NTAC response to AOAC charge regarding MPLS and SONET technical options for circuit delivery

Executive Summary

The AOAC charged the NTAC to evaluate the impact of moving the circuit network from the current SONET based system to a MPLS system. The paper below describes the specific charge, the assumptions and definitions, the pros, cons, and risks, and provides an alternative and conclusions.

The paper concludes that moving from the SONET based services to MPLS based services will allow for similar capabilities and will not degrade any existing application of which the NTAC is aware. Also included is a recommendation that the alternative of an Ethernet VLAN approach be evaluated.

Charge

Some of the Internet2 network circuits use MPLS and some use SONET. If Internet2 moves their circuits to a MPLS environment, is the Internet2 community comfortable with the level of service that such an environment provides? Please provide a response to this by November 20th.

Assumptions and Definitions of Circuit Services

DWDM Wavelength

A circuit implemented as a framed DWDM wavelength is the lowest layer and most dedicated circuit that can be provided over a pair of shared fiber. It is completed dedicated to the users at the endpoints and cannot be shared without placing higher layer equipment atop it.

Ethernet over SONET (SONET-based)

A SONET framed circuit rides atop a bundle of STS frames that are assigned to the circuit out of an aggregate pool of available bandwidth. The STS frames assigned to a circuit are dedicated exclusively to that circuit and cannot be used for more than one circuit. Performance should be similar to a dedicated wavelength, though multiple circuits can be provisioned over the same lambda. SONET-based circuits are generally point-to-point.

The current ION service uses Ethernet over SONET. It is a physically separate service from the Internet2 IP service. The Ciena Core Directors provide a SONET-based core with Ethernet encap/decap done on Ethernet line cards in the CD. Customer terminations can be Gigabit Ethernet or 10 Gigabit Ethernet. Circuits are all point-to-point, and bandwidth can be assigned in 50 Mbps increments.

Ethernet over MPLS (MPLS L2VPN)
A MPLS L2VPN circuit is carried over a shared IP substrate as an MPLS LSP. A MPLS L2VPN circuit provides no inherent bandwidth protection, though vendors may provide the ability to manage bandwidth via QoS, MPLS traffic engineering or other proprietary means. If no bandwidth protection is applied, performance should be similar to what can be achieved on a shared IP network (e.g. performance will be relative to the other shared traffic on the backbone interconnect).

Two types of MPLS L2VPN services that should be considered are Virtual Wire (also known as Pseudo Wire) and Virtual Private LAN Service (VPLS). A Virtual Wire provides point-to-point service between two endpoints, similar to a SONET circuit. VPLS can provide multipoint service between multiple endpoints, analogous to a native Ethernet VLAN service.

The MPLS-based service can be provisioned on existing customer-facing Ethernet interfaces that provide IP service, or on separate Ethernet interfaces that are specifically for this circuit service. The circuits are created using MPLS LSPs, and can use a number of different encapsulation methods to get traffic to and from customer terminating equipment.

Pros, Cons and Risks

Key Aspects

MEF Service Types - Two service types are defined by the Metro Ethernet Forum: E-Line and E-LAN[1]. These service types roughly translate to point-to-point (E-Line) and multipoint-to-multipoint (E-LAN) capable services.

- **SONET** - ION as configured supports E-Line.

---

[1] See http://metroEthernetforum.org/PDF_Documents/metro-Ethernet-services.pdf for details.

- **MPLS** - A new circuit service using MPLS could be engineered to support E-Line and/or E-LAN.

Resiliency - In the context of this discussion, resiliency refers to the ability of a circuit to recover from the failure of the underlying transport upon which the circuit is built (due to a link or node failure). Between two points on the backbone (ideally the ingress and egress points) a resilient circuit would have the capability to restore service between those two points after detecting a failure by avoiding the failed link or node. The generally acknowledged acceptable circuit restoration time is 50 ms.

- **SONET** has an Automatic Protection Switching (APS) feature that would provide 50 ms. restoration time. However, the current ION service does not implement APS. In order to guarantee end-to-end protection, the ION circuit customer would need to provide a *redundant* circuit and provide a mechanism for fault detection and cutover in case of a failure Current ION circuits will get the default Ciena Mesh protection, which may restore a circuit within seconds assuming bandwidth is available- but there are no guarantees.
- **MPLS** could be engineered to utilize MPLS Fast Reroute (FRR), which is capable of 50 ms. restoration times. FRR signals an alternate path in advance and ensures that bandwidth will be available when needed by either reserving it, or by preempting other lower priority traffic during an outage.

Traffic Engineering - Traffic Engineering generally refers to the ability to control the path that traffic takes through a network. In the context of this discussion regarding circuit services, the important aspect is the ability to affect the path a circuit takes in order to efficiently manage overall network resources.

- **SONET** uses Virtual Concatenation (VC) to allow for adjustments in circuit bandwidth utilization.

- **MPLS** has a rich set of traffic engineering features such as constraint-based routing and preemption that provide flexible mechanism for managing routing policy and traffic prioritization.

Operations, Administration, Maintenance (OAM) - In comparing SONET and MPLS it is important to keep in mind that SONET is a layer one protocol while MPLS is a layer 2.5 protocol which places it in the data link domain. The relevant detail is that both protocols achieve similar goals using completely different means. In our discussion, the end user experience is the critical metric. Since SONET is so low in the protocol stack it cannot improve performance of overlying protocols such as Ethernet. As an example, even though SONET has very strict transmission rates tied to stratum one clocking, the overlying protocol of Ethernet does not, so the user experience is much more dependent on Ethernet performance. SONET must function well, but primarily provides the path. Similarly, MPLS provides the path but the user experience will have the same sort of dependency on Ethernet. To be clear, both protocols must function well within their own domains but the end result is no better or worse than the weakest link in the protocol stack.

