284 23-3129 Uen Rev A | August 2009

white paper

Evolution to optical packet transport

Transport networks are evolving from TDM to carry packets more efficiently. This evolution takes advantage of a range of options to deliver all-IP services, leading to Full Service Broadband: the delivery of multiple services over any bearer, to any client device. This paper outlines the role of transport options.

ERICSSON

TAKING YOU FORWARD

Contents

1	Introduction: IP networks – a lot more than just IP	3
2	Connectivity options and convergence across network layers	4
3	Connection oriented networks: Ethernet evolves	6
4	Ethernet packet protocol support options	8
5	Transport for Ethernet via MSPP	9
6	Mobile backhaul	10
7	Network synchronization	11
8	Bulk switching of transport capacity	12
9	Ring and mesh networks: automation by ASON control plane	14
10	Linking transport with IP/MPLS	15
11	Keeping things simple to minimize opex	16
12	Summary	17
13	Glossary	18
14	References	19

1 Introduction: IP networks – a lot more than just IP

Network operators have consistently delivered major network re-engineering to cope with the challenge that traffic grows much faster than revenue. As proof of their progress, if a Boeing 747 aircraft had reduced in price, since its commercial launch, at the same rate as bandwidth has done in the same period, the 747 would now cost the same as a pizza.

Data traffic has driven that change for more than two decades; now the pace has greatly increased, with the growth of video downloads and peer-to-peer (P2P). Through a phone line, the private user can now access bandwidth that would have supported a large company a decade ago. For business users, growth is boosted by online services and potentially soon by cloud computing.

In order to meet the goal of broadband everywhere, dominated by internet-based services, operators have moved to define other services also in terms of the internet packet protocol, for example IPTV. The resulting network transformation takes them into a next generation network (NGN). This is defined by the International Telecommunications Union – standardization sector (ITU-T) as providing services based on internet protocol (IP), with a network that supports multiple levels of Quality of Service (QoS) and is independent of the transport technologies that provide connectivity. This "all-IP" approach means there would ideally be no service-specific networks or "stovepipes" and this represents a major change.

In using NGN to separate services from connectivity, operators can grow their networks by building more transport capacity, based on broadly the same type of connectivity for every service, instead of building a new and different network for each service. This aligns with what Ericsson call Full Service Broadband [1], to support convergence, all-IP evolution, network transformation and new business models for broadband everywhere.

Transport is essentially about design for medium to long-term capacity – months to years – but its flexibility can be invoked to operate in less than a tenth of a second to overcome network faults. It covers a wide range of technologies, from IP platforms through laser optics and Ethernet switches, to microwave radio links. Although the focus of this paper is IP and packet transport, there is a substantial network requirement – often at highly valuable tariffs – to carry traffic that may not be identified as IP.

Note that "all-IP networks" refers to all services in IP packet format, but several other network technologies are used to transport the packets, for example Ethernet and wavelength division multiplexing (WDM) to name just two. Despite their name, all-IP networks need a lot more than just packet technology, and a lot more than just IP.

2 Connectivity options and convergence across network layers

So exactly what type of transport is right for supporting packet connectivity? This is a major question for operators, and the subject of this paper, which describes options and the reasons for each one. The problem is not really the lack of a solution, but insufficient understanding of the solutions that exist. Several exist because networks vary so much in terms of geography, traffic mix and existing infrastructure.

An important component of transport is WDM, which multiplies the basic capacity of optical fibers and is described in section 8. On top of WDM, there are multiple options for carrying packets, frequently offered in Ethernet format. Irrespective of the transport bearer used, Ethernet is increasingly supported by packet protocol options for connectivity and management, notably the still-evolving multi-protocol label switching (MPLS), which aids Ethernet operation in the wide area network (WAN) environment. Key transport options are listed below and examples are shown in Figure 1.

- SDH – The existing circuit-based synchronous digital hierarchy (SDH) optical network is seen by some as legacy, but in its reincarnation as next generation SDH (NG-SDH) with a rich Ethernet feature set, it is still key to sustaining high performance and maintainability and continues to be expanded to carry packet traffic. This applies also to its North American version, synchronous optical networking (SONET). Ethernet on SDH/SONET (EoS) is widely used for high QoS services.
- OTN – Above 1Gbps, optical transport network (OTN) to G.709 – see section 11 – is intended to take over from NG-SDH for packet transport on time division multiplexing (TDM).
- Native Ethernet – Ethernet is carried directly on WDM or fiber or on microwave radio.
- “IP on WDM” or “IP on glass” – This avoids Ethernet networking, but irrespective of the bearer used, IP is always mapped via a physical layer framing structure – for example, packet on SDH/SONET (PoS) or increasingly Ethernet – to provide network functions such as clock extraction and rapid fault detection; usually this is all done inside an IP router.

