Network Working Group
Internet Draft
Category: Informational
Created: February 25, 2009
Expires: August 25, 2009

N. Sprecher, Ed.
Nokia Siemens Networks
A. Farrel, Ed.
Old Dog Consulting

Multiprotocol Label Switching Transport Profile Survivability Framework

draft-sprecher-mpls-tp-survive-fwk-01.txt

Status of this Memo

Abstract

Network survivability is the network's ability to restore traffic following failure or attack; it plays a critical factor in the delivery of reliable services in transport networks. Guaranteed services in the form of Service Level Agreements (SLAs) require a resilient network that detects facility or node failures very rapidly, and immediately starts to restore network operations in accordance with the terms of the SLA.

The Transport Profile of Multiprotocol Label Switching (MPLS-TP) is a packet transport technology that combines the packet experience of MPLS with the operational experience of transport networks like SONET/SDH. It provides survivability mechanisms such as protection and restoration, with similar function levels to those found in established transport networks such as in SONET/SDH networks. Some of the MPLS-TP survivability mechanisms are data plane-driven and are based on MPLS-TP OAM fault management functions which are used to trigger protection switching in the absence of a control plane. Other

survivability mechanisms utilize the MPLS-TP control plane.

This document provides a framework for MPLS-TP survivability.

Table of Contents

1. Introduction

Network survivability is the network's ability to restore traffic following failure or attack; it plays a critical factor in the delivery of reliable services in transport networks. Guaranteed services in the form of Service Level Agreements (SLAs) require a resilient network that very rapidly detects facility or node failures, and immediately starts to restore network operations in accordance with the terms of the SLA.

The Transport Profile of Multiprotocol Label Switching (MPLS-TP) [RFC5317], [MPLS-TP-REQ] is a packet transport technology that combines the packet experience of MPLS with the operational experience of transport networks such as SONET/SDH. MPLS-TP is designed to be consistent with existing transport network operations and management models and provide survivability mechanisms, such as protection and restoration, with similar function levels to those found in established transport networks (such as the SONET/SDH networks which provided service providers with a high benchmark for reliability).

This document provides a framework for MPLS-TP-based survivability. It uses the recovery terminology defined in [RFC4427] which draws heavily on [G.808.1], and refers to the requirements specified in [MPLS-TP-REQ].

Various recovery schemes (for protection and restoration) and

processes have been defined and analyzed in [RFC4427] and [RFC4428]. These schemes may also be applied in MPLS-TP networks to re-establish end-to-end traffic delivery within the agreed service level and to recover from 'failed' or 'degraded' transport entities (links or nodes). Such actions are normally initiated by the detection of a defect or performance degradation, or by an external request (e.g., an operator request for manual control of protection switching).

[RFC4427] makes a distinction between protection switching and restoration mechanisms. Protection switching makes use of pre-assigned capacity between nodes, where the simplest scheme has one dedicated protection entity for each working entity, while the most complex scheme has m protection entities shared between n working entities (m:n). Protection switching may be either unidirectional or bidirectional; unidirectional meaning that each direction of a bidirectional connection is protection switched independently, while bidirectional means that both directions are switched at the same time even if the fault applies to only one direction of the connection. Restoration uses any capacity available between nodes and usually involves re-routing. The resources used for restoration may be pre-planned and recovery priority may be used as a differentiation mechanism to determine which services are recovered and which are not recovered or are sacrificed in order to achieve recovery of other services. In general, protection actions are completed within time frames of tens of milliseconds, while restoration actions are normally completed in periods ranging from hundreds of milliseconds to a maximum of a few seconds.

The recovery schemes described in [RFC4427] and evaluated in [RFC4428] assume some control plane-driven actions that are performed in the recovery context (such as the configuration of the protection entities and functions, etc.). As for other transport technologies and associated transport networks, the presence of a distributed control plane in support of MPLS-TP network operations is optional, and the absence of such a control plane does not affect the ability to operate the network and to use MPLS-TP forwarding, OAM, and survivability capabilities.