- **SONET** has a rich set of OAM features, including excellent problem isolation down to the node level. Additionally, it has good data integrity tools enforced at the per bit level. It has very strict bandwidth allocation, which guarantees a protected level of throughput. In practice this can result in wasted bandwidth as unused bandwidth cannot be repurposed. Although it is not necessarily pertinent in the context of ION, it does have good resiliency, but at the cost of halving ones effective available bandwidth.
- **MPLS** similarly has excellent OAM tools either directly imbedded as a part of the protocol or in adjunct protocols specifically developed for MPLS such as RSVP. While SONET was developed primarily as a static path protocol (ION and IDC etc try to make it dynamic by automating

provisioning), MPLS was intentionally designed to be dynamic, although in practice, it is often used in a semi-static mode in most carrier networks. The tools supporting MPLS are much more familiar to router operators, such as traceroute and ping. Other operational parameters, such as data integrity are a part of the general router interface OAM, so although they are not directly a part of MPLS, they are a part of the general operating environment. MPLS-TE has good bandwidth guarantees. It does have an advantage; it uses statistical multiplexing which allows unused reserved bandwidth to be utilized by contending needs as long as the reservation is underutilized.

Quality of Service (QoS)

- **SONET** as mentioned before has very strict timing requirements and bandwidth allocation. A result of these attributes is very good QoS. Although, as mentioned above it does not inherently improve the user experience due to the protocol dependencies previously discussed.
- **MPLS** does natively have some QoS characteristics, but its real strength is in the ability to couple it to other router functions such as particular queuing strategies to set delivery and jitter bounds etc. Complementing techniques such differentiated services (diffserv) or policy based forwarding can enforce reasonably well-defined performance parameters.

Fault Detection

- **SONET** has excellent fault detection, providing operators with well defined failure type and quickly identifiable location information. However, again in the context of the current dynamic services portfolio, the native recovery mechanisms are not implemented. Hence, they are not relative to the current considerations.

- **MPLS** uses the RSVP protocol for path failure notifications and natively can use fast reroute (FRR) for fast route restoration. It can operate in one of two modes. In the first case one can have a predetermined backup path, in the second case RSVP will signal for a new path. If adequate resources are available it will dynamically reroute traffic. If resources are not available the reroute request will fail. Independent of MPLS, but in practice, another fault detection protocol is commonly used in MPLS environments, Bidirectional Forwarding Detection (BFD). It provides speedy link fault detection on the order of 10 ms.

Summary

Aspect	***Pros (of switching to MPLS)***	***Cons (of switching to MPLS)***	***Risk (of switching to MPLS)***
MEF Service Types	Adds multipoint service	None	None
Resiliency	Adds 50 ms restoration via FRR	None	None
Traffic Engineering	Greater Flexibility in bandwidth reuse	None	Depending on implementation, may impact other traffic
OAM	Tools are much more familiar to router operators	Lacks native per node reporting; can be mitigated by robust network monitoring	None
QoS	Provides a more hierarchical approach i.e. parameters are independently enforceable	Parameters such as jitter are bounded but not as strict as TDM, unlikely to be an issue	Little to none
Fault Detection	None	Does not have native per node reporting per se, however general router reporting and robust network reporting will mitigate this	Little to none

Alternatives

Metro Ethernet equipment providers have come to market with equipment that can provide QoS and traffic engineering capabilities. When looking at a solution as expensive as MPLS or SONET, an evaluation of an Ethernet VLAN approach is warranted and recommended.

A VLAN-based circuit provides point -to-point or multipoint services via 802.1q Ethernet VLAN tagging of individual Ethernet frames on a shared Ethernet link. A VLAN provides no inherent bandwidth protection, though vendors may provide the ability to manage bandwidth via QoS or other proprietary means. If no bandwidth protection is applied, performance should be similar to what can be achieved on a shared IP network (e.g. performance will be relative to the other shared traffic on the backbone interconnect)

Ethernet VLANs have been used on a nationwide scale in other networks to provide an inexpensive circuit based infrastructure.

Emerging standards for carrier-based Ethernet transport such as Provider Backbone Bridging with Traffic Engineering (PBB-TE, also referred to as PBT) and MPLS Transport Protocol (MPLS-TP, also called T-MPLS) are intended to address some of the limitations of using “native” Ethernet in a carrier environment. These standards should be included in the evaluation as these features become available in vendor implementations.

Cautions

Collapsing the Internet2 network to a single platform based around MPLS could cause Internet2 to lose a key differentiator from commodity providers. Ethernet over MPLS is a common service available from commodity providers. By moving to this strategy, Internet2 runs the risk of diluting its image as a provider of innovative services. Internet2 should continue to provide new and creative ways of providing circuit services to its members.

This may increase the likelihood of purchasing officers seeing Internet2 as providing a service similar to commodity services and requiring that service to go to bid.

Conclusions

Assuming proper engineering and bandwidth management, there is little technical risk in moving from a SONET infrastructure to a MPLS infrastructure. The QoS and traffic engineering capabilities available in MPLS allow for similar capabilities to SONET. The NTAC is unaware of any current application running on the ION SONET service that would be degraded on a MPLS provisioned service.