Nortel Technical **Journal**

Issue 1, February 2005

FOCUS ON **CONVERGENCE**

This inaugural issue of the Nortel Technical Journal highlights some of the key technology initiatives under way at Nortel to address Convergence, one of the most significant challenges — and opportunities — facing the industry today.

Contents

Issue 1, February 2005

WELCOME to the inaugural issue of the Nortel Technical Journal. This first issue highlights some of the technologies that our teams are developing to make convergence and its deployment by our customers a competitive reality. Upcoming issues will feature articles on advances in wireless, security, network architectures, and other leading-edge technologies.

Innovation for the convergence challenge

by Peter Carbone and Steven Foster

The growing costs and complexity of networking must be addressed before the industry can realize the full advantages of convergence. Nortel is addressing these challenges with technology innovation at all levels in the network, from common software platforms and new network engineering practices to integrated services delivery, a common wireline/wireless switching core, integrated optical/packet transport, and simplified network management systems.

For almost three decades, the industry has attempted to implement various forms of network convergence – from early integrated voice and data satellite-based services in the '70s, to ISDN and broadband ISDN in the late '80s/early '90s, to integrated on-demand services in the late '90s.

These early efforts at convergence, however, met with only limited success, primarily because all of the "pieces of the pie" were not in place – either the market was immature, the technology or standards were still evolving, or equipment from different vendors didn't interoperate easily.

In the last few years, the technology and marketplace have made dramatic leaps forward. Foundation technologies in IP, optical, and wireless are evolving to broadband, and new architectures are emerging for carrier-grade service delivery and peer-to-peer communications. At the same time, end-user demand for broadband access is exploding.

To position for these opportunities, providers are taking advantage of today's advanced technologies to converge their various domain- and technology-specific networks onto one next-generation network – an intelligent packet network that provides the platform for mobile and personalized multimedia services. This converged network will enable providers to significantly reduce network costs, and create the platform for a dramatically different and much richer services environment.

Ultimately, network convergence will result in the "virtualization" of the network, whereby the underlying network will be invisible to users. Services will no longer be associated with a specific device or connection; rather, they will be accessed via a common applications interface and the network will figure out how to deliver a service to the user, regardless of whether access is via a cell phone, PDA, desktop, television, camera-phone, text messaging device, or some other type of device.

Drive down costs, reduce complexity

To achieve this vision, the network itself must be transformed. Specifically:

- the ties that today tightly integrate services and applications to network-specific equipment and devices must be separated to provide new flexibility and functionality;
- the boundaries between networks need to be removed to give users access to any service over any network environment;
- protocols and platforms must be open and standardized;
- the converged network must be flexible, scalable, open, carrier-grade, secure, multivendor, and services- and applications-aware; and
- the services themselves must be simple, convenient, and natural to use.

In making this transformation, one of the most urgent and significant priorities facing the industry today is to address the growing cost and complexity of networking.

The challenge, and opportunity, for the technology community is to innovate, simplify, and lower costs, enabling enterprises and carriers to, for instance:

- lower capital expenditures through scaling of a single network versus multiple networks – through use of lower-cost IP ports rather than TDM, ATM, or frame relay ports, and through deployment of service nodes where they are most cost-competitive;
- lower overall operations costs by reducing the need for network-specific operational support systems and staffing; and
- create new revenue streams, by developing and deploying new service capabilities that would have been either impossible or cost-prohibitive to deploy on previously separate networks.

Not surprisingly, approaches to the convergence challenge are varied, and different companies are defining the challenge in different ways. Some, for example, are adopting a more narrow definition ('convergence is about IP') to fit their current portfolios. Others, such as Nortel, recognized early on that convergence is much broader and that it can, and will, occur at many points in the network – between layers, services, networks, devices, and platforms.

Convergence requires an end-to-end approach

Indeed, bringing networks together is a complex process and requires extensive knowledge of different products and technologies – data, voice, optical, and wireless. For example, deploying Voice over IP (VoIP) requires, at a minimum, solid expertise in access technology, IP data networks, soft switches, gateways, end-to-end network engineering, and quality of service. For a solution to perform at the network level, it is not enough to specify each individual piece part. The converged network must be designed and engineered as a complete end-to-end solution, with clear network performance and behavior expectations.

This first issue of the *Nortel Technical Journal* highlights some of the activities under way – across all points in the network – that will help change operating cost structures and add significant new functionality and simplicity. For instance, we are:

- exploring different possibilities for how to architect and manage next-generation converged networks to achieve greater simplification, improved performance, and lower costs (page 3);
- driving packet cost points into the optical domain by collapsing Layers 0 to 3 to provide a more scalable, manageable, lower-cost, carrier-grade capability to deliver broadband services more reliably than just using IP (page 9);
- developing a common switching core between wireless and wireline, and bringing mobility and broadband to wireline and wireless customers using an access-independent converged core based on 3G IMS/MMD standards (page 14);
- working to converge on a few common software platforms by drastically reducing the number of software bases we use today, which will shorten interoperability testing and verification cycles, lead to greater reliability, drive significant cost reduction, and improve time to market (page 20);
- demonstrating ways to drive more intelligent networking by creating the ability for applications to more effectively use all of the network's capability, rather than just the transport pipe (page 23);
- focusing on convergence at the services edge – a first point of connectivity. With our Multimedia Communication Server (MCS), we are providing an integrated service environment for personal services (page 27);
- developing innovative engineering practices, now a part of our network planning and engineering process, to translate the end-user quality of experience (QoE) into measurable network parameters and targets (page 31); and
- simplifying network management on a modern converged network (page 37).

In addition, Nortel is one of the few vendors with a converged network operating in its own corporate network – using real, working products in an integrated environment. This network is saving millions of dollars each year in expenses, and is becoming a model network for our customers (page 42).

These initiatives represent only a handful of technology-focused activities that are currently under way, as we focus on creating the key convergence building blocks that can be assembled, via a solutions model, to give customers a variety of ways to make cost-effective, revenue-generating network convergence a competitive reality.

Peter Carbone *is leader of Strategic Planning for the office of the CTO organization.*
Steven Foster *leads the Business Planning team within the Network Strategy group of the CTO Office.*

Enterprise leading the way for future converged networking

by Philip Edholm

For the past 15 years, changes in technology, traffic, and computing have helped dramatically reduce the cost and complexity of enterprise networks. Three factors, in particular — inexpensive bandwidth, bursty data, and different traffic types — have combined to make local area networks very simple and user-independent at the core. This article anticipates that these same factors will transform carrier networks, by shifting the current balance of bandwidth, complex control, and management. The result will be a rethinking of how we architect and manage next-generation converged networks to achieve greater simplification, improved performance, and lower cost.

The advent of highly variable data-rate edge traffic of multiple classes over a consistent end-to end Ethernet/IP structure – together with the explosion of cost-effective bandwidth – will enable a fundamental redefinition of the relationships between complex and scarce network elements. The key now is to use these changes to simplify carrier networks, in much the same way as has already happened in enterprise networks.

In the past, the primary driver of communications systems was voice traffic. Today, the primary driver is data (Figure 1). As well, bandwidth has become relatively inexpensive – a change from the past where bandwidth was a significant limitation and required the design and development of complex systems to allocate expensive bandwidth to traffic flows. Furthermore, traffic is now more and more mesh-oriented, and the new data traffic is very bursty in its need for bandwidth.

This voice/data transition has fundamentally altered the traditional traffic model, as has a new class of traffic – termed "best-effort" – which is becoming a significant, if not dominant, part of the overall traffic.

Prior to the bandwidth explosion that began in 2000, networks were operated to maximize bandwidth fill. Significant activities were undertaken to ensure that, wherever possible, networks were operating at 80% to 90% of capacity, in order to maximize the return on investments in scarce resources. Starting in 2000, however, the bandwidth explosion began to change fill requirements. As the cost structures of dense wavelength division multiplexing (DWDM), 10-Gigabit Ethernet, and next-generation Multi-Protocol Label Switching (MPLS) systems continue to impact the network, the cost of increased bandwidth is decreasing dramatically.

> By taking advantage of key changes — inexpensive bandwidth, bursty data, and different traffic types — it may no longer be necessary to manage networks as end-to-end connections.

As we move toward 2010, the gap between network capacity and demand is increasing, or will increase, where such new technologies as Optical Ethernet are available – with important implications for the network. When average utilization in the LAN is less than 30% – in the typical LAN, average utilization is less than 10% of available bandwidth – the process of redundancy and backup engineering is dramatically simplified. For example, in the LAN, Nortel's MLT (MultiLink Trunking) technology enables simple active-active redundancy and failure reconfiguration without having to engineer for capacity.

These changes – inexpensive bandwidth, bursty data, and different traffic types – are driving a rethinking of how applications and services can be delivered and provisioned, and how the network can be managed.

Systems such as Asynchronous Transfer Mode (ATM) and Resource ReSerVation Protocol (RSVP), for instance, were defined when the relative balance of traffic types and bandwidth fill required extensive management and reservations to deliver cost-effective solutions.

The future will be dramatically different. By properly classifying the different traffic types, which have different requirements, and then managing the operational parameters of the infrastructure, we can easily deliver a service level agreement (SLA) that meets the customer requirements.

For example, by assigning the small percentage of voice traffic to a traffic class that has absolute priority of transmission and no discard, the voice traffic will be transmitted through the network without any impact on the volume of traffic in other classes. As Figure 1 shows, the emergence of, first, mission-critical data and, next, best-effort data – combined with the cost effectiveness of operating the network at a lower fill percentage – will make the management of diverse services easy.

This ability to improve service management is evident if we look at the trend over the past five years in enterprise campus LANs, where exactly these mechanisms for classification and management using simple Class of Service (CoS) without guarantees is standard. The failure of ATM and RSVP in campus LANs is a result of the ease of managing multiple service types (including "real-time") over a network that has over-provisioned bandwidth headroom and uses simple CoS classification and queue management operation. As shown in Figure 1, the advent of simple Differentiated Services (DiffServ) in the Internet Engineering Task Force (IETF) anticipated these changes.

Figure 1. Projection of traffic types

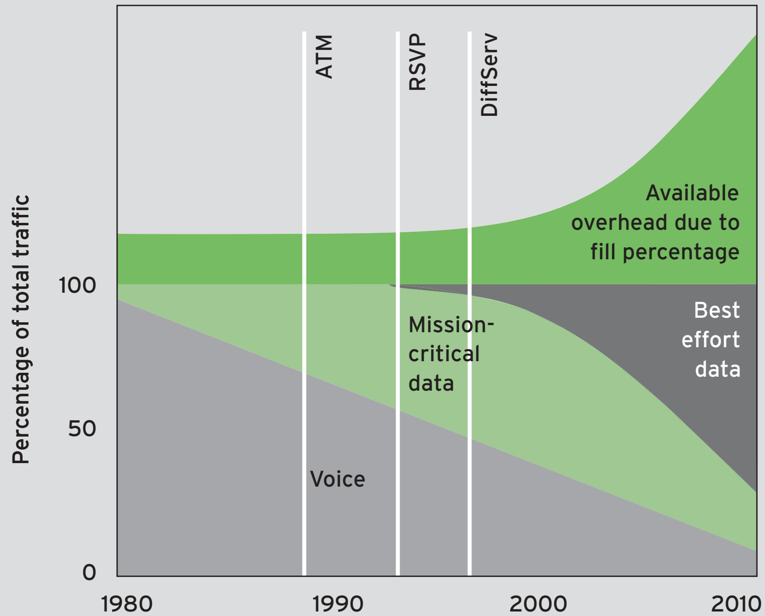

Inexpensive bandwidth, bursty data, and new traffic types have fundamentally altered the traditional traffic model. Data is now the primary driver of communications systems. In 1980, for example, 98% of all traffic was voice; by 2010, less than 10% of traffic will be voice, including videoconferencing. At the same time, bandwidth has become relatively inexpensive, more traffic is mesh-oriented (peer networking), and the new data traffic is very bursty in its need for bandwidth. As well, prior to 1995, virtually all data traffic was mission-critical. The advent, in 1995, of the web browser, e-mail, and the WWW changed that paradigm. Casual web browsing and attendant bandwidth uses created a new class of traffic commonly known as "best effort," fast becoming a significant if not dominant part of overall traffic. The combination of the emergence of, first, mission-critical data and, next, best-effort data, when combined with the cost effectiveness of operating the network at a lower fill percentage, will make the management of diverse services easy.

Reducing complexity

By taking advantage of these changes, it may no longer be necessary to manage networks as end-to-end connections. In fact, managing the network as an access edge and a separate core will be a major change that will be enabled.

Figure 2 shows the demarcation of the customer edge, the provider edge, and the transport core. At the edge of the carrier network, there must be a User Network Interface (UNI) that defines the service characteristics on the interface to the customer. The UNI defines the forwarding mechanism (Layer 2 or Layer 3) and the SLA characteristics (bandwidth, delivery and latency guarantees, and availability). The SLA must be managed and provisioned on an individual port/customer basis. However, once the traffic has left the UNI and entered the transport core, traffic can be managed using simple techniques.

As shown in Figure 3, this approach creates two distinctive management systems. The provider edge system configures and manages the UNIs based on the services purchased and the associated SLAs. Traffic is classified according to the SLA and input to the transport core. Within the transport core, individual flows are not the critical element; rather, the aggregate rates in any given class on any given link become the key criteria.

Therefore, while the edge is managed based on individual customers, the core is managed based on link operation and predictive intelligence that will anticipate resource utilization. Management may include systems that can predict the impact of a new customer entering the network based on location and SLA parameters, enabling checks of operational impact prior to allowing the new customer on the network or allowing requests to change SLA parameters.

A significant advantage of this approach is that, while each customer entity is managed at the edge, the core is managed as an overall service, eliminating the need to manage each path on an end-to-end basis. Again, experience in the enterprise LAN has shown that this management approach is practicable in the presence of reasonable bandwidth. Furthermore, this approach can reduce the complexity of provisioning and managing a network by two orders of magnitude.

Protocol convergence

While building the next generation of networks will involve rethinking assumptions, we must do this within a framework of resource allocation. One critical element of this rethinking is the relationship between metro and long-haul networks. The emergence of MPLS and GMPLS (Generalized Multi-Protocol Label Switching), along with high-speed optical, appears to define the core of next-generation networks. The challenge is how to intelligently connect these cores to the customers. Optical Ethernet, along with delivering services – either Layer 2 or Layer 3 – over Ethernet connections seems to be the preferred solution for next-generation edges.

In the enterprise, this solution allows traffic, which is 98-99% Ethernet in the campus today, to remain Ethernet end to end. At the same time, packet-processing decisions will be made at multiple layers.

In the highly divergent networks of the past, where the layers implemented different standards and protocols (Figure 4), systems had to be built separately in order to perform functions at each layer. As protocols and standards consolidate, decisions can be made simultaneously between the layers. For example, a forwarding decision can be made at Layer 2, while a quality of service (QoS) function can be performed at Layer 3 using DiffServ.

This blurring of the layers dramatically simplifies networks and will lead to the next generation of products – the specialized system. When we had to match and manage complex dissimilar protocols, the key was devices that could deal with the maximum in protocol and interface diversity. As we move forward, specialized devices – optimized for a specific point in the network with functional requirements limited to functions in Ethernet and IP at that point – will dominate. This shift

Figure 2. Managing the network as a separate access and core

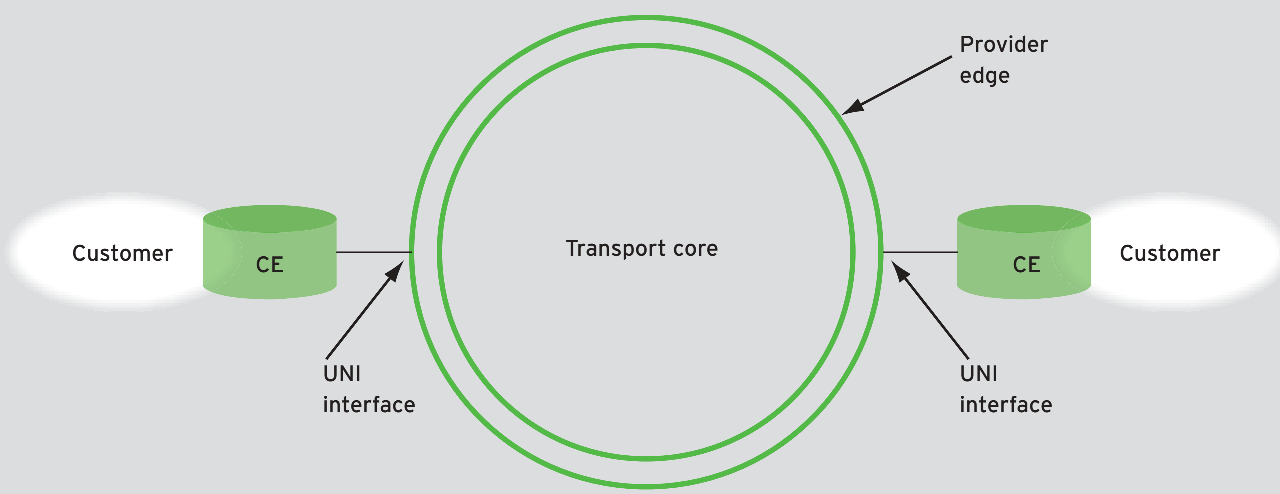

By taking advantage of bandwidth and traffic changes, it may no longer be necessary to manage the network as end-to-end connections. Experience in enterprise networks has shown that managing the network as an access edge and a separate core can help reduce the complexity of provisioning and managing a network by two orders of magnitude.

has already happened in the enterprise campus and it is now happening in such areas as VPN nodes and web switching.

Bandwidth for architectural advantage

A key challenge in reducing complexity is to allocate the intelligence to deliver services in a way that optimizes both cost and speed. In this context, cost involves both acquisition costs (CapEx) and ongoing operational costs (OpEx). As discussed earlier, trading bandwidth for complexity can significantly reduce operational costs.

With these concepts in mind, an architecture for the future can be postulated that distinctly separates the metro transport and the long-haul transport domains.

Within the metro domain, an Ethernet transport network yields significant benefits, because it enables bandwidth to minimize the complexity of path and QoS decisions at the low-cost points generated from the enterprise LAN segment. In the long-haul environment beyond the metro, however, the simplicity of the Ethernet transport mechanism does not offer the necessary benefits of scale and traffic management.

Figure 3. Two traffic management systems

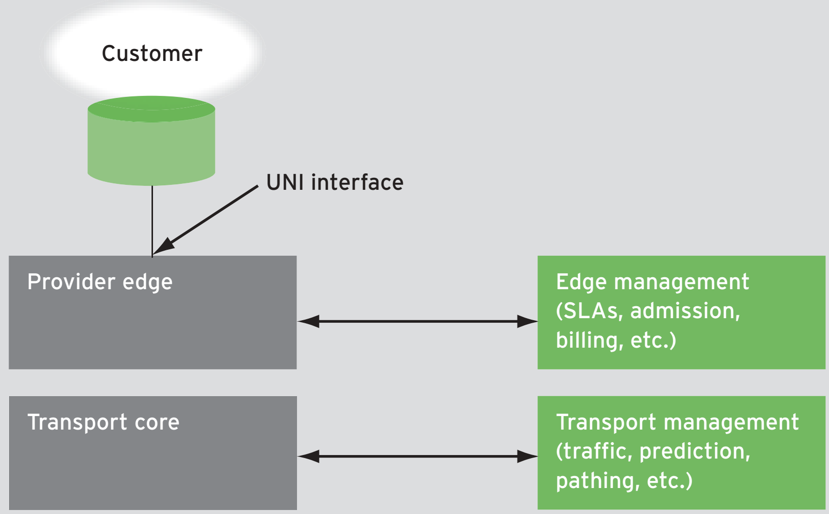

Managing the network as an access network and separate core creates two distinctive management systems: the provider edge system; and the transport core.

