ASON And GMPLS - The Battle Of The Optical Control Plane

An overview of the ongoing work of the IETF and ITU to standardize optical control plane protocols.

Nik Larkin, MPLS Architect, npl@metaswitch.com

First issued August 2002

Contents

1. Introduction

This paper provides an overview of the ongoing projects of two standards bodies in the area of the optical control plane.

Optical control plane standards will provide two benefits.

- They will allow automated optical networks to be constructed out of devices from a mixture of vendors.
- They will specify a minimum set of features that all conforming devices will need to support.

These benefits will reduce the cost of deploying and operating optical networks for the consumers of optical switches: the operators and service providers. At least, that's the theory.

At the time of writing, multi-vendor automated optical networks have not been deployed. However, single-vendor optical networks have been deployed, whether controlled by centralized management systems, or partially distributed. Distribution is achieved either using proprietary signaling and routing protocols or, in some cases, early versions of the standards. The timing and completeness of standards has a crucial impact on the industry, because they will impact an operator's business case for upgrading their existing networks and deploying next-generation switching technology.

The work to define the standards for the optical control plane is being done independently by two standards bodies

- the International Telecommunications Union (or ITU), which is developing the Architecture for Automatically Switched Optical Networks (or ASON)
- the Internet Engineering Task Force (or IETF), which is developing Generalized Multi-Protocol Label Switching (or GMPLS).

This paper discusses their interactions, and the challenges that these groups will need to overcome in order for the industry to end up with an effective and implementable set of standards.

The third group discussed in this paper is the Optical Internetworking Forum, or OIF, which blends together some of the work of both the IETF and ITU.

This paper is aimed at those who are unfamiliar with this standards work, or who are strongly affiliated to one body and are interested in an overview of the others.

The author has contributed to all three bodies in several of the areas covered by this White Paper, and Metaswitch (the author's company) actively participates in, but has no particular political affiliation to, a large number of groups in different standards bodies. This reflects our role as a developer of software implementations of many different standards, including a suite of optical control plane components. You can find more information about us in the final section of this paper.

The websites of the IETF, ITU and OIF can be found at http://www.ietf.org, http://www.itu.int and http://www.oiforum.com respectively.

2. Two worlds collide

Why does the title of this paper talk of a battle? Certainly, it is not inevitable that the two groups would come into conflict over the optical control plane, as the relationship between the ITU and IETF in this area is clearly defined, with no clash of responsibilities.

- The ITU produces requirements and architectures based on the requirements of its members. An explicit aim is to avoid developing new protocols when existing ones will do.
- The IETF produces protocols in response to general industry requirements, potentially including those coming out of the ITU.

Thus ASON and GMPLS should not be competitors, but instead should be complementary pieces of work. The desired outcome for both standards bodies is for ASON simply to reference the various protocol specifications in the GMPLS suite. Not only that, but processes have been created to ensure that the two bodies work well together. Liaisons are in place between the two, and each recognizes and values the work of the other. Plenty of individuals and companies are committed to ensuring that the two organizations are in sync.

The problem is that the devil is in the details. Although, in the ideal world, the relationship would be smooth, there are in fact a number of significant challenges for the two standards bodies to overcome, arising from a mixture of technical, cultural and "political" differences.

The next few sections describe these areas of conflict, comparing and contrasting the approaches of the two standards bodies. In each case, the intention is not to advocate one side or the other, but rather to present a balanced discussion of each.

3. The birth of GMPLS and ASON

This section gives some background on the two areas of optical control plane standardization in their respective standards bodies, and then discusses the intermediary role of the OIF.

3.1 GMPLS development in the IETF

3.1.1 History

GMPLS grew out of vanilla MPLS, a packet-switching technology designed to improve the efficiency of data networks. A flavor of MPLS, known as MPLS-TE, provided for provisioning of end-to-end connections using signaling with constraint-based routing. It was a natural move to generalize and extend it to cover circuit-oriented optical switching technologies such as time- and wave-division multiplexing (TDM and DWDM). Our white paper "MPLS In Optical Networks" describes the extensions that were made to MPLS-TE to support optical switching.

Within the IETF, the MPLS protocol development was carried out by the MPLS Working Group. As GMPLS appeared on the scene, the CCAMP Working Group (for “Common Control and Management Plane”) was created to provide it with a home.

Naturally, as an IETF protocol, GMPLS uses an IP-based control plane.

3.1.2 Philosophy

CCAMP’s charter is to define a set of protocols that will allow implementation of a wide range of interoperable electrical and optical switches. As well as the protocols themselves, the group provides informational architecture documents describing how the tools in this protocol “kitbag" are used together, and this architecture is described in such a way as to provide a large amount of flexibility to implementers.

Flexibility does not obviously go hand in hand with interoperability. It is left to interested companies or groups to perform interoperability testing of subsets of GMPLS. This testing is crucial for the industry to form a consensus on which of the optional portions of the standards are required in a given application.

The IETF is, in general, fast on its feet with respect to getting new ideas “out there” as protocols. GMPLS is no exception, and as a result has already been implemented on a number of different vendors' devices. Reflecting the priorities of its supporters, GMPLS is strongly focused on delivering features that are needed now.

GMPLS is designed to be used end-to-end in a "GMPLS everywhere" network. While it is certainly possible to deploy into existing back-level networks, it is left to the reader of the standards to figure out how to do this rather than covered explicitly.

3.1.3 Anatomy

The term "GMPLS" is colloquially used to refer to a set of protocols that, when complete, will work together to provide interoperable end-to-end provisioning of optical (as well as other) networks.

