OAM Mechanisms in MPLS Layer 2 Transport Networks

Rahul Aggarwal, Juniper Networks

ABSTRACT

This article describes OAM in MPLS layer 2 transport networks. MPLS networks used to transport layer 2 traffic are referred to as MPLS layer 2 transport networks. They may be used to connect legacy layer 2 networks and/or provide layer 2 service to a user over a MPLS network. As legacy layer 2 networks migrate to use MPLS for transport, the role of MPLS OAM mechanisms is becoming increasingly important. This is because the converged network must offer the same OAM functionality as existing layer 2 networks. This article emphasizes the importance of end-to-end OAM, while emulating existing layer 2 services using MPLS transport. End-to-end fault detection is described in the context of various layer 2 over MPLS transport network models. The article focuses on state-of-the-art MPLS label switched path and pseudo wire OAM mechanisms being developed by the IETF. This includes fault detection and isolation mechanisms such as LSP-Ping, bidirectional forwarding detection, and virtual circuit connectivity verification. The applicability of each of these mechanisms is given. In some cases it may be possible to carry layer 2 OAM cells end-to-end, while in other cases this may not be possible. The relationship between segment-based OAM mechanisms and end-to-end OAM is described for each of these cases.

INTRODUCTION

Next-generation network designs are increasingly using MPLS as a means for building "converged" networks with a common infrastructure for delivering various services. There is a strong business case for this as MPLS makes it possible to build a converged network that can provide both layer 2 and layer 3 services. This reduces operational and capital expenses. MPLS uses IP-based control protocols, and this fits well with an IP-based network design. It can be used for traffic engineering using constraint-based routing [1]. It provides subsecond protection using fast reroute mechanisms [2]. This makes it possible to transport voice and other traffic that has stringent availability requirements over MPLS. Virtual private network (VPN) services can also be provided using MPLS. Border Gateway Protocol (BGP)-MPLS layer 3 VPNs can be used by the service provider to offer layer 3 VPN services [3]. Also, existing layer-2-based VPNs can be supported over a MPLS network by transporting layer 2 ATM, frame relay, and Ethernet traffic over MPLS. Label Distribution Protocol (LDP) can be used as a control protocol for providing layer 2 point-to-point service over MPLS [4]. BGP can be used as a control protocol for building layer 2 VPNs over MPLS [5].

There is a large installed base of layer 2 asynchronous transfer mode (ATM) and frame relay networks. These networks are used by service providers and regional Bell operating companies (RBOCs) for providing VPN, voice over IP, broadband, and other services. For the reasons mentioned above these networks are increasingly migrating to use MPLS for transport. This migration enables service providers and RBOCs to move to a common MPLS-based network and at the same time leverage their existing investment in layer 2 networks. This is shown in Fig. 1. The mechanism that carries layer 2 traffic over a public switched network is called a *pseudo wire* (PW). An MPLS PW uses an MPLS label for demultiplexing [4]. It is also possible to build layer 2 VPNs using PWs.

Layer 2 transport over MPLS must provide the same functionality and service characteristics as legacy layer 2 networks. One such critical attribute is operations, administration, and management (OAM). Existing layer 2 networks use OAM mechanisms for fault detection. It is imperative that similar OAM functionality be available when transporting layer 2 traffic over MPLS. This is important from an operational viewpoint.

This article describes state-of-the-art OAM fault detection and isolation mechanisms for layer 2 transport over MPLS. Various MPLS OAM tools and their applicability are explained. The article focuses on the Internet Engineering Task Force (IETF) initiatives in the area of MPLS OAM, which are primarily restricted to fault detection and isolation. Mechanisms such as Y.1711 [6], being discussed in the International Telecommunication Union (ITU), are not

