Ethernet vs. MPLS-TP in Access Networks

Presented by:

Yaakov (J) Stein

CTO

RAD data communications

Unique Access Solutions

What is this talk about ?

Ethernet is the packet technology
that dominates access networks

MPLS-TP is threatening to replace Ethernet in these networks

Is MPLS-TP up to the task ?
Is MPLS-TP ready ?

I start with a brief review of
- characteristics of access networks
- characteristics of Ethernet and MPLS-TP

Then I present a direct technical comparison of
Ethernet vs. MPLS-TP

Access networks ?

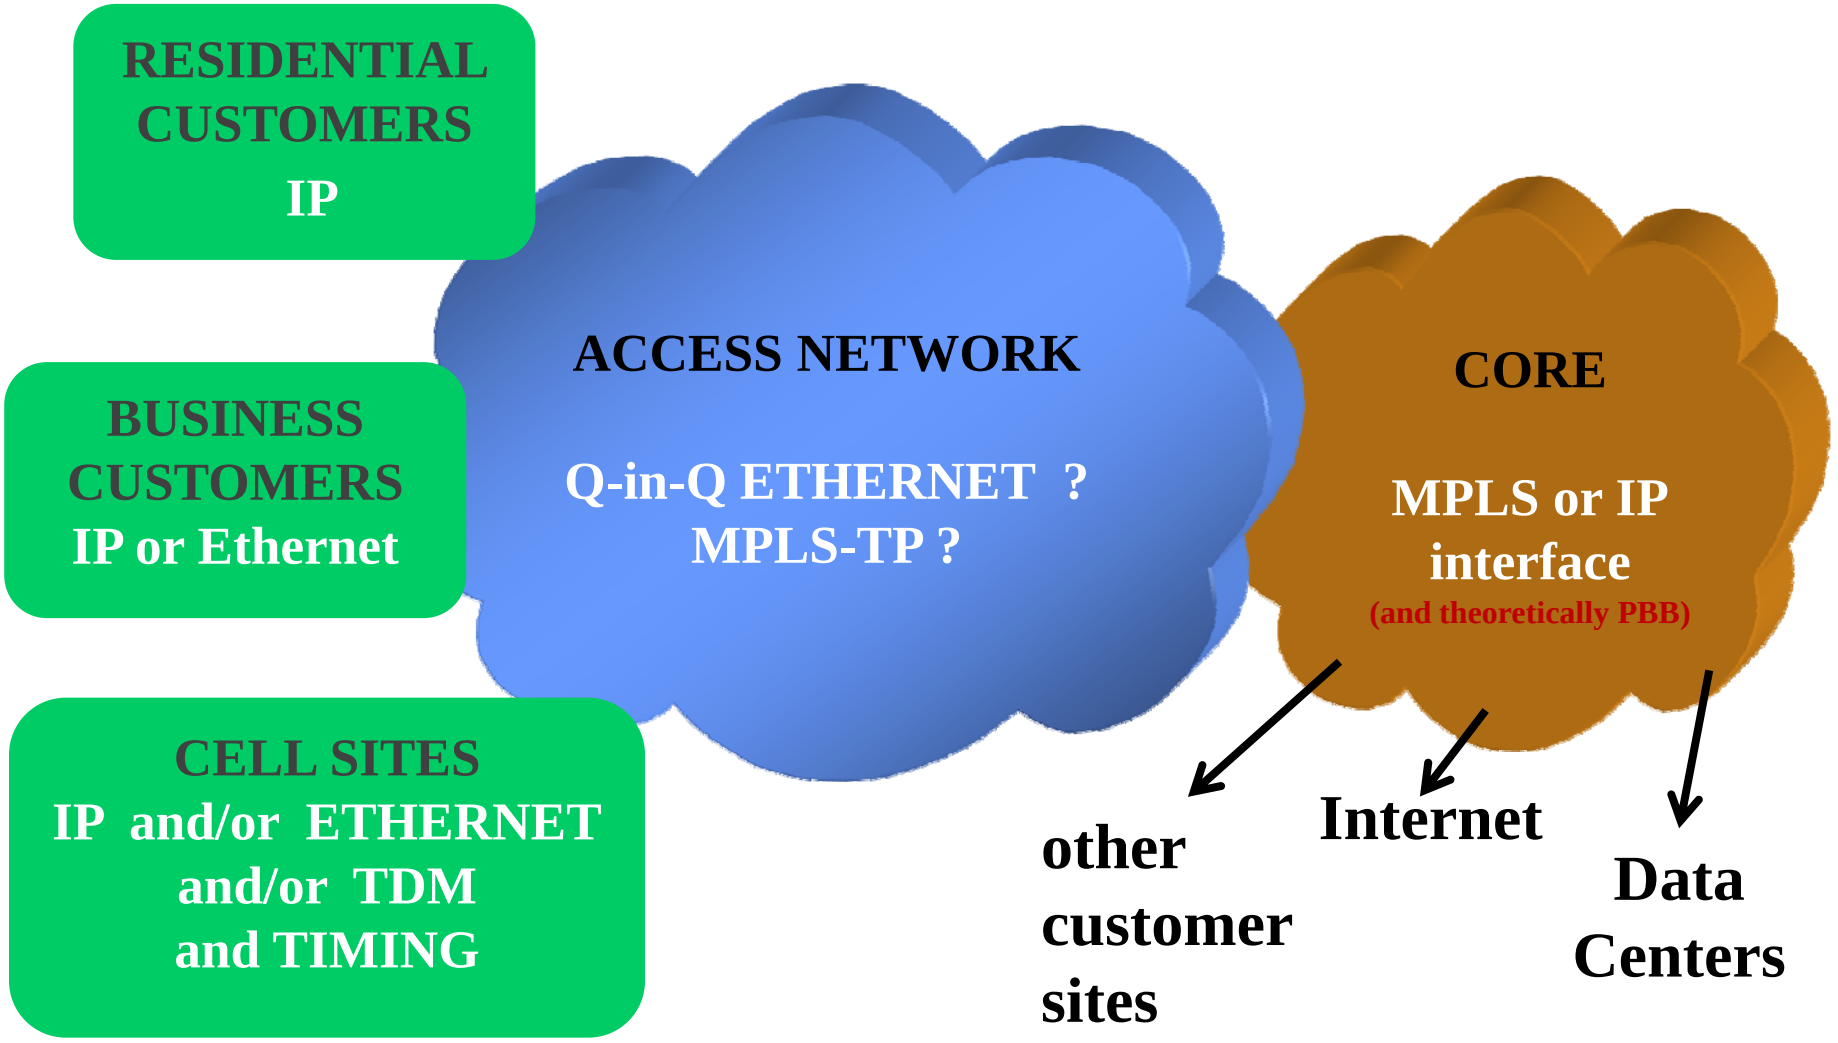

Why Ethernet and MPLS-TP ?

Ethernet started in customer network (LAN)
and for many years has moved into the access network (MEF)
MPLS started in the core network
and is now trying to conquer the access network

Access network segmentation

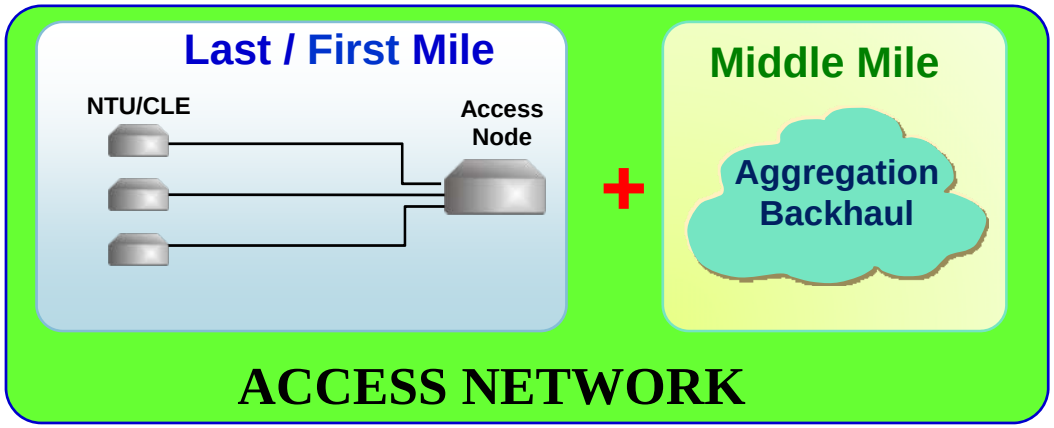

A recent trend is to segment the access network into :

