GMPLS: How Far We Have Come!

Kireeti Kompella

Juniper Fellow

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Incentives for Developing MPLS

The initial drivers for MPLS were:

1. Higher performance: replace IP lookups (longest prefix match) with switching
2. Lower cost: replace routers with MPLS switches (as had been tried with ATM)
3. Connection-orientedness in IP networks
4. Traffic Engineering (as with ATM switches)

Evaluation

1. Wire-speed IP lookups are possible without the help of (MPLS or ATM) switching
2. High-speed routers can be built at reasonable cost
3. Connection-orientedness in IP networks was an interesting idea, but why?
4. Traffic Engineering was possible to some extent simply by manipulating IGP metrics

Real Drivers for MPLS Deployment

Ultimately, however, SPs deployed MPLS for solid business reasons:

1. New features (IP VPNs) = new revenue
2. Fast Reroute = lower CapEx and OpEx
3. Network convergence = lower CapEx and OpEx

The simple equation at work:

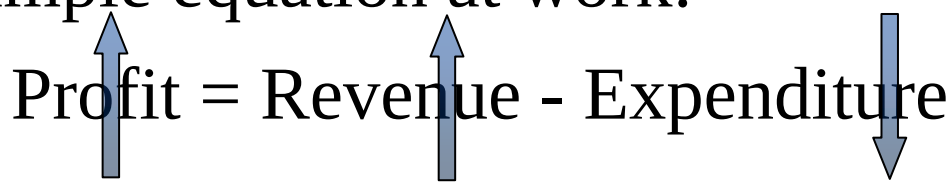

And What About GMPLS?

Well, the initial drivers for GMPLS were:

1. Provision new optical circuits faster (faster revenue recognition)
2. Offer new services (Bandwidth on Demand, Dynamic Circuit Provisioning, Optical VPNs) (more revenue)
3. Replace core routers with optical switches
4. Use a common control plane for the IP/MPLS network and for the underlying optical network (lower OpEx)

Evaluation

1. Realization: good potential here; on the other hand, provisioning is more than signaling
2. Realization: most of these services were predicated on unrealistic bandwidth growth; however, there is growth; there is a business case for O-VPNs, but it needs to be developed
3. Realization: both optical switches and routers are necessary in networks
4. Realization: *this* is the pragmatic, defensible reason for deploying GMPLS

Result

- GMPLS is not just cool technology, it has a real purpose
- So now, the details needed for deployment must be worked out and implemented
  - Multi-vendor interoperability, usability, manageability are now important concerns
- As you will see in the next few slides, the work currently being done in the CCAMP WG shows the transition to this new direction

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Latest Technical Progress in CCAMP

- GMPLS interoperability
  - Protocol issues (under-specification)
  - Addressing
  - Hints on compatible CSPF computations
- MPLS-GMPLS migration
- GMPLS-based protection and restoration
- Multi-layer/region GMPLS networks
- Use of GMPLS for ASON signaling and routing

Latest Technical Progress in CCAMP

- Inter-domain signaling and routing
  - This includes disjoint and protected paths
  - Done in cooperation with PCE WG
- Development of new features and services
  - Primarily L1VPNs
  - Done in cooperation with L1VPN WG
  - Some discussion on LCAS, to support bandwidth on demand
- GMPLS OAM work, still preliminary
- MIBs

Change in the Nature of Work

- Initially, the work in CCAMP was: “let’s take the paradigm of *label switching* and apply it to other networks: lambda switching, TDM switching, etc.
- Now, the work is: we know how to route and signal lambda paths -- now let’s figure out how best to use this in conjunction with IP/MPLS
- This signals a readiness to move on from just playing with concepts to making it real: the first step in the maturation of a technology

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Issue: Deployment Scenarios

- You may remember the (endless) debates over whether GMPLS supports a “peer” model or an “overlay” model
  - OIF defined a UNI with overlay semantics
  - Many claimed GMPLS only supported a peer model
  - CCAMP WG showed how GMPLS can be a UNI
- However, the real question was not which is better, or how it can be realized: the question is, What would allow GMPS to be deployed?
  - to overcome administrative and technical obstructions

Scenario: Network with Routers and OXCs

- Imagine a network that has “ordinary” routers, GMPLS-enabled routers and Optical Cross-Connects
  - Question: how to interconnect and organize all of these?