Thus, some of the MPLS-TP recovery mechanisms do not depend on a control plane and rely on MPLS-TP OAM capabilities to trigger protection switching across connections that were set up using management plane configuration. These mechanisms are data plane-driven and are based on MPLS-TP OAM fault management functions. "Fault management" in this context refers to failure detection, localization, and notification (where the term "failure" is used to represent both signal failure and signal degradation).

The principles of MPLS-TP protection switching operation are similar

to those described in [RFC4427] as the protection mechanism is based on the ability to detect certain defects in the transport entities within the protected domain. The protection switching controller does not care which monitoring method is used, as long as it can be given information about the status of the transport entities within the recovery domain (e.g., 'OK', signal failure, signal degradation, etc.).

An MPLS-TP Protection State Coordination (PSC) protocol may be used as an in-band (i.e., data plane-based) control protocol to align both ends of the protected domain.

The MPLS-TP recovery mechanisms may be applied at various nested levels throughout the MPLS-TP network, as is the case with the recovery schemes defined in [RFC4427] and [RFC4873]. A Label Switching Path (LSP) may be subject to any or all of MPLS-TP link recovery, path segment recovery, or end-to-end recovery, where:

- MPLS-TP link recovery refers to the recovery of an individual link (and hence all or a subset of the LSPs routed over the link) between two neighboring label switching routers (LSRs).

- Segment recovery refers to the recovery of an LSP segment (i.e., segment and concatenated segment in the language of [MPLS-TP-REQ]) between two nodes which are the boundary nodes of the segment

- End-to-end recovery refers to the recovery of an entire LSP from its ingress to its egress node.

Multiple recovery levels may be used concurrently by a single LSP for added resiliency.

It is a basic requirement of MPLS-TP that both directions of a bidirectional LSP should be co-routed (that is, share the same route within the network) and be fate-sharing (that is, if one direction fails, both directions should cease to operate) [MPLS-TP-REQ]. This causes a direct interaction between the recovery levels affecting the directions of an LSP such that both directions of the LSP are switched to a new MPLS-TP link, segment, or end-to-end path together.

The recovery scheme operating at the data plane level can function in a multi-domain environment; it should also protect against a failure of a boundary node in the case of inter-domain operation.

The MPLS-TP recovery schemes apply to LSPs and PWE3. This document focuses on LSPs and handles both point-to-point (P2P) and point-to-multipoint (P2MP) LSPs.

This framework introduces the architecture of the MPLS-TP recovery domain and describes the recovery schemes in MPLS-TP (based on the recovery types defined in [RFC4427]) as well as the principles of operation, recovery states, recovery triggers, and information exchanges between the different elements that sustain the reference model. The reference model is based on the MPLS-TP OAM reference model which is defined in [MPLS-TP-OAM].

The framework also describes the qualitative levels of the survivability functions that can be provided, such as dedicated recovery, shared protection, restoration, etc. The level of recovery directly affects the service level provided to the end user in the event of a network failure. There is a correlation between the level of recovery provided and the cost to the network.

This framework applies to general recovery schemes, but also for schemes that are optimized for specific topologies, such as mesh and ring, in order to handle protection switching in a cost-efficient manner.

This document takes into account the timing co-ordination of protection switches at multiple layers. This prevents races and allows the protection switching mechanism of the server layer to fix a problem before switching at the MPLS-TP layer.

This framework also specifies the functions that must be supported by MPLS-TP (e.g., PSC) and the management and/or the control plane in order to support the recovery mechanisms. MPLS-TP introduces a tool kit to enable recovery in MPLS-TP-based transport networks and to ensure that affected traffic is recovered in the event of a failure. Different recovery levels may be used concurrently by a single LSP for added resiliency.

Generally, network operators aim to provide the fastest, most stable, and the best protection mechanism available at a reasonable cost. The higher the levels of protection, the greater the number of resources consumed. It is therefore expected that network operators will offer a wide spectrum of service levels. MPLS-TP-based recovery offers the flexibility to select the recovery mechanism, choose the granularity at which traffic is protected, and also choose the specific types of traffic that are to be protected. With MPLS-TP-based recovery, it is possible to provide different levels of protection for different classes of service, based on their service requirements.

