

Moving to IP-based Telco Applications

October 2001

intel®

Moving to IP-based Telco Applications

Abstract

IP-based telecom applications offer many exciting new opportunities as the convergence of circuit-switched, wireless and packet switched networks accelerates. With these opportunities comes the challenge of migrating telco applications from deterministic signaling architectures (such as PCI) and time division multiplexed (TDM) data buses like H.110 to packet-based architectures and, more pointedly, IP networks. This paper will examine a variety of system architectures being used to develop packet switched platforms to meet this challenge.

Audience

The primary audience for this document is technical decision-makers and engineers using Intel® IA/IXA/XScale hardware and application programming interfaces (APIs) to design solutions around packet switched products and platforms. Developers of telecommunications infrastructure equipment, with an eye to providing lower cost, standards-driven IP-based services in order to enhance or compete with traditional, higher-cost circuit-switched services, will also be interested.

Readers will benefit from a basic familiarity with telecom switching systems, computer telephony, TDM buses (such as H.100 or H.110), CompactPCI*, switched fabrics such as Ethernet or InfiniBand*, and system and application design.

Introduction

Moving to IP-based telecom applications offers many new possibilities. In an open IP-based network, it is possible to insert new services at a variety of points, from the core to the edges. The industry expects arrival of a myriad new applications that can be added seamlessly to standard voice or data services, based on the interoperability of products supporting IP protocols. As new providers experience greater flexibility to integrate services, there will be greater demand on bandwidth of packet switched networks and, hence, enhanced revenue for network and service providers alike.

Network bandwidth has become something akin to a growing natural resource. Fiber networks have been built out with a vengeance over the past several years and must be provisioned with infrastructure equipment to meet bandwidth demand, while generating profitable

revenue for both network and service providers. Consequently, infrastructure equipment will see an evolution over the next several years, as the technologies enabling network convergence advance to meet equipment demands. This trend will continue as application performance and the availability of very large-scale, IP-based systems at the network core become reality.

Current State

Since its development in 1994, CompactPCI has become the de facto bus architecture for telecommunications equipment manufacturers. CompactPCI is an enhancement to the PCI bus specification, created to provide a rugged interface for applications demanding features such as hot swap/hot add and high reliability. Today, more than 400 companies design and/or manufacture boards and systems based on this standard.

In current telecommunications platforms, it is typical for a single board computer (SBC) to carry out the task of controlling any number of non-intelligent line cards, attached as peripherals on a CompactPCI bus. The SBC provides deterministic control over the CompactPCI bus by driving interrupts, shared by a number of like line cards. The line cards typically support I/O interfaces to telephones or other terminal equipment. They may also handle tasks such as voice processing using digital signal processing (DSP) technologies, or protocol conversions to attached Public Switched Telephone Networks (PSTN). Communications among line cards, for the purpose of data movement from one line interface to another is handled over a TDM bus (H.110) providing near-zero latency.

Moving Toward IP Networks

As we move toward IP-based architectures we recognize some of the challenges ahead. First and foremost, an IP network is neither deterministic nor near-zero latency. This impacts some applications, but not all. Here are some of the other challenges.

A Litany of Protocols

The convergence of PSTN, wireless, and IP-based networks (either private or the Internet), offers service providers the promise of a convenient way to add a variety of new, revenue-generating services. Today, these interface with proprietary systems utilizing more than a dozen "standard" protocols. Many of these protocols must utilize a hand-off of control information from the applications on packet switched networks to the existing, circuit-switched network signaling systems.

Moving to IP-based Telco Applications

Therefore, long-term agreements must be reached regarding which protocols will survive in route to a fully converged network. Simultaneously, protocols for international standards are being defined to determine how IP services and new networks can communicate seamlessly, worldwide.

Centralized or Edge Services

Further, agreement must be reached as to whether applications should be more centralized, or moved out to the edges of the network. It is not unusual for newly developed applications to first appear at the edges of the network, even out in customer premise equipment (CPE). This is often the only way for a customer to have a desired feature set. However, once fully developed, these applications tend to quickly move back to service providers or to centralized services at, or near, the core of the carrier's network. The motivation to centralize these applications, once they become widely available, is that service providers know how to cost-effectively provide services to tens of thousands of customers while maintaining extremely high service levels. However, the platform that worked for the end user will most likely not meet the requirements of traditional service providers, and the applications will not survive in a service provider's environment, where five-nines availability (99.999% up-time) is expected. Therefore, new architectures are needed to migrate these applications from the edge to the center of the network.