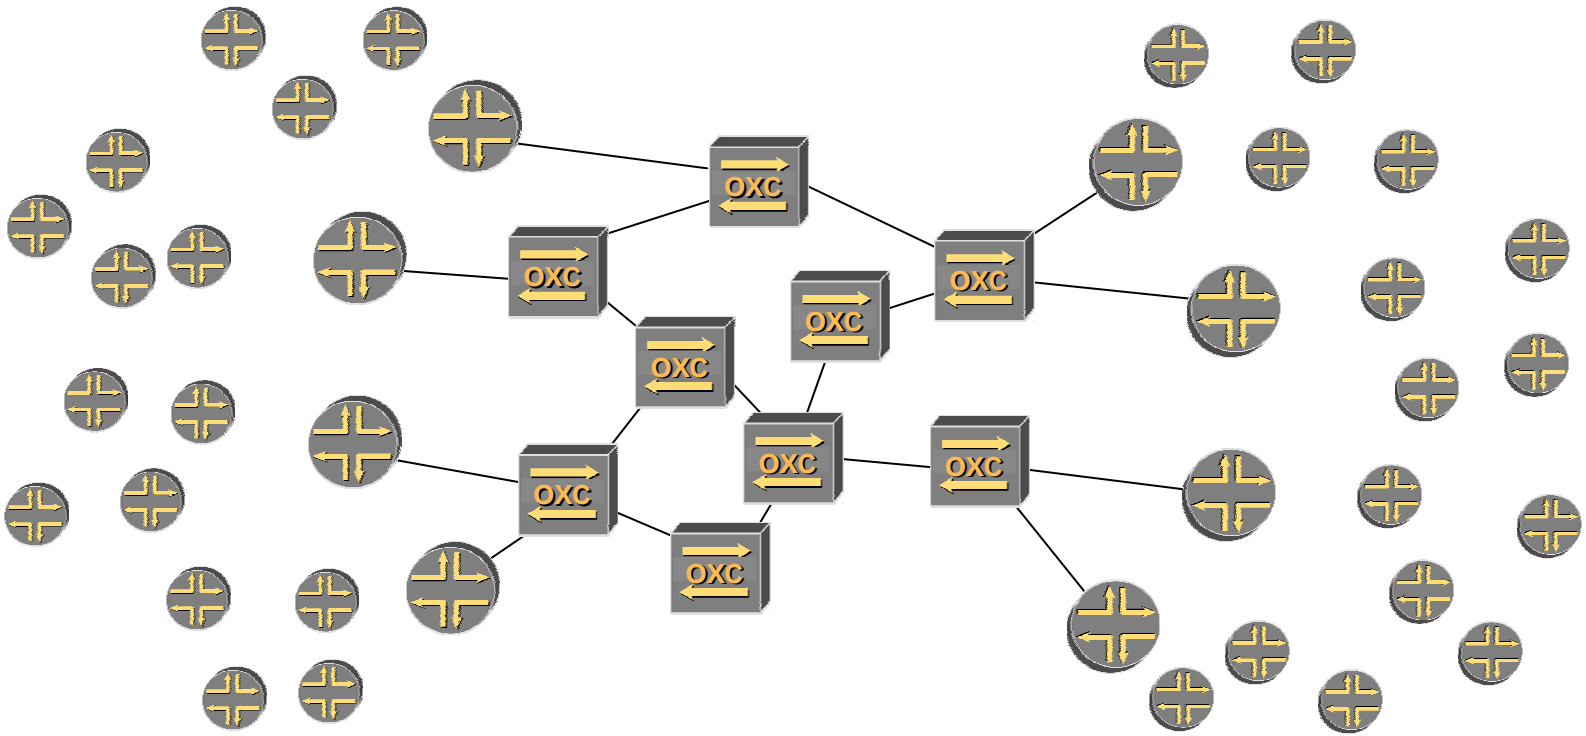

Peer: One Big Happy Family

- Happy?
  - The OXCs see a big IP/MPLS topology and are not happy
  - “Ordinary” routers see optical TE info and constraints and are not happy
  - Only the GMPLS routers are happy!