The protocols are as follows, although they do not all have the same level of take-up.

- Generalized RSVP-TE for signaling
- Generalized CR-LDP, also for signaling
- OSPF with TE extensions for intra-area routing
- ISIS with TE extensions, also for intra-area routing
- LMP and LMP-WDM for assorted link management and discovery functions.

RSVP-TE and CR-LDP are alternative protocols that effectively do the same thing. These rivals were inherited from MPLS-TE, where due to conflicting business interests of their employers, IETF members failed to agree on a single signaling protocol, to the dismay of much of the industry. (For a politically-neutral technical comparison between the two protocols, see our white paper "MPLS Traffic Engineering: A Choice of Signaling Protocols".)

As of July 2002, the MPLS Working Group, which originally developed CR-LDP, began to discuss whether work on that protocol should be suspended, due to the fact that a far larger number of companies were interested in implementing RSVP-TE. Whatever consensus is reached here is likely to ripple through to generalized CR-LDP too.

The ISIS and OSPF TE extensions are also functionally equivalent. Here, however, there are strong historical (as opposed to political) reasons for keeping both protocols, namely that the non-TE versions are both already widely deployed in data networks.

Inter-area optical routing has not been defined in detail at the time of writing, and the IETF is considering a number of options.

3.2 ASON development in the ITU

3.2.1 History

ASON was developed by Study Group 15 of the ITU-T, the ITU's telecoms standardization sector.

The work was initiated in response to a demand from ITU members to create a complete definition of the operation of automatically switched transport networks, management, control, data plane and all.

ASON is not a protocol or collection of protocols. It is an architecture that defines the components in an optical control plane and the interactions between those components. It also identifies which of those interactions will occur across a multi-vendor divide, and therefore require standardized protocols. Other areas are intentionally not standardized in order to allow vendors or operators to provide "value add".

As with most ITU projects, ASON was (and continues to be) developed in a top-down fashion, starting with a full and explicit list of requirements, moving on to high-level architecture and then individual component architecture. Only when component architecture is defined in detail are protocols held up to the architecture to see if they fit. Any protocol that fits the requirements of the component architecture can potentially get the ASON "stamp of approval".

3.2.2 Philosophy

Unlike in the IETF, where the optical control plane standards evolved out of a set of existing protocols, the ITU sat down to design the architecture from scratch. And whereas GMPLS was developed in a community strongly associated with IP-based data networks, ITU members primarily come from a telecoms background. Thus, while GMPLS inherits IP concepts and protocols, ASON draws on concepts from protocols used heavily in telecoms transport networks, such as SONET/SDH, SS7 and ATM.

As a generic reference architecture, ASON is intended to be complete, future-proof, highly scalable and highly resilient to faults, specifically targeted at transport networks, which are by their nature expensive to run. Operators need to know that when they add automation to such networks, this will provide the existing function at a lower cost. So, before they even get out of bed in the morning, they expect a clear description of how their requirements are met by the protocols.

ITU-T Study Groups meet at nine-monthly intervals compared to the four-monthly meetings of the IETF, making for a slower and steadier style of standards development.

ASON cannot be directly implemented, as it is a reference architecture. When complete, it will enable developers of existing protocols to identify any areas where ITU requirements are not being satisfied and enhance the protocols to fix the gaps.

3.2.3 Anatomy

The key ASON-related standards are as follows.

- Architecture for Automatically Switched Optical Networks (G.8080, formerly known as G.ason)
- Distributed Call and Connection Control (G.7713, formerly known as G.dccm), which covers signaling
- Architecture and Requirements for Routing in the Automatic Switched Optical Networks (G.7715, formerly known as G.rtg)
- Generalized Automated Discovery Techniques (G.7714, formerly known as G.disc).

Various protocols have been held up to the ASON architecture to see how well they fit, and alongside the core ASON specifications, ITU is also working on defining protocol profiles that will be ASON-compliant.

- PNNI based signaling (G.7713.1)
- Generalized RSVP-TE based signaling (G.7713.2)
- Generalized CR-LDP based signaling (G.7713.3)
- Discovery for SONET/SDH, incorporating some aspects of LMP (G.7714.1).

So, when it comes to selecting ASON-compliant protocols, the ITU currently suffers from the same curse as the IETF, except on a greater scale—too many signaling protocols all meant to do the same thing. As with the IETF, this is for perfectly healthy commercial reasons. Members have vested interests and loyalties to particular technologies, and it is a fact of life that a company with a significant investment or belief in one technology is not going to withdraw their support for that technology. At least, not without a fight.

3.3 The role of the OIF

The OIF is effectively located in the demilitarized zone between the ITU and the IETF. It numbers among its members both ITU and IETF exponents and has therefore been the crucible where compromises between ASON and GMPLS have been struck.

The mission of the OIF is to accelerate the uptake of optical networking technology, and therefore the two key outputs of its work are published implementation agreements, and interoperability demonstrations showing those agreements in action.

On the one hand, the OIF is in a unique position to stage the debate between the ASON and GMPLS protagonists, as it provides a forum where they are forced to explain their terminology and arguments to the other side. On the other, there are strong forces pulling it in different directions, and not all participants end up happy with the outcome.

The main output of the OIF's control plane work so far is the "User Network Interface 1.0 Signaling Specification" (OIF implementation agreement OIF-UNI-01.0), a fusion of high priority ASON requirements with a profile of various GMPLS protocols (RSVP-TE, CR-LDP and LMP). The OIF conducted a successful interoperability demonstration of an interim version of this specification based on RSVP-TE at SuperComm 2001.

