# Nortel Technical **Journal**

Produced by Nortel's R&D community

## FOCUS ON
## NEXT-GENERATION ETHERNET AND NETWORK INTELLIGENCE

This issue of the *Nortel Technical Journal* focuses on key technology disruptions taking place in today's network infrastructures. In particular, this issue highlights several Nortel-developed innovations – in Ethernet and optical technologies as well as new levels of network intelligence – that will enable operators and enterprises to enhance the value of their networks, reduce costs, create new sources of revenue, and deliver a superior quality of experience.

## Contents

Issue 4, August 2006

The ***Nortel Technical Journal*** aims to provide a vehicle for Nortel's global R&D community to share with one another, with customers, and with other external audiences the wealth of technology innovation under way across our network of laboratories.

**Published by:**
Office of the Chief Technology Officer

**Guest editor for this issue:**
David Hudson

**Executive editor:**
Peter Carbone

**Managing editor:**
Velma LeBlanc

The articles in this issue do not necessarily reflect a plan of record or commitment to product. The *Nortel Technical Journal* is intended simply as a vehicle through which we can exchange information, share views and perspectives, and expand meaningful dialog within our R&D community and with external audiences.

Back issues of the *Nortel Technical Journal* can be found on www.nortel.com/corporate/news/publications or by request at the email address below.

We welcome your feedback and look forward to your comments, suggestions for improvement, and ideas for future topics. Please send them by email to: norteltechnicaljournal@nortel.com

Printed in Canada on recycled and recyclable paper using vegetable-based inks.

# Next-generation Ethernet and network intelligence: Foundations for future networks

by David Hudson

This issue of the *Nortel Technical Journal* focuses on some of the key technology disruptions taking place in the underlying infrastructures of today's networks. Specifically, today's physical connectivity infrastructure – the "plumbing" – is being transformed to take advantage of optical and Ethernet technologies, as well as being equipped with new levels of network intelligence. These enhancements are critical for networks to support the coming wave of converged broadband traffic. Nortel is driving these transformations with innovative technologies and architectures that enable operators and enterprises to enhance the value of their networks, reduce costs, create new sources of revenue, and deliver a superior quality of experience to their customers.

Markets of all types – from consumer and enterprise, to healthcare, government, and education – are experiencing a rapid rush toward exciting new broadband multimedia IP-based applications and services.

Of note, the popularity of mobile videostreaming, real-time networked gaming, IP television (IPTV), and home networking is on the rise and new sources of broadband traffic are about to hit the mainstream – traffic that is largely video-driven, multimedia, mobile, and personalized. And, although video will dominate traffic volume growth, there will also be a massive growth in the number of network-attached devices, such as sensors, always-on 4G devices, and gaming devices.

Anticipating the coming bandwidth demand, operators for some time have been upgrading the capacity of their local loops, either by introducing the latest access technologies or implementing techniques to squeeze more value from their existing access "pipes," including copper, coax, and digital subscriber line (DSL). Operators worldwide, for example, will soon deploy WiMAX technology to bring improvements in speed, throughput, and capacity, as well as extended coverage for mobile subscribers. In the optical arena, passive optical network (PON) technologies are being used as a low-cost, simple way to deliver optical connectivity to homes.

> Innovations in the access portion of the network alone are not enough to support the coming wave of multimedia services and applications. Also needed is a more sophisticated network infrastructure, one that will increasingly lower costs, provide higher capacity and carrier-grade management and control of services, enable resource sharing in real time and on the fly, and help operators deliver new value.

While important, innovations in the access portion of the network alone are not enough to support the coming wave of new multimedia services and applications.

Also needed is a more sophisticated network infrastructure, one that will not only provide increasingly higher capacity at lower costs, carrier-grade management and control of services, and the sharing of resources in real time and on the fly, but one that will also help operators deliver new value.

Nortel's research teams are developing a suite of technologies and solutions to help operators transform their network infrastructures and ensure they continue to evolve along Moore's Law. These teams are also playing key roles in, and making important technical contributions to, all major telecommunications standards bodies working to define the network architectures of the future.

### Understanding the technology challenges

Over the years, as operators have deployed various technologies – including multiple physical layers and a host of new protocol extensions – to enable the interworking of services across different access technologies and consumer devices, they have as a consequence introduced massive levels of network complexity and cost.

To dramatically lower operations and management costs, reduce network complexity, and set the stage for future services, most operators today are moving to converge their multiple defined-purpose networks onto single next-generation networks based on Internet Protocol (IP).

As they do so, however, operators must approach convergence with the realization that such networks – where all traffic and service types are carried on the same packet-based infrastructure – cannot simply be a re-creation of the best-effort "commons" environment of traditional IP networks.

Traditional IP networks work well, as long as most of the time most of the users don't need most of the resources. Indeed, IP was not originally designed to achieve more than "best-effort" performance; it was certainly never architected to meet the quality-of-service requirements of such applications as real-time voice and high-bandwidth video; and for the most part, it did not include extensive monitoring capabilities – the "gas gauges and speedometers" that are indispensable to operators in determining whether a service has been delivered successfully. Nor was IP designed to ensure that bandwidth is shared equitably among these services; as with any ungoverned commons, a user could theoretically consume more than his or her fair share.

Until now, the flexibility and extensibility of IP has been used to mitigate these issues, but at a cost. While various technologies, such as Multiprotocol Label Switching (MPLS), were patched on to cope with growing utilization and to provide traffic management, these solutions are very processing-intensive. Although MPLS certainly does have a place in the wide area network (WAN), this technology can be far too costly in other areas of the network, notably the metro.

Alternative solutions are needed for tomorrow's converged network infrastructures.

Without doubt, as multimedia IP-based applications and services take hold, users will increasingly demand a service experience far superior to the one they currently get on today's legacy and IP networks.

For example, a mobile business user conducting an important voice call with a key client who happens to be simultaneously downloading a data file will not accept his conversation suddenly becoming clipped and unintelligible. An emergency call, such as a 911 call, must absolutely be delivered intact and with the highest priority. And soccer fans watching the World Cup live on their mobile devices will most certainly complain if they miss the winning goal.

Ensuring that converged next-generation networks are able to support these multiple services with reliability and quality of service poses several technical challenges:

• What is the best way to overcome the best-effort networking limitations of the traditional IP network, and to be able to equitably share network resources – such as bandwidth – among revenue-generating services, particularly with next-generation services that consume unpredictable and changing amounts of bandwidth?

• How do we meet the various quality-of-service requirements of different traffic types, such as real-time voice and high-bandwidth video?

• What is the best technology solution for driving carrier-grade reliability, robustness, security, quality, and determinism – the long-established hallmarks of dedicated circuit-switched networks – into the new converged infrastructures?

• How do we make the network easily and cost-effectively scale across the carrier's entire subscriber base?

• What technology innovations are needed to support the rising capacity demands of bandwidth-hungry consumer applications, particularly video?

• How do we support the increasing user requirements for mobility and personalization?

> Operators must approach convergence with the realization that future networks cannot simply be a re-creation of the best-effort "commons" environment of traditional IP networks.

### Network infrastructure for the future

Nortel's view of the network infrastructure required to address these challenges is shown in the figure on page 3. Briefly, this "roadmap" recognizes the following:

**Access:** There will continue to be a variety of access technologies for years to come. Over time, though, access will shift toward optical for buildings and wireless for everything else. This shift will accelerate as optical assumes the plug-and-play attributes of Ethernet and as wireless achieves broadband throughput (as discussed in Issue 2 of the *Nortel Technical Journal*, which focused on broadband wireless access).

**MANs/WANs:** As broadband access evolves, more bandwidth will be required in both metropolitan area networks (MANs) and wide area networks (WANs). In both, packet-optical connectivity will play a key role because it provides huge capacity at a lower and lower cost per bit delivered. Moreover, as the technologies for the metro packet-optical network become increasingly similar to those used in the wide area, common packet-optical platforms can provide the ability to service both network areas and provide greater flexibility in network configuration and capacity.

**Services edge:** The services edge – the new value point in the network – is the focal point for implementing network policy. In practice, this means matching the service that is delivered to the agreement between the service provider and end user. That agreement includes quality of service, security, community of interest (virtual private network), mobility, and other attributes. Media gateways, mobility gateways, and other subscriber- and services-aware devices are instances of the services edge capability.

**Figure. Functional architecture for next-generation networks**

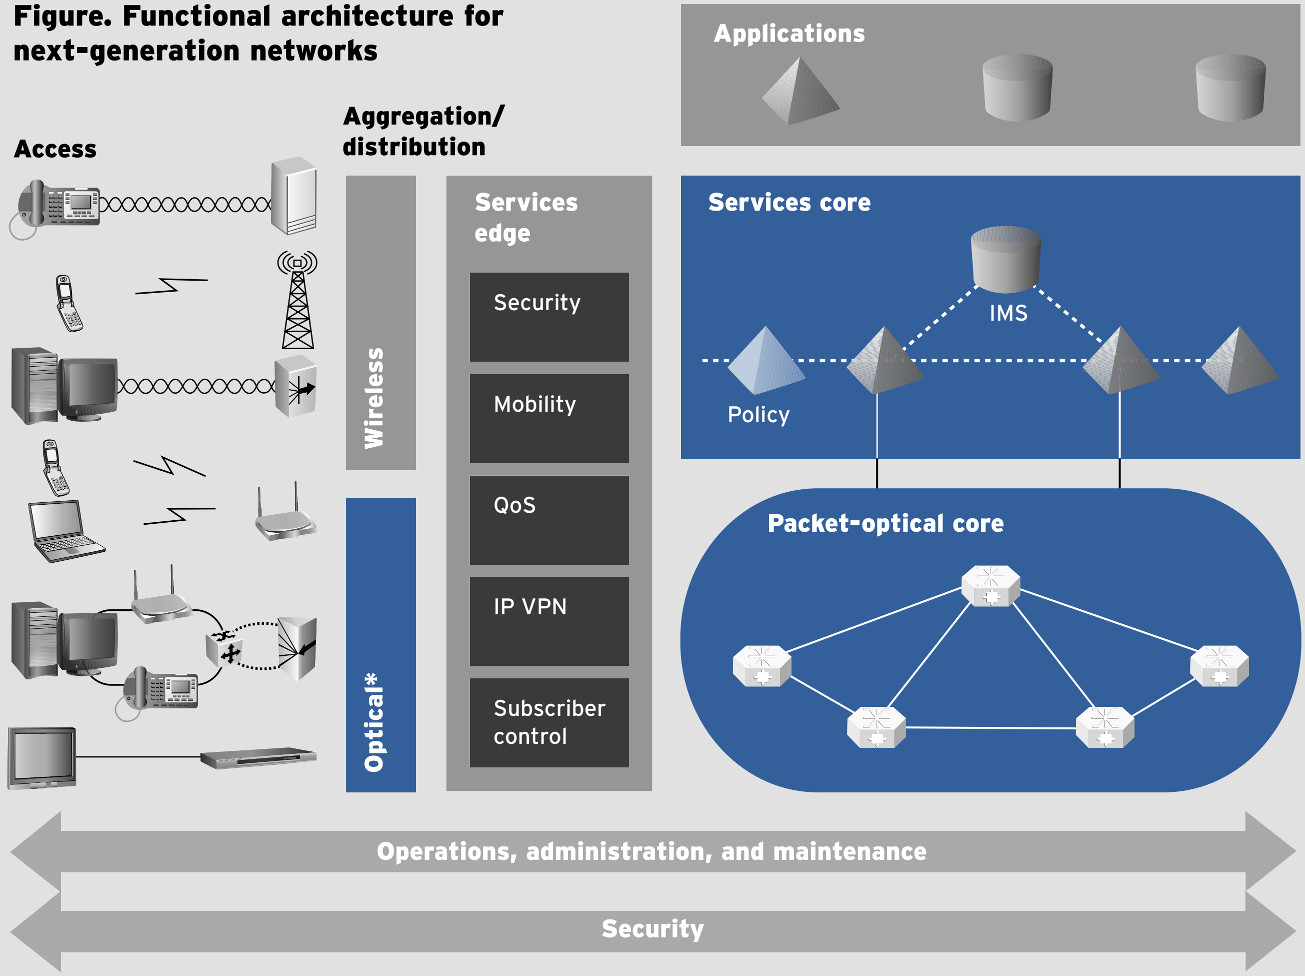


Nortel is working to transform today's infrastructures into the vastly simplified, high-capacity, and intelligent foundations that will be needed to support the coming wave of high-bandwidth converged multimedia services and applications. Shown above is Nortel's functional architecture for the next-generation network. The areas in blue indicate the key infrastructure areas that are highlighted in this issue. Specifically, in the metropolitan and wide area portions of the network infrastructure, Nortel is leveraging innovations in optical technologies as well as enhancements to Ethernet to significantly increase network capacity, simplify operations, and reduce costs. At the same time, Nortel is driving a new services delivery environment based on the IP Multimedia Subsystem (IMS) architecture, and is putting in place the key policy control capabilities needed for operators to support a richer set of applications, create new sources of revenue, and deliver a superior quality of experience for end users.

* Metro Ethernet is a technology option in this space

**Services core:** In the services core, the IP Multimedia Subsystem (IMS) architecture is rapidly being embraced by all industry players – standards bodies, customers, and vendors alike – and is emerging as the dominant architecture for delivering real-time services. IMS was conceived in the GSM/UMTS network environment as a way to deploy new broadband multimedia services and take advantage of the high speed and increased capacity that these networks were beginning to offer. IMS has been swiftly adopted by all standards bodies – across the wireless, wireline, and cable domains – and is universally seen as a powerful framework for tapping into the full capabilities of converged networks to create a rich set of applications and services that deliver new end-user value.

IMS, without a doubt, brings to the network new levels of intelligence, providing a single, standardized services control infrastructure that is separate from the underlying network and independent of access technology and device type. As a result, IMS enables applications to operate seamlessly, simply, and independently across the different access domains and service types. It also enables common user identification and billing services, and allows rapid service innovation with either operator-developed applications or via third-party applications.

Nortel is working across all of these areas. This particular issue of the *Nortel Technical Journal* focuses on innovations that take advantage of Ethernet and optical technologies as well as developments that drive new levels of intelligence into the network.

### Innovations in Ethernet

With increasing focus on the latest multimedia services and applications, it might be tempting to think that we're done with the lower layers of the networking infrastructure and that we are simply in a period of optimization and fine-tuning.

That is certainly not the case. The underlying connectivity infrastructure – the "plumbing" – is still a hotbed of innovation for our research teams as they work to address the key challenges of network convergence.

Specifically, as operators converge their many networks onto a single infrastructure for simplification, interworking, and cost reduction, they also need to preserve the ability to differentiate among the different services and to operate reliably. The challenge then for the connectivity, or transport, infrastructure is to deliver traditional reliability, still higher capacity, and even lower cost per bit in both equipment and operating costs.

Several technology choices exist today, including packet over SONET, ATM, pure SONET, dark fiber, or IP over MPLS, to name a few. But these address the above requirements only partially. IP over MPLS, for example, is very software- and configuration-intensive, and requires expensive memory and processor technologies in each device.

One technology that meets all the stated needs is Ethernet, which has emerged as the dominant connectivity technology in the enterprise and offers a low-cost and ease-of-use value proposition.

> **With increasing focus on multimedia services and applications, it might be tempting to think that we're done with the lower layers of the networking infrastructure and that we are simply in a period of optimization and fine-tuning. That is certainly not the case: "plumbing" is still a hotbed for innovation.**

### The metro Ethernet opportunity

The innovations Nortel is developing, in Ethernet and optical, as well as in driving new levels of network intelligence for services delivery, are essential building blocks needed for powerful network solutions across a wide variety of market spaces.

Most recently, Nortel launched a new business, called Metro Ethernet Networks, to focus on designing the solutions that enterprises and service providers will need to effectively support and manage the coming tidal wave of bandwidth demand in the metro portions of the network. This demand is being driven particularly by such applications as IPTV, video-on-demand, music videos, and video-enabled personal communications, which will tend to be cached locally in metro areas. Some analysts, in fact, predict that Internet-based video applications like these will drive bandwidth demand in excess of 60 percent per year for the next five years.

This new Metro Ethernet Networks business brings together Nortel's expertise in optical networking, its strong knowledge base and footprint in carrier data, and its experience in broadband access.

---

Ethernet, in fact, holds enormous potential as the networking foundation – the *lingua franca* – for tomorrow's carrier infrastructures, particularly in the metro area of the network (see sidebar on this page). As well, the packet versatility and sophistication of Ethernet marry well with the strong carrier-grade attributes of traditional optical transport, which delivers enormous bandwidth over huge distances with zero downtime and full management visibility. Ethernet, for its part, allows for fine granularity and efficient utilization of that bandwidth, inherently supports bandwidth on demand, and is universally plug and play.

The current vintage of Ethernet, however, must be enhanced as it steps out from the enterprise and into the carrier mainstream.

The second article in this issue (page 7) discusses Nortel's initiatives in developing enhanced, carrier-grade Ethernet by leveraging its leadership in "campus" Ethernet solutions and working to integrate Ethernet switching with more and more multicast, robustness, Layer 3 routing, security, and other upper-layer functions.

This article also examines the evolution of the optical fiber transport layer to increasingly greater packet capability and highlights how deeply network processors, digital signal processors, and packet fabric technology have reached into our optical platforms, including the Optical Multiservice Edge 6500.

Our development teams are also working on Ethernet-based solutions that can be applied for wireless backhaul applications. As 3G and then 4G mobile networks are deployed to drive more and more high-bandwidth data – particularly video to and from mobile devices – higher capacity backhaul that also delivers carrier-grade reliability becomes a priority (page 21). Ethernet over fiber is an attractive alternative to the traditional leased-line approach because of the technology price points, but only if stringent timing and synchronization requirements can also be met.

As Ethernet is readied for its role in the carrier environment, it must also provide end-to-end operations, administration, and maintenance (OAM) and protection to meet operators' requirements for visibility, control, and delivery of service level agreements (SLAs). In this area, which involves the evolving control plane, our teams are developing a new management paradigm that matches the high capacity and flexibility that Ethernet and optical provide.

The control plane is the subject of intense debate in the industry today. While IP is clearly the Layer 3 technology of choice, the solution for traffic engineering, customer separation, and operator control at scale in the metro network is being contested. Many early deployments literally used enterprise Ethernet switches "out of the box" and ran into issues when designed-for-the-campus solutions needed to meet metro requirements, such as the need for scale in the virtual local area network (VLAN) and the need for robust OAM functions and quality of service. And, while MPLS is the WAN technology of choice, it is certainly not the only option for metro networks.

As the article on page 25 explains, the best management approach for the next-generation network is not necessarily one that is based on what can be considered artifacts of technologies used in the past. Given an Ethernet transport solution in the metro, there are control plane options that are far more Ethernet-friendly and hence more opex-efficient. These options include MAC-in-MAC or Provider Backbone Bridging (PBB), and Provider Backbone Transport (PBT). In fact, Nortel recently announced its innovative PBT technology, which brings the deterministic, carrier-grade connection management characteristics of SONET/SDH into Ethernet.

## Innovations in optical technologies

Along with meeting the challenges of enhancing Ethernet, there is also still the basic challenge of meeting the rising bandwidth demand that new services and applications, particularly video, will need while driving down infrastructure costs. Our advanced technology teams are continuing their track record of pushing the limits of physics to deliver the largest number of bits over the greatest distance, using the least spectrum and power at the lowest cost.

Several new Nortel innovations enable operators to get more life out of the glass they already have in the ground. Nortel's electronic Dynamically Compensating Optics (eDCO) technology (page 10) is just one example. eDCO, which functions much like a traditional modem, is able to sense the quality of a fiber link, dynamically adjust the modulation of individual wavelengths to overcome the physical limitations of the fiber, and enable fiber spans of some 3,000 kilometers without dispersion compensation – a capability that dramatically simplifies network engineering, installation, and management.

> Ethernet holds enormous potential as the networking foundation for tomorrow's carrier infrastructures, particularly in the metro area of the network.

Another example is the Nortel-developed innovation to encrypt information at the lambda level. Nortel was the first in the world to demonstrate, in September 2005, integrated data packet encryption for high-speed 10-Gbit/s optical networks. (For more information, see Issue 3 of the *Nortel Technical Journal*, page 67.)

## New levels of network intelligence

Along with a transformed connectivity infrastructure, future networks will need added levels of network intelligence to properly control and manage the new multimedia services. From a technical perspective, this requirement is a difficult one to address, for several reasons.

First, converged multimedia services have highly variable bandwidth and quality-of-service requirements. The most innovative new applications will seamlessly blend voice calls with video and with interactivity, such as in a networked game.

Second, the growth of high-bandwidth services, particularly video, is causing bandwidth to flow through the network in much different ways than in the past. Networks can no longer be dimensioned on the assumption that users consume far more bandwidth than they deliver to the network or that they are on the network only infrequently. Peer-to-peer, place-shifting, and other capabilities allow any device to be an always-on content server. We will see enormous innovation and business opportunity around such capabilities.

Third, mobile and personalized services need a network that can match these services to the identity of each user, to their personal preferences, and to the devices they are using.

The network, therefore, needs to be highly adaptable to meet the range of service needs that are being converged onto a common platform and to deliver to end users a rich quality of experience, in a way that allows operators to successfully bill for that service.

Policy control addresses this opportunity. Essentially, the function of policy control is to govern the admission of users and applications onto the network and then monitor network resources once a call or session is admitted. It is the key function that communicates between the underlying network and the services and applications that ride on top. A policy decision may consider attributes such as quality of service, security, mobility, and even the type of device, among others.

The historic purpose of policy control was to help simplify the operator's job of provisioning the network – adding a subscriber or card, deleting a route, or pre-determining resources, for instance.

Some sophisticated policy control functionality has been developed for specialized applications that need, say, terabytes of bandwidth on demand. A movie studio, for example, would need such bandwidth to transfer large video files between production houses. Nortel's Dynamic Resource Allocation Controller (DRAC) is a form of policy control designed for these purposes. DRAC essentially allows such applications to interact through an interface to the optical network and provision the resources each application needs. (For more on DRAC, see Issue 1 of the *Nortel Technical Journal*, page 23.)

Similar policy control functionality is needed for next-generation networks, but at a much finer scale, in real time, and agnostic to devices and access technologies. Future networks must be able to handle the varying communication modes that individuals will use, often in a single session, and it must do this in real time. For instance, a user may start off a session with an instant message, and then switch to voice, then to video. That user may then hand off the session from a handheld device to a large screen, move from simple conversation to a secured transaction, or even shift from personal time to company time and the business VPN. All of these changes need to be understood and reflected in the policy that the network implements on behalf of that user.

All of these responsibilities are the domain of a new class of policy control being defined in various standards bodies, with the European Telecommunications Standards Institute (ETSI) being furthest along. Here, Nortel is making significant technical contributions and helping ensure that a common, standardized function is applied across all network domains. The article on page 33 discusses this policy control architecture in detail.

The policy control function being defined is based on the original work done by the Third Generation Partnership Project (3GPP) to define its new IMS services delivery architecture, of which policy control is a part.

The article on page 42 describes IMS in detail, touches on the upcoming challenges as IMS evolves, and highlights both the impact that IMS has on network infrastructures and the expertise Nortel is leveraging from its IMS experience.

The combination of this longstanding expertise and experience, across all technology areas, has created a unique packet-optical skill set in Nortel that is specifically suited to understanding the infrastructure requirements of tomorrow's services and applications, as well as the impact they will have on the underlying infrastructure.

Going forward, Nortel intends to continue to be at the forefront of driving the technology advances needed to lower network operations costs even further, meet the ever-rising demands for capacity, and ensure a highly intelligent, flexible network foundation for delivering multimedia-rich communications experiences, with extremely high reliability, a superior quality of experience, and unprecedented levels of personalization. ■

***David Hudson*** *is Network Architect, Mobility and Converged Core Networks.*

# Next-generation Ethernet: The key to infrastructure transition

by Phil Edholm and Paul Littlewood

Spurred by the huge growth in packet traffic on the network, carriers are seeking to transition their circuit-based infrastructures onto a common packet-based network to cut costs and efficiently deliver emerging high-bandwidth services, such as IPTV, mobile video, and multimedia business services. With its low cost and networking efficiency, Ethernet, already the dominant networking protocol in the enterprise space, offers an ideal solution for effecting this transformation. To ready it for this role, switched Ethernet – historically a "best-effort" technology – is being enhanced with the standards-based OAM, reliability, security, and scalability attributes that carriers have come to expect from their networks. Nortel has a number of industry-leading Ethernet technologies available for wavelength- and SONET/SDH-based solutions and is now at the forefront of the effort to standardize new switched Ethernet capabilities in the IEEE and ITU-T. Nortel continues to invest in all network Ethernet technologies and is significantly increasing investment in metro Ethernet functionality in its Optical Multiservice Edge 6500, Ethernet Routing Switch 8600, Metro Ethernet Routing Switch 8600, and Optical Metro 3500 product portfolios.

With much of the industry's attention focused on IP, Ethernet – even with its installed base of hundreds of millions of ports in the enterprise and broadband residential space – has in many ways been the underdog of the information age. That is about to change.

While Ethernet has been a staple in the enterprise local area network (LAN) for two decades – as the networking protocol used from the desktop, from the wiring closet, in the campus core, and increasingly for interconnections between sites – it is just now beginning to transform carrier networks.

There, its potential impact is even more profound, particularly as carriers seek to converge their existing circuit-based infrastructure and packet-service networks onto a single common network to dramatically reduce the cost and complexity of supporting services on multiple overlaid technology platforms.

A next-generation metro Ethernet network – with enhanced operations, administration, and maintenance (OAM), reliability, security, and scalability added to Ethernet's natural simplicity and efficient networking – offers an ideal solution for this transformation, promising to deliver the manageable, cost-effective converged networks that carriers need in order to thrive in today's very competitive environments.

As an original and leading advocate of carrier-class Ethernet-based solutions, Nortel is not only driving the standardization of these next-generation Ethernet capabilities with the Institute of Electrical and Electronics Engineers (IEEE) and International Telecommunication Union (ITU-T), but is also implementing them in its Ethernet Routing Switch 8600 (ERS 8600), Metro Ethernet Routing Switch 8600 (MERS 8600), Optical Multiservice Edge 6500 (OME 6500), and Optical Metro 3500 (OM 3500) product portfolios (see page 8).

Offering innovative Ethernet and optical portfolios that fit in both enterprise and carrier networks, Nortel is uniquely positioned to help customers build robust and efficient networks to fulfill these opportunities. Indeed, common chips, control protocols, security solutions, and other elements make Nortel's investments in Ethernet readily extensible to both the enterprise and carrier markets. Nortel also has the proven ability to deliver truly carrier-grade capabilities, as well as considerable experience in deploying Ethernet and optical technologies into a large number of carrier networks and an even larger number of enterprise networks.