- last/first mile
  - provides connectivity from customer site to first access node
  - leverages physical layer technologies such as DSL, active/passive fiber, microwave, HSDPA+, LTE, ...
- middle mile
  - collects and aggregates traffic from multiple access nodes
  - provides *backhaul* towards core

access / core differences (1)

Differences between core networks and access networks may translate to protocol requirements differences

core has relatively few **N**etwork **E**lements (routers, LSRs, switches)
access has many NEs (CPEs, NTUs, DSLAMs, aggregators)

- strong pressure on access NE price levels
- access needs to be as touchless as possible

core runs high*er* data-rates
access runs low*er* data-rates (including DSL, PON, wireless)

- core may guarantee QoS by resource *over*provisioning
- access needs QoS mechanisms

access / core differences (2)

core is richly connected

access topology is simple (usually trees or rings)

- fault in access network affects fewer people
  but fewer bypass options
- core can *get away* with fast rerouting
- access network requires OAM and planned APS

core NEs are well guarded

access NEs are easily accessible

- core can be considered a *walled garden* from a security PoV
  strong security to and from the outside world
  loose security on the inside
- customer networks too are considered walled gardens
- but it is impractical to protect the entire access network

Ethernet / MPLS-TP differences (1)

While both Ethernet and MPLS are commonly used to carry **IP** there are some fundamental protocol differences:

Ethernet defines from L0 to L2 (but may run over MPLS)
MPLS requires a server layer to transport it (which may be Ethernet)

Ethernet frames are inherently self-describing
MPLS packets do not contain a PID

every Ethernet frame contains a global non-aggregatable destination address
MPLS packets have only locally-meaningful labels

every Ethernet frame contains a unique source address
MPLS packets contain no source identifier

both Ethernet and MPLS-TP can transport IP and other clients
both Ethernet and MPLS-TP can transported over SDH and OTN

Ethernet / MPLS-TP differences (2)

both Ethernet and MPLS-TP define FM/PM OAM and APS

Ethernet does not define a routing protocol (neglecting TRILL, etc.)
but defines a number of L2CPs
MPLS leverages the entire IP suite of protocols

Ethernet does not tolerate loops
MPLS has a TTL field

Ethernet and MPLS both define 3-bit priority (DiffServ) marking
S-tagged Ethernet also supports Drop Eligibility marking

Carrier grade Ethernet supports bandwidth profiles (bucketing)

Ethernet defines timing (1588) and security (MACsec) protocols

A single entity claims to *hold the pen*
for both Ethernet (IEEE) and MPLS (IETF)
but in practice multiple competing SDOs work on development

Face - off

We can now compare Ethernet and MPLS-TP for access networks

We will consider the following criteria :

1. **F**ault **M**anagement functionality
2. **P**erformance **M**anagement functionality
3. **A**utomatic **P**rotection **S**witching mechanisms
4. **Q**uality **o**f **S**ervice mechanisms
5. **Traffic** - handling diverse client types
6. **Timing** –high accuracy time and frequency distribution
7. **Integration** with surrounding networks
8. **CAPEX**
9. **OPEX**
10. **Security**

Each will be scored for :

1. **suitability** 2 points
2. **coverage** 4 points
3. **maturity** 4 points

FM – the arguments

Access networks require strong FM capabilities
in order to minimize down-time

Ethernet, once without OAM now has two (Y.1731/CFM and EFM)

Having a unique source address
Ethernet is particularly amenable to trace-back functionality

QinQ is not true client-server, but this is covered up by MEL

Y.1731 is full-featured – comprehensive set of FM TLVs
EFM is more limited, but adds dying gasp critical for CPEs

Interop issues of both OAMs have finally been resolved
and implementation agreements (e.g. MEF-30) resolve details

MPLS had no true full-featured OAM
but had basic heartbeats (BFD) and diagnostics (LSP-ping)

The IETF designed MPLS-TP FM based on the GACh and

- BFD for CC
- LSP-ping for on-demand diagnostics
- new frame formats to fulfill specific requirements