Figure 5 outlines an architecture that optimizes all necessary elements. Service management is at the provider edge, where it belongs, both managing actual service delivery performance and enabling SLAs that can deal with the very bursty nature of data. At the same time, the metro transport network is optimized for speed without being burdened by the complexity and intelligence necessary to provide large-scale national networks. The "smart" layer at the interface between the metro and long-haul networks translates from the simplicity and

Figure 4. Protocol convergence

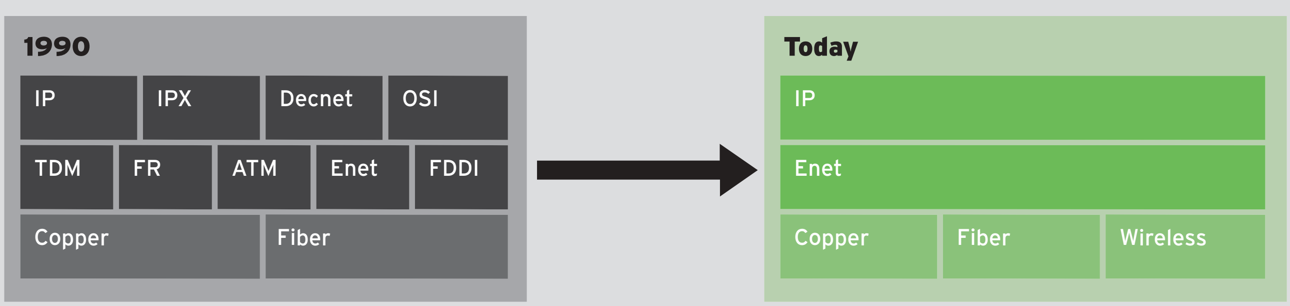

The explosion in the 1990s of creativity and invention in networks created a large number of competitive (or at least co-existent) protocols and media/control. During the late 1990s, organizations converged on IP as the common Layer 3 protocol, and within enterprise campuses Ethernet became the dominant solution. The same convergence is happening at Layer 2 where Ethernet is fast becoming the converged standard. This protocol convergence is based on enterprise preference, the end-to-end values of commonality, and component costs, and has a significant impact on the way we process and move data.

Figure 5. Architecture for the future

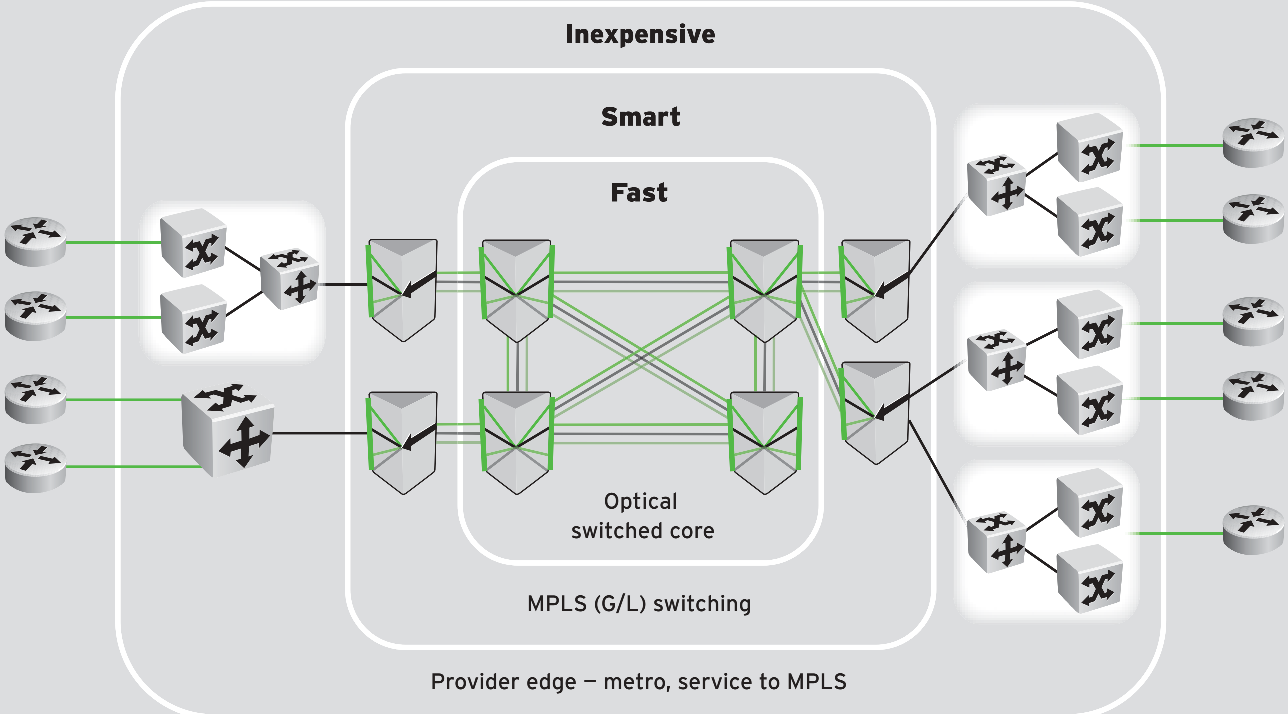

This diagram depicts an architecture for the future that distinctly separates the metro transport and the long-haul transport domains. In this architecture, service management is at the provider edge, while the metro transport network is optimized for speed without being burdened by the complexity and intelligence necessary to provide large-scale national networks. A "smart" layer lies at the interface between the metro and long-haul, with the intelligence necessary to deal with the potentially thousands of paths and connections necessary to build such networks. The "fast" layer, or the long-haul optical switched core, is based on a combination of MPLS and optical.

cost-effectiveness of the metro to the more complex and scalable core. This layer has the intelligence to deal with the potentially thousands of paths and connections necessary to build such networks. Finally, the long-haul core is very fast, based on a combination of MPLS and optical. By leveraging the underlying SONET/SDH transport, not as a circuit layer but as an optical bandwidth layer, we can achieve a new paradigm that provides the high-speed capabilities needed to minimize complexity at reasonable cost.

The benefits of such an architecture are significant. In the metro, it is possible to deploy low-cost devices and inexpensive transports derived from the millions of Ethernet enterprise ports and devices. The edge devices can do the policing and management for the SLAs, including customer separation through simple encapsulation and auto-discovery of paths to other nodes.

Between the metro and long-haul environments, traffic is mapped into MPLS – the logical alternative because MPLS enables traffic engineering and management across the core for new services, while enabling existing services to coexist through GMPLS.

A large number of carriers are committing to MPLS as the way to manage core backbones. [The MPLS Label Switched Paths (LSPs) need only be configured between metro environments because, with traffic already separated by customer, these paths can be shared among multiple customers.]

The impact of this change in path management is striking – instead of having to build LSPs either for each customer connection or between all edge devices, LSPs are required only between the metro core devices. As the number of paths and complexity expand exponentially with the number of points (10 interconnected points require 10x9 or 90 paths; 100 interconnected points require 99x100 or 9,900, or 100 times as many), moving this point to the metro/long-haul edge and having 20 to 50 metro edge devices connected reduces the complexity by 100 to 10,000 times.

With much higher bandwidth throughout the network and varied traffic types, the network can be managed as an edge admission network and a core transport. As discussed earlier, the edge can be managed for SLAs and customer requirements, while core links can be sized and managed for presented traffic, and the different traffic types can manage any load variation on a real-time basis. The result is a network that is very easy to manage and provision, both at the customer and the overall transport level.

Conclusions

While the network transformations anticipated in this article will not happen overnight, they are in fact beginning. On enterprise campuses, a bandwidth explosion has already transformed networks, and DiffServ-based CoS is used for IP telephony and other forms of real-time traffic. Optical Ethernet, being delivered in some markets and under evaluation by virtually every carrier, will extend this bandwidth into the metro environment. Finally, the capacity explosion from DWDM and the emergence of MPLS/GMPLS will deliver the core long-haul. While this solution will be packet-based, it will use the underlying SONET/SDH infrastructure for optical bandwidth, provisioning, and management.

Nortel is well-positioned to drive this transformation, with our work in optical transport, Optical Ethernet, and the Multiservice Provider Edge, as well as our ability to build new networks and capabilities that capitalize on this network transition. The key will be working with our enterprise and service provider customers to help them move through the transformation.

Philip Edholm *is Enterprise CTO and vice-president of network architecture for Enterprise Networks.*

Technology for optical/packet convergence

by Dave Hudson, Paul Littlewood, and Chris Chartrand

Today's carrier networks comprise multiple service-specific overlay solutions based on different networking technologies — a mix that is inefficient, costly, and inflexible to changing service demands. A converged edge and core network based on recent technology advancements in optical and packet data networking will enable simpler and more agile networks that support multiple services and multiple applications on single platforms. At Nortel, we are supporting customers in this transition to convergence with the Multiservice Provider Edge 9000 (MPE) and the Optical Multiservice Edge 6500 (OME). While MPE and OME on their own have many independent deployment and application scenarios, this article discusses the commonalities in their design intent and technology choices, as well as how they could be deployed together for a packet-optimized metro access and converged services edge/core infrastructure.

In order to meet customer demand, carriers have deployed multiple overlay networks that were originally designed to carry service-specific traffic – such as TDM for traditional voice, frame relay for business data, and IP for Internet traffic. The access and transport networks that support these, as well as traditional leased-line services, are largely based on installed SONET/SDH equipment.

The mix of network traffic, however, is changing from TDM to voice – or even multimedia – over IP, while data connectivity services, such as VPNs, metro Ethernet, and broadband access, are experiencing tremendous growth. The result is a large and increasing amount of packet-based traffic running over an inefficient mix of existing single-service networks. Of even more concern to service providers is that although data traffic is growing, revenue per bit is declining and will likely continue to do so at an increasing rate.

This reality leaves carriers facing an industry that is in an unprecedented state of flux and under continual competitive attack. To defend their turf and meet the increasingly complex needs of their customers, carriers must simplify network operations, reduce capital costs, and make it easier to deploy and bundle new services. They are looking for simplification through the integration of multiple functions onto single products and tight optical/packet functional interoperation, both commonly known as convergence.

Specifically, today's carriers are looking for the flexibility to quickly increase or redeploy capacity in their networks to minimize the time and cost it takes not only to add new services, but also to add value to existing services or even to "sunset" services that are no longer economically viable. They are also seeking an efficient way to knit together emerging technologies with a variety of legacy technologies that aren't about to disappear overnight – a situation that likely will become even more complicated as mergers, acquisitions, and inter-carrier partnerships bring together what are often very different network systems and operational procedures (Figure 1).

As a result, service providers are now starting to specify solutions that integrate several networking layer functions into single products that simplify and reduce the number of nodes and networking protocols. We are seeing not only a drive to converge many legacy networks onto a common infrastructure, but also a high degree of commonality in how functions are distributed around the network in order to facilitate the introduction of new services across multiple access technologies. For example, service providers are looking to converge service adaptation, aggregation, and grooming into packet-enabled transport network elements (NEs) that can be deployed at a number of locations across networks.

These requirements for flexibility and integration are reflected in the technology choices we made in designing two key Nortel products – the Multiservice Provider Edge 9000 (MPE), and the Optical Multiservice Edge 6500 (OME):

- The MPE 9000 collapses multiple service-specific network overlays between Layer 2 (frame relay, ATM, Ethernet) and Layer 3 (MPLS/IP) of the OSI stack.
- The OME 6500 is a flexible SONET/SDH transport network element that collapses multiple service-specific overlays among Layer 0 (photonic/transparent), Layer 1 (SONET/ SDH), and Layer 2 (packet/Ethernet). The OME is also designed for multiple applications throughout metropolitan networks, enabling service providers

to collapse the number of NE types they need.

In this context, the MPE, defined to be at the packet services edge, can be thought of as the point of service unification and personalization, and is the gateway into the wide area network (WAN). The OME can be thought of as a feeder or fan-in device to the MPE, efficiently switching and aggregating packet and circuit traffic (Figure 2). The key requirements in these products to support convergence are a very high availability and maintainability, a flexible architecture, and ease of functional enhancement.

The MPE and OME represent new classes of highly integrated products capable of supporting and interworking multiple network protocols and supporting multiple applications in the network. Used in combination, they provide a very efficient and scalable means to build services-rich networks that offer optical broadband, private line, Ethernet, virtual private wire, and IP services. A key attribute is the ability not only to support traditional circuit and packet architectures, but also to migrate toward carrier-class Ethernet as the emerging Layer 2 of choice (see sidebar on page 12).

Technology enablers

In building these products, we based our hardware and software designs on the latest industry standards, and on the innovative application of leading-edge technologies, including:

- a flexible, expandable, and robust software architecture;
- a hardware architecture that provides rigorous segregation of functions to enable independent evolution and enhanced resiliency;
- programmable data processing and transport processing components, such as network processor units (NPUs) and field programmable gate arrays (FPGAs); and
- flexible optical I/O based on small form-factor pluggable (SFP) interfaces.

On the software side, modularity is the key, enabling the platforms to be effectively used in multiple applications or markets. In addition, a

Figure 1. Today's network complexity

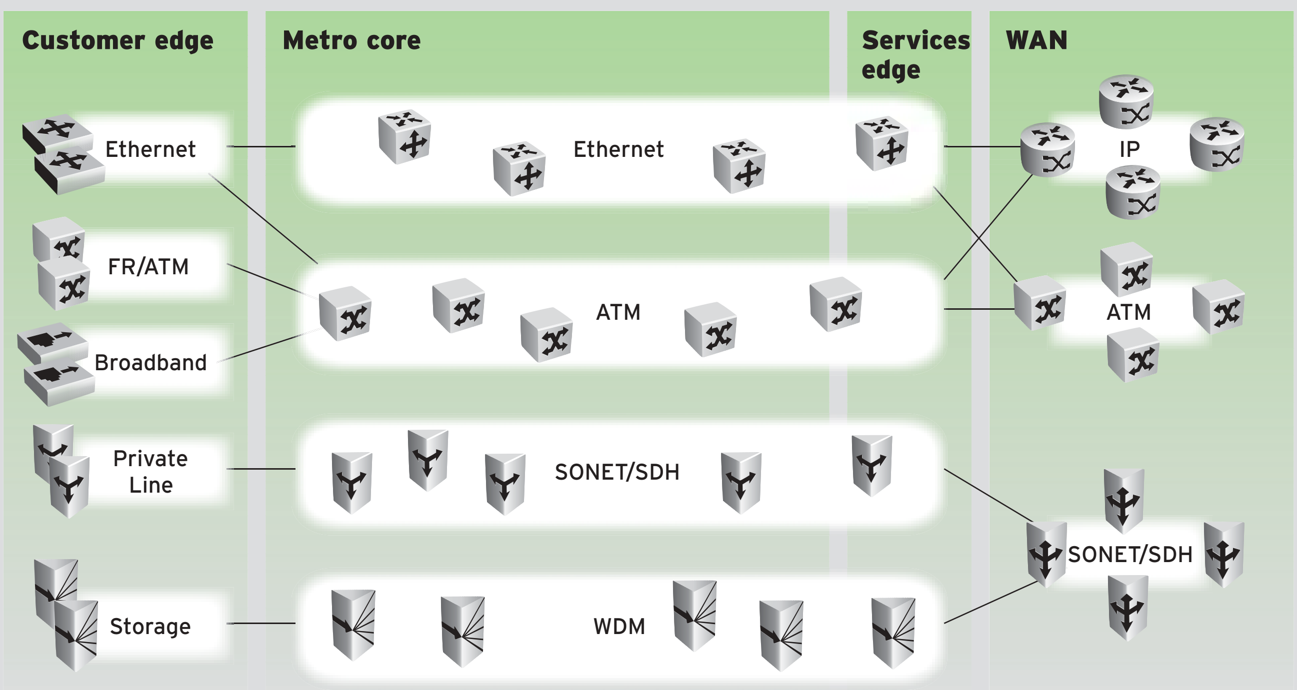

Service providers have responded to the demand for new service types by building new networks of dedicated capability to carry traffic picked up at the customer edge to equipment at the services edge. For example, service providers currently offer Optical Ethernet services over an overlaid infrastructure of interconnected Ethernet switches. Frame relay (FR) or ATM services require another separate network. Private line services are supported by the SONET/SDH network, and storage connectivity and managed wavelength services are typically supported on a WDM underlay. Overlays have also been used for switching in long-distance wide area networks (WANs), although today they are being converged onto Multi-Protocol Label Switching (MPLS) networks. These parallel, defined-purpose networks increase operating costs for service providers and create an unwieldy architecture that makes it difficult to bundle services and respond quickly to changing customer or market demand.

modular software architecture provides partitioning and protection for independent software processes. By confining software upgrades or problems to one process while leaving other processes up and running, our services platform is much more robust, providing much higher levels of service and product availability.

Our modular software implementation also enables logical routers and software sparing that ensures fairness, logical isolation, and protection between customers, services, and processes – attributes that are especially important in today's world of on-line service attacks, viruses propagating through the web, and other unforeseen events. These attributes not only provide a measure of security, but they also give service providers the tools to converge services that have vastly different characteristics on the same network or node without fear of having low-profit, low-priority services impact premium services.

From a hardware perspective, the OME architecture employs a common bus and backplane structure for each line card slot, which enables any card slot to be used with any line card. This architecture ensures the platform will remain evergreen as network requirements evolve. Similarly, the MPE's mid-plane architecture provides complete separation of control plane and data plane processing and traffic from the line cards, allowing all cards to be swapped out without disrupting service. Since any port (or interface) can be coupled to any service, flexible line and equipment sparing configurations are also enabled – even among cards of different types and capacities. This capability helps convergence by eliminating stranded investments and by simplifying service migration and interworking.

The use of programmable components – NPUs in the MPE and OME, and FPGAs in the OME – is critical to the application and service flexibility we are trying to provide. These reprogrammable components enable service parameters to be changed

Figure 2. Convergence architecture

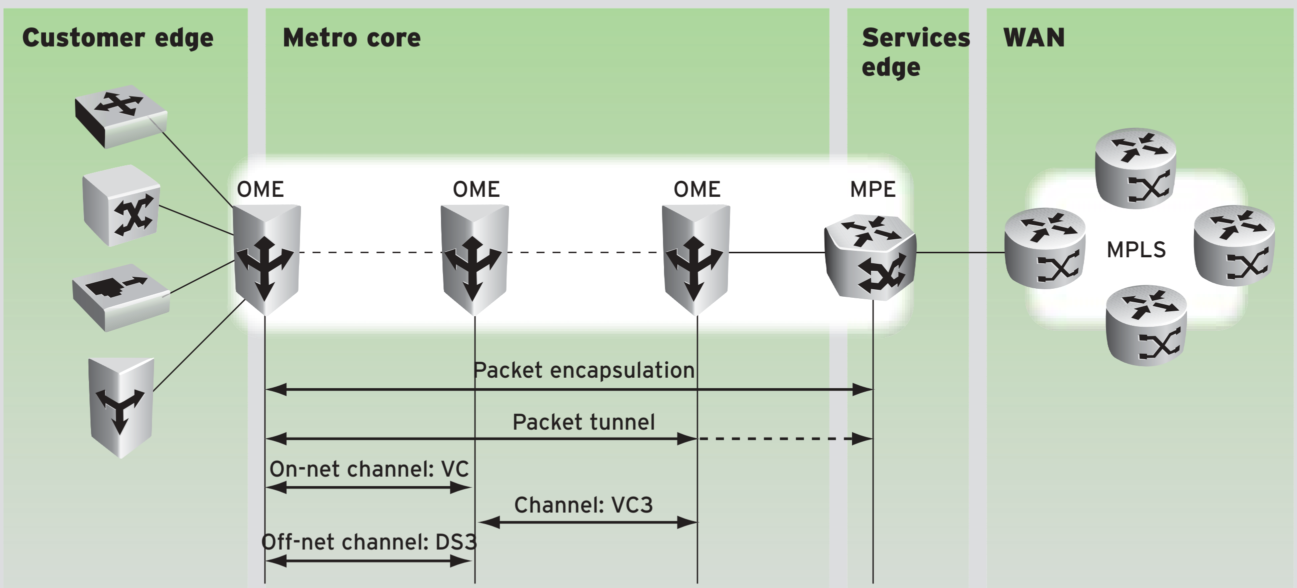

Multiservice Provider Edge (MPE) and Optical Multiservice Edge (OME) are designed for multiservice networks supporting current and future services. Although services will terminate directly on the MPE, for those services accessing the MPE over an optical network, careful selection of technologies supporting secure aggregation of traffic is required. A standard packet encapsulation, such as VLAN and/or Pseudowire Emulation (PWE), will be used for both Ethernet and non-Ethernet payloads. The encapsulation provides separation, service OAM, and security. Networking is accomplished using a packet tunnel, such as an Ethernet Virtual Circuit or MPLS label switched path, which connects the customer-located equipment and the appropriate MPE at the services edge.