2. Terminology and References

The terminology used in this document is consistent with that defined in [RFC4427]. That RFC is, itself, consistent with [G.808.1].

However, certain protection concepts (such as ring protection) are not discussed in [RFC4427], and for those concepts, terminology in this document is drawn from [G.841].

Readers should refer to those documents for normative definitions. This document supplies brief summaries of some terms for clarity and to aid the reader, but does not re-define terms.

In particular, note the distinction and definitions made in [RFC4427] for the following three terms.

- Protection: re-establishing end-to-end traffic using pre-allocated resources.

- Restoration: re-establishing end-to-end traffic using resources allocated at the time of need. Sometimes referred to as "repair".

- Recovery: a generic term covering both Protection and Restoration.

Important background information can be found in [RFC3386], [RFC3469], [RFC4426], [RFC4427], and [RFC4428].

3. Requirements for Survivability

MPLS-TP requirements are presented in [MPLS-TP-REQ]. Survivability is presented as a critical factor in the delivery of reliable services, and the requirements for survivability are set out using the recovery terminology defined in [RFC4427].

These requirements are summarized below. This section may be updated if changes are made to [MPLS-TP-REQ], and that document should be regarded as normative for the definition of all MPLS-TP requirements including those for survivability.

General:

- Must support protection and restoration.
- Must be applicable at various nested levels, including link, LSP segment and LSP end-to-end path, PW segment and end-to-end PW.
- Should be equally applicable to LSPs and pseudowires.
- Must provide appropriate recovery times.
- Should support the configuration of the recovery objectives (such as BW and QOS) per transport path.
- Must scale when many services are affected by a single fault.
- Must support management plane control.
- Must support control plane control.
- Must be applicable for any topology.
- Must provide coordination between protection mechanisms at

different layers.
- Must provide mechanism to prevent recovery operations thrashing.
- Must support Physical layer fault indication as a trigger to the recovery operation.
- Must support OAM based triggers to the recovery operation.
- Must support administrative commands as triggers to the recovery operation (e.g. force switch, etc).
- Must support a mechanism to allow the distinction of recovery actions that are initiated by administrative commands from those that are initiated by other means.
- Should support control plane triggers when a control plane is available.
- Must support the management plane configuration of timers used for the recovery operation.
- Must support the management plane configuration of the elements of controls (triggers for recovery).
- Must support the control plane configuration of the recovery entities and functions (if the control plane is present).
- Must support the control plane signaling of an administrative commands if the control plane is present).
- Must support the control plane signaling of the protection state, in order to synch the protection state between the edges of the protection domain.

Restoration:

- Must support soft re-routing (Make-before-break).
- Must support pre-planning of restoration resources.
- May support computation of restoration resources after failure.
- May support shared mesh restoration.
- May support hard LSP restoration (break-before-make).
- Must support restoration priority (under operator configuration)
- Must support preemption priority during restoration (under operator configuration).

Protection:

- Must support bidirectional 1+1 protection switching (which should be the default behavior) and 1+1 unidirectional protection switching for P2P paths.
- Must support bidirectional 1:n protection switching (which should be the default behavior) for P2P paths.
- Must support 1:1 and 1+1 unidirectional protection switching for P2MP.
- Must support protection ration of 100%.
- Should support 1:n shared mesh protection. (*** contradict the Must support above) .
- Must support shared bandwidth.

- Must support the definition of shared protection groups (to allow coordination of protection actions).
- Must support sharing of protection resources.
- Must support revertive (which is the default behavior) and non-revertive behavior.
- Must support the management plane configuration of the protection path and the protection group.
- Must provide clear indication of the protection state of the transport path.
- May provide different mechanisms optimized for specific topologies (such as ring topologies). Such mechanisms must interoperate with the mechanisms that are defined for the arbitrary topology). For the specific requirements for ring topologies, see Section 4.5.2 on rings.