### Ethernet evolution

Together with IP, Ethernet now defines communications. It has moved from its inception as a small-office printer and disk attachment system in 1982 to become the dominant transport technology in the enterprise today, displacing over the years such technologies as Token Ring, Token Bus, Fiber Distributed Data Interface (FDDI), and Asynchronous Transfer Mode (ATM). It has evolved from its humble 3-megabit-per-second (Mbit/s) beginnings in the original Blue Book (the original Digital Equipment Corporation, Intel Corporation, and Xerox Corporation specification for Ethernet) to speeds of 10 gigabits per second (Gbit/s) and beyond. It has moved from coax to twisted pair and fiber, to being the basis for such advanced wireless technologies as IEEE 802.11 (WLAN) and IEEE 802.16 (WiMAX). And virtually every packet in today's networks starts and ends its existence as Ethernet. Even the next generation of cellular networks will use this simple, ubiquitous, and extensible technology.

Through this entire transition, the true genius of Ethernet remains the design of its packet structure: a variable length from 64 to 1,530+ bytes for payload efficiency; a simple, self-describing framing format applicable to a variety of transmission media; and a 48-bit globally unique address space with an optional 12-bit provisioned network identifier.

The address structure, with 24 bits assigned to a component vendor and the other 24 assigned by that vendor, can support a huge number of devices [$2^{48}$, or about 280 trillion network elements (NEs)].

At its core, Ethernet is a networking protocol: the inclusion of source and destination addresses in the packet enables the network to make node and connectivity discovery and packet-forwarding decisions. These capabilities are the keys to the future of Ethernet.

With its inherent networking capability, Ethernet is very flexible in terms of use. As shown in Figure 1, Ethernet can be deployed from a logical level in three ways:

- as a point-to-point link-level framing protocol, where addresses are not used for forwarding;
- as a many-to-many switching protocol, where the addresses are used to forward the packets; and
- as a many-to-one aggregation protocol, where addresses also are used to forward the packets.

One of the key questions now facing network designers is how best to use Ethernet in converged networks. While some would argue that the future of Ethernet is merely as a link protocol between Layer 3 (L3) Internet Protocol (IP) routing nodes, this approach fails to take advantage of Ethernet's extremely low-cost packet aggregation and switching capabilities in the Layer 2 (L2) networking domain.

### Metro network technology options

For many years, the carrier infrastructure has been based on SONET/SDH and TDM networking in the form of DSx and Ex links in the digital domain, and virtual tributaries (VTs) and virtual containers (VCs) in the optical domain. However, the growing dominance of packet traffic and packet services – including such emerging high-bandwidth video services as IPTV, mobile video, and multimedia business services – has exposed TDM switching's inflexibility and inefficiency for handling packet traffic.

To be compatible with TDM switched networks, campus protocols have been mapped to the telecommunications protocols (e.g., DSx, Ex, OCx, STMx, frame relay, and ATM) at the edge of campus networks. While this approach worked well when the majority of enterprise traffic was voice and the packet-to-circuit mapping demands were small, the growth in enterprise packet traffic (including Voice over IP as well as data) and inter-campus traffic is now creating complexity and inefficiency for the enterprise at the metro edge.

## Nortel's Enterprise and Metro Ethernet portfolios

To support next-generation Ethernet, Nortel offers a market-leading portfolio, including: for enterprise applications, the Ethernet Routing Switch product family; and, for the carrier environment, the Metro Ethernet Routing Switch 8600 family, the Optical Metro 3500, and the Optical Multiservice Edge 6500 family.

The **Ethernet Routing Switch (ERS)** portfolio comprises families of fixed and stackable Ethernet switches, including the ERS 8600, that are designed to cover the whole range of enterprise needs, from wiring closets to large data centers. Switches can be configured to perform both aggregation and switching to support the next-generation Ethernet architecture.

The **Metro Ethernet Routing Switch 8600 (MERS 8600)** is the carrier network version of the ERS 8600. It includes features designed specifically for the carrier network to provide the robustness, security, and interfaces required for next-generation Ethernet applications. The Metro Ethernet Services Module 8668 in the MERS 8600 is a network processor unit (NPU)-based carrier-class solution that provides an Ethernet user-to-network interface (UNI) and services demarcation point, which is used to encapsulate traffic for scalability and customer separation and to enable measurement of end-to-end service performance.

The **Optical Metro 3500 (OM 3500)** multiservice platform is a SONET/SDH platform designed for a very broad range of applications, including Ethernet service transport and aggregation. The OM 3500 is a compact, flexible platform that supports a custom mix of DS1, DS3, DS3/VT Transmux, EC1/STS1, storage (Fibre Channel and FICON), Ethernet (10/100 Base-T, 100 Base-FX and 1000 Base-SX/LX), OC-3, OC-12, and OC-48, as well as a family of Ethernet Packet Edge Modules equipped with Ethernet UNIs and network-to-network interfaces (NNIs).

The **Optical Multiservice Edge 6500 (OME 6500)** is a next-generation optical convergence platform that blends multiple services and network layers (Layers 0, 1, 2) onto a single platform. This platform gives carriers the flexibility to select not only which functions the platform will provide – optical transponder, TDM switch, or packet switch – but also any mix of them in the required proportion. The OME 6500 supports an integrated Layer 2 service switch (L2SS) capability via its NPU-based design.

More information on these products can be found at www.nortel.com.

Moreover, with Ethernet supplanting the multitude of campus data protocols and now predominant in the enterprise, service providers are being asked to provide simpler on-ramps to their networks and to simplify their networks around a common Ethernet transport protocol.

In truth, service providers have offered Ethernet services to enterprises for more than five years. These services allow enterprise customers to avoid the cost of packet-to-circuit mapping at the edge of their networks – replacing it with a simple Ethernet handoff to the service provider's network.

With the rapid and continuing growth of Ethernet transport and ser-

*continued on page 13...*

**Figure 1. Ethernet applications**

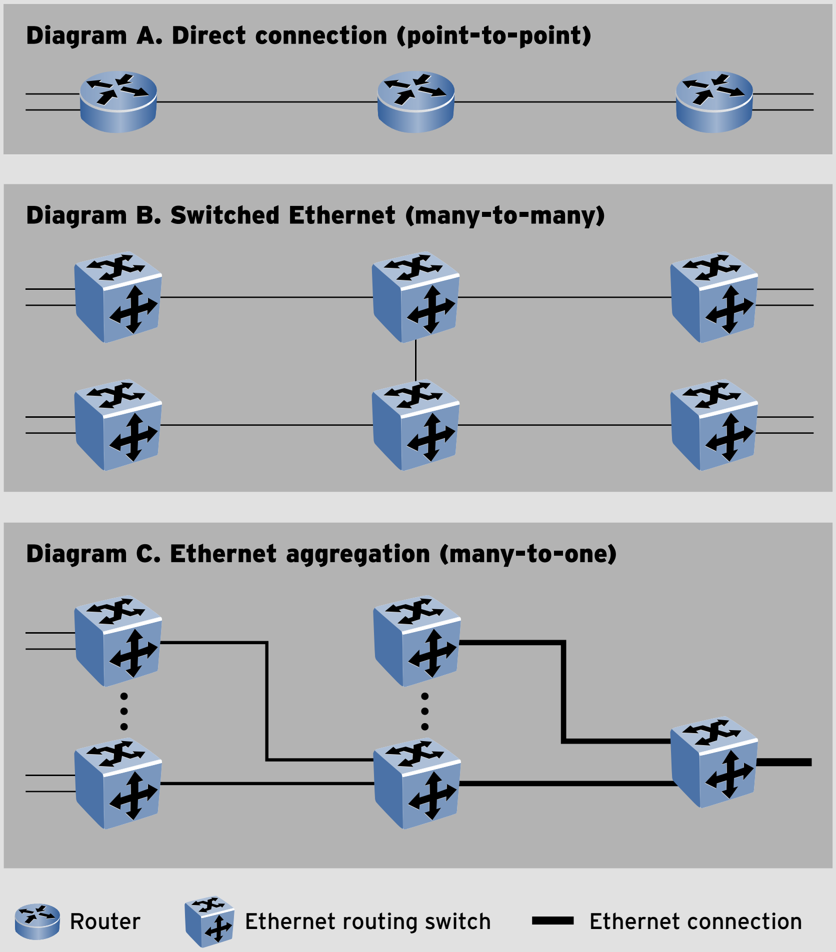


Ethernet's flexibility enables it to be utilized in numerous ways when building a network. In its most basic application, Ethernet is used to interconnect routers to provide very low-cost connectivity (Diagram A). In its most comprehensive application (Diagram B), Ethernet forms a network layer in its own right to provide low-cost connection-oriented or connectionless services. Another application of Ethernet (Diagram C) is to aggregate traffic from multiple edge network nodes to a smaller number of core network nodes.

**Direct connection:** As shown in Diagram A, Ethernet can be used to connect routers directly without using its networking capability. In this application, basically only the framing capability of Ethernet is used (to appropriately order information). Ethernet in this case provides an efficient means to multiplex higher-layer packets and measures transmission accuracy using its error-checking capability. This application would be used in an all-IP network where packet-routing decisions are made only by Layer 3 network elements.

**Switched Ethernet:** As shown in Diagram B, Ethernet routing switches, such as Nortel's Ethernet Routing Switch 8600 (ERS 8600), can be connected together to provide very low-cost packet delivery to a multitude of nodes. This application is commonly used in both enterprise campuses and carrier networks, where Ethernet switches are used to interconnect other network elements (such as computers or servers) to routers, or digital subscriber loop access multiplexers (DSLAMs) to content servers. Switches use the Ethernet addresses assigned to the client ports to build the location information needed to forward packets appropriately. The location of these addresses is discovered by the network of Ethernet switches using traditional broadcast learning.

**Ethernet aggregation:** Diagram C shows Ethernet routing switches being used to aggregate traffic (and distribute it toward the network edge) from remote nodes to core nodes, for example. (A dual-star architecture would likely be used for resiliency. However, to simplify the diagram, the architecture is not shown here.) This configuration is simpler to administer because it is inherently loop-free and therefore does not require the spanning tree protocol for resiliency. In carrier networks, this configuration would be used, for example, to connect customer routers or Ethernet switches to core routers. The network Ethernet switch in this instance could be a blade on a multiservice provisioning platform, such as Nortel's Optical Multiservice Edge 6500, or could be a Nortel Metro ERS 8600.

# Nortel's electronic Dynamically Compensating Optics (eDCO): An industry first

by David Donald

Developed by Nortel, electronic Dynamically Compensating Optics (eDCO) is a groundbreaking, industry-first method that uses electronic rather than optical technology to compensate for dispersion in optical networks. Applied directly at the signal source – that is, on the transmitting optical interface – rather than on the fiber link itself, eDCO dynamically adjusts the modulation of individual wavelengths to overcome the inherent physical limitations of fiber and enable fiber spans of approximately 3,000 kilometers (km) without dispersion compensation – a capability that dramatically simplifies engineering, installation, and management of the network (Diagram A).

Dispersion is one of the main impairments incurred by an optical signal as it propagates between the transmitter and receiver in fiber-optic transmission systems. Dispersion manifests itself in the "broadening" of the optical pulses in the time domain, which results in intersymbol interference that corrupts the transmission beyond recovery. Due to the nature of the optical fiber itself, the amount of broadening incurred is directly proportional to the distance transmitted. For lower-rate signals, where the duration of the optical pulse is longer, the impact of this broadening is relatively minimal.

As the transmission rate increases, however, this pulse broadening can significantly impair performance. At 10 gigabits per second (Gbit/s) and beyond, transmission without compensation has historically been limited to less than 110 km. Until very recently, this compensation has been provided by inserting coils of additional optical fiber into the transmission fiber link. Commonly referred to as dispersion compensating fiber (DCF), this additional optical fiber has propagation characteristics that are complementary to the transmission fiber, which counteracts the pulse broadening effect.

However, DCF has also created several complications for the operator:

• The compensation provided is based on the length of the piece of DCF; hence different coils are required to compensate for different values of dispersion. These packaged coils of DCF are called dispersion compensating modules (DCMs).

• The DCMs themselves introduce significant optical loss, which requires additional optical amplifiers within the transmission link.

• Because different types of commercially available single-mode transmission fiber have different dispersion values for given distances, different coils are required to support each type – all of which have to be ordered, installed, and spared by the operator.

Given these drawbacks, many equipment providers and suppliers have been pursuing alternative methods of compensating for dispersion over the past several years. For example, there are now commercially available application-specific integrated circuits (ASICs) that provide receiver-based compensation for about 110 km of transmission fiber.

## Nortel's approach

Nortel has pursued another approach to solving this problem by applying electronic dispersion compensation technology. eDCO is unique in the industry and a world first in offering two distinct characteristics:

• it can provide dispersion compensation for up to approximately 3,000 km of transmission fiber; and

• compensation is applied to the signal transmission, in the optical transmitter, rather than on the transmission fiber itself or at the receiving end.

Briefly, some of the commercial reasons for Nortel adopting this innovative approach are:

• It significantly reduces the cost and complexity of optical transport networks (whether metropolitan, regional, or long-haul) by eliminating the need for both dispersion compensation modules (DCMs) and the associated optical amplifiers needed to overcome the loss created by the DCMs (for any transmission distance of up to approximately 3,000 km).

• It improves the reliability and availability of the optical network by reducing the number of active components involved in the transmission of the signal.

• It enables the evolution of optical networks to agile, flexible networks where service initialization and restoration can be achieved at the optical layer without significant re-engineering or re-characterization – in a matter of minutes rather than days or weeks.

In developing the technology, Nortel took advantage of the often ignored fact that an optical signal is an electromagnetic field, and therefore has a phase response in addition to a magnitude response as a function of frequency. The broadening of the optical pulses – the dispersion that they experience in the time domain – is a result of the phase response of the transmission fiber.

Knowing this, the solution to compensating for this response is to manipulate the phase of the optical signal such that the complex conjugate of the signal, multiplied by the response of the transmission fiber, leads to the desired optical pulse at the receiver. Fairly straightforward signal-processing techniques can be used to compute how much pre-distortion of the signal is required to provide this compensation.

The technology required to implement this manipulation is where Nortel has truly broken ground. Commercial-grade, real-time digital signal processors can operate on signals running at perhaps several hundred megabits per second (Mbit/s). However, to provide the necessary functions to operate on a signal with a data rate of 10 Gbit/s, six trillion arithmetic operations per second are required. Nortel developed a custom ASIC – referred to within the development team as the “WARP” chip – to achieve this speed.

The WARP ASIC is implemented in standard silicon BiCMOS and accepts the 10-Gbit/s client data stream. It manipulates the signal to compensate 100 percent for the dispersion present in the optical fiber, via a complex finite impulse response (FIR) filter operating at twice the data rate.

The other significant technological advancement within the WARP ASIC is the 6-bit resolution, 21-giga-sample-per-

**Diagram A. Simplifying the network with eDCO technology**

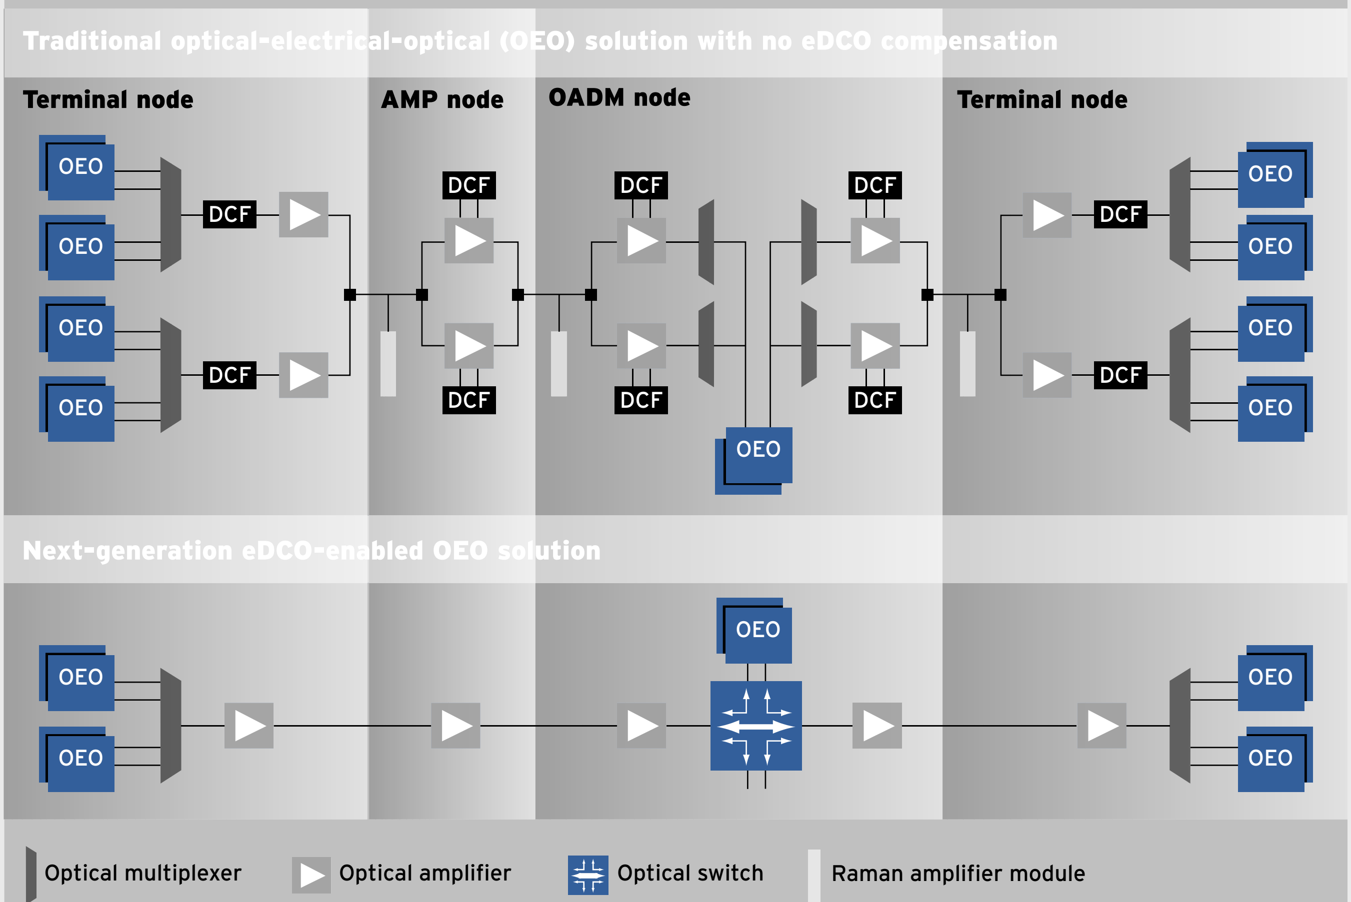


Nortel’s electronic Dynamically Compensating Optics (eDCO) technology reduces the cost and complexity of optical transport systems by eliminating 100 percent of the optical compensators and nearly 50 percent of the optical amplifiers in the photonic layer.

The upper half of the diagram shows a traditional optical system. Electrical-to-optical signal conversion is first performed by optical-electrical-optical (OEO) circuit packs. Optical signals are then multiplexed, amplified, and compensated using traditional dispersion compensating fiber (DCF). Transport beyond 100 kilometers or so requires additional amplification and optical compensation at intermediate sites at amplifier (AMP) nodes and optical add-drop multiplexer (OADM) nodes. At the terminal node, the optical signal is converted back to an electrical signal using OEO.

The lower half of the diagram shows an optical system equipped with eDCO-enabled OEOs. Electrical-to-optical conversion of the signal is again the starting point, but now eDCO is applied at the origination of the transmission prior to multiplexing. By eliminating the equipment needed for optical compensation, eDCO reduces both the complexity of deployment and the initial capital investment required.

second digital-to-analog converter, two of which are used to create the complex optical signal. Traditionally, only the amplitude of the optical signal has been modulated to transfer information from transmitter to receiver. With eDCO, the phase of the transmitted signal can be altered in addition to its amplitude – and altered in a manner that pre-compensates for the unique characteristics of a particular fiber link (see Diagram B).

The digital-to-analog converter (DAC) combines 21-GHz digital logic and high-speed, high-linearity, low-noise analog signals – a combination that was considered unachievable by Nortel's competitors.

To achieve this rate, Nortel designers developed a number of key components,

**Diagram B. How eDCO works**

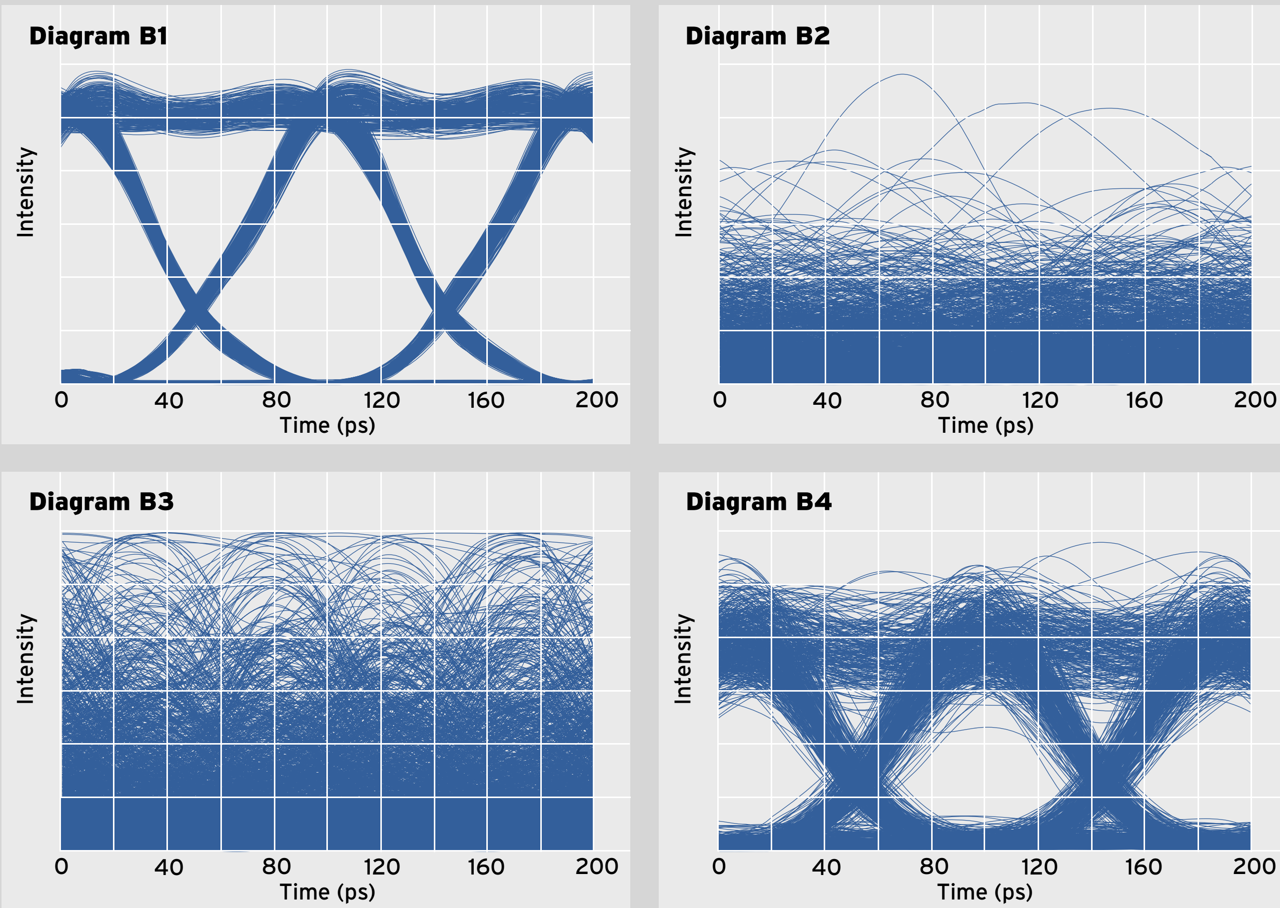


Nortel's electronic Dynamically Compensating Optics (eDCO) technology alters the phase of a transmitted signal by pre-compensating for the unique characteristics of a particular fiber link.

Diagram B1 – commonly known as an eye diagram (which traces the optical intensity vs. time for combinations of symbol patterns) – shows a traditional modulated 10 gigabit-per-second (Gbit/s) signal in the time domain, representing the 1s and 0s in the data transmission and the transitions between them. When the signal is propagated through fiber (in this case 1,600 kilometers of fiber) and no compensation is applied, the information in the signal at the receiving end, shown in Diagram B2, is completely destroyed and no decision can be made on whether a 1 or 0 was transmitted.

Diagram B3 shows the signal with eDCO compensation applied. eDCO pre-distorts the original transmitted signal. Although there is no discernable 1 or 0 obvious within the transmitted eye, this pre-distortion is such that it compensates for the propagation characteristics of the fiber, resulting in a signal at the receiving end (Diagram B4) that has an obvious eye where the 1s and 0s and the transitions between them can be clearly discerned.

including an array of 128:1 multiplexers to generate the high-speed digital input for the DAC, an on-chip voltage controlled oscillator, and a phase locked loop (PLL) frequency synthesizer to provide the necessary low-jitter 21-GHz system clock, as well as the actual DAC block and a high-linearity output amplifier. Design at such a high clock rate required precise design with little margin for error, and involved optimizing the operation of each transistor, performing electromagnetic simulations of interconnects and passive components, and handcrafting the circuit layout to control delays and parasitics.

Based on Nortel innovations that have resulted in several patent applications, the eDCO technology is now fully productized and available on the Optical Multiservice Edge 6500 (OME 6500), Nortel's next-generation packet-optical convergence platform. Nortel has shipped more than 1,000 transponders with eDCO in OME 6500 broadband deployments with such major customers as Comcast Corporation, SURFnet b.v., and RISQ (Réseau d'informations scientifiques du Québec).

eDCO is truly a groundbreaking technology that is changing the way Nortel's customers deploy, manage, and maintain their optical networks – delivering the agile photonic networking that provides the flexibility they have long sought.

***David Donald*** *is Leader, Optical Terminal Systems and Architecture, Optical Networks.*
*Acknowledgement:* ***Peter Schvan,*** *Team Leader, High-Speed Circuit Design, for the section on the digital-to-analog converter technology in the WARP ASIC.*

vices in the enterprise, however, service providers must now decide on the most effective network solution to efficiently support this traffic. To date, VLAN-based switched solutions have dominated but they will not be adequate for meeting the future requirements of massive scale, traffic engineering, remote network operation, and precise service definitions.