Compatibility with SONET and SDH networks is achieved by mapping packet tunnels into channels, such as SONET/SDH virtual containers (VCs) or DS3s (for off-net access). OMEs in the metro core switch and aggregate the packet tunnels from the customer edge to the services edge. At the services edge, packet tunnels are extracted from SONET/SDH and are mapped to Ethernet connections between the OME and the MPE.

Ethernet as a convergence technology

by Liam Casey

The low cost, simplicity, and pervasiveness of Ethernet – the dominant protocol for data connections in the local area network (LAN) – make it an attractive technology for converged optical/packet solutions.

Nortel has a strong history of involvement in Ethernet technology. We were one of the founding members of the Metro Ethernet Forum – an industry consortium promoting the adoption of Optical Ethernet as the technology of choice in metro networks worldwide. We were also the first to prototype 10 GigE and the first to develop an Optical Ethernet product – the Optical Multiservice Edge (OME).

From a technical perspective, Ethernet can be used in different ways in different convergence areas, exploiting specific Ethernet properties. The sections below briefly discuss our view of both the advantages and the challenges of using Ethernet in each of these different convergence areas.

Broadband physical interface

The ubiquity of Ethernet has ensured that it is the least expensive broadband interface for user access devices. In fact, no matter what the service, almost all packet devices today have an Ethernet 10/100Base-T physical interface. As a result, Ethernet has become the *de facto* standard for connecting to the converged network.

Using Ethernet as the interface between a customer and a carrier, however, does present some challenges. Ethernet traffic management, for example, is very weak and today no standard User Network Interface (UNI) maintenance protocol exists for Ethernet. A specific traffic management challenge is that Ethernet is implemented in customer equipment with no expectation of requirements to shape packets to a service rate less than the interface line rate, even though metro Ethernet service definitions assume this can be done.

Optical link technology

In converged networks where all services are transported as packets, Ethernet offers a simpler and cheaper alternative to SONET/SDH as the optical link. Where once it was assumed Ethernet was for use in LANs inside buildings, 100Base-FX and 1000Base-LX Ethernet links can now use optics that work over distances of up to 80 kilometers (50 miles). On the plus side, Ethernet defines its own framing and error detection, while ATM and frame relay rely on SONET/SDH to do this for them, meaning they cannot operate without a SONET/SDH layer. Furthermore, the plug-and-play nature of Ethernet means that networks self-configure and do not require extensive and expensive configuration provisioning like SONET/SDH systems.

While the primary reason for using Ethernet optics is for the cost reduction enabled by convergence, maintenance challenges must be solved before deployment moves beyond the edge of the metro. SONET/SDH, for instance, carries path and span overhead in each frame, which simplifies fault detection and diagnostics. While 802.3ah OAM standards have defined similar capabilities for Ethernet, these are limited to Ethernet First Mile (EFM) links. Nortel is proposing that these standards be extended into metro networks.

Common Layer 2 packet transport

IP packets can be carried over any Layer 2 transport network. Until recently, however, frame relay and ATM were viewed as the preferred technologies for transporting IP packets in the metro and wide area networks (WANs), while Ethernet was the overwhelming choice in enterprise LANs. Today, however, increased bandwidth available in the metro network has eliminated the rationale for not using Ethernet as the Layer 2 transport everywhere. When link bandwidths are greater than 10 Mbit/s, the 14- to 18-byte header of Ethernet is not nearly as limiting as it is on the low-speed 56-kbit/s links typical of frame relay deployment. In addition, while ATM is required on asynchronous digital subscriber loop (ADSL) uplinks (typically 340 kbit/s) to reduce the jitter of voice packets interleaved with data packets, the jitter from full Ethernet packet interleaving at very high speed digital subscriber loop (VDSL) and optical speeds is not a problem.

As well as being the most widely understood Layer 2 packet transport, Ethernet brings resiliency options that are not available with ATM and frame relay. MultiLink Trunking (MLT), defined in 802.3ad, and split MLT, a Nortel value-add, allow for load sharing across Ethernet links, improving network reliability.

However, as Ethernet moves from the LAN environment to access and metro packet networks, separation of customer traffic is required. Standard Ethernet, therefore, needs to be augmented with some form of customer label, while maintaining transparency to customer traffic. Although there are different industry approaches, Nortel is a leading advocate of MAC in MAC encapsulation (MAC stands for Media Access Control), where the customer Ethernet packet is encapsulated in a provider's Ethernet packet to provide this separation.

Ethernet as a switching technology

Standardization of Ethernet frame formats and forwarding behavior has made it very attractive for chip vendors to produce complete Ethernet switches on a chip, unlike other protocols, where packet switching and forwarding requirements require a combination of more expensive Network Processor Unit (NPU) and fabric chips. Using Ethernet-formatted packets in the metro should enable the use of lower-cost, higher-capacity network elements.

These Ethernet chips, however, have been designed to switch according to the 802.1 D and Q standards applicable to LANs, and this is often not the form of switching required in metro networks.

Ethernet services

One question swirling around packetization and Layer 0-2 convergence is "what will replace the TDM private line?" The Metro Ethernet Forum *(http://www.metroethernetforum.org/)* envisages that Ethernet Private Line (EPL) and Ethernet Virtual Private Line (EVPL) will be the basic next-generation transport services. EPL primarily involves adapting Ethernet packets for transport over a photonic or SONET/SDH path and, as such, would be used where a single customer can fill an optical fiber pipe. EVPL, on the other hand, has additional capabilities and would be used where multiple customers need to share a transport pipe in an efficient and secure manner.

Liam Casey *is architect for the converged core, working in the Office of the CTO.*

without having to change the hardware. In the MPE, any service can be supported on any port or any channel, allowing service providers to quickly provision new services alongside legacy services without having to accurately forecast the demand for cards and ports. Interface, bandwidth, protocol, or service of choice can be offered on a site-specific or application-specific basis, with easy migration between technologies (from frame relay to Ethernet or IP, for example).

With standards in the convergence area still evolving, it is important to employ flexible transport processing components in order to achieve quick time to market. That was one of the reasons we chose FPGAs to do much of the traffic processing in the OME. For example, GFP (Generic Framing Protocol) is an important technology that enables new service convergence onto existing infrastructures by providing efficient transport for Ethernet and Fibre Channel OBS (Optical Broadband Services) over SONET/SDH. In order to take an early lead in this emerging market, we introduced a version of GFP ahead of the standards, which have only recently been ratified by the International Telecommunication Union (ITU). Because we are using programmable FPGAs, we are able to adapt the existing hardware as the standards are refined.

In the OME, NPUs were also used in architecting the packet switch cards [Layer 2 Service Switch (L2SS)] to allow for functional flexibility as customers adapt and modify their networks. The L2SS will be used to not only fulfill such transport applications as aggregation of upstream traffic for bandwidth and port-fill efficiency, but also to support Ethernet services as defined by the Metro Ethernet Forum (MEF) industry consortium (see sidebar).

This flexibility and future-proofing is crucial in the area of metro Ethernet. Carrier plans in the metro are increasingly turning toward Ethernet for the next generation of network build, and standardization is under way to add the requisite features to Ethernet to address the challenge. This includes well-specified OAM, auto-discovery, and hierarchy (see sidebar).

Another technology choice – small form-factor pluggable (SFP) interfaces – delivers optical rate and reach flexibility on a per-port basis. SFP technology enables service providers to contain costs by reducing higher-priced spare cards inventory and deploying only the I/O needed at that time. Additional SFP modules can be added later in response to service demand. Pluggable technology is important in convergence to tighten the degree of physical integration between wavelength division multiplexing (WDM) and optoelectronic NEs. The flexibility to selectively deploy a range of wavelengths on a card reduces both sparing requirements and the likelihood of stranding bandwidth over time.

By incorporating these technologies in our latest generation of products, such as the MPE and OME, we will be able to provide the required breadth of capability and the functional and service flexibility that operators need to offer new services and realize operating efficiencies – while maintaining the familiar carrier-grade reliability that users will continue to demand from future services.

Dave Hudson *is Leader, Wireline & Optical Product Strategy.*
Paul Littlewood *is Network Architect, Wireline & Optical Architecture.*
Chris Chartrand *is Senior Product Manager, Services Edge Portfolio, Wireline & Optical.*

Wireline and wireless operators converging on 3G IMS/MMD architecture

by Liam Casey

With both wireline and wireless operators aligning on the 3G IP Multimedia Subsystem (IMS)/Multimedia Domain (MMD) network architecture, the way is now clear to deliver a personalized set of service features and call behaviors to individual subscribers at a single number, no matter what type of device they use to access the network. When fully defined, an extended IMS/MMD architecture will enable convergence on a common wireline/wireless switching core, and services delivery across the two domains. Nortel, through its Converged Multimedia Services (CMS) thrust, is developing the key network elements and capabilities necessary to implement this IMS/MMD architecture.

Call it serendipity. Wireline carriers – seeing their revenue streams under attack from cellular, Internet, cable, and emerging VoIP providers – are searching for ways to retain customers. Wireless operators, meanwhile, looking to add new broadband multimedia services to boost their stalling revenue growth, have already begun work on developing an open, industry-wide, standards-based multimedia services delivery framework. This work is being done in the 3rd Generation Partnership Project (3GPP) and has been labeled IP Multimedia Subsystem, or IMS. 3GPP covers GSM, GPRS, and UMTS wireless networks, while another partnership, 3GPP2, covers CDMA wireless networks. (For more on 3GPP, see page 16.) For the rest of the article, we will refer to IMS, but it can be assumed the discussion also incudes the 3GPP2 equivalent, called MMD (Multimedia Domain).

Seeking a multimedia communications offer to solve their customer retention challenge, wireline operators realized that the IMS strategy already embarked upon by wireless providers offered a suitable framework for the scalable, distributed, IP-based environment they would need to deliver multimedia services to intelligent subscriber devices anywhere in the network.

This fortuitous confluence of wireless and wireline interests around the IMS architecture makes it possible to converge on an access-agnostic common wireless/wireline switching core. Once in place, this converged core will provide a common services environment with a powerful set of service enablers that will allow operators to rapidly develop and deploy new revenue-generating multimedia services and deliver a seamless user experience across the wireline and wireless domains.

Nortel's thrust to develop a converged IMS for both wireless and wireline applications is called the Converged Multimedia Services (CMS) solution. In CMS, our intent is to provide an end-to-end IMS solution, with in-house development around the core competencies of session control, database, media control, and interworking with current legacy voice (TDM) and VoIP networks, both wireless and wireline. We are also actively participating in Next Generation Networks (NGN) standards working groups to incorporate into the wireless-initiated IMS standards the capabilities we determine are needed for wireline operation.

Nortel's thrust to develop a converged IMS for both wireless and wireline applications is called the Converged Multimedia Services (CMS) solution. In CMS, our intent is to provide an end-to-end IMS solution.

The new communications environment

Stepping back, we can see that the IMS standardization comes at an opportune time to exploit a number of technology developments that are already well under way. These include:

- *Ubiquitous IP* – An underlying IP packet transport network can accommodate different multimedia streams that require substantially different bandwidths. By contrast, the voice telephony network uses a single uniform 64-kbit/s bandwidth for media streams, making it an impractical solution for these new multimedia services.

• *Wireless personal terminals* – Wireless phones and a plethora of personal devices (such as PDAs, gaming machines, and BlackBerry devices) that have had wireless or WLAN communications added to them have changed the telephony model. A wireless phone is a personal phone – people can use it and are reachable wherever they can get coverage. The PSTN, on the other hand, is not designed to exploit the opportunities for personalization that this development provides.

• *Affordable terminal displays* – Today's end devices (such as cell phones, BlackBerry devices, and PDAs) have not only powerful processing capabilities but also graphical display screens to enable features and applications, such as presence and location, that are impossible to provide using the man-machine interface of a POTS phone.

IMS architecture

The IMS architecture is not only about new multimedia services, but also about creating a more powerful and uniform framework for delivering such services. The IMS architecture covers all aspects of session control, application invocation, roaming, policy management, and security (Figure 1). Briefly, the key building blocks include:

• *Session Management (Call Control)* – IMS uses Session Initiation Protocol (SIP) to provide call establishment, call participant management, and feature invocation. SIP is the key technology that integrates media streams, applications, and clients (terminals) into a

Figure 1. Major components of the IMS architecture

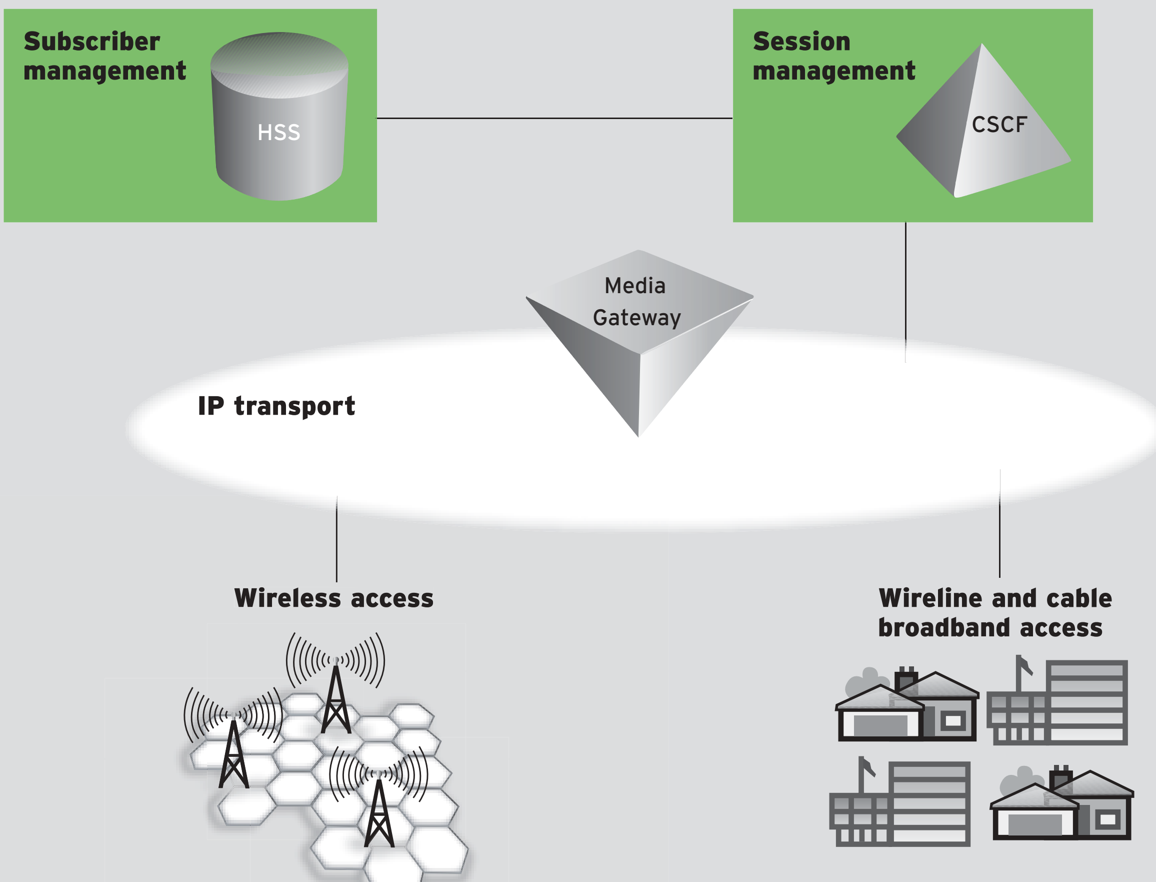

The IMS architecture is key to the convergence of wireline and wireless networks. The main functions covered by the IMS architecture are: Subscriber Management, which enables personalized services and identity management using subscriber profiles maintained in the Home Subscriber Server (HSS); Session Management, performed by Call State Control Functions (CSCFs), which handle SIP signaling for call establishment, call participant management, and feature invocation; and Media Gateways, which manipulate and transport media streams to and from access networks (wireless, and broadband wireline and cable). All of the above are interconnected by an IP transport network.

single, unified services model that is much more powerful than that of the telephony network. Moreover, SIP supports the establishment of bidirectional media streams between users and between an application server and user. Unlike web-style invocation of an application server, which must be client initiated, SIP enables push services, where an application server makes the call to a client.

In the IMS architecture, SIP signaling is handled by functional entities called Call Session Control Functions (CSCFs). To accommodate roaming and internetwork operation, CSCFs can assume different roles, including: serving CSCFs, which perform the basic session originating or terminating treatment for SIP messages; proxy CSCFs, whose main job is to relay SIP signaling between a user's terminal and a serving CSCF, and provide the access layer abstraction; and functions, such as a Breakout Gateway Control Function and Media Gateway Control Function, for interworking between IMS and the legacy TDM and VoIP worlds. (For more on SIP, see page 29.)

- *Personalization and Identity Management* – In keeping with the trend toward the use of personal communication devices, IMS changes the location-to-location communication model of the PSTN to a person-to-person communication model. Within the IMS architecture, the Home Subscriber Server (HSS) is the key repository for subscriber information and profiles. These profiles can be personalized to the individual subscriber's unique preferences for such attributes as feature behavior and call behavior. The HSS delivers this personalized profile to different application servers in the network so that the subscriber has a single identity no matter what device they use to access the system.
- *Quality of Experience (QoE) and Policy Management* – Subscriber QoE will be the critical attribute in the success of IMS. An enhanced QoE is provided by such IMS mechanisms as single sign-on, centralized presence and group management (e.g., buddy lists), and service delivery across device and access types. In addition, operators need to offer subscribers the same – or even better – levels of reliability, predictability, and quality of service (QoS) that users take for granted in the PSTN voice network. That level of QoS is not currently associated with IP networks. IMS is seen as the driver that will finally lead to policy-based QoS mechanisms being deployed in IP networks.

For the time being, however, the focus in IMS is on resource management for the access network – for instance, controlling how many sessions use a particular radio segment, and reserving the resources needed to guarantee that media streams will be carried at their specified rates. The assumption is that the core IP network, perhaps using DiffServ QoS mechanisms, has sufficient bandwidth not to be a bottleneck.

SIP, by itself, is not a resource reservation or resource management protocol. However, an inherent part of SIP messaging, called Session Description Protocol (SDP), provides the application-level trigger points for the CSCFs to request the assignment of the resources for the required QoS for a session. The IMS architecture defines a Policy Decision Function to act as an intermediary between the CSCF and the "network

3GPP and 3GPP2 standards development

The 3rd Generation Partnership Project (3GPP) and 3rd Generation Partnership Project 2 (3GPP2) are collaborations involving several standards bodies from around the world to develop technical specifications and a framework for third-generation wireless networks. Nortel is active in both groups.

3GPP specifications, which are being developed for GSM/UMTS wireless networks, have advanced through a number of releases (R99, R4, R5, and R6). Release 5 defined IMS (IP Multimedia Subsystem), which provides the platform for an open, industry-wide, standardized multimedia service delivery architecture covering all aspects of session control, roaming, quality of service, policy management, and enhanced IP addressing. The R6 release is enhancing IMS to include wireless local area network (WLAN) access.

3GPP2 specifications are being developed for CDMA networks. 3GPP2 has closely aligned its MMD (Multimedia Domain) standards to IMS and, in fact, the working groups have stated their intent to make them the same.

Although wireless operators provided the original impetus for IMS, wireline carriers also want to employ the IMS architecture. Regional bodies, such as the European Telecommunications Standards Institute (ETSI) and the US-based Alliance for Telecommunications Industry Solutions (ATIS), have standards efforts under way to add wireline broadband access to the 3GPP/3GPP2 standards. This work is being fed into the International Telecommunication Union (ITU) Next-Generation Network (NGN) focus group, which is overseeing the development of specifications for the fully converged multimedia communications network of the future.

