Building IP networks using Advanced Telecom Computing Architecture

Dharmaraja Rajan, *Nortel,* Research Triangle Park, North Carolina, USA.

Abstract

The second generation of Advanced Telecom Computing Architecture (ATCA) based on PCI Industrial Computer Manufacturers Group (PICMG) specification has evolved to a live deployment phase. Products from world renowned telecom vendors in the Wireless (CDMA, GSM, and UMTS), IMS, Circuit core and Packet Core to build scalable carrier grade IP and converged networks are being deployed. These networks provide IP telecommunication, IP multimedia communication, scalable web collaboration, unified communication as well as triple and quadruple play services.

This paper analyzes the emergence of ATCA compliant hardware and software platform from a global leader in telecommunication. The paper discusses the ease of building applications using the high speed interconnect technology, redundancy capability, multi-switching fabric usage, high availability carrier grade middleware and simplified network management for next generation networks. The CAPEX and OPEX benefits to service providers of using ATCA are discussed.

1. Introduction

Proprietary platforms are proving to be highly expensive to develop, maintain, extend, and provide expected return on R&D in an era where open standards have proved to be achievable, cost effective and welcomed by customers. The need for increased interoperability across network elements, high availability even beyond 'six nines', centralized simplified network management in complex solutions like IP Multimedia Systems (IMS) and fixed-mobile convergence has increased the need for an open computing architecture in the next generation of networks. The Advanced Telecommunications Computing Architecture (ATCA) [1] provides the architecture to meet the next generation of standardized platform for carrier grade communications equipment.

2. ATCA Architecture

ATCA based on the PCI Industrial Computer Manufacturers Group (PICMG) 3.0 and 3.1 series[2], incorporates the latest trends in high speed interconnect technologies for data transport that defines the next generation architecture for building high end, "carrier grade" equipment. The specifications are oriented around switch fabric technology instead of a conventional parallel bus. The specifications allow open multi-vendor architecture based on open non-proprietary standards for board, backplane, and chassis vendors to independently develop products that meet the rigorous Network Equipment Building System (NEBS), European Telecommunication Standard Institute (ETSI) compliancy coupled with High Availability (HA), robust system management and high bandwidth needed for next generation communication equipment.

ATCA is more advanced than Compact Peripheral Component Interconnect (CPCI) in terms of offering carrier-grade availability, flexibility, expandability up to 16-slots, larger and more powerful boards (or blades) with advanced system management and offering standardization in a range of switching technology. ATCA architecture defines a platform with major components including the Front Board assembly, Rear transition module (RTM) assembly, Backplanes, Shelf and Shelf management, power distribution, Thermal and data transport framework.

The Front Board is 8U high form factor (3U=100mm) and includes mechanical faceplates, electronic functions and three connector zones. Zone 1 is for power connection and Shelf management, Zone 2 is for Data Transport Interface and Zone 3 for user defined Input/Output interconnects. RTMs are 8U high optional assembly units, providing user defined input and output connectivity to the companion Front Board from the rear using Zone 3 connections. The Front Board can support up to four Peripheral Mezzanine Cards (PMCs).

The Backplane is a passive circuit board providing the Zone 1 and Zone 2 connectors for Front Board Slots connected via high speed signal pairs. The Backplane supporting 2 to 16 slots, acts as power

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE

distributor, test bus and Shelf Management signal distributor through the Zone 1 connectors. Other Base Interface, Fabric Interface, Update Channel Interface, and Synchronization Clock Interface signals are distributed through the Zone 2 connectors. The Backplane supports a variety of system fabric topologies for connecting boards together. The need for highly available systems dictates redundant Backplane fabric schemes such as the Dual Star Topology. Where distinct fabric interfaces are used for unique purposes, a Dual-Dual Star is a prominent system of choice. In systems with a small number of nodes, where direct data path between nodes is scalable, a Full Mesh topology is selected. The fabric interface provides connectivity among the boards. Multiple fabrics can be supported in a single system. Dual Star Topology backplanes require two dedicated slots called Hub Slots (located in Logical Slots 1 and 2) for Hub Boards to be inserted. The remaining 14 slots (maximum), called Node Slots used by applications are routed to Hub Slots. Each Hub Slot has a Channel connection to each Node Slot in the backplane.

