Multi-Service Computing Platforms for Telecom Applications

by Robert G. Pettigrew, P.Eng,
Motorola Computer Group, Ottawa, ON

Abstract

The next-generation telecom applications being built at the edge of the network - SS7 gateways, media/VoIP gateways, media gateway controllers and edge routers - rely on multi-service computing platforms with powerful processing and storage capabilities that support multiple protocols such as IP, ATM and Frame Relay packet data. The flexibility inherent in the architecture of multi-service computing platforms protects the industry-wide investment in circuit-switched technology and provides telecom OEMs and third-party vendors with an easier path to the newer all packet network. Where time-to-market was once the primary motivator for outsourcing of platform solutions, in today's more volatile market telecommunications environment OEM's are relying on outsourcing to stay focused on leveraging their resources and unique expertise. The name of the game is ensuring they are well positioned to address the challenges and opportunities of the converged telecommunications network when the market rebounds.

Sommaire

Les applications en télécommunication de la nouvelle génération, construites à l'extrémité du réseau (passerelles SS7, passerelles media/VoIP, les contrôleurs de passerelles media et les routeurs d'extrémité), dépendent d'unités de traitement possédant de grandes capacités de traitement et de stockage et qui permettent une compatibilité avec des protocoles de communication de type "paquet" tels que l'IP, l'ATM et le "frame relay". La flexibilité inhérente dans l'architecture des unités de traitement multi-service protège les investissements dans les technologies orientées "circuit" et fournit des OEMs dans le domaine et des vendeurs tiers possédant une meilleure flexibilité vers les nouveaux réseaux de type "paquet". Dans le passé, le temps de commercialisation était la première motivation pour sous-traiter les unités de traitement ; dans le monde actuel, plus instable, la sous-traitance est utilisée par les constructeurs d'OEMs afin de pouvoir conserver leur ressources ainsi que leur expérience unique. L'objectif est de s'assurer qu'ils restent bien positionnés pour aborder les défis et les opportunités de la convergence des données lorsque la situation économique repartira.

1.0 Introduction

The explosive growth in outsourcing platform design to standards-compliant systems in the telecommunications equipment infrastructure industry was originally driven by incredible time-to-market pressures faced by the industry. Rather than design everything from the ground up, telecommunications equipment vendors concentrated on their key differentiating technologies, and outsourced as much of their systems as possible.

Recent industry volatility has not alleviated pressures to outsource. Engineering design teams are now focussed on ensuring they are well positioned to take advantage of the inevitable resurgence of the market. Today, shrinking development budgets have replaced time-to-market as the primary motivator for outsourcing.

As the telecommunications market evolves, the increasing convergence of voice, data and multimedia applications onto telecom networks has involved deployment of various access devices to the edge of the network to provide new applications or services. The proliferation of these devices has lead to interoperability and configuration management issues at the network edge. These issues have escalated the need for a platform that can provide powerful processing and storage capabilities for control applications and a high throughput bearer architecture that supports multiple services on circuit-switched and packet networks. The ability to accommodate different protocols such as Internet Protocol (IP), asynchronous transfer mode (ATM), and Frame Relay without unnecessary conversion and overhead represents a significant step forward in network convergence.

2.0 Standardization of Platform Architectures

Technologies such as CompactPCI® have been adopted by telecommunications OEM's to provide standards-based platforms. Governed by PICMG, CompactPCI systems adapt PCI technology ubiquitous in desktop computers to something suitable for deployment in industrial and telecommunications applications. Specifically, CompactPCI offers a rugged form factor, with gas-tight pin-in-socket connectors, support for hot-swap of boards, rear access for cabling, and excellent shock and vibration characteristics.

CompactPCI systems suitable for deployment in telecommunications central office environments must also meet stringent survivability (earthquake and fire) and Five Nines availability requirements. Systems that meet these requirements have been available for several years, from companies such as Motorola. These systems adhere to PICMG standards such as PICMG 2.0, the core CompactPCI specification, PICMG 2.1, which adds hot-swap of I/O boards, PICMG 2.5 (H.110), which adds a data path for Time Domain Multiplexed (TDM) voice channels, and PICMG 2.9, which defines shelf management.

PICMG 2.16 for Compact PCI, introduced by Performance Technologies Inc. and ratified in September 2001, overlays an embedded Ethernet switching fabric on the backplane in CompactPCI systems. It complements standards such as PICMG 2.0 and 2.5, making it possible to implement systems that incorporate CompactPCI bus, H.110 and Ethernet fabric communication mechanisms.