The first generation of technology, Multiprotocol Label Switching (MPLS), was designed and optimized to fit the requirements of the network core, where modest scaling and internetworking were required. However, while adequate for this role, MPLS faces a different set of requirements in the metro. Although it addresses some of the requirements, MPLS must contend with major challenges in specific areas, such as scale and operations simplicity. Moreover, MPLS adds significant deployment and operations cost because of its use of a complex control plane and the interdependence between the control plane and data plane for basic OAM (see article on page 25).

Originally, the role of MPLS was as an accelerated transport technology controlled by an IP control plane, but it has subsequently acquired a distinct role as a service infrastructure, exemplified by Pseudo Wire Encapsulation (PWE) or RFC2547 IP virtual private networks (VPNs). As such, it has become a complex technology employing many protocols, while at the same time not offering simple network maintenance tools.

Moving this complexity to the network edge results in a proliferation of expensive and high-maintenance routers. With network edges typically serving 200-2,000 end customers, a network with 500,000 end customers, for example, would require more than 250 routers. Other challenges involve IP control plane scaling and MPLS trunking management. Even without individual service paths, a network with 250 provider edge (PE) routers and four service classes (one for signaling) would have $250^2/2$, or 31,250 paths. If individual customer paths must be managed, that number can explode into the tens of millions.

A new generation of Ethernet technology offers carriers the opportunity to simplify the edge of the network. As with enterprise deployments, Ethernet's efficiency in aggregating traffic from groups of customers enables a more efficient architecture, either eliminating PE router complexity or restricting it in the hierarchy to the core metro or regional points of presence (POP) level. For example, with a ratio of 20 access POPs to each metro or regional POP on average, the number of PE routers could be reduced from 250 to 25 and the number of MPLS paths could drop to less than 1,250 – a reduction in path complexity of 99.975 percent. In addition, an Ethernet-based solution shifts management complexity from unmanned locations to key core POPs where it can be handled more effectively.

Ethernet has been an infrastructure technology to IP for decades. Indeed, it has already established itself as the common physical layer used between routing devices in provider offices and is frequently used as an interface to the end customer. And, for a while now, the IEEE has been enhancing Ethernet's networking layer with carrier network features. With its rapidly expanding capabilities, Ethernet has now emerged as an attractive option that is well-suited to future enterprise and metro networks.

### The simplicity of L2 Ethernet networking

The advantage of interconnecting groups of devices using L2 Ethernet forwarding versus an all-routing (all-IP) L3 infrastructure is based on fundamental differences between the two technologies, especially in how addresses are assigned and managed.

Ethernet is based on purely random address allocation: the address of each Ethernet device – the Media Access Control (MAC) address – is permanently assigned at manufacture and is globally unique. The MAC address is made hierarchical with an optional provisioned component, the VLAN, which is part of the forwarding decision and was created originally to improve performance and security for sub-groups of LAN users.

In IP, on the other hand, the structure of the network is based on assignment of specific parts of the address to different regions or "subnets" of the network. These regions, whether physical or logical, are not associated with a device at manufacture. In the IP world, addressing is something that must be dynamically managed, along with the tables for forwarding in routers.

Within a domain of Ethernet devices, forwarding through a network is accomplished by discovering the location of a specific MAC address, learning how to reach that address, and then using the destination address of a launched packet to decide how to forward it through a node. L2 devices, such as Nortel's Metro Ethernet Switch 8600, are generally much simpler than L3 devices, such as routers. While L2 devices today support the features of L3 class-of-service (CoS) prioritization, they are not tasked with the expense of building the full IP forwarding environment, which requires knowledge of the network topology. As a result, L2 Ethernet domains in the network architecture can dramatically simplify the operation of the overall network.

The metro network can be sub-divided into two domains: aggregation nearer to the edge and switching at the core. L2 Ethernet can be employed in both places (Figure 2). Ethernet edge aggregation removes the complexity of L3 IP forwarding (routing tables, address management, etc.) from the multitude of edge equipment that exists to simplify network configuration and operation, while L2 Ethernet forwarding in a metro core can yield advantages through the low-cost interconnection of routers without having to employ the complexity of a router. It is important to understand, however, that Ethernet is compatible with a metro network model based on either MPLS in the core or Ethernet in the core.

## Ethernet requirements for service provider networks

To be deployable as a carrier network infrastructure, the industry must add to Ethernet the scale and management capabilities that carriers have come to expect from SONET/SDH and TDM – while building on the Ethernet capabilities (simplicity and low cost) already proven in enterprise deployments.

Ethernet, therefore, is being enhanced to support a number of concepts fundamental to carrier transport networks, including:

• the scalability to accommodate hundreds of thousands of services;
• independence between the services offered to a network user and the services used to maintain and manage the network infrastructure itself (for efficient network operation);
• separation between the carrier infrastructure and services (for secure and robust services);
• independence of the networking technology from the physical medium (i.e., the flexibility to carry traffic over copper, fiber, and wireless media);
• separation of connectivity from protection (i.e., the ability to support subnetworks for maintenance and protection, as well as a wide variety of resiliency techniques, such as redundant links and switched rings);
• high availability and resiliency (low service downtime);
• specified and guaranteed service attributes (limits on delay, jitter, bandwidth, etc.) plus service performance data for service level agreement (SLA) assurance;
• the ability to flexibly carry any service (i.e., through Pseudo Wires); and
• support for remote operations (such as configuration, network maintenance, performance monitoring, and fault isolation).

Nortel has gained considerable experience in these areas since the late 1990s, through helping carriers build service networks based on many aspects of this model using Nortel's ERS 8600, OM 3500, and OME 6500 portfolios – and the solutions have always encompassed robustness features required for infrastructure use.

## Provider Ethernet

While standards for Ethernet services (e.g., E-Line, E-LAN, and E-Tree) have been defined by the Metro Ethernet Forum (MEF) – a group of more than 80 carriers, equipment vendors, and end users that is driving the definition of and applications for a truly carrier-class Ethernet – and have existed for a number of years, carrier infrastructure requirements had not been expressly considered from a network viewpoint by any standards body until recently.

To address the carrier infrastructure requirements, Nortel has taken a leadership role in a number of IEEE and ITU-T standards committees (such as IEEE 802.1ah and ITU-T Y.1711) in defining work programs and in making many key contributions to new standards.

At the same time, we have chosen to continue to invest in products and standards consistent with the Provider Ethernet architectural models being described in the IEEE. To extend enterprise bridging technology for application in carriers' networks, IEEE Committee 802.1ah is working on Provider Backbone Bridging (PBB), often called MAC-in-MAC. PBB is designed to cleanly separate the end-customer network from the carrier infrastructure, enabling secure customer and content separation and bringing massive scalability (up to millions of service instances per metro) to Ethernet. Nortel has been a driving force in the standardization of PBB, and an early implementation of PBB has been available on Nortel's Metro Ethernet portfolio for several years, while a fully standardized version will be delivered later this year as the standard is ratified.

PBB is applicable to Ethernet both at the customer edge and in the metro core. It will interwork with current Ethernet access techniques based on VLAN

**Figure 2. Metro Ethernet model**

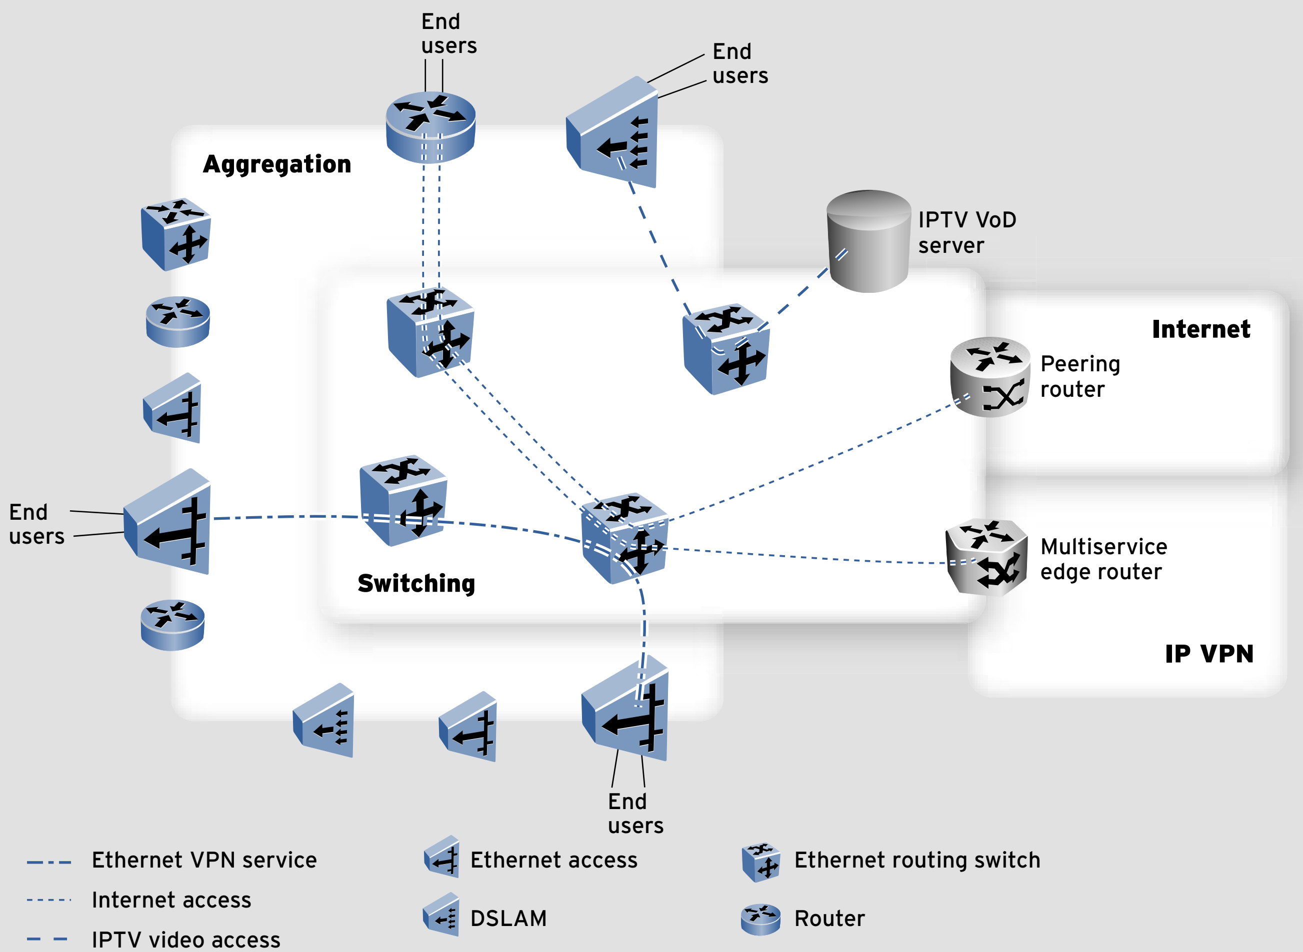


The metro Ethernet infrastructure model has two basic components: an aggregation subnetwork and a switching subnetwork.

The aggregation subnetwork combines incoming traffic from multiple end-user sources and requires only limited connectivity. End-user packets are statistically multiplexed together, with forwarding priority (in the case of two packets contending for a common port) based on a set of programmable rules that ensure high-priority or delay-sensitive traffic is transmitted more quickly through the network. For outbound traffic (toward the end user), the aggregation subnetwork distributes an Ethernet packet to the appropriate location based on the destination address assigned by the source node. Prioritization of traffic is less of an issue here, but the egress port of the Ethernet switch must ensure that the traffic is shaped (traffic capacity peaks are leveled off using first-in, first-out memory) to comply with the capacity of the downstream network.

The switching subnetwork has greater connectivity than the aggregation subnetwork but must be carefully designed to eliminate or minimize loops that have a detrimental effect on traffic throughput. As shown by the different dotted lines in the diagram, the switching network connects the traffic from the end user via the aggregation subnetwork either to:

- another end-user site on the same or different segment of the aggregation subnetwork in the case of an Ethernet VPN service;
- a core IP node, such as a multiservice edge router or an Internet peering router; or
- a video server for IPTV services, such as an IPTV video-on-demand (VoD) server.

In the switching subnetwork, information used at an individual switch for forwarding a packet to the appropriate output port is either learned using a technique called bridge learning (the technique deployed in networks today) or provisioned by an external operations support system or control plane. The latter approach has been developed by Nortel in conjunction with a major carrier and is called Provider Backbone Transport. This technology allows connections to be engineered between end nodes with defined transmission parameters to ensure, for example, that sufficient capacity is available between nodes or that transmission delay specifications are not exceeded. As a result, stringent and contractual service level agreements can be defined and maintained.

separation and VLAN stacking (known as QiQ) that have been defined by IEEE 802.1ad, adding value to deployed networks through service and network scaling by eliminating the MAC explosion issue inherent in QiQ networks. This issue arises because QiQ networks rely on the end-customer MAC address for switching. These networks may suffer because many client MAC addresses could be flooded to the network at each network port and propagated through the metro. As a result, client MAC addresses permeate the service provider network, where only a single MAC address representing the client service endpoint needs to be used. The additional carrier MAC encapsulation defined in PBB solves this problem by encapsulating the multitude of client MAC addresses at the edge of the network (Figure 3).

One of PBB's strengths is its support for multipoint-to-multipoint services (E-LAN in MEF terminology). Point-to-point services (E-Line) are supported, but PBB lacks specific mechanisms for traffic engineering (TE) of individual services. Recognizing this limitation and understanding the scaling and operations issues associated with TE-capable MPLS cores, Nortel, in conjunction with a major European customer, has created an innovative architecture called Provider Backbone Transport (PBT). Developed to enable Ethernet-based metro networks to work seamlessly with carriers' existing MPLS core networks, PBT incorporates into Ethernet the deterministic, carrier-grade connection management characteristics of SONET/SDH.

PBT leverages PBB but allows paths between service endpoints to be explicitly engineered across a network. With PBT, the network can be designed to support rigorous Ethernet SLAs. Furthermore, the use of PWE over PBT (a form of "Dry Martini" as specified by IETF) allows the determinism of PBT transport to be extended to carry such legacy L2 technologies as frame relay and ATM. (For more detail on PBT, see the article on page 25).

Although PBB and PBT standards define much of the technology, there are components of Provider Ethernet that are worth highlighting, and some of these go well beyond the current standards work. Specifically:

• user-to-network interfaces (UNIs) encapsulate end-user traffic using techniques, such as MAC-in-MAC, that are designed with end-to-end data links to enable the derivation of, for example, service performance data (jitter, packet loss, etc.);
• network-to-network interfaces (NNIs) provide a means to efficiently hand off traffic between NEs without having to demultiplex it to a multitude of lower-rate signals or de-encapsulate it to the payload, thereby enabling overhead continuity to be maintained across the network;
• hierarchy, which is added by sub-dividing the network geographically, provides scale, regional independence, and management simplicity;
• tunneling provides its own hierarchy by "wrapping" other paths and services to simplify the interconnection of NEs, dramatically reducing the number of operations required; and
• service labeling provides an ability to explicitly inventory services and make them visible throughout a network so that clear service billing records can be built and problems affecting the network can be clearly related to service impairments.

Providers planning to offer Ethernet MAN and WAN services using the MEF service definitions will offer E-Line, E-Tree, and E-LAN options for both retail and wholesale services, while providing higher bandwidth (replacing or complementing E1/T1 digital private line) and managed SLAs. Backward compatibility with TDM networks is ensured using such adaptation technologies as Generic Framing Procedure and Virtual Concatenation sets of protocols (see page 18). Ethernet is envisioned to support a range of other service types, such as Internet access, voice access, VPN access, and content delivery. Similar to today's TDM network, operational procedures for in-service and out-of-service testing and reporting are being developed by the industry – a process into which Nortel is providing significant input.

As in enterprise networks, Ethernet in the next-generation carrier metro network will be used for efficient forwarding through the core. While MPLS was designed as an L2.5 protocol, it was also designed for simple forwarding, and lends itself to operating directly on an L2 infrastructure. The infrastructure can also take advantage of some of Nortel's leading resiliency differentiators, such as the Split Multi-Link Trunking and Resilient Packet Ring (RPR) protocols.

The distribution of synchronization across the network is another key component that must be incorporated in next-generation Ethernet. Today, network elements are synchronized to a hierarchy of very accurate clocks (an attribute know as traceability), and SONET/SDH and T1/E1 links are used to distribute synchronization. Because this mechanism may not be available in future networks, strategies must be developed to synchronize Ethernet systems to this hierarchy of very accurate clocks and provide clock traceability so that far-end systems or client systems (such as wireless base stations) can extract synchronization (see page 21).

## Network and service management

In addition to these networking capabilities, next-generation Ethernet network elements must also support other key transport network features, including the necessary instrumentation to measure network and service performance at key locations in the network – a capability that has been an indispensable feature of TDM networks (Figure 4). These measurements are processed locally into meaningful information and sent to network operations centers (NOCs) at secure locations, where carriers use the information to manage their networks and services.

Moreover, because carriers' customers

**Figure 3. Ethernet end-user and carrier address separation**

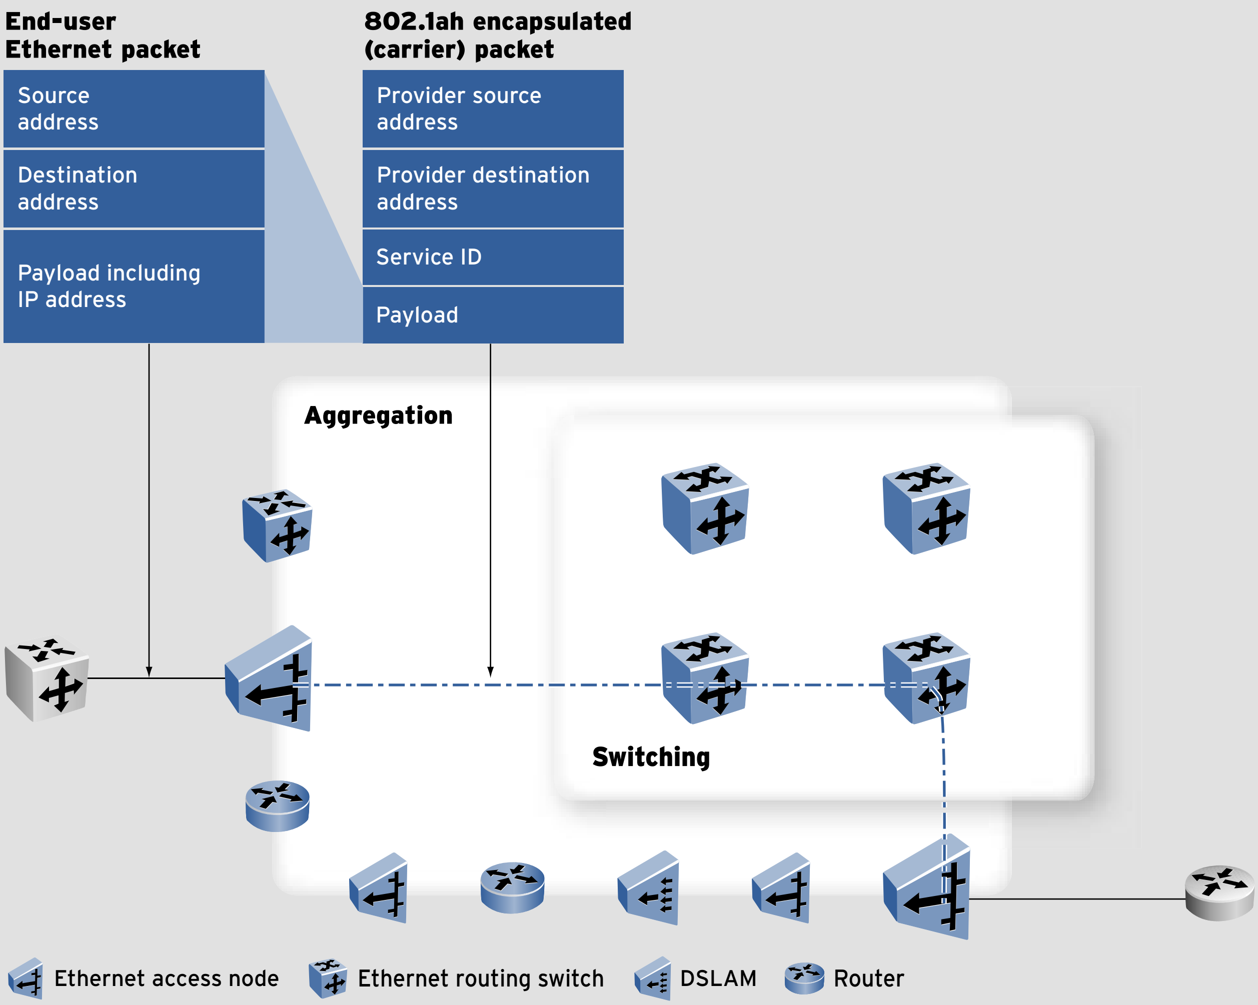


Nortel is working with such standards organizations as the IEEE and ITU-T to define standards that will enable packet infrastructures to be built that are as robust as today's TDM networks. One of the key concepts in building a robust infrastructure is a clear separation between the end-users' address spaces and the carrier's. In the case of an Ethernet Virtual Private Network (EVPN) service, this separation requires the encapsulation of the end-user's packet in a new packet created by the carrier's edge node (the Ethernet access node in the diagram). The method also ensures address separation between end users and results in a more secure network. Nortel has a number of products – namely the Optical Metro 3500, Optical Metro 1400, and the Metro Ethernet Routing Switch 8600 – that are capable of this function today using a pre-standard version of the technology.

The methodology behind this concept is being standardized in IEEE work group 802.1ah. As shown in the diagram, the end-user's Ethernet packet becomes the payload for the new 802.1ah encapsulated packet (the carrier packet). By provisioning the network address in the carrier packet, the carrier can determine and control how traffic is forwarded (as well as to where it is not forwarded) through its network, providing benefits for troubleshooting and security. In addition, the network is made more scalable since the number of addresses the carrier's core switching equipment must tabulate is limited to only the carrier's provisioned addresses. Only the carrier's edge node potentially has visibility to the end-users' MAC addresses. This node encapsulates many end-user addresses (in some cases thousands) to a far fewer number of carrier addresses.

modify their service requirements regularly, it is imperative that carriers are able to manage their networks automatically and remotely from the NOC without having to dispatch technicians. To enable these operations features, the IEEE 802.1ag working group is defining the key network and service maintenance points, the functionality at these points, the end-to-end communication mechanisms (intra-network to determine end-to-end performance data), and the message sets to communicate the data to the NOC. The hierarchy defined in

*continued on page 20...*

# Integrating packet technologies on optical transport products

by Craig Suitor

With the growing predominance of Internet Protocol (IP)-based services, access and transport networks are migrating toward Ethernet to take advantage of its low cost and networking efficiency. As a result, packet processing technologies (such as Ethernet switches, network processors, and traffic managers) that were primarily seen only in enterprise networks are starting to appear directly on transport equipment.

Nortel, in fact, is integrating best-in-class Ethernet technology and intelligence into its optical transport portfolio, primarily on its Optical Multiservice Edge (OME) 6500 offering. Integrating Ethernet technologies directly on transport equipment rather than deploying separate Ethernet and transport solutions allows customers to benefit from lower capital and operational expenditures.

At the same time, one must recognize that migrating these technologies onto legacy networks creates a hybrid network that generally will normalize to one technology. Since transport and access networks are still dominated by Synchronous Optical Networks/Synchronous Digital Hierarchy (SONET/SDH) and Plesiochronous Digital Hierarchy (PDH) interfaces, it is therefore still prudent to map packet technologies such as Ethernet into these interface types.

To maintain the attributes of a transport network in this transition, payloads generally need to be aggregated to provide more efficient use of bandwidth and segregated to keep customer or service separation. Other key attributes of a transport network are the ability to offer out-of-band management interfaces for operations, administration, and maintenance (OAM), distribution of network synchronization, fault sectionalization, and standards-based mid-span meet.

The most basic offering for Ethernet services on a transport network is Ethernet Private Line (EPL), which supports 10/100/1000 BaseT, Gigabit Ethernet (GE), and 10GE point-to-point services used for connecting two routers or for offering leased line service to enterprises. While EPL does not support packet aggregation, it can be used to bring Ethernet payloads to an aggregation point. Customer and service separation is achieved via TDM channel separation.

To support packet aggregation on a transport platform, packet processing functions need to incorporate encapsulation of packet payloads for transparency, as well as solid traffic management to effectively separate customer and service traffic flows so that they do not adversely affect one another.

## L2SS, RPR cards

To provide these capabilities, Nortel has added new cards to its optical portfolio, including the Layer 2 Service Switch (L2SS) card and its derivative, the L2SS PDH card, as well as the Resilient Packet Ring (RPR) family of cards.

Like EPL, the L2SS card supports point-to-point services but adds Ethernet switching capabilities. The card terminates multiple EPL circuits on its SONET/SDH interfaces and multiple Ethernet ports on its faceplate. These cards enable packet aggregation and switching between all available ports, with ingress classification, metering, marking and policing, queuing, shaping, and scheduling on the egress. Queuing, shaping, and scheduling are key capabilities supported by the L2SS's network processor unit (NPU)-based design, enabling separation of customer traffic so that one customer does not impact another.

A derivative of this card, the L2SS PDH, allows Ethernet-based services (such as L2 and L3 VPN services) to be delivered to a customer over existing PDH links (DS1, E1, E3, DS3). This capability enables operators to leverage their installed leased line access network so that they do not have to deploy an overlay network to support these services. The L2SS PDH card terminates EPL circuits originating at a customer premise, where the Ethernet frames are mapped into PDH signals and carried over a traditional PDH network. Adding this capability to transport platforms enables the delivery of a lower-cost L3 VPN service by eliminating the need to use channelized interfaces on the routers, which are expensive and generally underutilized. This capability also enables the delivery of an L2 VPN directly off the transport infrastructure, which allows carriers to generate a revenue stream from what was traditionally a cost center.

The RPR family of cards has all of the L2SS Ethernet functionality but replaces the point-to-point infrastructure with a resilient carrier-grade packet ring infrastructure. RPR is a metropolitan area network (MAN) technology that supports data transfer among stations interconnected in a dual-ring configuration. RPR combines the best capabilities of SONET/SDH and Ethernet, providing SONET/SDH's efficient support for ring topology and fast recovery from link failures (such as fiber cuts), as well as offering Ethernet's data transport efficiency, simplicity, familiarity, and cost advantages. RPR provides sub-50 millisecond (ms) ring protection and restoration times for each service carried over the ring.

Nortel began delivering RPR capabilities on its Optical Metro 3500 product in 1999 and has continued development through to the latest delivery of a 5-gigabit-per-second (Gbit/s) RPR on the OME 6500 in 2006. Nortel's latest offering on OME is 802.17-compliant and provides for up to two separate rings per card, with a total ring capacity of 10 Gbit/s.