0163-6804/04/$20.00 © 2004 IEEE

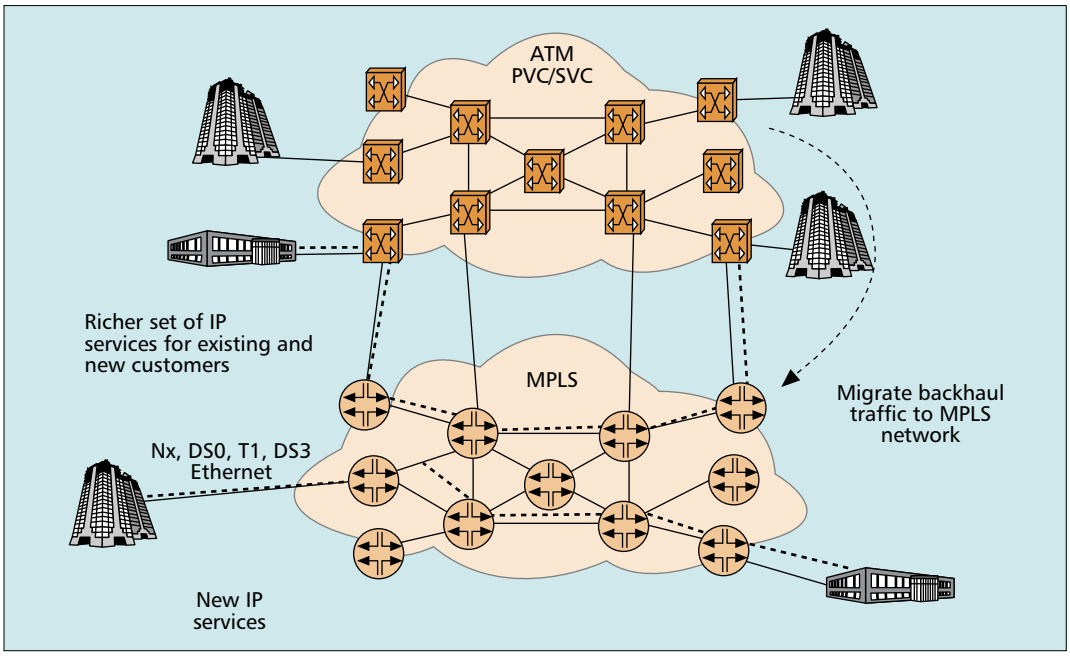

Figure 1. *ATM core offload.*

In the overlay model the MPLS network transparently transports the Layer 2 traffic. Layer 2 signaling protocols are transported transparently by the MPLS network to establish end-to-end Layer 2 connections. In this model PWs are used to inter-connect Layer 2 trunks which carry a number of Layer 2 connections.

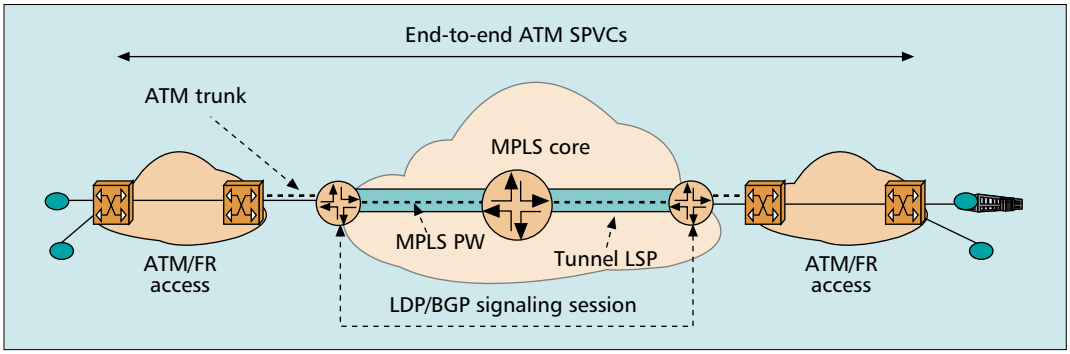

Figure 2. *Overlay model.*

described herein. The following section introduces the concept of an MPLS PW. We then describe different models for layer 2 transport over MPLS and end-to-end fault detection. MPLS PW OAM mechanisms are described next, and a brief look at existing layer 2 OAM mechanisms is provided. The article goes on to describe the relationship between end-to-end fault detection and the segment-based OAM mechanisms.

MPLS PW

An MPLS PW is the mechanism used to carry layer 2 traffic over MPLS. It is a point-to-point entity that interconnects two attachment circuits (AC). An AC is a layer 2 circuit being transported over MPLS. For instance, it may be an ATM permanent virtual circuit (PVC), or a frame relay or Ethernet port. Figure 2 shows an MPLS PW used to interconnect two ATM trunks. An MPLS PW is signaled using mechanisms described in [4, 5]. A PW label is assigned by the egress of the PW to the ingress of the PW and signaled using one of the signaling mechanisms. Multiple PWs can be carried over one MPLS LSP that exists between the ingress and egress nodes. This is shown in Fig. 2. The ingress of the PW encapsulates the layer 2 traffic in MPLS using a two-label stack. The upper label is the MPLS LSP label, while the inner label is the PW label. The PW label is used by the egress of the PW to identify the PW. Between the layer 2 payload and the MPLS label stack the ingress may insert a PW control word (CW). The PW CW can be used for sequencing and fragmentation.