PICMG 2.16 could provide an all-packet infrastructure with IP as the core packet transport. The architecture provides up to 1Gb/s per board of bandwidth for a system capacity approaching 20Gb/s in large systems. However, this architecture only supports IP transport, limiting its scope largely to Metropolitan Area Networks and Local Area Networks. Although there is movement towards Gigabit IP transport, generally the core of the network uses TDM and ATM for public switching.

3.0 Meeting the Needs of the New Network

A flexible platform capable of accommodating a variety of protocols and services integrated into a base platform represents an optimum solution for the ever-changing needs of the telecom industry. The infrastructure of the platform requires sufficient bandwidth to meet current and future needs at the network edge where the industry will experience significant growth in the foreseeable future.

To meet the high bandwidth and speed requirements of the edge and access portions of the network, telecom OEMs must efficiently process large amounts and different types of traffic. A key strategy to meeting these high bandwidth and speed requirements - while also leveraging the move toward open standards - is to add a high-speed mesh data fabric to augment the PICMG 2.16 Ethernet packet switching backplane standard. This fabric could be used for distributed processing and storage, as well as for bearer plane applications. The high-speed mesh data fabric fills a void for bandwidth, protocols and Quality of Service (QoS) that are not supported by IP, creating a multi-computing environment that can compete with higher-end symmetric multiprocessing systems.

Combining these features in a single data transport platform allows the entire spectrum of edge applications to be addressed. The resulting multi-service platform delivers optimum performance in data plane, control plane or integrated data and control applications.

Motorola Computer Group recently introduced a system that incorporates this full mesh fabric to address the evolving needs of the packet world. The highly flexible Multi-Service Packet Transport Platform (MXP) is built to adhere to open industry standards regulated by PICMG, provides fast throughput (more than 700Gb/s aggregate bandwidth), and the flexibility and scalability to deliver multiple services that can connect different networks - IP and ATM.

Motorola's packet transport platform addresses the critical areas OEM's face when considering the packet-based network:

- Leveraging open standards,
- Addressing multiple protocols,
- Meeting increasing speed and throughput requirements, and
- Connecting to traditional circuit networks.

3.1 Leveraging Open Standards

Standards-based systems provide access to compatible products from a broad range of vendors, maximizing OEM choice. Standards can only be generated from existing technology and in fast-evolving markets need to be augmented with forward-looking new ideas and technologies. The MXP design approach offers a desirable combination of standards-based design and forward-looking innovation.

The MXP platform has an open, extensible network architecture that adheres to all relevant PICMG CompactPCI standards - PICMG 2.0 Core Specification, PICMG 2.1 Hot Swap Specification and PICMG 2.16 Packet Backplane Specification for packet switching architectures. In addition, the MXP adheres to PICMG 2.9 Intelligent Packet Management Initiatiative for optimal system resource management.

IPMI is an industry standard used in Intel® based enterprise computers for management of peripherals by a PC. The IPMI protocol defines the data structure of the messages passed over the IPMI bus (IPMB) and Inter-Chassis Management Bus (ICMB), based on the I2C serial bus standard.

The IPMI Peripheral Management (PM) chip on each board in the system provides device discovery and control of power to the rest of the board. PICMG adopted IPMI for use in CompactPCI shelves or chassis as the PICMG 2.9 standard. Many telecom equipment suppliers have developed support for the PICMG 2.9 standard in both CompactPCI chassis and boards.

3.2 Addressing Multiple Protocols

Defined as a true multi-service platform, the MXP provides the flexibility and scalability needed to deliver multiple services that can connect to different networks. The platform architecture adds new features, such as a high-speed mesh data fabric and packet processor, which provides much higher data throughput and also supports multiple packet protocols. This feature removes the need for protocol conversion and opens the way for much higher data throughput rates. It also enables upgrade paths to users whose evolving needs cannot be met by architecture based entirely on existing standards.

3.3 Meeting Speed, Throughput and Application Requirements

Three key features of the MXP series architecture enable it to achieve the high speeds and throughput necessary to support applications both on the control and data plane:

- Non-shared, multi-gigabit fully meshed data fabric enabling a high-speed data environment,
- Fibre Channel fabric to support distributed storage applications,
- IPMI interconnection for combined system and resource management that increases the available payload by eliminating the need for PCI bridging and PCI system slots.