For more information, visit the 3GPP website at *http://www.3gpp.org/*. The article at *http://world.us.nortel.com/WorldOnline/main.cfm?articleID=34863&language=English* outlines where Nortel has taken a leadership position.

gates" that control the basic underlying network and radio resources.

With its wireless heritage, the IMS architecture defined in 3GPP requires some extensions and modifications for use in wireline networks. The most obvious of these are extensions to new access network types, such as asynchronous digital subscriber loop (ADSL) and wireless local area network (WLAN). In addition, while IMS currently focuses on supporting roaming individual users, it is likely that CSCF roles and HSS functionality will need to be modified in IMS to support wireline residential and enterprise users as well.

Subscriber Quality of Experience (QoE) will be the critical attribute in the success of IMS. An enhanced QoE is provided by such IMS mechanisms as single sign-on, centralized presence and group management, and service delivery across device and access types.

The road ahead

Given the different motivations of wireless and wireline operators for its adoption, it is expected that the rollout of IMS will be different for each group. For both network types, however, there will be a high degree of synergy between the IMS elements and elements already deployed in packet and wireless networks (Figure 2).

Wireline operators can currently deploy Nortel's Multimedia Communication Server (MCS) as a standalone service offering and gain experience with the delivery of multimedia services and evaluate customer response and market demands. However, full-scale deployment of these services will require our CMS solution. For most wireline operators, CMS will be positioned as an overlay to their Succession packetized network. This overlay will require new HSS and CSCF platforms, but our existing Packet Voice Gateway (PVG) media gateways will handle both legacy and IMS calls. Our Communication Server 2000 (CS2K), while continuing to be the call server for legacy (PSTN) calls, will also have the role of a media gateway control function, providing the proper signaling and controlling bearer interworking between the wireline and wireless domains. Over time, IMS is expected to replace the

Figure 2. The path to the Converged Multimedia Services Solution

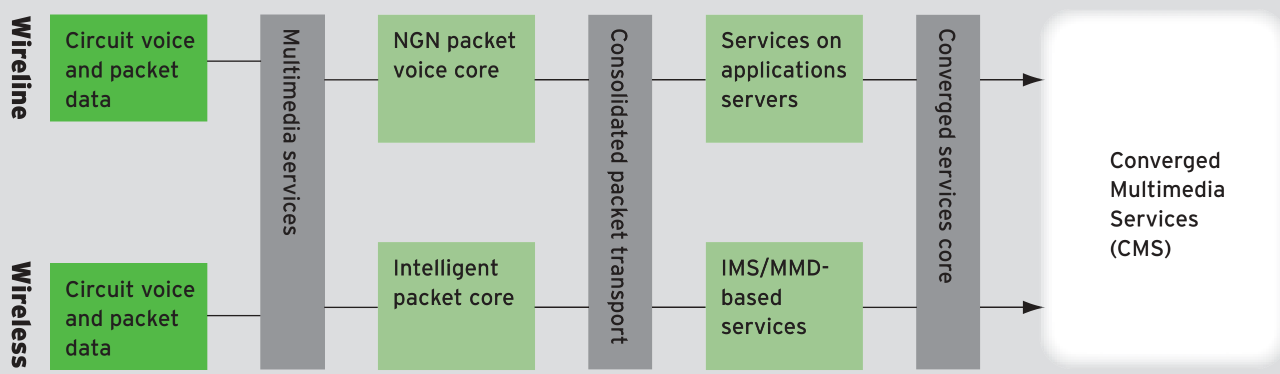

Wireline/wireless convergence is being enabled by packetization of the network and the growing prevalence of broadband access. Both wireline and wireless operators will begin to add multimedia services to the circuit voice and data services that they currently offer. At the same time, wireline networks are evolving to a next-generation network (NGN) packetized core, and wireless networks are evolving to an intelligent packet core, which is similar but based on different standards. When both networks have packet cores, it will be possible to consolidate on a single packet transport network. As well, we are seeing an evolution in wireline networks, as services delivery moves from Class 5 offices onto separate application servers. In the wireless network, the emergence of broadband access that can carry media streams other than voice will drive the deployment of IMS/MMD-based services.

These trends lead to a converged services core, with the same service sets offered on both networks. Ultimately, this evolution will lead to our Converged Multimedia Services (CMS) solution, where multimedia services can be delivered over a consolidated packet transport core and converged services delivery core.

AdvancedTCA standards speed development, increase pace of innovation

by Alan Hurren and the ATCA Design Authority Team

The Advanced Telecom Computing Architecture (AdvancedTCA) – which is a series of hardware specifications that defines a modular architecture for designing telecom equipment – is receiving significant attention and investment from the industry, due in large part to the non-proprietary nature of the specifications and the support by many leading vendors, including Nortel, Lucent, NEC, and Siemens.

AdvancedTCA is seen as a way to speed time to market, reduce costs, and increase the pace of innovation to support new services and a competitive equipment environment. AdvancedTCA, for example, is enabling Nortel and other solutions providers to integrate best-in-class third-party AdvancedTCA hardware and software products from a rapidly growing number of suppliers into their offerings, and focus their internal R&D investment on differentiating capabilities. As industry volume increases, AdvancedTCA adopters will also be able to enjoy the resulting economies of scale.

Defined by the PCI Industrial Computer Manufacturers Group (PICMG), the AdvancedTCA specifications cover board, backplane and shelf mechanicals, system management, power distribution, backplane input/output (I/O), connector zoning, shelf thermal dissipation, regulatory guidelines, fabric technology, and advanced mezzanine cards (AMCs). (More information on PICMG and AdvancedTCA can be found at *http://www.picmg.org/* and *http://www.picmg.org/newinitiative.stm,* respectively.)

The benefits of migrating to AdvancedTCA are well recognized by service providers, several of which have provided guidance to Nortel and other vendors in defining priority requirements for AdvancedTCA-based products. The key benefits for service providers include:

- from inception, recognition of unique telecom requirements for minimum system downtime and stable grade of service under a broad range of operating conditions;
- improvements to product performance, footprint, and operations compared with previous commercial technologies;
- an increased supplier base for hardware and middleware, leading to increased supplier/telecom solution provider competition and a broader product scope; and
- improved technology time-to-market, accelerating the availability and penetration of features and the revenue that these technologies enable.

AdvancedTCA is changing our industry, moving it from yesterday's proprietary technologies to today's non-proprietary industry standards – and in the process is encouraging competition in the telecom marketplace by dramatically reducing the cost of hardware development, which is lowering the barriers to market entry for many smaller firms. The "open" nature of the AdvancedTCA specification is also allowing some industry players to move up the value chain. For example, companies that once specialized in products at the semiconductor level are now able to move beyond that to integrate several of those products at the board level for use in an AdvancedTCA chassis.

Another major area of industry activity is the development of software to support AdvancedTCA-standard hardware. Many third-party software vendors, for example, will soon offer Linux and high-availability middleware products that support AdvancedTCA hardware. Industry forums also are working to create non-proprietary specifications for this software, which are expected to drive further investment and innovation in software components. This should provide Nortel with a wider choice of standards-compliant third-party software products, enabling us to focus our R&D on developing the capabilities that will differentiate our products and solutions from the competition for the benefit of our customers.

At Nortel, the ATCA Design Authority supports AdvancedTCA technology reuse across all Nortel leadership categories. The ATCA Design Authority comprises individuals from organizations throughout Nortel who have the knowledge and expertise to provide guidance to product design teams on technology selection and implementation for AdvancedTCA.

Nortel's first adoption of AdvancedTCA technology will be as a platform in the Converged Multimedia Services portfolio to support a number of applications in wireline and wireless markets. Over time, other Nortel products may also be developed to the AdvancedTCA specifications.

Alan Hurren *is part of the R&D Investment Strategy team in the office of the CTO.*

existing circuit-switched domain with an extended set of peer-to-peer session-based services. The CS2K is expected to expand the roles it currently performs to include the role of application server (for many business features).

GSM/UMTS wireless operators, on the other hand, have their evolution path options mapped out for them by the 3GPP. They can packetize their voice core (called the circuit side or CS) as defined by 3GPP Release 4 (R4) and/or introduce IMS in 3GPP Release 5 (R5). CDMA operators will have similar choices, as defined by 3GPP2.

The focus of IMS/MMD deployments will be on the new services, based on presence and location, which SIP and the IMS architecture will enable. It is not believed that there will be much VoIP on the actual radio link, primarily because the circuit transport of voice over the air is too efficient to be displaced by VoIP when spectrum is the scare resource that it is today. In the GPRS/UMTS arena, Nortel's proposal for SIP Circuit Bearer (SCB) technology allows wireless terminals to obtain IMS services (by signaling with SIP on their packet side), as well as retain the circuit radio link efficiency by carrying voice and video via the legacy radio link transport format.

Cable operators are another set of players that are expected to deploy IMS. IMS is a close fit with the PacketCable Multimedia architecture currently being defined by Cablelabs, the standards body for cable operators. Some cable operators have already deployed the earlier PacketCable standard to deliver PSTN-like voice services to their customers, while others are considering going straight to PacketCable Multimedia and IMS. Again, for cable operators that have deployed the Succession PacketCable voice solution, the PVG media gateways and CS2K will be part of the IMS solution.

Longer-term, it is expected that access technology evolution will eliminate the differences between networks in effectively supporting real-time packet media streams, such as VoIP and video. Operators will no longer be categorized as wireless, wireline, or cable, but will all be IMS operators, competing to offer a whole suite of true multimedia applications and provide a seamless subscriber experience anywhere, anytime, on any device.

Liam Casey *is an architect for the converged core, working in the office of the CTO.*

Software platform convergence to support common features, next-generation products

by Laurence Beaulieu

Today, Nortel supports more than 100 software bases specific to its various product portfolios. With convergence comes the opportunity to converge the number of software bases onto just a few software platforms that will support common features across multiple product portfolios and lines of business, and on which we can build Nortel's next-generation products — enabling us to achieve R&D efficiencies, faster time to market for new applications, and improved software quality, robustness, and security. Currently, Nortel is working on four software platforms: the Multiservice Packet Platform (MPP), Management Services Platform (MSP), Optical Multiservice Edge (OME), and a software platform in the Converged Multimedia Services (CMS) portfolio.

Ten years ago, the majority of Nortel's products were built on fewer than 20 software bases. In fact, most of the company's carrier and enterprise products were built on the DMS software base, and development of new features was leveraged across numerous applications. However, faced with challenging time-to-market pressures during the boom era of the late 1990s, many product groups built their own software bases or inherited software bases from the companies Nortel acquired during that period. As a result, the number of software bases has mushroomed to more than 100. Development of the same or similar features – Multi-Protocol Label Switching (MPLS), Internet Protocol (IPv6), and security, for example – across many software bases drives up development costs and slows the release of common features across Nortel's solutions.

Today, most of Nortel's products are either built on individual software bases or share a common software base across just a few products in the same product family – for example, the BayStack Operating System Software (BOSS) is shared across the family of Ethernet switches, and the Passport Carrier Release (PCR) software is shared across the family of Multiservice switches. Typically, these software bases are architected or designed to support a specific suite of applications in a single product area, and they do not have the support infrastructure for consideration as a common building block or software platform that would support multiple products across the different lines of business. (For more on what constitutes a software platform, see page 22).

The benefits of software convergence are expected to be significant, enabling Nortel to achieve substantial R&D efficiencies by reusing common technology, collaborating on common development, and reducing development duplication.

Now, however, network convergence is providing the opportunity to converge the many software bases onto just a few software platforms, and in the process drive significant efficiencies and other benefits.

To take advantage of this opportunity, Nortel is planning to drive down the number of active software bases from more than 100 to about 40 within the next three years – and eventually to just a handful. R&D investment will be focused on active software bases that remain and will continue to have new feature development. Non-active software bases will be maintained in sustaining mode with little, if any, feature development, until they are discontinued.

It is planned that the majority of Nortel's investment in this area of software development will be directed toward common software platforms, and products that are early in their life cycle will be considered for migration to one of the new software platforms. Products built off software bases that are late in their life cycle will not be migrated to one of the new software platforms.

For example, Nortel's Services Edge R&D team is targeting to reduce the number of software bases/platforms from ten to as few as one or two over the next three years (Figure 1).

The benefits of this software convergence are expected to be significant, enabling Nortel to achieve substantial R&D efficiencies by reusing common technology, collaborating on common development, and reducing development duplication. For example,

- Verification cost can be reduced because common test cases do not need to be repeated across different products.
- Software quality and robustness can be improved as the common software is stressed across many applications, which will help to flush out software bugs more quickly. Indeed, analyses done by Nortel and the industry have found that open source software, such as Linux, is very robust because the code is reviewed by many eyes and stressed across many applications.
- Improved security can be implemented across Nortel products when all products are built on a software platform that uses the same security mechanisms, ensuring full interoperability.
- Faster time to market for new applications can be achieved as product teams leverage the software platform's common architecture, capabilities, and support infrastructure, allowing greater focus on developing the application features.
- The economies of using third-party or open source code can be leveraged within our products by aligning the software platforms with the various commercial off-the-shelf (COTS) technology initiatives and standard interfaces.
- Designer mobility as well as flexibility in staffing development programs may be increased, and new team members can be expected to ramp up to speed faster when they are working with a common technology across different projects.

To achieve these benefits, Nortel is developing or deploying four next-generation software platforms to support new applications and converge some existing applications. These software platforms are:

- Multiservice Packet Platform (MPP), on which several products and applications are being built, including Multiservice Provider Edge (MPE) and our next-generation Contivity Secure LAN product;
- Management Services Platform (MSP), an OAM framework that has been leveraged by the GSM wireless, UMTS wireless, and other products;
- Optical Multiservice Edge (OME), which was released last year and is being used in Optical to converge numerous products and applications; and
- A platform in the Converged Multimedia Services (CMS) portfolio, on which several products are being built, including our next-generation call server [which will merge the existing Multimedia Communication Server (MCS) product and next-generation Session Initiation Protocol (SIP)], and the recently announced Nortel Call Session Controller.

In 2005, we will be exploring opportunities to develop other software platforms, especially for wireless access and enterprise products, where our current network-element-focused software platforms (MPP, the software platform in the MCS portfolio,

Figure 1. Services Edge Software Convergence

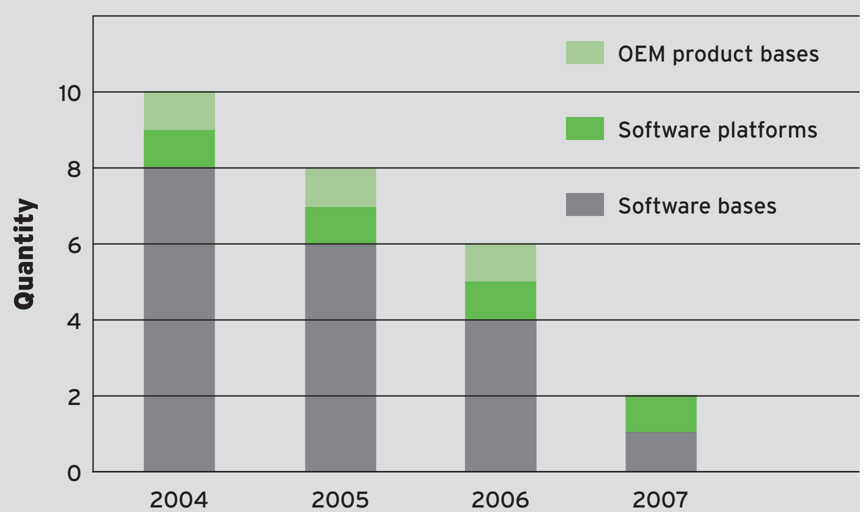

Nortel's Services Edge team is targeting to reduce the number of software bases/platforms from 10 to 2 over the next three years. Currently, there are five network element software bases, one software platform [Multiservice Packet Platform (MPP)], and four OAM software bases. The current goal is to reduce to one software platform (MPP) for network elements, and migrate three of the OAM software bases to Multiservice Data Manager (MDM), which could potentially migrate to the Management Services Platform (MSP) software platform in the future.

What is a software platform?

At Nortel, a software platform is a common set of managed, integrated, and tested software functions and capabilities that meets the needs of a suite of software products and applications, and allows them to be developed and delivered independently. It is designed to be portable across various hardware platforms through the use of well-defined and abstracted hardware interfaces.

Software functions and capabilities – the software architecture – include an underlying operating system, carrier-grade extensions (as required), well-defined hardware and application interfaces, and operations, administration, and management (OAM) elements, as well as high-availability middleware and application middleware.

Key components of high-availability middleware are message service, event notification service, high-availability services, and availability management frameworks. Application middleware consists of many of the components that are typically contained in individual applications, and the components common to two or more applications are centralized in the application middleware for use by all applications. Application middleware may provide protocols, web services, directory services, security services, database services, overload controls, session management, and OAM components.

In addition to functions and capabilities, a software platform consists of a support infrastructure of people, organizational structure, and processes. Key elements of the support infrastructure are:

- Functional roles – design authorities, project office, architects, and designers;
- Feature development – plan of record/ roadmap, feature capture process, collaborative development models, feature review and acceptance, and feature testing;
- Operation – formal agreements, process for bug tracking and fixes, and a software release mechanism; and
- Technical support – technical documentation and platform test suite.

At Nortel, we have established support infrastructure requirements that Nortel software platforms will be required to meet in order to be considered a full software platform. Over the next few years, we are planning to evolve Nortel's software platforms to meet the support infrastructure requirements and provide the full suite of technical features.

and OME) do not fully address these product segments. This focus may involve expanding an existing software base to become a common software platform or developing a totally new software platform. For example, the MSP OAM software platform will address wireless access, and there may be an opportunity to leverage the MSP platform for enterprise.

Another area of focus is the use of these software platforms in the convergence of the wireline and wireless core switching networks – a major Nortel, as well as industry, initiative. This convergence is being driven by the carriers to leverage common applications and reduce cost. The specifications for these networks are being defined by the 3rd Generation Partnership Project (3GPP) and 3rd Generation Partnership Project 2 (3GPP2). (For more on 3GPP/3GPP2 and wireline/wireless convergence, see page 14.)

A current view is that Nortel could build all its next-generation products for the converged wireline/ wireless core network on as few as three of the software platforms (the OME is not applicable in this particular area):

- MPP, which could be the software platform for all the network elements, such as GGSN [Gateway GPRS Switching Node (GPRS - General Packet Radio Service)], PDSN (Packet Data Serving Node), SGSN (Serving GPRS Switching Node), and Media Gateways;
- A software platform in the CMS portfolio that could be used for all the service core network elements, such as the Nortel Call Session Controller and Nortel Policy Controller; and
- MSP, which could be the software platform and framework for all the OAM elements.

To facilitate adoption of the software platforms for new product developments – or for the migration of existing products that are early in their lifecycle and require significant investment over the next few years – Nortel has established a Design Authority (DA) team with members having expertise in specific platforms.

Laurence Beaulieu *is an R&D Analyst in the CTO Office's Applied Technology Leadership and R&D Investment Effectiveness team.*

Project DRAC: Creating an applications-aware network

by Franco Travostino, Rob Keates, Tal Lavian, Inder Monga, and Bruce Schofield

Intelligent networking and the ability for applications to more effectively use all of the network's capability, rather than just the transport "pipe," have been elusive. Until now. Nortel has developed a proof-of-concept software capability — service-mediation "middleware" called the Dynamic Resource Allocation Controller (DRAC) — that runs on any Java platform and opens up the network to applications with proper credentials, making available all of the properties of a converged network, including service topology, time-of-day reservations, and interdomain connectivity options. With a more open network, applications can directly provision and invoke services, with no need for operator involvement or point-and-click sessions. In its first real-world demonstrations in large research networks, DRAC is showing it can improve user satisfaction while reducing network operations and investment costs.

