ERICSSON WHITE PAPER

284 23-3141Uen Rev A | August 2010

ericsson.
com

PACKET NETWORKS'

SWISS ARMY KNIFE

DECIDING WHEN AND WHERE TO USE MPLS TO IMPROVE END-TO-END PACKET PERFORMANCE

Multi-protocol label switching (MPLS) was created to improve packet performance in the core of the internet, but has been adapted for many other uses. The growth of IP-based services requires end-to-end control of packet performance, so MPLS is being extended into optical packet transport networks. The use – or not – of specific MPLS features raises issues in network design and planning.

ERICSSON

ADDING VALUE TO NETWORKS

MPLS was created to improve packet performance in the core of the internet, and is widely used for that purpose, but has also been adapted for other uses. In particular, MPLS is being used more and more in optical packet transport networks to provide the required end-to-end packet performance control for the growing number of IP-based services. This raises several questions, including which of its features to use and where.

Transport networks continue to evolve. The current landscape of fixed-mobile convergence, packet-based services, rapidly increasing traffic and pressure on revenues is forcing operators to maximize the efficiency of each network layer to minimize capital expenditure (capex) and reduce operational expenditure (opex). The metro segment of the packet transport network is particularly important in ensuring high-quality, reliable and resilient networks.

MPLS has a broad range of applications to enhance the performance of packet networks [1], so many that it has been described as the “Swiss Army knife” of packet networks. Best known for use in business-oriented and performance-critical services, MPLS offers security features and enhanced quality of service (QoS).

Transport is about planning for medium- to long-term network capacity – static for months to years – but it can switch paths in less than a 10th of a second to overcome network faults. It covers a wide range of technologies, from wavelength division multiplexing (WDM), synchronous digital hierarchy (SDH), synchronous optical networks (SONET) to Ethernet switches and microwave radio links. To an Internet Protocol (IP) packet, a router is just another transport box. A transport network includes a hierarchical architecture, plus network management and resilience. Transmission refers to the engineering of point-to-point links.

MPLS encapsulates data packets, adding packet headers that enable a variety of features, and provides a rich set of control and monitoring functions. Multi-protocol refers to its original design, which supported three older protocols: asynchronous transfer mode (ATM), IP and frame relay. Its applications are now mainly for IP and Ethernet [2]. Label switching refers to MPLS packet header address information being altered, node by node, across a packet network along a label-switched path, guided by label-switched routers. Compared with the relatively fixed headers of IP, this provides an extra degree of route control. MPLS is the key to the network transformation provided by next-generation networks (NGNs). In these networks, applications that need different QoS, such as different technical performance factors, are all delivered via the IP format. Despite this single delivery format, they are carried by a range of transport technologies in a converged network, shared across services [3].

PLANES OF OPERATION

MPLS operates on these three notional planes:

1. Data plane – MPLS has a flexible format for adding information to data packets, which carry revenue-earning payload and represent the data plane.
2. Management plane – MPLS has a message set that carries network information for configuration and performance monitoring in the management plane.
3. Control plane – MPLS has a set of rules for signaling the required network connectivity via messages carried in the control plane.

MPLS has evolved separately in each plane for different applications and has extended its influence beyond packet-only networks.

Consider the data plane, where payload is carried. MPLS has evolved a range of packet header definitions that provide features, such as the management of extra addresses (beyond what Ethernet can support), and the emulation of other protocols, such as ATM. The data plane can be configured via a dynamic control plane or a relatively static management plane.

The rapid growth of packet-based services has had the biggest impact on the management plane. Operators want to manage packet-based equipment in much the same way as they manage transport networks based on SDH/SONET and WDM. MPLS transport profile (MPLS-TP) [13], which adds transport-specific messages while ignoring some packet-specific ones, was developed to meet this challenge. MPLS-TP is both a "lightweight" sub-set and an extension of today's set of MPLS management messages. It operates seamlessly with MPLS, and is due for standards completion in 2011.

