White Paper

Ethernet vs. MPLS-TP

An In-Depth Review of Their Strengths and Weaknesses in the Access Network

By **Yaakov (J) Stein,** Chief Technology Officer
RAD Data Communications

RAD data communications
The Access Company

Abstract

Ethernet is the packet technology that now dominates access networks. Recently, a version of MPLS known as MPLS-TP has been proposed as an alternative. Can MPLS-TP replace Ethernet in access networks? Is MPLS-TP ready to do so? We present below an analysis of the strengths and weaknesses of both protocols in order to make the required comparison.

Contents

Access Networks and Core Networks

Core networks have relatively few network elements (routers, LSRs, switches), while access networks have many (CPEs, NTUs, DSLAMs, aggregators). This means that there is stronger pressure on access network CapEx, and that access network elements need to be as "touchless" as possible.

The core runs higher data rates while the access runs lower data rates (including DSL, PON, wireless). Thus the core may guarantee Quality of Service (QoS) by resource overprovisioning, while the access needs true QoS mechanisms (such as token bucketing).

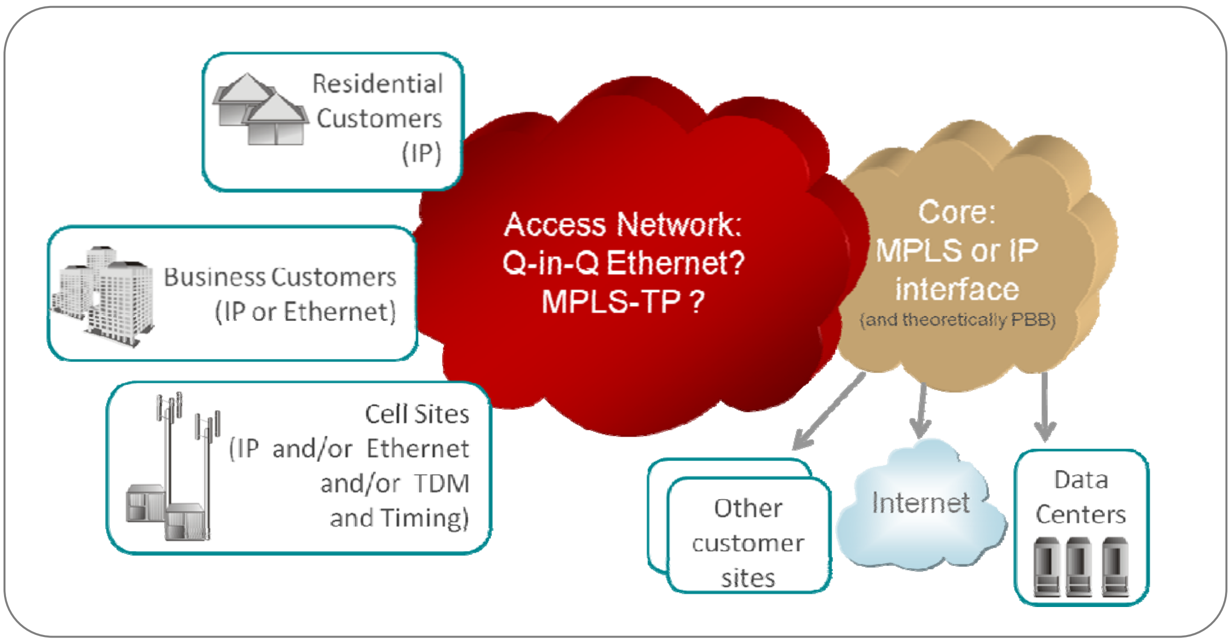

The core is richly connected, while access network topologies are typically tree or ring configurations. Thus a fault in the access network affects fewer people than a fault in the core, but there are fewer bypass options. The core can get away with fast re-route (FRR) while the access network requires OAM and planned Automatic Protection Switching (APS).

The core network elements are in well-guarded networking installations while access network elements are often readily accessible to outsiders. Thus, the core can be considered a "walled garden" from a security point of view, since it features strong security to and from the outside world but loose security on the inside. While customer networks can also be considered walled gardens, it is impractical to protect the entire access network.

Ethernet and MPLS-TP

While both Ethernet and MPLS are commonly used to carry IP, there are many fundamental protocol differences between the two. Ethernet is defined from Layer 0 to Layer 2 (but may run over MPLS), while MPLS always requires a foreign server layer to transport it (which may be Ethernet).