The full 1.0 version, which is fairly close to the interim version (but not as close as anticipated), has not yet been publicly interop tested, though such testing is happening in private. The specification is still in early stages of deployment, partly because vendors have only recently begun to add UNI 1.0 capability to their devices and partly because of differing views over its priority.

The OIF is also working on a second version of the UNI specification that adds new features requested by its carrier members, and an E-NNI implementation agreement. (The UNI and E-NNI concepts are both discussed later in the paper.)

4. Differences in methodology

The most important cultural difference between the IETF and the ITU is in the way that they go about developing standards, where the key theme is the balance between pragmatism and perfectionism.

4.1 IETF methodology – Darwinists?

The methodology of the IETF is highly pragmatic. The process by which protocol features are added to GMPLS is that an individual submits an Internet Draft describing the new feature, and if the draft gets sufficient consensus, the feature is incorporated into the protocol. It is left as an exercise to the reader of the draft to figure out whether a given feature is useful and whether their company should implement it in their device or network. For a protocol feature to become standardized, the IETF requires that interoperating implementations of that feature already exist.

This process can result in temporary windows where there are duplicate protocol features, features that are inconsistent in design, or features being used in ways that they weren't designed for. The IETF tackles this by creating architecture and framework documents that look at the big picture. The process of natural selection combined with review is effective at paring down the number of features and improving the specifications.

Few GMPLS drafts contain full FSM (finite state machine) descriptions for their protocols, rigid descriptions of all the possible types of errors and how they are handled, or abstract models showing information flows between components in the network. Any important omissions are expected to be found in review or in interoperability testing, and corrected. Again, a Darwinistic approach ensures that specifications evolve in such a way that they will serve the community well. (Consistent with evolution, connoisseurs of GMPLS RSVP-TE will recognize that the RSVP protocol certainly comes with its fair share of vestigial tails and appendices.)

For a purist ITU-er, the whole process can appear as over-hasty—how does the IETF know that the protocols will work if there is nowhere that defines the function of the network? And what hope do implementers have, if a protocol specification doesn't have a complete explanation of what it means for an implementation to conform to it?

Others would argue that without this fast-moving pragmatic approach, the Internet itself might not have been so successful.

4.2 ITU methodology – Creationists?

By contrast with the IETF, the ITU is a lot more careful—some would say perfectionist. The goal of ASON is to specify in detail how an optical transport network operates, including some of what the IETF would consider implementation details. This requires a full and thorough requirements process, and can result in features that look useful on paper, but which no-one knows for sure whether there is a business case in implementing.

A key difference between the ITU and the IETF is that the ITU avoids the iterative, “suck it and see” approach of implementing and deploying a feature before standardizing it. Instead, they perfect the standard so as to minimize problems during deployment.

Unlike in the IETF, there is no natural selection to act as a brake to feature creep, just the judgment of the people creating the specification. However, this holistic (if you like, Creationist) approach means that the ITU specs convey a clear and consistent vision earlier in their development cycles than the IETF.

To a hardline IETF-er, the ITU may appear to operate in an ivory tower, creating a perfect architecture that could never get implemented. One thing that particularly inflames the CCAMP mailing list is when ITU-ers criticize an Internet Draft for missing a requirement that is perceived as not fully defined or low priority in the IETF, and as a result, delay its progress along the path to standardization. This has been resolved for the signaling protocols by starting a “GMPLS for ASON” Internet Draft (currently known as draft-lin-ccamp-gmpls-ason-rsvpte-oo) that resides alongside but independently of the main GMPLS drafts. However, at the time of writing, similar struggles (and very heated ones) are underway related to LMP and its discovery function.

5. ASON versus GMPLS business model

By looking at the history and methodology of the protocol development and the line-up of companies that supports each set of standards, it is possible to characterize the differences in business models being employed by exponents of ASON and GMPLS. A convenient analogy here is the fable of the race between the tortoise and the hare, though in this case the winner is far from clear.

5.1 ASON – the tortoise

We know already that, as far as possible, ASON supporters want to get everything right first time. It follows that companies who support ASON had better be large and stable enough to be able to live without it for a while. A generous estimate at the time of writing is that ASON may be complete in late 2003-2004.

Large incumbent network operators are likely to view ASON favorably, as in times of economic slowdown, there is little to be gained, and much to be lost, from spending money on building new networks, so why not invest in research that will make those networks better when the climate is right? ASON's strong focus on maintaining compatibility with existing transport network protocols and providing a smooth upgrade path is also essential reassurance for operators with large centralized or proprietary optical networks.

Similarly, equipment vendors who did not get caught up in the GMPLS "gold rush" at the turn of the millennium and whose target market is the large operators are also likely to support ASON.

5.2 GMPLS – the hare

In 1999-2000, there was a move towards faster, higher-bandwidth provisioning. This was particularly attractive for small, fast-moving competitive network operators who aimed to deploy cheap new networks and woo users away from the traditional incumbent operators. In response, a flood of start-up optical equipment vendors were created, who were under extreme time pressure to get their devices on the market in order to stay alive. At the same time, a number of existing large equipment vendors also saw customer demand for optical switching equipment, and invested in GMPLS early to stay ahead of the competition.

Both types of companies provided the drive in the IETF for rapid development of signaling and routing protocols suitable for optical switching. Many of the GMPLS protocol features were proposed to address specific needs of vendors or operators and those requirements were often identified during deployment.