The EPL, L2SS, and RPR cards exploit several technologies that have been introduced to support the transport of Ethernet traffic, as well as other 8-bit/10-bit (8B/10B)-encoded interfaces – such as Fibre Channel (the gigabit-speed network connection technology primarily used for storage area networks) – over SONET/SDH and PDH networks. These technologies include Generic Framing Procedure, Virtual Concatenation, and Link Capacity Adjustment Scheme.

**Generic Framing Procedure (GFP)** is defined in ITU-T G.8040 and provides a means of encapsulating variable-length higher-layer client signals over a transport network. GFP significantly reduces the amount of overhead associated with point-to-point transport, and provides a more efficient transport mechanism than the technologies that preceded it. A primary use of GFP is for mapping Ethernet packets into a SONET payload. A companion standard, ITU G.8050 defines a means to map Ethernet payloads into PDH circuits.

Both framed and transparent mappings into GFP are supported, for Ethernet and 8B/10B-encoded interfaces, respectively.

Framed mapping is generally used for Ethernet encapsulation because it enables a sub-rate payload to be allocated through the network. It does this by preserving the Ethernet frame from destination address to frame check sequence, while discarding the interpacket gap, pre-amble, and idle frames and utilizing pause frames toward the Ethernet interface for flow control. This mapping is more efficient than X.86 or Point-to-Point Protocol (PPP) and is based on standards that support multivendor interoperation.

Transparent mapping is used for block-code-oriented 8B/10B-encoded client signals, which are generally transmitted at line rate. Unlike the framed mapping mode where the Ethernet MAC layer is stripped before mapping into GFP, transparent mode maps all client bits, including the MAC, through the WAN. Using this mapping mode, a GE or FC100 signal would consume 21 STS-1s (or 1.09 Gbit/s) of network bandwidth. Although this mapping mode generally is carried full rate through the network, Nortel offers sub-rate capabilities and reach extension on its FC100 and FC200 cards by using the Fibre Channel protocols before mapping into the WAN. This capability allows Fibre Channel signals to be transported thousands of kilometers, whereas the protocol is generally limited to about 200 kilometers.

**Virtual Concatenation (VCAT)** is used to split up SONET/SDH payloads into "right-sized" groups to support different customers and services, improving network utilization significantly by effectively spreading the load across the whole network. Defined in ITU-T G.707 and G.783, VCAT enables SONET/SDH payloads that are larger than the base rate (51 Mbit/s for SONET and 155 Mbit/s for SDH) to be directed to where there is available bandwidth in the network in increments of the base rate. The payload can use available timeslots within the same SONET/SDH line or can be diversely routed through many different paths within the network. At the receiving end, VCAT payloads can be re-assembled with as much as 128 ms of differential delay between the arrival of the fastest channel and the slowest channel.

**Link Capacity Adjustment Scheme (LCAS)** is defined in ITU-T G.7042 and provides a means of adding and deleting members of a VCAT group while in-service. This capability enables a network provider to increase bandwidth as customer demand increases, by offering the customer a GE interface but selling bandwidth across the network in 51-Mbit/s increments. With LCAS, the network automatically adjusts the payload size if one or more members of the VCAT group is lost during a network failure, and then automatically re-adjusts it to add them back when the network recovers. This mechanism can be used in packet networks as a rudimentary protection mechanism by diversely routing the VCAT connections over many links and adjusting the available network bandwidth as failures occur.

***Craig Suitor*** *is Technical Advisor, Optical Systems and Architecture, Metro Ethernet Networks.*

802.1ag makes it suitable for service, tunnel, transport, and link OAM, providing a simpler, more functional, and more elegant OAM solution than that available from MPLS.

### Ethernet in the enterprise

Many of the application principles discussed earlier have been derived from the long history of deployment of Ethernet technology in enterprise networks. However, while Ethernet has become dominant in the enterprise, there is discussion, as there has been in carrier network deployments, about whether all packet control should reside at the L3 IP layer, to the exclusion of the L2 Ethernet layer. The all-IP approach, however, runs counter to the rule of thumb that has been used in networking for many years: *switch where you can and route where you must.*

The typical campus network consists of a single core with many wiring closets (for simplicity, redundant cores have been omitted from this discussion). The complexity of deploying and maintaining L3 IP devices is about an order of magnitude greater than that of L2 devices, for the reasons previously described. Add in the very large numbers of L3-capable edge nodes, and network complexity would effectively be increased by two orders of magnitude over an Ethernet approach, resulting in much

*continued on page 24...*

**Figure 4. Network and service management**

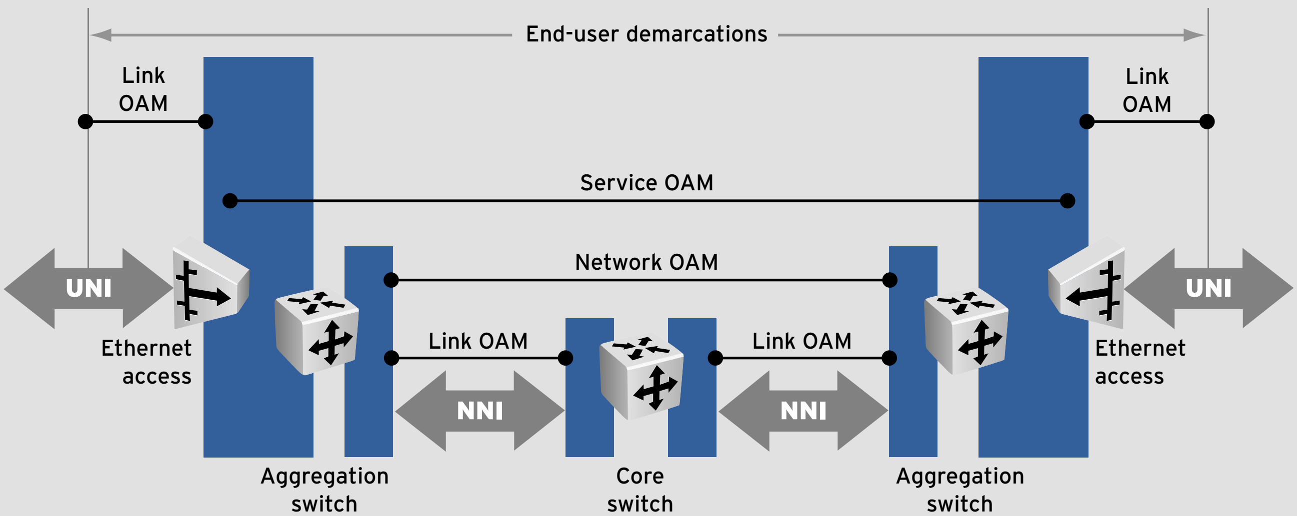


Measurement of service and network performance through embedded instrumentation is a critical capability in carrier networks. This instrumentation enables automated measurement and diagnosis of the carrier network, and continuous measurement of the quality of end-users' services. Specific operations, administration, and maintenance (OAM) packets are created at defined locations in the network and are examined at complementary locations in the network. Fault and performance data are derived and reported to the network operator.

Work is ongoing in IEEE Work Group 802.1ag and the ITU Ethernet OAM study group to define these network locations, the type of measurements performed, and the reporting of the data derived from the measurements. Nortel has led much of this standardization work and has ongoing development projects to implement this technology in its Optical Multiservice Edge (OME), Multiservice Switch (MSS), and Metro Ethernet Routing Switch (MERS) product portfolios.

In the figure, some key locations for network and service instrumentation are indicated by the blue bars.

For network OAM, maintenance packets are sent between peer pairs (between pairs of aggregation switches, for example) to verify path connectivity, along with a number of more detailed measurements, such as packet loss count or packet loss ratio. Information gathered from these locations can be used to determine overall network performance and to assist with fault finding.

For service OAM, service encapsulation points, indicated by the user-to-network interface (UNI) designation, are critical maintenance locations where service performance is measured, similar to the way that network performance is measured at network peering points. In this instance, the measurements (for packet jitter or round-trip delay, for example) are made to determine service performance and adherence to contractual service level agreements.

Link OAM provides measurements that allow the performance of an individual link in the network to be monitored. These measurements ensure that a service provider has performance visibility to the physical edge of its network (the end-user demarcation points in the figure) and aid in fault isolation within the network itself. The IEEE 802.3ah standard specifies the network architecture and types of measurements at the demarcation points.

# UMTS wireless backhaul synchronization requirements in metro Ethernet network environments

by Michel Ouellette

UMTS wireless operators are increasingly viewing an evolution to metro Ethernet networks as a solution to reduce complexity and cost in their backhaul radio access networks (RANs). These backhaul networks currently account for the largest portion of their operational expenditures due to the cost of leasing circuits from third-party wireline operators.

While their installed third-generation UMTS equipment is based on an ATM infrastructure (and their earlier second-generation GSM equipment on a TDM infrastructure), next-generation UMTS equipment will be IP-based and equipped with native Ethernet interfaces. With the continued penetration of Ethernet, metro Ethernet is becoming a viable and economical option for wireless backhaul.

Whether the backhaul RAN network is TDM, ATM, or Ethernet, there will always be one common aspect that is fundamental in a UMTS network: *All UMTS network elements are synchronized to a hierarchy of very accurate clocks within the network, an attribute known as traceability*. Traceability provides the basis for reliable UMTS network operation with limited operational and equipment cost.

The migration of the RAN toward a metro Ethernet network environment will require UMTS operators to evaluate different synchronization strategies for their UMTS network elements (NEs). As with today's ATM-based RAN, synchronization and clock traceability will not always be under the UMTS operator's direct control, but rather subject to service level agreements (SLAs) with third-party wireline operators across intermediary portions of the network.

Even with metro Ethernet networks, UMTS operators will continue to expect that the stringent synchronization requirements be met. Contrary to some beliefs in the industry, those requirements are not expected to be relaxed, and might even become more stringent. Careful network synchronization planning and design, backed up with extensive testing, will be of the utmost importance.

Nortel is well-positioned to help operators deliver superior UMTS service dependability and quality – including fewer dropped calls, higher speech quality, less speech clipping, faster data downloads, increased mobility, and spectrum reuse – by meeting the UMTS requirements for network synchronization (slip-free operation and traceability), node synchronization (estimation of timing differences between NEs), and radio interface synchronization (e.g., radio frame transmission and stable air interface frequency), the latter being of utmost importance to provide seamless soft handoffs between cells.

Nortel, for example, has taken a lead role in establishing requirements and specifications for modeling, building, and testing proper equipment, including its involvement in such standards work as the ITU-T G.8261 Timing and Synchronization Aspects in Packet Networks, a new recommendation released in April 2006, which has an initial focus on Ethernet transport. In addition, Nortel has made significant contributions to conferences run by the National Institute of Standards and Technology (NIST), Alliance for Telecommunications Industry Solutions, Inc. (ATIS), and the International Telecommunications Synchronisation Forum (ITSF).

In its leadership role in bringing synchronization to the packet world, Nortel is leveraging its expertise and experience in synchronization technology that extends back nearly 30 years in the circuit-switched world, beginning with the introduction of its DMS switching portfolio and continuing with its SONET/SDH transport and GSM wireless portfolios, where synchronization provided the heartbeat of these products. This leadership is now being extended to its Metro Ethernet Network portfolio, which consists of Optical Metro, Carrier Ethernet, and Multiservice Switching products. In addition, Nortel operates one of the industry's most comprehensive test labs and is using it to develop, test, and validate new synchronization solutions proposed within the industry.

## Today's environment

As shown in the upper portion of the diagram (see next page), the UMTS RAN network today is composed primarily of PDH circuits (T1/E1) to the NodeBs (base transceiver stations), and SONET/SDH circuits (OC-n/STM-n) to the Radio Network Controller (RNC) and UMTS switching centers, which include such network elements as Nortel's Mobile Switching Center (MSC), Gateway MSC (GMSC), Serving GPRS Support Node (SGSN), and Gateway GPRS Support Node (GGSN).

A NodeB today generally has from one to four T1/E1 circuits, but could grow to 16 T1/E1 circuits and even more to achieve the capacities required for deploying data technologies such as High Speed Downlink Packet Access (HSDPA). Currently, ATM services such as Inverse Multiplexing over ATM (IMA) are used to interconnect the UMTS NEs (NodeB - RAN - RNC) over an ATM backhaul network. In the future, the ATM backhaul will slowly transition to a carrier-grade Ethernet backhaul to provide a more efficient topology in metro environments, enabling the evolution to a hybrid backhaul network.

Inherent in PDH and SONET/SDH technologies are embedded network clock synchronization and distribution capabilities within the physical layer. As shown in the diagram, the larger NEs, such as the RNC, often have access to their own source of synchronization from primary reference clocks (PRCs), such as global positioning system (GPS) receivers, atomic Caesium frequency standards, or SONET/SDH circuits that provide traceability to other PRCs located within the network.

Direct access to PRC signals, however, is generally not available to the UMTS NodeB when it operates in frequency division duplex (FDD) mode, and the only alternative is to derive synchronization from the T1/E1 circuits coming from the ATM RAN backhaul network. As shown in the lower portion of the diagram where the backhaul is based on a future Ethernet backhaul infrastructure, the synchronization will be derived from separate T1/E1 circuits using ATM-Ethernet or TDM-Ethernet interworking and emulation switches, or directly from the physical layer of Gigabit Ethernet (GE) or Fast Ethernet (FE) switches and links. It is also expected that synchronization will have to be delivered over such Ethernet backhaul even when the underlying medium to the NodeBs consists of, for instance, digital subscriber loop (DSL) and microwave radio links.

Although today the RAN network is well-synchronized (primarily driven by the Layer 1 transport of SONET/SDH), problems due to synchronization can still arise as a result of loss of traceability to a PRC, impairments such as jitter and wander accumulation introduced by the transport and switching equipment of the backhaul RAN network itself, or impairments due to SONET/SDH mapping process and pointer activity.

As a result, most NodeB equipment manufacturers use sophisticated clock recovery engines and precision oscillators (e.g., ovenized quartz oscillators) to maintain the stringent synchronization requirements that are expected at the NodeB (e.g., the standard 50 parts-per-billion air interface requirement specified by Third Generation Partnership Project TS 25.402) and the E1/T1 jitter and wander network limits as defined in ITU-T G.823/G.824. Nevertheless, these solutions continue to be much more cost effective than having to deploy atomic clocks within a NodeB product or to deploy GPS receivers on a network-wide scale.

If the quality of the synchronization is unsuitable, the NodeB and mobile device may lose "lock" and cease to function for

### UMTS wireless backhaul synchronization

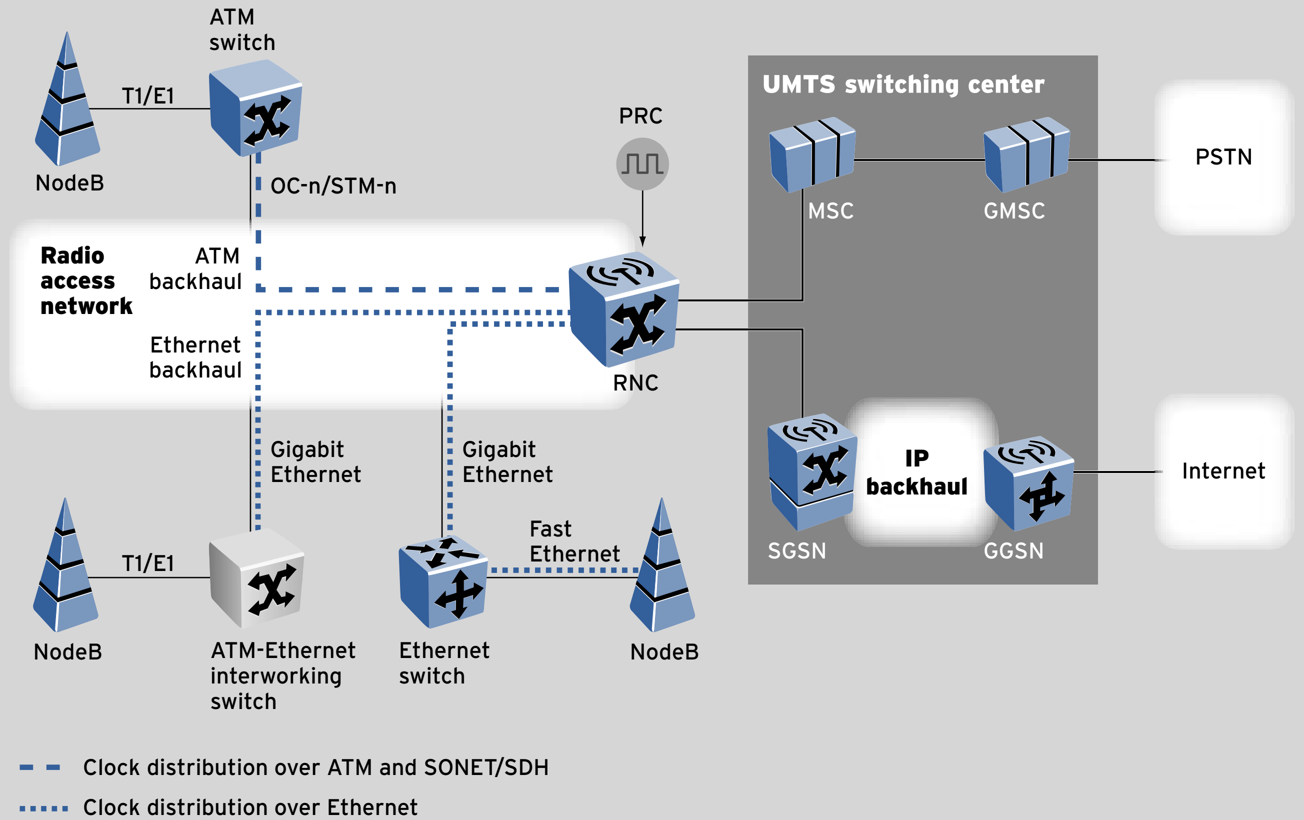

a period of time or inter-cell handoffs may fail, increasing the number of dropped calls and affecting the voice quality or decreasing the data throughput. Under extreme conditions, the NodeB may even go “off air” and provide no service at all. The result is loss of revenues and poor quality of service and experience. That said, PDH and SONET/SDH synchronization solutions in use today generally are very robust and are at the heart of most wireless and wireline operators’ networks.

**Tomorrow’s environment**

It is well known that Ethernet networks are asynchronous, meaning they do not provide any accurate clock synchronization and distribution capabilities. This limitation has led in the last few years to the development of technologies to transport synchronization over a packet network (e.g., Ethernet) to address such applications as UMTS wireless backhaul. However, most large standardization bodies have focused on protocols to adapt and interconnect traditional TDM/ATM networks to packet networks, and have overlooked the issue of maintaining network synchronization and distribution.

About a year ago, the International Telecommunication Union (ITU-T) was asked to look into this issue, and in April 2006 released a new recommendation called ITU-T G.8261 that deals with aspects of synchronization in packet networks, with an initial focus on Ethernet transport. Two different approaches to distribute synchronization over a packet network are being discussed.

The first approach is to distribute timing at Layer 2 (the datalink layer), using timestamps embedded within the Ethernet frame to carry timing information from a transmitter clock (e.g., a GPS receiver at the RNC) to the NodeB. Using very sophisticated algorithms, phase locked loops (PLLs), and temperature-compensated or ovenized quartz oscillators, the NodeB can reconstruct very accurately the frequency of the PRC transmitter clock.

One of the greatest strengths of this approach is its transparency to the underlying datalink and physical layers, which enables synchronization distribution across different technologies, such as DSL, microwave radio, WiMAX, and Ethernet. However, this technique is sensitive to impairments common to packet networks, such as delay variation, loss, and traffic congestion.

The use of novel Ethernet technologies proposed by Nortel, such as Provider Backbone Transport (PBT), minimizes the impact of such impairments on the quality of synchronization. PBT, for instance, provides deterministic traffic engineering and controllable performance guarantees to facilitate the use of Layer 2 synchronization techniques. In addition, these synchronization distribution techniques are often proprietary and are typically embedded into separate devices known as interworking functions. Some of the protocols being defined to transport timing information are IEEE1588-PTP, Timing over Packet (ToP), and modified circuit emulation.

The second approach is to distribute timing at Layer 1 (the physical layer), synchronizing the physical layer signals and bit stream of Ethernet links such as FE, GE, and 10GE in a point-to-point fashion. Synchronization at the physical layer is not new, and leverages well-known principles and experiences gained in SONET/SDH networks, including the use of synchronization status messages (SSM), and applies them to Ethernet. The IEEE 802.3 standard specifies that oscillators have a frequency deviation of no more than +/- 100 ppm (parts per million) for 100Base-TX links, but nowhere does it say that this signal cannot be accurate to 0 ppm, or traceable to a PRC clock. This technique is now part of ITU-T G.8261 and is sometimes referred to as Synchronous Ethernet. It is mainly driven and supported by large network operators and equipment vendors, and is applicable primarily to large carrier-grade Ethernet systems and new network builds.

Using a comprehensive test plan, Nortel is studying both approaches over a live production network that is maintained by Nortel’s Information Services (IS) group. The network consists of more than 10 network elements encompassing such technologies as Ethernet, RPR, SONET, and DWDM. Quality-of-service mechanisms can also be enabled, and packet network impairments (such as the addition and removal of background traffic load, link failure, and asymmetric paths) can be created in real time to stress-test the synchronization solutions under study.

Testing has verified that in general the Layer 2 approach – where timing information is carried within the Ethernet frames – is transparent to the underlying medium (RPR, SONET/SDH, etc.), while some solutions proposed within the industry were not fully transparent to the underlying medium and therefore could not meet the synchronization requirements in a predictable fashion.

In addition, Nortel is currently testing and verifying wireless backhaul over its PBT technology against a set of synchronization requirements (e.g., jitter, wander, UMTS air interface stability, and TDM/ATM connectivity). It has been observed that in non-PBT network environments (e.g., a native Ethernet network) those requirements can be difficult to guarantee due to additional packet impairments that are imposed and that cannot be easily controlled. PBT, for instance, allows symmetric paths to be constructed within an Ethernet network, eliminating one of the most challenging issues for proper Layer 2 synchronization distribution, which arises because the latency and bandwidth in asymmetric paths exhibit differences in each direction.

In comparison, testing has shown that the Layer 1 approach, while it works only

## UMTS *continued*

for the Ethernet medium, delivers remarkable performance in terms of jitter and wander. The Layer 1 approach is totally immune to any packet impairments because the timing information is carried at the physical layer within the bit-encoding signal. Test results within ITU SG15/Q13 have shown that Ethernet PHY (physical layer) chipsets are extremely fast at acquiring and locking to a timing reference signal and frequency offsets. It has also been shown that in a simple network configuration composed of a few Ethernet switches, the Layer 1 approach will produce two to three orders of magnitude better wander performance compared to some of the best-known Layer 2 solutions. For instance, Layer 1 will produce wander in the order of a few to tens of nanoseconds, whereas Layer 2 will produce wander in the order of a few hundreds of nanoseconds up to a few microseconds.

Eventually, it can be expected that both approaches (Layer 2 and Layer 1 synchronization) will co-exist, likely in a very complementary fashion serving different applications and network environments. By adopting these two approaches for UMTS, as well as for other applications, Nortel will be well-positioned to serve the different needs of its various wireline and wireless customers.

***Michel Ouellette** is a System Design Engineer with the Advanced Technology Research group within the CTO Office.*

higher capex and opex costs and reduced service availability.

The alternative to locating all packet control in the L3 IP layer is to retain L2 Ethernet forwarding and enhance it with IP Type of Service (ToS) and Differentiated Services Code Point (DSCP) remarking and queue management, keeping a simple network edge structure and centralizing in the core the more costly functions, such as routing and sophisticated security management.

The key to next-generation Ethernet implementation in the enterprise, therefore, is to restructure the hierarchy by limiting the L2 function in the wiring closet to that of efficiently moving packets to the core. By changing the function of Ethernet switches in the wiring closet from any-to-any switching to many-to-one aggregation, control at an individual flow level is managed in the core switch.

This hierarchical approach will enable networks to have the best of both worlds: the simplicity of L2 Ethernet at the edge, with individual L3 IP device and flow management at the core. The Nortel Secure Network Access (SNA) solution is an example of simplifying security intelligence with a centralized model. The Nortel SNA solution is tightly integrated with the converged IP network infrastructure and supports authentication, compliance testing, and policy-driven user controls across the complete range of user devices. (For more on Nortel SNA, see Issue 3 of the *Nortel Technical Journal*, page 31.)

Furthermore, L2 Ethernet can be employed to interconnect the core routers typically used in a large co-located campus or a campus distributed across a metropolitan area. In these campus networks, technologies such as dense wavelength division multiplexing (DWDM) are primarily intended for point-to-point interconnection, while technologies such as Resilient Packet Ring (RPR) employ an L2 fabric to efficiently interconnect routers via a ring configuration with high network availability.

This concept of using intelligent L2 Ethernet fabrics to connect a number of L3 core devices enables much simpler configuration, while retaining the dynamic bandwidth and path management that is missing in point-to-point schemes. In fact, Nortel is using these capabilities in its own enterprise network architectures to limit complexity and create a new paradigm in network simplicity and manageability. In Nortel's network, employing RPR has led to a 70-percent reduction in cost, while bandwidth has increased fourfold.

With this architecture, the next-generation enterprise campus will centralize traffic management and switching but, more importantly, offer a controllable, truly cost-effective solution.

With ongoing work across the broad range of required technologies, Nortel is positioned to lead the transformation that will enable its customers and their customers to create dramatically more manageable information and communications environments. In enterprise campuses, this transition has begun with the movement to Ethernet switching and Gigabit Ethernet switching and now needs to add the management capability to protect against increasingly sophisticated security threats.

The key is using the simplicity and networking efficiency of Ethernet technology together with the robustness of optical technology to deliver manageable, scalable, and cost-effective networks. With strong positions across both these technologies and its Ethernet and Metro Ethernet portfolios, Nortel is delivering on the promise to bring next-generation Ethernet to both enterprise and service provider customers. ■

***Phil Edholm** is CTO, Enterprise Networks.*
***Paul Littlewood** is Leader, Carrier Architecture, Metro Ethernet Networks.*

# Taking control: The evolving role of the control and data planes

by Dave Allan and Nigel Bragg

With new innovations in Ethernet that promise to significantly enhance both its scalability and OAM capabilities, a path is being opened for carriers to converge their existing circuit-based infrastructures onto a common metro Ethernet solution to simplify operations, save costs, and deliver the bandwidth needed to support new broadband and IPTV service deployments. To take advantage of these new capabilities, Nortel has unveiled a new technology called Provider Backbone Transport (PBT). Designed specifically to meet the needs of carrier environments, PBT will bring the determinism, OAM capabilities, quality of service, traffic engineering, and resiliency attributes that carriers are familiar with in the circuit world to a converged next-generation, standards-based metro Ethernet solution. With PBT, Ethernet can take on the role of a fully functional and fully featured data plane which, when combined with an enhanced Generalized Multiprotocol Label Switching (GMPLS) control plane, will enable the common operational base that carriers are seeking.