Ethernet frames are inherently self-describing, while MPLS packets do not contain a protocol ID. Every Ethernet frame contains a global non-aggregatable destination address, but MPLS packets have only locally-meaningful labels. Every Ethernet frame contains a unique source address, but MPLS packets contain no source identifier.

Both Ethernet and MPLS-TP can transport IP and other clients. Both Ethernet and MPLS-TP can be transmitted over SDH/SONET and OTNs (Optical Transport Networks).

Both Ethernet and MPLS-TP define fault management and performance management OAM, as well as APS mechanisms. Ethernet does not define a routing protocol (neglecting TRILL and similar recent proposals) but defines a number of Layer 2 control protocols (L2CPs). On the other hand MPLS leverages the entire IP suite of protocols.

Ethernet does not tolerate topology loops, while MPLS, having a Time To Live (TTL) field, can survive transient loops. Ethernet and MPLS both define 3-bit priority (DiffServ) marking. S-tagged Ethernet also supports Drop Eligibility marking. Carrier-grade Ethernet supports bandwidth profiles (token bucketing). Ethernet defines timing (1588) and security (MACsec) protocols.

For both protocols a single entity claims to hold the pen for the specification – the IEEE for Ethernet and the IETF for MPLS. However, in actuality, multiple competing Service Data Objects (SDOs) work on development of both.

Ethernet vs. MPLS-TP

We can now commence comparing Ethernet and MPLS-TP for access networks. We will consider ten criteria: fault management functionality, performance management functionality, APS mechanisms, QoS mechanisms, handling of diverse client traffic types, high accuracy timing (time and frequency) distribution, integration with surrounding networks, CapEx, OpEx, and security. Each will be scored for: Suitability, with a maximum of 2 points; Coverage, with a maximum of 4 points; and Maturity, with a maximum of 4 points.

Comparison Criteria

Fault Management

The Arguments

Access networks require strong fault management capabilities in order to minimize down-time. Ethernet, which began life without any OAM, now has two different OAM protocols (Y.1731/CFM and 802.3-EFM). Having a unique source address, Ethernet is particularly amenable to trace-back functionality. QinQ Ethernet does not provide true client-server separation (because of the common addresses), but this is ameliorated by Y.1731's "MEL" level. Y.1731 provides a comprehensive set of fault management functions. EFM is more limited, but does add "dying gasp" functionality, which is often critical for CPEs. Interoperability issues regarding both OAMs that were initially a concern have finally been addressed, and implementation agreements (e.g. MEF-30) resolve the details. MPLS had no true full-featured OAM but did have basic heartbeats (BFD) and diagnostics (LSP-ping). The IETF designed MPLS-TP FM based on generalizing the pseudowire Associated Channel into the GACh and using BFD for CC, LSP-ping for on-demand diagnostics, and various new frame formats to fulfill specific requirements.

Suitability

Ethernet, having a source address, is highly suited. MPLS, having no true address, requires extra work. The bottom line: Ethernet (2 points) is more suited than MPLS-TP (1 point).

Coverage

Y.1731 is full featured, so EFM fulfills its requirements. MPLS-TP FM was designed to be similar to connectivity fault management (CFM) but is missing dying gasp. The bottom line: it's almost a tie (4 points for Ethernet to 3 points for MPLS-TP).

Maturity

Y.1731 and EFM are interoperable and widely deployed. Some MPLS-TP features are undergoing initial trials. The bottom line: Ethernet wins a wide margin (4 points to 1 point) over MPLS-TP.

Fault Management Subtotal: 10 points for Ethernet, 5 points for MPLS-TP.

Performance Management

The Arguments

Performance Management is a useful tool for maintenance and diagnostics of the access network. The ITU's Y.1731 (but not the IEEE's CFM), supports performance management (loss, delay, PDV measurement) using a request-response model. Y.1731 is also used as the base for commissioning procedures (Y.1564). Widespread vendor interoperability has been demonstrated.

RFCs 6374 and 6375 define a set of performance management functions based on the GACh. These were designed to be hardware-friendly, yet flexible, supporting byte or packet counters, 1588 or NTP style timestamps, and traffic-counters or synthetic loss. Implementations, however, have yet to be announced.

Suitability

Neither protocol has an inherent advantage or disadvantage. The bottom line: a tie (each receives 2 points).