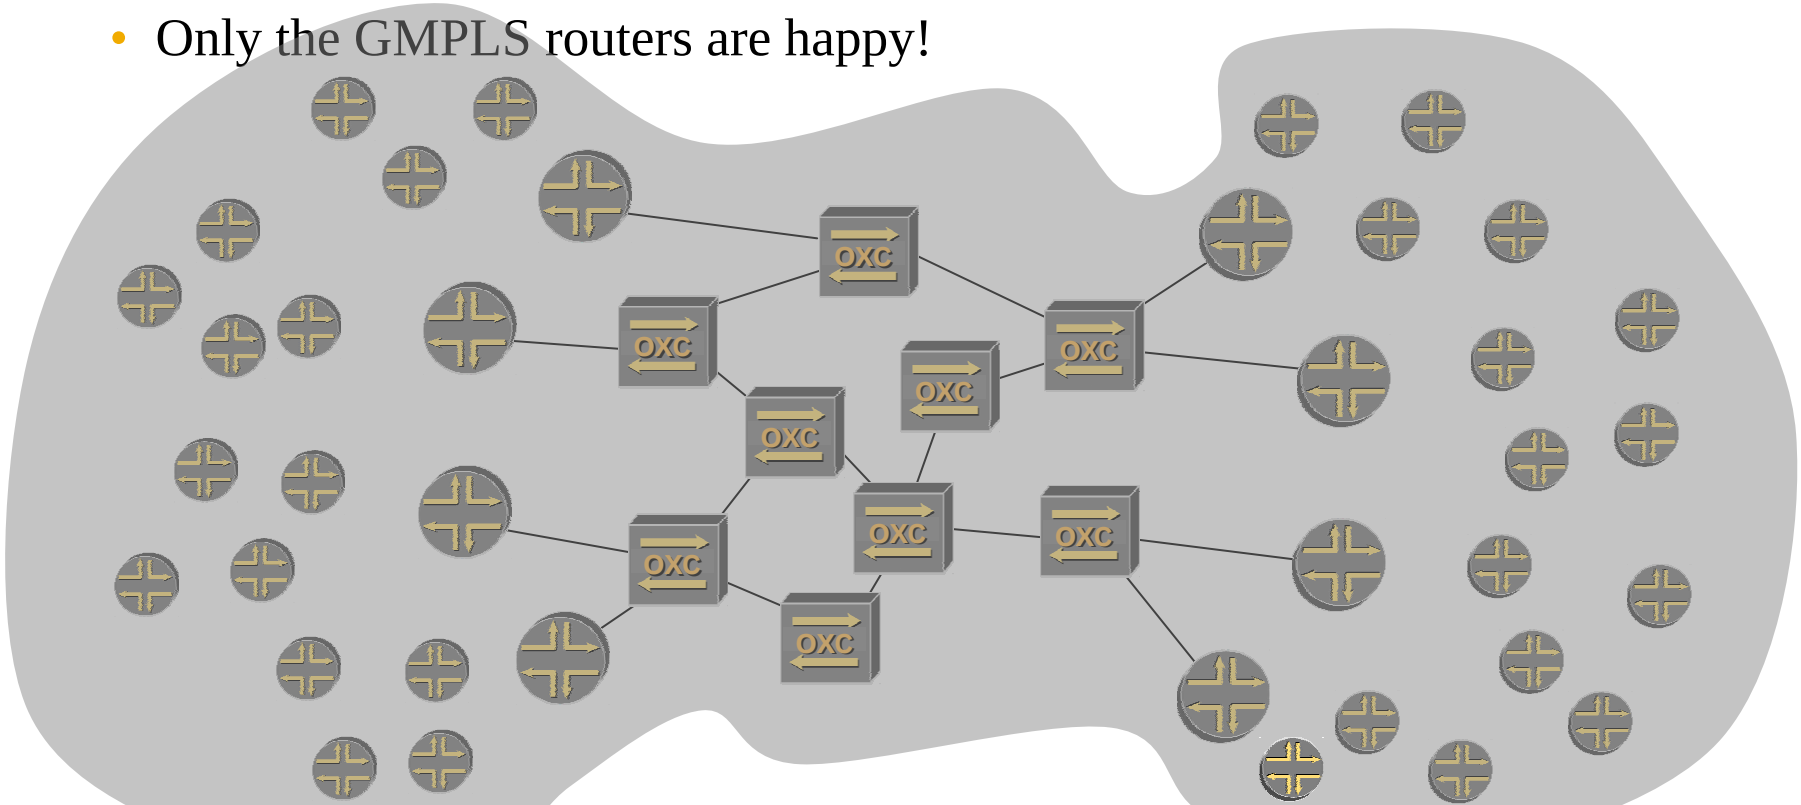

Overlay: Isolate Routers from OXCs

- Now, the OXCs are in their own domain, and routers in a separate domain -- but now, the two are not connected in the control plane!