The Update Channel interface is a direct high speed interface between two similar function boards where large traffic volumes need to be exchanged in very short time to meet specific application requirements.

Subrack provides the various mechanical structures that connect the Front Board, Optional RTMs and Backplane. A Shelf (also known as Chassis) consists of subrack, Backplane, Front Boards, Cooling devices, RTMs, Cable and Power supply. A frame can support multiple ATCA shelves. Each shelf can support from 2 to 16 slots.

The shelf management system focuses on low level hardware management. It monitors, report anomalies, controls, takes corrective actions when needed, manage power, cooling, and interconnect resources to assure proper operation of boards and other shelf components. An Intelligent Platform Management Interface (IPMI) infrastructure provides communications, management, and control among the distributed controllers and to an overall system manager. A system manger may be integrated into the shelf or a common system manager can manage multiple shelves. The ATCA specification outlines architecture for high-speed management services but does not outline the In-band application management. It recommends redundant shelf managers but does not standardize on the mechanisms of fail-over process and selection of backup etc.

Dual 48V DC power from independent, distributed filtered power feeds to minimize fault dominance and increase reliability is specified for all active component of the shelf. The maximum front board power dissipation of 200W per single slot board and 5 to 25W per RTM excluding fans and other shelf elements is specified. Proper air cooling to limit front board, RTM power dissipation, slot cooling, air filters to meet NEBS compliancy is mandated. Fans are mandated to be designed to operate with a single fan failure and continue operating under the specified thermal environment.

The most important components in the ATCA architecture are the application blades that go into node slots to differentiate product offerings. The support of multi-vendor blades within the high capacity, high performance, high reliability, low cost ATCA framework enables true interoperability that will be the natural choice for telecom equipment manufacturers and solution integrators. Conforming to ATCA standards is Nortel's Versatile Service Engine (VSE) platform. VSE is a true carrier grade ATCA platform that is built to meet the high capacity and reliability needs of Next Gen Wireless access, media gateways, soft switches, and IP based applications.

3. Nortel® ATCA based Versatile Service Engine platform Architecture

Nortel Versatile Service Engine [3] is a fully configured carrier-class system and common platform of choice built for use across various network solutions for CDMA, GSM, and IMS etc. The platform hardware consist of three ATCA Shelves, each with 16-slot ATCA system, two optional thin servers that support OAM function and a central office alarm unit, all housed in a Nortel standard ETSI-2000 frame (600mm wide, 600mm deep, 2125mm high) meeting NEBS level 3, ANSI and ETSI standards. The VSE shelf is standard 12U high with 16 front and 16 rear hot-swappable compliant slots. Two slots are dedicated as Hub slots and 14 Node slots are available for applications. Hot-swappable Power Entry Modules (PEMs) with four independent feeds, provide power supply for the chassis fans, boards, Field Replaceable Unit (FRU) inventory storage device, power filtering and circuit protection. Each shelf has nine fan modules, allowing the shelf to operate indefinitely with a single fan failure. A simplified cooling system consumes less then 200W power per board. This provides greater tolerance to partial failures of platform functions. Each shelf consists of a set of fully redundant Common Boards and a set of Processing boards. The

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE

Common Boards provide secure Networking, Switching, Data management and Storage as detailed in section 3.1. The Processing boards provide solution specific carrier-grade HA applications. Eight slots in the bottom shelf are used by the common boards. Two slots in the upper shelves are required for a pair of Hub Boards. All other slots are available for application boards.

Figure-1 shows a typical VSE ATCA platform system.

Figure 1. Three shelves VSE Platform

VSE Interconnect Architecture supports Layer-1 physical interconnects with backplane and inter-shelf connectivity. Backplane base fabric connectivity for OAM and Data fabric connectivity for IP payload are physically separate switching planes that use dual-star configured redundant links. Layer-2 interconnects (Inter-shelf connectivity between Hub boards) is a 'full mesh' configuration that helps to maximize throughput, where every slot has a dedicated link to every other slot. This supports a potential aggregated terabit/second data transfer within a single shelf.