Coverage

Both protocols support all features, although MPLS may be more flexible. The bottom line: a tie by design (each receives 4 points).

Maturity

Y.1731 is fully interoperable. MPLS performance management is not (widely) implemented. The bottom line: Ethernet wins by a wide margin (4 points to 0) over MPLS-TP.

Performance Management Subtotal: 10 points for Ethernet, 6 points for MPLS-TP.

APS

The Arguments

APS is a complex subject and requires careful protocol work and proper configuration. In general, we need solutions for both linear protection and ring protection. Ethernet, not allowing loops, has a particular problem with rings. Numerous open loop protection mechanisms (e.g., G.8032) have been proposed and deployed, but these are not compatible with QoS mechanisms. MPLS in the core exploits Fast ReRoute (FRR) instead of APS, but FRR requires rich interconnection, and so is usually not applicable to access networks. The IETF has standardized RFC 6378 for MPLS-TP linear protection and there are proposals for ring protection (but no RFC yet).

Suitability

Ethernet is not suitable for ring protection, while MPLS has no particular strengths or weaknesses. The bottom line: MPLS wins easily (2 points) over Ethernet (0 points).

Coverage

G.8031/G.8032 fulfill current requirements. RFC 6378 provides for linear protection, but there is still no RFC for ring protection. The bottom line: Ethernet narrowly wins (3 points to 2) over MPLS-TP.

Maturity

G.8031/G.8032 have been extensively debugged and were updated more than once (for good or for ill).

MPLS-TP APS has only been partially finalized and has not yet been deployed. The bottom line: Ethernet wins (4 points to 1) over MPLS-TP.

APS Subtotal: 7 points for Ethernet, 5 points for MPLS-TP.

QoS

The Arguments

While in general two types of QoS need to be considered: hard QoS (IntServ and traffic engineering) - Connection Admission Control (CAC) and Resource Reservation; and soft QoS (DiffServ and traffic conditioning) - priority marking, discard eligibility, queuing, and bucketing. PBB-TE (PBT) defines hard QoS for Ethernet, but is not widely implemented. Ethernet has P-bits for prioritization marking and S-tagged Ethernet has discard eligibility marking. The MEF's bandwidth profile defines a token bucketing algorithm. Ethernet headers are self-describing, supporting Traffic Awareness (TA). MPLS-TE supports hard QoS resource reservation, but "TE" is not considered relevant for access networks. Traffic Class aware (and L-LSPs) support DiffServ prioritization, but MPLS packets are not self-describing and require DPI for traffic analysis.

Suitability

Ethernet supports all QoS types. MPLS does not define for (bucket-based) traffic conditioning. The bottom line: Ethernet narrowly wins (2 points to 1) over MPLS-TP.

Coverage

MEF standards have been proven. Without bucketing, MPLS is at a disadvantage. The bottom line: Ethernet narrowly wins (4 points to 3) over MPLS-TP.

Maturity

Ethernet bandwidth profiles are standardized and certification programs exist. MPLS-TP offers nothing special. The bottom line: Ethernet wins a wide margin (4 points to 0) over MPLS-TP.

QoS Subtotal: 10 points for Ethernet, 4 points for MPLS-TP.

Traffic Types

The Arguments

No transport protocol is useful if it cannot transport the required client traffic. Ethernet differentiates traffic types via Ethertype marking or LLC, and can directly carry IPv4, IPv6, MPLS, Ethernet, Fibre Channel, and low-rate TDM (MEF-8). Ethernet does not directly carry other legacy traffic types (e.g., ATM, Frame Relay), but can indirectly carry them by using a pseudowire without MPLS label. MPLS can carry IPv4, IPv6, MPLS, and pseudowire, and pseudowire, for its part, carries Ethernet, Fibre Channel and all legacy types. Defining a new pseudowire type would require IETF consensus, but the new packet-pseudowire provides more freedom. Neither is universal, but existing mechanisms can be extended to cover new cases.

Suitability

Ethernet supports arbitrary clients via Ethertypes. MPLS supports arbitrary clients via pseudowires. The bottom line: a tie (2 points each).

Coverage

Ethernet does not support all legacy traffic types (e.g., ATM, Frame Relay). MPLS, via pseudowires, supports most traffic types. The bottom line: MPLS wins (3 points) over Ethernet (2 points).

Maturity

Both Ethertypes and pseudowire have been widely deployed. The bottom line: a tie (4 points each).