The customers of both types of vendor are interested in standards as they provide the promise of future interoperability. This is an important safety net because it means that operators are not “locked into” a particular supplier and can shop around when expanding their network.

The common mantra is "keep it simple". They believe in an incremental solution rather than trying to build a Ferrari before building a Ford Model T. Or that is the way they would like to see it anyway.

It is news to no-one that many start-ups fell by the wayside in the recent economic slowdown, but that does not change the fundamental rationale behind this business model, namely to get new technology deployed as soon as possible, and to keep at the front of the standardization game.

6. “GMPLS everywhere” versus ASON reference points

This section leaves behind the cultural differences and moves onto the technical differences. However, all of the technical differences can be traced back to the history of the standards development and the priorities of the developers, in other words, the factors we have discussed in the previous sections.

6.1 GMPLS everywhere

GMPLS switches are seen as operating in a GMPLS-only cloud of peer network elements. Nodes at the edge of the cloud are capable of accepting non-GMPLS protocol data and tunneling it across the GMPLS cloud to other edge nodes.

All the nodes and links that constitute the GMPLS network share the same IP address space and information is shared freely between nodes. In other words, GMPLS implies a trusted environment.

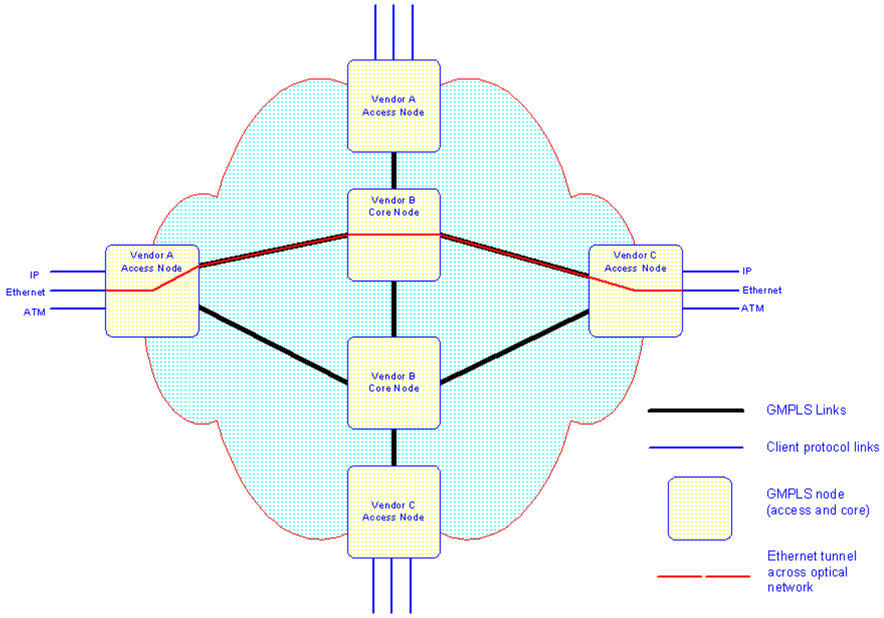

Figure 1 - Simple GMPLS network showing Ethernet tunnel

When full data plane interoperability is achieved, any of the network elements in the cloud may be swapped for a different vendor’s network element. Until then, GMPLS can be used to interface between groups of network elements from different vendors.

6.2 ASON reference points

By contrast with the "GMPLS everywhere" approach, a key principle of ASON is to build in support for legacy network devices explicitly into the architecture. Full multi-vendor interoperability is seen both as a low priority and unrealistic to achieve in the near term, not least because of data plane compatibility issues.

ASON views the network as composed of domains which interact with other domains in a standardized way, but whose internal operation is protocol-independent and not subject to standardization. The interface between such domains is known as the exterior node-to-node interface, or E-NNI. E-NNIs can also be usefully classified into "intra-operator" and "inter-operator".

The I-NNI (interior NNI) is the vendor-specific, proprietary interface used within a single-vendor domain.

The conception of the network is also extended more widely than in GMPLS, to allow users to participate in the automated control plane. Here, the "user" is an endpoint device that requests the services of the transport network rather than provides them. In ASON, users can request connection services dynamically over a user-network interface, or UNI. In GMPLS, the closest thing to an ASON user is simply a GMPLS edge node, but this is not an exact mapping of the ASON concept.

The ASON way of looking at the network is not all that different from the GMPLS picture, once you

- relax the definition of a GMPLS "node" so that it does not always correspond to a single network element, but can instead be a group of network elements, or a proxy operating on their behalf
- redraw the boundaries of the network clouds to illustrate UNI, I-NNI and E-NNI interfaces.