The telecommunications industry is on the cusp of a significant network transformation as packet technology supplants the traditional TDM circuit networking world. This transformation has largely been driven by the success of Internet-based technologies in the enterprise and residential broadband space, where the current applications mix is most efficiently and economically supported by a packet infrastructure.

This trend will be reinforced by the impending scaling required for universal broadband, video-on-demand (VoD), and IPTV.

Although the network layer is seeing a transformation, however, the management and operations paradigm has not yet evolved to address the simplification required for performance and reliability in a ubiquitously packet world.

Accordingly, carriers are seeking to simplify their networks by converging on a common operations solution that combines the low cost and networking flexibility of packet processing with the determinism, the operations, administration, and maintenance (OAM) capabilities, and the operational attributes (such as resiliency, quality of service, and scalability) to which carriers are accustomed in their circuit infrastructures (see article on page 7).

> The combination of emerging Ethernet standards, the near universal view that Ethernet will be the predominant link layer of the future, and innovations in Ethernet Layer 2 networking capabilities position Ethernet as a converged and fully functional carrier infrastructure for the future.

Another generation of innovation in transport, therefore, is required to align the industry and fully leverage Moore's law. This article describes a significant step forward in that direction.

To date, the packet network has existed side-by-side with phone and private-line traffic as an application overlay on the SONET/SDH network, with a comfortable division of functionality between the circuit world and the packet overlay. Going forward, however, the SONET/SDH network will become increasingly absent as carriers de-layer their networks onto a purely packet infrastructure over copper, fiber, and dense wavelength division multiplexing (DWDM).

This de-layering will require a rationalization of data plane functions that used to be distributed across Layer 0/Layer 1/Layer 2 (L0/L1/L2) – the physical and datalink layers – and Layer 3 (L3) – the IP layer – as well as the creation of a common control plane that will tie these functions together as a single system. (The data plane comprises network elements and the links between them that carry customer traffic; the control plane includes routing, signaling, link management, and other network protocols that are used to set up, maintain, and terminate data plane connections.)

The important questions are: "How far can the network infrastructure be de-layered and simplified without losing the expected circuit-based attributes?" and "Is there an emerging single layer that if combined with a state-of-the-art control plane meets carrier requirements?"

While conventional wisdom suggests there is no single layer ready to step into this role, the combination of emerging Ethernet standards, a near universal view that Ethernet will be the predominant link layer of the future, and innovations in Ethernet L2 networking capabilities is changing that view and positioning Ethernet as a converged and fully functional carrier infrastructure for the future.

### Network evolution

While there have been numerous trends in circuit and packet networking technology over the years, and frequent cross-pollination of capabilities as layers and sub-layers absorb functionality from their neighbors as they attempt to displace them, the essential starting points are as follows:

- Circuit devices bundled a lot of service and resource attributes into a "connection" and had a fully self-contained data plane that supported such attributes as resiliency and OAM capabilities. Only recently has a control plane been added to the circuit world, since the complexity of using a control plane in a circuit world outweighed its value, due to the relative inflexibility and long holding times of transport connections.
- The packet data world, on the other hand, typically has had only limited data plane OAM capability and has relied primarily on the control plane for resiliency. However, the relative flexibility of packet communication has meant that resiliency mechanisms that maintain a level of service – albeit a degraded one – under fault conditions have become an option.

Indeed, a succession of packet and cell technologies, such as Multiprotocol Label Switching (MPLS), have already incorporated a control plane as part of their architectures. In the data world, the control plane is implemented in-band with the bearer path and, as such, has accreted a number of data plane OAM functions, such as proxy fault detection, fault propagation, and alarm management. The MPLS control plane inherited this paradigm, and its data plane was accordingly simplified.

When experience demonstrated that the MPLS control plane could not proxy all aspects of operations, its data plane was retrofitted with increased OAM functionality to take on new roles. When not coupled to deterministic data plane behavior and operational practices, however, inconsistent duplication of functionality between the MPLS data plane and control plane becomes simply a significant source of operational complexity.

To handle the new roles, the MPLS control plane has been augmented with several control protocols, all of which have reached a level of complexity that has become a significant source of network unreliability. Indeed, the MPLS control plane may no longer be tasked with the first level of fault management, and loss of control connectivity with a peer is assumed to be a "crash" first, and a data plane problem only if the peer is not observed to restart.

> When looking at the requisite attributes for a packet infrastructure, the increasingly ubiquitous Ethernet link layer is rapidly being equipped with well-specified OAM and Layer 2 networking functionalities that enable it to assume a larger role beyond that of a simple router interconnect.

In one approach to rectify the control plane dependency in MPLS, MPLS data plane state and fault management has been entirely decoupled from the control plane. This decoupling has frequently made synchronization of both data plane and control plane OAM functionalities an intractable problem. Faults detected via the overlaying of instrumentation on the data plane cannot be correlated with control plane indications, leading to multiple sources of alarms.

The latest evolution of MPLS, Generalized MPLS (GMPLS), has added new functionalities associated with scaling the operations of large networks and is rapidly becoming a universal solution for adding a control plane for path maintenance to circuit technologies that previously did not have one. The GMPLS control plane, for example, has been applied to such transport technologies as SONET/SDH that previously did not use a control plane and, over many years, had accreted functionality associated with scaling the operations of large networks. In this way, GMPLS has been augmented with such functionalities as bidirectional connections and the ability to signal protection attributes that, although not required by MPLS, are useful in the overall scheme of things – in particular, features that acknowledged requirements for the scaling of transport networks.

In Nortel's view, when a fully functional Ethernet data plane is combined with the GMPLS control protocol, configuration of the carrier infrastructure can be fully automated, the rich set of resiliency options already available extended, and the evolutionary path for an Ethernet infrastructure continued – advancing from "turn it on and hope for the best" to a comprehensive carrier infrastructure that marries the best of the packet and circuit worlds. [For more on applying the GMPLS control plane to Ethernet, see *GMPLS Control of Ethernet PBT switches* (www.ietf.org/internet-drafts/draft-fedyk-gmpls-ethernet-pbt-00.txt) originally published in the Internet Engineering Task Force (IETF) standards organization.]

### The Ethernet solution

When looking at the requisite attributes for a packet infrastructure,

*continued on page 29...*

## Figure 1. Control plane evolution

**Diagram A. Transport world**

Pre-planned resiliency against all failures

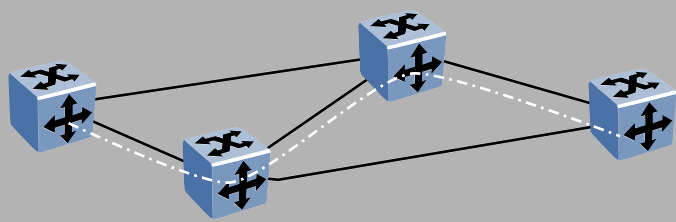

— Physical link  -·- Data plane OAM

**Diagram B. Data world**

Dynamic resiliency against limited set of failures

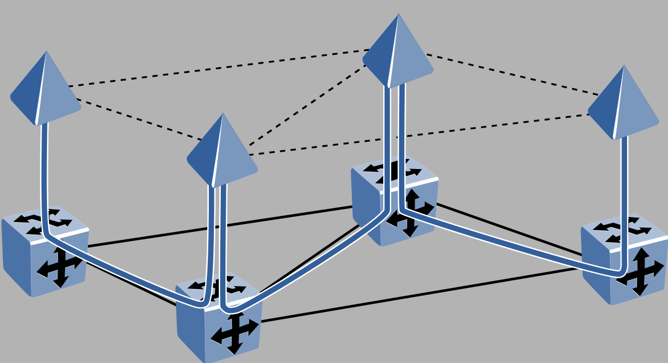

— Physical link  — Proxy OAM path

--- Control plane connectivity

**Diagram C. Evolved world**

Full resiliency with independent data and control planes

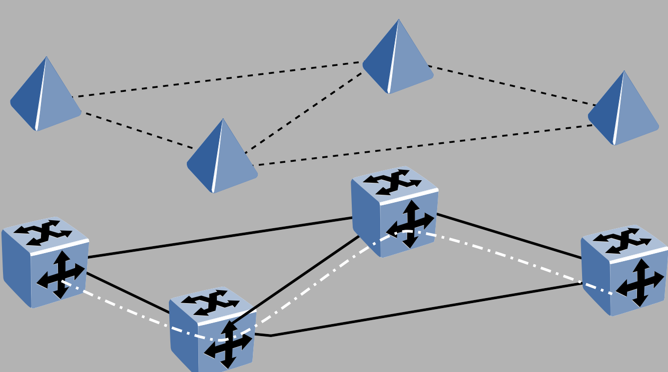

— Physical link  -·- Data plane OAM

--- Control plane connectivity/OAM

Switch  Control functions located in switch

These diagrams show how the control plane is evolving as best-in-class capabilities are added from both the transport and data worlds.

In the traditional transport world (Diagram A), switches are connected via physical links (copper, fiber, RF), and all executive actions (i.e., decisions with respect to fault management that can be made locally) are delegated to the data plane. Only a limited ability to respond to faults exists because recovery actions are pre-planned and cannot accommodate multiple failure scenarios.

In the data world (Diagram B), control functions (routing, signaling, link management, and other network protocols used to set up, maintain, and terminate the data plane connections) are contained within the switch, and the control plane acts as proxy for the data plane operations, administration, and maintenance (OAM), handling all the executive actions. The control plane functions are able to communicate with each other to manage traffic flows over the data plane. However, the control plane has over time accreted so much functionality that it has become overly complex and a major source of instability in the network.

In the evolved world (Diagram C), the data plane and control plane are decoupled, with clearly defined boundaries between the executive actions for fault management that are delegated to the data plane and to the control plane. As a result, data plane survivability is completely independent of the control plane. With this decoupling, the data plane can offer a quick response to any single failure while the control plane can be used to address complex or multiple failures.

# Ethernet OAM and resiliency: Making Ethernet suitable for carrier operations

by Raymond Allard

Critical to any carrier network infrastructure, operations, administration, and maintenance (OAM) and resiliency capabilities are fundamental tools required to run large, reliable networks. Accordingly, Nortel is working with such standards bodies as the Institute of Electrical and Electronics Engineers (IEEE) and the International Telecommunication Union (ITU-T) to not only provide Ethernet with OAM and resiliency capabilities similar to what the TDM network delivers today, but also to capitalize on Ethernet's own inherent advantages.

The TDM capabilities that Ethernet will need to duplicate include, for example, the ability to protect traffic within 50 milliseconds (ms) of the occurrence of a failure, and the ability for operators to manage the flow of traffic with comprehensive performance-monitoring counters and such connectivity verification functions as path trace and link trace.

To these familiar capabilities, Ethernet adds key functions of its own, such as the ability to prioritize different classes of traffic, operate in a point-to-point or multipoint topology, and provide far more granular management of traffic – unlike TDM, which treats all traffic with the same priority, (i.e., higher-priority traffic cannot preempt lower-priority traffic). Without the multiplexing and aggregation capabilities of Ethernet, TDM transport pipes are often underutilized, which leads to inefficient management of bandwidth and results in higher costs. With the added capabilities it provides, an Ethernet solution promises to not only reduce carriers' operations costs but enable efficient delivery of services.

Introducing new Ethernet functions, however, will require a strong OAM offering beyond the current OAM set. To this end, Nortel is actively leading Ethernet OAM standards activities. For instance, Nortel is co-chairing the IEEE 802.1ag committee that is defining those standards. Nortel is also implementing these new Ethernet OAM functions across its Metro Ethernet portfolio, including the Metro Ethernet Routing Switch 8600 (MERS 8600) and Optical Multiservice Edge 6500 (OME 6500) products.

A substantial part of the attraction of an Ethernet-based infrastructure is its relative simplicity and low cost compared to such alternative technologies as Multiprotocol Label Switching (MPLS) that have attempted to provide a carrier-grade operations environment, but often result in a complex design and therefore much higher capex and opex costs. MPLS, for example, supports Label Switched Path (LSP) pings to verify connectivity to other network elements. Such packets must be handled by software, however, which means that they either are supported only at a low rate or require a large processor infrastructure. Similarly, Fast Re-Routes (FRRs) – a tunnel protection mechanism in MPLS – require the use of a control plane in order to provide resiliency. Use of a control plane involves more complex protocols and requires more memory and processors, especially in large metro deployments where the control plane needs to keep track of a large number of states on other network nodes and where, as a result, scalability becomes an issue.

To endow Ethernet with a world-class OAM suite, its mechanics must be simple, supportable in hardware, standards-based, and ubiquitous in the network. Beyond facilitating carrier operations, another goal is to enable network resiliency without the implied use of a control plane.

## Standards activities

The definition of these Ethernet OAM capabilities is being driven in two main standards bodies: IEEE 802.1ag Ethernet Connectivity Fault Management (CFM), and ITU-T Y17.ethoam.

IEEE 802.1ag defines the following OAM packet types:

- continuity check messages (CCMs), which are unidirectional Ethernet packets that allow the user or network to determine connectivity to other network elements' ports;
- loopback messages (LBMs), which are bidirectional Ethernet packets that allow the user or network to provide a ping-like function where the terminating node bounces the packet back to its originator; and
- linktrace messages (LTMs), which are bidirectional Ethernet packets that allow a user to determine which path the traffic is taking through the network. An LTM issued by an originating node will receive a reply from each of the nodes along the path up to the terminating node. The reply contains the network address, which is used for determining the proper network topology and for path management.

All of these OAM messages follow the same data path as the user traffic, and by appropriately setting the Ethernet priority bits (P-bits), they can also be used to provide coverage within a particular class of service (CoS). By processing these packets in hardware, it is possible to support very high rates, which makes these functions suitable for troubleshooting when there is a partial loss of packets due to network failure or congestion.

The standards also allow for vendor-specific extensions to provide added value. Nortel, for example, is working to provide a means for measuring latency and jitter through the network to improve performance.

A companion standard of IEEE 802.1ag, IEEE 802.1ah is defining Provider Backbone Bridging (PBB), often called MAC-in-MAC. PBB is designed to cleanly separate the end-customer network from the carrier infrastructure, thereby enabling service scaling and secure customer and content separation. Provider Backbone Transport (PBT), a Nortel innovation created in conjunction with a major customer, leverages PBB but also allows paths between service end-points to be explicitly engineered across a network. Defined specifically to suit carrier infrastructure operations, PBT is a highly scalable technique for configured operation of Ethernet switches. PBB and PBT are introducing a hierarchy of services and tunnels within Ethernet that does not exist in a TDM network. This hierarchy is an important aspect of scaling Ethernet for carrier metro operations (for more on PBB and PBT, see the main article).

However, IEEE 802.1ag is not specific to PBB/PBT and can therefore also be operated in native Ethernet, Q, and QiQ networks [IEEE 802.1q (Virtual Bridged LAN) and 802.1ad (Provider Bridges)].

The second set of standards activities – ITU-T Y17.ethoam – introduces other familiar concepts, such as the ability to issue an alarm indication signal (AIS) on a particular tunnel or service if a local failure has occurred. The AIS can then be used to trigger protection immediately.

### Achieving network resiliency via Ethernet

A multitude of resiliency schemes already exist in the network today, including link protection schemes (e.g., Ethernet Link Aggregation – IEEE 802.3ad), equipment resiliency (e.g., SONET/SDH 1:1 protection), and path resiliency (e.g., MPLS FRR, and SONET/SDH Unidirectional Path Switched Ring).

With its new Ethernet hierarchy of services and tunnels, IEEE 802.1ag is adding the ability to peer with other Ethernet bridges over eight distinct logical levels – allowing peering within the link, tunnel, service, and customer frame domains. By assigning one of these levels to a tunnel, continuity check messages (CCMs) can now be used to trigger tunnel protection and achieve excellent protection times (e.g., 50 ms).

Such CCMs must obviously be issued at a high rate (approximately every 10 ms) in order to provide a suitable response time. In a multipoint topology, the different levels can be used to segment individual branches or areas of the network. Since high-capacity ports, such as 10-gigabit Ethernet, may support a fairly large number of Ethernet tunnels, the protocol needs to be supported mainly in hardware. Hardware implementation also provides other benefits, such as the capability to continue operating without interruption during software resets and system upgrades.

As discussed, Ethernet is acquiring a world-class carrier-grade infrastructure in terms of OAM and path resiliency. Moreover, such capabilities are being introduced without compromising the basic simplicity and low cost of Ethernet, while delivering a value proposition that rivals today's TDM network.

The challenge ahead is to make this new infrastructure common throughout the network – on equipment located both at the customer premises and in the metro core and access network. Nevertheless, there is no doubt that this new OAM capability is a key and makes Ethernet an attractive infrastructure play within the carrier market.

***Raymond Allard** is Senior System Architect, Packet Switching, Optical Networks.*

the increasingly ubiquitous Ethernet link layer is rapidly being equipped with well-specified OAM and Layer 2 networking functionalities that enable it to assume a larger role beyond that of a simple router interconnect.

Recognizing this trend, Nortel has developed a new technology called Provider Backbone Transport (PBT). Nortel is working with some of the world's largest service providers and standards bodies to facilitate broad adoption of PBT as a metro technology. Nortel, for example, has introduced PBT to the IETF, Institute of Electrical and Electronics Engineers (IEEE), International Telecommunication Union (ITU-T), DSL Forum (DSLF), and TeleManagement Forum (TMF).

PBT transforms Ethernet technology – traditionally restricted to small-scale local networks – into a more reliable, scalable, and deterministic technology, making it suitable as the basis for fixed and mobile carrier networks to deliver video and broadcast, multimedia, broadband data, and voice services.

Defined specifically to suit carrier infrastructure operations, PBT is a simple point-to-point tunneling technology that enables service providers to specify the path that an Ethernet service should take across the network. PBT allows for quality-of-service (QoS) guarantees by reserving bandwidth for real-time services, and provides for 50-millisecond service recovery times should a connection fail – matching the benchmarks set by today's existing SONET/SDH optical transport standards.

PBT is now available on Nortel's Metro Ethernet Routing Switch (MERS) 8600, with development under way to incorporate PBT into the Optical Multiservice Edge 6500 (OME) 6500 and other Nortel Ethernet-capable platforms.

With the increased point-to-point scalability that PBT brings, Ethernet can take on a substantially larger role than conventional wisdom suggests, and the architectural boundaries and partitioning of functionality between the control plane and data plane can be redrawn to achieve a superior result.

This redrawing can be successful only if we are completely clear about the scope of the task and the criteria for success: *we are asking Ethernet to take on alone the transport role that until now was performed by a combination of SONET/SDH and MPLS with its IP control plane*. It is precisely this expanded role that must be implemented, however, to achieve a comprehensive de-layering of the network along with increased network agility and automation, while sacrificing none of the circuit-based infrastructure OAM and resiliency values that carriers have come to both expect and depend upon.

While this may sound like a major challenge, the "new" Ethernet has been or is being equipped with distinctly different attributes – such as strict hierarchy, absolute addressing, and data plane OAM – many of which are directly analogous to SONET/SDH capabilities.

In addition, IP-derived control plane components are available to provide the required provisioning functionality. The challenge, then, is actually to configure the GMPLS control plane attributes and parameters so that they fully exploit the Ethernet data plane capabilities, not ignore them or duplicate them unnecessarily.

To meet this challenge, PBT offers a number of attributes that are important from the control perspective:

• The only infrastructure connectivity type is point-to-point. This attribute, along with admission control and edge policing, gives the carrier complete control of traffic flow in the network and the degree (if any) to which different customers' traffic interacts. This control is becoming increasingly important as carriers focus more on traffic-engineered applications governed by service level agreements (SLAs) and on the concurrent support of legacy circuit applications over packet networks.

• PBT uses IEEE 802.1ah (Provider Backbone Bridging) encapsulation, where there is mapping between the client layer and the server layer at the boundary of the transport network domain. This mapping isolates customer address spaces from the carrier address space. This structure provides complete transparency of the network to customers and complete control of the domain to the carrier, which offers numerous operational benefits, enhances scalability, and secures the network.

• PBT uses absolute destination-based forwarding in the core [Destination Address + VLAN Identifier (DA + VID)]. The primary benefit is scaling: as the network expands, the core forwarding states grow in linear proportion to the number of destination addresses alone, rather than in proportion to the number of source/destination address relationships, which has a square law relationship to the number of nodes. An important secondary benefit is that the globally unique identifiers in PBB enable the immediate self-identification of data plane configuration problems (i.e., misdirected traffic self-identifies) and separation of services (important for security reasons), unlike label-swapping techniques where a mis-swapped or mis-forwarded label may seem to be perfectly valid at the next switch to which it has been inadvertently sent.

• Although PBT employs destination-based forwarding in the core, all PBT packets also carry their source address, preserving end-to-end visibility of individual connections, and these connections can be instrumented by very simple data plane OAM heartbeats (verification messages) and performance monitoring (PM) flows. By contrast, label-swapping technologies that use label merging mean that when a packet arrives at its destination, any indication of where it came from has been lost.

• The combination of the PBT data plane OAM and alarm management permits fault correlation and performance management to be pushed back into the network, dramatically reducing the telemetry and associated retrieval, storage, and processing required in the network to provide a comprehensive view of operations.

From the perspective of a control plane, PBT has one overwhelmingly important property: point-to-point connections built using absolute addressing and strict hierarchy can be reliably monitored by simple, continuously "on" data plane OAM, the failure of which can be used to trigger fast protection switching actions. This capability relieves the GMPLS control plane of the need to deliver the practically impossible combination of properties expected of incarnations of the classical IP control plane, namely:

• achieving good use of resources under the usually fault-free regime in which today's networks are generally operated; and

• at the same time, being capable of very rapid restoration to mitigate the effects of the occasional faults that arise.

By contrast, when used with PBT, the GMPLS control plane can be viewed as primarily a provisioning application for the setup of Ethernet switched paths. The control plane uses centralized route-optimizing calculations, undertaken with comparatively relaxed performance requirements, to compute and install both working and protection paths.

**The "new" Ethernet has been or is being equipped with distinctly different attributes – such as strict hierarchy, absolute addressing, and data plane OAM – many of which are directly analogous to SONET/SDH capabilities.**

The supporting OAM functions offer complete fault coverage, rather than simply responding to link cuts. The actual fault detection and switching are delegated to the connection endpoints, which enables a quick response. What's more, because PBT uses 1:1 or head-end switching to conserve protection route data path resources until the route is required, it is also efficient.

With PBT, autonomous operation of the GMPLS control plane for restoration is now used only for less-critical services that can accept longer recovery times (seconds to tens of seconds or even longer) and do not warrant protection, and for recovery from multiple faults that have resulted in simultaneous failure of some diverse-routed paths (which are presumed to be rare).

This arrangement also means that the network does not need to have all affected endpoints contend immediately, in an unsynchronized fashion, for the remaining network resource base, which reduces race conditions and retries. To elaborate on this point, every successful restoration of a path in the network modifies the resource database that the other impacted endpoints depend upon for their computation of the new path. When the resource database is undergoing churn, many path computations are made using out-of-date or soon-to-be-out-of-date information. When the backup path can sustain the SLA, the reconvergence of the network on a protected state can be performed in a more leisurely fashion by having nodes randomly defer path setup so that resource contention is minimized.

A fully instrumented Ethernet data plane also relieves the GMPLS control plane of having to proxy for the data plane for fault detection. While the control plane can fate-share with the data plane, it does not need to. More importantly, the data plane can continue to operate according to the last-provisioned state in the temporary absence of the control plane, while autonomously continuing to offer resilient connectivity. This capability completely decouples data plane availability and survivability from control plane outages.

## The redefined network

In a collapsed and converged infrastructure, the functional lines are redrawn: the Ethernet data plane assumes a more consistent and complete role, so that it can offer service and the first level of resiliency against failure, independent of the presence of a control plane; and the GMPLS control plane is operationally harmonized with the required data plane capability.

The control plane provides multivendor interoperability of provisioning, a higher degree of consistency in provisioning and configuration because the control and data paths are congruent, and additional restoration options, especially in the presence of multiple failures.

> Extending the transport paradigm to Ethernet with PBT and a control plane positions Ethernet as the networking technology against which other metro solutions will be measured.

To complement the emerging PBT Ethernet data plane attributes to produce a full-capability packet infrastructure solution, the GMPLS control plane already has a number of useful attributes, including:

- the ability to associate simple connections in order to provide protected paths; and
- mechanisms to recover state from its peers after an internal failure, such that the control plane can re-boot without interrupting data plane operation, a capability known as "graceful restart."

What GMPLS needs, however, is the ability to distribute Ethernet "labels" and populate Ethernet forwarding tables, as well as the associated path-computation mechanisms to compute and install diverse-routed working and protection paths in a near-optimal configuration from a traffic engineering perspective.

Implementations of GMPLS will also need the scalability to complement the linear data path scaling properties of PBT. Today's switches running PBT can comfortably support tens of thousands of network elements, corresponding to full meshes of 50 million point-to-point connections in the extreme case, with likely deployments in the hundreds of thousands to millions of connections.

With these additions, a basic scalable and instrumented point-to-point infrastructure can be achieved, and GMPLS and PBT then used to automate the creation and installation of connections as building blocks for connectivity services, such as private line, and for backhaul to higher-layer service nodes.

## Services construction

To construct these connectivity services, there are two aspects to consider: the first, supporting the requisite adaptations; the second, supporting the required connectivity.

Through the use of MPLS, adaptations already exist for Layer 3 VPN and legacy Pseudo Wire (PW) services, and the existing service infrastructure can be mapped directly onto Ethernet in the role of a packet-switched network and lower layers. Through the use of IEEE 802.1ah MAC-in-MAC, an Ethernet-optimized solution also emerges.

To support the required connectivity, services construction is focused on three connectivity primitives:

- scaling large numbers of point-to-point service instances, such as Internet access and router interconnect for businesses;
- scaling large numbers of sparse (small) multipoint-to-multipoint communities

of interest, serving them with both unicast and broadcast/multicast, for such services as business E-LAN (L2 any-to-any services); and

• providing a small number of relatively dense source-specific multicast trees, for such applications as IPTV distribution.

The first primitive – point-to-point connectivity – is directly supported by PBT together with the GMPLS control plane as described earlier.

Building on this capability, the second primitive – multipoint-to-multipoint connectivity – can be implemented with simple hub-and-spoke topology overlays and the use of Ethernet VLANs as switch partitioning mechanisms. Many service instances can share a common, simplified topology. In addition, data plane protection switching provides resiliency options while permitting a loop-free VPN topology to be maintained.

The third primitive – dense multicast point-to-multipoint connectivity – is Ethernet's original forte, and can operate concurrently with PBT operation under the common control plane (a simple VID configuration vs. VID/MAC).