A perfect example, today, of an application showing this strain is cellular voice mail. Sometimes it can take days for a new voice mail notification to reach a customer traveling through several outlying locations, i.e. hopping airport-to-airport. Is this a centralized application, an edge application or a distributed application? Where is the failure? Why does it take so long to rectify?

Distributed Application Environments

It is widely believed that neither core nor edge-located services will fill all needs, and that new services must be developed around architectures that can support distributed application environments throughout the network, coupled with highly available services platforms. Distributed applications are highly scalable, with no forklift upgrade requirements, and the addition of hardware must transparently provide new or more powerful resources.

Stable Platform-Management Standards

Highly available systems require integrated management from the platform to the application. All services

applications require stable and well-managed hardware platforms, achieved by providing health and performance metrics to top-layer management services. In turn, these metrics must be communicated through upper-management layers to the applications. Therefore, the platform management system must be built around standards, including standard APIs, which will allow it to interface with existing billing systems, management systems, management tools, and top-level services applications.

Fault Domains

Fault domain refers to a defined portion of a system that will be affected by a catastrophic failure, i.e. an event that cannot be managed by a system's redundancy capabilities. Further, the size of the domain is established by the risk mitigation of a failure to a defined number of users/subscribers. From the perspective of a service provider, services platforms must be capable of containing a fault to a small enough domain that, while a platform may degrade in its ability to provide services, it does not fail altogether.

Self-Repair and Scaling

Well-developed architectures, based on distributed applications running on next-generation platforms, must be capable of not only mitigating risks but, to a very high degree, they must be capable of self-repair. With large-scale telecommunications equipment, and from the perspective of a services application, there should be little perceived difference between a failed hardware module that halts service, an over-taxed hardware system that must be scaled to meet service level guarantees, and simply a new application that periodically halts. In any case, additional system resources, either on that platform or on another network-attached platform, must be able to pick up and maintain user services, or take on workload generated by additional users. Without this capability it would be impossible to manage systems and operate a profitable services business.

Packet-Based Architecture

A Different Approach

Getting Started

Today, we can begin rapid development of IP-based applications on equipment destined for the network edge—in central offices (CO), wireless base stations, business tenant service facilities (BTS), service providers (xSPs) of all types, and also customer premise equipment (CPE).

Moving to IP-based Telco Applications

We will look at a basic CompactPCI telecom system and the evolution to support IP-based services. We'll look at the integration made possible by packet switched backplane (PSB) architectures and the advances in standards-based building blocks that provide a logical migration path to the next generation network architecture. We will also discuss when certain applications and services will migrate.

First Generation CompactPCI Telecom Systems

In a first-generation CompactPCI telecom system (Figure 1) a single board computer called a System Master is the PCI bus controller. In this example, a PCI bus acts as the control plane for a set of telephone line cards and a PSTN interface card. Data transfers between line cards and the PSTN interface travel over an H.110 bus (CompactPCI version of the H.100 specification). The H.110 bus is the data plane. Traffic between local voice channels travels across the H.110 bus without touching the PSTN. External traffic also transverses the H.110 bus, but then moves out through the PSTN interface.




Figure 1: First Generation CompactPCI Telecom System

The CompactPCI bus is interrupt-driven and provides deterministic control of the attached CompactPCI peripheral line cards. The TDM, H.110 data bus gives time slices to data traffic across the bus. This low-latency bus provides excellent capabilities for real-time applications. An IP network interface (Ethernet) can act as a management access point or administration point for the telecom system. Such a system might be used for basic computer telephony applications such as a business PBX. Additionally, the IP interface might be used to attach new value-added services.