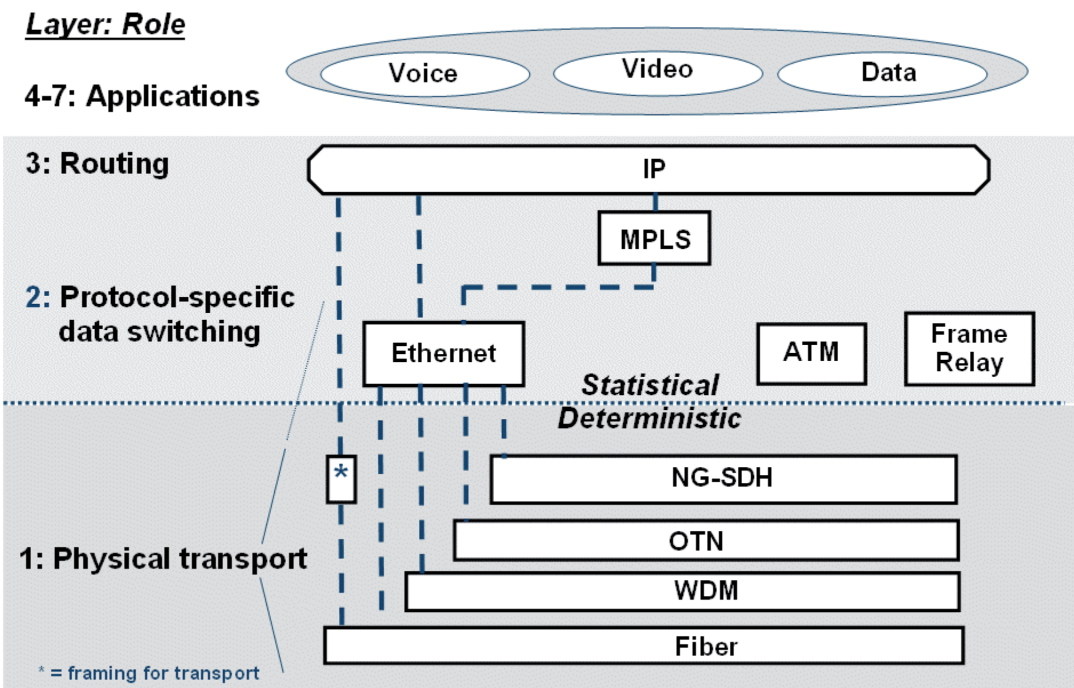

Figure 1: Examples of Ethernet in public network technology layers

In theory, the concept of layers in telecoms networks means that any layer can be changed without affecting others. In practice, IP traffic growth requires capacity increases in every layer, from the upper layers 4-7, which analyze packet contents for network and service policies, through layer 3, which routes packets, layer 2, which switches and aggregates them, and layer 1, which carries them transparently on copper, radio and fiber.

Organizing a telecoms network into logical layers eases scalability, technology evolution and network management as demands change. Scalability in transport means that as the network grows, its costs per user do not rise, and network partitions for maintenance and replacement remain viable. Typically it requires equipment to have progressively higher port speeds.

For the change to NGNs, the issues operators face vary greatly, depending for example on whether they have both fixed and mobile networks, have absorbed other networks with different architecture, or are incumbents or recent entrants. Regulatory and organizational structures are also key factors. As a result, operators vary in how they use transport, but most use at least two or three of the options above.

Start-up networks typically begin with IP on WDM towards the core, plus native Ethernet for the outer parts of the network. Some later add NG-SDH, in order to offer a richer mix of services or to carry Ethernet without concern over interoperability between operators, or end-to-end delay, jitter and packet loss. NG-SDH and native Ethernet are both attractive because as node count rises, total cost of ownership (TCO) is typically reduced by introducing lower layer network equipment for less complex functions. One operator study showed TCO was reduced by more than 97 percent when moving down by two layers. Examples are Ethernet packet switches for aggregation, to replace routers, and optical or NG-SDH switches for restoration and provisioning of bulk network capacity, to replace packet switches. For a given port capacity, each such replacement by equipment acting one layer below can show a capital expenditure (capex) saving in the range 65 percent - 85 percent, based on historical trends. There is also a significant saving in operational expenditure (opex), as explained in section 11.

The change to NGNs involves several aspects of convergence:

- “Network” merges mobile networks with the fixed networks used for residential broadband, Ethernet business services and even core transport
- “Equipment” integrates different network layer functions – for example, packet processing, bulk provisioning of capacity, optical switching – into one type of network equipment where formerly there were several different types
- “Service” renders multiple services accessible to any device or screen – for example, phone, TV, PC; Ericsson calls this FSB.

Across all these aspects, there are multiple and occasionally conflicting standards, industry pressure groups and emerging technologies. The winners will be those based on transport being the lowest common denominator, like the electricity grid, on which everything else relies for a low cost base and ultimate reliability. Convergence is a long process, not a fast result. BT began its own version – 21CN – in 2003, connected the first customers in 2007, and expects the process to continue for the first phase into 2011.

3 Connection oriented networks: Ethernet evolves

The Ethernet brand has been almost confusingly successful, being used to name physical interfaces, network switching technology and customer services. In this paper we focus on network technology, where Ethernet has become preferred for connecting servers, for operator packet aggregation – to improve the fill of lightly loaded ports – and increasingly for the provisioning of broadband service capacity.

A common driver, especially for metro applications [2], has been flexibility and statistical gain, so that different users and services can dynamically share bandwidth. The telecoms market is much more demanding than the enterprise market, in terms of reliability, QoS, standardization and management of services – all defined by the Metro Ethernet Forum (MEF). This needs improved design and scalability of products and networks, summarized as "Carrier Ethernet."