With the availability of these solutions, all three required connectivity configurations for service delivery can be condensed into a common infrastructure with a common control plane to deliver unified resource management and operational simplicity. The application of control planes to traditional transport networks offers increased automation, as well as the pre-integration of multivendor solutions.

Ethernet has added two key functionalities that suit it to assume the role of the transport network: the first is OAM capability in the same class of completeness that characterizes the existing transport infrastructure; and the second is the ability to have forwarding configured directly by a control plane, which enables new connectivity capabilities not normally associated with Ethernet operation.

These advances have been achieved with minimal changes to Ethernet as specified, making it possible to leverage Ethernet's high degree of specification and commoditization for low-cost, efficient operations. Extending the transport paradigm to Ethernet with PBT and a control plane positions Ethernet as the networking technology against which other metro solutions will be measured. With its leadership in such solutions as PBT, Nortel has been instrumental in bringing these innovations both to standards bodies and the marketplace. ■

***Dave Allan*** *is Senior Advisor for Strategy and Architecture in the CTO Office.*
***Nigel Bragg*** *is a Consultant in the EMEA CTO organization.*

# The importance of policy-based resource control in future networks

by Nalin Mistry

**Tomorrow's converged networks need to be highly adaptable to support a wide variety of service needs over a common IP-based network infrastructure. Key to meeting these service needs is policy-based resource control, which provides mediation between applications and the underlying network layer to intelligently manage network resources – dynamically and in real time. For operators, policy control is critical for delivering a wide variety of high-value services with guaranteed quality of service across fixed, wireless, and cable access technologies. Standards bodies are currently working to define policy control for Next-Generation Networks (NGNs). Nortel – building on its strengths in real-time networking and leadership in deploying Voice over IP networks – is at the forefront of shaping the policy-based resource control functionality needed for tomorrow's networks.**

As network operators move to transform their existing infrastructures to a converged IP-based network, they need to focus on more than putting in place intelligent, high-capacity, reliable "speeds and feeds." They also need to be able to offer guaranteed quality of service (QoS) to their end users in accordance with service level agreements (SLAs), as well as to control and properly charge for existing and new services.

Achieving QoS in next-generation networks differs dramatically from how QoS is achieved in traditional networks. Specifically, today's public switched telephone network (PSTN) and most public land mobile networks are largely based on circuit-switched technology, whereby a bearer channel (for user traffic) is dedicated for the duration of a session. By dedicating an entire channel, the network can guarantee a given level of bandwidth for each session.

This approach of allocating resources has been the case for many years, although it is notable that there have been a number of innovations in the signaling and control plane [e.g., Signaling System No. 7 (SS7), Intelligent Networks (INs), and media gateway controller-/media gateway-type architectures] that made it possible for call signaling and bearer traffic to traverse different paths and even different nodes in the network.

Traditional IP networks, by contrast, typically provide service on a "best-effort" basis, whereby the bandwidth obtained at a given time for a given session is not guaranteed. While this best-effort approach has worked thus far for Internet-delivered multimedia traffic (such as web browsing, email, and peer-to-peer applications), it is clear that going forward the variability of best-effort IP service will hamper the delivery of many new applications with real-time characteristics. These applications include Voice over IP (VoIP) and video, as well as applications that blend several features of different applications.

Furthermore, as these applications evolve from being PC-based to being targeted at consumer electronic devices (TVs, mobile phones, and PDAs, for instance) and marketed as services to a non-technical mass market, then the network itself will need to evolve to provide a framework in which application and content providers can obtain appropriate network resources for their services, in accordance with SLAs.

Addressing this challenge is the job of an emerging new network capability called policy-based resource control.

Essentially, policy-based resource control is concerned with controlling the admission of bearer flows derived from services/applications into the network and then monitoring network resources once admitted.

This article focuses on one critical aspect of policy control in next-generation networks: that of the policy control engine required to manage the network resources and interact with the underlying network infrastructure. Policies dealing with specific user-level requirements [such as authentication, authorization, and accounting (AAA), security, location, and presence] as well as application-level requirements are managed elsewhere in the network and are not the focus of this article.

Policy-based resource control is the key function that communicates between the underlying network infrastructure and the applications that ride on top. By linking the applications' needs with the connectivity network, policy control provides the critical intelligence to translate the needs of various services and applications into the language the connectivity infrastructure will understand (e.g., queue depth, congestion, and packet drop rates). In this way, policy control acts as the "glue" between applications and the underlying network infrastructure, and is vital to enabling next-generation networks to admit and manage sessions,

*continued on page 36...*

## The relationship between NGN and IMS architectures

Next-Generation Networks (NGNs) and IP Multimedia Subsystem (IMS) are terms often used loosely in the industry, leading to confusion and difficulty in describing the impact of NGNs and IMS on underlying transport/network technologies.

For this reason, it is important to understand how NGNs relate to IMS, and where NGN standards are being defined.

As Diagram A shows, a number of different standards bodies are involved in the standardization of NGNs – organizations that all adopted the work begun during the late 1990s in the 3G mobile community (shown in grey) and in the Internet Engineering Task Force (IETF, shown in light grey) to define a new multimedia services delivery framework that is today known as simply IMS (see article on page 42). Specifically:

• the Third Generation Partnership Project (3GPP) and 3GPP2 began to define, respectively, the IMS core architecture (UMTS domain) and the Multimedia Domain architecture (CDMA domain);

• in close collaboration with 3GPP, the Open Mobile Alliance Ltd. (OMA) focused on defining wireless services and third-party applications for the new architectural environment;

• the GSM Association concentrated on defining requirements for global wireless operators; and

• the IETF focused its attention on standardizing IP interfaces, including Session Initiation Protocol (SIP), Diameter [an authentication, authorization, and accounting (AAA) protocol] and the Common Open Policy Service (COPS) protocol for supporting policy control over quality of service (QoS) signaling protocols.

By the early 2000s, several wireline-focused standards organizations had adopted the core IMS architecture as the basis for their NGN frameworks and began extending the IMS specifications for interworking across all types of access technologies (wireline, wireless, and cable). These organizations (shown in light blue) include:

• the European Telecommunications Standards Institute (ETSI), through its Technical Committee TISPAN (which stands for Telecoms and Internet converged Services and Protocols for Advanced Networks). ETSI TISPAN, following significant pressure from a number of European PTTs, was the first

**Diagram A. IMS/NGN standards landscape**

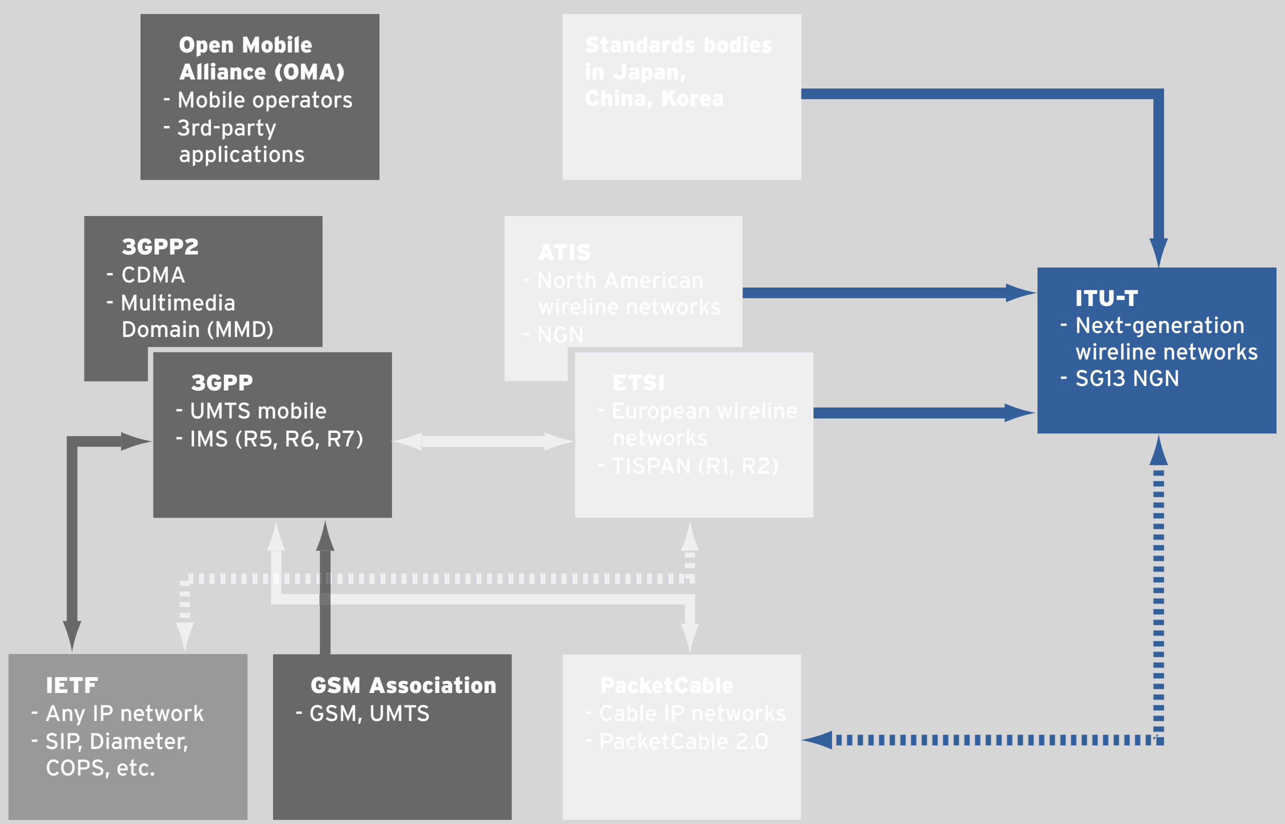

standards organization to agree to a Release-based approach to developing NGN specifications. TISPAN has continued to maintain its lead in many areas of NGN development, including policy. In January 2006, TISPAN published its Release 1 specifications. ETSI is working closely with 3GPP as that body evolves the IMS standards. In particular, TISPAN and 3GPP have agreed on the concept of "one IMS," meaning that any changes that TISPAN makes to IMS that are also deemed by 3GPP to be useful for 3G operators will be incorporated directly into IMS specifications. Alignment between the two organizations is driven by a combination of liaison statements and contributions from member companies. TISPAN also maintains liaison relationships with many other standards development organizations and forums (e.g., ITU-T, ATIS, DSL Forum, and OMA).

• the Alliance for Telecommunications Industry Solutions (ATIS), which is concentrating on the requirements of NGN networks in North American wireline environments and is working closely with ETSI, particularly in the area of new topics, such as IPTV, for Release 2. ATIS members are contributing North American wireline network requirements (often first developed within ATIS) into the ITU-T;

• the PacketCable project within CableLabs, which defines IP-based standards for the cable television industry. CableLabs has been making contributions on IMS topics in 3GPP and will likely soon contribute its work to ITU-T;

• members of standards organizations in Japan, China, and Korea, who are feeding their requirements directly into the ITU-T; and

• the International Telecommunication Union Telecommunications Standardization Sector (ITU-T). The role of the ITU-T in the development of the NGN is to obtain global consensus on international standards, through a combination of liaison discussions with other regional bodies, such as ETSI TISPAN or aggregations thereof (e.g., regional body partnership projects such as 3GPP), and through the active participation of global companies such as Nortel. Nortel continues to recognize the importance of ITU-T in achieving global standards and indeed in meeting the needs of its potential customers (particularly from China, Japan, and Korea) who are not as active in other standards development organizations,

**Diagram B. TISPAN NGN architecture**

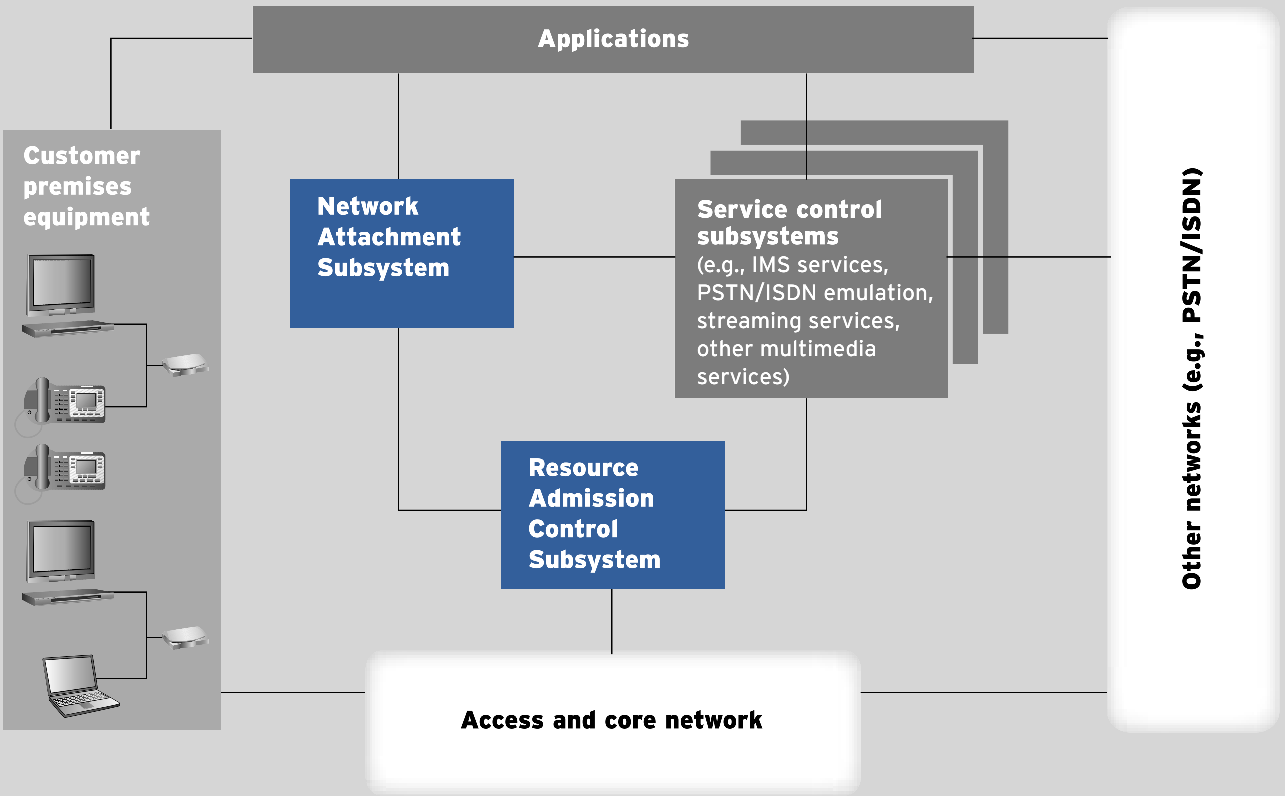

such as TISPAN, ATIS, and 3GPP.

While IMS is at the heart of all emerging NGN standards, it is just one of a number of service control subsystems.

Diagram B (page 35) illustrates the TISPAN NGN architecture. (The two areas in blue make up the TISPAN policy control architecture, discussed in this article.) In this architecture, the call control components of IMS are considered to be part of the IMS service control subsystem. Examples of other service control subsystems include PSTN emulation and streaming services. Elements in the transport plane, application of policy, and storage of subscriber data are all considered to be common functions in the NGN and not particular to any service control subsystem. Given the importance of IMS within the NGN, TISPAN has been working closely with 3GPP to ensure consistency of direction.

no matter what device is being used or where an individual is located.

Ultimately, the goal of policy-based resource control is to ensure that the different types of traffic – voice, video, and instant messaging as examples – are given appropriate QoS treatments, such as meeting bandwidth and latency requirements. Policy control is also key to application transparency across existing and emerging access technologies, and therefore to permitting applications to be agnostic of the networking layer.

### Intense standards activity

The area of policy control is the subject of intense standards development activity in the industry today, and various organizations are working to specify resource control functionality as part of Next-Generation Network (NGN) architectures.

Currently in the policy control area, various architectures are being developed by a number of standards bodies (see page 34), including:

- the European Telecommunications Standards Institute (ETSI), specifically the Telecoms and Internet converged Services and Protocols for Advanced Networks (TISPAN) Technical Committee;
- the International Telecommunication Union (ITU-T);
- the Third Generation Partnership Project (3GPP); and
- the PacketCable organization.

Most of these policy control architectures build on the early work of the 3GPP to define a new services delivery framework for the wireless domain – an architecture it termed the IP Multimedia Subsystem (IMS) (see article on page 42).

As a result of these various standards development activities, emerging policy control architectures differ subtly from each other. Industry opinion is quickly gaining momentum, however, and it is anticipated that many aspects of policy control architectures will soon be harmonized.

Nortel is at the forefront of these standards efforts and is applying its expertise in real-time networking both to develop the right policy control capabilities that tomorrow's networks will need and to help drive standards development toward a common approach.

As well, as standards continue to evolve and mature, Nortel continues to build its expertise in the policy area. Already, for example, Nortel supports an open interface to the policy functions in its IMS solution and has performed lab interoperability testing with several third-party policy server platforms.

### ETSI TISPAN

In January 2006, ETSI TISPAN published a set of standards[1] specifying the requirements, architecture, and protocols for Release 1 of the NGN (see page 34 for an overview of the TISPAN NGN architecture).

The TISPAN architecture defines how policy interacts with the connectivity network to enable dynamic policy-based resource control as a tool for delivering the QoS required by applications.

Specifically, the architecture comprises two key subsystems, which have interfaces to the transport network(s). These subsystems are:

- the Resource Admission Control Subsystem (RACS)[2], which provides policy-based resource control to applications as well as addresses mediation and border control capabilities; and
- the Network Attachment Subsystem (NASS)[3], which provides attachment control, such as authentication, authorization, and assignment of IP addresses.

The remainder of this article focuses specifically on the RACS portion of the ETSI TISPAN architecture, primarily because the TISPAN Technical Committee is leading the thinking in many aspects of policy-based resource control and how that control applies to both fixed and mobile environments.

Indeed, the RACS architecture addresses a number of unique requirements that result from the need to provide policy-based resource control in wireline environments. In the ITU-T, many aspects of the TISPAN work have been recognized as part of the development of the ITU-T documents.

### Key requirements driving RACS design

The development of the RACS architecture was driven by several key requirements that differ significantly from past assumptions regarding policy-based resource control architectures developed in other forums. Specifically, the following key requirements were imposed during the design of the RACS.

• Resource control is of paramount importance to the access and aggregation parts of the network – the areas where network providers envisaged that the largest bottlenecks would take place during the Release 1 timeframe. RACS is not restricted, however, to use only between access and core networks; it also provides policy-based resource control capabilities between core networks.

• Resource control must be able to be deployed in a multi-network environ-

**Figure 1. TISPAN RACS architecture (Release 1)**

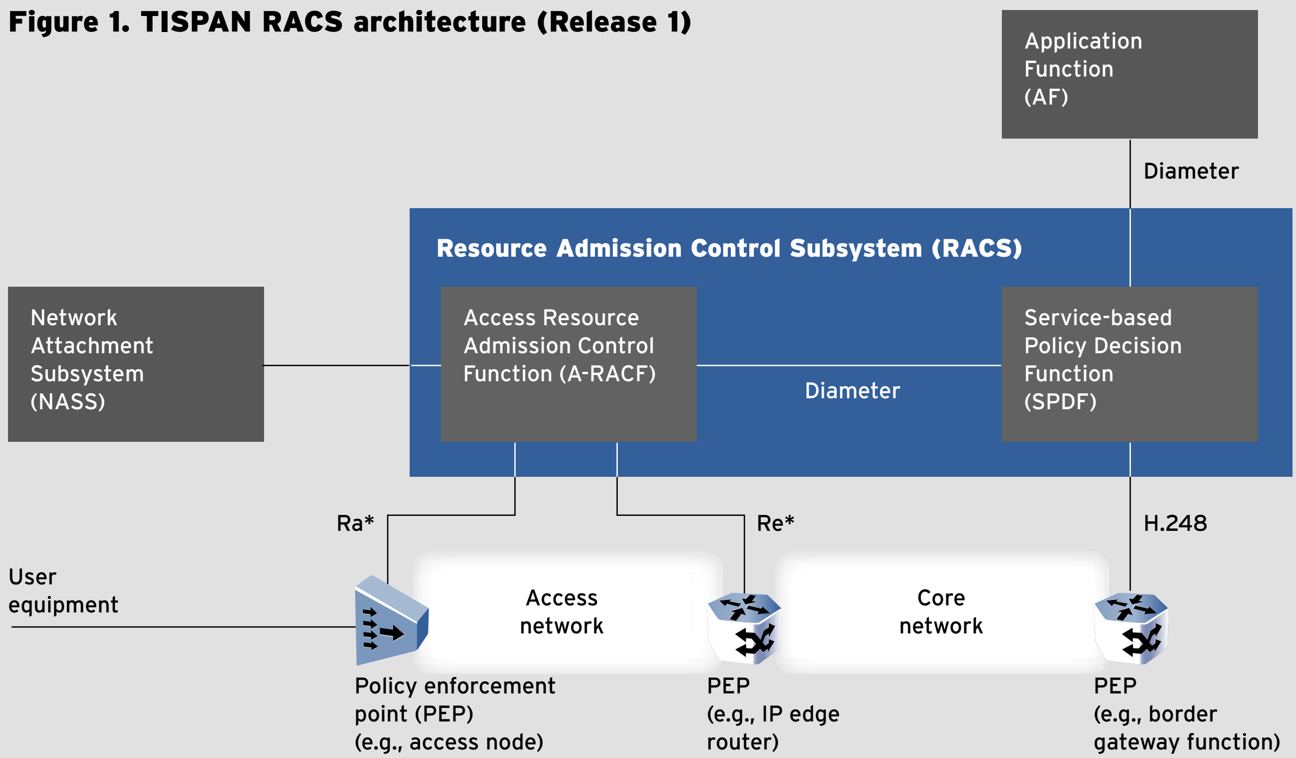


* Ra and Re interfaces are currently undergoing standardization

Shown above is the Resource Admission Control Subsystem (RACS) Release 1 architecture as defined by TISPAN (Telecoms and Internet converged Services and Protocols for Advanced Networks), which is a standardization body of the European Telecommunications Standards Institute (ETSI) that has developed a set of specifications for Release 1 of the Next-Generation Network (NGN).

The two key components of the RACS Release 1 architecture, which are discussed in greater detail in the article, include the:

• Access Resource Admission Control Function (A-RACF); and

• Service-based Policy Decision Function (SPDF).

The A-RACF and the SPDF perform similar functions, such as making bandwidth management decisions, but they do this at different points in the network. An A-RACF would typically be assigned to a specific access network, while an SPDF would interact with one or multiple A-RACFs. The A-RACF and the SPDF interact using the Diameter protocol, which is an AAA (authentication, authorization, and accounting) protocol defined by the IETF.

The SPDF, in turn, interacts with a Policy Enforcement Point (PEP) in the underlying core network via a profile of H.248, the media gateway controller signaling protocol defined jointly by the ITU-T and the IETF. In the core network, this PEP could be, for instance, a border gateway function.

In the access network, the A-RACF interacts with PEPs located in the underlying access (and metro) networks, such as access nodes and IP edge routers, via the Ra and Re reference points. These reference points have not yet been standardized to stage 3 (protocol level) in TISPAN; however, they will likely be addressed in the timeframe of TISPAN NGN Release 2 (due mid-2007).

**Figure 2. Push call model (as applied in a wireline network)**

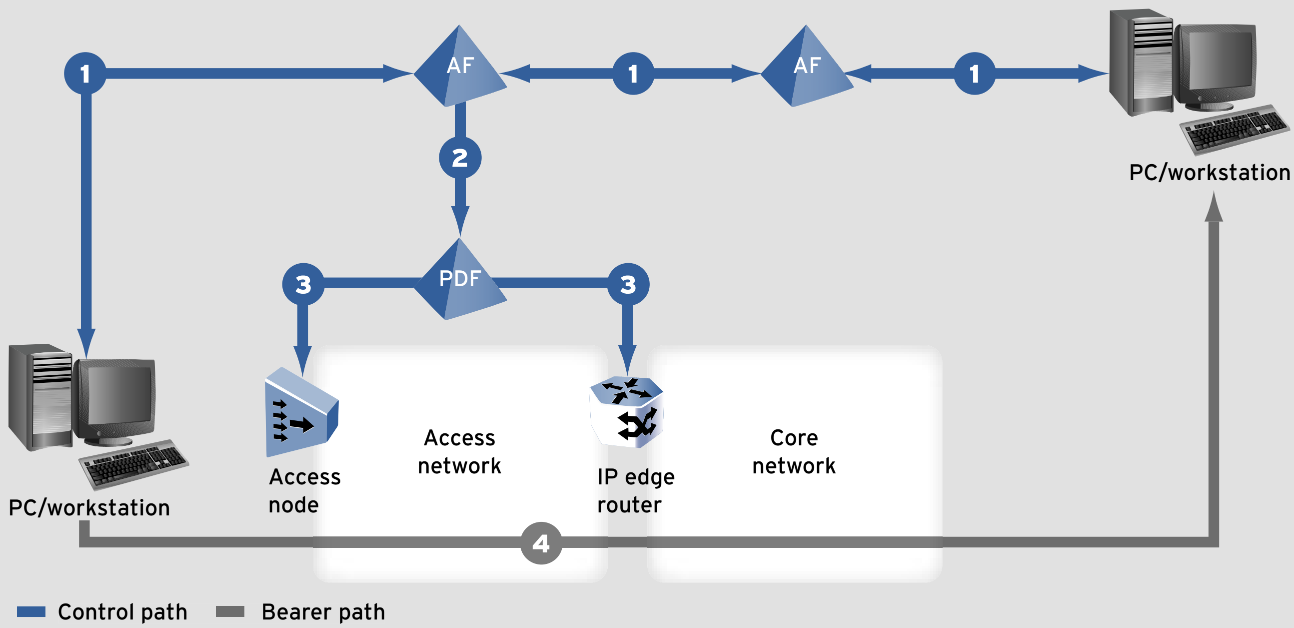


This figure shows a simplified view of the Push call model, as it would be applied in a wireline network. For simplicity, a single network is shown. In reality, however, the call flow would traverse multiple Application Functions (AFs), each of which would typically be associated with a separate network. An operator's network would typically include many AFs. An example of an AF is the Call Session Control Function (CSCF) element in an IMS subsystem.

**Step 1:** To initiate a communication session, the user device sends control signaling information to the applicable AF, requesting that a session and the bandwidth necessary to support that session be admitted into the network.

**Step 2:** The AF issues a request for bandwidth to the Policy Decision Function (PDF). The PDF then performs route computation based on a number of factors, including the quality of service (QoS) required, network topology, available bandwidth, path diversity, and static subscriber-related policy parameters, such as information entitlements. Based on these factors, the PDF determines whether the session can be admitted and computes the optimal path through the portion of the network for which it is responsible.