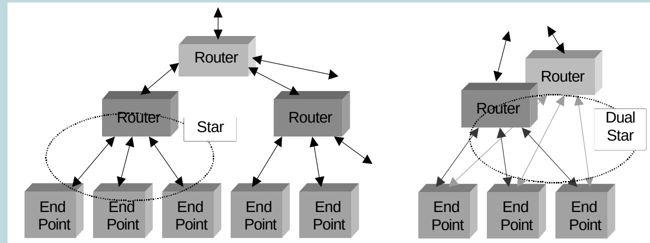

Figure 1: Star and Dual Star Topologies

Star topologies use a 'point to point' configuration where each device uses a dedicated link to send and receive data from a central resource. This resource provides the data distribution for the system. Ethernet networks are a hierarchy of star networks. Each sub-network is a leg on a star of the next layer in the hierarchy. Because reliance on a single central resource can result in a loss of all elements below a failure point, star topologies require redundancy for reliability. The PICMG 2.16 specification topology is a dual star configuration.

There are many topologies for wiring boards and systems to transport information. They differ in factors like cost (pins to connect, speed, logic costs, etc.), reliability and complexity. Motorola's MXP Multi-service Packet Transport Platform combines the advantages of star and full mesh technologies giving OEMs flexibility to use the topology most suited for various network applications and services.

In the MXP architecture, all the slots follow the core specification (PICMG 2.0) for power distribution and form factor and use the physical connection layer protection for hot swap (PICMG 2.1). The traditional CompactPCI bus has been eliminated, removing the extra complexity of dedicated system slots and PCI bridges. The slots are connected with three different networks: IPMI for shelf management; switched Ethernet; and the full mesh fabric.

The PICMG 2.9 standard defines the IPMI network, which is controlled by an alarm board resident in a dedicated alarm slot containing the IPMI Bus Management Controller. The PICMG 2.16 standard defines the dual-star switched Ethernet network. Each slot in the system includes two Ethernet interfaces, each capable of speeds up to 1Gb/s, routed to one of two dedicated switch slots. Ethernet traffic injected onto these

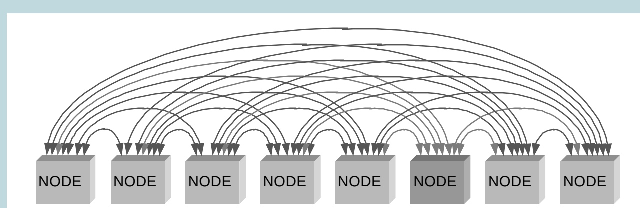

Figure 2: Mesh Topologies

Mesh Topologies are a superset of star topologies and also use 'point to point' connections. As networks add interconnects to eliminate 'dead branches' in a star topology, they reach a point where all nodes have connections to all other nodes. At this point, the hierarchy disappears. Each node can be an endpoint, a router or both. This figure illustrates how a node can act as a star point for all other nodes in the system.

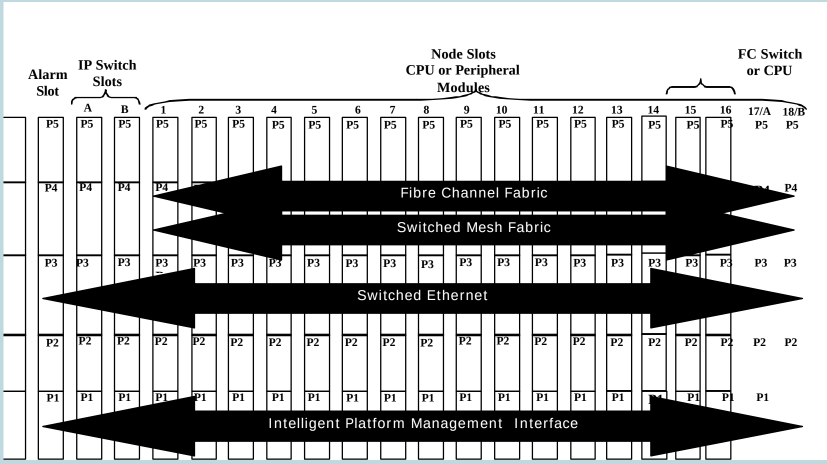

Figure 3: MXP MESH Architecture

The MXP multi-service architecture takes advantage of the mesh topology to offer a more resilient solution that needs no dependence on a central resource. The mesh architecture, which is being submitted to PICMG for consideration as an open industry standard, enables a more scalable solution than that presented merely by a star network. The backplane architecture is depicted here.

interfaces can be switched to any other slot, or out external interfaces on the switch boards.