The enterprise origins of Ethernet result in a more dynamic environment for standards and for vendors and operators compared with traditional telecoms technologies such as SDH and asynchronous transfer mode (ATM). As a result, the MEF, Institute of Electrical and Electronic Engineers (IEEE) and ITU-T each take the lead in different aspects of Ethernet: services, technology and transport networking. One result has been connection-oriented (CO) Ethernet.

The key to predictable performance in packet networks is the use of CO working, in partnership with the connectionless working of IP for flexibility. IP's connectionless nature is dynamic and autonomous, but with CO, an end-to-end path is defined before packets flow. The path is constrained by several possible factors, depending on the bandwidth and level of resilience required, in order to achieve the required QoS for each application type and to deliver the service level agreement (SLA) agreed with the customer. This gives a better match to how operators design and manage their transport networks for lowest opex, as in ITU-T G.805/G.809. Examples of extra functionality are the switching of go and return directions together for provisioning and for restoration, plus fast protection switching (50mS). Connection in the data plane can be point-to-point or point-to-multipoint, set up by management commands or by semi-autonomous signaling according to rules held in a control plane (similar to signaling in the telephone network).

CO working has long been used in core networks and in the internet, at first based on ATM switches. Later, routers took on the role, initially with layer 2 tunneling protocol (L2TP), which creates virtual layer 2 tunnels through a layer 3 IP-routed network. Although still sometimes used, this is not a good match to the needs of operators in terms of multi-vendor operations, administration and management (OAM), and it adds some of the complexity of layer 3. Later, routers used MPLS to provide CO.

MPLS acts between layers 2 and 3 of the network. In principle it is independent of both but shares some functions. It is combined with Ethernet in Carrier Ethernet switches and with IP in routers. It adds one or more packet headers and began as a way to simplify packet processing in IP routers, by having simpler rules for those headers. Soon its ability to accept successively more headers meant that it gained functions. It displaced ATM, initially in the network core and gradually outwards.

For IP, MPLS adds label switched paths (LSPs) to give connection-oriented features such as resilience, scalability and stability against faults and traffic incidents. An IP network needs these features for good QoS. MPLS also provides isolation between groups of potentially overlapping IP addresses and so is valuable for additions to enterprise networks.

For Ethernet, MPLS adds service features, such as for QoS in virtual private networks (VPNs), and for security. In this case Ethernet is still used for physical interfaces, but end-to-end connection-oriented management relies on information in MPLS headers added in the Ethernet frames.

MPLS is well established in the network core as “IP/MPLS,” but its progress to the outer network is variable. This is partly because of perceived complexity, plus variations in standards and between vendors, because in the aggregation network, multiple vendors are more commonly used than in the core. Although not acceptable to every operator, single vendor networks have obvious advantages here. Because of the limited progress of MPLS to the whole network, the need for CO packet transport has resulted in other methods.

4 Ethernet packet protocol support options

From several candidates, two methods emerged to adapt or support Ethernet for telecoms networking: provider backbone bridge – traffic engineering (PBB-TE) and MPLS transport profile (MPLS-TP). Both allow networks to be less autonomous and more connection-oriented. “Traffic engineering” refers to putting traffic where the capacity is.

There are key differences in operations and standards between the two methods.

- PBB-TE enhances Ethernet functionality, based on an extra Ethernet header for use by the network operator in provisioning and related tasks. Its IEEE designation is P802.1Qay.
- Instead of enhancing Ethernet, MPLS-TP adds an MPLS header to the Ethernet frame. It borrows some functionality from MPLS, with the emphasis on simplicity compared to full MPLS. In return, MPLS has gained features, notably for OAM, from the creation of MPLS-TP. The Internet Engineering Task Force (IETF) leads MPLS standards work.

The two are sometimes positioned as competitors, with PBB-TE claimed to have lower opex and MPLS-TP claimed to be more flexible. In practice they are complementary, in that individual vendors and operators are prepared to support both, for different roles.

PBB-TE relies only on switching, namely layer 2 processing rather than layer 3, leading to the claim of lower opex and potentially lower capex. For ease of integration, its OAM approach is modeled on that of existing transport networks and differs from the common approach of MPLS and MPLS-TP. On the other hand, the use of a common MPLS-based approach throughout the network is claimed to give cost savings. The balance between these two claims depends on opex forecasts which can be hard to verify, so each operator must decide on its own processes. For example it can depend on how planning and maintenance for each technology is allocated between departments. PBB-TE is considered by some operators to be more appropriate for backhaul and data center connectivity, with MPLS-TP used for wider roles. MPLS-TP appears to have more supporters. The choice is likely to be based on operators’ OAM preferences rather than on other merits of each method.

In standards, PBB-TE is complete and MPLS-TP is proceeding. PBB-TE emerged from an earlier proposal, provider backbone transport (PBT). Similarly, MPLS-TP replaced transport MPLS (T-MPLS) – that is, a profile of MPLS replaced a variant – and in the process avoided potential conflict between standards bodies.

CO packet transport, however achieved, is vital for operators to control resilience and QoS and so offer credible and individual SLAs to customers. This gives improved differentiation for services and customers, and enables a focus on services and software, moving beyond “dumb pipes.”