4. Functional Architecture

This section presents an overview of the elements of the functional architecture for survivability within an MPLS-TP network. The intention is to break the components out as separate items so that it can be seen how they may be combined to provide different levels of recovery to meet the requirements set out in the previous section.

4.1. Elements of Control

Survivability is achieved through specific actions taken to repair network resources or to redirect traffic onto paths that avoid failures in the network. Those actions may be triggered automatically by the network devices (detecting a network failure), may be enhanced by in-band (i.e. data-plane based) OAM fault management or performance monitoring, in-band or out-of-band control plane signaling, or may be under direct the control of an operator.

These different options are explored in the next sections.

4.1.1. Manual Control

Of course, the survivability behavior of the network as a whole, and the reaction of each LSP when a fault is reported, may be under operator control. That is, the operator may establish network-wide or local policies that determine what actions will be taken when different failures are reported that affect different LSPs. At the same time, when a service request is made to cause the establishment of one or more LSPs in the network, the operator (or requesting application) may express a required or desired level of service, and this will be mapped to particular survivability actions taken before and during LSP setup, after the failure of network resources, and upon recovery of those resources.

The operator can also be given manual control of survivability actions and events. For example, the operator may force a switchover from a working path to a recovery path (for network optimization purposes with minimal disturbance of services, like when modifying protected or unprotected services, when replacing network elements, etc.), inhibit survivability actions, enable or disable survivability function, or induce the simulation of a network fault. In some circumstances, a fault may be reported to the operator and the operator may then select and initiate the appropriate recovery action.

4.1.2. Failure-Triggered Actions

Survivability actions may be directly triggered by network failures. That is, the device that detects the failure (for example, Loss of Light on an optical interface, or failure to receive an OAM continuity message) may immediately perform a survivability action. Note that the term "failure" is used to represent both signal failure and signal degradation.

This behavior can be subject to management plane or control plane control, but does not require any messages exchanges in any of the management plane, control plane, or data plane to trigger the recovery action - it is directly triggered by data plane stimuli. Note, however, that coordination of recovery actions between the edges of the recovery domain may require message exchanges for some qualitative levels of recovery.

4.1.3. OAM Signaling

OAM signaling refers to message exchanges that are in-band or closely coupled to the data channel. Such messages may be used to detect and isolate faults, but in this context we are concerned with the use of these messages to control or trigger survivability actions.

OAM signaling may also be used to coordinate recovery actions within the network.

4.1.4. Control Plane Signaling

Control plane signaling is responsible for setup, maintenance, and teardown of LSPs that are not under management plane control. The control plane can also be used to detect, isolate, and communicate network failures pertaining to peer relationships (neighbor-to-neighbor, or end-to-end). Thus, control plane signaling can initiate and coordinate survivability actions.

The control plane can also be used to distribute topology and resource-availability information. In this way, "graceful shutdown" of resources may be effected by withdrawing them, and this can be used as a stimulus to survivability action in a similar way to the reporting or discovery of a fault as described in the previous sections.

4.2. Elements of Recovery

This section describes the elements of recovery. These are the quantitative aspects of recovery; that is the pieces of the network for which recovery can be provided.

Note that the terminology in this section is consistent with [RFC4427]. Where the terms differ from those in [MPLS-TP-REQ] a mapping is provided.

4.2.1. Span Recovery

A span is a single hop between neighboring MPLS-TP LSRs in the same network layer. A span is sometimes referred to as a link although this may cause some confusion between the concept of a data link and a traffic engineering (TE) link. LSPs traverse TE links between neighboring label switching routers (LSRs) in the MPLS-TP network, however, a TE link may be provided by:

- a single data link

- a series of data links in a lower layer established as an LSP and presented to the upper layer as a single TE link

- a set of parallel data links in the same layer presented either as a bundle of TE links or a collection of data links that, together, provide data link layer protection scheme.

Thus, span recovery may be provided by:

- moving the TE link to be supported by a different data link between the same pair of neighbors

- re-routing the LSP in the lower layer.