**Step 3:** Once the PDF decides to admit the call, it decrements the available bandwidth and then "pushes" the policies out to the Policy Enforcement Points (PEPs) in both the access network (to, say, an access node) and the core network (to an IP edge router, for instance).

**Step 4:** The user session – via the bearer path – then progresses through the network to its destination.

ment, with each network (access, aggregation, core network, etc.) potentially being owned/operated by different business entities. This requirement is reflected in market trends whereby wholesale broadband is offered to Internet Service Providers (ISPs) by incumbent operators, while broadband subscribers pay the ISPs for their broadband pipes.

• TISPAN assumes that the end user/subscriber may also have an in-home network, such as a small WiFi network used to connect a number of PCs or other IP-capable consumer electronic goods.

• As a further consequence of the potential multi-network/multi-owner nature of the NGN (which may result in networks on the call or the bearer path using network address and port translators, or NAPTs), the RACS architecture must also ensure that services can work in the presence of such devices.

## Meeting the requirements

Figure 1 highlights the two functions that comprise RACS – namely the Service-based Policy Decision Function (SPDF) and the Access Resource Admission Control Function (A-RACF).

Essentially, both entities can be termed Policy Decision Functions (PDFs), in that they receive application requests for network resources and make policy decisions based on (operator-provisioned) policy rules and on dynamic information about the network resources within the network. Note that the RACS architecture does not specify how such dynamic information, specifically as it relates to network topology, is obtained from the network or how it is stored and subsequently provided to the RACS for the purpose of bandwidth management decisions (e.g., connection admission control, or CAC).

The result of a policy may lead to further action being taken in terms of requiring entities in the transport layer to enforce transport policies. These entities, which include border gateway functions (BGFs) and IP edge routers, are termed Policy Enforcement Points, or PEPs. Typically, the BGF can perform such functions as gate control, per-IP-flow admission control, packet marking, traffic policing, network address translation (NAT), hosted-NAT, and usage metering.

The use of two PDF entities is driven by the multi-network/multi-owner requirement discussed earlier, and recognizes that each network domain owner (operator) will likely have its own policies rather than trusting another operator to directly manipulate transport resources in its domain. Thus, the A-RACF specifically looks after and applies policy decisions on behalf of the access network (operator), while the SPDF applies policies on behalf of the core network (operator). The use of distinct PDF entities also means that the interface between the SPDF and the A-RACF can be an inter-domain interface.

Another consequence of the multi-network/multi-owner requirement is that applications can be owned and/or operated by a business entity that does not necessarily provide the network(s). As a result, the interface between the Application Functions [e.g., an IMS Call Session Control Function (CSCF) or a PSTN/ISDN emulation call server] and the SPDF can be inter-domain.

In recognition of the fact that any of the networks involved in a particular communication may introduce NAPTs, the RACS architecture allows the SPDF to request a BGF to perform hosted-NAT. The SPDF uses this function of the BGF to obtain NAT bindings, request media latching, and report its findings to applications. In turn, the applications may correct any addressing information embedded within the application signaling, thus enabling the service (signaling and media) to get through the NAPTs.

Lastly, an SPDF can communicate with multiple A-RACFs and an A-RACF can communicate with multiple SPDFs. In terms of a particular RACS session (to reserve or modify resources for a specific service), this "multiplicity relationship" means that an SPDF "talks to" an A-RACF but, in making its policy decision, the A-RACF needs to take into account resources already assigned to this line in support of other services (which may have been made by the A-RACF as a result of resource requests emanating from other SPDFs).

### Putting it all together

To better understand how the various components of the RACS architecture work together, the following example is a walk-through of a RACS session.

Prior to any interaction with services (such as IMS), an IP connection must first be established. To obtain an IP pipe, the user connects to the network through another subsystem called the Network Attachment Subsystem (NASS), which is responsible for authentication and authorization of the subscriber to the access network. The NASS also provides the IP address and binds this IP address with such parameters as a line identifier, which identifies the access type.

Once attached to the access network, the NASS uploads a set of static policy parameters to the A-RACF in the RACS. These static policy parameters are derived from the user network profile, which is stored in the profile database in the access network. Static policy parameters may include such items as the maximum bandwidth that the access line can support.

At this point, the user can begin to interact with the services. To do this, some application signaling will take place between the user and an Application Function (AF) in order to register the user with the application. When a user starts to use a service (e.g., through the initiation of a SIP INVITE to a CSCF in the case of an IMS service), it is at this point that the AF begins to interact with the RACS – specifically to communicate the bandwidth and potentially other resource control requirements to the SPDF. Note that the AF discovers which SPDF to talk to either by querying the NASS – specifically the Connectivity Session Location Function (CLF) – to obtain the IP address or Fully Qualified Domain Name (FQDN) of the SPDF, or via some other means, such as static configuration.

The SPDF then applies policies (which may include simple rules such as "is this application allowed to make this request for network resources?") and, if appropriate, initiates a command to the BGF to enforce particular traffic policies. The SPDF may also initiate a further resource request to the A-RACF to request reservation of resources in the access network. This subsequent request may result (after application of policy at the A-RACF) in a command to the IP edge router to apply a particular traffic policy. Note that RACS Release 1 does not specify a protocol for the interface between the A-RACF and the IP edge router, although limited stage-2 guidance is provided in ES282003.

Given that it acts as the fulcrum for policy-based resource control in the network, the SPDF also supports the coordination of resource control requests/responses for a particular RACS session. For example, the SPDF may wait for a response from the A-RACF as to whether or not resources can be assigned before replying to the AF with the overall status of all resources that have been assigned.

### Push versus Pull

The approach to taking a policy decision can be described as either "Push" or "Pull," depending on whether the applications push policies down to the PDFs (and subsequently to the PEPs, such as the BGF in the transport network) or whether the policy decision is solicited from the transport network itself (via signaling in the transport layer). More specifically:

• the Push model – so named because the PDF itself "pushes" policies out

**Figure 3. Pull call model (as applied in a wireless network)**

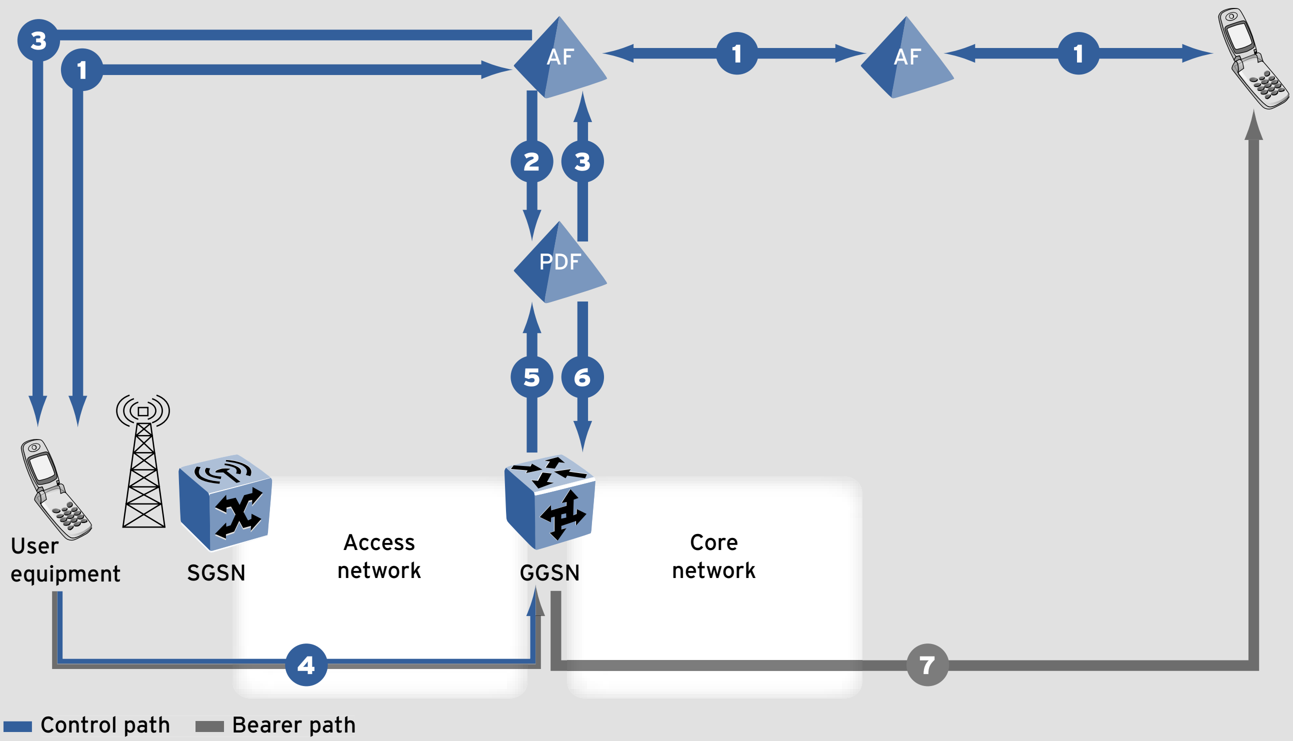


This figure illustrates a simplified view of the Pull call model, as it would be applied in a wireline network. As in Figure 2, a single network is shown.

**Steps 1 and 2:** The Pull call model follows the same first two steps as in Figure 2, where the user equipment sends a session request to the applicable Application Function (AF), and the AF, in turn, sends the request to the Policy Decision Function (PDF).

**Step 3:** Once the PDF receives the session request, it sends a response – in the form of a token – authorizing the session. The token is sent back to the AF, which in turn sends that token back to the user equipment.

**Step 4:** The user equipment receives the token and recognizes that the session request has been affirmed. At this point, the user equipment sends a request for bandwidth and simultaneously begins sending the session data (bearer traffic, indicated on the diagram by the combined blue and grey arrow). The traffic flow is intercepted by the Policy Enforcement Point (PEP), which in this example is a Gateway GPRS Support Node (GGSN).

**Step 5:** The GGSN questions the PDF to obtain authorization for the new token it received. The PDF then checks the available bandwidth.

**Step 6:** The PDF makes the bandwidth management decision and informs the GGSN, allowing the GGSN to update its policies.

**Step 7:** The session information (bearer traffic) continues to its destination.

to the access network – may be used regardless of whether the user equipment does or does not itself support native QoS signaling. In this model, which is described in detail in Figure 2, the AF issues a request for bandwidth to the PDF. The PDF then checks the available bandwidth and authorizes QoS for the service if there is sufficient bandwidth. The PDF then decrements the available bandwidth and notifies the PEP function(s), so that the PEP can update its policies; and

• the Pull model – so named because the access network itself "pulls" policies from the PDF – requires that the user equipment is capable of using Layer 2 or Layer 3 signaling to signal QoS requests necessary to meet its QoS needs. An example of Layer 3 (IP) signaling is the Resource Reservation Protocol (RSVP). In this Pull mode, which is described in detail in Figure 3, the PEP detects a new media stream for which it has not received authorization. The PEP questions the PDF to obtain authorization. The PDF then checks the available bandwidth and consults with the AF (such as a CSCF in IMS-based subsystems) before authorizing the new stream. The PDF then decrements the available bandwidth and informs the PEP, allowing the PEP to update its policies.

The RACS Release 1 supports only a Push-based approach to policy control. Other policy-based resource control architectures, such as the 3GPP Service-Based Local Policy, support a Pull-based approach. It is not yet clear whether future policy-based architectures will require support of both approaches.

Nortel recognizes, however, that both models will co-exist at least in the short term and has plans to support both in its policy control solutions.

### Reservation model

In addition to supporting both Push and Pull modes, a PDF can support a single-stage or two-stage bandwidth reservation operation. In a two-stage operation, the bandwidth management function initially reserves the bandwidth after being informed of a call request. When the call is answered, the PDF confirms the bandwidth reservation. If the call is not answered, the bandwidth may be released. The RACS architecture supports both single-stage and two-stage resource reservation.

As with the Push and Pull models, Nortel recognizes that both reservation models will co-exist in the short term and plans to support both in its policy control solutions.

### Future directions

Policy-based resource control architectures will continue to evolve based on:

• the requirements of new services, such as support of multicast for video/IPTV-type services;
• the need to support new business models; and
• the need to further clarify how the information that supports policy-based resource decisions is obtained, stored, and used by the policy control infrastructure.

While architectures in various standards development organizations remain misaligned at present, it is likely that pressure from operators for fixed-mobile convergence, together with pressure from vendors to cut development costs, will drive architectural alignment.

Nortel will continue to be an active player in all these discussions, through its strong technical contributions and participation in all of the major standards bodies, as well as through its work with customers and partners.

In addition, Nortel continues to build on its significant experience in policy-based resource control (pre-IMS) and, moving forward, will bring that wealth of experience to its IMS development activities.

Through all of these efforts, Nortel will help to address the current and emerging technical challenges with policy control, to ensure that policy-related architectures are aligned going forward and that its network solutions support all existing and emerging standards for policy control in general and for policy-based resource control in particular. ■

***Nalin Mistry*** *is Senior Architect, Mobility and Converged Core Networks.*
*Acknowledgement: The author would like to note the contributions of* ***Chris Hogg*** *of Nortel's Carrier NGN International Standards and Strategy team.*

**References:**

[1] For the list of TISPAN Release 1 Specifications, see http://portal.etsi.org/docbox/TISPAN/Open/NGN_Published/PUBLISHED_NGN_SPECIFICATIONS.doc

[2] ES282003: Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); Resource and Admission Control Sub-system (RACS); Functional Architecture

[3] ES282004: Telecommunications and Internet converged Services and Protocols for Advanced Networking (TISPAN); NGN Functional Architecture; Network Attachment Sub-System (NASS)

**Note:** ETSI Specifications can be downloaded at: http://pda.etsi.org/pda/queryform.asp (ETSI on-line account required).

# IMS: Addressing the infrastructure challenges of converged multimedia services delivery

by John Boden

The IP Multimedia Subsystem (IMS) is emerging as the multimedia services delivery architecture of choice around the world. Yet several technology challenges need to be addressed to optimize the IMS architecture for a truly converged communications environment and to ensure that the network infrastructure can provide the quality of service, scalability, and reliability that operators need. Nortel, through hundreds of pre-IMS deployments and more than a dozen IMS customer trials – the only vendor, in fact, with experience across all access technologies – has gained unique insight into the challenges of building a packet infrastructure for services delivered by a next-generation network.

Over the past few years, a new multimedia services delivery architecture, called IP Multimedia Subsystem (IMS), has emerged as one of the hottest topics in the telecom industry today. IMS is being considered – by a diverse set of operators and standards development organizations worldwide – as the architecture of choice for delivering real-time multimedia services over next-generation converged infrastructures.

Even though IMS cuts the ties between services and their specialized networks, the IMS components (see page 44) must still interwork closely with the underlying network.

In fact, effective interworking with the network infrastructure, or "plumbing," is the most important factor to successful commercial deployment of IMS. The underlying network must be able to support IMS deployment at scale, provide bullet-proof reliability, support real-time performance, dynamically apply varying quality-of-service (QoS) requirements, and handle the associated messaging load.

Engineering the network and the IMS functionality for these requirements, therefore, requires a deep understanding of how IMS-based applications specifically impact the network.

While the industry is largely still in the early trial phase and IMS standardization activities are ongoing (see page 46), at Nortel, we have already amassed significant experience in deploying and trialing IMS-based solutions with lead customers around the world and across all network environments. In fact, Nortel is currently the only vendor engaged in IMS trials with leading operators across all access domains – wireline, cable, CDMA, and GSM/UMTS. These trials include, for example, an evaluation in Verizon's lab in Waltham, Massachusetts, of Nortel's IMS multivendor interoperability and full support for voice services over fiber to the premise (FTTP).

In addition, Nortel has already gained many years of experience in deploying IMS-like capabilities with its MCS (Multimedia Communication Server) and Communication Server 2000 (CS 2000) products. While these products are considered pre-IMS due to IMS standards being immature, they are based on SIP (Session Initiation Protocol) – the signaling protocol at the heart of IMS. Experience with deploying the MCS and CS 2000 provided our development teams with significant insight into the kind of network requirements that IMS-based applications drive.

The teams then fed these insights and technical contributions back into the various standards development bodies and used the learnings to build Nortel's IMS offer. For instance, the initial SIP standard specifications defined only a basic voice call and a few basic features, such as call forwarding. Nortel, through its work with several lead customers, went beyond the SIP standards to develop an extensive and rich suite of sophisticated voice features, including automatic call distribution features, and contributed these to the SIP standards to help drive the next level of open SIP-based applications and to set the stage for other real-time services, such as video. These specifications were also provided free through a simple download from Nortel's corporate web site, at www.nortel.com.

Building on this experience, Nortel is now applying its expertise to address the next set of IMS challenges, specifically to enable IMS to work seamlessly across the diverse next-generation environments that are emerging.

### Key technology challenges moving forward

Despite the widespread acceptance of the basic architecture (see page 44), there are still many challenges with extending and adapting IMS so that it becomes the converged architecture for multimedia services over multiple access types.

Indeed, while deploying the IMS

framework as a specialized functional overlay is relatively straightforward, utilizing IMS as the basis for a converged network – one that operates seamlessly across the different access domains and service types – requires that several important challenges be addressed, including:

- policy control;
- maintaining architectural alignment as the various IMS specifications are adapted and standardized across multiple access technologies and organizations;
- parity with existing services; and
- moving from a functional architecture to deployable solutions.

### Policy control

Understanding the policies required in an IMS-based network and the design of a policy control capability that operates consistently end-to-end are critical to successful deployment of IMS across multiple access types.

At a high level, IMS networks involve two types of policies: user-level policies and application-level policies.

- User-level policies are defined as the rules that govern the type of service – amount of bandwidth, quality of service, location, presence, and security levels, for instance – to which users are entitled, according to the service level agreements (SLAs) they have with their service providers.
- Application-level policies are those that define the specific behaviors of the various types of applications – voice, video, or data – as well as the requirements that each needs, such as the amount of bandwidth and priority level. For instance, video requires high bandwidth but is less susceptible to packet loss and jitter than voice, which for its part requires significantly less bandwidth.

In the past, these policies were an inherent quality hard-wired into each dedicated network. A voice network, for instance, had certain policies about traffic and user behaviors that were embedded into the basic design of that voice network.

In a next-generation converged network, however, where multiple types of traffic flows ride on the same core network, the network simply cannot be designed for only one type of data flow. What is needed instead is an intelligent policy engine that is able to direct the network to behave appropriately and apply the right network resources for each type of traffic flow.

The IMS standards currently being defined include such a policy engine, called the Policy Decision Function (PDF). The PDF provides the key linkage between the applications and the underlying transport network, and is responsible for governing the admission of applications traffic and monitoring network resources for that traffic. The PDF is vital to the network's ability to recognize and behave according to the various policies, as well as to ensuring that the right amounts of bandwidth are provisioned, in real time, for the differing quality-of-service requirements.

> Through hundreds of pre-IMS deployments and more than a dozen IMS customer trials, Nortel has gained unique insight into the challenges of building a packet infrastructure for services delivered by a next-generation network.

Indeed, policy control will enable operators to deliver and guarantee quality of service, as well as operate their networks in the most cost-effective fashion, versus over-provisioning bandwidth. Nortel has a significant development activity under way in this area, and policy-based resource control is the subject of a separate article in this issue, on page 33.

At the same time, IMS offers the potential to solve a number of challenges with which the industry has been wrestling.

Security policy is just one example, and specifically understanding the policies required – and how to implement them in a next-generation converged network – to protect against piracy in IPTV applications, which has long been a difficult challenge for the digital rights management industry. Here, IMS has strong authentication and identity management capabilities upon which to build. Nortel is currently very involved within these communities, contributing its technical expertise to develop optimal technology solutions, as well as to understand the policy-related requirements on the IMS and underlying infrastructures.

Even though such solutions are in their infancy, Nortel continues to remain at the forefront of the various standards bodies, to ensure these solutions become embedded into evolving IMS and IPTV standards for the benefit of the industry as a whole.

### Architectural alignment

The second challenge involves aligning the different IMS variations that are emerging in the industry. While there is alignment around the basic IMS architecture at a high level by the different standards development bodies, IMS must nevertheless undergo adaptations to suit the specific behaviors of individual networks. These unique requirements are driven by the nature of the access domain (e.g., hybrid fiber coax, DSL, WiFi, and WiMAX) and the associated established policy management of each.

For instance, the differences in QoS requirements in wireless versus wireline networks are significant. Wireless operators, who must operate within an inherently bandwidth-constrained environment, are looking not only to ensure that a minimum amount of bandwidth can be provided to each subscriber in order to deliver acceptable QoS, but also to balance that against the need to conserve spectrum and share bandwidth

*Continued on page 46...*

## Key components and values of the IMS architecture

The IP Multimedia Subsystem (IMS) architecture that is today being adopted by operators and standards bodies worldwide promises to revolutionize the way next-generation communications services are created, deployed, and delivered. It will also provide a real-time services delivery platform for creating disruptive new services and applications concepts.

Fundamentally, IMS transforms services delivery by separating the services control capabilities that were previously tightly coupled in specialized purpose-built networks. For example, voice services required a specific network to handle the real-time behavior of voice traffic. Data services needed networks optimized for symmetrical data transport. And video-on-demand services required networks customized for the asymmetrical bandwidth requirements of video delivery.

With next-generation converged infrastructures able to support all of these types of services and applications, the multiple services control constructs of the past are both unnecessary and costly. IMS provides a single, standardized services control infrastructure that is separate from the underlying network and independent of access technology and device type. IMS logically breaks out common communications functions – including session control, common media resources such as group management lists, tones, and announcements, as well as voice, video, and data applications, to name a few – into individual building blocks that can be easily and quickly mixed and matched.

The idea is that by interfacing with the underlying network in a standard way, new applications concepts can be deployed much more rapidly and cost-effectively, by mixing and matching established individual building blocks (such as authentication, accounting, common data store, and session control) rather than develop each service separately for each individual network type, as was the case in the past.

By contrast, in traditional services delivery, all of these various functions are tightly coupled to both the application and to the access technology.

Take, for example, a server that performs recorded announcements or tones for a wireline voice service. There is currently no standard way to reuse that server to perform announcements in a wireless environment, and the functionality would have to be developed from scratch, even though the announcements are identical. With IMS, service developers can simply reuse one standardized, common communications function – for any service, any access technology, and any device type.

### Core IMS architecture

The heart of the IMS architecture (see diagram) is the distributed services control capability, which consists of three main elements:

• the *Call Session Control Function* (CSCF) provides session control between end users using the standardized signaling control protocol called Session Initiation Protocol, or SIP. To accommodate roaming and inter-network operation, CSCFs can assume different roles, including:

– the Proxy CSCF (P-CSCF), which will most likely be located close to where the end user is physically located. The P-CSCF provides the connection from the device into the network, and its main job is to relay SIP signaling between a user's terminal and a serving CSCF;

– the Interrogating CSCF (I-CSCF), which provides the interface into the subscriber's home network (e.g., between peer carrier networks); and

– the Serving CSCF (S-CSCF), the "brains" of the network, which coordinates all authentication and services, and performs the basic session origination or termination treatment for SIP messages.

• The *Home Subscriber Server* (HSS) houses all of the subscriber data and the rules that govern which services a subscriber has access to. The fact that subscriber data is held centrally has several advantages over traditional fixed networks, where subscriber data was typically provisioned on the line switch where the subscriber was located. If the subscriber physically moved, then the provisioning and routing information also had to move. In addition to enabling mobility and nomadicity, and aiding in the scaling of the network – and since subscribers don't have to be provisioned each time a person moves or is moved when a network expands – centralizing the data also means that common data can be stored centrally. This centralization allows multiple application servers or instances of an application server to access the stored data, thereby reducing the complexity of provisioning a new service.

• *application servers* host the specific services – video, voice mail, call waiting, IPTV, Push-to-Talk, for example – and interact with the S-CSCF using SIP.

In addition, the IMS architecture includes several other important elements, including:

• a *Policy Decision Function* (PDF), which provides the interface to the different access types and the underlying network. The main job of the PDF is to ensure that the right network resources (such as bandwidth) are provisioned for a given quality of service, in real time and as the service requirements change dynamically, such as with a session that begins as a voice call and then adds a video application. The PDF also allows the IMS service that's being delivered to the end user to be controlled and charged, and to be provisioned with the proper level of quality. (This function is described in more detail in the article on page 33.);

• a suite of *media resources*, which house

## IMS architecture

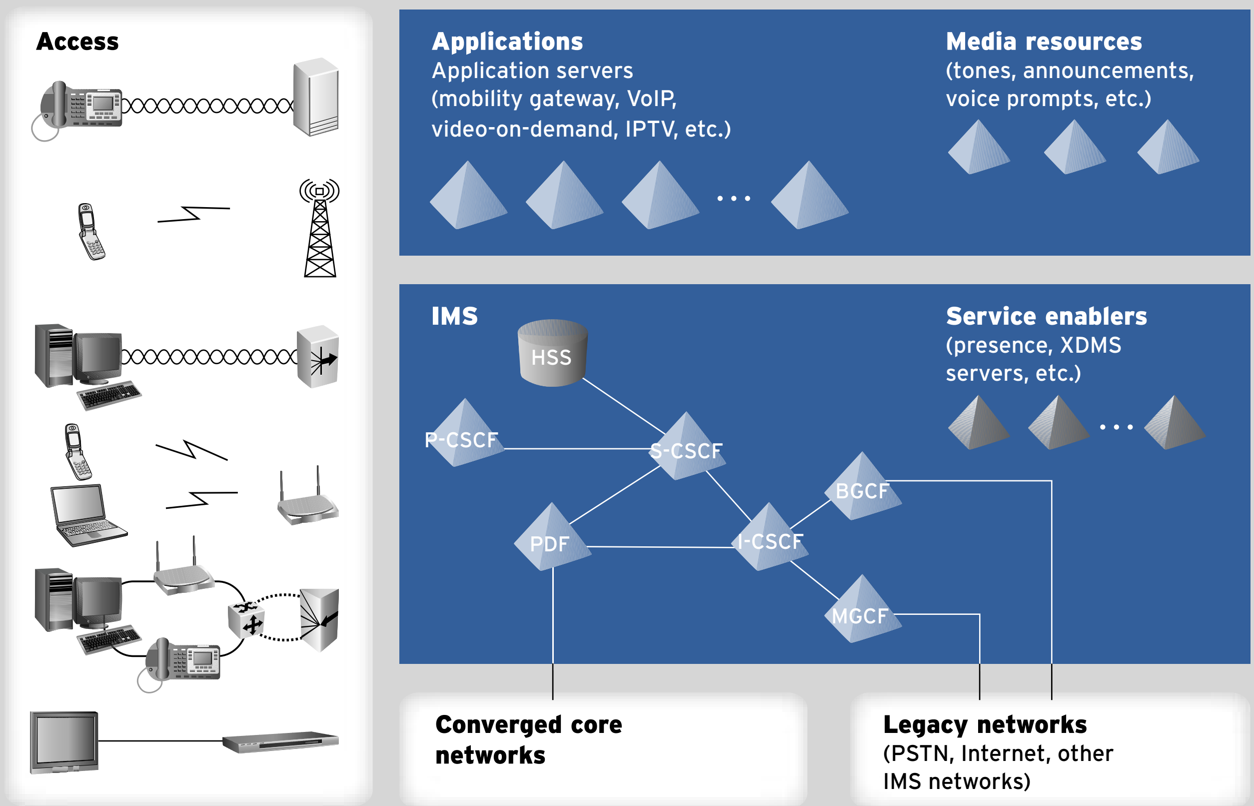


such common functions as tones, announcements, and voice or video prompts;
• a suite of *service enablers*, which include common capabilities such as presence servers and common data servers; and
• *call servers* – called the Media Gateway Control Function (MGCF) and Breakout Gateway Control Function (BGCF) – that allow the SIP-based applications to interwork with the PSTN and other legacy networks.