5 Transport for Ethernet via MSPP

Ethernet can be transported by several methods: on copper, fiber, WDM or microwave radio, or on SDH via generic framing procedure (GFP); a mapping format that allows also for non-Ethernet data interfaces. SDH was created primarily for data transport, not for voice traffic as sometimes claimed. Initially a variety of data interface options existed; Ethernet finally dominated and its features became a defining part of NG-SDH. This continues to be deployed for new data networks because of its proven robust performance, protection and restoration options, plus low opex and low cost integration with existing networks and their operations support system (OSS).

Once mapped into SDH VC-12s, data traffic can be given an arbitrary allocation of bandwidth in increments of 2Mbps by using virtual concatenation (VCAT). Use of link capacity adjustment scheme (LCAS) allows in-service changes of that allocation to be hitless and allows protection bandwidth to be used for lower priority data traffic, even between end points with different vendors. A key benefit is that all of this involves only the terminating equipment, without change to the existing SDH core network. SDH/SONET and NG-SDH provide massive network capacity. Ericsson alone deploys over 450,000 nodes, and many operators have tens of thousands.

For 10 years operators have said they will stop deploying SDH/SONET "within 5-10 years" and they still say that!

In terms of sales, today's dominant optical product is the multi-service provisioning platform (MSPP) – see Figure 2 – which is based on NG-SDH, including Ethernet, and can also have WDM features. MSPPs can have other names, such as multi-service transport platform (MSTP). They appeared from 2002 to provide Ethernet switching and aggregation alongside TDM capability and they enable the phased transition that many carriers want, where the shift to Ethernet is handled customer by customer to maximize the utilization of existing assets. Products vary in their Ethernet/TDM balance and in size, from "pizza box" to subrack, with capacity from one to many 100s of Gbps. The MSPP name is often associated with most ports being in TDM format, but some recent MSPPs – sometimes called mini-packet optical transport platforms (POTPs) (section 8) – can have 100 percent Ethernet ports. Typically MSPPs are used at the first point of packet aggregation as traffic enters the network, and they also provide grooming; that is, switching direction according to service. They are widely deployed for business and mobile networks.

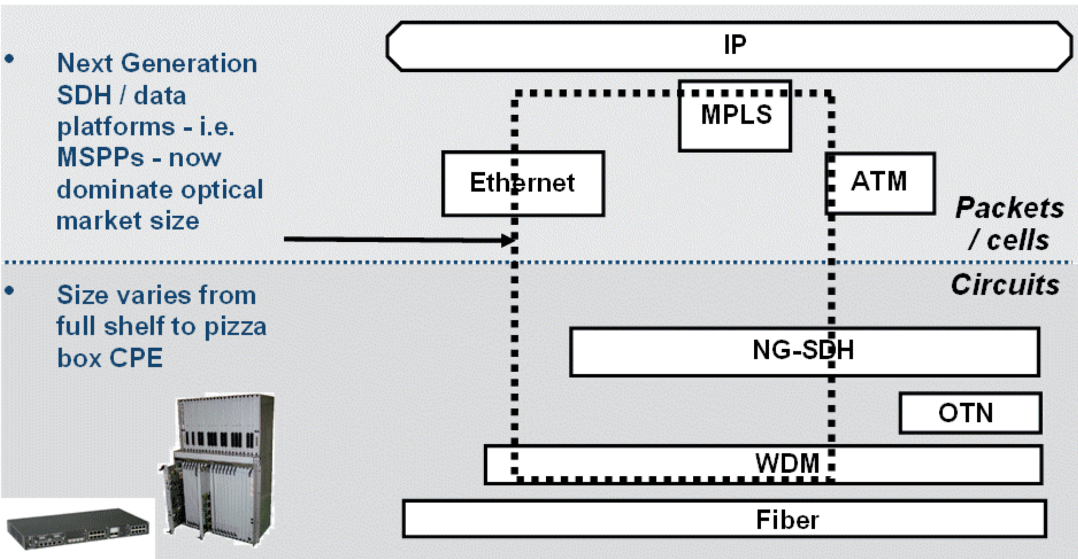

Figure 2: MSPP functions

6 Mobile backhaul

To connect base stations in the radio access network (RAN) to the core network, mobile backhaul (MBH) [5] uses mostly radio (60 percent worldwide) and fiber links. With the spread of fixed rate tariffs, MBH has to grow very fast [4], as fixed broadband did after AOL triggered a similar tariff change in 1996. In some markets, more than 75 percent of all new broadband subscriptions are for mobile broadband.

Growth of data traffic is putting severe pressure on MBH, which can already absorb 25 percent of opex for mobile operators. Any increase in costs, such as for leased lines, is not acceptable. Many operators now build their own backhaul capacity or else move their leased capacity from TDM to packet, which can reduce cost to 1/3 (except where high QoS is required, then cost may hardly change, because of extra network engineering required). Where practical, the best arrangement is to share capacity within the same operator's fixed network, whose traffic is normally many times higher.

Long Term Evolution (LTE) [7] and high speed packet access (HSPA) [6] will magnify the issue. For the next few years, HSPA and HSPA+ are expected to dominate mobile broadband. Data traffic is expected to overtake voice traffic in mobile networks by 2010. Packet operation allows local aggregation (statistical gain) in the RAN, traditionally done with ATM, and so reduces backhaul costs.