Traffic Types Subtotal: 8 points for Ethernet, 9 points for MPLS-TP.

Timing

The Arguments

Distribution of highly accurate timing (frequency and Time of Day) is crucial for some access network applications, notably cellular backhaul. Two protocols have become standard for this purpose: Synchronous Ethernet (SyncE) is an Ethernet-specific (MPLS does not define a physical layer) physical layer mechanism for frequency distribution; and IEEE 1588-2008(defined for Ethernet and UDP/IP) for Timing over Packet.

For best performance 1588 relies on on-path support elements (boundary clocks or transparent clocks) that have only been defined for Ethernet. The IETF TICTOC working group is presently developing 1588oMPLS.

Suitability

Ethernet supports ToP and defines a physical layer to support SyncE. MPLS may be able to support 1588 but no physical layer mechanism. The bottom line: Ethernet wins (2 points to 1) over MPLS-TP.

Coverage

Ethernet meets all requirements with SyncE, 1588, boundary clocks, and transparent clocks. 1588oMPLS to support timing over packet may be on its way. The bottom line: Ethernet wins (4 points to 1) over MPLS-TP.

Maturity

ITU-T has defined profile(s) for 1588 use. MPLS presently has no timing support. The bottom line: Ethernet wins a wide margin (4 points to 0) over MPLS-TP.

Timing Subtotal: 10 points for Ethernet, 2 points for MPLS-TP.

Integration

The Arguments

The access network needs to integrate both with the core and with customer networks. Cost and complexity will be minimized by a smooth hand-off (i.e., access protocol compatibility with other network protocols). Customer networks may have Ethernet or TDM interfaces (IP over Ethernet, Ethernet over TDM, Ethernet over SDH/SONET), so Ethernet in the access is a perfect match while MPLS is a reasonable match, since these protocols can be tunneled over MPLS. Core networks are usually MPLS (IP over MPLS, MPLS over Ethernet, MPLS over SDH/SONET). MPLS-TP reuses existing MPLS standards, thus maximizing compatibility (although the issue of stitching vs. seamless connection needs to be considered). Ethernet cannot seamlessly interface with an MPLS core.

Suitability

Ethernet is a perfect match for customer network, but not for the core. MPLS-TP is the best match for the core, but not for the customer premises. The bottom line: a tie (1 point each).

Coverage

Ethernet QinQ and MACinMAC are perfect for the hand-off to the customer. MPLS-TP does not require a gateway for forwarding to the core but control protocols may not interconnect. The bottom line: neither is perfect (3 points for Ethernet to 2 points for MPLS-TP).

Maturity

Ethernet QinQ is widely deployed at present, while seamless MPLS is still in its infancy. The bottom line: Ethernet wins by a wide margin (4 points to 1) over MPLS-TP.

Integration Subtotal: 8 points for Ethernet, 4 points for MPLS-TP.

CapEx

The Arguments

Access network providers need to keep their costs down. Due to the large number of network elements, the access network is extremely CapEx-sensitive. Ethernet switching fabrics are inherently non-scalable, since their long global addresses cannot be aggregated. However, due to their popularity, Ethernet switches are inexpensive (as a result of high volume and large R&D investment in cost reduction). Carrier-grade Ethernet switches do need extra functionality. Ethernet also supports CapEx-saving architectures (e.g., EPON).

LSRs, on the other hand, are complex and expensive. Reducing the price of network elements (an MPLS switch in place of an MPLS router) was the unstated motivation for MPLS-TP. Pure MPLS network elements have simple forwarding engines, and thus should be less expensive than Ethernet switches, but they still require Ethernet, SDH/SONET or OTN interfaces.

Suitability

Ethernet is inexpensive, but cannot scale forever. MPLS-TP allows for significant cost reduction vs. full LSR. The bottom line: close (1 point for Ethernet, 2 points for MPLS-TP).

Coverage

R&D and huge volumes have driven down Ethernet CapEx. MPLS-TP-specific devices can be low cost. The bottom line: a tie (4 points each).

Maturity

MEF certification programs are available for carrier-grade Ethernet switches. Many trials are using (perhaps downgraded?) full LSRs. Chip sets are starting to come out to address this. The bottom line: Ethernet has the advantage (4 points to 2) over MPLS-TP for now.

CapEx Subtotal: 9 points for Ethernet, 8 points for MPLS-TP.