Layer 2 over MPLS Network Models and OAM

In this section we describe two possible network models for transporting layer 2 traffic over MPLS: the overlay and peer-to-peer models. We then look at end-to-end fault detection approaches in the context of these models.

Network Models

Overlay Model — In the overlay model the MPLS network transparently transports the layer 2 traffic. This is shown in Fig. 2. Layer 2 signal-

In the peer-to-peer model, the Layer 2 network peers with the MPLS network. In this case the MPLS network is aware of the individual Layer 2 connections. A PW is used to carry a Layer 2 connection. Thus PWs are used to inter-connect Layer 2 circuits. These Layer 2 circuits may be configured circuits or may be signaled SPVCs.

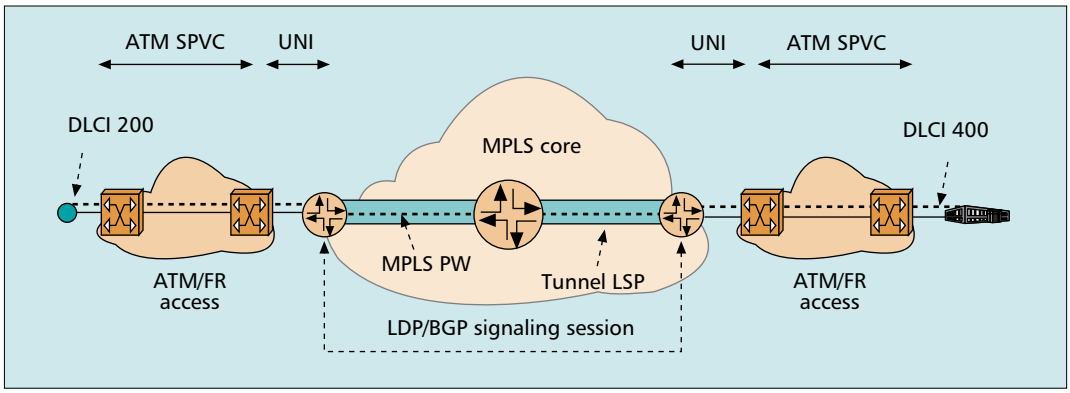

Figure 3. *Peer to peer model.*

ing protocols (e.g., private network–network interface, PNNI) are transported transparently by the MPLS network to establish end-to-end layer 2 connections. In this model PWs are used to interconnect layer 2 trunks that carry a number of layer 2 connections. The MPLS network is not aware of the individual layer 2 connections. For example, all traffic coming in on an ATM port on a provider edge (PE) device can be transported over an MPLS PW, which can be terminated on an ATM port on a remote PE device. This model has the obvious advantage of leaving the layer 2 intelligence with the layer 2 endpoints. Using a PW to transport layer 2 circuit aggregates also has favorable scaling properties.

Peer- to-Peer Model — In the peer-to-peer model, the layer 2 network peers with the MPLS network. In this case the MPLS network is aware of the individual layer 2 connection, as shown in Fig. 3. A PW is used to carry a layer 2 connection. Thus PWs are used to interconnect layer 2 circuits. These layer 2 circuits may be configured circuits or signaled SPVCs. An MPLS PW maps to an individual layer 2 circuit, as shown in Fig. 3. The disadvantage of this model is that the MPLS network is no longer transparent to individual layer 2 circuits. Furthermore, the number of PWs is more than that in the overlay model. However, in certain cases the peer-to-peer model may be useful. One case is when one layer 2 endpoint is on a layer 2 (e.g., ATM) network and the other one is on the MPLS network. Also, if the layer 2 endpoints are on different media and the MPLS network is required to perform the interworking function, the peer-to-peer model can be useful.

End-to-End Fault Detection

Layer 2 transport over MPLS must provide end-to-end fault detection functionality equivalent to ATM networks. This section takes a brief look at end-to-end fault detection in existing ATM networks. Fault detection on the end-to-end path when transporting a layer 2 circuit over MPLS is then described in the context of both overlay and peer models.

Existing OAM mechanisms in ATM networks enable fault detection on an end-to-end layer 2 circuit. The fault detection interval depends on the service characteristics, and may range from under a second to a few seconds. It is to be noted that such fault detection is used to detect faults in the end-to-end path between the ATM endpoints. Hence, a failure anywhere in the end-to-end path is recognized as a fault. Alarms can be issued to alert an operator who can diagnose the fault. Fault detection can also result in the layer 2 connection being rerouted to another path between the endpoints of the layer 2 connection.