Moving the protected LSP to another TE link between the same pair of neighbors is known as segment recovery and is described in Section 4.2.2.

[MPLS-TP-REQ] refers to a span as a "link".

4.2.2. Segment Recovery

An LSP segment is one or more hops on the path of the LSP. In some MPLS-TP documents LSP segment is referred as LSP Tandem Connection (Note that recovery of pseudowire segments is discussed in Section 6.)

Segment recovery involves redirecting traffic from one end of a segment of an LSP on an alternate path to the other end of the segment. This redirection may be on a pre-established LSP segment, through re-routing of the protected segment, or by tunneling the protected LSP on a "bypass" LSP.

Note that protecting an LSP against the failure of a node requires the use of segment recovery, while a link could be protected using span or segment recovery.

[MPLS-TP-REQ] defines two terms. A "segment" is a single hop on the path of an LSP, and a "concatenated segment" is more than one hop on the path of an LSP. In the context of this document, a segment covers both of these concepts.

4.2.3. End-to-End Recovery

End-to-end recovery is a special case of segment recovery where the protected LSP segment is the whole of the LSP. End-to-end recovery may be provided as link-diverse or node-diverse recovery where the recovery path shares no links or no nodes with the recovery path. Note that node-diverse paths are necessarily link-diverse, and that full, end-to-end node-diversity is required to guarantee recovery.

4.3. Levels of Recovery

This section describes the qualitative levels of survivability function that can be provided. The level of recovery offered has a direct effect on the service level provided to the end-user in the event of a network fault. This will be observed as the amount of data lost when a network fault occurs, and the length of time to recovery connectivity.

In general there is a correlation between the service level (i.e., the rapidity of recovery and reduction of data loss) and the cost to the network; better service levels require pre-allocation of resources to the recovery paths, and those resources cannot be used for other purposes if high quality recovery is required.

Sections 6 and 7 of [RFC4427] provide a full break down of protection and recovery schemes. This section summarizes the qualitative levels

available.

4.3.1. Dedicated Protection

In dedicated protection, the resources for the recovery LSP are pre-assigned for use only by the protected service. This will clearly be the case in 1+1 protection, and may also be the case in 1:1 protection where extra traffic (see Section 4.3.3) is not supported.

Note that in the bypass tunnel recovery mechanism (see Section 4.4.3) resources may also be dedicated to protecting a specific service. In some cases (one-for-one protection) the whole of the bypass tunnel may be dedicated to provide recovery for a specific LSP, but in other cases (such as facility backup) a subset of the resources of the bypass tunnel may be pre-assigned for use to recover a specific service. However, as described in Section 4.4.3, the bypass tunnel approach can also be used for shared protection (Section 4.3.2), to carry extra traffic (Section 4.3.3), or without reserving resources to achieve best-effort recovery.

4.3.2. Shared Protection

In shared protection, the resources for the recovery LSPs of several services are shared. These may be shared as 1:n or m:n, and may be shared on individual links, on LSP segments, or on end-to-end LSPs.

Where a bypass tunnel is used (Section 4.4.3) the tunnel might not have sufficient resources to simultaneously protect all of the LSPs to which it offers protection so that if they were all affected by network failures at the same time, they would not all be recovered.

Shared protection is a trade-off between expensive network resources being dedicated to protection that is not required most of the time, and the risk of unrecoverable services in the event of multiple network failures. There is also a trade-off between rapid recovery (that can be achieved with dedicated protection, but which is delayed by message exchanges in the management, control, or data planes for shared protection) and the reduction of network cost by sharing protection resources. These trade-offs may be somewhat mitigated by using m:n for some value of m <> 1, and by establishing new protection paths as each available protection path is put into use.

4.3.3. Extra Traffic

A way to utilize network resources that would otherwise be idle awaiting use to protect services, is to use them to carry other traffic. Obviously, this is not practical in dedicated protection (Section 4.3.1), but is practical in shared protection (Section

4.3.2) and bypass tunnel protection (Section 4.4.3).