The control plane provides semi-autonomous operation of network-wide functions, similar to the end-to-end signaling that occurs when a telephone number is dialed. The MPLS control plane can work closely with the IP control plane, but typically covers fewer nodes. Its key roles are often taken over by the management plane. MPLS led [3] to generalized MPLS (GMPLS): a control plane for WDM, SDH/SONET and optical transport networks (OTNs), and into the future, for Ethernet. GMPLS is the control mechanism for automatically switched optical networks (ASON) and automatically switched transport networks (ASTN), which are widely used for restoration.

NETWORK LAYERS

MPLS allows network paths to be specified to optimize performance and avoid network problems, such as congestion or excessive delay, and is a key feature in IP routers and Ethernet switches. Unlike most other protocols, such as IP, ATM or SDH/SONET, MPLS generally does not have its own equipment, but shares with other protocols. It is therefore sometimes described as a "shim layer" or "layer 2.5" as it fits between layers 2 and 3 of the conventional layered view of telecom networks. This view, shown in Figure 1, is loosely based on the open systems interconnection (OSI) layered model of data networks.

Organizing a telecoms network into layers makes it easier to maintain, grow or replace portions to respond to changing demands. Each layer is designed to carry traffic transparently from the layers above, and therefore can be changed with minimal impact on the other layers. Regular capacity increases in each layer are needed to keep up with the rapid growth of packet-based services. The upper layers, 4-7, analyze packet contents for network and service policies, layer 3 routes packets; layer 2 switches and aggregates them and layer 1 carries them on copper, radio and fiber. Convention dictates that the term transport primarily includes layer 1, but as packet traffic grows, the term has been extended to involve layers 2 and 3.

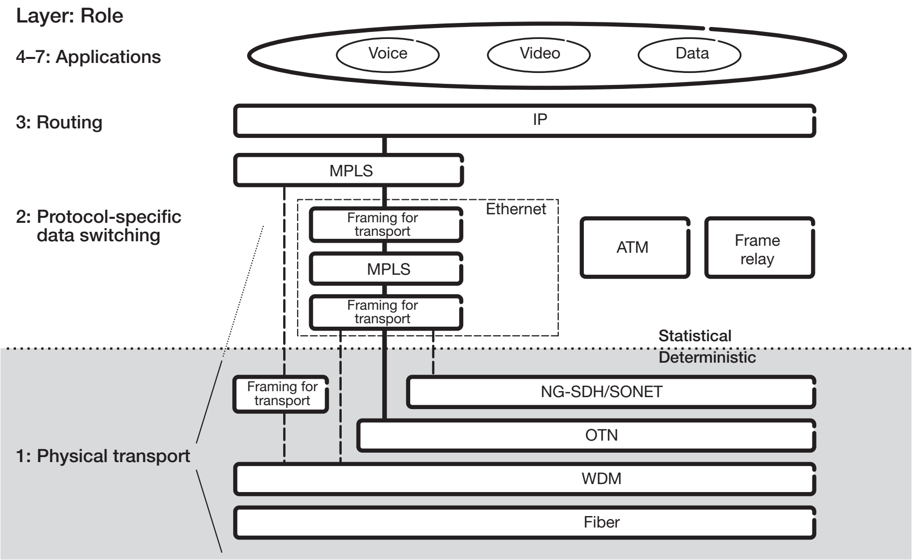

Figure 1: Examples of inter-layer connections in the data plane: the solid line from IP to OTN represents a long-term trend

MPLS shares equipment with other layers – it is independent of other layers, but relates to them. It has some attributes of IP in layer 3, notably in control plane signaling, which leads some observers to confuse the two. MPLS is also used in packet equipment that provides pseudo wire emulation (PWE) [4] of layer 1, notably to mimic time-division multiplexed (TDM) services, mostly at rates of 1.5Mbps and 2Mbps (T1 and E1 respectively).