In many base stations, LTE will be co-located with 2G or 3G, leading to shared backhaul. The industry's response is in two forms: migrate gradually from TDM to Carrier Ethernet packet, through a hybrid / dual-standard stage, or deploy just Carrier Ethernet packet transport [10].

- A typical hybrid plan is for voice traffic to remain on TDM until traffic exceeds 4xE1, the typical maximum of voice circuits per base station, and then add Ethernet. A hybrid backhaul solution takes advantage of flexibility in radio base stations and in optical MSPPs, to mix TDM and packet traffic and to change that mix with time. This approach works especially well when there is an integrated network management system across fiber and microwave radio systems. Migration through a hybrid stage minimizes disruption to synchronization and management systems.
- E1 for GSM voice can be carried by Ethernet via a circuit emulation service (CES) for TDM, to the MEF 8 standard, or alternatively via additional MPLS headers to IETF RFC 4553 and RFC5086 standards, to form pseudo-wires. These are alternatively named pseudo-wire emulation edge-to-edge (PWE3) and, like CES, involve structuring a packet stream so that it can transport non-packet payloads, with certain limitations such as imperfect transfer of synchronization. Emulation for E1s faces a choice between several different "standards" and can involve extra investment, notably for third-party network management and for synchronization (see section 7), but is appropriate when only a few E1s remain. It allows an all-packet backhaul solution for the lowest cost of leased line bandwidth.

Within an MBH network, the relatively small address range allows layer 2 Ethernet switching to be used in place of more costly layer 3 IP routing. Ethernet packet traffic may be carried on various bearers: microwave radio, fiber, WDM or NG-SDH. Even access networks may be used; for example, copper or cable TV/multi-service operator (MSO) networks, or Gigabit passive optical network (GPON) [8] of 2.5 or 10Gbps total capacity.

7 Network synchronization

Mobile base stations need a reliable source of highly accurate synchronization [3], complying with ITU G.8261 and G.823/G.824, for several reasons, including smooth handover between cells. Synchronization through the MBH network is mainly supplied by extracting clock information from E1 links at 2Mbps (or potentially from SDH in larger sites), as is done for the fixed network and for legacy services such as integrated services digital network (ISDN). The global positioning satellite system (GPS) is used sometimes but can be expensive or difficult in some locations because of antenna issues.

Synchronization can be done via packet networks, either by using dedicated synchronization packets in layer 1, or from the bit stream of an Ethernet signal (synchronous Ethernet or Sync-E). Both methods have had some deployments.

- The main advantage of synchronization packets is that end-to-end connection can be arranged via any transport technology, and even potentially include very precise phase and timing data. There are two new requirements:
  - The need for equipment to provide very accurate clocks for base stations; these clocks are updated by timestamps in the packets.
  - Network engineering for QoS to minimize lost packets and delay variation in packet arrival times; this requirement favors layer 2 switching over layer 3 routing.
- Once the standardization process for variants precision time protocol (PTP) (IEEE 1588v2) and network time protocol (NTP) (refined and based on SoIP) is complete, this method will be used more widely.
- Advantages of using synchronous Ethernet – for example, at 1 or 10Gbps – are that a standard exists, G.8262, for it to be used in much the same way as plesiochronous digital hierarchy (PDH) or SDH; for example it includes status messages to aid OAM. Disadvantages are that it provides only frequency synchronization not phase; few vendors or deployed systems currently have the capability in Ethernet products to allow an end-to-end Sync-E link; and it can currently be used only within a single operator domain.

A typical operator scheme is to introduce the different methods in the order: E1, packet, Sync-E.

8 Bulk switching of transport capacity

WDM expands the capacity of optical fiber to overcome fiber cost and scarcity. It also enables optical switching – for resilience and provisioning – at lower cost per bit and potentially faster response than any other technology. WDM systems include:

- DWDM – dense WDM, to 160 channels for metro and core applications
- CWDM – coarse WDM, typically below 16 channels for metro.

With optical amplifiers along the line, distances of over 4000km can be achieved at 10Gbps per wavelength, down to around 1200km at 100Gbps.

Most systems are point-to-point – but deployment of more complex networks based on reconfigurable optical add-drop multiplexers (ROADM) for DWDM has been extensive, pioneered in Europe by Ericsson's Marconi operation, and driven globally by video traffic in metro networks. ROADMs allow low cost access to a portion of through traffic and are used in both metro and long-haul systems. Individual wavelengths can be dropped, added, or multicast, a flexibility achieved without the need for in-line regenerators, which can add cost and block future upgrades in system capacity. If regenerators are omitted, many DWDM systems designed for 10Gbps per wavelength can be upgraded to 40Gbps. Based on recent system trials, further upgrades are planned to 100Gbps [9], the latest interface rate for Ethernet, as defined in IEEE 802.3ba.

Recently the use of all-optical wavelength selective switching (WSS) devices, usually based on moving mirrors, has allowed the number of DWDM ports (or "degree") per ROADM to increase from 2 or 3 up to 8 or 9, as shown in Figure 3, a progression from ring or chain operation to an optical switching node in a mesh. This needs better planning tools because optical budgets are now for multiple analogue hops, but it supports more flexible network planning and better resilience. Some WSS can provide additional flexibility:

- Directionless switching removes connectivity limits ("blocking") of some early designs by allowing any wavelength on any DWDM port to be dropped locally, and added back in.
- Colorless card slots allow any shelf slot to access any add-drop wavelength.