With these components, IMS provides a distributed services control architecture that enables applications to be decoupled from both the underlying network and the applications that sit on top. In this way, wholly new "blended" services – voice combined with video and text messaging, for example – can be created that will further enrich end-users' lives.

For instance, with IMS, users will no longer need to interact with multiple separate networks – all accessed and managed using different user IDs and passwords and separate devices, each with their own learning curve. With IMS, a user would be able to receive all of his or her communications services – whether voice, data, mobility, television, radio, cable, or gaming – on a single device, and manage them with a single user ID and phone number and with integrated voice mail.

While this simplified access and management helps immensely in reducing the required number of access devices, IMS also provides identity management, which enables the network to intuitively know how to find the user and direct the service to the device that is most convenient to that individual. Identity management allows the introduction of a personalized service based on the user and not on the particular device he or she is using at the moment.

IMS will enable operators to greatly reduce the time required to deliver new, innovative – and revenue-generating – services, while both reducing operating costs and ensuring a much simpler and higher quality of experience for real-time applications than much of today's Internet-based applications provide.

equitably across the subscriber base.

In the wireline domain, or DSL space, on the other hand, where bandwidth is more abundant, operators are looking to deliver much higher levels of bandwidth in order to meet the strict SLAs they have with their subscribers, many of whom expect a rich quality of experience.

In a converged service environment like IMS, however, where an end-user service may need to operate across both wireline and wireless access networks, these two differing bandwidth policies would result in a network unable to provide uniform quality of service end-to-end.

As these adaptations are developed to support the unique requirements of the different access types, the challenge is to ensure that they evolve in a consistent architectural fashion and do not diverge in significant ways. Otherwise, these networks will not interwork and the result would be a separate IMS architecture and implementation per access type.

Addressing this challenge requires industry collaboration to define and agree on a truly common IMS core architecture for multimedia services across all network types.

This agreement and cooperation among standards bodies is fundamental to delivering a common architecture and network that supports ubiquitous services delivery across all access technologies and ensures "any service, anywhere, any time." The Third Generation Partnership Project (3GPP) is playing a central role in this effort, by expanding its scope to include primary voice service support and multiple access technologies (e.g., WiFi) and by incorporating requirements from both the Telecoms and Internet converged Services and Protocols for Advanced Networks (TISPAN) and PacketCable standards organizations.

Nortel is also a major contributor to this effort. Through its experience in trialing IMS across all access technologies, Nortel has gained unique understanding

## Evolution of the IMS architecture

The IP Multimedia Subsystem (IMS) architecture was developed originally in the GSM/UMTS network environment to deliver IP-based multimedia services over a wireless packet access network. [The initial IMS concept was conceived in 1998 and put forward to the Third Generation Partnership Project (3GPP).] IMS was seen as a way to take advantage of the high speed and increased throughput that these wireless networks were beginning to offer, to introduce new services – voice, video, presence, messaging, and so on – to wireless subscribers. In this domain, IMS started quite modestly as an IP overlay to GPRS (general packet radio service) and UMTS networks.

About the same time, the standards development organizations of other domains – CDMA wireless, wireline, and cable – were beginning to define their own next-generation multimedia architectures. These organizations recognized the potential of the 3GPP's new IMS architecture, not only because of its established head start but also because it addressed such key carrier concerns and value differentiators as security, quality of service, scalability, and mobility. These organizations subsequently adopted IMS as the basis for their new multimedia core architectures. Specifically:

• the Third Generation Partnership Project 2 (3GPP2) body adopted IMS as the basis for its Multimedia Domain (MMD) in CDMA;

• the European Telecommunications Standards Institute (ETSI), through its Technical Committee TISPAN (Telecoms and Internet converged Services and Protocols for Advanced Networks), picked up IMS as part of its next-generation multimedia standard capability;

• the Alliance for Telecommunications Industry Solutions (ATIS) Packet Technologies and Systems Committee (PTSC) is building on the IMS framework for wireline networks in the U.S.; and

• the PacketCable standards organization, in the cable domain, adopted IMS as the basis for its multimedia services moving forward.

As a result, the term 'IMS' has today come to mean a whole class of related architectures and standards that have a common look and feel.

Even though the different organizations are all working on variations of IMS to accommodate the various characteristics of their specific network types, the core architecture remains common to all. This architecture includes several fundamental components, including a Call Session Control Function (CSCF), Home Subscriber Server (HSS), and application servers (see page 44).

These IMS components interact with each other and with the various access devices via SIP, the standardized signaling protocol used to establish "sessions," where any number of different traffic types and services can be linked together into a single seamless session. A user, for example, could begin a session as a typical voice call, and then add or drop any number of different types of communications – from a video call, to a text message, to a collaborative "whiteboarding" capability – at any time during a single session.

into where architectural discrepancies can occur and how to ensure that the needs of all access requirements are met. Nortel is already addressing such challenges in its broad range of IMS trials and deployments with its IMS product solution.

Indeed, Nortel made a fundamental decision several years ago, at the start of its IMS development, to build a single IMS core for the GSM/UMTS, CDMA, cable, and wireline spaces, even in advance of the TISPAN, PacketCable, 3GPP, and 3GPP2 convergence on IMS.

That choice has proved prescient. In the more than a dozen IMS trials that Nortel is currently running or has completed, it has been able to use the same core components across all networks, varying only the applications that run on top of that core based on customers' needs.

Through its IMS trials, Nortel has come to understand how to best accommodate, in its IMS core, the diversity of access network behaviors.

For instance, the degree of user authentication tends to vary in different access networks, from largely unauthenticated in the case of wireline access, to authenticated in GSM environments. As well, the time required to transmit information to, and receive a response from, the various devices ranges from nearly instantaneous in wireline environments to much longer in wireless networks.

Nortel has enabled its IMS solution to address this diversity in behavior by abstracting the common session, service, and media control functions into a reusable multimedia core, based on open and standard SIP and Diameter (an IETF-defined authentication, authorization, and accounting protocol) interfaces. Access-, network-, and application-specific components are decoupled from this common core and are deployed around the edges of the IMS infrastructure.

This decoupling is an important architectural principle within Nortel's IMS solution, because it enables truly converged services that span different network environments and operate across them.

At the same time, Nortel is playing an important role in all the key standards development bodies to ensure that the various organizations cooperate to address these cross-access technology issues and is sharing the lessons it has learned from its IMS trials with the appropriate standards bodies.

### Services parity

A third key challenge for broad IMS-based deployments involves the service set. For IMS to become the basis of a truly converged network, it must support at least parity of the existing legacy services, such as PSTN voice, wireless voice, and short message services. The level of parity ranges from acceptable equivalency to complete and transparent emulation.

Nortel is currently the only vendor engaged in IMS trials with leading operators across all access domains – wireline, cable, CDMA, and GSM/UMTS.

This challenge involves defining a "super-set" of services, and then ensuring that these services offer equivalent functionality and operate with equivalent reliability and quality across all domains. What makes this aspect challenging is the strong desire to not re-implement today's PSTN and to not embed legacy PSTN concepts such that they impair future services and deployment models.

For subscribers to be willing to "switch over" to IMS services, the services offered must deliver the same functionality that users have come to expect in their traditional networks. Dialing *76 to return a call, for instance, could be replaced with "click a button to return the call."

The service set must also include traditional E.164 (the international public telecommunications numbering plan defined by the ITU-T) and national dial plan support, along with fundamental regulatory support of such services as emergency services (e.g., E.911) and Lawful Intercept. While these services are not as compelling as some of the more exciting – and revenue-generating – features that IMS enables, they are nevertheless gating factors to deployment. Without them, the IMS solution cannot be effectively deployed.

What's more, these requirements are challenging to implement in a next-generation network environment. For example, a mobile subscriber from, say, Texas, happens to be traveling in Hong Kong and makes an emergency call. The network can't very well route the call to the local 911 center back in Texas. Rather, the network has to be able to identify the subscriber, match that subscriber information to his or her current location, and rapidly deliver the call to the emergency center closest to the subscriber.

As Nortel did with its CS 2000 softswitch products, it viewed these requirements, especially emergency services and Lawful Intercept, as fundamental for the first release of its IMS product solution designed for commercial scale and deployment. Nortel leveraged the services capabilities already built into both its Voice over IP (VoIP) and SIP-based products to ensure that services parity was factored into the design from inception.

### Solution deployment

Another challenge involves implementing the IMS functional architecture – where each function is logically broken out as an individual component – in a fashion that is both compliant with the emerging standards and as efficient as possible.

Certainly, viable deployment models need to be defined that meet the requirements of different carrier types, with their different network topologies and technologies, and highly varied business models. These IMS deployment models also need to be flexible and evolvable in order to grow and scale, as well as adapt to new business paradigms.

While each function within the architecture could, theoretically, be deployed as a separate product, such a deployment model is highly unlikely and very costly. Instead, several functional components could be housed together in the same node, offering a more cost-efficient approach.

To provide customers with the deployment flexibility they will need, Nortel designed its IMS solution to enable the combination and distribution of various components, such as the Serving-CSCF, Proxy-CSCF, and Interrogating-CSCF.

Based on its experience with more than 100 live VoIP networks that support SIP and the additional nearly 50 live MCS 5200 networks (which, like IMS, are pure SIP-based), Nortel understood that in large-scale networks there are instances where integrating functional components is necessary to minimize messaging volume and inter-component latency, and to ensure that the time it takes for an end-to-end call to traverse the various IMS components is not longer than the time it took in previous-generation networks for a call to traverse the various switching elements.

Of course, making this challenge easier to address was Nortel's understanding of the nature of traffic flows on an infrastructure that is based on optical and packet (specifically Ethernet) technologies. For instance, traffic requires less time in transit in an optical Ethernet infrastructure, both because this network relies on fewer "boxes" – hubs and routers – and because new technologies, such as eDCO, can extend fiber's reach between two endpoints. (For more on eDCO, or electronic Dynamically Compensating Optics, see page 10.) Faster, more efficient transport, in turn, means that more time can be devoted to processing in such IMS nodes as the CSCF.

As a result of their infrastructure knowledge, combined with their deployment and trial experience, Nortel teams have gained valuable expertise in knowing how to group functional components for cost-efficiencies, as well as which components should be separated to meet operators' evolving requirements for scalability.

**In the more than a dozen IMS trials that Nortel is currently running or has completed, it has been able to use the same core components across all networks, varying only the applications that run on top of that core based on customers' needs.**

### Nortel's IMS solution

Nortel's IMS solution is based on a highly distributed and flexible architecture, which allows customers to select best-in-class multivendor services and infrastructure components within their IMS deployments. This flexibility represents a shift in the way customers deploy networks and requires a software architecture that facilitates and enables multivendor interoperability, integration, and testing. Nortel's IMS solution has been designed to address this deployment flexibility and multivendor interworking, and includes:

• open standards compliance to key service-enabling interfaces (e.g., session control and user-profile access);
• an application-centric software architecture that allows flexible deployment and customization via configuration, as well as aligns with industry software practices rather than with the nodal models traditionally used in the telecom industry;
• integration tools – such as real-time message tracing and message injection – to enable rapid interoperability testing; and
• software architecture flexibility to allow rapid adaptation, often through configuration, to multivendor interfaces.

The flexibility and interoperability strengths of Nortel's IMS solution have been demonstrated via participation in multiple customer deployments as well as industry interoperability forums, such as the Open Mobile Alliance (OMA) TestFest. These activities have included interoperability of the IMS core with many different vendors' products, including:

• other vendors' IMS core infrastructure products [CSCF, Home Subscriber Server (HSS)];
• a wide variety of SIP application vendor products;
• multivendor access network and QoS/bandwidth-control products; and
• many different SIP client products, including PC-based clients, wireless devices, and dual-mode handsets (e.g., CDMA and WiFi).

The success of these deployments and interoperability forums was based on Nortel's standards-compliant and mature IMS platform, which consists of both purpose-built components and several existing industry-leading products.

Among the key elements of Nortel's IMS portfolio are:

• the Call Session Controller 1000, offered on Nortel's Versatile Services Engine (VSE) platform, which includes the CSCF functionalities and provides SIP session management and control;
• the Service Capability Manager 1000, which orchestrates the way application servers are brought together to deliver services to end users;
• the Home Subscriber Server 1000, which leverages decades of technology leadership and operational experience gained through Nortel's Home Location Register (HLR) portfolio to enable a

new level of subscriber identity management;
• the Media Gateway Controller 2000, which is based on Nortel's CS 2000 product, to provide for gateway control and interfaces to the PSTN;
• the PSTN Emulation Server, which supports full residential and business services and TDM and IP clients;
• the Application Server 2000 and 5200, which provide a complete set of integrated voice, multimedia, and mobility services over SIP; and
• the Policy Decision Function, which provides the brokering function between applications and the underlying network to ensure quality of service.

This solution builds on Nortel's significant leadership in developing and deploying both wireless packet networks and wireline VoIP networks. Nortel, in fact, was the first to deploy SIP voice and multimedia services in the wireless, wireline, and cable markets, and has extensive expertise deploying its SIP-based MCS 5200, with more than 40 SIP-based multimedia customers globally. In addition, Nortel has more than 600 carrier VoIP softswitches carrying live traffic in networks around the world – more than double the VoIP ports of its closest competitor.

Nortel's IMS solution also offers one of the industry's most advanced IMS session and control architectures, and utilizes the latest carrier-grade Advanced Telecom Computing Architecture (AdvancedTCA) technology to bring high reliability and interoperability.

We have also worked to ensure our IMS solution is fully standards-compliant and open to third-party development. In fact, Nortel collaborates with more than 100 developers, partners, and applications providers through its seven Open SIP Interoperability labs, six IMS Enablement and Innovation Centers in North America, Europe, and Asia, and two Joint Customer Innovation Centers.

An example of one of these IMS Enablement and Innovation Centers is a collaborative initiative with IBM in Montpellier, France, where both companies are working to showcase a live, integrated proof-of-concept IMS solution that can help customers realize dramatically shortened IMS service deployment times, simplified introduction of new services, and faster routes to next-generation service revenue opportunities.

As IMS further evolves, Nortel will continue to build on its significant leadership and experience to ensure network operators and service providers can fully leverage the intelligence, flexibility, and richness that IMS offers, while achieving robustness, scalability, and carrier-grade performance needed for commercial deployment. ■

***John Boden** is Leader, IMS and VoIP Strategy Team, Mobility and Converged Core Networks.*

# Newsbriefs

## Nortel Technology Award of Excellence program recognizes achievements in hardware, software, and solutions R&D

Nortel recently recognized three of its R&D teams as part of its global Technology Award of Excellence program – a program that rewards individuals and small teams whose innovations deliver clear customer value, exhibit exemplary behaviors, and contribute to Nortel's overall industry leadership.

The awards recognized a cross-section of achievements spanning hardware, software, and solutions development, including, respectively, the winning OFDM-MIMO team and two other teams that progressed to the final round of consideration: the HSDPA (High Speed Downlink Packet Access) team and the Converged Mobility Solution for Healthcare team.

### OFDM-MIMO team

OFDM-MIMO combines advanced multiple-input multiple-output (MIMO) antenna technology with orthogonal frequency division multiplexing (OFDM), a frequency-division-based modulation technique that uses orthogonal sub-carriers to minimize interference. It is becoming recognized industry-wide as a best-in-class solution for making efficient use of available spectrum, boosting throughput, significantly reducing operating costs, and enabling a better user experience. (For more information on OFDM-MIMO, see Issue 2 of the *Nortel Technical Journal,* page 26.)

The OFDM-MIMO team led the industry in researching, prototyping, and trialing OFDM-MIMO, accumulating an impressive list of industry firsts while being instrumental in embedding significant patented technology into key industry standards. Nortel's OFDM-MIMO technology, in fact, forms a key component of the standards for OFDM-MIMO in WiMAX, Third Generation Partnership Program (3GPP) Long Term Evolution, and the 3GPP2 OFDM-MIMO evolution.

More recently, Nortel led the breakthrough with other major CDMA vendors to agree upon an OFDM-MIMO framework for EV-DO Revision C standardization. This agreement is a significant wireless broadband industry milestone since all next-generation wireless broadband technologies (WiMAX 802.16e, 3GPP LTE/3GPP2, and CDMA EV-DO Revision C) are now based on OFDM-MIMO. The first commercialization of OFDM-MIMO for a high-power macrocellular system from Nortel will be with WiMAX 802.16e.

OFDM-MIMO will be the foundation technology for much of Nortel's wireless innovation, and indeed the industry's, for years to come. As a result of this team's innovation and drive, Nortel is in a strong leadership position in this key technology.

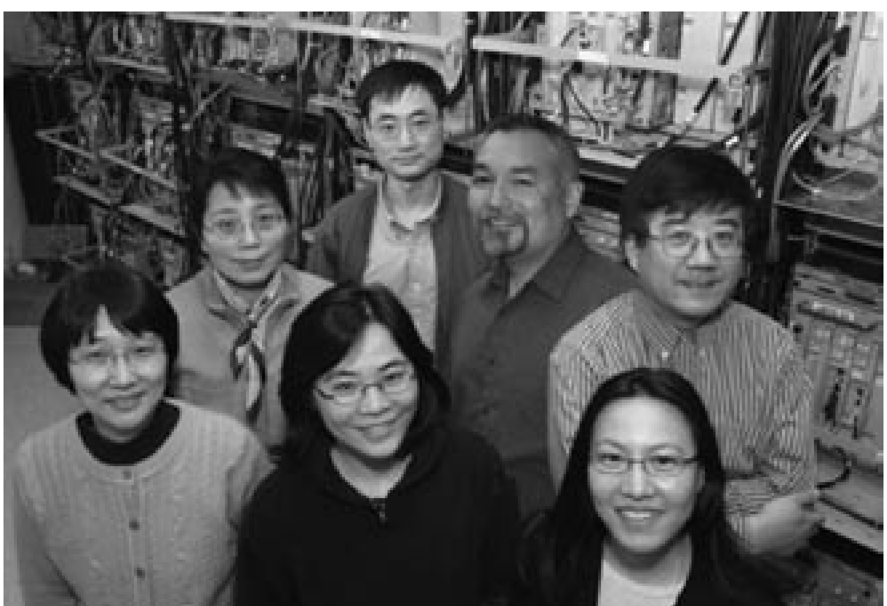
**OFDM-MIMO team: From left to right, front row: Jianglei Ma, Peiying Zhu, Mo-Han Fong; back row: Hang Zhang, Ming Jia, Charlie Martin, and Wen Tong.**

### HSDPA team

High Speed Downlink Packet Access (HSDPA) is a new Universal Mobile Telecommunications System (UMTS) packet air interface that delivers higher data rates, lower latency, and higher capacity to support mobile broadband services such as web browsing, video-streaming, mobile gaming, and mobile music and, in the future, Voice over Internet Protocol (VoIP). (For more information on HSDPA, see Issue 2 of the *Nortel Technical Journal*, page 13.)

This team developed an innovative solution that enables Nortel's UMTS wireless customers to upgrade to HSDPA via software-only upgrades to their existing Nortel UMTS base transceiver station (BTS) infrastructures – without sacrificing quality or performance. This software-based solution made it possible for Nortel to bring HSDPA to market ahead of its competitors, most of whom are building new

**HSDPA (High Speed Downlink Packet Access) team: From left to right, Philippe Bustin, Sarah Boumendil, Bastien Massie, Aurélie Gervais Jean-Marie, Luc Noiseaux, and Christophe Bejina.**

hardware to deliver this capability.

Now in full commercial deployment, HSDPA has achieved a number of world firsts:

- In July 2006, KTF, one of South Korea's leading cellular providers, launched a next-generation, ultra-high-speed 3.5G wireless network in Seoul and other cities across South Korea using HSDPA technology from the LG-Nortel joint venture. The network's capability to support high-speed, real-time, high-definition video content is a key differentiator over current technologies.
- In February 2006, Orange, using Nortel infrastructure, became the first operator in France to launch a live HSDPA network. The deployment took place in Lyon.
- And in November 2005, Nortel and Qualcomm were the first to achieve a call at rates up to 3.6 Mbit/s in the downlink with a category 6 mobile handset.

### Converged Mobility Solution for Healthcare team

This team integrated several Nortel technologies [including the Multimedia Communication Server (MCS), IP telephony and data, wireless local area network (WLAN), and security] with context-aware components (such as sensors and location systems) from third parties to develop an innovative solution for the healthcare market.

This development is an excellent example of technology innovation at the solutions level. Indeed, this team has extended Nortel's product value proposition by integrating all of the elements required to create a complete medical mobility solution.

This team also developed a user-based scenario to demonstrate how the solution could be applied in a healthcare setting. The demo has been showcased at a number of tradeshows and to targeted audiences, including customers and industry analysts. For more on Nortel's mobile medical solutions, see www.nortel.com/healthcare. ■

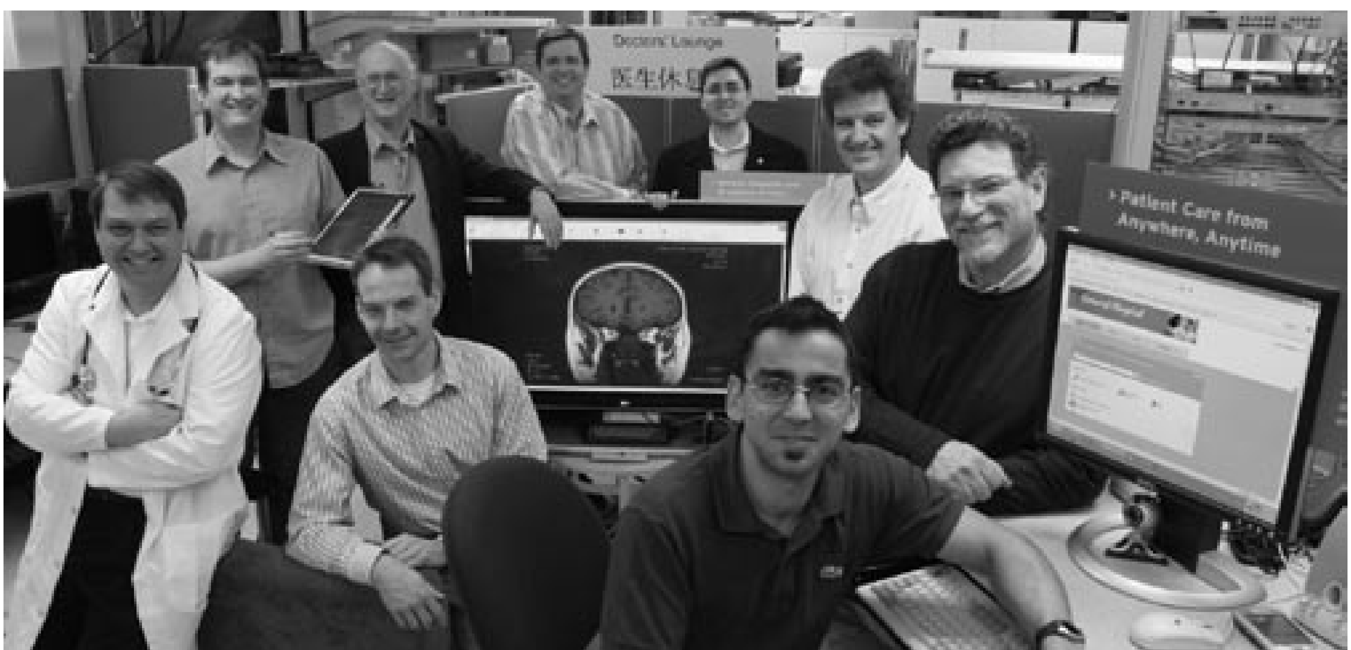


**Converged Mobility Solution for Healthcare team: From left to right, front row: Denis Plante, Jeff Fitchett, Guyves Achtari, Kent Felske; back row: Andrew Paryzek, Alan Graves, Bruce Wallace, Eric Bernier, and Guy Duxbury.**

## Nortel patent award ceremonies celebrate innovative spirit of employees

Reflecting the depth and breadth of its technology innovations, over the last several months Nortel has recognized some 653 Nortel inventors with awards for recently issued patents and patent applications filed.

The recent award ceremonies were held in eight Nortel R&D locations around the world, including: Ottawa, Canada; Richardson, Texas; Chateaufort, France; Harlow, U.K.; RTP, North Carolina; Billerica, Massachusetts; Santa Clara, California; and Galway, Ireland.

A number of the recipients at each ceremony were also honored with Cumulative Patent Awards for each multiple of five patents issued.

At year-end 2005, Nortel held approximately 3,700 issued U.S. patents and 1,800 in other countries worldwide. Since 2000, Nortel has consistently been ranked among the Top 50 companies globally for newly issued U.S. patents, as reported by IFI Claims, an established source of U.S. patent data.

Nortel's patent portfolio reflects the breadth of its R&D investment. Many of the patents cover standards-essential, standards-related, and other fundamental and core solutions in such areas as code division multiple access (CDMA), Universal Mobile

In Chateaufort, France

In Ottawa, Canada

In Richardson, Texas

Telecommunications System (UMTS), Global System for Mobile Communications (GSM), Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) transport technologies, Voice over Internet Protocol (VoIP), Ethernet, and Multiprotocol Label Switching (MPLS), to name a few.

Reflecting a strong customer focus, current R&D and associated patent activity encompasses innovations not only in hardware technologies, such as electronic Dynamically Compensating Optics (see page 10), but also increasingly in software, services, and applications, including innovations that focus on how to put solutions together and create value for customers, such as Nortel's solutions in the healthcare field (page 51). ■

Photographs on pages 50, 51 (bottom), and 52 (bottom) by Mike Pinder; page 51 (top) by Veronique Delacourt; page 52 (top) by Dominique Arestan; and page 53 by Andrew Thillainathan.