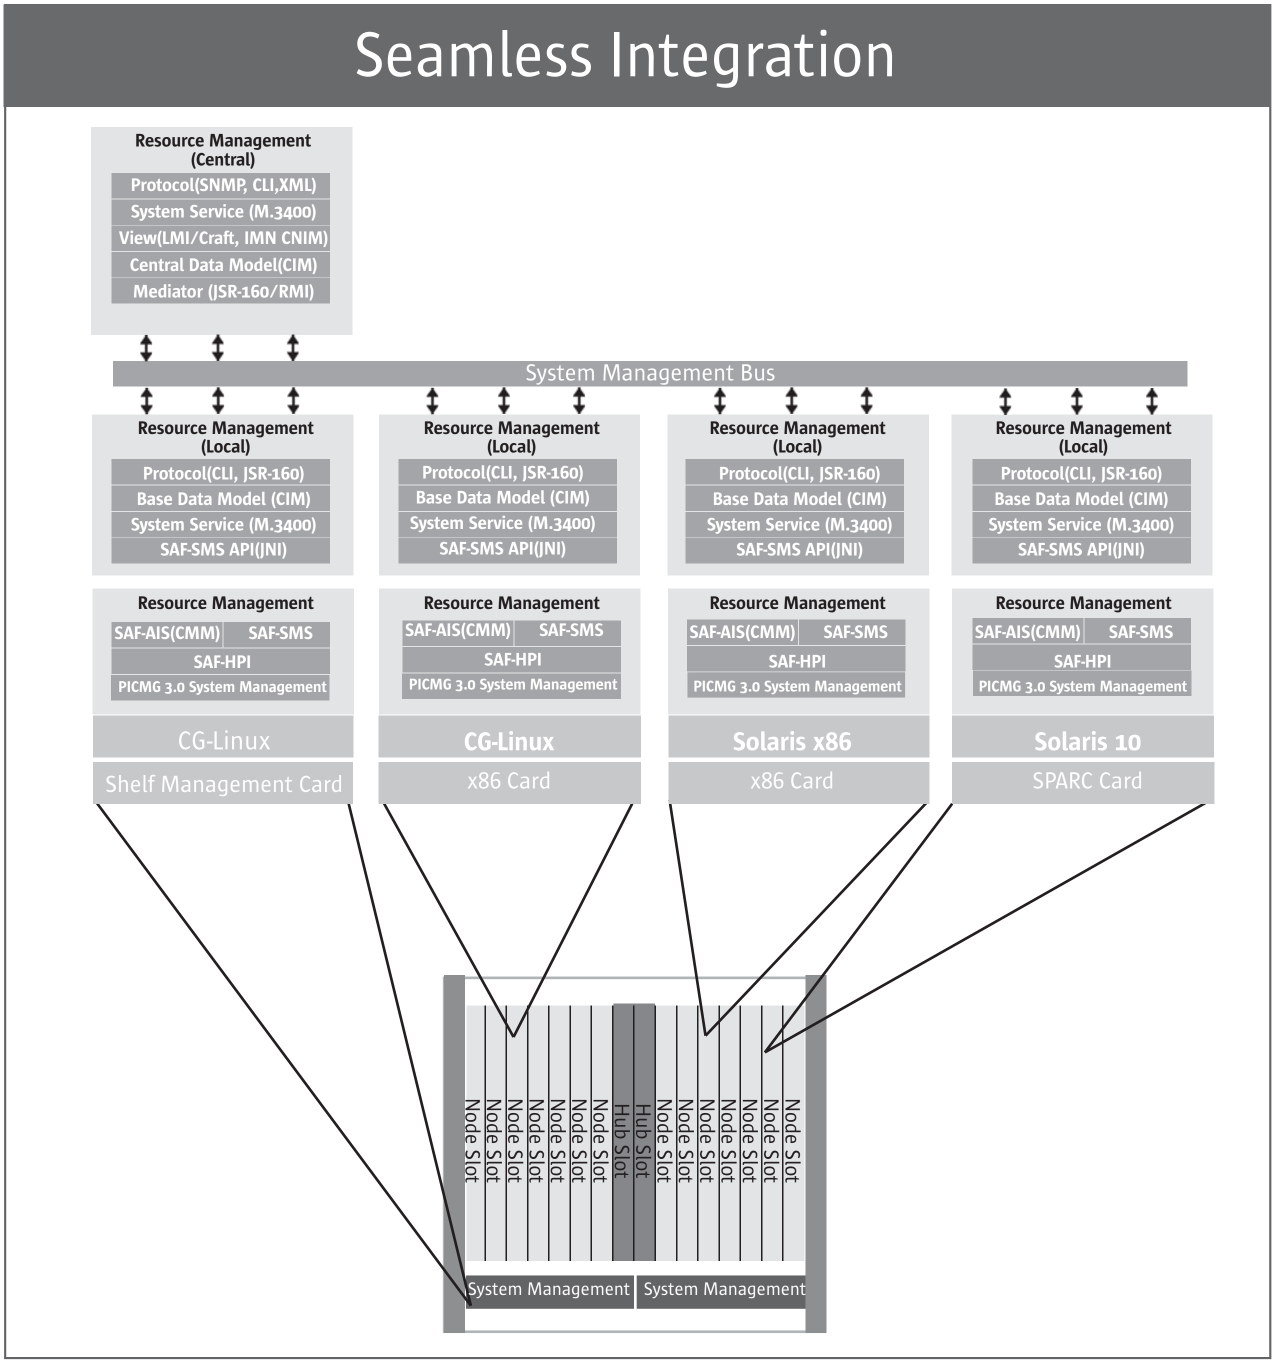

Figure 14: An example of the common information modeling (CIM) process for Sun ATCA system management.