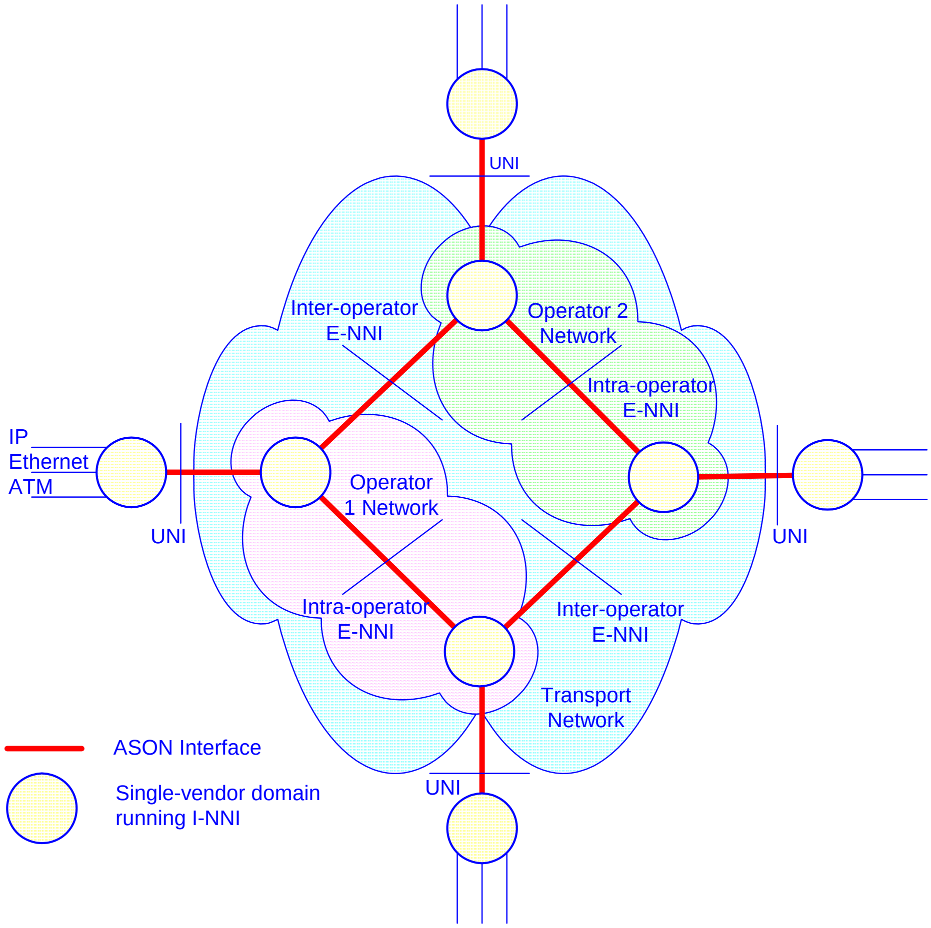

Figure 2 – simple ASON network showing UNI and E-NNI reference points

The UNI, E-NNI and I-NNI are known as “reference points”, and the UNI and E-NNI indicate the locations in the network where standardized protocols will need to be used. Each reference point has different requirements on the degree of information hiding that occurs at that reference point.

- The UNI is an untrusted reference point, and hides all routing and addressing information pertaining to the interior of the network from the user. ASON is very clear on the fact that users should belong to a different address space from internal network nodes, and this means that when GMPLS is mapped onto the ASON UNI reference point, the usual IP address cannot represent a user.
- The I-NNI is a trusted reference point. Full routing information can be flooded.

- The inter-operator E-NNI is a semi-trusted reference point. Some degree of routing information is exchanged to allow routes to be calculated across the network, but network internals are hidden to avoid leakage of confidential information between operators.
- The intra-operator E-NNI is either trusted or semi-trusted, depending on the administrative structure of the particular operator.

6.3 Where do the conflicts arise?

The UNI requires new features that are not provided in core GMPLS.

First, new addresses need to be assigned to users of the network in order to maintain complete separation of the user and the network addressing spaces. This is a security requirement of the operators who are supporting ASON. Next, because no routing information is allowed to flow across the UNI, the user cannot calculate suitable routes itself. Instead, it must pass its requirements across to its neighbor in the network. Finally, the user needs to have an expectation of what requirements the network can actually satisfy in advance, which creates the need for a start-of-day service discovery process.

The initial work to define the UNI profile of GMPLS has been done by the OIF in the UNI 1.0 specification mentioned earlier. This involves creating a profile of the two GMPLS signaling protocols that satisfies the signaling requirements above, and also enhancing the LMP protocol to include service discovery. The ITU has both influenced and drawn heavily on the OIF work in this area.

Another gap between the ASON architecture and the current GMPLS protocol definition is the ASON requirement to allow call setup signaling, as distinct from connection setup. An ASON “call” is an association between two user endpoints. The concept of a call, which is inherited from telephony protocols, is problematic to map onto GMPLS because

- GMPLS does not have “users” in the ASON sense of the term
- GMPLS signaling already has a built-in association between endpoints, so an ASON call looks like duplication of function.

There are proposals on the table to extend GMPLS signaling to include ASON call setup, which will give the ITU-ers the support they need, but are likely to meet resistance from pure GMPLS vendors who perceive them as unnecessary.

Moving onto routing, it is clear from the above that an ASON network will have a requirement to flood user address reachability that will not be supported by unmodified GMPLS routing protocols. Apart from that, to a casual observer, it might look as if trusted E-NNI routing requirements can be met by intra-area protocols such as OSPF-TE, and semi-trusted E-NNI routing requirements can be met by an inter-area protocol such as BGP.

This would certainly be a fairytale ending for the IETF, as it would prove that they were right all along, and ASON was finally getting around to concluding the same thing. However, this is to miss the most fundamental technical differences between the two groups, which relate to

- the ASON layer model
- hierarchical routing.

These topics are covered in the next two sections.

7. How layered is layered?

A key concept in optical networks is that of layering. In this paper, by "layering", we are referring to the ability of a network to nest finer-granularity, lower-bandwidth connections over coarser-granularity, higher-bandwidth connections using a multiplexing function. (In GMPLS terminology, these are known as hierarchical LSPs.)