FM – the verdict

Suitability

- Ethernet, having a SA, is highly suited
- MPLS, having no true addresses, requires extra work

BOTTOM LINE - Ethernet is more suited (2 points 1 points)

Coverage

- Y.1731 is full featured, EFM fulfills its requirements
- MPLS-TP FM was designed to be similar to CFM but missing dying gasp

BOTTOM LINE – almost tie (4 points 3 points)

Maturity

- Y.1731 and EFM are interoperable and widely deployed
- some MPLS-TP features are seeing initial trials

BOTTOM LINE - Ethernet wins a wide margin (4 points 1 point)

TOTAL 10 points 5 points

PM – the arguments

Performance Management is a useful tool for maintenance and diagnostics of the access network

The ITU Y.1731, but not the IEEE 802.1ag supports PM (loss, delay, PDV, ...) using a request-response model

Y.1731 is used as the base for commissioning procedures (Y.1564)

Widespread vendor interoperability has been demonstrated

RFCs 6374 and 6375 define a set of PM functions based on the GACh

These functions were designed to be HW friendly, yet flexible

- support byte or packet counters
- 1588 or NTP style timestamps
- traffic-counters or synthetic loss

Implementations have yet to be announced

PM – the verdict

Suitability

- neither protocol has an inherent advantage or disadvantage

BOTTOM LINE – tie (2 points 2 points)

Coverage

- both protocols support all features
- MPLS may be more flexible

BOTTOM LINE - tie by design (4 points 4 points)

Maturity

- Y.1731 is finally interoperable
- MPLS PM is not (widely) implemented

BOTTOM LINE - Ethernet wins a wide margin (4 points 0 points)

TOTAL 10 points 6 points

APS – the arguments

Automatic Protection Switching 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 has a particular problem with rings
There are many *open loop* ring protection (e.g., G.8032)
but these are not compatible with QoS mechanisms

MPLS in the core exploits Fast ReRoute (RFC 4090) 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

APS– the verdict

Suitability

- Ethernet is not suitable for ring protection
- MPLS, has no particular strengths or weaknesses

BOTTOM LINE – MPLS easily wins (0 points 2 points)

Coverage

- G.8031/G.8032 fulfill current requirements
- RFC 6378 for linear protection, no ring protection RFC yet

BOTTOM LINE – Ethernet narrowly wins (3 points 2 points)

Maturity

- G.8031/G.8032 have been extensively debugged and have been updated more than once (good or bad?)
- MPLS-TP only partially finalized and not yet deployed

BOTTOM LINE - Ethernet wins (4 points 1 points)

TOTAL 7 points 5 points

QoS – the arguments

Two types of QoS need to be considered

1. hard QoS (IntServ, Traffic Engineering)
   Connection **A**dmission **C**ontrol and Resource Reservation
2. soft QoS (DiffServ, traffic conditioning)
   priority marking, discard eligibility, queuing, bucketing

PBB-TE (PBT) defines hard QoS, but is not widely implemented
Ethernet has P-bits for prioritization marking
and S-tagged Ethernet has discard eligibility marking
MEF's BW profile defines a bucketing algorithm
Ethernet headers are self-describing – support **T**raffic **A**wareness

MPLS-TE supports resource reservation
but TE may not be relevant for access networks
and **T**raffic **C**lass (and L-LSPs) support DiffServ prioritization
application awareness – MPLS packets are not self-describing
MPLS packets are not self-describing, require **DPI** for TA

QoS – the verdict

Suitability

- Ethernet supports all QoS types
- MPLS does not define for (bucket-based) traffic conditioning

BOTTOM LINE – Ethernet narrowly wins (2 points 1 point)

Coverage

- MEF standards have been proven
- w/o bucketing MPLS is at a disadvantage

BOTTOM LINE – Ethernet narrowly wins (4 points 3 points)

Maturity

- Ethernet BW profiles are standardized and certification programs
- MPLS-TP – nothing special

BOTTOM LINE - Ethernet wins a wide margin (4 points 0 points)

TOTAL 10 points 4 points

Traffic – the arguments

No transport protocol is useful
if it can not transport the required client traffic