These together reduce site visits and simplify provisioning. ROADMs with WSS now dominate optical switching in public networks. This contrasts with fiber-to-fiber switching where signals are nominally at the same wavelength (usually "grey" at 1310nm), a system researched intensively since the 1980s and now used mainly in specialized networks such as broadcasting and internet exchanges.

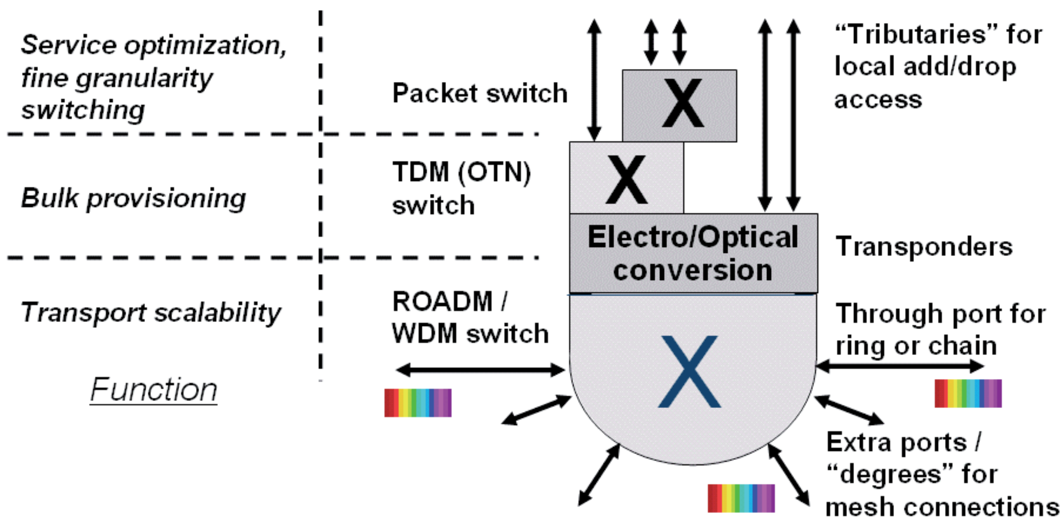

Figure 3: ROADM evolves into POTP

All-optical switching suits the progression to higher transport speeds, 40Gbps and 100Gbps, needed to cope with traffic growth. For typical packet traffic, the alternative of parallel 10Gbps paths results in poorer utilization of transport capacity, because of router internal processes. In the US, 40Gbps is heavily deployed; 100Gbps is progressing through standards and expected to be in mature deployment from 2011. On the client side, they both reduce the number of connections needed in server farms and internet exchanges, where connection management is a serious challenge. On the network-facing side they increase capacity per fiber and per wavelength. That is important where there are limits on fiber count or on the space and power available for remote amplifier clusters. It is sometimes appropriate to use the higher speed on only one side, facing network or client, with parallel connections on the other, and so use a muxponder or an inverse multiplexer, respectively, to perform the conversion. For the network-facing side, in order to fit existing wavelength channel plans despite the higher rates and long distances, vendors of 40Gbps and 100Gbps systems offer various complex modulation formats and receiver technology, such as coherent detection to extract extra signal information.

As WDM capability grew it became clear that better performance monitoring was needed. Also needed was finer switching granularity, but not as detailed as for packets or conventional TDM/SDH. ITU-T defined the OTN to G.709 – also known as the optical transport hierarchy (OTH). This gives a range of optical interfaces at speeds beyond SDH, in a format to contain TDM and Ethernet packet traffic, for example the Optical Transport Unit 4 (OTU-4) at 112Gbps, to carry 100Gbps Ethernet. Compared with others, this format gives improved error handling and management features, such as for multi-operator transits and carriers' carrier operation, which are not supported in WDM standards. OTN as an interface format has gained wide acceptance for 10Gbps router ports, because of its performance monitoring and enhancement features. Its component optical data units (ODU) are now emerging as a "sub-lambda" mesh switching format; that is, for traffic within wavelengths or "lambdas." Use of this format can be especially effective when combined in the same equipment with switching at SDH and/or at layer 2 Ethernet. In contrast to SDH, OTN does not need synchronization by a central clock. It can carry clock signals for other functions, such as SDH (G.8251) or synchronous Ethernet.

MSPPs result from the combination of NG-SDH with Ethernet switching and aggregation. In a further stage of transport integration, the addition of ROADM and optionally OTN features results in another product category, the POTP, which emerged in 2007. The addition of electronic switching on top of WDM switching adds flexibility – for example, for hair-pinning – to simplify operations and traffic migration. It adds "opaque" switching via an electrical switch in between optical ports (OEO architecture) alongside the ROADM's "transparent" switching via an optical switch in between optical ports (OOO architecture). One role of the OEO switch is wavelength translation. This can remove wavelength blocking, where there is contention between two or more signals of the same wavelength, seeking the same outgoing wavelength slot.

Like the MSPP, the POTP has alternative names, such as packet optical transport system (P-OTS), and is likely to appear in a range of capacities from hundreds of Gbps to several Tbps, to serve broadband fixed and mobile data, Ethernet services, and IPTV. It has features optimized for bandwidth-hungry applications such as IPTV; for example, multicast management by internet group management protocol (IGMP) snooping. Just as the MSPP has come to dominate sales in the optical sector, the POTP is expected to play a major role in the evolution of packet transport and there is likely to be a grey area of definition between these two product types.

