Packet Optical Networking: Evolving the Metro Infrastructure with Fujitsu

● Paul Havala ● Sam Lisle ● Jun Tanaka ● Kazuyuki Miura

(Manuscript received April 10, 2009)

Network providers have relied on optical networks as scalable, low-cost, general-purpose aggregation and transport infrastructures. The emerging communications services environment is placing significant stresses on metro networks, including the aggregation and transport infrastructure. As the existing one evolves to relieve these stresses, it must address three critical challenges: providing significantly more bandwidth at lower cost, delivering a scalable Ethernet infrastructure, and managing the simultaneous growth of time-division-multiplexing and Ethernet traffic. Fujitsu is helping its customers to evolve their aggregation and transport infrastructures to face these challenges. In 2008, it introduced FLASHWAVE 9500, the first in its planned family of Packet Optical Networking Platforms (Packet ONPs). Packet ONPs represent an emerging class of optical networking equipment that integrates several important new photonic, electronic, and software technologies. They meet the three key challenges while retaining the "low-touch" operational capabilities of traditional optical networking equipment, thereby allowing service providers to evolve their metro infrastructure to profitably deliver the new generation of Internet protocol services.

1. Introduction

Network providers have relied on optical networks as scalable, low-cost, general-purpose aggregation and transport infrastructures. Optical elements provide significant economic value and deliver operational efficiency by embodying "low-touch" operational capabilities including simple, in-service software upgrades, software backward compatibility, robust management interfaces, and network element fault tolerance.[1)-4)] These operational capabilities have enabled the massive geographical scaling necessary to connect the tens of thousands of access points in a metropolitan (metro) area with the handful of sophisticated service elements located in the metro core, as depicted in **Figure 1**.

The emerging communications services environment is placing significant stresses on metro networks, including the aggregation and transport infrastructure. To meet the demands of multimedia applications, network providers are delivering unprecedented amounts of bandwidth for residential, enterprise, and mobile services. Although these services are increasingly Internet protocol (IP) services, the transport bandwidth comprises a complex mix of traditional time-division-multiplexing (TDM) and Ethernet traffic, both of which are experiencing significant growth. As the existing metro optical aggregation and transport infrastructure evolves in response to these stresses, it must address three critical challenges:

- Providing significantly more bandwidth at less cost. By many estimates, bandwidth is projected to grow at over 50% per year for the next several years.
- Delivering a scalable Ethernet infrastructure.

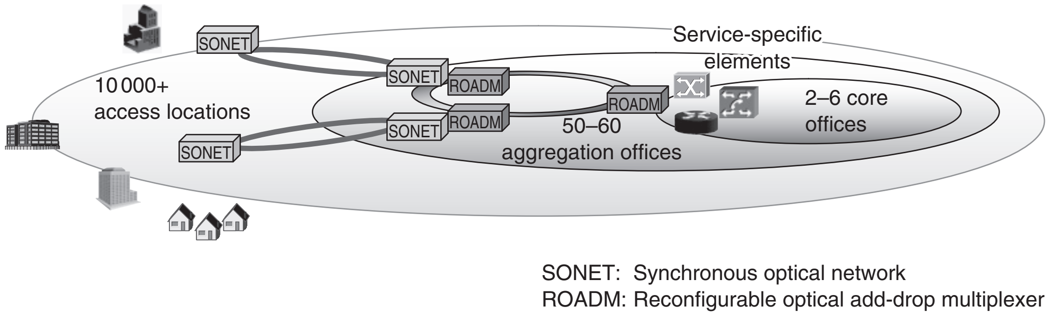

Figure 1
Classic metro networking: leveraging optical networking for distributed aggregation and transport.

Ethernet has considerable attractiveness as an inexpensive and flexible interface, and Ethernet-based services are growing rapidly. Since Ethernet has evolved from a local area network (LAN) to a wide area network (WAN) technology, service providers must build an infrastructure that not only transports Ethernet, but also efficiently aggregates Ethernet traffic to create efficient handoffs to the IP service edge.