COMPLEMENTS EACH TECHNOLOGY

Widely deployed Carrier Ethernet equipment has enhanced features, compared with office-style Ethernet, that improve network availability and reduce opex. MPLS adds key features to Carrier Ethernet, such as greater network availability, multiple QoS streams per port and value-added services, such as virtual private networks (VPNs) [3], which can generate higher revenues than can be justified by high availability alone. The underlying Ethernet technology itself reduces costs, allowing operators to meet the challenge of falling revenue per bit carried and at the same time, offer lower service tariffs.

Packet aggregation, based on statistical gain, is increasingly important in mobile backhaul and metro networks, as it improves the fill of transport pipes and traffic ports on routers (see Figure 2). Vendor and operator preference for the aggregation roles is divided between Ethernet switches and small routers; both typically include MPLS. An alternative, also supported by Ericsson, is to combine different approaches, with Ethernet aggregation, MPLS, edge routing and service functions in a single enhanced router.

As shown in Figure 1, if IP/MPLS packets are carried through Ethernet/MPLS aggregation, two sets of Ethernet framing bytes can be needed for transport, simply to pass through Ethernet ports. This only happens if MPLS is used for both IP and Ethernet. Instead, IP/MPLS packets can be passed directly to layer 1, avoiding Ethernet aggregation, which minimizes the accumulation of header bytes and the subsequent cost in both bandwidth and processing power.

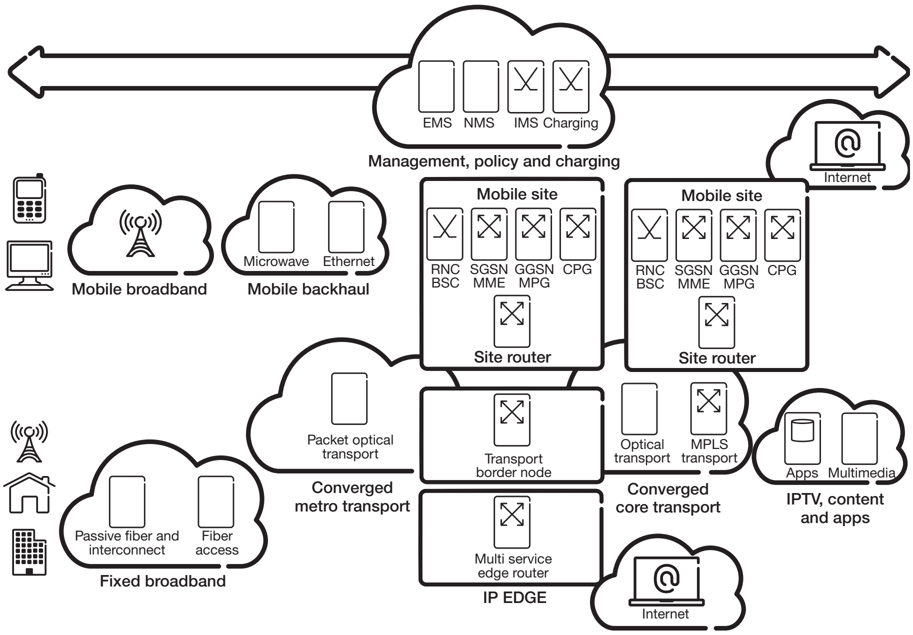

Figure 2: Example of network reference diagram for end-to-end IP-based services

Ethernet products and vendors offer different MPLS features, especially in the case of control plane functions, which enable a larger numbers of nodes. These functions can require significant processing power. If MPLS is being used primarily for managing transport, rather than services, it is possible to omit some functionality to reduce cost.

MPLS AND OPTIONAL NETWORKING TECHNOLOGIES

MPLS has been described as a potential replacement for SDH and SONET, as it can match network resilience and the fine granularity of bandwidth, and therefore support a wide range of business-critical services. It lacks, however, some of the vendor interoperability and mature operations, administration and maintenance (OAM) functions that characterize SDH and SONET and are demanded by operators. The improvements introduced by MPLS-TP will help fill that gap.