9 Ring and mesh networks: automation by ASON control plane

Evolution of WDM networks into a mesh through the use of WSS ROADMs echoes the growth of SDH in core networks, from multiple rings to a mesh for lower overall capex (20-30 percent saving shown in operator studies) and a better match to router networks. Rings, whether structured as "self-healing" or "path-protected," remain essential in those metro networks where a mesh could need too many cable ducts, and this is recognized in the new standard for Ethernet rings, G.8032.

In a mesh, recovery time after a fault can be longer than the 50mS typically achieved in rings and in path switching (1+1) across a network, but few applications are disturbed by the longer time; 500mS is a common figure. This is still much better than the time taken by service layer restoration and minimizes disruption. In a mesh, multiple paths exist for restoration and can be selected automatically by a transport control plane. For the architecture of this plane, ITU-T G.807 and G.8080 describe automatic switched transport/optical network (ASTN / ASON). These use the generalized MPLS protocol suite (GMPLS) which was developed from the control plane of IP/MPLS. Based on these origins, the transport control plane is commonly referred to as "ASON," or occasionally as GMPLS or ASTN.

ASON improves several network aspects:

- Restoration (maximizes availability via multiple alternative paths in a mesh)
- Provisioning (lower opex for setting up paths through multiple nodes)
- Resource management (automated discovery of topology and resources, as in IP routers, to minimize data entry errors).

ASON supports multiple network-sharing algorithms, allowing leased (or "colored") sub-networks rather than just leased lines, and is proving valuable for mobile backhaul. When ASON is applied, "optical switching" replaces the former name "cross-connection" for the action of high capacity circuit switches based on SDH/OTN. This form of "optical switching" is misnamed and done electronically, not to be confused with true optical domain switching, done typically by WSS mirrors in ROADMs. The combination of ASON, tunable lasers and advanced ROADM switches minimizes site visits and is sometimes called Zero Touch Photonics.

10 Linking transport with IP/MPLS

There is much research interest and potential opex benefit for services, in linking the optical transport control plane (often called ASON) and the IP/MPLS control plane, but operators remain cautious about exposing their robust circuit-based transport to the greater volatility of router networks. However, they will accept that ASON be used to automate the configuration of PBB-TE and MPLS-TP connections on Ethernet, because those can replace SDH. This eases the evolution to packet-based technology in transport.

The flexibility of mesh switching in WDM, as well as in SDH/OTN, is increasingly used for bypassing those IP/MPLS routers that play a largely transit role. The role arises because it is uneconomic to fully interconnect a router mesh network of many nodes. Therefore much packet traffic flows semi-transparently through intermediate transit routers, adding cost and latency (delay) to what is basically only a regenerator function. Extra delay is often unacceptable in video broadcasts. Bypass with WDM or SDH or OTN allows total router capacity, and network cost, to be significantly reduced. The capex saving – estimated by one operator at around 20 percent – depends on many variables, usually not specified in published papers, and study continues.

Just as important, after using router bypass, operators have reported reduced packet loss, shorter convergence time after cable breaks, and improved network availability. In general, all-router networks – even those with duplicated routers and interconnections – struggle to beat 99.9 percent availability, 100 times worse than the 99.999 percent ("five nines," downtime of 6 minutes per year) typically achieved by lower layer transport.

Such bypass needs integration of network design and operation between routers and packet transport, for example to identify suitable LSPs for bypass, something that is new to many operators, and it benefits from the recent appearance of layer 2 switching features within optical switch/cross-connect platforms, such as POTPs.

11 Keeping things simple to minimize opex

Several opex components increase, as well as capex, for each packet processed, when signal processing capability progresses from layer 1 transport through layer 2 switching, to layer 3 IP routing and into higher layers for applications support, such as deep packet inspection (DPI). Opex is key because it tends to increase much faster, per year, than capex. Therefore it pays to perform tasks in as low a network layer as possible. Although tasks such as packet aggregation, network resilience and drop-and-continue can be done in routers optimized for layer 3 and above, or in multi-service edge devices optimized for even higher layer roles such as video services, in many cases there is much lower opex if these tasks are done in Ethernet switches or in MSPPs/POTPs, optimized for layers 1 and 2. (Some Ethernet switches are based on slimmed down routers, rather than being purpose-built, so are not optimized for low TCO and may not deliver the expected benefits.) The main opex/TCO components that decrease towards lower layers are listed.

- Energy consumption for equipment and its air conditioning. Lower layer equipments require silicon devices that are much less powerful. In addition, their energy needed per bit carried has been reduced typically by over 60 times in the last 15 years.
- Equipment size – and therefore rent for building space. It is not unusual to see figures such as 5:1 claimed for equipment size ratios between layer 1 and layer 3, for the same traffic capacity.
- Depreciation. Operators write off the cost of routers, for example, over far fewer years than the cost of optical transport, so router cost per year is much higher.
- Staff grades and time. There are far fewer configuration options towards the lower layers, therefore staff of lower grades and expense can be used, spending less time on such tasks.
- Unplanned service disruptions. Simpler operations mean fewer configuration errors and incompatibilities, so less effort is needed to fix problems.