- Managing the simultaneous growth of both TDM and Ethernet. While Ethernet-based access to Ethernet and IP services continues to grow rapidly, TDM traffic—especially TDM circuits used as access for IP/MPLS virtual private network services—also continues to grow (MPLS: multiprotocol label switching). Ethernet and TDM traffic both require a lot of aggregation, but have traditionally relied upon diverse technologies. Network providers can ill afford to deploy overlay aggregation networks.

Fujitsu is helping its customers to evolve their aggregation and transport infrastructures to face these challenges. In 2008, it introduced FLASHWAVE 9500, the first in a planned family of Packet Optical Networking Platforms (Packet ONPs). Packet ONPs represent an emerging class of optical networking equipment. They meet the three key challenges while retaining the "low-touch" operational capabilities of traditional optical networking equipment, thereby allowing service providers to evolve their metro infrastructure to profitably deliver the new generation of IP services. What distinguishes Packet ONPs from previous-generation network equipment is the integration of several important new photonic, electronic, and software technologies.

2. Important new technologies

2.1 Enabling photonic technologies

Photonic technologies enable metro networks to scale in bandwidth and geography. The flexible photonic layer based on reconfigurable optical add-drop multiplexers (ROADMs) can eliminate 90% of the capital cost of passing traffic through intermediate offices and has uniquely enabled service providers to deploy metro bandwidth of hundreds of gigabits per second to support high-bandwidth services. ROADMs provide significant operational benefits by enabling a service provider to deploy new wavelengths by remote provisioning rather than by dispatching a technician to each site along the network to manually patch the traffic. Innovations such as dynamic transient response amplifiers were a critical building block in establishing the photonic layer as the transport convergence layer for all services.[5]

Stand-alone ROADMs must currently be stacked alongside stand-alone synchronous optical network, synchronous digital hierarchy

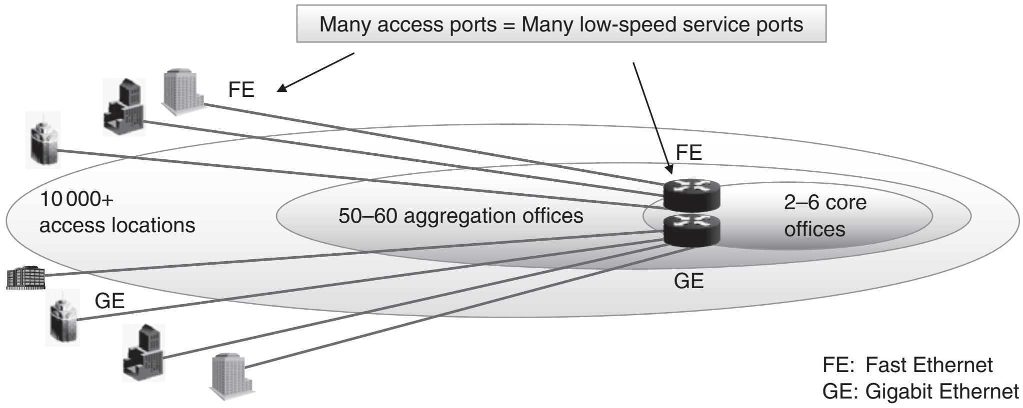

Figure 2
Ethernet aggregation challenge.

(SONET/SDH) and Ethernet elements to provide sub-wavelength grooming and aggregation and make efficient use of each wavelength. As bandwidth continues to grow, it becomes increasingly important to tightly integrate this ROADM-delivered photonic layer with TDM and Ethernet sub-wavelength grooming.

One critical enabling technology is the "pluggable ROADM". Ongoing density enhancements are now enabling the integration of amplifier units and wavelength-selective-switch-based optical fabrics onto smaller form factors. As these form factors improve, the ROADM hardware becomes a pluggable option in networking platforms that perform sub-wavelength grooming. This tight integration removes substantial capital cost and operational complexity. It is no longer necessary for network providers to deploy separate transponder hardware or other interfaces between stand-alone ROADM elements and subtended elements.