Where appropriate, additional services may be added to systems using IP connections. Voice mail, caller ID, call forwarding, and database services, linked to caller ID for customer resource management (CRM) systems, are a few examples. These services, when added externally, usually as stand-alone systems, are interconnected via an external network and switch. (Figure 2)

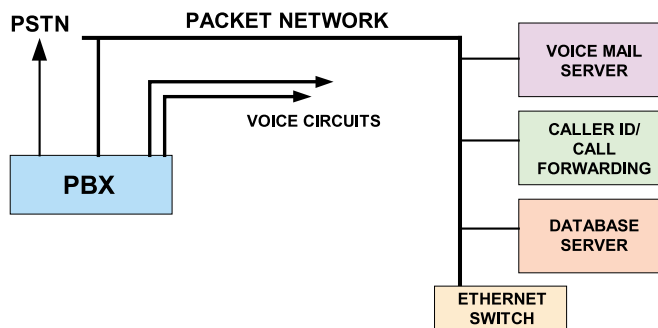


Figure 2: Value-Added Services over IP

Other new services such as Voice over IP (VoIP) require integration inside the telecom system, if acceptable performance is to be achieved. A gateway board is added as a CompactPCI peripheral attached to the CompactPCI bus and the H.110 bus. Now, with additional hardware (gateway board) and the additional software for VoIP, it is possible to take advantage of VoIP capabilities. (Figure 3)

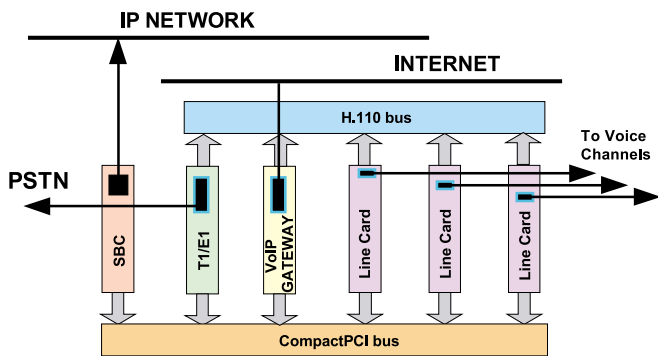


Figure 3: Integrating a VoIP Gateway

This CompactPCI system has begun to grow its usefulness, but a look at the overall solution (Figure 4) shows plenty of opportunity for further integration by bringing not only the extra servers inside the box, but the interfaces as well. The packet switched backplane (PSB) specification initiates this transition.

Moving to IP-based Telco Applications

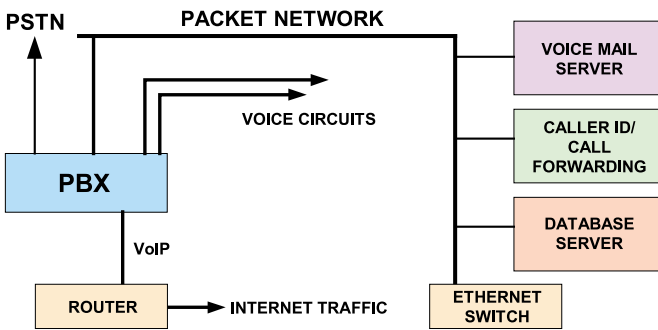


Figure 4: Expanding Capabilities Allows Many External Interfaces

Packet Switched Backplane Specification

In October 2000 the PCI Industrial Computer Manufacturers Group (PICMG*) set out to define a CompactPCI Packet Switched Backplane Specification, describing a fabric and a fabric switch slot on a CompactPCI backplane (Figure 5). While this has possibilities for the types of edge solutions described so far, the long-range ambition is to focus the capabilities of the various bus architectures and interfaces down to a single, packet-based architecture.

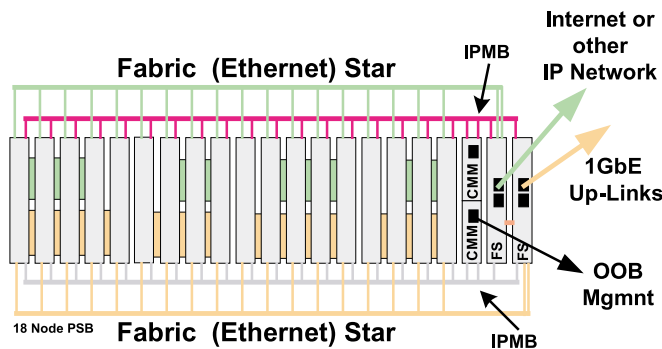


Figure 5: A Packet Switched Backplane Architecture

The configuration in Figure 6 shows what the solution, originally developed around several different boxes, might look like in an integrated packet switched backplane solution. Built around a CompactPCI PICMG 2.16-compliant platform, this solution takes advantage of the integration of Ethernet switches, with fabric on the backplane, to provide physical connectivity between any slot (node) and the Ethernet switch slots. This is accomplished while maintaining use of the CompactPCI control plane and the H.110 data plane for certain cards, namely the line cards, which could be voice processing cards, SDSL concentrators, or a VoIP gateway board. (For clarity, additional CompactPCI bus segments, as well as associated H.110 bus segments not in use, have been omitted in Figure 6.)