7.1 Common aspects of layering

In layered networks, a connection is set up at a lower layer (n -1) in order to provide a link at a higher layer (n). This is to say that the connection endpoints at layer n-1 become directly adjacent at layer n.

This is best understood with a diagram (where the "MUX" depicts the adaptation and termination functions that allow traffic from a higher layer to be multiplexed over a lower layer).

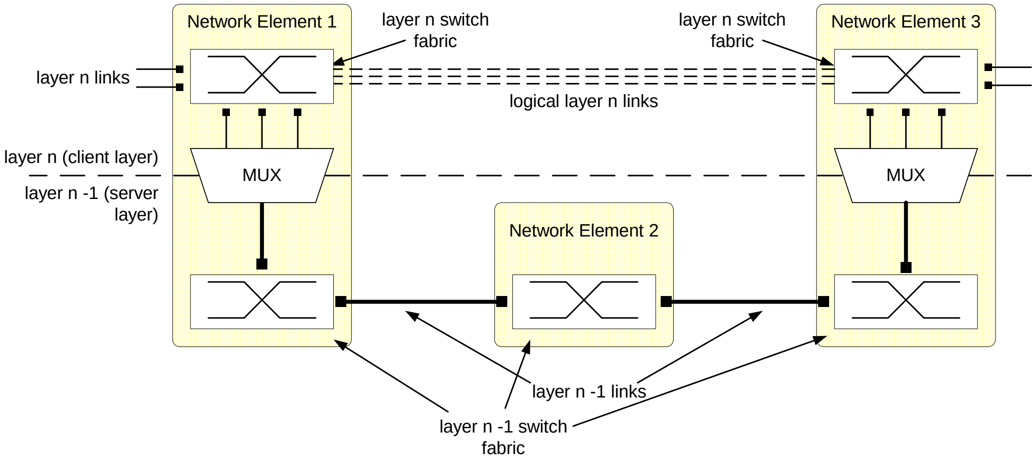

Figure 3 – layering

Thus, connection setup and teardown operations at layer n-1 are used to modify the network topology at layer n.

7.2 Links

It is reassuring to note that both the ITU and IETF agree that a network is a set of nodes connected by a set of links. However, the agreement more or less ends there. The sticking point is the function of a link, in terms of the types of traffic it can carry.

In GMPLS, a link is defined to be capable of supporting multiple different layers of switched traffic. For example, in GMPLS routing, a node can indicate whether it is any combination of lambda-switch capable, TDM capable or packet-switch capable for a given link. A higher-layer link realized over a lower-layer connection is known in GMPLS as a “virtual link”.

In ASON, a link is defined to be capable of carrying only a single layer of switched traffic. A link realized over a real physical medium is indistinguishable from one realized over a lower-layer, higher-bandwidth connection from the point of view of signaling, routing and discovery.

This allows and requires each layer of the network to be treated separately. “Treated separately” means that for each layer, there is a layer-specific instance of the signaling, routing and discovery protocols running.

(Note that with hierarchical routing, there are actually several instances of the routing protocol operating within a single layer: one instance for each routing hierarchy level. Routing controllers may maintain and advertise a separate topology for each switching layer in the network. Then, at a given layer, they may also structure that topology information into more or less abstract levels prior to distributing it. Hierarchical routing is discussed in more detail in the next section.)

The differences between the ITU and IETF here can be partly attributed to the fact that IETF routing protocols have only traditionally been required to deal with a single layer—the IP layer, whereas the ITU has defined a number of layered transport plane technologies and the terminology to go with them.

7.3 Layered signaling

Signaling is uncontroversial in this area, as both groups view it as intrinsically single-layer. This is because the purpose of signaling is to set up a switched connection, and connections are between endpoints at the same switching layer.

It is certainly possible in the course of signaling for lower-layer operations to be invoked on demand, but this is best seen as a case of multiple instances of signaling at different layers, rather than a single instance of signaling that spans layers. The most likely location for this kind of invocation is at the UNI, in cases where the network uses a larger switching quantity than is used over the UNI link. However, operators are understandably nervous about allowing high-cost connections to be set up automatically on demand in this way.

7.4 Layered routing

In GMPLS, a real physical fiber might be represented by OSPF-TE as a single logical link with multiple switching capabilities. By contrast, in ASON, the multiple logical links supported by the fiber must be advertised at their respective layer in the routing protocol.

The ITU see this strict requirement on routing layering as crucial to allowing scalable administration of large networks, as it allows each layer to operate independently of any other layer. Adding more layers does not increase the complexity of route calculations or information flooding within a particular layer, only the entity that arbitrates between the layers at each node. ITU-ers see the IETF solution as a “munge” of layers, forcing the inter-layer complexity to be resolved either by human operators or by route computation algorithms, neither of which come cheap in their different ways.

By contrast, many in the IETF see this requirement as over-engineered and actually unscalable. Each new layer adds many logical adjacencies and links compared to the “munge” solution (for want of a better word), creating the specter of bloated memory requirements for network elements and greatly increased traffic in the control network. Furthermore, each link and node at each layer requires its own unique identifier, so there is a need for a large address space capable of accommodating multiple layers.

While there is a significant conceptual mismatch here, there are ways that the GMPLS routing protocols can be used in a strictly layered application like ASON. There are two broad options.

- Run an instance of a GMPLS routing protocol for each switching layer.
- Find a way to multiplex information about multiple switching layers over a single instance of the routing protocol using the existing support for multiple switching types, and then separate it out again prior to constructing the routing database at each routing controller entity.