Pluggable ROADM hardware improves ROADM simplicity by integrating more functionality on fewer hardware units and greatly reduces the amount of fiber that must be installed. Further automation enhancements, such as integrated power and connectivity tests and in-situ span measurements, can greatly reduce the operational cost associated with extending the photonic footprint.

2.2 Enabling electronic technologies

The challenge of the simultaneous growth of TDM and Ethernet is unique in the optical era. Ethernet is attractive as a customer interface because it is elastic and enables the end user to increase bandwidth or add more connections without requiring the network provider to visit the site. Historically, switched Ethernet services networks have been deployed with little or no aggregation—with customer ports being tied directly to service element ports, as depicted in **Figure 2**. Access is sometimes provided over a Layer 1 infrastructure that efficiently connects each customer port to the service element. This approach alleviates the significant fiber costs of a pure "home-run" architecture, but because Layer 1 infrastructures lack Ethernet aggregation, this results in "port overload" on IP and Ethernet service elements. Widely distributing service elements for aggregation and transport can result in high operational costs because these platforms do not provide the previously mentioned low-touch operational capabilities of optical elements.

TDM continues to be attractive because of its high reliability, ubiquitous availability, and clean carrier-to-carrier handoff. In particular,

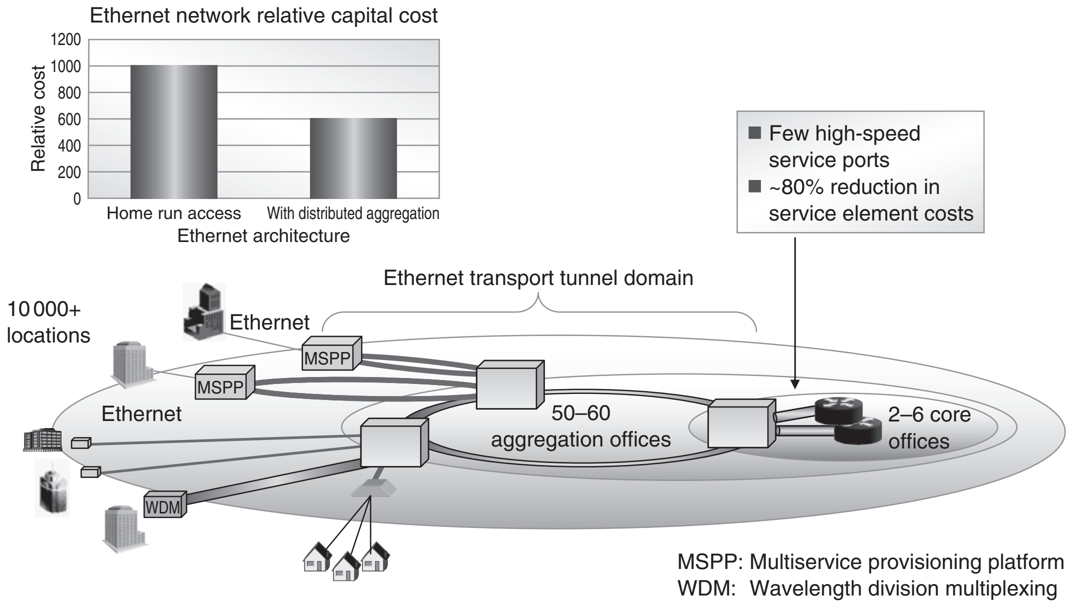

Figure 3
Distributed Ethernet aggregation using connection-oriented Ethernet tunnels.

circuits and revenue for TDM-based access to IP/MPLS virtual private network services continue to grow.