Figure 6: An Integrated Packet Switched Backplane Solution

Application servers, now integrated on the platform, utilize the PSB to provide connectivity, eliminating the need for external chassis and associated network cabling and switches. Bringing these services onto the robust CompactPCI platform further enhances reliability and availability of services, creating a flexible and manageable standards-based platform for a variety of IP-based solutions.

More Than High-Density, Cableless Interconnect

As important as they may be, there is much more to packet switched architectures than integrating high-density servers and eliminating cabling.

Ethernet Switches

Currently, fabric performance is limited by switch speed. While the PSB fabric, defined by the PICMG 2.16 specification, calls for supporting up-to-1 Gigabit full-duplex Ethernet speeds, current silicon for switches and associated physical interface chips demands more physical space and more power than is available for a CompactPCI form-factor board.

Presently, fabric switch cards will support 10/100 Megabit Ethernet to each node, with two 1 Gigabit Ethernet uplink ports. Yet-to-be-released silicon will produce fabric switch cards capable of supporting Gigabit Ethernet to each node across a non-blocking internal fabric. Naturally, adding faster uplinks will be welcomed next. However, it is expected that internal node-to-node communications over the PSB will demand more of telecom applications, such as voice recognition and voice portals, than the faster uplinks speeds.

How does this bode for current CompactPCI architectures that utilize a CompactPCI bus for control and an H.110 bus for data? Will they give way to fabrics like

Moving to IP-based Telco Applications

Ethernet or InfiniBand? Except for very small systems, 100 Megabit Ethernet or even Gigabit Ethernet will not meet the requirements of large, scalable systems that need the deterministic capabilities of a PCI bus or the near-zero latency of a TDM bus. This means that, for first-generation platforms with a packet switched backplane, CompactPCI and H.110 buses will often remain. SBCs, as servers capable of "drone mode," can disconnect from the PCI bus when installed in peripheral slots of a CompactPCI system, thus allowing multiple servers on the PSB.

Platform for Development

PSB Architecture Solutions

Figure 7 shows the backplane configuration of an Intel® NetStructure™ ZT 5090 4U General Purpose Packet Switched Platform. This 7-slot, CompactPCI development platform meets the packet switched backplane specifications of PICMG 2.16. It features a dedicated slot for an integrated Ethernet switch, two CompactPCI bus segments—each with an H.110 telephony bus—and an IPMI-based (Intelligent Platform Management Interface) chassis management module.

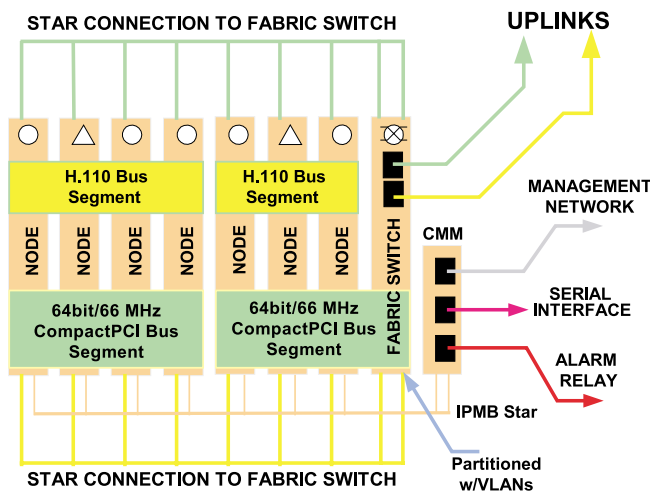


Figure 7: Sample Packet Switched Backplane System Configuration

Chassis Management Module


The chassis management module (CMM) is the first element of a comprehensive, standards-based, platform-to-application management infrastructure that will provide cost-effective, scalable, and high-availability (HA) services. The chassis management module (CMM) provides health and performance metrics from the chassis fans and power supplies, and from any node

board supporting the PICMG 2.9 CompactPCI System Management specification. The front panel provides a telco alarm relay interface, a Serial RS-232 interface and a 10/100 Megabit Ethernet interface to the CMM's on-board processor.