Across a number of different fields – health care, manufacturing, and research and education, for example – users are demanding more and more network resources, such as bandwidth, quality of service, and security, in order to collaborate and share data around the globe. To meet these demands, the CIOs for these institutional, commercial, and research pools have employed several solutions, including supercomputing technologies, grid computing, and peak provisioning and manual provisioning of network resources. None of these methods, however, is optimal.

Supercomputing technology, for example, is expensive to scale and limited to those institutions and researchers that can afford it; high-performance grid computing, while it enables multiple users to share resources to boost processing power, is difficult to achieve due to the cost of interconnecting processors with low latency and high bandwidth; and statically and peak-provisioned network setups can result in over-provisioning of network resources, which is expensive and results in CapEx and OpEx inefficiencies.

At the same time, operators must optimize their networks to meet diverse user requirements, which can range from single-user multi-Gbit/s data transfers (as with grid applications), to best-effort many-to-many kilobytes (such as e-mail). How can the network provide the cost/performance profile of lower-layer (e.g. optical) technologies while achieving the dynamic capabilities of higher layers? As well, how can the network adapt to the evolving applications?

Manual provisioning of network resources has also been known to present challenges. For example, processing their massive data files requires researchers to first contact a network administrator to set up and provision the appropriate network resources – typically a manual, error-prone task accomplished through a point-and-click session operated by the network administrator, which can lead to delays and potential failures.

Moreover, as data travels through the end-to-end network and across different networks, it typically encounters different types of network technologies – from packet, circuit, wireless, and wireline to various access environments – each with its own separate topologies, protocols, and features, again leading to missed opportunities or high CapEx/OpEx costs.

Dynamic Resource Allocation Controller

To address these challenges, Nortel has developed a proof-of-concept capability, called Dynamic Resource Allocation Controller (DRAC, pronounced d-rack).

Essentially, DRAC acts as an agent of the various applications, brokering and configuring on an end-to-end basis all the necessary pieces of the network, regardless of the type of network – circuit or packet, wireless or wireline. DRAC enables applications to control their share of network resources, yet without requiring them to interface directly with a wide range of diverse and constantly evolving network protocols, features, and devices. Put another way, DRAC lays the tracks ahead of the train, adjusts the network resources that an application needs, and steers the data through the network – and it does this either dynamically in real time or on a time reservation basis.

DRAC is implemented as software that is designed to be portable to any Java platform. This middleware sits between applications and the network (whether management planes, control planes, or individual network elements) and provides applications with an abstracted view of the underlying network (Figure 1). The DRAC southbound interface features a framework for the instantiation of multiple "drivers," corresponding to the different signaling techniques encountered in legacy networks (Figure 2).

Three degrees of coupling

With the DRAC concept, Nortel envisions three degrees of coupling between applications and networks.

- *First degree:* In hybrid optical and packet networks, such as SURFnet6 (discussed later in this article), DRAC provides "cut-through" capabilities across network layers by steering very large flows of packets or low-latency applications dynamically over Layer 1 instead of Layer 3. For example, instead of dedicating routing resources to multi-Gigabit point-to-point file sharing applications or alternatively setting up a dedicated and costly high-bandwidth optical connection, DRAC simply sets up and takes down "ephemeral" optical circuits as they are needed, minute-by-minute, hour-by-hour. By bypassing the routing layer for this type of traffic, DRAC enables a higher-performance network experience for both routed and bypassed traffic and reduces the total number of routers required in the network. In fact, in one real-world design, DRAC reduced the number of required routers from 20 to 2. The same thesis can be applied to other environments featuring diverse technologies, such as between wireline and wireless.
- *Medium degree:* DRAC is capable of recognizing the network footprints of a given application (through deep packet inspection or direct signaling from the application, for example). DRAC makes sure that the network reacts appropriately to an application's behavior. For

Figure 1. DRAC core framework

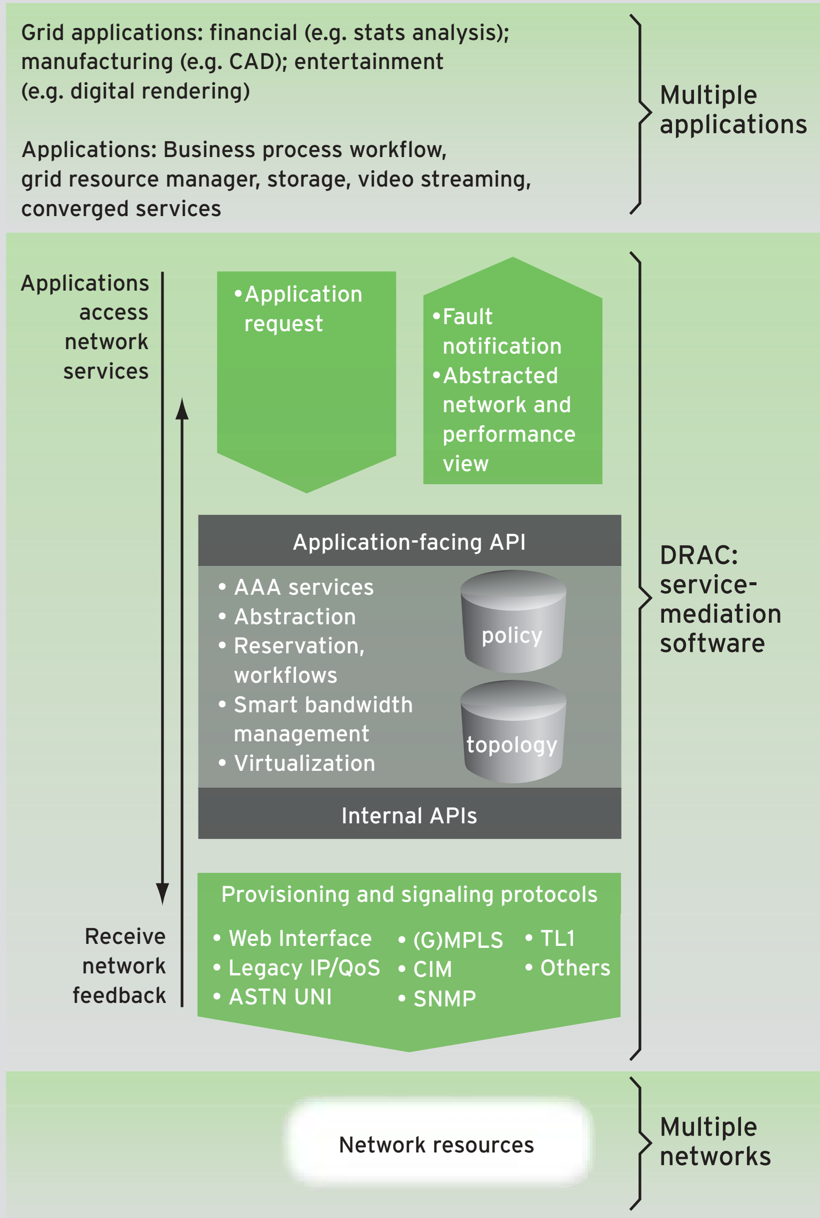

The DRAC core framework includes AAA (authentication, authorization, accounting) services, policy engine, topology discovery engine, resource reservation services, workflow utilities, interdomain routing facilities, and smart bandwidth management fixtures. The interface to applications is bi-directional, enabling network performance and availability information to be abstracted upward toward the application. The DRAC provides applications with the means to directly drive their share of network resources within a policy-defined envelope of flexibility. Network resources include bandwidth, quality of service (QoS), security, acceleration appliances, sensors, and more. As well, the DRAC strategy is to use existing standards and toolsets for interfaces, which greatly simplifies deployment in multivendor, multi-technology environments.

example, when a critical storage restore operation is initiated due to disaster recovery, DRAC ensures that the network dedicates a large fraction of its resources to expedite that operation.

• *High degree:* DRAC becomes privy to the overall "flight plan" of an application. For instance, DRAC learns how a particular workflow unfolds among peering instances of a distributed application. That way, DRAC can anticipate the network requirements, evaluate what-if scenarios, and enact failure-recovery strategies that are cognizant of the workflow. These are the defining properties of what we call "workflow-engaged networks" (WENs). In all cases, DRAC enables much more efficient use of network resources, leading to operational and capital savings.

Application value

Currently, the DRAC value propositions have been validated within four vertical market segments:

• Hybrid optical and packet networks validate the DRAC cut-through capabilities by steering very large data flows across an ephemeral optical circuit, and allowing smaller flows that are more tolerant to latency and/or congestion loss to communicate via Layer 3.

• Within data centers, DRAC will help storage partners realize large bandwidth savings across metropoli-

Figure 2. DRAC middleware

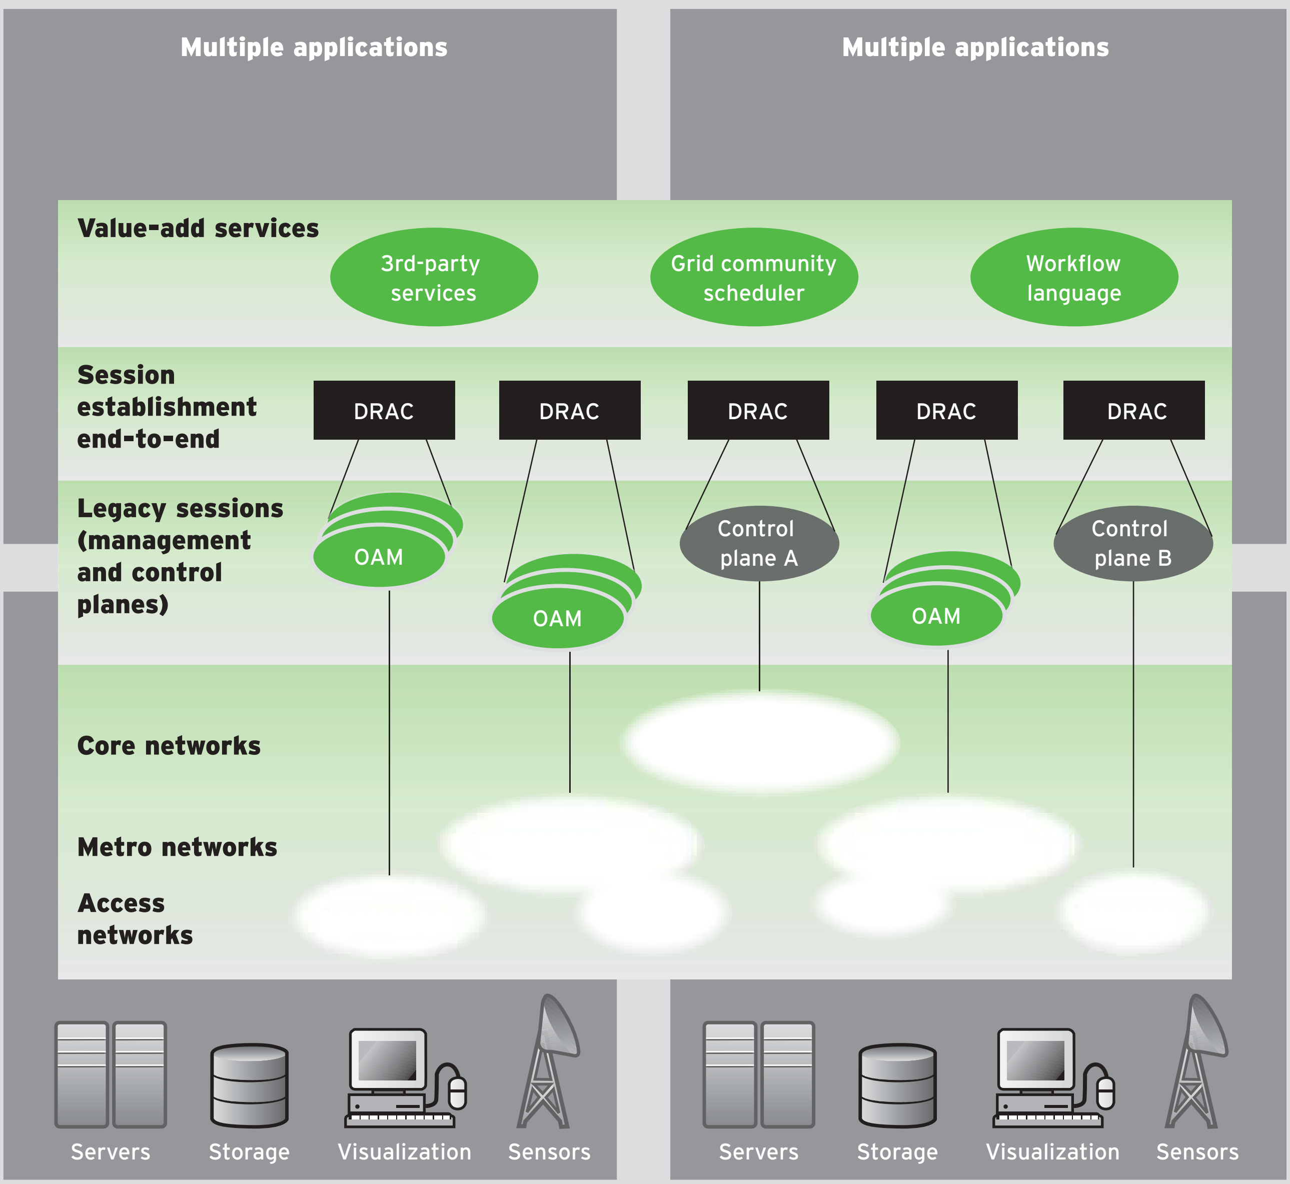

DRAC sits between applications and the various networks that connect end systems, such as servers, storage, visualization devices, and sensors.

What is SURFnet?

SURFnet is a national high-performance computer network in the Netherlands that connects more than 150 institutions in higher education and research to each other, as well as to other networks around the world.

Operated by a Netherlands-based company, SURFnet is among the leading research networks in the world, collaborating closely with national and international organizations and serving as an advanced test environment for investigating new technologies that will continually improve the reliability, security, and speed of its network.

Among its many ongoing projects, SURFnet is responsible for the realization of the GigaPort Next-Generation Network – a project of the Dutch government, trade and industry, and educational and research institutions – that aims to strengthen the national knowledge infrastructure. Research on optical and IP networking and grids is a prominent part of the project. For more information, visit *http://www.surfnet.nl/.*

tan area networks (MANs) when operating replication, business continuance, and disaster recovery applications. During trials completed with some of Nortel's storage partners, the control and monitoring of storage and network management functions were consolidated into one unified "cockpit" versus two, to command and control these functions, with DRAC discovering the topology and composing a complete system view inclusive of the storage topology.

• Healthcare workflows (such as in radiology practices) are a natural fit for the DRAC's ambitions in WENs. Beyond the bandwidth savings seen in data center scenarios, the WEN can improve on dependability, while optimizing the expenses in network and storage setups. [For further detail, refer to Schofield's paper "Workflow Engaged Networks for Radiology in Metro Regions," presented at the 90th RSNA, Radiological Society of North America, Chicago, November 28, 2004. This is the study of a network (wireline, LAN+MAN) utilized for Filmless Radiology, with and without a DRAC framework.]

• Within grid computing communities, we are working to elevate the network to a primary grid-managed resource, akin to CPU and storage resources. DRAC can tame the complexity and diversity of network elements to open the way for e-utilities.

Demonstrating the values

DRAC has demonstrated compelling values in a wide range of applications. Detailed studies by Nortel have shown large cost savings in MANs and wide area networks (WANs), providing an appealing alternative to the old approach of static, over-provisioned networks.

As well, DRAC is undergoing validation in real-world network deployments with high-performance computing networks. For instance, the Netherlands-based SURFnet is currently deploying DRAC at the heart of a hybrid optical and packet network, called SURFnet6, which is being realized in the context of the GigaPort Next-Generation Network project (see sidebar). (SURFnet 6 optimizes Layer 1 network resource utilization based upon end-user requirements.) In this implementation, applications driving multi-Gbit/s transfers bypass the packet layer and are steered directly onto wavelengths between end points across temporarily assigned optical links.

In addition to the SURFnet project, Nortel recently demonstrated DRAC features at the Supercomputing 2004 conference in November 2004. This demonstration was done in cooperation with recognized research and education leaders, including SURFnet, Netherlight, the University of Amsterdam, Internet2, Canarie, iCAIR, and Starlight. Encouraged by the results of the DRAC preliminary results, Nortel is currently looking to a broader deployment market opportunity.

Franco Travostino *is leader of an Advanced Technology Team that is exploring applications-engaged networks and grid infrastructures. This team includes* Tal Lavian, Inder Monga, and Bruce Schofield.
Rob Keates *is senior manager, optical networks marketing.*

Note: This article has been adapted from a Nortel Applications Brief (www.nortel.com/drac*), published in November, 2004, and distributed at Supercomputing 2004, held in Pittsburgh, Pennsylvania.*

Converging innovative multimedia services onto one product solution

by Doug McGregor

Since their introduction in the 1980s, services such as calling-line ID, call waiting, and voice mail have generated significant revenue for traditional carriers. Since then, however, carriers have offered little new from a "value-added" services perspective, other than a recent focus on providing narrowband and broadband Internet access. But that's changing. Innovative multimedia services and service bundling from non-traditional carriers — including wireless, cable/MSO, and even Internet providers — are spurring traditional carriers to focus on introducing new value-added services. Nortel is playing a lead role in this evolution by helping carriers identify these revenue-generating service opportunities, and by converging those multimedia services onto single product solutions, like the Multimedia Communication Server, or MCS 5200.

In recent years, traditional carriers have focused almost exclusively on providing network infrastructure – the "big pipes." While concentrating on leveraging that capacity to provide subscribers with Internet access, many carriers have been reluctant to invest in other value-added services that ride over those pipes. This is particularly true in the Voice over IP (VoIP) arena, where traditional carriers fear the cannibalization of their existing revenues (primarily derived from voice connectivity) and the need to prematurely write down the value of their TDM network assets. Nevertheless, market forces are creating new opportunities – if not imperatives – for these carriers to invest in services, and Nortel is playing a key role in enabling these new capabilities. These market forces include:

• *Competition through convergence*: Wireless operators have taken a lead role in developing new services in recent years, including text and multimedia messaging. They definitely know how to compete and how to identify services that subscribers will value and pay a premium for. These new services – coupled with consumer behaviors (such as embracing instant, text, and video messaging) and emerging technologies (including WiFi) – are creating a number of scenarios where wireless providers are now competing head-to-head with their wireline brethren. Likewise, cable operators are becoming equally innovative in the war for subscribers, offering – in addition to cable TV service and high-speed Internet access – telephony and multimedia services over IP in combination with compelling new video propositions, such as IPTV.

• *No monopoly on network coverage:* The Internet now provides a very viable alternative to the telephone network for global communications. The pervasiveness of the conventional telephony network is no longer a competitive differentiator in much of the world.

• *Not just voice:* The Internet makes multimedia communication possible and, for that reason, has triggered a major change in the telecom marketplace, which previously had to settle for "just voice." Primary research conducted on behalf of Nortel shows that this technological discontinuity is coinciding with rising consumer and business demand for a rich interaction of services. Some public carriers, like Bell Canada, as well as cable operators, are seizing the opportunity to offer multimedia services that address a growing consumer/business appetite – and gain first-mover advantage in the emerging services market.

• *Reduced barriers to entry:* Gateway technology in conjunction with VoIP leads to a further interesting dynamic. By leveraging gateway technologies, carriers – with minimal additional capital investment – can compete outside of their normal business areas. For example, a carrier that provides service mainly in Ontario and Quebec – two of Canada's eastern provinces – could deploy a gateway and offer service to customers on, for instance, Canada's west coast. The corollary to this, of course, is that some carriers are seeing their traditional customer base attacked by these new competitors, which means they are seeking new ways to hold onto their existing clientele.