With both TDM and Ethernet network attachments growing significantly, network providers are forced to either deploy two aggregation networks using existing disparate technologies or to look for alternatives that enable a single infrastructure to economically aggregate both Ethernet and TDM traffic. Two key technologies address this difficult challenge: connection-oriented Ethernet transport and the universal electronic switching fabric.

2.3 Connection-oriented Ethernet transport tunnels

Ethernet as a connectionless technology cannot by itself deliver the connection performance and reliability needed to deliver a full range of services. It must rely on some underlying connection-oriented transport tunnel. Layer 1 transport of Ethernet is widely deployed and meets connection performance requirements, but it does not provide Ethernet aggregation or present an efficient Ethernet handoff to the IP service edge. Packet service platforms provide Ethernet aggregation, but they are complex elements with a high-touch operational profile that do not lend themselves to highly distributed aggregation infrastructure deployment and are not optimized for tunnel transport.

Several Layer 2 transport tunnel technologies exist or are under development including routed MPLS with Martini encapsulation,[6] provider backbone bridging with traffic engineering (PBB-TE),[7] and MPLS transport profile (MPLS-TP).[8] Non-routed, circuit-based approaches such as PBB-TE and MPLS-TP minimize the software complexity of aggregation elements and fit within the current circuit-based operational environment of the aggregation network.

Connection-oriented Ethernet transport technologies can realize a distributed aggregation and transport network for metro Ethernet that reduces network costs while achieving the necessary quality of service (QoS) and circuit protection required for a general-purpose scalable infrastructure. This type of network is illustrated in **Figure 3**.

2.4 Universal electronic switching fabrics for managing TDM and Ethernet growth

Universal electronic switching fabrics that switch both circuit and packet traffic in their native formats are vital to delivering significant TDM and Ethernet circuit growth on a single converged network.[9] Universal switching fabric technology enables a level of network element integration and convergence that has not been possible in previous generations of equipment. ROADM and SONET elements alike have included interface cards to adapt Ethernet to the SONET or ROADM environment.[10] This has been termed card-level integration. Card-level integration is an excellent fit for spoke site applications, but because there is no centralized packet fabric there can be no Ethernet traffic aggregation or grooming across the interface cards. Therefore, network solutions still require subtended Ethernet elements with large fabrics to provide the necessary aggregation at network hub sites.

By contrast, a single universal electronic switching fabric offers an elegant solution to this problem. With recent advances in application-specific integrated circuit (ASIC) technology, it is now possible to develop an integrated fabric unit that 1) identifies input traffic as either circuit or packet, 2) switches the incoming traffic natively, and 3) assigns the switched traffic to an outgoing set of interfaces. Since this type of fabric is built from highly integrated ASIC technology, the cost can be low. This chassis-level integration results in a modular system that does not penalize the network provider for unwanted functionality or force the network provider into predicting the ratio of Ethernet to TDM traffic. A universal fabric lets a network element become technology neutral, enabling a single infrastructure to manage both TDM and Ethernet growth.

Tightly integrating the universal electronic fabric with the pluggable ROADM lets a single chassis not only provide ROADM-based photonic transport, but also eliminate subtended stand-alone SONET and Ethernet aggregation equipment, as illustrated in **Figure 4**.

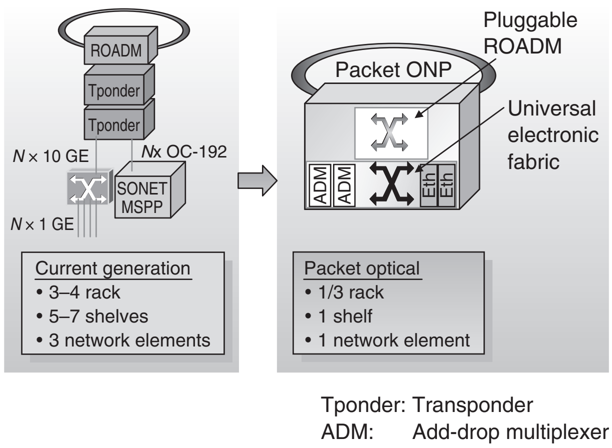

Figure 4
Packet-ONP-based scheme for integrating ROADM networking with universal sub-wavelength grooming.