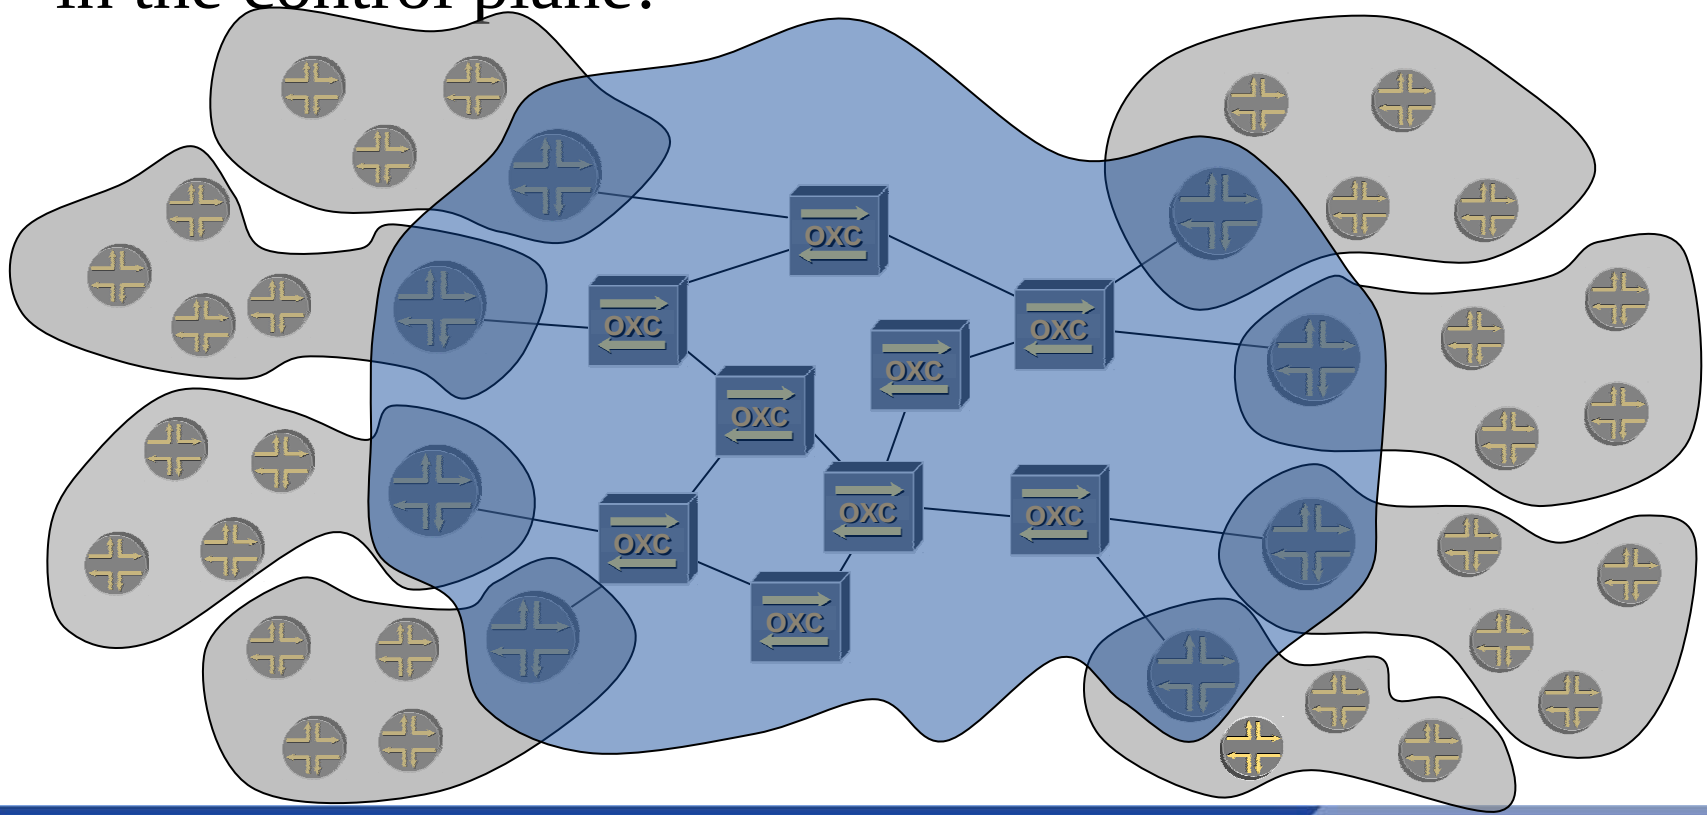

Integrated/Augmented: Peer *Plus* Overlay

domain #1: GMPLS routers + OXCs (peer)
domain #2: GMPLS + other routers (peer)

domain #1 is overlaid over domain #2

GMPLS routers (circled in red) run two parallel control planes: GMPLS and IP/MPLS