When transporting a layer 2 circuit over an MPLS network, the end-to-end path comprises:

- A layer 2 circuit, trunk or SPVC that connects the layer 2 ingress to the MPLS network (i.e., the ingress layer 2 path)
- MPLS PW, which is transported over a MPLS LSP
- Layer 2 circuit, trunk or SPVC that connects the MPLS network to the layer 2 egress (i.e., the egress layer 2 path)

This can be seen in Fig. 3. Each of these segments may have their respective OAM mechanisms. The relationship between these segment OAM mechanisms and end-to-end fault detection depends on whether it is possible to carry layer 2 OAM cells end-to-end between the layer 2 ingress and egress.

In the overlay model described earlier, layer 2 OAM cells are carried end-to-end between the layer 2 ingress and egress. They are transported transparently by the MPLS network. In this case end-to-end fault detection is performed by the layer 2 OAM cells. In the case of ATM, for example, these may be continuity check (CC) cells. Segment OAM mechanisms can be used to verify the status of a particular segment if a fault is detected in the end-to-end path. However, segment OAM mechanisms are not used for end-to-end fault detection.

In the peer-to-peer model with like layer 2 media at the ingress and egress, layer 2 OAM cells may be transported between the layer 2 ingress and egress if the PEs support OAM cell transport. If so, as in the overlay model, end-to-end fault detection does not rely on segment OAM. However, the PEs may not be capable of transporting OAM cells. Also, in the case of disparate layer 2 media at the ingress and egress, the MPLS network may perform the interworking. If the payload is IP, such interworking will result in the MPLS PW carrying IP traffic. The layer 2 header is stripped off at

the PEs. Hence, OAM cells sourced by the layer 2 ingress terminate at the ingress PE. In these cases end-to-end fault detection cannot be performed using end-to-end layer 2 OAM cells. Instead, segment-based OAM mechanisms are used to perform end-to-end fault detection. As a result, interworking between the segment-based OAM mechanisms is also required [7].

In the next two sections we describe the various segment-based OAM mechanisms. In particular, we focus on the OAM mechanisms for MPLS LSPs and PWs. We then describe the role of these segment-based OAM mechanisms when layer 2 OAM cells are carried end-to-end. Interworking between segment-based OAM mechanisms when it is not possible to carry layer 2 OAM cells end-to-end is also described.

MPLS PW OAM Mechanisms

The IETF is developing OAM mechanisms for fault detection and isolation in MPLS networks. This includes OAM mechanisms for MPLS LSPs and PWs. When an MPLS PW is transported over an MPLS LSP, OAM mechanisms for the MPLS LSP can be used to detect failures in the data plane of the transit label switched routers (LSRs). In addition, MPLS PW OAM can be used to detect failures in the forwarding plane on the egress of the MPLS PW. This verifies that the MPLS PW label is indeed present in the MPLS forwarding table of the egress PE and is bound to the PW on which the ingress is sending packets. There are two parts to an MPLS PW OAM solution:

- A mechanism for detecting failures in the data plane and verifying the data plane against the control plane
- A mechanism for identifying OAM packets at the egress of the PW

MPLS LSP Ping [8] is a basic OAM building block in MPLS networks. This is used for providing the first mechanism above. In addition, bidirectional forwarding detection (BFD) can also be used in conjunction with LSP-Ping. There are different means of providing the second mechanism, and one of them is virtual circuit connectivity verification (VCCV) [9]. Let us look at these mechanisms in further detail.

LSP-Ping

When an LSP fails to deliver user traffic, the failure cannot always be detected by the MPLS control plane. An example of this is corruption of the MPLS forwarding table. This issue has been seen in real operational networks. LSP-Ping is a tool that makes it possible to perform data plane fault detection and also verifies the MPLS control plane against the data plane. It enables users to detect traffic "black holes" or misrouting and provides a mechanism to isolate faults. It is modeled after the ping/trace-route paradigm: ping (ICMP echo request) is used for connectivity checks, and trace-route is used for hop-by-hop fault localization as well as path tracing. LSP-Ping specifies a ping mode and a trace-route mode for testing MPLS LSPs. The basic idea is to test that packets belonging to a particular forwarding equivalence class (FEC) actually end their MPLS path on an LSR that is an egress for that FEC. LSP-Ping does this by sending MPLS echo request packets along the same data path as other packets belonging to this FEC. An MPLS echo request also carries information about the FEC whose MPLS path is being verified. This echo request is forwarded just like any other packet belonging to that FEC. In ping mode (basic connectivity check), the packet should reach the end of the path, at which point it is sent to the control plane of the egress LSR, which then verifies that it is indeed an egress for the FEC. In trace-route mode (fault isolation), the packet is sent to the control plane of each transit LSR, which performs various checks that it is indeed a transit LSR for this path; this LSR also returns further information that helps check the control plane against the data plane (i.e., that forwarding matches what the routing protocols determined as the path). LSP-Ping can be used for Resource Reservation Protocol (RSVP) or LDP signaled LSPs. It can also be used for layer 3 VPN FECs and MPLS PWs.