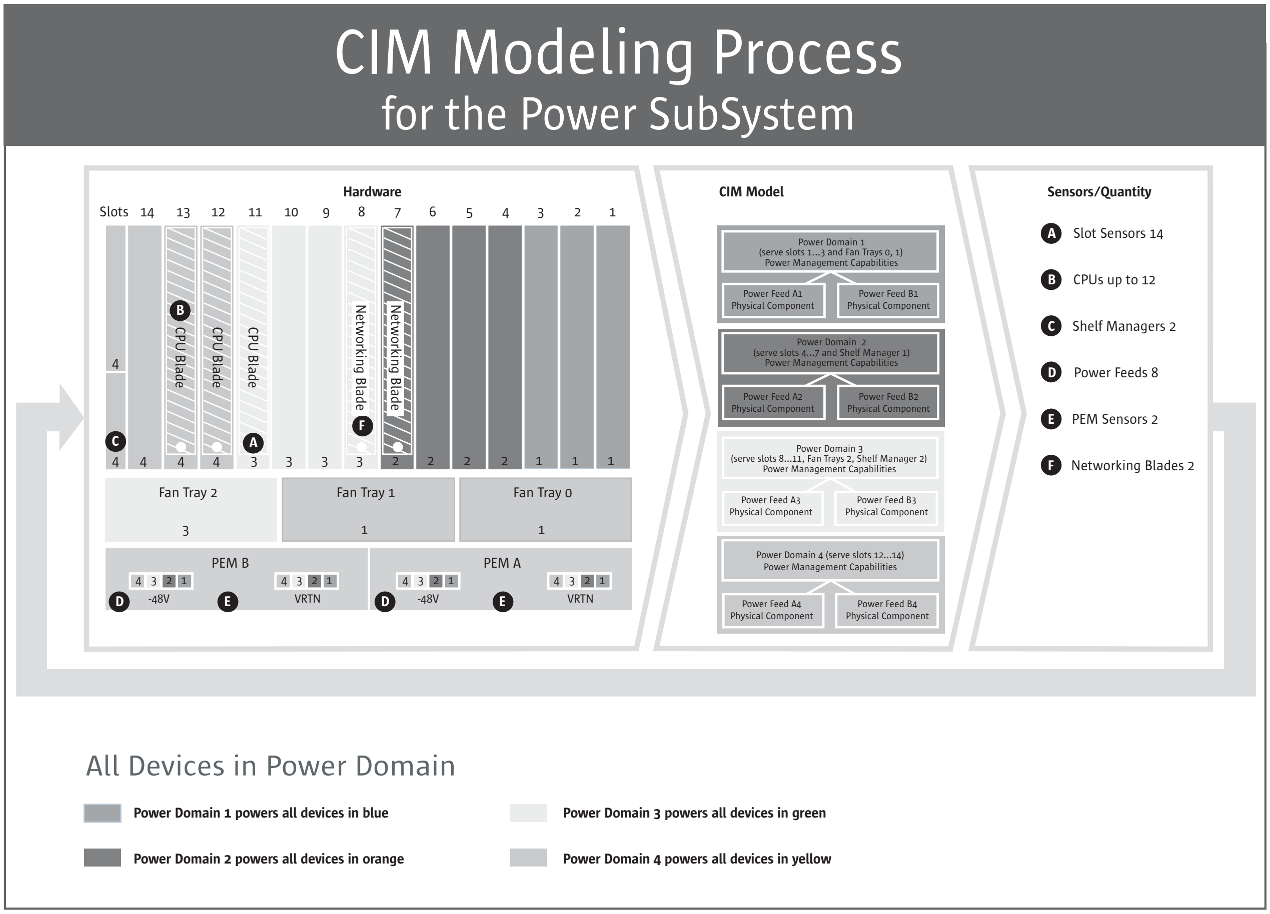

Because Sun implements its system management stack in modular units at the most appropriate layer, NEPs gain the flexibility to use individual pieces of Sun technology—such as the Netra High Availability Suite (NHAS) or the Solaris Operating System—or the entire range of Sun technology, as their needs demand. And although emerging standards do not currently address this requirement, Sun's management middleware ensures that vital information is provided to both high-availability middleware for service failover and to the service management interface for availability monitoring. Further, Sun's innovative instrumentation and middleware can eliminate up to 75 percent of the heuristic code NEPs would otherwise be required to develop—speeding time to market for new services and reducing implementation costs by up to 30 percent per application. This, in turn, frees valuable resources for use in developing a more competitive edge.

Sun also developed its Managed Object Hierarchy (MOH) as a Java-based framework NEPs can run on top of any OS with a well-defined set of carrier-grade APIs in either rackmounted or blade-based implementations. In the MOH framework each component is defined with Java Management Extensions (JMX) technology, with each layer reporting to the next, providing data critical to diagnosing and isolating faults and even offering detailed replacement instructions. Java's portability and flexibility eliminates OS lock-in, speeding and simplifying the integration of mixed operating systems and hardware. With standards-based APIs semantically and syntactically consistent across the infrastructure, Sun's advanced management stack enables NEPs to maximize the value of COTS components.

By overlaying SAF APIs when they emerge on top of the existing Sun management stack—including the MOH and Processor Management Services (PMS)—Sun can offer NEPs proven, integrated information models that bridge the standards gap. These valuable mapping layers will further speed NEPs time to market and reduce development costs by up to 40 percent. In sharp contrast, competing vendors are foregoing such mapping, leaving the time-consuming development to NEPs and integrators, or are scrambling to catch up by building a management stack from scratch.

Sun's integrated management stack will provide a system-level solution that will allow NEPs to manage groups of racks instead of individual servers. Sun can achieve this advanced system management through its proven API mapping and by requiring that components integrated into its ATCA offerings are SAF compliant at every layer instead of only at the lowest layer, as being adopted by many competitors.