Ethernet carries traffic via Ethertype marking or LLC
and can directly carry IPv4, IPv6, MPLS, Ethernet,
fiber 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 via PHP'ed MPLS PWs

MPLS can carry IPv4, IPv6, MPLS, and PWs
and PWs carry Ethernet, Fiber Channel and all legacy types
Defining a new PW type requires IETF consensus
but the new *packet-PW* provides more freedom

Neither is universal
but existing mechanisms can be extended to cover new cases

Traffic – the verdict

Suitability

- Ethernet supports arbitrary clients via Ethertypes
- MPLS supports arbitrary clients via PWs

BOTTOM LINE – tie (2 points 2 points)

Coverage

- Ethernet does not support all legacy traffic types (ATM, FR)
- MPLS, via PWs, supports most traffic types

BOTTOM LINE – MPLS wins (2 points 3 points)

Maturity

- both Ethertypes and PWs have been widely deployed

BOTTOM LINE – tie (4 points 4 points)

TOTAL 8 points 9 points

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

1. Synchronous Ethernet (SyncE)
   is Ethernet-specific (MPLS does not define a physical layer)
2. IEEE 1588-2008 (defined for Ethernet and UDP/IP)
   for **T**iming **o**ver **P**acket
   on-path support elements (**B**oundary **C**locks or **T**ransparent **C**locks)
   have only been defined for Ethernet

The IETF TICTOC WG is presently working on 1588oMPLS

Timing – the verdict

Suitability

- Ethernet supports ToP
  and defines a physical layer to support SyncE
- MPLS may be able to support 1588 (but what about SyncMPLS?)

BOTTOM LINE – Ethernet wins (2 points 1 point)

Coverage

- Ethernet meets all requirements with SyncE, 1588, BC, TC
- 1588oMPLS to support ToP may be coming

BOTTOM LINE – Ethernet wins (4 points 1 point)

Maturity

- ITU-T has defined profile(s) for 1588 use
- MPLS *presently* has no timing support

BOTTOM LINE - Ethernet wins a wide margin (4 points 0 points)

TOTAL 10 points 2 points

Integration – the arguments

The access network needs to integrate with

- the core network
- the customer network

Cost and complexity will be minimized by smooth hand-off
i.e., access protocol compatibility with other network protocol

Customer networks may have Ethernet or TDM interfaces
(IP over Ethernet, Ethernet over TDM, Ethernet over SDH)

So Ethernet in the access is a perfect match

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)

MPLS-TP reuses existing MPLS standards
thus maximizing compatibility (stitching ? seamless ?)

Ethernet can not seamlessly interface with MPLS core

Integration – the verdict

Suitability

- Ethernet is a perfect match for customer network, but not for core
- MPLS-TP is the best match for core network, but not for customer

BOTTOM LINE – tie (1 point 1 point)

Coverage

- Ethernet QinQ and MACinMAC perfect customer hand-off
- MPLS-TP does not require GW for forwarding to core
  but control protocols may not interconnect

BOTTOM LINE – neither perfect (3 points 2 points)

Maturity

- Ethernet QinQ presently widely deployed
- seamless MPLS still in its infancy

BOTTOM LINE - Ethernet wins a wide margin (4 points 1 point)

TOTAL 8 points 4 points

CAPEX – the arguments

Access network providers need to keep their costs down
Due to the large number of NEs
access networks are CAPEX sensitive

Ethernet switching fabrics are inherently nonscalable
since its long global addresses can't be aggregated
Due to popularity Ethernet switches are inexpensive
(high volumes, large R&D investment in cost reduction)
However, carrier-grade Ethernet switches need extra functionality
Ethernet supports CAPEX-saving architectures (e.g., EPON)

LSRs are complex and expensive
Reducing the price of NEs (MPLS *switch* instead of MPLS *router*)
was the unstated motivation for MPLS-TP
Pure MPLS NEs have simple forwarding engines
and thus should be less expensive than Ethernet switches
but still require Ethernet or SDH or OTN interfaces

CAPEX – the verdict

Suitability

- Ethernet is inexpensive, but can not scale forever
- MPLS-TP allows for significant cost reduction vs. full LSR (vs Eth ?)