When using LSP-Ping for MPLS PWs the echo request packets are encapsulated with the inner PW label and outer LSP label. The FEC identifying the PW is carried in the packets. The packets are carried to the egress along the LSP path transporting the PW. At the egress they are identified as PW packets, and the control plane verifies that it is indeed the egress for the PW and has allocated the same label carried by the incoming packet. It then responds to the ingress with an MPLS echo reply. An LSP-Ping failure implies a failure in the MPLS PW path. This may be a data plane failure, or the data plane and control plane may be out of sync at a transit LSR or the egress PE. In such a case LSP-Ping can be used on the MPLS LSP carrying the MPLS PW to determine whether the failure is at a transit LSR. If not, the problem can be isolated to the egress of the PW.

BFD

BFD provides a low-overhead short-duration detection of failures in the path between adjacent forwarding engines, including the interfaces, data link(s), and to the extent possible the forwarding engines themselves [10]. In addition, BFD can provide failure detection on any kind of path between systems, including virtual circuits and tunnels (as long as there is some return path, of course).

BFD can be used to detect an MPLS LSP data plane failure. As described in the previous section, LSP-Ping can be used to detect MPLS data plane failures and verify the MPLS LSP data plane against the control plane. BFD can be used for the former, but not for the latter. However, a combination of LSP-Ping and BFD can be used to provide faster data plane failure detection and/or make it possible to provide such detection on a greater number of LSPs. The LSP may be associated with an RSVP session, LDP prefix, layer 3 VPN FEC, layer 2 VPN FEC, or MPLS PW. LSP-Ping is a more compu-

BFD provides a low-overhead short-duration detection of failures in the path between adjacent forwarding engines, including the interfaces, data link(s), and to the extent possible the forwarding engines themselves.

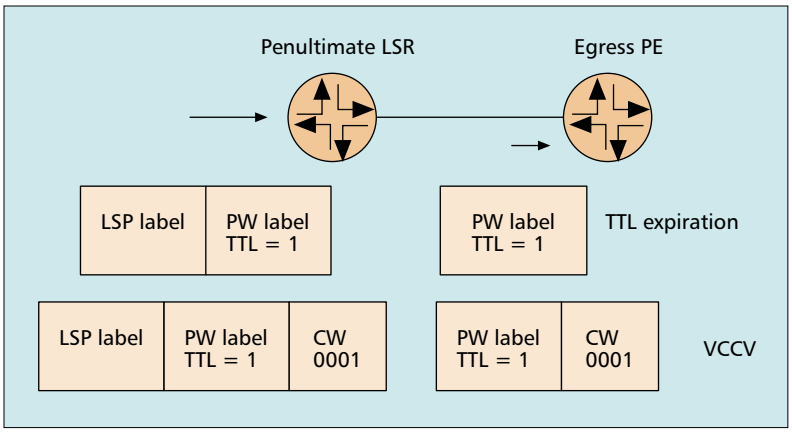

Figure 4. *Two mechanisms for identifying PW OAM packets.*

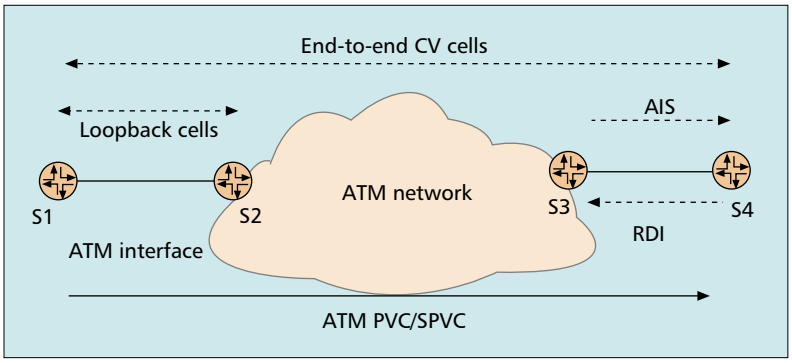

Figure 5. *ATM OAM mechanisms.*

tationally expensive mechanism than BFD for detecting MPLS LSP data plane faults. The control plane processing required at the egress/ingress LSRs for processing LSP-Ping echo request/reply messages is greater than that required for BFD control packets. Furthermore, BFD is a fixed format protocol; hence, it is possible to implement BFD in hardware or firmware. Thus, the use of BFD for detecting MPLS LSP data plane faults has the following advantages:

- Support for fault detection for a greater number of LSPs
- Subsecond fault detection granularity

A BFD session is established for the LSP under consideration. LSP-Ping is used for bootstrapping the BFD session [11]. BFD control packets are then used for fault detection at the required detection interval. BFD control packets are encapsulated in the MPLS LSP. Thus, when BFD is used on a MPLS PW, BFD control packets are encapsulated in the MPLS PW and follow the same data path as the MPLS PW. They are processed at the MPLS PW egress following normal BFD control packet processing. Also, LSP-Ping is used to periodically verify the control plane against the data plane by verifying the MPLS LSP and FEC mappings.

VCCV

One of the issues in MPLS PW OAM is to define a mechanism that enables the egress of the PW to identify the OAM packet. One of the ways to do this is for the ingress to set the time to live (TTL) = 1 in the inner label of an LSP-Ping or BFD packet. The egress of the PW receives an MPLS packet with only the PW label if the outer LSP label was popped by the penultimate LSR. Otherwise, it will receive an MPLS packet with a two-label stack that has an outer LSP label and an inner PW label. After popping the outer label the egress node will process the PW label. In either case the egress node will decrement the TTL of the PW label from one to zero. This will result in TTL expiration; hence, the egress node will kick the packet out of the forwarding path. The packet will then be sent to the OAM processing module. This module will process the LSP-Ping or BFD OAM packet.

However, it may not always be possible to ensure that the penultimate-hop LSR will not overwrite the TTL field in the inner label when penultimate-hop-popping (PHP) is used. If this is seen as an issue, VCCV [9] provides an alternate mechanism for identifying the OAM packet at the PW egress. A special bit is used in the PW control word that follows the inner PW label in order to identify the OAM packet. Thus, the egress of the PW can kick the packet out of the forwarding path on recognizing this special bit. It can then continue with OAM packet processing. The TTL expiration and VCCV-based schemes for identifying PW OAM packets are illustrated in Fig. 4.

Layer 2 OAM Mechanisms

This article does not attempt to describe existing layer 2 OAM mechanisms in detail. However, let us take a brief look at these for completeness. In existing ATM networks alarm indication signal (AIS), remote defect indication (RDI), continuity check (CC), and loopback (LB) cells provide fault management functions. F4 OAM cells are used to monitor a virtual path, F5 OAM cells to monitor a virtual channel. AIS is issued by the local end when an error is detected. RDI is issued by the remote end when an error is detected. CC cells can be used end-to-end for detecting the liveliness of the end-to-end path. LB cells are used for liveliness detection on a segment. Figure 5 shows some of the OAM mechanisms in an ATM network. OAM mechanisms in frame relay networks are not as advanced as ATM networks. They provide OAM between links using LMI or link integrity verification messages.

End-to-End Fault Detection and Segment-Based OAM Mechanisms

As mentioned earlier, the relationship between end-to-end fault detection and segment-based OAM mechanisms depends on whether it is possible to carry layer 2 OAM packets end-to-end between the layer 2 ingress and egress. We will study this relationship when it is both possible to do so and not.