The Solaris Operating System (OS)

The Solaris 10 OS offers users unprecedented flexibility, and is available for both AMD Opteron processor-based and SPARC-based platforms. The Solaris Operating System has roots in the telecommunications industry—having originated from Berkeley UNIX®, which emerged from Bell Labs—and is virtually an industry-standard OS for mission-critical telecom applications. Tools including Solaris Dynamic Tracing (DTrace) enable NEPs to maximize the number of users per system and improve response times for increased customer satisfaction, and the Solaris OS advantage of Predictive Self-Healing in real-time maximizes system and application service availability by automatically diagnosing, isolating, and recovering from system faults as they occur. Additionally, Solaris Containers allocates additional resources to applications as needed to meet peak demand, enabling NEPs to provide predictable service levels even when multiple applications are on the same system. Network downtime is minimized with Solaris OS security enhancements such as buffer overflow protection at the application and server level, an integrated IP filter, a new cryptographic framework, and process and user rights management. Solaris Flash Archives and other new management features enable fast installation of new OS and application software on servers over a network, reducing system administration costs.

Netra High Availability Suite (NHAS)

By reducing the need for developing custom software, NEPs can accelerate the deployment of highly available services at lower costs. NHAS Foundation Services software provides this ability. NHAS enables NEPs to reduce OS and hardware integration costs by using preintegrated, Netra-based platforms, focus valuable resources on more strategic activities, and reduce development time and cost by embedding high availability software into solutions based on the Sun computing platform.

NHAS is a modular suite of scalable, reliable software services for deploying Carrier-Grade Linux and the Solaris OS in a highly available environment, meeting telecoms' fast failover requirements. NHAS is flexible enough to create a dynamically scalable cluster of distributed nodes—or augment existing high-availability frameworks. NHAS delivers zero recovery time for network link failures and ensures continuous availability of core services such as fault management, reliable IP transport and NFS, reliable boot and DHCP services, and reliable cluster communication and management. Finally, NHAS offers support for the dynamic addition of new nodes and applications streamlines horizontal scaling, and automated software installation and configuration allow quick and easy deployment.

Chapter 7
Sun Services

Customers receive end-to-end support with a comprehensive range of Sun Services, including a global presence in over 170 countries and the Sun Customer Ready Systems program. One example of Sun Services for telecom is SunNet Connect.

SunNet Connect

Administrators are able to reduce management costs and increase system availability for SPARC-based Sun systems with SunNet Connect services, an integrated set of online system management tools. With SunNet Connect, administrators use a single common, portal-based interface to monitor key system metrics, receive automatic notification of system events, view detailed configuration and performance data, and request remote servicing from Sun professionals if desired. SunNet Connect services also provide clear, actionable service messages that identify repair part numbers and the appropriate repair procedure, enabling NEPs to detect and resolve problems faster.

Chapter 8

Conclusion

The telecommunications industry is one of the world's most dynamic fields in new technological growth and development, and the specifications of ATCA enable NEPs and service providers to meet its demands in a standardized, streamlined manner. The benefits of AdvancedTCA are driving the industry toward new economies and new innovations. Sun is integral to ATCA: Both for helping create these cutting-edge standards, and for developing the technology to match and exceed them. Sun products have guaranteed lifecycles, including advanced change notification and extended system availability to reduce the risk of adopting new solutions. The Sun ATCA solution is a proven one from an industry leader, offering strong price/performance, unified system management, interoperability and flexibility, investment protection, exceptional RAS, and more. The result is lower TCO and the ability to leverage economies of scale in the telecom industry now and in the future. Sun provides the easiest, most effective solution for adopting ATCA.

© 2006 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, CA 95054 USA

All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California.

Sun, Sun Microsystems, the Sun logo, AdvancedTCA, Netra, Solaris, StorEdge, Sun Cluster, Sun Java, and SunNet are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the United States and other countries.

UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a). DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS HELD TO BE LEGALLY INVALID.

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

©2006, Sun Microsystems, Inc. All rights reserved. Sun, Sun Microsystems, the Sun logo, AdvancedTCA, Netra, Solaris, StorEdge, Sun Cluster, Sun Java, and SunNet are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc., in the United States and other countries. AMD, the AMD Arrow logo, AMD Opteron, HyperTransport, and combinations thereof are trademarks of Advanced Micro Devices, Inc. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd. Linux is a registered trademark of Linus Torvalds.

SunWin #465271