As indicated above, finding an addressing scheme that allows clean isolation of layers could be the biggest sticking point here.

7.5 Layered discovery

The original LMP draft does not cover the case where multiple instances of LMP are used at different switching layers. However, the LMP-WDM extensions show that the IETF is envisaging running LMP at two layers, and this could in theory be extended to a fully layered ASON model. As with routing, the issues will not be so much in using the protocol in a layered environment as finding a structured addressing scheme that will allow each layer to have its own address space.

It should also be noted that the functions fulfilled by LMP do not map exactly onto the ITU conception of discovery as described in G.7714, and there is a heated debate currently underway in CCAMP about whether LMP should be enhanced to include technology-specific fully automated discovery.

Some in the ITU also feel that control channel management does not belong in the same protocol as LMP's other functions of fault localization and link property correlation.

The outcome of this debate could either be that additional extensions are included in the IETF LMP draft, or (probably more likely) that the required extensions will be progressed independently, whether in the IETF, OIF or ITU. The OIF has already developed SONET/SDH neighbor discovery extensions to LMP as part of UNI 1.0.

8. Routing hierarchy in optical networks

8.1 Hierarchical routing

In order to provide a complete description of a network layer, it is necessary to provide information about all of its nodes and links. However, flooding a complete topological network description in a routing protocol becomes impractical once the network grows beyond a certain size (low hundreds of nodes is often quoted for IP networks), due to the frequency of updates and the large number of consumers of those updates.

In order to scale networks that use distributed routing beyond a certain size, it is necessary to reduce the amount of information being flooded. First, the network is administratively partitioned into routing areas. Then, routing databases are populated with more detailed information about the local routing area, and less detailed information about remote routing areas. Routing areas can themselves be partitioned recursively, creating a hierarchy of routing information that varies in its level of summarization. A routing protocol instance runs at each level of this hierarchy.

8.2 Link state and path vector routing

There are two broad approaches to hierarchical routing already in use in packet-switching networks: using path vector routing at the top level of the hierarchy, as featured in BGP, and using fully hierarchical link state routing, as featured in PNNI.

- BGP floods path vector information rather than link state information. In order words, it advertises routes to destinations, not network topology. Where multiple destinations are reachable via the same route, they are aggregated, so that only one route is advertised. When a single destination is reachable via multiple routes, the least costly route is retained and the others are discarded. A link state protocol such as OSPF or ISIS runs below BGP, creating a typically two- or three-level routing hierarchy.

  It goes without saying that BGP is pretty well field-hardened (if you downloaded this white paper from our website, then you just used it).

- PNNI creates a hierarchy of routing controllers all with a link state view of the network and can be run recursively at each level of the hierarchy, unlike BGP, which just runs at the top level of the hierarchy. Higher level routing controllers have a wider view of the network but more abstract information about the nodes and links. Lower level routing controllers have a narrower view of the network but detailed information about the nodes and links. PNNI is not limited to the two or three hierarchy levels found in IP networks.

  PNNI is a proven, mature and highly scalable protocol, but its multi-hierarchy routing features have not been widely deployed, especially in multi-vendor networks.

The crucial difference between these two methods of routing abstraction is that it is not possible to calculate routes using path vector information. This is for the simple reason that the path vector information already *is* a pre-calculated route.

Here we come to the core of the problem with using path vector information for optical networks. Whereas in IP routing it does not particularly matter which links a particular packet traverses to get to its destination, in circuit-switched networks, an attempt to set up a connection over the "wrong" set of links will either simply fail, or worse, could be an extremely costly mistake for the network operator. For example, if the operator is going to be penalized for any service interruptions, it had better be sure that its connections use protected links.

Path vector protocols advertise pre-calculated routes. How can the initiator of an optical connection be guaranteed to find a pre-calculated route satisfies the constraints of a particular connection? The number of potential combinations of constraints is large, meaning that it is highly complex to create a strategy for publishing several routes to the same destination, each calculated using a different set of constraints.

The conclusion in the ITU is that path vector information will not be sufficient for large-scale end-to-end optical network routing and that a fully hierarchical link state protocol is required. The IETF seems divided on the subject, although consensus may be further off because CCAMP's multi-area research has up till now received much less attention than the higher priority single-area work.

While it is a lot clearer how constrained path computation will work in a fully hierarchical routing scheme, the complexity here lies in the process of abstracting and summarizing a lower level in the hierarchy to present a meaningful and useful topology at a higher level—there is skepticism from some in the IP community about whether this is practicable at all and also about whether more than three levels of hierarchy are actually needed.

8.3 OIF DDRP

In early 2002, a proposal appeared in the OIF to enhance OSPF-TE to turn it into a hierarchical link state routing protocol, known as a DDRP (for "domain-to-domain routing protocol"). The choice of OSPF was fairly arbitrary, and largely down to its familiarity and widespread use in the IP world.

Subsequently, the DDRP work was broken up into two strands: a protocol-independent description of requirements and architecture; and two protocol-specific documents, one based on OSPF and the other based on ISIS. When this work is complete, the OIF will make a decision about which of the two DDRP flavors to adopt for its E-NNI implementation agreement.

The protocol changes in each case are fairly minor. However, the decision to use a DDRP is of major consequence, as any body that adopts DDRP is effectively giving up on the IP routing model for optical network provisioning, and moving to a fully hierarchical model.