The third infrastructure is a high-speed mesh data fabric that provides connectivity for every slot to every other slot in a 17 x 18 way interconnect. These connections, implemented with a four-wire interface capable of speeds up to 2Gb/s, are ideal for high-bandwidth bearer data applications. Because each interface is a point-to-point connection, they can be used independently of each another, making possible heterogeneous architectures supporting a variety of protocols.

Offering this mesh fabric technology to the industry as a candidate for an open standard should also ensure a full spectrum of boards will be available for a variety of other applications and interfaces. OEMs could also choose to implement their own interfaces for custom applications and product differentiation.

Motorola provides a distributed, flexible storage mechanism by implementing a dual-star Fibre Channel network topology, by dedicating two nodes in the mesh as hub slots for Fibre Channel aggregation modules.

Fibre Channel offers lower overhead, strict network determinism and higher throughput for networking and data distribution. In addition, the Fibre Channel network can connect CPUs with external HA Storage systems. A single HA Storage system can support multiple chassis. For systems with smaller storage needs, the MXP architecture supports plug-in disk modules. They are connected to the same networks as the CPUs and are served by the same Fibre Channel aggregation boards.

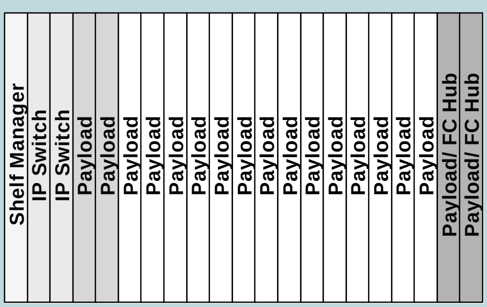

Figure 4: Flexible Chassis Configurations

The MXP topology can be applied to any chassis configuration. In a 21-slot configuration targeted for high availability applications, the system consists of an alarm slot, two fabric switches and 18 slots for task processors, network processors or storage elements. All 18 payload slots can be used for payload resources. If the Fibre Channel network is needed, two of the slots are used for the Fibre Channel aggregation hubs. The processors and disks interface to two predefined channels in the mesh.

4.0 Application Examples

The carrier-grade, high availability MXP series offers a flexible platform that can address next-generation applications such as SS7 gateways, media gateways, voice over IP (VoIP) gateways, media gateway controllers, edge routers, soft switches and access gateways. The MXP series can be scaled up or down for performance and throughput depending on the application and number of users. These options make the platform suitable for applications built at the edge of the network involving both control and bearer plane traffic. MXP support of high availability storage area networks (HA SAN) and distributed Fibre Channel storage also enables the MXP to support storage-and server-based applications including home location registers (HLR), billing servers and feature servers.

OEMs can configure the MXP with powerful combinations of distributed task processors, network processors, digital signal processors or storage elements to address many different applications and provide common development and deployment solutions.

5.0 Conclusion: Satisfying Solution to Manage Network Evolution

Highly flexible, multi-service architectures are strategic levers for reusing and driving down development and deployment costs in the telecommunications industry. Open standards and use of off-the-shelf standard products enable OEMs to more quickly develop applications that can leverage a multitude of third-party content and software. A single platform that achieves the highest levels of performance while also being flexible enough to be easily configured for a variety of applications anywhere in the network can help to ensure OEMs are well equipped to meet the needs of a changing marketplace. The packet transport platforms offer the necessary open architectures, speed and flexibility requirements to satisfy the needs of today's network while also providing the foundation to evolve into next-generation solutions.

About the author

Robert G. Pettigrew is Canadian Engineering Manager with Motorola Computer Group (MCG), and is responsible for ensuring the requirements of Canadian customers are reflected through innovative product design. To a large degree, Mr. Pettigrew's role involves managing partnerships with original equipment manufacturers (OEMs) such as Nortel, Newbridge, and Glenayre that reduce development time and result in breakthrough products for the telecommunications sector. Mr. Pettigrew manages teams of engineers across Canada who ensure OEMs in the telecommunications sector succeed in the intensely competitive global telecom market.