Intel NetStructure ZT 5090 4U General Purpose Packet Switched Platform

For this platform (ZT 5090), a star configuration of IPMB buses connects each node slot to the CMM slot on the backplane. This ensures that CMM functionality will not be affected by a failure on any given IPMB bus segment connected to an individual node slot. In addition, this prevents security breaches that could allow any given node to gain access to management of another node or chassis component.



Intel NetStructure ZT 7101 Chassis Management Module

Management Software

Figure 8 shows what the management layers might look like, applied to a PICMG 2.16-compliant CompactPCI platform. At the lower levels, system

Moving to IP-based Telco Applications

health and performance metrics are delivered via platform enhancements that communicate between the platform's field replaceable units (FRUs) and middleware. An HA management software layer provides communications between health and performance metrics and top-level management applications. Applications with this knowledge can then respond to changes in platform health or performance, thereby maintaining services to users.

Middleware and Upper-Layer Management Applications

Development of middleware, defined as supporting a standard, extensible set of APIs for upper layers of management, is a critical element of creating an HA services platform. Below are six additional management functions that could be enabled by a standards-based management middleware (**Figure 8 "Management Applications"**).

Deployment

Deployment is, in most cases, the first management application required. This allows deployment of services to application servers, and configuration of in-chassis boards to establish a services platform. It can also be used to scale the platform by deploying additional services on newly installed or idle boards.

Fault Detection and Recovery

When a fault alert is provided by the platform's health- and performance-monitoring mechanisms, additional layers of management software should acknowledge the fault and determine the best action to recover or restore services.

Performance Monitoring and Scaling

Performance monitoring will trigger deployment of additional resources when a system approaches a performance threshold, warning of impending system overload.

Facilities Management

Monitoring certain chassis environmental conditions, such as thermals and power loads, is required to alert facilities management teams to conditions that might threaten services, including services of adjacent equipment. Likewise, facilities environmental conditions, such as an air conditioning system malfunction, may instruct equipment to limit services, throttle processor speeds, or temporarily shutdown unneeded equipment until the conditions are rectified.

Network Management

Since networks are typically managed somewhat independently from services platforms, the local network operations center (NOC) will require removal of a