VSE Shelf Manager and Alarm Module (SAM) board communicates through IPMI Bus to all the active components in the shelf and has external interfaces to the switch boards and SAM on other shelves. The SAMs work with the switch boards to provide complete control of the shelf using Shelf Manager Controller (SMC).

3.1. Common boards

3.1.1. Hub Board - The redundant Hub board (also called Shelf controller in VSE) are placed in slots 1 and 16 in each VSE shelf, physically separated to prevent accidental removal or damage. The board supports dual star topology and provides data fabric switching for a wide variety of applications, as well as shelf-wide and/or frame-wide management. The hub boards are configured in a load shared arrangement. The data fabric in each shelf supports Gigabit Ethernet (GigE) links to all slots, with two physical slots supporting four GigE, six slots supporting two GigE and six slots supporting one GigE each. The board implements IPMI, and supports a full suite of standard external interface from Front Board and RTMs.

3.1.2. Network Gateway (NGW) Board – This board routes packet traffic between the subnets contained within the internal VSE network as well as to and from the WAN domain. For reliability and load balancing, VSE hosts redundant NGW boards in active-active configuration. Three interfaces are supported. A four GigE WAN interface considered as the VSE network demarcation point connects to customer's packet network. It supports packet filtering and acts as the secure routing interface. A two GigE 'BASE' network interface in a 1:1 active/active load sharing mode is used for network activation and maintenance. An Eight (four + four) GigE FABRIC network interface in active/active load sharing mode is used for data payload routing.

3.1.3. Data Manager – This board hosts applications that manage the various Operations, Administration, Maintenance and Performance (OAMP) functions. These include security, configuration, software management, reliable file system, Command Line Interfaces (CLI) and Machine-Machine Interfaces (MMI) for the VSE. It operates on an active/standby

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE

redundancy model. It uses an RTM to communicate to the Storage Board, located in adjacent slot.

3.1.4. Storage Board – The storage board contains a pair of standard SCSI disk drives that provide the high storage needs for various applications. It is located in the slot adjacent to data manger and connected through the RTM. The SCSI storage board does not connect to the base and data fabric interfaces. It contains no on-board processors or file system capability. Data redundancy is provided by configuring a second pair of Data Manager and Storage board. Data is mirrored on each Storage Board and replicated between Data Managers.

3.1.5. Application Boards - These boards provide a processing environment for a wide variety of carrier grade application software. In VSE these include Call Processing applications that can handle millions of Busy Hour Call Attempts (BHCA) and provide very high carrier grade availability. Call Processing boards using PMCs provide the specialized functions. IMS applications using standard Internet protocols (SIP, RTP) are implemented in a holistic architecture for conversational services (voice or video) over IP. IMS call handling applications like Call/Session Control Function (CSCF), which is split into three parts - the proxy-CSCF (P-CSCF), interrogating-CSCF (I-CSCF) and serving-CSCF (S-CSCF), are configured in separate redundant boards. Subscriber information management applications like Home Subscriber Service (HSS) for IMS, Packet Mobile Switching Center (controlling many remote radio networks), and Home Location Register (HLR) application for CDMA and GSM are all configured on VSE application boards. Application boards can be configured in a 1:1 Active/Hot-Standby mode and/or N+1 load-shared redundancy model. Redundant application boards need not be adjacent nor on the same shelf.

4. VSE Operating system and Middleware

VSE platform uses Carrier Grade Linux (CGL). CGL confirms to Open Source Development Labs (OSDL) [4] compliant Linux Kernel with carrier grade enhancements and Service Availability Forum (SAF) [5] compliant HA framework. The HA middleware supports many Nortel value added enhancements in checkpointing, memory synchronization, cluster management and failover, Ethernet link redundancy and failover, in-service software management and various OAMP functionalities. Figure-2 shows VSE architectural layers

Using SAF compliant Hardware Platform Interface (HPI) for HW management allows VSE platform to set and retrieve configuration or operational data about the hardware components, and to control the operation of those components. HPI provides hot-swap capabilities that allow a hardware component to be added to, or removed from the system with minimum downtime. HPI discovers the capabilities of the hardware platform, and maps those capabilities into the model, which the platform maintains and presents to the applications and other middleware. HPI reads sensor values and operates watchdog timers to monitor the health of the hardware components.