2.5 Enabling software technologies

Central to managing and operating a converged element that supports both Layer 1 and Layer 2 technologies is the software concept of operational virtualization. Operational virtualization enables network providers with organizational and operations system boundaries between Layer 1 and Layer 2 management to simultaneously and cooperatively access and manage an integrated aggregation and transport infrastructure network.

Layer 1 and Layer 2 elements have traditionally been managed by Transaction Language-1 (TL-1) and the simple network management protocol (SNMP), respectively. It is therefore important for any Layer 2 functionality to be addressable via SNMP, while Layer 1 functionality continues to be addressable by TL-1. Some previous next-generation SONET elements supported multiple management protocols. What is new is the granting of authority to management users on the basis of the technology or layer for which they are responsible. Historically, network management authority

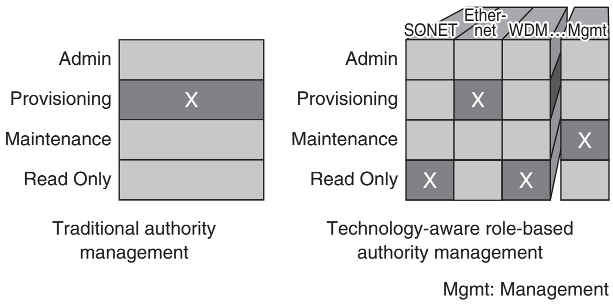

Figure 5
Traditional and role-based authority management approaches.

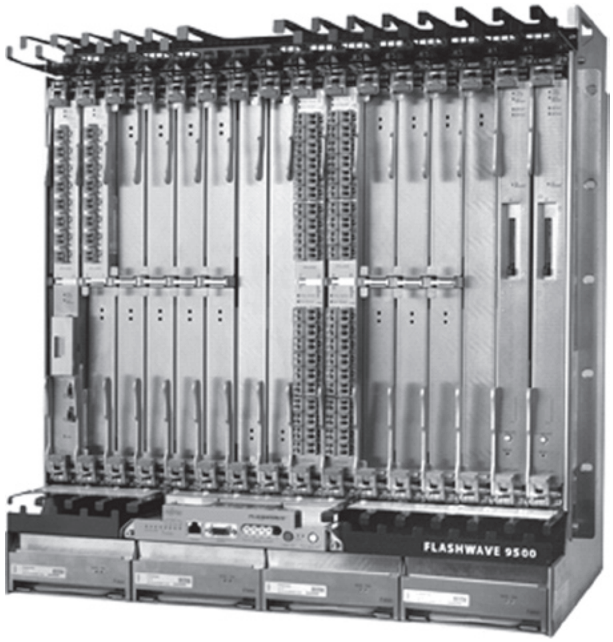

Figure 6
FLASHWAVE 9500 HDS.

has been granted only one-dimensionally. If a user had provisioning privileges on the element, then that user had provisioning privileges for all aspects of the element. In converged network elements with highly granular object models that clearly delineate between technologies, a role-based authority scheme allows individual users to be granted authority based on their sphere of responsibility, as depicted in **Figure 5**.

With a role-based authority scheme, it is now possible to support a range of network management applications that account for the diversity in network provider environments. As the example application in Figure 5 shows, users in a Layer 2 network operations center can be granted provisioning privileges for Ethernet transport while having only read-only privileges for the Layer 1 entities. In this example, the Layer 2 network operations and provisioning center can assume full responsibility for provisioning and maintaining Ethernet connections while receiving alarm and fault management information from the underlying Layer 1 network.