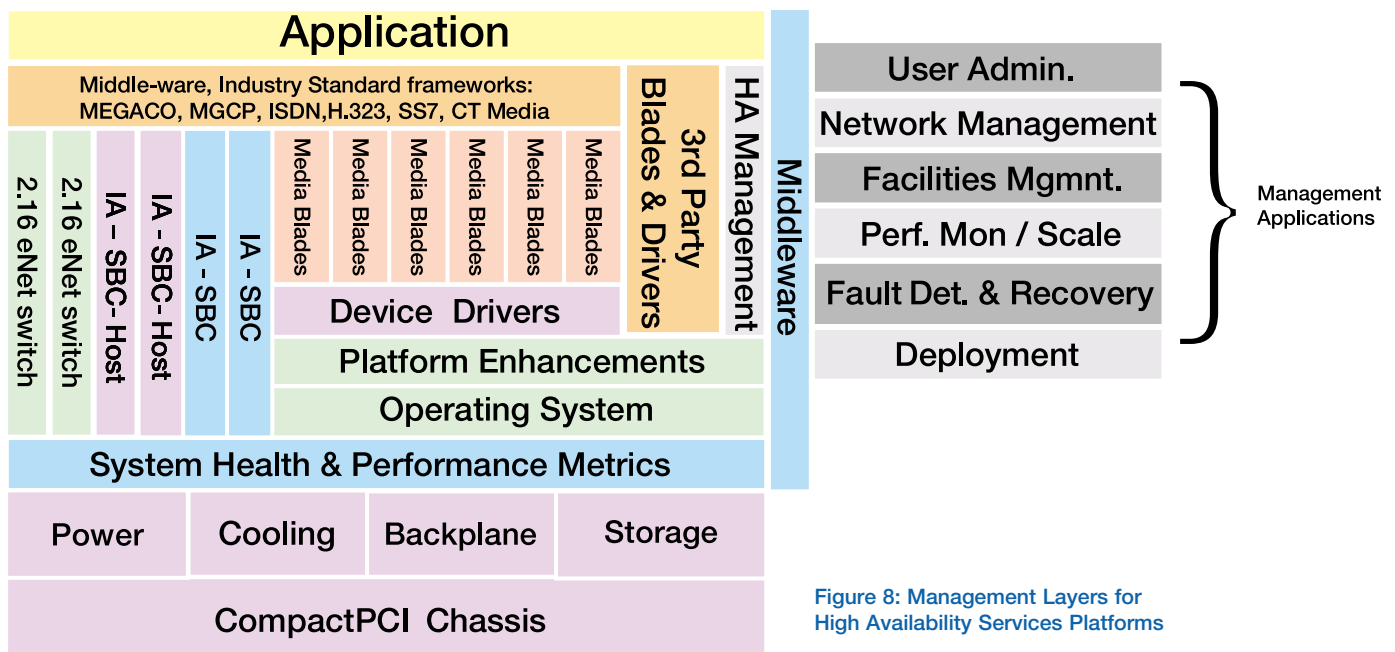


Figure 8: Management Layers for High Availability Services Platforms

Moving to IP-based Telco Applications

system (network disconnect) that is potentially disruptive to other network operations. Definitions of network parameters, i.e., virtual LANs (VLANs), are typically the responsibility of network administrators, as is monitoring of network performance. The ability to read network health and performance metrics from a service platform allows this level of management.

User Administration

In order to reduce costs, it is advantageous to limit in-house, service-provider management tasks that could well be handled by an end customer. Using a secure Web-based interface, customers could conveniently administer additions, edits and deletions to meet their own needs. These tasks, linked to automatic billing services, could include the addition/deletion of new services, subscribing to an existing service, or adding users.

Hardware Failover

CompactPCI Redundant System Slot (RSS)

Since we expect to see continued use of the CompactPCI bus with a System Master providing control to a number of telecom line cards, a solution for hardware failover of the System Master continues to be extremely important.

The purpose of a redundant system slot (RSS) is to provide seamless backup of service, should a System Master fail. If this should occur, a second (redundant) System Master board would affect a hardware take-over of the line cards. Not only is this fast—around 10 milliseconds for hardware take over—but with HA-aware application support, the application can continue, virtually seamlessly. This saves the high cost of additional line cards and associated external switching required to manage a hand-off, in the event of a failure.

Because of the adoption of the PICMG 2.9 specification, which drives health and performance metrics from every board in a system, much of the fault detection originally defined in another PICMG specification (2.13) is no longer required. RSS will be supported on all future PICMG 2.16- and 2.9-compliant CompactPCI System Master boards from Intel, with a simple firmware upgrade.

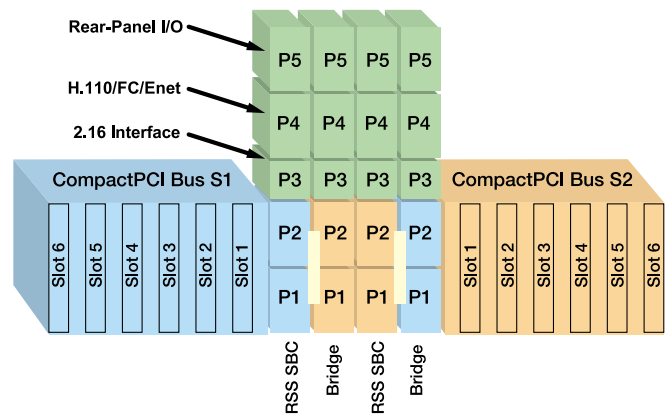


Figure 9: Dual CompactPCI Bus RSS System Configuration

Figure 9 is an example of a dual CompactPCI bus RSS System. RSS-capable PICMG 2.16 System Master boards (RSS SBC) provide RSS functionality. An additional mezzanine board (bridge) provides dual CompactPCI bus RSS capability, supporting up-to-12 line cards. Without the mezzanine board, the RSS SBC will support a single CompactPCI bus segment.