While traditional carriers are feeling compelled to offer new services just to remain competitive – especially when up against non-traditional competitors such as the cable/MSO operators – new multimedia services can also offer carriers a lucrative payback. By introducing the right services, carriers can turbo-charge top-line revenue growth. To understand the revenue potential of new services, consider the wireless carrier industry in Western Europe. The emergence of data services, including

WAP (Wireless Applications Protocol) and SMS (Short Message Service) text messaging, demonstrates that carriers can profit from the introduction of new capabilities. In Western Europe alone, some 150 million SMS subscribers in 2001 generated almost EURO10 billion in revenue for service providers, and the market continues to grow (Source: Global Information Inc. *http://www.gii.co.jp/press/iu10823_en.shtml*)

If rich multimedia services represent the next wave of revenue opportunity for carriers, many of you may be wondering – as I did a number of months ago – how the Internet fits in. After all, aren't many of these services available free of charge right now via the Internet? Yes, many of them are, but you need look no further than Microsoft MSN Messenger to see the challenges these services must still overcome. For example, working through the menu on my own desktop, I found the option to "make a phone call." However, when I invoked it, I was blocked by firewall issues associated with my home router. Bypassing my home router allowed me to overcome this hurdle, but then I was presented with screens that asked me to establish a commercial relationship with a VoIP provider in order to place the voice call. Clearly, Microsoft hasn't fully integrated the old (the traditional voice network) with the new (the PC multimedia experience).

At Nortel, we recognize that the key to widespread adoption of rich multimedia services is their integration with existing voice services and desktop gear. In fact, recent market research sponsored by Nortel shows that both enterprise and consumer customers want to maintain the familiarity and investment in their current voice services, while still enjoying the productivity, mobility, and simplicity advantages of integrated multimedia services. In short, they want to have their cake and eat it too.

That's why the MCS 5200 allows service providers to offer new value-rich multimedia services that complement rather than replace existing voice services. Nortel's MCS 5200 integrates multiple media (voice, video, text, and data), allowing users to talk, see, and share information from any location as simply and easily as they pick up the phone and make a call today. For both enterprises and consumers, the MCS 5200 removes the boundaries associated with existing multimedia services, providing seamless communications across the Internet, as well as traditional wireline and wireless networks.

Figure 1. Multimedia Communication Server (MCS) 5200: voice, video, text, and data integrated into one seamless communication session

All-in-one tool
- Voice and video calls
- Collaborate and share files
- Instant message

Contact and call management
- Find personal and work contacts
- Track incoming/outgoing calls

Availability control
- Screen and route calls according to predefined rules
- Redirect priority callers to where you can be reached

Our services value proposition

The MCS 5200 provides a rich multimedia experience for the user (Figure 1) because Nortel understands some of the key ingredients for success in the still-maturing services market.

• *Integration of services:* One of Nortel's unique value propositions is that we understand the benefit of integrating multimedia services with respect to one another. For example, the MCS 5200 provides a "Friends Online" directory that allows you to check on a person's availability, or "presence." This, in turn, makes it possible to determine whether the person you are trying to reach is "active and available," on the phone or away from their desk. This "presence" information makes it simple to determine the most efficient means of communications – whether it's best to instant message them, place a regular voice call, set up a point-to-point video call, or transfer a file.

Personally, I find the instant messaging capability with presence a significant help in my work day. Furthermore, the integration of instant messaging with wireless clients, such as the BlackBerry device, means that I get call logs and instant messages even when traveling on the far side of the globe. [Note: Many of our startup competitors have different strategies primarily around reinventing voice with some call control capabilities, such as follow-me/find-me services. Often a separate screen is invoked with little to no real integration of these services.]

The MCS 5200 is a prime example of the innovation and value our internal R&D capability brings to the table. That said, "no company is an island" and Nortel isn't capable of single-handedly creating all the innovation that this marketplace demands and can deliver. Alliances with third-parties are creating a rich portfolio of devices (e.g., telephone sets) and soft clients (e.g., PC- and mobile terminal-resident clients), which means that carriers can better customize and tailor the services they offer their customers.

• *End-to-end solutions*: A large part of Nortel's value is in being able to provide an "end-to-end solution." It's a term that can roll too easily off the tongue. However, the complexity of integrating services into IP and conventional telephony networks, as well as a carrier's back-office systems, can be a huge challenge that few vendors can meet. Nortel can. Many competitors, especially the startups, provide only software solutions and leave this heavy network and office integration work to the carrier.

• *The importance of marketing new services:* Many carriers appreciate support in selling value-added services to their customers. Years ago, when "calling-line ID" services were being launched, Nortel worked in close collaboration with carriers to develop marketing materials that they could use to position and sell these new features. Fast forward to today

Session Initiation Protocol

by Liam Casey

Session Initiation Protocol (SIP) is fast becoming a dominant signaling protocol for voice and multimedia communications in the converged network.

Early approaches – rooted in the telecom world – to packetize voice and multimedia communications focused on adapting the existing International Telecommunication Union (ITU)-defined ISDN and ISUP (ISDN User Part) signaling protocols to run on IP. (ISDN and ISUP together form the basis for current PSTN operation.)

These solutions, however, were considered complex, so those in the IP world (the so-called "net heads"), working through the Internet Engineering Task Force (IETF), started developing SIP almost a decade ago as a back-to-the-basics reaction to the perceived complexity of the ISDN and ISUP standards.

SIP started off simply enough as a peer-to-peer protocol with just four messages, which enabled users to place a basic phone call over an IP network. Today, however, it would be hard to argue that SIP is any less complex than ISDN and ISUP. Indeed, as development on the SIP standard has progressed, it has evolved to include many of the features of ISDN and ISUP – features such as call forwarding, transit network selection, and caller ID display that provide the important functionality that users have come to expect from their phone service.

So, why does SIP have the edge? The belief is that SIP will enable more features and applications to be developed much more rapidly, for two primary reasons:

• Like HTML, SIP syntax is text-based, making it easier for developers to comprehend and debug SIP applications; and

• SIP provides an open way in which to add features to call processing. Instead of AIN (Advanced Intelligent Network) trigger points, the SIP model simply passes, or "daisy chains," the SIP messages through a series of proxies or agents, where each agent can perform functions that modify the call. This model makes it easier for a third party to develop a new feature or service – a developer needs only to produce an agent that deals with standard SIP messages rather than having to develop code that is specialized for a specific telephony switch or Service Control Point (SCP).

Even as the SIP standard evolves to support voice capabilities beyond simple connection, it is this flexibility that is driving strong momentum behind its adoption as the signaling base for all future multimedia converged networks.

By adopting SIP for VoIP call processing as well as for applications development, Nortel will be well positioned to create new multimedia applications for its customers – either by itself or in partnership with third-party applications developers.

For more information on SIP, start with *http://en.wikipedia.org/wiki/Session_Initiation_Protocol*

Liam Casey *is architect for the converged core, working in the office of the CTO.*

and Nortel has a staggering amount of material to help position services to a carrier, as well as materials that carriers can use with their customers. From what I've been able to see in the marketplace, no other vendor, large or small, has taken this collaborative approach to helping carriers market and sell new services.

• *The benefit of selling to many different customer segments:* Because our customers use Nortel's portfolio to address the needs of many different market segments – including large enterprises, small and medium-size businesses, and consumers – we are in an excellent position to share learnings across segments. If we find out why a feature or service is successful in one segment (e.g., the enterprise market), we can apply that knowledge to address the needs of another segment, or use that information to drive future developments. Within Nortel, we have both carrier- and enterprise-focused MCS organizations. We meet regularly and align our strategic plans so that we can mutually benefit from what we have learned in dealing with our different customer segments.

Evolving MCS 5200

Inside Nortel, the MCS 5200's value is one of both technological reuse of the features and repartitioning of capabilities so that MCS 5200 applications (such as "presence") can be extended into other parts of the network for integration with other services. Today's MCS 5200 solution traces its roots back several years and combines different functions and products, from both Nortel and best-in-class third parties. On the horizon, we see a number of factors shaping the continuing evolution of the MCS 5200, including:

• *Emergence of IMS:* Both wireline and wireless carriers are beginning to converge on the nature of an IMS (IP Multimedia Subsystem) network for the delivery of voice and multimedia services in the future. The 3GPP Forum, in which Nortel is a participant, is specifying these network architectures (page 16).

• *Emergence of converged wireless/ wireline services:* With network convergence, service providers can open new business opportunities through the delivery of services that cross traditional networking lines. Personal Communicator is an example of a service that provides a new level of convergence between cellular and wireless LAN networks. With Personal Communicator, users can take advantage of their home wireless LANs and broadband connections and a dual-mode (cell and WiFi) phone to place calls that seamlessly migrate from one network to another as the user's location changes. Utilizing the user's presence information, the network can determine if inbound calls should be routed to the user's home via the broadband network, or to the user's cell phone. In addition, if the handset detects signal degradation on one network, it can initiate a call handoff to the other network

• *Influence of "back-office systems:"* Perhaps more profound will be the impact on MCS of "back-office systems" (systems in the traditional IT domain, such as billing systems). Back-office systems are often the single largest impediment to getting a new service to market. Historically, back-office investment was on par with the investment required in the network. Today, back-office investment can be as much as ten times the cost of the network equipment. This side of the industry is also working on advanced architectures to facilitate communication among the many subsystems within the back-office environment. Some customers are telling us that in the future some aspects of real-time processing, particularly in the signaling path, will move into the IS domain (the back office) and away from network-centric operation.

• *Customization of services:* Customers who are already launching new services, especially in the carrier space, find a large degree of "service customization" is required to support their own brand and marketing strategies. Fortunately, there are third parties that offer service-creation capabilities that enable peer-to-peer networking over SIP (Session Initiation Protocol, which allows the separation of the application from the network) and allow our customers to undertake the customization they want. (For more details on SIP, see page 29.)

As you can see, the very nature of the carrier business model is changing – it's no longer about providing simple connectivity. That's why, for carriers, the true value of the MCS 5200 is the ability to bring together different worlds, enabling service innovation and eliminating communication boundaries. For end users, the benefit is the communications services and solutions that enable them to interact with one another on a variety of levels, creating a virtual sense of community. For service providers, the combination of value-rich services, service bundling, and convergence allows them to meet the current and future voice, video, and data networking needs of their consumer and enterprise customers, building their brand and providing top-line growth and bottom-line improvement.

Doug McGregor *is General Manager, Multimedia Communication Server.*

Quality of experience as an integral part of network engineering

by Kathy Bharrathsingh

Nortel is a recognized world leader in the engineering of voice networks that meet end-users' Quality of Experience (QoE) expectations, and is taking on the challenge of doing the same for converged networks. This article introduces an innovative engineering process, termed QoE engineering, that makes QoE requirements an integral part of network engineering. The article also looks at the parameters required for different service domains — voice, data, video — as well as end users' expectations for various services in a converged network.

A key challenge in the successful realization of network convergence is to ensure that all applications running over the network perform well, regardless of whether they are voice, video, data, real-time, or non-real-time applications. The converged network must be able to efficiently carry all traffic types without degrading any of them, and it must meet the combined requirements for all services at a level equal to or even better than what the user has become accustomed to from a single-service (voice-only or data-only) offering.

Addressing this challenge means that at each level of network convergence, performance parameters must meet the strictest requirement of all those defined for the individual services, while simultaneously delivering an overall acceptable quality of experience (QoE) that includes security, reliability, and availability.

QoE is the users' perceptions of how well a system, application, or network interaction performs relative to their expectations, as well as how intuitively they can use an application or service to accomplish a task in a timely and efficient manner, without concern for the underlying network elements.

QoE has sometimes been used interchangeably with quality of service (QoS). However, while QoE and QoS may be related, they are not the same. For example, it is possible to have excellent QoS but poor QoE, as with the flawless transmission of garbled packets.

Furthermore, while QoS is measured objectively, QoE is a subjective measurement that generally requires translation into quantitative data. QoE can be objectively quantified using standardized statistical procedures and various analysis methods (ratings, frequency counts, response times, cluster analyses, probit analyses, etc.). Another approach to quantifying QoE is systematic modeling consistent with that used for performability evaluation of fault-tolerant systems.

Traditionally, providers have focused on QoS to ensure service performance. QoS involves such measurable parameters as service availability, delay, delay variation (or jitter), throughput, packet loss rate, bit error rate, and signal-to-noise ratio. QoS helps operators determine the levels of quality to use for different services, as well as understand how to configure services in order to differentiate them. At the same time, operators must balance this against the need to minimize cost and maximize link utilization.

Focusing on and incorporating QoE as part of our engineering methodologies ensures a customer-centric perspective and helps us move beyond traditional service level agreement (SLA) metrics and exclusive network performance QoS metrics. With QoE, customer needs and expectations become central to product design and business processes.

Indeed, understanding QoE and being committed to ensuring quality of user experience can play a major role in increasing customer loyalty and enhancing a company's reputation and business performance.

QoE engineering methodology

To help drive QoE into the network planning and engineering processes within Nortel – and to achieve both QoE and QoS service requirements in a converged network – Nortel has created an innovative QoE engineering methodology.

With this methodology, the interrelationship of QoE/QoS and traffic engineering is defined based on a top-down approach, starting at the end-user level. By ensuring that end-users' preferences, tolerances, and requirements are considered early in the design and engineering of networks and solutions, Nortel can offer its customers an enhanced ability to differentiate their service offerings with services that fully leverage the power of converged networks.

The objective of the methodology is to facilitate the selection of effective QoS mechanisms that satisfy the end-user QoE of a given application. By creating an application layer human-factor model and considering the underlying network transport layer characteristics, we can eliminate non-practical end-user solutions, deliver up-front assurance to our customers that we understand the needs of their customers, and reduce SLA non-conformance. This methodology can also be used to evaluate quantitatively the effectiveness of QoS mechanisms against user-defined QoE requirements.

The QoE engineering methodology is shown in Figure 1. This process captures a set of key performance indicators that have been used to objectively analyze the behavior and performance characteristics of a QoS-enabled network from a user perspective.

In order to deliver acceptable service quality in a converged network, QoE should be part of the traffic engineering process. It is important to understand the distinction between QoE engineering and traffic engineering. Traffic engineering is the prevailing technique for mapping traffic flows to ensure optimally utilized bandwidth and

Figure 1. QoE engineering methodology

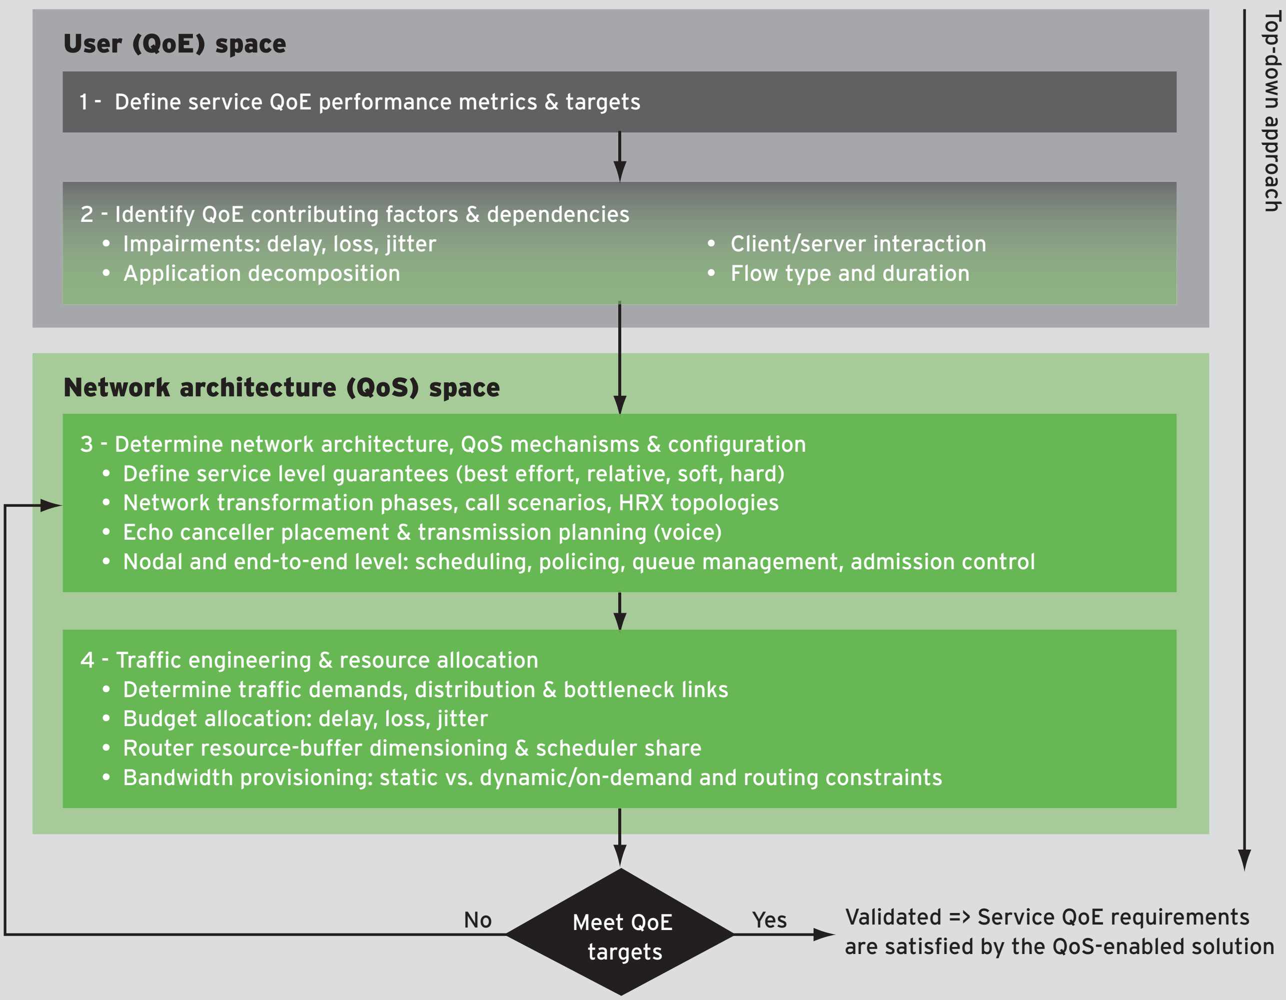

Quality of experience (QoE) is the overall performance of a system from the point of view of the users.

Quality of service (QoS) refers to a set of technologies (QoS mechanisms) that enable the network administrator to manage the effects of congestion.

Nortel has created an innovative methodology as part of its engineering practices to achieve both quality of experience (QoE) and quality of service (QoS) requirements and to assist in the successful deployment of converged networks.

Figure 2. Triple play QoE / QoS metrics and targets

Services	QoE / QoS performance targets	
	Metrics	**Targets**
Conversational voice	R-factor	ΔR PSTN - Packet < 3R
	End-to-end delay	< 150 ms
	Distortion	Ie < 3R
	DSL access jitter	< 25 ms
	Packet loss	less than 10^{-3}
Broadcast TV video	Channel change time end-to-end delay	750 ms +/- 250 ms
	VoD control end-to-end delay	< 200 ms
	Distortion	TBD
	DSL access jitter	< 100 ms
	IP packet loss	less than 10^{-6}
Data	Interactive application response time	< 200 ms
	Responsive application response time	< 4s (<2 preferred)
	Packet loss	< 1% (0.1% preferred)

QoE

QoS

QoE / QoS metrics and targets are required to judge, validate, and engineer triple play solutions

prevent congestion build-up. Traffic engineering is one of the necessary steps of QoE engineering but alone it is not sufficient. In order to meet specific service quality targets and user QoE, traffic engineering needs to be supplemented by additional considerations highlighted in the QoE engineering methodology.

Triple play scenario

To demonstrate the benefits of the QoE methodology, the Nortel team examined the QoE requirements for triple play services (voice, data, video) – a key multimedia service offering that takes advantage of converged networks and one that many service providers are considering and/or are beginning to deploy.

To understand the considerations for triple play services, the following objectives were used as a baseline:

- Determine the potential impairments and mitigations of real-time VoIP, video, and non-real-time data applications;
- Develop a modeling framework to predict the behavior of packet networks carrying triple play services and compare these against requirements defined by the user QoE;
- Validate quantitatively the performance of various QoS-enabled architectures and associated mechanisms, as well as their effectiveness in delivering end-user QoE; and
- Determine what QoS mechanisms to use and when to use them.

Figure 2 summarizes the metrics and targets for triple play solutions. The metrics were divided into the following classes: QoE and QoS; QoE focused on user requirements; and QoS-targeted network level L2-L3 characteristics. It should be noted that while these metrics and targets serve as an example in a triple play scenario, similar conceptual frameworks can be applied to security, reliability, and so on.