One gap that does remain is the difference in resource consumption. MPLS, like other packet-based technologies, needs far more silicon elements for a given bandwidth than TDM-based technologies, such as SDH/SONET. Consequently, more floor space, power consumption and cooling are required, leading to costs that are only justified when the extra functionality of MPLS is needed. For transport with simpler needs, operators are deploying a new generation of TDM-based switching – OTN [12] – which operates above the bit rate domain of SDH/SONET but still provides layer 1 functions that support the functions of MPLS in layers 2 and 3.

SDH and SONET, however, are still being deployed by large and small operators, typically in the form of multi-service provisioning platforms (MSPPs). These represent next-generation versions of SDH and SONET, and add Ethernet ports and switching to the established SDH/SONET features. Increasingly, MSPPs are being enhanced with limited MPLS functions.

The latest trend is to combine selected MSPP features with OTN or configurable WDM in packet- optical transport platforms (P-OTPs), alternatively known as packet-optical transport systems (P-OTSs) [3]. WDM features may include optical or "lambda" switching for mesh and ring networks in reconfigurable optical add-drop multiplexers (ROADMs). These can have several multi-wavelength ports for massive switching capacity. ROADMs can be controlled by a GMPLS-based control plane, and also potentially by MPLS-TP in the management plane.

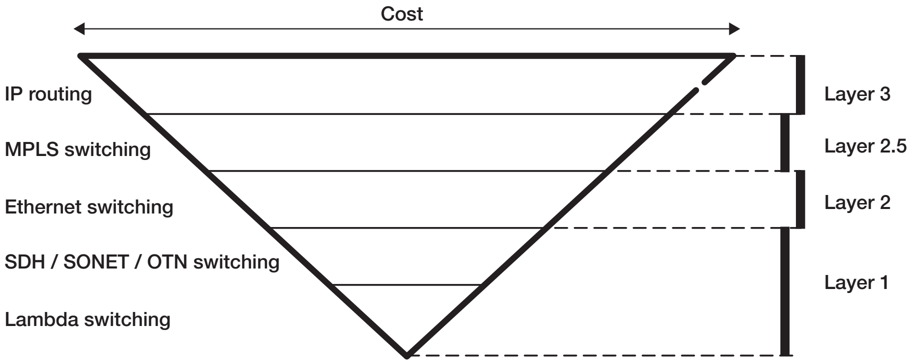

Figure 3: Relative cost of network layers (After Stu Elby, Verizon)

For bulk switching, lambda switching by ROADMs gives the lowest cost, but also the fewest features. Figure 3 shows an operator's view of the relative costs of switching in different layers. In reality, the differences are larger than suggested, and of course, networks require functions from several layers.

Although a lot of traffic is expected to migrate to MPLS, it may be better for some to bypass it. Deciding what traffic goes where is a key network planning issue for operators. Depending on how it is applied, MPLS can either boost or limit performance in terms of network availability or latency. In packet networks, latency – known as "ping time" or the round-trip time from user to user – is increased by propagation delays. Low latency is vital for many services, such as financial and storage-area networks and for high-packet throughput in transmission control protocol/internet protocol (TCP/IP), where latency of 200ms can limit speed to below 3Mbps. For less demanding services, a typical allowance across an operator's backbone network would be 50ms. At 4.9 microseconds/km, the fiber delays alone in a metro network can reach 5ms, which is the typical latency limit for automated protection switching of electricity supplies. For lowest latency, packet processing – including MPLS – needs to be minimized or managed carefully. On the other hand, MPLS can help by constraining traffic to pre-defined paths of lower latency. These paths can stay within the layer 2 or layer 3 packet network or include bypass paths via layer 1 switching.