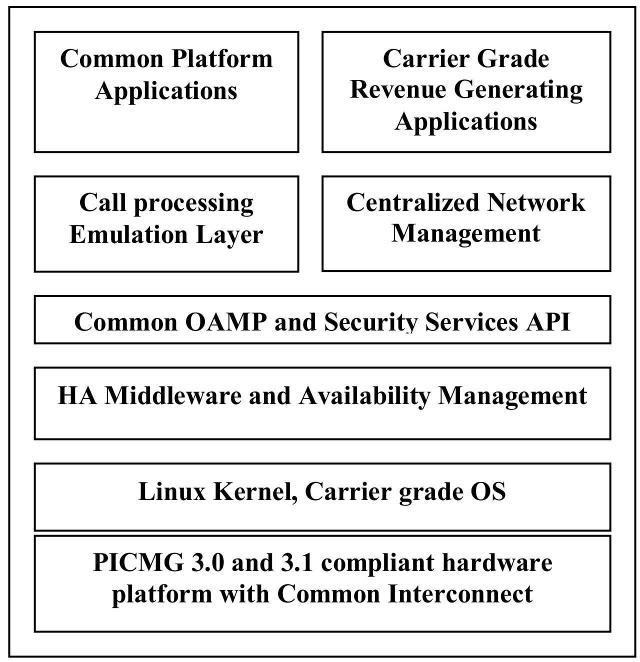

Figure 2. VSE Architecture

VSE uses SAF compliant Application Interface Specification (AIS) based middleware for availability modeling, registration and de-registration, health monitoring, availability management, cluster membership service, checkpoint service, event service, message service, resource locking service, protection group management and error reporting. The availability management framework coordinates redundant resources within a cluster and enables highly available applications. It determines the readiness state and the HA State of a component, and checks the health of a component by invoking callback functions of the component. Nortel's value added functionality to hardware platform and application software management is a

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE

key to meeting various carrier-grade needs for next generation telecommunication applications.

5. Network Management in VSE

Operations, Administration Maintenance and Provisioning (OAMP) services are provided within the VSE platform, to ensure consistent treatment of product applications. Standard APIs provide access to the middleware. The core services are fault management, configuration management, accounting, performance management, software management, maintenance and security. These applications reside on the Data Manager Board and use the disks on Storage Board for data persistence.

Fault Management services enable superior fault containment by providing fault isolation, correlation of faults at multiple layers in the hierarchy and high levels of coverage for hardware faults, interconnects and software faults. Higher availability is achieved by mitigating faults instead of taking more severe actions (e.g. board fail over to standby). The Fault Management system supports an active alarm list, local and centralized event history log with configurable retention period. Fault management system also supports XML and SYSLOG data formats for alarm and log information transport to north bound systems.

The Configuration management subsystem provides ability to do various configurations and provisioning activities using a secure Command Line Interface (CLI). Configuration of the shelf, boards, NTP, DHCP, network addressing, routing and other device configuration is done as part of installation and commissioning. Various services are provisioned as the deployment progresses and system is activated to provide the service. VSE resident applications can be configured, managed and monitored from a northbound Element Management System (EMS) Graphical User Interface (GUI) or Network Management System (NMS) or Operation Support System (OSS). The CLI supports both a human-machine-interface (HMI) and a machine-machine-interface (MMI) which includes additional attributes to ensure high reliability for bulk processing. Auto-Discovery of boards and dynamic provisioning capabilities simplify operation and reduce provisioning errors. Removal and insertions of FRUs are automatically detected and managed. The Configuration data model allows applications to express their internal configurable components in an XML representation using the NETCONF standard (RFC 4741) [6].

Accounting functions are performed by the redundant HA server system that sits on the same PTE-2000 frame above the ATCA shelves for CDMA and GSM solutions. For IMS, billing event messages are transmitted using the DIAMETER protocol to a third party system for further processing and transmitting to the downstream billing systems. Redundant backup of all billing events and retransmission capability (in event of failure) ensures that no billing information is lost. This provides the highest possible reliability.