System processing power in lower layers has fewer potential diversions – intentional or not – and therefore traffic performance is more predictable. All of the factors above help to explain why the market for dedicated optical transport in packet networks continues to grow and is bigger than the combined markets for IP Edge and Core, plus Metro Ethernet aggregation.

Of course there are situations where it makes sense to leverage the full range of router functions as much as possible and to minimize the addition of dedicated optical transport equipment. This typically occurs in start-up or smaller networks such as for utilities, but when the equipment count reaches hundreds or thousands, the TCO benefits of a properly layered network become clear.

Operators consider all of the above factors when deciding network architecture and how much to do in each layer, and whether to aggregate and groom locally or to transport lightly filled traffic pipes to central sites for more concentrated processing of packets. There is a trend for more traffic to be back-hauled to internet exchanges and server farms (including peer-to-peer traffic), increasing the advantage of layer 2 aggregation-only options, over full local switching or layer 3 routing capability. Each case has to be costed. Much depends on geography, traffic patterns, existing infrastructure and staff distribution. The varying balance between all these factors leads operators to different solutions, often using more than one in the same network.

12 Summary

Transport networks are continuously evolving. The current trends towards fixed-mobile convergence, packet-based services and rapidly increasing traffic but slowly rising revenues are forcing operators to leverage the best of available options in each layer in order to minimize capex and especially opex. There are clear roles for each layer in the packet network; “all-IP” networks need a lot more than just IP. While NG-SDH retains a key role, the latest generation of Packet Optical Transport Platforms is integrating cost effective wavelength switching, extended reach optics, connection-oriented packet standards and OTN/ODU networking. The great diversity of network contexts results in there being no single solution that is optimum for all cases and this paper has described the main options and their selection criteria.

13 Glossary

ASON	automatic switched optical network
ASTN	automatic switched transport network
ATM	a synchronous transfer mode
capex	capital expenditure
CO	connection oriented
CPE	customer premises equipment
E1	2048 kbit/s PHH E-carrier
GMPLS	generalized multi-protocol label switching
HSPA	high speed packet access
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
ITU-T	international telecommunications union – standardization sector
IP	internet protocol
LSP	label switched path
MBH	mobile backhaul
MEF	Metro Ethernet Forum
MPLS	multi-protocol label switching
MPLS-TP	multi-protocol label switching, transport profile
MSPP	multi-service provisioning platform
NGN	next generation network
NG-SDH / NGS	next generation SDH
OAM	operations, administration and management
opex	operational expenditure
OTN	optical transport network
PBB-TE	provider backbone bridge – traffic engineering
POTP	packet optical transport platform
QoS	quality of service
RAN	radio access network
ROADM	reconfigurable optical add-drop multiplexer
SDH	synchronous digital hierarchy
SLA	service level agreement
TCO	total cost of ownership
TDM	time division multiplexing
WAN	wide area network
WDM	wavelength division multiplexing
WSS	wavelength selective switching

References

[1] Ericsson AB., June 2008. *Full Service Broadband Architecture.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/fsb_architecture.shtml [Accessed June 9 2009].

[2] Ericsson AB., November 2007. *Full Service Broadband Metro Architecture.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/FSB_metro_architecture.shtml [Accessed June 9 2009].

[3] Ericsson AB., November 2008. [Ericsson.com]. *Synchronization in packet-based networks: challenges and solutions.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/synchronization-in-packet-based-networks.shtml [Accessed June 9 2009].

[4] Ericsson AB., April 2008. *MOBILE BROADBAND. A focused offering within Full Service Broadband*. [internet]: Lightreading. Brochure (EN/LZT 108 9699 R2): Available at: http://downloads.lightreading.com/docs/ericsson/mobile_broadband_brochure.pdf [Accessed June 9 2009].

[5] Ericsson AB., October 2008. *High-speed technologies for mobile backhaul.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/high-speed_mobile_backhaul.shtml [Accessed June 9 2009].

[6] Ericsson AB., May 2007. *HSPA, the undisputed choice for mobile broadband.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/HSPA_and_mobile_Broadband.shtml [Accessed June 9 2009].

[7] Ericsson AB., October 2007. *Long Term Evolution (LTE): an introduction*. [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/LTE_introduction.shtml [Accessed June 9 2009].

[8] Ericsson AB., 2008. *Current and next-generation PONs: A technical overview of present and future PON technology.* [internet] Ericsson Review No.2. Available at: http://www.ericsson.com/ericsson/corpinfo/publications/review/2008_02/files/3_PON.pdf [Accessed June 9 2009].

[9] Ericsson AB., December 2008. *100Gb Ethernet transport: preparing the highway for high bandwidth services.* [internet] white paper: Ericsson. Available at: http://www.ericsson.com/technology/whitepapers/broadband/100gb-ethernet.shtml [Accessed June 9 2009].

[10] Jonathon Olsson, Ericsson AB., December 2008. *Carrier Ethernet Backhaul Services for Mobile Operators.* [internet] article: Ethernet Academy. Available at: https://www.ethernetacademy.net/index.php/2008120686/Ethernet-Academy-Articles/carrier-ethernet-backhaul-services-for-mobile-operators.html [Accessed June 9 2009].