OSPF-based DDRP was shot down in flames when it appeared in the OIF in early 2002. People believed that it was too drastic, too soon and written without sufficient consensus from the rest of the body. However, as of the July 2002 OIF meeting, the DDRP concept returned with a vengeance, and the OIF agreed to hold a public interoperability demonstration of an "interim" intra-operator E-NNI, including a limited form of OSPF-based DDRP at OFC 2003. This is an important milestone, as it will show optical vendors and operators beginning to get to grips with hierarchical optical routing, albeit in a simplified network topology.

Following the precedent of the other OIF work, it seems likely that the ITU will adopt the OIF's DDRP work as the basis for ASON-compliant routing protocols (as specified in G.7715). It is not nearly so clear whether or not the IETF will be dislodged from the view that a three-level hierarchy is all that is required in the near future.

If you would like more details of the software function required for the OIF OFC demo, please contact Metaswitch (protocols@metaswitch.com).

9. Conclusion

Overall, optical control plane standardization is an important and exciting area, which presents some complex technical challenges intermingled with plenty of political pitfalls.

Although this paper has depicted the ITU and IETF as warring parties by way of illustrating their very different approaches and priorities, the real battle to be won is not for control over the standardization process, but the battle for compromise and consensus amid these pitfalls.

This process naturally gives rise to controversy and the occasional skirmish, but such a lively debate is a healthy sign for the industry, provided it can, so to speak, balance the yin of ASON with the yang of GMPLS. It is crucially important to the success of the optical control plane that these debates result in constructive compromise between the ITU, IETF and OIF.

10. Background reading

10.1 ITU recommendations

ITU-T Rec. G.807	Requirements for the Automatic Switched Transport Network (ASTN)
ITU-T Rec. G.8080/Y.1304	Architecture for the Automatically Switched Optical Network (ASON)
ITU-T Rec. G.805	Generic Functional Architecture of Transport Networks

10.2 IETF drafts

Note that all Internet Drafts are work in progress and may be subject to change or may be withdrawn without notice.

draft-ietf-ccamp-gmpls-architecture	Generalized Multi-Protocol Label Switching (GMPLS) Architecture
draft-ietf-ipo-ason	Automatic Switched Optical Network (ASON) Architecture and Its Related Protocols
draft-ipo-optical-inter-domain	Optical Inter Domain Routing Considerations
draft-lin-ccamp-gmpls-ason-rsvpte	Generalized MPLS (GMPLS) RSVP-TE Usage and Extensions For Automatically Switched Optical Network (ASON)
draft-aboulmagd-ccamp-crldp-ason-ext	CR-LDP extensions for ASON

10.3 OIF documents

OIF-UNI-01.0	User Network Interface (UNI) 1.0 Signaling Specification

10.4 Metaswitch White Papers

Metaswitch has published various other White Papers on MPLS and other topics. These can be downloaded from our website, at http://www.metaswitch.com.

11. About Metaswitch

Metaswitch is a privately owned technology company based in London, UK. We have US offices in Alameda, CA, Reston, VA, and Boxborough, MA.

Our Network Protocols Division is the leading developer and supplier of (G)MPLS, OSPF(-TE), ISIS(-TE), BGP, VPN, RIP, PIM, IGMP, MLD, ATM, MGCP, Megaco, SCTP, SIP, VoIP Conferencing, Messaging, Directory and SNA portable products. Customers include Alcatel, Cisco, Fujitsu, Hewlett-Packard, Hitachi, IBM Corp., Microsoft, Nortel and Sun.

Our company culture focuses on building software of consistently high quality, developed and supported by engineers who are with Metaswitch for the long term.

- Founded in 1981, we have over 450 employees, of whom 280 are engineers. The average length of service of engineers at Metaswitch is 8 years, and the annual attrition rate is 3%.
- Throughout this period, Metaswitch has been consistently profitable with profits exceeding 15% of revenue. 2007-2008 revenues were $118m with $22m profit.
- Over 90% of revenue is generated from exports and 80% is from customers in the US (so we are very used to working with American companies).
- The company is privately held by top-tier investment firms Francisco Partners and Sequoia Capital, as well as the Employee Benefit Trust (EBT). As part of this ownership structure, Metaswitch distributes a share of profit to all employees, equitably rewarding them for their contribution and encouraging long-term commitment.
- As a private company with an emphasis on long-term stability, we are not driven by the short-term requirements of quarterly profit statements. This means that we can concentrate on providing software as we would like – that is, developing high quality implementations of complex technologies.

Our routing protocols are designed from the ground up to address next generation networking issues such as massive Internet scalability, optical routing at multiple layers, virtual routing, MPLS and TE/CSPF, and VPNs.

DC-MPLS, DC-VPN Manager, DC-BGP, DC-OSPF, DC-ISIS, DC-IGMP, DC-PIM and DC-LMP provide a complete set of solutions for optical and packet control plane requirements. These include integrated VPN solutions for BGP/MPLS VPNs and Martini.

All of the Metaswitch protocol implementations are built with scalability, distribution across multiple processors and fault tolerance architected in from the beginning. We have developed extremely consistent development processes that result in on-time delivery of highly robust and efficient software. This is backed up by an exceptionally responsive and expert support service, staffed by engineers with direct experience in developing the protocol solutions.

Nic Larkin was the senior architect for Metaswitch's UNI implementation. He has contributed to several IETF, ITU and OIF documents, and plays a key role in product architecture and standards-based development in Metaswitch's Network Protocols Group.

Metaswitch Networks
100 Church Street
Enfield
EN2 6BQ
England
+44 20 8366 1177
http://www.metaswitch.com