Performance management services collect, process, report, store and deliver VSE Operational Measurements (OM) to North bound EMS and/or OSS systems. OMs for various boards and applications can be configured using VSE CLI. Accumulated measurements can be processed and XML reports transferred to remote destinations securely, using SFTP.

VSE System and application level security services support Authentication, Authorization and Accounting (AAA) as well as applicable compliance with GR-815 standards. Pluggable Authentication Module (PAM) based Centralized authentication using RADIUS protocol support secure shell, strong password management, single sign-on capability. Role based access controls limit the extent of operations that the user can perform. An audit log records user activity, configuration data change, raises logs and alarms when a user attempts to exceed privileges. VSE supports secure install and commissioning, secure remote emergency access, and controlled EMS and OSS communication using a secure protocol, to ensure end-to-end security.

A Unified Software Management (USM) service provides integrated software delivery for the platform, including software installation, patching, software upgrade, backup, and restore. USM coordinates and manages upgrades for all software and firmware on the system. USM supports software authentication checking using digital signature. USM recognizes and selects dependent software loads. This simplifies user action when upgrading and applying in-service patches. Users can view loads available and select a single application target (for load evaluation) or an entire system. Simple CLI commands permit users to begin, suspend, commit, or abort an upgrade. All software (installation loads and patches) are managed as software parcels.

Quick recovery from failure scenarios including dead office recovery and emergency recovery is supported with reliable call level trace and trace points that are triggered on instruction or event. A flight recorder function collects meaningful software debug data that survives board restart and non-recoverable failures. Data is captured at multiple

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE

levels from interrupt handlers to system loggers for off board analysis and to enable quick system recovery.

6. Results and Conclusion

ATCA drives the convergence of network elements in the Core, Switching, and Access portfolio into a board/blade based architecture that reduces real-estate and CAPEX. It supports interoperability, thereby reducing spares and OPEX cost, when compared to proprietary systems. Introducing any new platform in telecom computing architecture space requires care to ensure that the stringent reliability, high capacity and quality needs are met.

Nortel, a global leader and Tier-1 telecom vendor in telecommunications, has released the next generation of ATCA platform, the Versatile Service Engine, to the CDMA, GSM and IMS markets worldwide. Results of the deployments are very encouraging and show ATCA as a promising architecture for next-generation telecom and enterprise IP networks.

VSE ATCA architecture with its robust carrier grade solution leveraging Nortel's carrier-grade design practices and past COTS experience has helped us to advance the standards for next generation of telecommunication computing platforms. Nortel offers a total solution with a single point of support of the total solution. Thus various non-technical factors like ATCA vendor selection, developing partnership strategy, OEM agreements, supplier engagement and support model to meet various deadlines and avoiding deployment delays causing revenue risk can be avoided by selecting a tier-1 total integrated solution provider like Nortel. The future directions:

Emergence of technologies such as high memory CPU with chip multi-threading architecture, 10-40 Gigabit Ethernet fabric-switching, modularity enhancements to ATCA (such as AMC carrier blades with multiple AMC modules) and MicroTCA advancement will enable many more telecom and enterprise applications to adapt ATCA based platforms in the next two years.

Middleware vendor consolidation, alignment to SAF standards, streamlined deployment of various services and decoupling of OS from hardware, and allowing multiple OS to operate virtually, on the same blade, are emerging software trends that will help in plug-and-play type software components on a standardized platform.

Vendors, service providers and solution integrators adapting to the ATCA architecture will seek common ground to meet the bandwidth hungry, high quality, high reliability, next generation application needs at lower and lower cost, while providing richer functionalities expected by the consumer.

7. Reference

1. Advanced Telecommunications Computing Architecture http://www.picmg.org/v2internal/newinitiative.htm

2. PCI Industrial Computer Manufacturers Group 3.1 series http://www.picmg.org/v2internal/specifications.htm

3. Nortel Versatile Service Engine http://www.nortel.com/vse

4. Open Source Development Labs http://www.osdl.org

5. Service Availability Forum specifications http://www.saforum.org/specification

6. NETCONF Configuration Protocol http://www.rfc-editor.org/rfc/rfc4741.txt

Third International Conference on Networking and Services(ICNS'07)
0-7695-2858-9/07 $20.00 © 2007 IEEE