Key targets for triple play services

Voice-related targets are derived mainly from the International Telecommunication Union (ITU) ITU-G.107 E-model standard. The E-model is a planning tool for estimating the overall quality in a telephone network. The basic premise for this model is that network impairments – such as noise, echo, delay, codec performance, and jitter – are always psychologically additive. In simple terms, the overall quality (termed the 'R' factor) is calculated by estimating the signal-to-noise ratio of a connection and subtracting the network impairments that, in turn, are offset by any expectations of quality by the caller. The objective rating of quality (R factor) is said to reflect the overall caller experience.

The targets from this standard have been created from extensive subjective studies by Nortel and other telecommunications industry participants. Defining targets is not a trivial task and requires the appropriate subjective evaluation expertise along with human factors expertise. A user experience-based approach implies that there might be multiple targets and that no single target may be acceptable in all situations.

For example, mobile users have different expectations than wireline users, as do users making international calls and those making local calls. As a result, different targets may be required depending on the call scenarios supported by the network. It has been determined that a difference of 3R is not noticeable by typical users and therefore packet networks could be engineered within this margin in order to provide an equivalent replacement technology.

As well, packet networks can introduce potential degradation to the voice/data channel. Sources of impairment can be classified into two distinct groups: non-controllable (or intrinsic) and controllable. Impairments are considered "controllable" if the network architect, the box designer, or the network operator can make choices that increase or decrease the impairment. Non-controllable impairments, on the other hand, are those where the design of the equipment or network has no, or very limited, influence.

Video-related targets are not as well defined as voice targets, although the process of defining them is under way both at Nortel and in the industry. Figure 3 illustrates the typical end-to-end video service transmission, and highlights some of the potential video traffic network impairments.

Because video delivery involves many network elements, in order to achieve acceptable QoE, it is important to monitor video quality both at the individual network element level and at the end-to-end connection level.

Components in the video transmission chain can be classified into three categories: video acquisition and encoding; video packetization and transport; and video decoding and display.

Video acquisition and encoding components include video source, video encoding, and rate shaping.

The source of video can be a film, analog tape, digital storage, or live event (analog or digital). The quality of original materials greatly affects encoding efficiency and overall quality. Noise in the source materials wastes encoding bits and can affect quality. In addition, the source materials may be of varying resolutions and therefore varying quality to begin with.

Video encoding is accomplished using video codecs suited for the particular transmission method and capacity. Depending on the type of application, several parameters of video encoding are defined, including bit rate, Group of Pictures (GOP) structure, constant or variable bit rate (CBR/VBR), and frame rate. Currently for broadcast applications, MPEG-2 is widely used. However, MPEG-4 AVC (also known as MPEG-4 Part 10, H.264, JVT) and/or SMPTE VC-1 (Windows Media 9) are expected to gain market share since they offer significant bit rate reduction (up to 2x) for comparable quality over MPEG-2.

Video rate shaping (also known as digital turnaround or grooming) is required in the deployment scenario where the access network data rate is lower than the original source video coded bit rate or when the access network has links with different bit rate capacity. Rate shaping sometimes includes a transcoding step where MPEG-2 is re-encoded in MPEG-4 AVC or VC-1 and can also convert a variable bit rate (VBR) stream to a constant bit rate (CBR) or capped VBR stream to facilitate network engineering and constrain bandwidth requirements.

Video packetization and transport components include those for packetization, and the packet network, access network, and home network. Packetization occurs at both the MPEG level and the network transport level.

At the MPEG level, video programs can be packaged individually as Single Program Transport Streams (SPTS) or in groups as a Multiple Program Transport Stream (MPTS), each with MPEG transport packets of 188 or 204 bytes. SPTS is used in telco IPTV applications where only a single channel per TV is sent to the home because of access network bandwidth restrictions. MPTS is used in digital cable and satellite applications where all content is broadcast to each home simultaneously.

At the network transport level, the MPEG SPTS or MPTS streams are then further packetized in the format required for transport. Typically in telco IPTV deployments, IP transport is used with 7 MPEG packets per IP packet from the video head-end, and IP over ATM (AAL5) or Ethernet is used in the access network. In packet networks, issues of delay, jitter, and loss must also be dealt with. For broadcast TV applications, delays are generally not problematic, since there typically is a buffer of about 100 to 1000 milliseconds in the set-top box (STB) at the customer premise. However, video quality degrades severely with packet loss, as well as with the type of MPEG information lost. The use of intelligent priority-marking algorithms can preserve video quality as video traffic passes through the network.

The access network may consist of coax, copper, fiber, or fixed wireless network elements. A minimum of 12 to 24 Mbit/s is required to offer video services to the home. Currently, xDSL technologies (bonded ADSL2+ and VDSL) seem to be the most economical means of deployment. DSL data rate is a function of the copper loop length, with limitations resulting from crosstalk from neighboring copper pairs in a bundle, as well as from other noise sources (e.g., RF interference and disturbance due to lightning). The main challenge is to ensure that the FEC inter-leaver depth is more than the duration of impulse noise to achieve satisfactory BER.

The home network is another potential source of video impairment and is less well controlled (from the telco perspective) than the rest of the network. Often, video distribution in the home will be done using a separate physical network to isolate other home traffic from the video stream. The existing coax cable used for analog video distribution in the home is typically targeted with Ethernet over coax and analog RF used. Any packet loss specifications set to ensure video service QoE must be

Figure 3. Typical end-to-end video service transmission

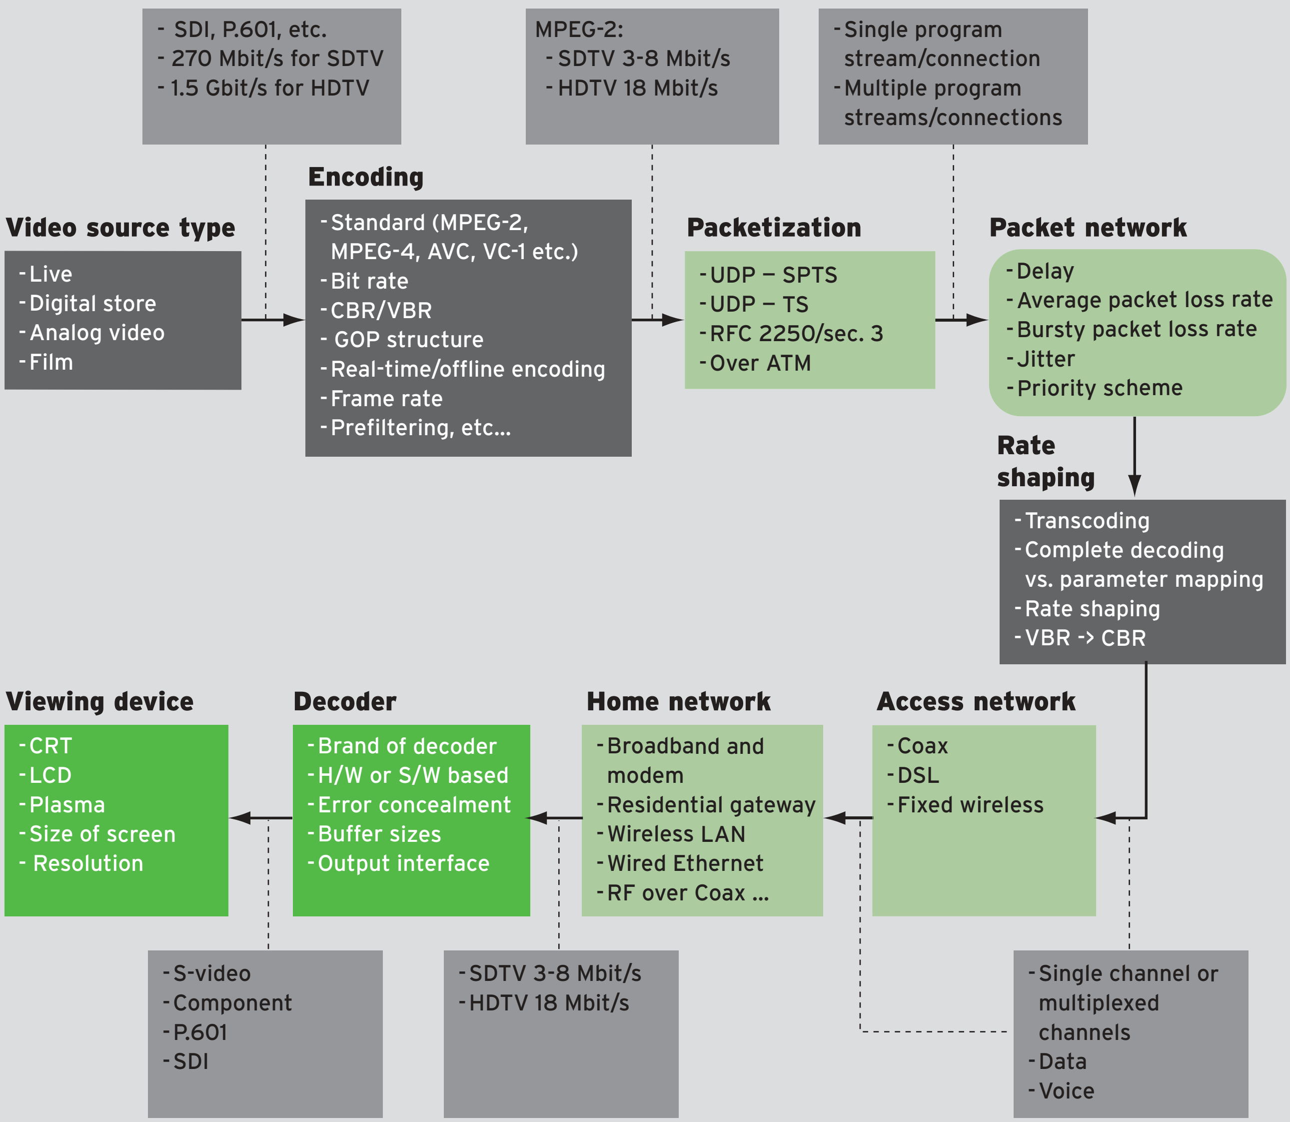

In the typical end-to-end transmission of video traffic, potential traffic impairments can occur at both the individual network element level and at the end-to-end connection level, including at the points of video acquisition and encoding (dark grey boxes), video packetization and transport (light green boxes), and video decoding and display (bright green boxes).

set from an end-to-end perspective – from the video head-end to the set-top box.

Video decoding and display components include the decoder and viewing device.

Video decoding is typically done by set-top box (STB) hardware, which also performs program stream demultiplexing and clock synchronization. Buffering in the STB (designed to compensate for network delay jitter), as well as error concealment algorithms employed at decoders, are important contributors to the resulting video quality at this stage. However, increased buffer size can negatively impact channel change time, while error concealment algorithms employed at decoders remain mostly proprietary.

Video display devices can be a significant factor in an end-user's perception of video quality. Issues such as type of screen (e.g., CRT, LCD, plasma), size of screen, and resolution can all affect perceived video quality. For example, the pixelization effect, as well as other impairments, are generally considered tolerable, if noticeable at all, on a standard TV, but tend to be more pronounced and objectionable when viewed on a large-screen, high-resolution TV.

Data-related targets have not been examined or defined as extensively as those for voice targets. While some attempts have been made to examine data application parameters in order to define targets (e.g., e-mail, web browsing, bulk data file transfer, and e-commerce), a more systematic taxonomy-based approach to defining data-related performance targets using a top-down model is required. Such an approach is required in order to deal with the increasing number, context, and usage scenarios of these applications. In addition, it is important to understand and be able to quantify the interaction effect of data applications, computing model, and devices. This measurement will become increasingly significant as deployment of converged networks gives rise to more devices and data applications. Data applications have also typically been delivered using best-effort services that limit a service provider's ability to meet target requirements. Converged networks will enable differentiation and priority treatment of different classes of data traffic.

The slower advancement of data target requirements may be due, in part, to the general perception that most data applications are delay-insensitive. From an OSI model technical standpoint, this may be accurate, but from a QoE perspective, the reality is quite the opposite, as users are increasingly expecting and demanding finite time-bounded data transactions. The QoE targets shown in Figure 2 were based on Nortel internal studies along with industry published results. Significantly more work is needed to ascertain acceptable targets for various data applications.

Interplay is critical

While it is imperative to understand these voice-, video-, and data-related targets, the understanding of the interplay of these various media and services in the network and how it impacts the end user's QoE will be instrumental to Nortel creating technology advantage and differentiation for its customers.

As well, convergence is being viewed from many different perspectives as increasing numbers of players enter the convergence space. Understanding these varied perspectives will enable Nortel to better position for and capitalize on the opportunities convergence presents.

What we are doing is essentially changing how people communicate. We are doing more than QoS or QoE; we are directly enabling an evolutionary change in how we interact with our world. To this end, Nortel continues to raise the bar of excellence through technological innovation, taking market-share and mind-share, and asserting global leadership in this era of network convergence.

Kathy Bharrathsingh *is a research scientist within the CTO Office, and is looking at the deployment and quality implications of future and evolving technologies, with a focus on end-user requirements.*

Transforming the way we manage networks

by Bill Bourne and Miguel Planas

As convergence becomes a reality and service providers begin to transform their networks, designers can finally address a long-recognized network management challenge and start to overhaul the way networks have traditionally been managed. Ultimately, as network management is driven into the network itself, this capability will become just another IT application rather than an entirely separate — as well as costly and cumbersome — infrastructure. With this approach, operators will be able to focus their operations support systems on managing their business and their customers, reduce operations costs, and enable profitable new services. Nortel is ideally positioned to lead the charge in this transformation.

Inside the network operations centers of typical large service providers sits a vast array of workstations that provide the software needed to manage their networks.

Traditionally, network management systems have been developed separately for each network, network node, network element – and often for individual functions within a single network element – as well as for each service and business application. These systems for the most part do not interwork – each has its own interface protocols, operational models, and "look and feel." A network provided by a single vendor, then, can involve many different network management systems, as well as many different workstations. Factor in multiple networks with gear from several vendors, and you can appreciate how service providers must deal with a tangled and expensive array of isolated, self-contained systems to manage their networks.

No wonder that today's network management offerings have earned the amusing moniker of "swivelware," a term that refers to the way network operators must use their swivel chairs to move between their console stations to operate, control, and monitor all the different management systems. It's even reported that one service provider went so far as to purchase sneakers for those of their employees who must physically travel between the many different stations in order to manage and operate the network.

Network convergence represents a long-awaited opportunity to fundamentally change the way networks are managed. Unifying and simplifying the network infrastructure opens the way to unifying and simplifying the way the network is managed.

Underlying this humor is an urgent and serious concern – managing the network in traditional ways has become far too complicated and costly. Service providers typically invest more heavily in managing the business and its services than they do in network infrastructure (Figure 1). In fact, the critical path – and majority of the cost – in offering any new service is not the deployment of the network infrastructure, but rather the modification to large complex operations environments required to support the new service.

While the industry has long recognized that network management needed to be transformed, addressing the challenge in the face of multiple networks and many different types of equipment was cost-prohibitive.

Until now. Network convergence represents a long-awaited opportunity to fundamentally change the way networks are managed. Unifying and simplifying the network infrastructure opens the way to unifying and simplifying the way the network is managed. Convergence also provides the opportunity to turn the traditional technology-centric view of network management on its head (see sidebar).

Already, providers have begun to position themselves for this opportunity. One key customer, for instance, is transforming its national network with a converged network and expects to save significant operations costs, including in network, services, and business management systems.

Indeed, many large service providers today have hundreds of different operations support systems (OSSs), many internally developed or custom built, and are expected to make substantial investments in both network infrastructure and OSSs to transform network management. In doing so, providers want to direct these systems to manage the business and customer relationships. Today's OSSs are not able to do so because they have to perform too many network management functions. Service providers do not expect equipment providers, such as Nortel, to provide customer-centric OSSs. These are built by vendors focused on such areas as business management, accounting, and customer relationship management (CRM) applications. However, network vendors can help tremendously in off-loading the management of the network infrastructure from these OSSs. In fact, service providers increasingly view this functionality as part of the network itself, and not as applications external to the network.

Nortel is introducing a common network management architecture for all its network solutions, where all the technology layers that have evolved over the years to manage the network collapse to just two: the network transport and data plane, and the network control and services plane.

Moving management into the network

To address this need, Nortel is introducing a common network management architecture for all its network solutions (Figure 2).

In this new architecture, all the technology layers that have evolved over the years to manage the network collapse to just two: a layer to manage the network's transport and data plane (connectivity), and one to manage the network's control and services plane.

Managing the transport and data plane involves functions that keep track of such information as network inventory and topology; network state (e.g., faults, alarms, equipment and link states, performance, utilization, and loading); management user security (authentication, authorization, accounting); network control (configuration and provisioning); and test and diagnostics.

These capabilities traditionally have been provided by the myriad of separate systems and worksta-

Figure 1. Traditional network management: complex and costly

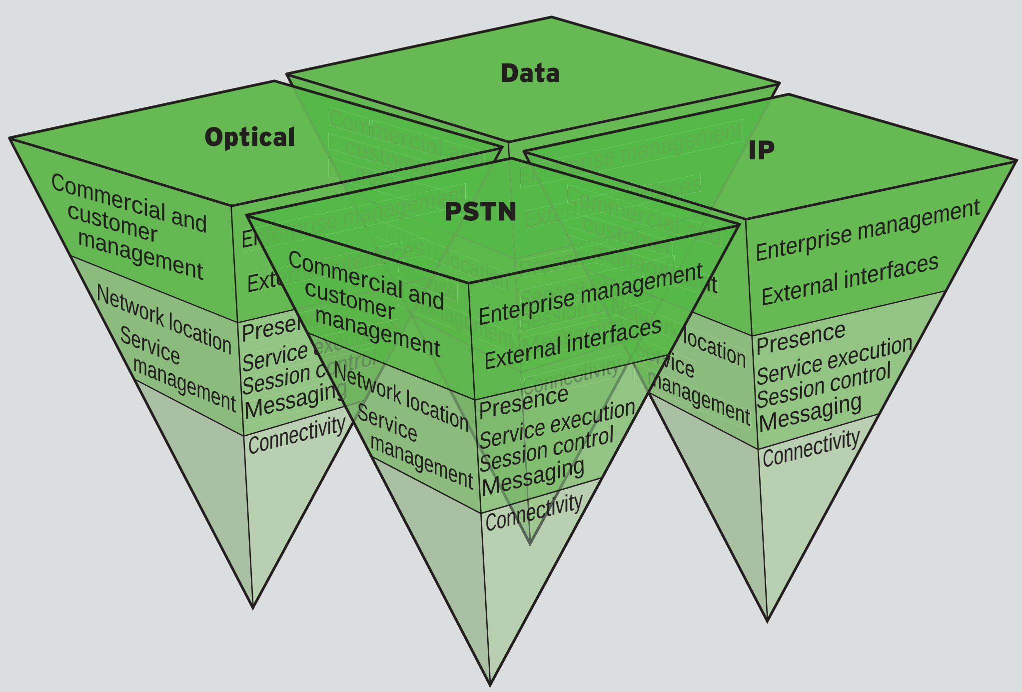

Today, operators must invest in and manage multiple network management systems, each one designed to manage a specific type of network.

tions described at the beginning of this article. Increasingly, though, these capabilities will move into the network itself. As technologies such as IP, along with ever-lower hardware costs, make bit transport more and more commoditized, automated, and competitive, vendors will need to seek new sources of differentiation. Moving network management functions into the network will enable equipment providers to add value beyond standard bit transport. These values could include, for instance, lower deployment costs, a higher level of automation, and increased efficiencies, as well as such functionality as automatic configuration.

Managing the network control and services plane is a key area that focuses on applications, such as setting up and controlling multimedia and conferences, and managing a user's presence across the network. These high-value software-based applications are key to unleashing the power of the transformed network and represent the next "big thing" in the network. These applications represent a new and growing area from which equipment providers such as Nortel are expected to derive an increasing percentage of their business in the future.