Packet Architecture Server Failover

Just like the RSS failover architecture, IP-based packet architectures need failover capabilities and scalability for distributed HA applications, and this typically requires clustering. Naturally, clustering that supports distributed applications is quite different from standby failover clustering architectures, or those that require a dedicated cluster manager node.

Standby/failover architecture (two-node cluster) lacks scalability while a dedicated cluster management node allows a single point of failure. Purpose-built clustering technology, on the other hand, is necessary in order for distributed computing architectures to achieve high availability for IP-based applications. In fact, next-generation packet switched system architectures handling packet processing, call routing and similar tasks, will require the same core distributed-application clustering technology.

Just as with new RSS architectures, clustered IP-based services platforms will benefit from the integration of standards-based management architectures. A platform's capability to manage fault detection, along with newly developed clustering technology and upper management layers, supports both failure-recovery and scaling for true carrier-class services.

Moving to IP-based Telco Applications

Next Generation PSB

When discussing the next-generation packet switched backplane architecture, it is important to understand how it fits with current CompactPCI system architectures and future high-speed fabrics.

Core CompactPCI specification

The CompactPCI core specification, PICMG 2.0, has provided excellent solutions to-date for robust telecom-grade platforms. This has been the first successful standards-based, telecom-like, Eurocard architecture to penetrate the telecommunications marketplace with embedded processing solutions that support a wide variety of applications. However this bus-based architecture, with a System Master and very loosely defined management, has limitations for large-scale systems based on future packet-based networks.

PICMG 2.16 and PICMG 2.9

Building on the core CompactPCI specification, the PICMG 2.16 specification adds a pair of switched star fabrics on the backplane that can each provide up-to-1 Gigabit full-duplex Ethernet speeds between a node board and an Ethernet switch board. Added to that is the capability for two (redundant, if required) Ethernet fabric switches. These two switches can provide either network port redundancy or independent public and private networks. Intelligent switch designs will support

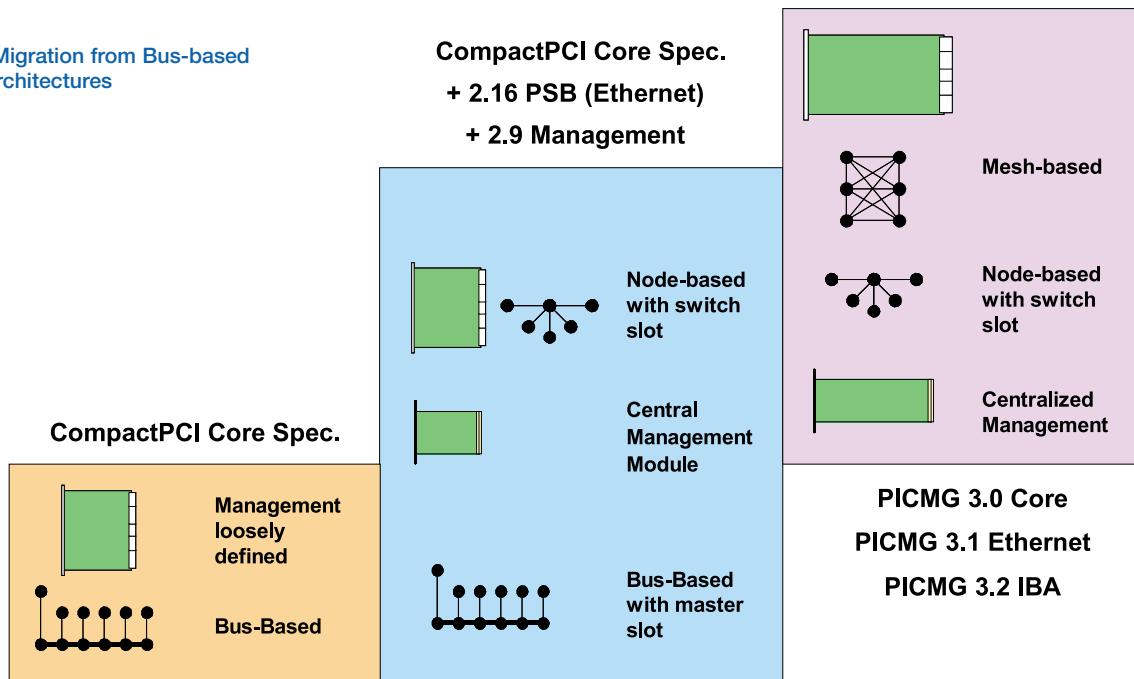
VLANs, intelligent packet routing, and Quality of Service (QoS) for traffic management and prioritization. High-speed uplinks of 1 Gigabit, with non-blocking wire speed switching, will enhance performance. As the silicon becomes available for higher speeds, we will see Gigabit node-link speeds, non-blocking wire speed switching and, finally, 10-Gigabit uplinks. Initially, the number of ports for Gigabit switches will be limited to 16 ports, however they will be non-blocking designs.