The connection-oriented (CO) operation of MPLS in routers can add value by steering traffic through layer 1 switching, rather than through a chain of transit routers. Switching in layer 1 with WDM or a TDM-based technology, such as OTN, delivers low latency. Such switching also gives high capacity and is complementary to routers. It is often used to bypass them, especially when carrying well-filled pipes of IP traffic. Unlike most packet technology, its delays are negligible compared with fiber delays in carrier networks. For transit traffic, the delay in each node of layer 1 is typically a few tens of microseconds, although even less than this matters in some financial transactions. The fastest routers can have similar delays under ideal conditions, but more often in reality, each one adds several milliseconds, hence the need for bypass.

OPTONS FOR CONNECTION-ORIENTED WORKING

Most IP traffic relates to "best effort" services, but some video and business services rely on predictable performance. The key to achieving this in packet networks is CO working, in partnership with the connection-less (CL) working of IP for flexibility. CO and CL can both be achieved by MPLS, but the main value of MPLS is in CO, where it dominates over older alternatives, ATM and layer 2 tunneling protocol (L2TP). Unlike CL, a path in CO is defined by network management or a control plane before packets flow. This echoes the familiar approach of SDH/SONET, based on International Telecommunication Union –Telecommunication Standardization Sector (ITU-T) G.805/G.809. It is done in a way that delivers resilience, scalability and stability against faults and traffic incidents. The path is constrained to achieve the required QoS for each application type and deliver according to service level agreements (SLAs).

MPLS is well established to deliver CO in the network core, in the form of IP/MPLS, but its progress to the access network has been variable. Part of the reason is the complexity required in the control plane when so many more nodes are involved, and part is because of variations in standards and between vendors. The latter is particularly important in the aggregation role, often based on Ethernet switches, where it is more common to find multiple vendors.

To offer CO everywhere, standards bodies have defined two possible approaches:

1. Provider backbone bridge – traffic engineering (PBB-TE): an extension of Ethernet defined by the Institute of Electrical and Electronic Engineers (IEEE). It builds on the CL provider backbone bridge (PBB) [5], which adds functionality via an extra Ethernet header used solely by the network operator for provisioning and related tasks. PBB-TE imposes constraints to achieve CO operation. PBB-TE replaced provider backbone transport (PBT) which is still widely referenced.
2. MPLS-TP uses an MPLS header added to the Ethernet frame. Its sub-set or profile of MPLS functions is defined jointly by the Internet Engineering Task Force (IETF) [6] and the ITU-T. The focus is on functions for transport rather than service and on simplicity compared to full MPLS.

MPLS-TP replaced transport MPLS (T-MPLS), which originated from the IETF's MPLS standard, but began to deviate from its parent when developed by a second standards body, the ITU-T [7]. PBB-TE standardization work was completed quickly because of its close relationship to Ethernet and to existing transport networks' OAM. Compared to MPLS-TP, PBB-TE has less in common with MPLS, is less flexible and is not as widely supported.

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

MAKING THE RIGHT CHOICES

Rapid traffic growth and the need for high QoS raises MPLS-related issues for network designers of mobile backhaul (MBH). While traffic in fixed networks still dominates, the number of mobile broadband subscribers exceeded digital subscriber line (DSL) customers for the first time in 2009 [8]. The number of mobile subscribers currently sits at 5 billion and the number of devices – static and mobile – connected through mobile networks is predicted to grow to 50 billion by 2020, enabled by high-speed packet access and 4G/LTE.

One network design issue is to what extent MPLS is included in traffic paths for MBH. The high proportional rate of growth of mobile data traffic is pushing operators to adopt Carrier Ethernet packet transport for lower costs in MBH [9] on both microwave and fiber links. MBH connects radio base stations (RBSs) via the radio access network (RAN) to the much smaller number of switching sites. The limited number of nodes in MBH means there is little need for the extra complexity of MPLS or IP. However, because of its scalability, MPLS is valuable when integrating MBH with a larger fixed network. MPLS can also be an advantage when high QoS is needed, such as for signaling channels, which can face heavy demands from increasingly popular smartphones.