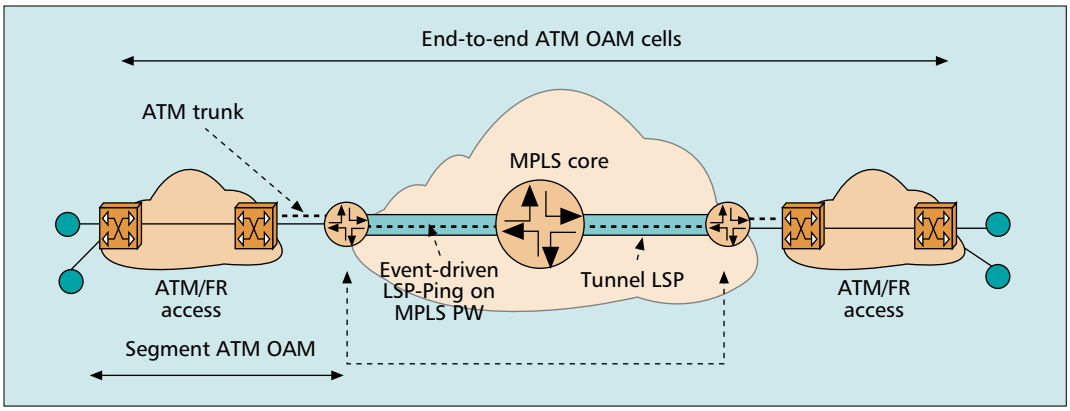

Figure 6. *End-to-end L2 OAM cells.*

Periodic fault detection is also carried out on the L2 ingress path and egress path. If a failure is detected it is conveyed to the other PE. This can be done by sending a BFD message with a special diagnostic code.

End-to-End L2 OAM Cells

As mentioned earlier, it is possible to carry layer 2 OAM cells end-to-end in:

- The overlay model
- The peer-to-peer model, with like end media when the PEs are capable of transporting the L2 OAM cells

When this is the case end-to-end fault detection relies on the end-to-end L2 OAM cells. This is shown in Fig. 6. The interval of fault detection depends on the interval at which the L2 OAM cells are transmitted. Thus, end-to-end fault detection does not depend on segment-based OAM mechanisms. However, these mechanisms are useful for isolating the fault once it is detected. This is shown in Table 1.

Thus, once the end-to-end path is detected as down, L2 segment OAM between the L2 ingress and the MPLS PW ingress can be used to check if the L2 ingress path is up. If L2 ingress segment OAM fails, the fault is in the L2 ingress path. Similarly, PW OAM can be used to check if the PW is up. LSP OAM can be used to check if the LSP on which the PW is carried is up. If the LSP is up but the PW is down, the fault can be isolated to the PW egress forwarding path. Akin to the L2 ingress path, the L2 egress path can also be verified using L2 OAM on it.

It is to be noted that in this case there is no need to periodically transmit PW OAM packets. Thus, LSP-Ping may be sufficient for diagnosing the MPLS PW. LSP-Ping will probably be triggered by an operator after the end-to-end L2 path is detected to be down.

L2 ingress OAM	L2 ingress path
LSP OAM	LSP path
PW OAM	PW egress forwarding path
L2 egress OAM	L2 egress path

Table 1. *Fault isolation with segment OAM.*

L2 OAM Cells Terminated at the PW Ingress PE

L2 OAM cells may have to be terminated at the PW ingress PE if:

- The PE is not capable of transporting the OAM cells.
- The MPLS network is performing layer 2 interworking.

In this case end-to-end fault detection relies on the segment-based OAM mechanisms in the L2 ingress path, the MPLS PW, and the L2 egress path. L2 OAM cells are not sent end-to-end in this case. Interworking between MPLS PW OAM and OAM on the L2 ingress path, and also between PW OAM and OAM on the L2 egress path is needed for end-to-end fault detection [7]. This is shown in Fig. 7.

Periodic fault detection is carried out on the MPLS PW. The use of BFD in conjunction with LSP-Ping is well suited to periodic PW fault detection. As described earlier, it is capable of supporting subsecond fault detection and has good scaling properties as far as the number of PWs that use OAM concurrently is concerned. If a fault is detected on the MPLS PW by either

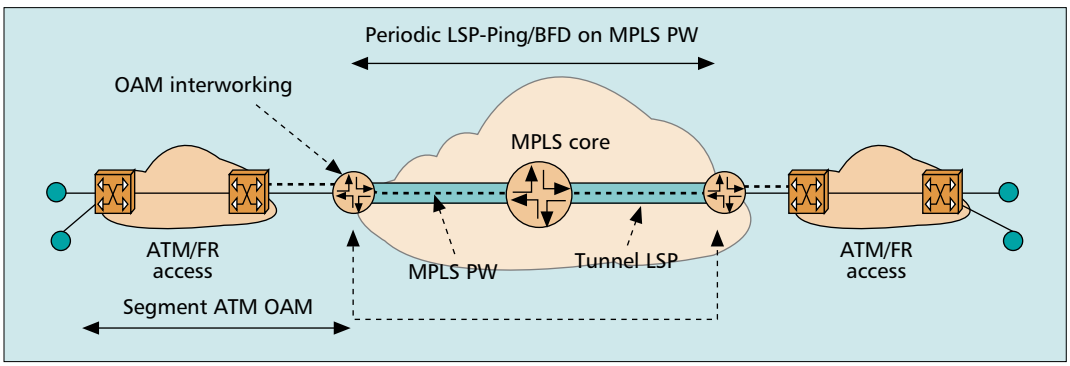

Figure 7. *ATM OAM cells terminated at the ingress PE.*

Layer 2 transport over MPLS is being increasingly deployed and is being used to build next generation multi-service networks. Such networks must support OAM mechanisms that are similar in function to those in existing Layer 2 networks.