OpEx

The Arguments

OpEx considerations must take into account direct operating costs, staffing and minimizing unchargeable overhead. The reduction of direct operating costs for networks with a large number of elements that must work reliably and be interoperable requires minimum touch (auto-discovery, zero-touch configuration, etc.), and the use of fault management and control plane or management plane protocols. Maintaining competent staff requires that employees be available, trained and retained. Overhead minimization applies to per-packet overhead (OAM, control plane/management plane packets).

Basic Ethernet is zero-touch by design, but carrier-grade Ethernet may require many configuration parameters. Ethernet has a large number of useful Layer 2 control planes (STP, ELMI, GVRP) but no universal control plane protocol. In addition to equipment certification, the MEF has initiated certification for carrier Ethernet engineers. Main Ethernet overhead is high, but tags add only a small increment. Basic MPLS relies on IP routing protocols, but MPLS-TP is designed to be able to function without a control plane (although a GMPLS-based control plane has been defined as an option). MPLS-TP can operate without IP forwarding (eliminating IP logistics). Control plane and management plane can be carried in GACh (although this is not yet fully developed). Specific vendors have expert certifications, but none specific to MPLS-TP. MPLS-TP's look and feel is similar to that of other transport networks. In an effort to minimize retraining, extensions to existing operations support system (OSS) may be leveraged.

Suitability

Metro Ethernet offers low OpEx. MPLS-TP is designed to be inexpensive to maintain. The bottom line: a tie (2 points each).

Coverage

Ethernet has (an inelegant) CP, benefits from available staff and produces moderate overhead. MPLS-TP learned from previous efforts. The bottom line: a tie (4 points each).

Maturity

Ethernet benefits from extensive experience and certification programs. Extensive MPLS operational experience is only partially applicable. The bottom line: Ethernet wins (4 points to 2) over MPLS-TP.

OpEx Subtotal: 10 points for Ethernet, 8 points for MPLS-TP.

Security

The Arguments

Security is perhaps the most important telecom issue today. OAM, APS and QoS mechanisms are powerless to cope with Denial of Service attacks. Access network elements are frequently physically unprotected, and so ports must be protected; packets must be authenticated and their integrity checked; confidentiality mechanisms may be required as well; and management and control protocols must be hard-state.

Ethernet packets carry unique authenticatable source addresses, while MACsec and its 802.1X extensions define mechanisms that can be used to protect carrier networks (although a hop-by-hop security model may not always be ideal). MPLS, however, was designed for core networks ("walled gardens"), with the assumption that there are no inside attacks. Forwarding plane attacks are based on lack of authentication/integrity. Control plane attacks are based on soft state protocols.

Suitability

Ethernet has an authenticatable, unique source address. MPLS has no source identifier and uses soft-state CPs. The bottom line: Ethernet wins by far (2 points to 0) over MPLS-TP.

Coverage

Ethernet has MACsec and 802.1X, but may need more. MPLS-TP has little positive support (but it does support attacks). The bottom line: Ethernet easily wins (3 points to 1) over MPLS-TP.

Maturity

MACsec is starting to appear in standard chipsets. The MPLS community is not addressing the TP security problem. The bottom line: Ethernet clearly wins (2 points to 0) over MPLS-TP.

Security Subtotal: 7 points for Ethernet, 1 point for MPLS-TP.

Final Scores

	Suitability	Coverage	Maturity	Total
Ethernet	16/20	35/40	38/40	89
MPLS-TP	14/20	27/40	11/40	52

So Ethernet wins. Of course, while all deployments have their own particular requirements, equal weight has to be given here to all ten considerations. The reader can easily recalculate the final score for his particular scenario.

Of course some coverage and all maturity scores will change over time. That being said, MPLS-TP lost 29 points due to its lack of maturity and nine points due to its lack of security. Add timing issues and we see that MPLS-TP will not be able to cope with all access network requirements in the near future.

www.rad.com

International Headquarters
RAD Data Communications Ltd.
24 Raoul Wallenberg St.
Tel Aviv 69719 Israel
Tel: 972-3-6458181
Fax: 972-3-6498250
E-mail: market@rad.com
www.rad.com

North America Headquarters
RAD Data Communications Inc.
900 Corporate Drive
Mahwah, NJ 07430 USA
Tel: (201) 529-1100,
Toll free: 1-800-444-7234
Fax: (201) 529-5777
E-mail: market@radusa.com
www.radusa.com

RAD
data communications
The Access Company

Catalog no. 802531 version 3/2012