A second network design issue for MBH is whether to include MPLS in synchronization paths. These commonly use TDM but synchronization is starting to be delivered via special packets instead, which need high QoS to minimize lost packets and delay variation in packet arrival times. This delay requirement favors layer 1 switching and simple paths, rather than the extra complexity of MPLS or IP routing. The resilience offered by MPLS however can be valuable – for example, when an operator uses leased paths and therefore cannot be certain that repairs will be made quickly.

OPERATIONS, ADMINISTRATION AND MANAGEMENT

One organizational issue for operators is how far to integrate OAM staff across packets and circuits and across service and transport. OAM relies on the reporting of status information about faults and performance. In SDH/SONET, this is carried in status bits that incur a low overhead in terms of bandwidth and transit processing, so they are sent frequently – at 125 microsecond intervals – and report many parameters, covering each segment of an end-to-end connection.

In contrast, in MPLS networks, OAM relies on a packet per item of status information. This has a higher cost in bandwidth and packet processing power. Depending on the internal design of network nodes, this can affect traffic performance. Therefore in normal operations, less information is sent compared with SDH/SONET. When a network problem occurs, interactive processes, such as ping and trace, are started.

Clearly, the OAM procedures and demands on the skills of operations staff differ significantly in the two cases.

MPLS uses conventional IETF-style packet reporting of status. Some operators, however, would prefer that OAM for connection-oriented packet working appeared as in traditional transport and so be familiar to "metro" or transport operations staff, as defined in [10]. The MPLS-TP standard will be completed in 2011 and its OAM, still under discussion, is expected to be closer to the IETF reports, ensuring compatibility with MPLS. However it will also include features familiar from traditional transport.

OAM includes end-to-end monitoring across multi-operator connections, but is impeded by the many options that exist in operator services and in the Metro Ethernet Forum's (MEF) external network-network interface (ENNI). This is a planning issue for operators and one reason why a high proportion of Carrier Ethernet services are still carried via SDH/SONET, rather than via MPLS.

NETWORK RESILIENCE

By predetermining paths, MPLS improves the management of load-sharing and network resilience and therefore improves network performance [11]. There are operational differences when compared to using SDH/SONET for resilience, and both may be needed for different aspects of the same network. In particular, fast switching requires fast delivery of status information, so specific configuration is needed if MPLS is to approach the SDH/SONET benchmark of 50ms recovery time. Being simpler, lower layers switch faster. Traditionally, such configurations in packet networks were proprietary, such as for fast rerouting (FRR), but standards are emerging in MPLS. For example, MPLS ring protection is to be defined in standards by 2011 to align with the layer 2 Ethernet ring protection of ITU-T G.8032.

Unless performance targets demand it, network redundancy should not be repeated at successive layers, as this costs bandwidth each time. For example, one layer – such as a long subsea WDM link – could have low availability that compromises the end-to-end availability figure. In that case, redundancy could be provided in WDM's layer 1, as well as redundancy in layer 2 or 3 to overcome faults in routers and packet switches. MPLS-TP could provide overall control of the redundant multi-layer paths.

CONCLUSION

A balance between carrier revenue and costs needs to be struck as traffic volumes grow. For best results, each layer of the network needs to be optimized. When used appropriately, MPLS can play a major role in achieving this, by providing essential functions for high-value, end-to-end, IP-based services, improving network performance and supporting legacy services through PWE.

Once MPLS is deployed, the addition of MPLS-TP, which applies a common methodology for provisioning and OAM, will help to minimize opex, as MPLS is extended across packet optical transport networks and especially metro areas. The result takes the industry towards high QoS, end-to-end, IP-based services across multiple operator and vendor domains.