In addition, the PICMG 2.16 specification is the first to suggest support for the highly desirable IPMI management specification (PICMG 2.9). IPMI management, via IPMB bus connections specified for CompactPCI boards, provides the interface to transport messages, sending management information from chassis or boards to a device such as a CMM for use by upper-management applications. In turn, it may also pass messages from upper-layer management applications down through a CMM to effect a change to the platform or board, prescribed by the management application or the system administrator. Consequently, the new standards provide not only a foundation for packet-based architectures, but a standards-based management architecture, as well.

PICMG 3.0

With the CompactPCI bus and the H.110 bus both reaching their limitations, a mesh fabric architecture is favored to support a long-term ability to scale. This

Figure 10: Migration from Bus-based to Fabric Architectures



Moving to IP-based Telco Applications

next-generation packet switched architecture, built on the PICMG 3.0 core specification, will continue with node-based switched Ethernet slots and centralized management, and eventually transition in favor of a mesh-based fabric on the backplane over a CompactPCI bus (Figure 10). Designing initially on 1X InfiniBand (IBA), then moving to 4X IBA and then 12X IBA, provides a long-term architecture to build upon and the ability to migrate applications with minimal or no change.

Summary

The PICMG 2.16 specification introduces packet switched backplane capability to CompactPCI and introduces PICMG 2.9 as a workable set of standards for management interfaces. This new standards-based platform bridges current and PSB architectures, allowing immediate development of packet-based applications and management capabilities that will carry us into the future with PSB architectures for IP-based applications.

The PICMG 2.16 PSB architecture is the vehicle to begin developing next-generation, IP-based solutions. Integral to this will be development of next-generation RSS capabilities to support real-time critical applications. Software development will expand to new clustering technologies, enabling scalable server-based IP applications, once serviced by large centralized systems. Development of upper layers of management software will be necessary to provide links between platform health and performance metrics and the applications they support. And finally, development of HA-aware applications will complete the system for IP-based solutions.

As new, higher-performance fabrics such as 12X IBA become available in next-generation platforms based on PICMG 3.X standards, development of high-performance packet processing engines will produce highly scaleable systems, ready for everything from the core of the network to the edge. Developments in HA-aware

applications and standards-based management software architectures, will provision seamless migration from the first-generation Ethernet PSB architecture to high-performance, 12X IBA fabrics for standards-based IP platform solutions.

About Intel CompactPCI Building Block Solutions

Intel supplies telecom OEM equipment designers with standards-based, carrier-grade, building block solutions for Ethernet connectivity and processing in the network. Solutions include packet switched backplane technology and support for hot swap, high-availability and multicomputing capabilities, ranging from low-cost single board computers to fully integrated systems for mission-critical applications. Modular, scalable, off-the-shelf solutions and support for major operating systems and real-time software speed application development.

www.intel.com/network/csp/products

Information in this document is provided in connection with Intel® products. No license, express or implied, by estoppel or otherwise, to any intellectual property rights is granted by this document. Except as provided in Intel's Terms and Conditions of Sale for such products, Intel assumes no liability whatsoever, and Intel disclaims any express or implied warranty, relating to sale and/or use of Intel products including liability or warranties relating to fitness for a particular purpose, merchantability, or infringement of any patent, copyright or other intellectual property right. Intel products are not intended for use in medical, life saving or life sustaining applications. Intel may make changes to specifications and product descriptions at any time, without notice.

Intel®, Pentium® III and NetStructure™ are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.
* Other names and brands may be claimed as the property of others.

© Intel Corporation 2001. All rights reserved.