PE, an error notification is sent to the L2 endpoint. This may be an RDI message if the L2 circuit is ATM or an LMI message for frame relay, indicating a fault. The PE detecting the fault stops transmitting BFD control packets to the remote PE. This causes the remote PE to send an error notification to the L2 endpoint to which it is attached.

Periodic fault detection is also carried out on the L2 ingress and egress paths. If a failure is detected it is conveyed to the other PE. This can be done by sending a BFD message with a special diagnostic code [9]. The other PE then issues an error notification to the L2 endpoint to which it is attached. The PE that detects the failure can also use PW signaling mechanisms such as withdrawing the PW label or a notification message to convey the error to the remote PE [4]. However, this typically involves longer latency than the use of BFD.

Interworking between segment-based OAM mechanisms is described in further detail in [7].

Conclusion

Layer 2 transport over MPLS is being increasingly deployed and used to build next-generation multiservice networks. Such networks must support OAM mechanisms that are similar in function to those in existing layer 2 networks. This article describes state-of-the-art MPLS OAM. It also described how end-to-end fault detection can be carried out for various layer 2 transport over MPLS network models. In particular, the role of MPLS PW OAM was described in detail. Router vendors are already shipping some of the OAM mechanisms described herein such as LSP-Ping. Other mechanisms such as BFD are in the process of being implemented. However, technology in this area is still evolving. Further implementation, deployment, and operational experience are needed in order to mature MPLS OAM for layer 2 transport over MPLS.

References

[1] D. Awduche *et al.*, "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209.
[2] P. Pan *et al.*, "Fast Reroute Extensions to RSVP-TE for LSP Tunnels," draft-ietf-mpls-rsvp-lsp-fastreroute-04.txt
[3] E. Rosen *et al.*, "BGP/MPLS VPNs," draft-ietf-l3vpn-rfc2547bis-01.txt
[4] L. Martini *et al.*, "Pseudowire Setup And Maintenance using LDP," draft-ietf-pwe3-control-protocol-04.txt
[5] K. Kompella *et al.*, "Layer 2 VPNs over Tunnels," draft-kompella-l2vpn-l2vpn-00.txt
[6] ITU-T Rec. Y.1711, "OAM Mechanism for MPLS Networks."
[7] T. Nadeau and M. Morrow, "Pseudo Wire (PW) OAM Message Mapping," draft-nadeau-pwe3-oam-msg-map-03.txt
[8] K. Kompella *et al.*, "Detecting MPLS Data Plane Failures," draft-ietf-mpls-lsp-ping-03.txt
[9] T. Nadeau and R. Aggarwal, "Pseudo Wire (PW) Virtual Circuit Connection Verification ((VCCV)," draft-ietf-pwe3-vccv-00.txt
[10] D. Katz and D. Ward, "Bidirectional Forwarding Detection," draft-katz-ward-bfd-01.txt, Aug. 2003.
[11] R. Aggarwal and K. Kompella, "BFD For MPLS LSPs," draft-raggarwa-mpls-bfd-00.txt

Biography

RAHUL AGGARWAL (rahul@juniper.net) is with the IP routing and MPLS engineering group at Juniper Networks and works primarily on the M and T series platforms. He contributes to the design and development of MPLS TE, layer 3 and 2 VPNs, multicast, and routing. He works closely with service providers. He also works closely with other vendors and is a significant contributor to Juniper's IETF standardization efforts. He is also a regular presenter at several MPLS related conferences. Prior to joining Juniper Networks he was with the SmartEdge IP routing group at Redback Networks. He was one of the leading architects and developers of the MPLS implementation on the SmartEdge 800 Router. Prior to joining Redback Networks, he was at Fore Systems between 1998 and 2000, where he worked on the development of OSPF, MPLS, and TE. He received a B.E. in electronics and communication from Indian Institute of Technology, Roorkee in 1996. He received his M.S. in computer science from the University of Minnesota in 1998. His graduate research was focussed on QoS and stored video delivery across resource constrained networks. His professional interests include routing and signaling development, IP and MPLS system and forwarding design, packet classification, VPNs, subscriber management, QoS, and traffic engineering.