When a network resource that is carrying extra traffic is required for protection, the extra traffic is disrupted - essentially it is pre-empted by the recovery LSP. This may require some additional messages exchanges in the management, control, or data planes, with the consequence that recovery may be delayed somewhat. This provides an obvious trade-off against the cost reduction (or rather, revenue increase) achieved by carrying extra traffic.

4.3.4. Restoration and Repair

If resources are not pre-assigned for use by the recovery LSP, the recovery LSP must be established "on demand" when the network failure is detected and reported, or upon instruction from the management plane.

Restoration represents the most cost-effective use of network resources as no resources are tied up for specific protection usage. However, restoration requires computation of a new path and activation of a new LSP (through the management or control plane). These steps can take much more time than is required for recovery using protection techniques.

Furthermore, there is no guarantee that restoration will be able to recover the service. It may be that all suitable network resources are already in use for other LSPs so that no new path can be found. This problem can be partially mitigated by the use of LSP setup priorities so that recovery LSPs can pre-empt other low priority LSPs.

Additionally, when a network failure occurs, multiple LSPs may be disrupted by the same event. These LSPs may have been established by different Network Management Stations (NMSs) or signaled by different head-end LSRs, and this means that multiple points in the network will be trying to compute and establish recovery LSPs at the same time. This can lead to contention within the network meaning that some recovery LSPs must be retried resulting in very slow recovery times for some services.

Both hard or soft LSP restoration may be supported. In hard LSP restoration, the resources of the LSP are released before the full establishment of the recovery LSP (i.e., break-before-make). In soft LSP restoration, the resources of the LSP are released after the full establishment of an alternate LSP (i.e., make-before-break).

Note that the restoration resources may be pre-calculated and even pre-signaled before the restoration action starts, but not pre-

allocated. This is known as pre-planned LSP restoration. The complete establishment/activation of the restoration LSP occurs only when the restoration action starts. The pre-planning may happen periodically to have the most accurate information about the available resources in the network.

4.3.5. Reversion

When a service has been recovered so that traffic is flowing on the recovery LSP, the faulted network resource may be repaired. The choice must be made about whether to redirect the traffic back on to the original working LSP, or to leave it where it is on the recovery LSP. These behaviors are known as "revertive" and "non-revertive", respectively.

In "revertive" mode, care should be taken to prevent frequent operation of the recovery operation due to an intermittent defect. Therefore, when the failure condition of a recovery element has been handled, a fixed period of time should elapse before normal data traffic is redirected back onto the original working entity.

4.4. Mechanisms for Recovery

The purpose of this section is to describe in general (MPLS-TP non-specific) terms the mechanisms that can be used to provide protection.

4.4.1. Link-Level Protection

Link-level protection refers to the paradigm whereby protection is provided in a lower network layer.

Link-level protection offers the following levels of protections:

- Full protection, where a dedicated protection entity (e.g. a link or span) is pre-established to protect a working entity. When the working link fails, the protected traffic is switched onto the protecting entity. In this scenario, all LSPs carried over the entity are recovered (in one protection operation) when there is a failure condition at the link-level. This is referred to in [RFC4427] as 'bulk recovery'.

- Partial protection, where only a subset of the LSPs carried over a given entity is recovered when there is a failure condition. The decision as to which LSPs will be protected and which will not depends on local policy.

When there is no failure on the working link, the protection entity

may transport extra traffic which may be preempted when protection switching occurs.

As with recovery in layered networks, the protection mechanism at the link-level needs to co-ordinate the timing for switchover, in order to avoid race conditions and to enable switchover to be performed at the link level before the upper level.

Note that link-level protection does not protect the nodes at each end of the entity (e.g. a link or span) that is protected. End-to-end or segment LSP protection should be used to protect against a failure of the edge node.

4.4.2. Alternate Paths and Segments

The alternate paths and segments refer to the paradigm whereby the protection is performed at the same network layer of the protected LSP/segment-LSP.

Different levels of protection may be provided:

- Dedicated protection, where a dedicated entity (e.g. LSP, segment LSP) is fully pre-established to protect a working entity (e.g., LSP, segment LSP). When there is a failure condition on the working entity, the normal traffic is switched over into the protection entity.