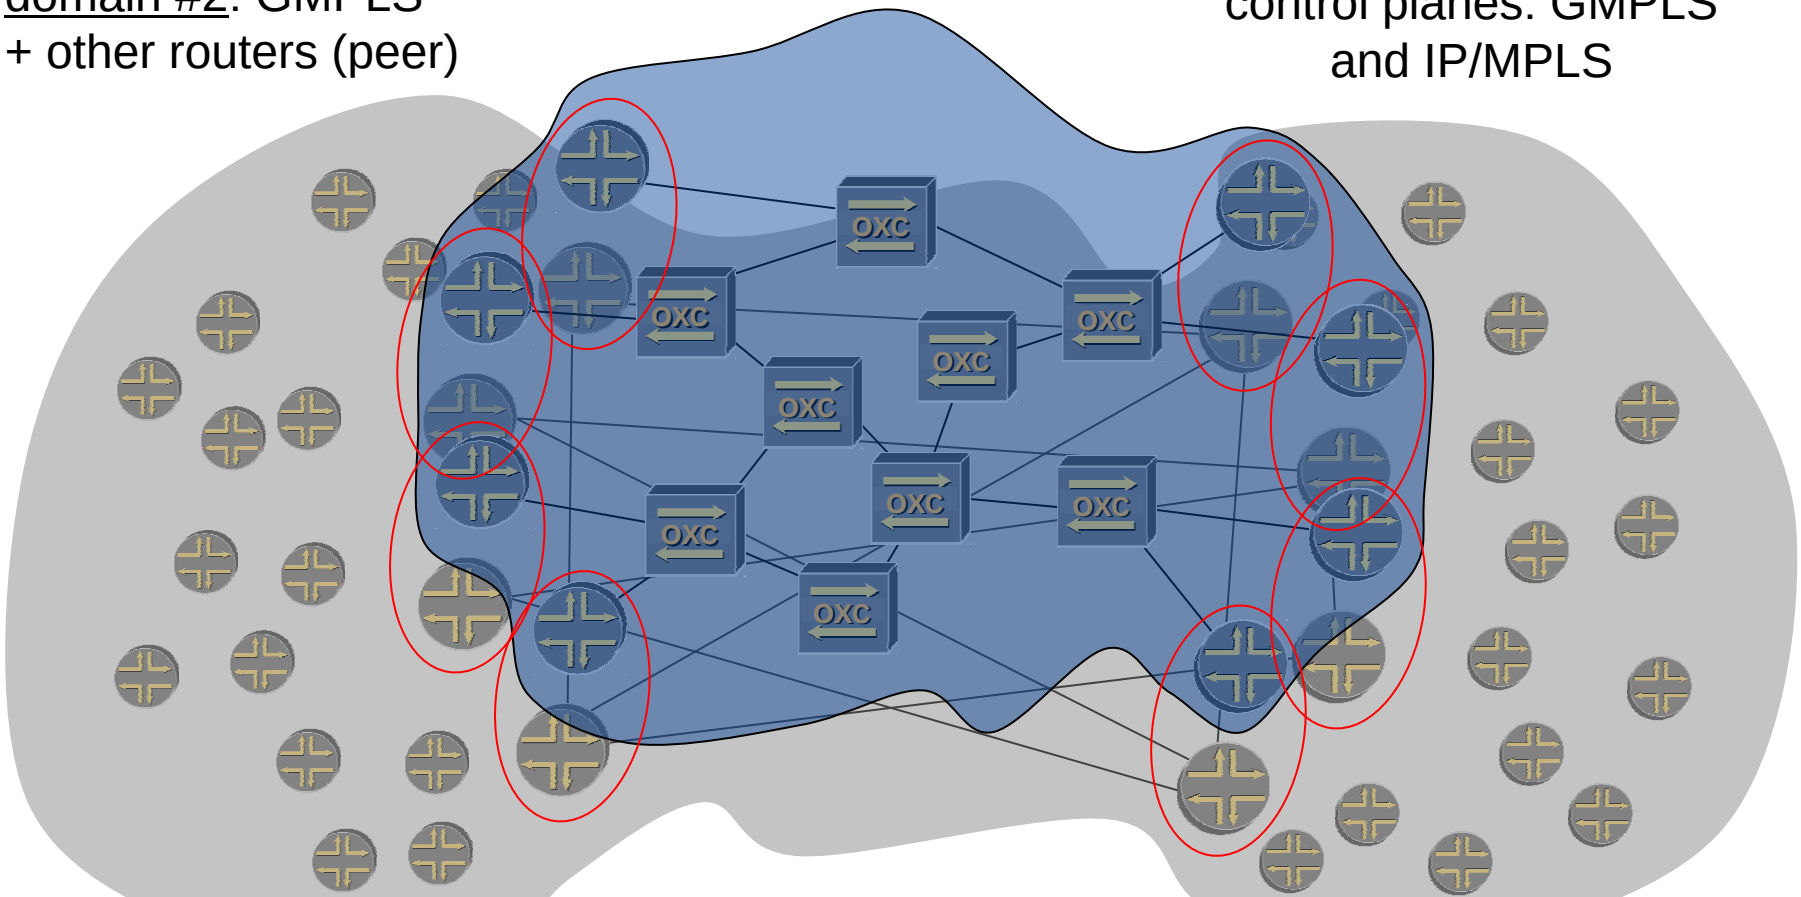

Maturity: Multi-layer Service Network

Here is a creative approach to recognizing the merits of both the peer and overlay models