BOTTOM LINE – tie (1 point 2 points)

Coverage

- R&D and volumes have driven down Ethernet CAPEX
- MPLS-TP-specific devices can be low cost

BOTTOM LINE – tie (4 points 4 points)

Maturity

- MEF certification programs for carrier-grade Ethernet switches
- Many trials are using (down-graded?) full LSRs
  chip sets are starting to come out to address

BOTTOM LINE – advantage to Ethernet (4 points 2 points)

TOTAL 9 points 8 points

OPEX – the arguments

OPEX considerations that we will take into account

- direct operations cost
- staffing
- minimizing unchargeable overhead

Reduction of direct operations costs
for networks with large number of NEs

- equipment must work reliably and interoperate
- requires minimum touch (autodiscovery, zero-touch config., etc.)
- use of FM, **C**ontrol **P**lane or **M**anagement **P**lane protocols

Maintaining competent staff requires

- finding (need to be available)
- training
- retaining

Overhead minimization applies to

- per packet overhead
- OAM, CP/MP packets

OPEX – the arguments (cont.)

Basic Ethernet is zero-touch by design
- but *carrier-grade* may adds many configuration parameters

Ethernet has a large number of useful L2CPs (STP, ELMI, GVRP)
- but no universal CP protocol

In addition to equipment certification
- MEF has initiated certification for carrier Ethernet engineers

Main Ethernet overhead is large, but tags add only a small *delta*

Basic MPLS relies on IP routing protocols
- but TP is designed to be able to function w/o CP

GMPLS CP has been defined as an option

TP can operate without IP forwarding (eliminating IP logistics)
- CP and MP can be carried in GACh (although not yet developed)

Specific vendors have expert certifications
- but none specific to MPLS-TP

TP is similar to other transport networks (*look and feel*)
- in an effort to minimize retraining
- may leverage extensions to existing OSS

OPEX – the verdict

Suitability

- Metro Ethernets have been shown to be low OPEX
- MPLS-TP is designed to be inexpensively maintainable

BOTTOM LINE – tie (2 points 2 points)

Coverage

- Ethernet has (inelegant) CP, available staff, medium overhead
- MPLS-TP learned from previous efforts

BOTTOM LINE – tie (4 points 4 points)

Maturity

- extensive experience and certification programs
- extensive MPLS operational experience only partially applicable

BOTTOM LINE – Ethernet wins (4 points 2 points)

TOTAL 10 points 8 points

Security – the arguments

Security is perhaps the most important telecomm issue today
OAM, APS, QoS mechanisms
are powerless to cope with **D**enial **o**f **S**ervice attacks

Access network NEs are frequently physically unprotected, so
1. ports must be protected
2. packets must be authenticated and integrity checked
3. confidentiality mechanisms may be needed
4. MPs and CPs must be hard-state

Ethernet packets carry unique authenticatable source addresses
MACsec and its 802.1X extensions define mechanisms
that can be used to protect carrier networks
(although hop-by-hop security model may not always be ideal)

MPLS was designed for core networks (walled gardens)
with the assumption that there are no inside attacks
Forwarding plane attacks based on lack of authentication/integrity
Control plane attacks based on soft state of protocols

Security – the verdict

Suitability

- Ethernet, has an authenticatable unique SA
- MPLS has no source identifier and uses soft-state CPs

BOTTOM LINE – Ethernet wins by far (2 points 0 points)

Coverage

- Ethernet has MACsec and 802.1X, but may need more
- MPLS-TP has little positive support (but it does support attacks …)

BOTTOM LINE – Ethernet easily wins (3 points 1 point)

Maturity

- MACsec is starting to appear in standard chipsets
- MPLS community is not addressing the TP security problem

BOTTOM LINE - Ethernet clearly wins (2 points 0 points)

TOTAL 7 points 1 point

The totals

The final scores :

	suitability	coverage	maturity	total
Ethernet	16/20	35/40	38/40	**89**
MPLS-TP	14/20	27/40	11/40	**52**

Caveats :

- Deployments have particular (non)requirements
  but we gave equal weight to all 10 considerations
- Some coverage and ***all*** maturity scores will change over time
  Note: MPLS-TP lost
  - 29 points due to lack of maturity
  - 9 points due to lack of security