3. Fujitsu products

3.1 FLASHWAVE 9500 Packet Optical Networking Platform

Packet ONPs embody the key photonic, electronic, and software technologies described above to meet the emerging challenges of the terabit (Tb/s) metro network, Ethernet scaling with QoS and protection, and universal support for both TDM and Ethernet aggregation. Fujitsu recognized the emerging market for this new kind of system several years ago and began a dedicated program to design, develop, and deliver the industry's leading Packet ONP product family. Last year, Fujitsu introduced the first member of that family, FLASHWAVE 9500 HDS (see **Figure 6**). The introduction of this system represents the culmination of a worldwide development, with contributions from teams in Japan, the USA (Texas, New York, and California), and India. The collaboration, ingenuity, and dedication of this worldwide planning and development team resulted in a world-class Packet ONP with the following market-driven technology innovations.

- To economically scale all types of bandwidth, FLASHWAVE 9500 HDS features the industry's most compact pluggable ROADM. A single 13-RU FLASHWAVE 9500 HDS shelf can support a full eight-degree ROADM, a two-degree ROADM with 360 Gb/s of drop capacity, or a variety of other combinations (RU: rack unit). This compact, flexible implementation lowers costs and simplifies network operations.

- FLASHWAVE 9500 HDS features a 480-Gb/s universal electronic switching fabric that aggregates and grooms both TDM circuits and Ethernet connections with equal facility and can support Optical Transport Network (OTN) switching in the future. This enables network providers to economically deploy a flexible transport infrastructure that can grow to address all types of traffic needs, depending on market conditions and end user requirements.
- With support for connection-oriented Ethernet transport tunnels, FLASHWAVE 9500 HDS provides distributed Layer 2 aggregation to greatly decrease the cost of Ethernet and IP services networks. The connection-oriented nature of these tunnels, along with the high availability of the FLASHWAVE 9500 system, allows service providers to achieve the QoS, survivability, and availability required for a general-purpose scalable infrastructure. This, in turn, enables service providers to deploy a wide range of offerings, from best-effort services to private-line equivalent services.
- FLASHWAVE 9500 HDS provides a number of software features, including role-based authority management, that allow service providers to leverage and evolve the enormous existing optical network base while utilizing familiar operational procedures.

3.2 Packet ONP product family

Packet ONPs will enable service providers to evolve their metro infrastructures to more profitably deliver existing and emerging IP services. Over time, the entire metro infrastructure—from access to aggregation to core—must evolve. Successful network equipment vendors will offer not just a single Packet ONP product, but a complete family of products. Fujitsu is meeting this challenge by developing new products and augmenting existing ones to deliver a highly integrated product family, as illustrated in **Figure 7**. In addition to FLASHWAVE 9500 HDS, this product family includes the following.

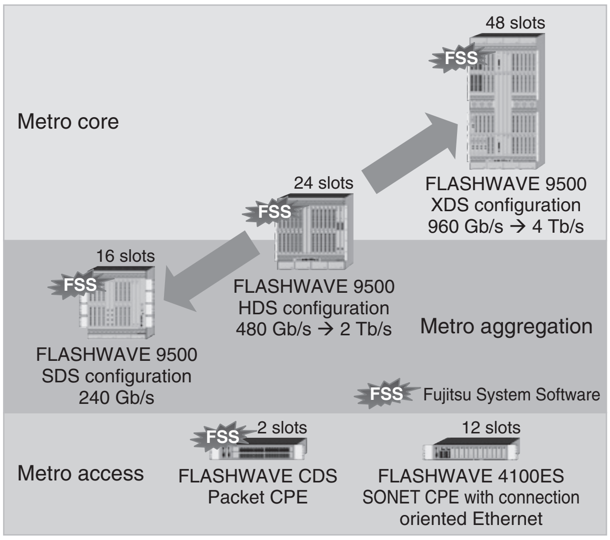

CPE: Customer premises equipment

Figure 7
Fujitsu's Packet ONP family.