1. Use multiple, loosely coupled control planes within a single router
2. In one control plane, run a “peer” model between routers and optical switches
3. In the another, run a service network that overlays over the optical network
4. Repeat for other service networks that use the same underlying optical infrastructure

Issue: Routers vs. Optical Switches

A. GMPLS-enabled optical switches will make core routers obsolete

B. No, routers are necessary -- however, do we really need optical switches?

Again, the real issue is lost in this debate. Both routers and switches serve important functions, and only when we realize this can we make progress …

Maturity: Live and Let Live

- Routers do fine-grained (per-packet) switching
- OXCs offer large bandwidth pipes (λ switching)
- Routers need an optical infrastructure
- A static optical infrastructure is sub-optimal, but reconfiguration requires traffic statistics

- So, work together:
  - Routers provide topological and bandwidth requirements to optical network via GMPLS signaling
  - OXCs give routers high bandwidth connectivity

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Standards Bodies

- Several standards bodies are creating standards for signaling in optical networks
  - The IETF, primarily the CCAMP, L1VPN and PCE WGs
  - The ITU-T, primarily Study Groups 15 and 13
  - The OIF
- In the past, there has been a lot of overlapping work, with multiple parallel efforts solving the same problems
- Communication between the SDOs was poor

Current Situation

- Things have improved greatly
- The ITU-T SG 15 and CCAMP have joint design teams to evaluate routing requirements and produce a solution satisfactory to both groups
  - Also on protection and restoration
- There is a lexicography to “translate” between ASON and GMPLS terminology
- The OIF and CCAMP are better aligned, and now communicate frequently, exchanging documents in progress

IETF WG Cooperation

- Even within the IETF, there is a need for coordination
  - Link-state IGP changes (ISIS-TE, OSPF-TE) need to be approved by the appropriate WG
  - Changes to RSVP-TE and to packet switching need to be approved by the MPLS WG
  - The CCAMP WG works closely with the PCE WG and the L1VPN WG
- A good example is the work on Point-to-Multipoint TE LSPs -- this is being done in the MPLS WG, but CCAMP is very much in the loop

Desired Future

- From my personal point of view, the following would be an ideal situation:
  - The ITU-T defines requirements for optical networks in general and ASON in particular
  - The IETF defines IP-based standards meeting these requirements (OSPF-TE, RSVP-TE, LMP, etc.)
  - The OIF produces implementation agreements based on these standards and organizes interop events
  - At every stage, all SDOs exchange communications and ensure they are always in sync

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Deployment Status

- GMPLS lags MPLS deployment by quite a lot
  - One reason is unfortunate timing: GMPLS was getting into its stride just as the Internet bubble burst
- Nevertheless, there has been significant progress, both in standards and in understanding why and how to deploy GMPLS
- There have been interop events and technology bake-offs, but even more, there are proof-of-concept demonstrations showing the efficacy of GMPLS in optimizing optical networks

Deployment Status

- There is a curious parallel with IPv6
- Both GMPLS and IPv6 are inevitable. The question is not *whether* but *when*
- In both of these, Japan has been a key figure, pushing the technology forward as fast as it can
- To my knowledge, both of these are ready to break through, IPv6 first and GMPLS shortly thereafter

Agenda

- MPLS/GMPLS development history
- Latest technical work
- GMPLS maturation
- Standards bodies
- Deployment status
- Summary

Summary

- GMPLS has achieved maturity in many ways: technology, standards work, relationships
- Much of this growth has been driven by work done by Japanese service providers and vendors, at the forefront of which are NTT and KDDI
- I commend both on taking an active role at the IETF as well as in their labs and testbeds
- I thank you all for not losing faith!
- I most of all urge you to keep up your good work. We have come far, but still have a long way to go

Thank you!