A service provider's customers will interact with this network services and control plane in two ways:

• Customers will sign up for services, and set up and maintain their business relationship with the service providers via service portals that interact with the provider's OSS. In turn, the OSS will send service control requests to the network via the network's control and services plane. These portals may be customer self-service web portals or service provider call centers.

Turning traditional network management upside down

Service providers cannot get value from network transformation without fundamentally changing the way the network is managed. To do that, the technology community must embrace a new approach to network management and discard the established architectural models that have been in place for decades. Diagram A shows the network management architecture of the Telecommunication Management Network (TMN) model, defined by the International Telecommunication Union (ITU-T).

Diagram A. TMN network management architecture

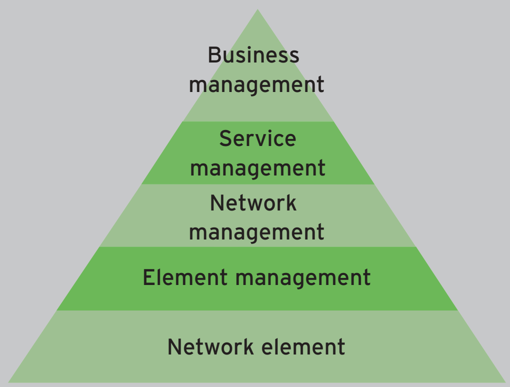

The model specifies five layers of management functionality, putting the network infrastructure (network element layer) as the foundation for all other layers and, therefore, of the greatest importance. In this technology-centric view, everything is centered on the network. This popular image also implies a world in which service providers invest most heavily in a network infrastructure (bottom layer) and progressively less on a series of management layers, culminating in a small amount of investment focused on managing the service provider's business.

This traditional network-centric approach does not represent today's reality for service providers.

The actual investment that service providers make, both in technology and labor, turns this model upside down (Diagram B). In reality, providers spend significantly more in managing the business and the services than they do on the physical infrastructure. And, since they must deal with fragmented network infrastructures, today's providers are faced with not one, but many, of these "upside-down" pyramids of investments, forcing them to view the network from a customer and business perspective, rather than from a network functionality perspective.

Diagram B. Today's network management reality

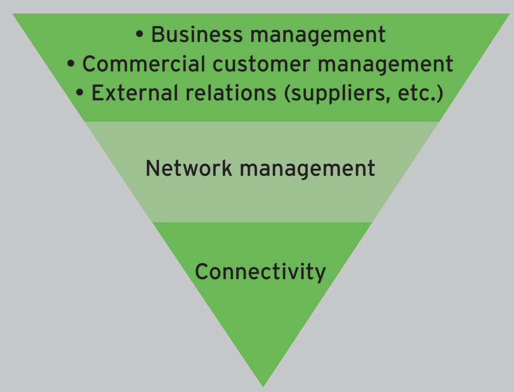

Figure 2. Nortel common network management architecture

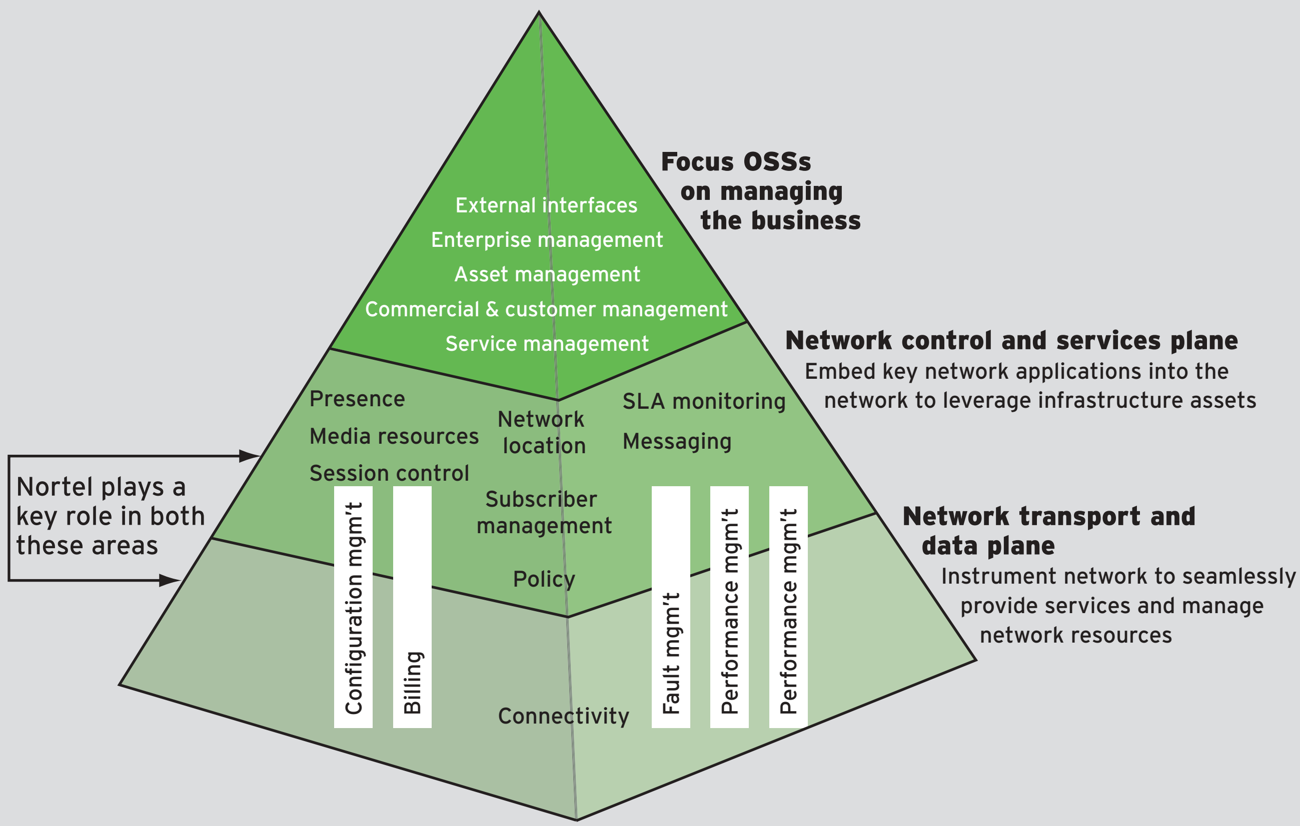

With the Nortel common network management architecture, all the technology layers that have evolved over the years to manage the network collapse to just two: the network transport and data plane; and the network control and services plane.

• To access and/or manipulate their individual services, customers will interact directly with the network service and control plane, rather than via portals to the OSSs. For instance, users would not expect to have to access a web portal to change the channel on their home theatres. In this case, the network services may need to interact with the OSS, to obtain service registration and authorization data, provide usage and billing data, and so on.

Nortel is addressing these two key areas, and working to ensure that Nortel equipment and solutions, and the applications that manage them, communicate in standard ways and implement common controls.

To this end, Nortel has commissioned several technical teams to define a number of common components, including a common alarm record, common performance record, common NE (network element) inventory model, and common human-machine interface standard. As well, Nortel is in the process of building and deploying the technology to support these components, several of which have already been implemented in various network solutions.

The intelligent network

As software functions – previously considered to be out-of-network applications – move down into the network, the network becomes more intelligent and "self-operating," allowing service providers to simplify their operations environments and reduce costs. For example, policy control and provisioning will be handled in-network through the 3GPP R5 IMS Policy Decision Function, which will be implemented by Nortel's Session Policy Controller (SPC) product.

As well, the migration of network management functionality from OSSs into the network opens a major opportunity to both simplify and strengthen the linkages between the network and the service provider's OSSs.

For instance, with the network able to perform network management itself, OSSs will not need to concern themselves with the technical details of correcting faults or extracting data from the network. Therefore, the number of commands from the OSSs to the network can be greatly reduced, and the messages themselves significantly simplified. In addition, data transmitted to the OSSs from the network can be abstracted to remove details that are no longer needed by the OSSs.

Nortel's opportunity

Nortel understands networks, including transport (Optical, Ethernet, ATM/frame relay, IP, MPLS, wireless) and network services and applications (e.g., soft switches and call servers, multimedia messaging, subscriber identity, and signaling). Nortel also understands network reliability and robustness. Since most high-value network management applications not only require an intimate knowledge of the network infrastructure, but also must be made highly reliable if they are to be allowed to directly control the network, Nortel is in a unique position to supply key functionality as it migrates from the OSS environment into the network.

For instance, we have significant expertise in effectively managing policy controls, session and messaging controls, subscriber identity, presence and location, services registries, dynamic congestion control and bandwidth management, lawful intercept, and hacking/fraud detection and prevention, to name a few areas.

The transformation of network management also provides an opportunity for Nortel to communicate to customers the significant values that network management brings, particularly in a converged network. To this end, Nortel could begin to offer various network management functionalities, previously part of individual products and platforms, as a variety of distinct software products, bringing new value to customers.

For services providers, the fundamental transformation of network management provides vast opportunities for them to offer new value to their customers. Moreover, a fully unified converged infrastructure is a key enabler for service providers as they seek to re-invent the way they manage their business to capitalize on these opportunities.

Bill Bourne *is the project leader for the evolution of the Carrier OAM applications architecture.*
Miguel Planas *is project leader for the NE OAM Design Authority, as well as for the Network Assurance and Operations Management Integrated Project Team.*

Nortel's network: Demonstrating the benefits and challenges of convergence

by Don Dixon

The Nortel Information Services (IS) organization has spent nearly two decades on the front lines of convergence, implementing Nortel solutions on the company's internal network and building a working demonstration of the advantages and opportunities that convergence brings. Throughout this time, the IS team has come to understand that convergence is not just about bringing together specific technologies, but also about the impact that these converging technologies have on the business, its operations and services, and on employees themselves. This experience has given Nortel valuable insight into the convergence challenges that many of its enterprise customers are facing today.

Convergence is not new, nor is it only about the integration of voice and data onto the same network. Indeed, throughout the industry, and even throughout Nortel today, you are likely to hear many different takes on convergence, including:

- "Firewalls are now doing routing, and routers can act as firewalls."
- "Voice switches now support IP traffic."
- "Computers now manage voice mail and telephones manage e-mail."
- "The line between the Nortel enterprise network and the public network is blurring."
- "Palm devices and PCs can do many of the same functions because both run converged applications."
- "Network access is now available anytime, anywhere, over many types of devices."
- "Network security and physical security are coming together."
- "Managing the network is as much about applications and end users as it is about devices."

Convergence is all of these, which is why Nortel's Information Services (IS) Technology organization defines convergence as the coming together of multiple services (operations), technologies (infrastructure), and organizations to enhance cost efficiencies, collaboration, mobility, productivity, network management, and security. This definition reflects the experience that the IS team has gained, through implementing new technologies and solutions as they have unfolded and addressing the challenges associated with that implementation.

Convergence in Nortel's own network

Running multiple services on the same network is not new; in fact, different traffic types began sharing the same systems decades ago. Although the business case at that time was based on voice – with data riding for free – the goal of leveraging technology to simplify the network and reduce costs was the same as it is today.

In the late eighties and early nineties, the introduction of ATM (Asynchronous Transfer Mode) and routing platforms gave rise to several challenges, particularly with understanding the complexities of new technologies and with bringing together disparate engineering and operations teams. Even simple terminology, for example – which still remains somewhat different today between the voice and data environments – often created confusion and misconceptions. As well, control over such issues as who would operate the equipment and carry the budget became often-asked questions.

Challenges during that era faced the sales force as well: whereas sales teams were previously concerned with meeting only one decision maker, they also faced the challenge of selling solutions that crossed a number of domains and that required the coming together of decision makers from many areas (voice, data, computing, security, and applications). (This phenomenon continues to this day – the convergence of systems and applications and the growing complexity of solutions means that our sales teams must operate in an increasingly multi-faceted environment.)

But with these challenges also came opportunity. Nortel's IS managers quickly realized the critical need to share knowledge and, where possible, resources. Organizationally, the voice and data teams were brought under one umbrella, both regionally and globally, in order that an entire view of IT networking costs, staff, and skills could be better understood. This teaming marked the beginning of an ongoing cycle of organizational changes brought on by technology convergence.

In the late 1990s and early 2000s, the introduction of Optical Ethernet (OE) advanced network convergence further. After trials and cost analyses, which included many discussions with carriers, OE was deployed to transport both voice and data traffic, first in North America and later in Europe. Since that initial deployment, OE has helped reduce the WAN cost per megabit by 68% and overall network costs by 78%, while quadrupling capacity and simplifying the network architecture between regions – from 559 virtual connections then to 164 today.

As well, the OE implementation reduced bandwidth costs and enabled consolidation opportunities in the applications and computing environments. Furthermore, OE and convergence eliminated many previous barriers – such as cost of data transport and speed of access – ultimately allowing for the elimination of some 9,000 servers and 1,200 applications instances in the past five years.

These results not only encouraged ideas on how convergence could dramatically reduce costs, simplify operations, and allow Nortel to take greater control of its own network, but also highlighted the challenges involved with converging disparate systems. For example, products that typically didn't interwork demanded interoperability, and standalone network management products demanded umbrella systems that gave an overall view of the complex network. And, while the cost of transporting voice and data may have been reduced, the skills and systems needed to engineer and operate the network grew.

The IS experience not only encouraged ideas on how convergence could dramatically reduce costs, simplify operations, and allow greater network control, but also highlighted the challenges involved with converging disparate systems.

By the year 2000, though, the potential for convergence to achieve greater cost reduction and efficiencies far outweighed these challenges, and was spurred by the cost challenges the company was facing at that time. However, network costs do not drop linearly with the reduction in employees. A network's fixed costs typify a step function, often related to product lines and geographies. The IS challenge, then, was to move from fixed to variable network costs, while allowing greater network access by partners, customers, and suppliers; meeting employees' mobility requirements; and fulfilling the growing demand for increased bandwidth. Today, there are some 500,000 customer users and 12,400 supplier users connecting into Nortel's network, along with 30,000 Nortel employees, many of whom are connecting remotely (see page 44 for more on the capabilities of Nortel's IS infrastructure).

The innovative thinking brought on by the necessity for cost reduction soon led to the evolution of new services and, in early 2000, the first Voice-over-IP trials began. Driven by the desire to reduce employee expenses and increase the mobility and productivity of the sales teams, this new service uncovered opportunities that were soon exploited in other areas; VoIP, for example, enhanced remote access capabilities and is now available to more than two-thirds of the population.

At the same time, the IS team recognized that, in order to better analyze the capabilities and potential of the converged network from a business perspective, they needed to pull together the thinking within the different IS technology teams. A core team of data and voice professionals were brought together into a converged New Product Introduc-

Running Nortel on Nortel: by the numbers

Today, Nortel's internal network spans 62 countries, links 247 locations around the globe, and supports approximately 42,000 desktop units and 2,500 servers. In addition to 30,000 employees, the network provides access to some 500,000 customer users, along with 12,400 users from the supplier community, representing more than 425 companies.

On a monthly basis, this network handles:
- 48.2 million e-mail messages
- 80,000 reservationless audio conferences
- 37.8 million telephone voice minutes (70% packetized; 30% public)
- 250 webcasts, live and on-demand
- 1.84 petabytes of routed data traffic

IT Infrastructure: doing more with less
- Reduced number of servers from 12,000 to 2,500
- Reduced number of applications from 2,700 to 1,500
- Since 2000, reduced annual IT budget by US$1.3 billion
- Over past three years, Optical Ethernet technology helped lower overall network costs by 78%, while increasing capacity by 390%

Security at every layer
- Perimeter and network layers include firewalls and routers, intrusion detection, and spam filtering
- Application and desktop layers include corporate directory, single password management, identity management, OS management, and remote automatic updates and patching for desktops
- Anti-spam infrastructure stops 90%-95% of inbound spam

Making work a thing you do, not a place you go
- More than 66% of the workforce has ubiquitous mobility
- Secured mobility environment includes remote access, VPNs, WLANs, IP telephony, and MCS collaboration services
- Teleworking program avoids US$22 million in annual real estate facility costs
- Flexible working capabilities help maintain productivity during disruptive events
- Nortel teleworkers report 15% increase in productivity

Wireless local area network
- Nearly 1,000 WLAN access points deployed to more than 100 locations
- Built-in WLAN for all new laptops (802.11 a/b/g)
- Large outdoor areas covered by wireless mesh networks
- 2%-3% productivity improvement thanks to WLAN
- Used by more than 8,000 Nortel employees
- Fewer moves, adds, and changes is saving an estimated US$950,000 per year

Rich multimedia environment
- More than 3,000 webcasts in 2004 (25% growth over 2003)
- 2.3 million webcast participants in 2004 (50% increase over 2003)
- Eliminated 95% of traditional videoconferencing rooms
- In-house multimedia studios save an estimated US$35 million each year
- eLearning capabilities are saving customers an estimated US$18 million in annual travel costs

Highly productive R&D environment
- Reduced software complexity through code consolidation into a single catalog, reusable between products
- Tools that integrate easily with suppliers while ensuring security at highest level
- Improvement in software compile and build (loadbuild) times – from 2.5 hours to 7 minutes – boosts designer produc-tivity

NPI: our network as a living lab

In 2004:
- 3,235 software loads introduced
- 116 new hardware items introduced
- 1,257 CSRs entered
- 127 trials completed; 167 ongoing

tion team, with the mandate to identify trial opportunities for voice and data products, as well as those for converged applications and services that could ride on the IP network. This team has since grown to include skills from all portfolios (wireless, optical, wireline, and enterprise) and additional technologies, such as computing, security, and network management. The team is responsible for more than 200 hardware and software product introductions annually.

Post year 2000, the converged network, per se, had almost become an accepted part of doing business. With IP as the convergence building block, it became easier to deploy converged applications quickly and efficiently.

Specifically, we saw an extension of the network itself, converging into the public domain. Employees started to expect ubiquitous access into the enterprise, from either public or private networks, and they wanted access to their corporate applications, messaging, and webcast services while traveling in other countries. These demands led to a new kind of convergence – a convergence of services, applications, and devices that has not only provided operation and cost benefits, but has also enhanced productivity.

Convergence challenges ahead

Today, Nortel's network environment is primarily IP-based, with more than 70% packetized private voice versus 30% public voice. While business needs are being met, however, we still face challenges. For example, the convergence of services, devices, applications, and networks gives rise to concerns around complexity and security. Managing the network, with its significant breadth and depth, takes sophisticated network management and security capabilities that can prevent, detect, and resolve failures and threats. The network must also allow Nortel to cultivate partner, supplier, and customer relationships, while providing appropriate protection of intellectual property.

To address these challenges, the IS team has once again repositioned the organization and created teams that combine the following expertise:

- voice and data, to reflect the converged network environment;
- collaboration tools, including conferencing and video, to focus on the convergence of multimedia applications onto the same systems;
- messaging and information exchange, to reflect the converged way (device-independent) in which information is now shared; and
- centralized wireless program management, to focus on private/public handoffs and wireless VoIP.

As we head into the latter part of the decade, convergence is reaching closer to the desktop and end-user device, and is expanding in such areas as voice and data in the LAN; communications and business applications (e.g. click-to-call); and computing and communications. Over the next 18 months, we expect to see other changes:

- For employees, convergence will bring greater transparency between public and private domains, as well as a "limitless" network – one that will enable employees to work at any pace, at any time. Simple user interfaces on any device will be in demand regardless of the application's complexity.
- For IS services, it will mean increased mobile access using VoIP, voice over WLAN, video over IP, whiteboarding, and conferencing capabilities.
- For the IS staff, it will mean the continued convergence of skills and teams. Information sharing among the various technology teams will be even more critical.
- For IS operations, it will mean a stepped-up focus on security and network management, along with engineering for QoS interoperability, and new technologies such as storage area networks over IP.

The applications that bring us the information that we demand so immediately will also converge, and will present opportunities to consolidate information in ways that make it easiest for end users to get what they want, when they need it. And at Nortel, the evolution of convergence will, once again, mean the evolution of our network, services, organizations, and processes.

Don Dixon *is Leader of Technology Services in the IS organization.*

>THIS IS **THE WAY**

>THIS IS **NORTEL**