- FLASHWAVE 9500 SDS and XDS. These shelf configurations broaden the market for FLASHWAVE 9500 to cover the metro core (XDS) and the lower end of the metro edge (SDS). All three systems (XDS, HDS, and SDS) run common software, called Fujitsu System Software (FSS), to ensure identical system behavior for a common feature set. Moreover, they are all supported by Fujitsu's craft tool (NETSMART 500), element management system (NETSMART 1500), and network design tool (NETSMART 2000).
- FLASHWAVE CDS. This compact metro access product extends TDM and connection-oriented Ethernet services to customer locations. It also runs the common FSS software and is supported by NETSMART

500 and 1500.

- FLASHWAVE 4100ES. This multiservice provisioning platform is deployed in the metro access networks of North America's largest service providers. It supports Ethernet over SONET today and a future release will add connection-oriented Ethernet services, which will complete its transition from multiservice provisioning platforms to Packet ONPs. FLASHWAVE 4100ES is supported by NETSMART 500 and 1500.

The high degree of commonality among the various products—common hardware units, software, and management system support—requires a tremendous amount of cooperation and collaboration from the worldwide planning and development team. This commitment to commonality provides significant financial benefits to Fujitsu: common hardware and software development and management system support reduce our overall development costs; moreover, the consistent system behavior that comes from common software lowers the costs of many related functions, including development of our own management systems, integration with third-party management systems (e.g., from Telcordia), and the production of technical publications and training materials.

Perhaps more importantly, the commonality across Fujitsu's Packet ONP family benefits our customers: a single element management system can be the building block for a service provider's entire metro transport network. Common hardware units lower inventory costs and identical system behavior streamlines operations and lowers associated costs such as training costs and the cost of integrating the product family members into our customers' operations support systems.

4. Market success

These customer benefits have resulted in early market success for Fujitsu's Packet ONP product family. In 2008, Fujitsu announced that Verizon had selected FLASHWAVE 9500 for wide deployment in North America. Verizon will use the platform to enhance its entire metro transport architecture by supporting and consolidating SONET, Ethernet, and wavelength service traffic for a variety of applications, including FiOS, Verizon's successful residential broadband offering.

Since then, FLASHWAVE 9500 has scored wins in North American service provider and enterprise networks, for applications covering TDM, ROADM, and connection-oriented Ethernet, with global deployments just around the corner. The future for Fujitsu's Packet ONP family looks promising.

5. Conclusion

As the existing metro optical aggregation and transport infrastructure evolves in response to the stresses placed on it by the increased volume and packet nature of emerging services, it must address three critical challenges: providing significantly more bandwidth at lower cost, delivering a scalable Ethernet infrastructure, and managing the simultaneous growth of TDM and Ethernet. Packet ONPs represent an emerging class of optical networking equipment that integrates several important new photonic, electronic, and software technologies to resolve these three network issues in a way that retains all the operational benefits of traditional optical networking equipment.

Fujitsu is offering a family of Packet ONP products, led by FLASHWAVE 9500, to help service providers migrate their entire metro infrastructure. A worldwide planning and development effort is delivering that family of products with a high degree of hardware, software, and management system commonality. This commonality translates into lower product introduction costs for Fujitsu and lower network deployment costs for its customers. Fujitsu's vision, ingenuity, execution, and spirit of collaboration have resulted in early market success and have laid the foundation for Fujitsu's continued industry leadership.