GLOSSARY

ATM	asynchronous transfer mode
BSC	base station controller
CL	connection-less working
CO	connection orientated
CPG	Converged Packet Gateway
EDGE	Enhanced Data rates for Global Evolution
EMS	Enterprise Multimedia Server
GGSN	Gateway GPRS Support Node
GMPLS	generalized multi-protocol label switching
IEEE	Institute of Electrical and Electronic Engineers
IETF	Internet Engineering Task Force
IMS	IP Multimedia Subsystem
IP	Internet Protocol
IPTV	Internet Protocol Television
ITU-T	International Telecommunication Union – Telecommunication Standardization Sector
MBH	mobile backhaul
MME	Mobility Management Entity
MPG	Mobile Packet Gateway
MPLS	multi-protocol label switching
MPLS-TP	multi-protocol label switching, transport profile
MSPP	multi-service provisioning platforms
NG	next generation
NMS	network management system
OAM	operations, administration and management
opex	operating expenditure
OTN	optical transport network
PBB	provider backbone bridge
PBB-TE	provider backbone bridge – traffic engineering
PWE	pseudo wire emulation
QoS	quality of service
RNC	radio network controller
ROADM	reconfigurable optical add-drop multiplexers
SDH	synchronous digital hierarchy
SGSN	Serving GPRS Support Node
SLA	service level agreement
SONET	synchronous optical network
TDM	time division multiplex/multiplexing
WDM	wavelength division multiplex/multiplexing

REFERENCES

[1] MPLS-RC, April 26 2006*The MPLS FAQ*. [Online] Available at: http://www.mplsrc.com/mplsfaq.shtml [Accessed August 3 2010].

[2] IEEE. IEEE 802 LAN/MAN Standards Committee. [Online] Available at: http://www.ieee802.org/ [Accessed August 3 2010].

[3 Ericsson AB., August 2009. *Evolution to optical packet transport.* (white paper). [Online] Available at: http://www.ericsson.com/news/090801_evolution_to_optical_packet_transport_20100510174715 [Accessed August 12 2010].

[4] IETF. IETF RFC4385, updated by RFC5586.

[5] Ericsson AB., January 2090. *Full Service Broadband Metro Architecture* (white paper). [Online] Available at: http://www.ericsson.com/news/090101_fsb_metro_architecture_20100510174711?idx=20&categoryFilter=white_papers_20100505131558 [Accessed August 12 2010].

[6] IETF. *The Internet Engineering Task Force (IETF)*. [Online] Available at: http://www.ietf.org/ [Accessed August 3 2010].

[7] ITU. *ITU-T Study Group 15 Study Period 2005-2008*. [Online] Available at: http://ties.itu.int/ftp/public/itu-t/tsg15opticaltransport/COMMUNICATIONS/ [Accessed August 3 2010].

[8] Infonetics Research, April 2010. *2G/3G and LTE Mobile Broadband Devices and Subscribers* (fourth quarter (4Q09) report). [Online] Available at: http://www.infonetics.com/pr/2010/4Q09-2G-3G-LTE-Mobile-Broadband-Device-Market-Highlights.asp [Accessed August 12 2010].

[9] Jonathon Olsson, Ericsson AB., December 2008. *Carrier Ethernet Backhaul Services for Mobile Operators* (article). [Online] Available at: https://www.ethernetacademy.net/index.php/2008120686/Ethernet-Academy-Articles/carrier-ethernet-backhaul-services-for-mobile-operators.html [Accessed August 3 2010].

[10] ITU-T Y.1730 and Y.1731.

[11] Bocci, M, et al., 2008. Network High Availability for Ethernet Services Using IP/MPLS Networks. *IEEE Communications Magazine*, March 2008, pp 90-96.

[12] ITU-T, December 2009. *Interfaces for the Optical Transport network (OTN)*, Recommendation G.709.

[13] IETF.org. *MPLS-TP Standard*. [Online] Available at: http://wiki.tools.ietf.org/misc/mpls-tp/ [Accessed August 3 2010].