References

1) S. Lisle: The Broad (And Surprising?) Future of SONET. 2003 Technical Proceedings from the National Fiber Optic Engineers Conference. September 8–11, 2003. pp. 1634–1643.
2) L. Steinhorst, S. Han, and E. Koopferstock: SONET and DWDM: Competing Yet Complementary Technologies For The Metro Networks. Technical Digest of the Optical Fiber Communication Conference (OFC)/National Fiber Optic Engineers Conference (NFOEC), Anaheim, CA, March 2005, Session NTuH1.
3) S. Han, S. Lisle, and G. Nehib: Economic considerations drive choice of triple-play transport. *Lightwave*, Vol. 24, Issue 6. pp. 1–35 (June 2007).
4) S. Lisle: Controlling opex, availability in packet networks. *Lightwave*, Vol. 24, Issue 11, p. 13 (2007).
5) C. Tian and M. Kinoshita: Analysis and Control of Transient Dynamics of EDFA Pumped by 1480- and 980-nm Lasers. *Journal of Lightwave Technology*, Vol. 21, No. 8, pp. 1278–1734 (August 2003).
6) L. Martini, E. Rosen, N. El-Aawar, and G. Heron: Encapsulation Methods for Transport of Ethernet over MPLS Networks. *RFC 4448*, April 2006.
7) IEEE 802.1Qay IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks—Amendment: Provider Backbone Bridge Traffic Engineering.
8) B. Niven-Jenkins, et al.: MPLS-TP Requirements. draft-jenkins-mpls-mpls-tp-requirements-01, IETF, October 31, 2008.
9) S. Gringeri, et al.: Optimizing Transport Systems to Integrate TDM and Packet Services. in Optical Fiber Communications Conference and Exposition and The National Fiber Optic Engineers Conference on CD-ROM (Optical Society of America, Washington, DC, 2007), NTuA4.
10) S. Han: Architectural and Economic Impact of the Integration of SONET and DWDM Platforms. Technical Digest of the Optical Fiber Communication Conference (OFC)/National Fiber Optic Engineers Conference (NFOEC), Anaheim, CA, March 2006, Session NThA5.

Paul Havala
Fujitsu Network Communications Inc.

Mr. Havala received B.S. and M.S. degrees in Electrical Engineering from Michigan State University, USA. He has over 19 years of experience in the telecommunications industry and has served in technical marketing, business development, product line management, and senior technical roles at Bellcore (now Telcordia), DSC (now Alcatel-Lucent), and White Rock Networks. He joined Fujitsu Network Communications Inc. in 1996 and is currently the senior director of product planning, leading product line management for the FLASHWAVE 9500 product family, product marketing for Fujitsu's packet optical networking portfolio, and overall management of Fujitsu's data initiatives.

E-mail: paul.havala@us.fujitsu.com

Sam Lisle
Fujitsu Network Communications Inc.

Mr. Lisle received B.S. and M.S. degrees in Electrical Engineering from the University of Iowa and Georgia Tech., USA, respectively. He worked for Bellcore (now Telcordia), specializing in the reliability analysis of SONET transport systems and fiber optic media. He joined Fujitsu Network Communications Inc. in 1993 and is currently a market development director, focusing on packet optical networking technology and applications. He holds U.S. patents in packet-aware optical networking and was instrumental in the definition of the Fujitsu FLASHWAVE 4000 series of MSPP platforms.

E-mail: sam.lisle@us.fujitsu.com

Jun Tanaka
Fujitsu Laboratories Ltd.

Mr. Tanaka received B.S. and M.S. degrees in Electrical Engineering from Tohoku University in 1987 and 1989, respectively. Since joining Fujitsu Laboratories Ltd. in 1989, he has been engaged in R&D of MPEG coders/decoders, AAL2 standardization, real-time transmission over the Internet, and Ethernet transport systems. He is currently a senior researcher focusing on system requirement specifications of FLASHWAVE 9500. His research interests include datacenter networking, the next-generation Internet, and quantum communication theory. He is a member of IEEE and IEICE.

E-mail: tanaka.jun.777@jp.fujitsu.com

Kazuyuki Miura
Fujitsu Ltd.

Mr. Miura received a B.E. degree in Electronics and Communication Engineering from Waseda University in 1982. He joined Fujitsu Ltd. in 1982. He has been engaged in the development of transport and telecommunication systems and is currently General Manager of the FLASHWAVE 9500 project.

E-mail: miura.kazuyuki@jp.fujitsu.com