  Dedicated protection may be accomplished by the 1:1 or 1+1 protection schemes. When the failure condition is eliminated, the traffic may revert to the working entity. This is subject to local configuration.

- Shared protection, where one or more protection entities are pre-established to protect against a failure of one or more working entities (1:n or m:n).

  When the failure condition on the working entity is eliminated, the traffic should revert back to the working entity.

4.4.3. Bypass Tunnels

A bypass tunnels is a transport entity (LSP) that is pre-provisioned in order to protect against a failure condition along a network segment, which may affect one or more LSPs that transmit over the network segment.

When there is a failure condition in the network segment, one or more of the protected LSPs are switched over at the ingress point of the

network segment and transmitted over the bypass tunnel. The natural way to realize this is using label stacking. Label mapping may be an option as well.

Different levels of protection may be provided:

- Dedicated protection, where the bypass tunnel has resource reservations sufficient to provide protection for all protected LSPs without service degradation.

- Shared protection, where the bypass tunnel has resources to protect some of the protected LSPs, but not all of them simultaneously.

4.5. Protection in Different Topologies

As described in the requirements listed in Section 2.8.5 and detailed in [MPLS-TP-REQ], the recovery techniques used may be optimized for different network topologies if the performance of those optimized mechanisms is significantly better than the performance of the generic ones in the same topology.

It is required that such mechanisms interoperate with the mechanisms defined for arbitrary topologies to allow end-to-end protection and to allow consistent protection techniques to be used across the whole network.

This section describes two different topologies and explains how recovery may be markedly different in those different scenarios. It also introduces the concept of a recovery domain and shows how end-to-end survivability may be achieved through a concatenation of recovery domains each providing some level of recovery in part of the network.

4.5.1. Mesh Networks

Linear protection provides a fast and simple protection switching mechanism and fits best in mesh networks. It can protect against a failure that may happen on an entity (element of recovery that may constitute a span, LSP segment, PW segment, end-to-end LSP or end-to-end PW).

Linear protection operates in the context of a Protection Domain which is composed of the following architectural elements:

- A set of end points which reside at the boundary of the Protection Domain. In this simple case of a 1:n or 1+1 P2P entity, exactly two endpoint reside at the boundary of the Protection Domain. In each transmission direction one of the end points is referred to as a

source and the other one is referred to as a sink.

In the case of unidirectional P2MP, three or more endpoints reside at the boundary of the Protection Domain. One of the endpoints is referred to as source/root and the other ones are referred to as sinks/leaves.

- A Protection Group which consists of a Working (primary) entity and one or more Protection (backup) entities. In order to guarantee complete protection, a dedicated Protection entity should be pre-provisioned to protect against a failure of the Working entity. Also the Working and the Protection entities should be disjoint entities, i.e., the physical routes of the Working and the Protection entities should have complete physical diversity. Note that resources of the Protection entity may be degraded from the Working entity. In such a case, the Protection entity may not have sufficient resources to protect the traffic of the Working entity.

  As mentioned above in section 4.3.2, the resources of the Protection entity may be shared as 1:n. In such a case, the Protection entity might not have sufficient resources to simultaneously protect all of the Working entities that may be affected by fault conditions at the same time.

Protection switching occurs at the protection controllers which reside at the edges of the Protected Domain. The working and protection entities reside between these endpoints.

[MPLS-TP-REQ] requires that both 1:n linear protection scheme and 1+1 protection schemes are supported. The 1:n protection switching, bidirectional protection switching should be supported. In 1+1 linear protection switching both unidirectional and bidirectional protection switching should be supported.

In bidirectional protection switching, in the event of failure, the recovery actions are taken in both directions (even when the fault is unidirectional). This requires the synchronization of the recovery state between the endpoints of the protection domain.

In unidirectional protection switching, the recovery actions are taken only in the affected direction.

1:1 linear protection:

- In normal conditions the data traffic is transmitted over the working entity. Normal conditions are defined when there is no failure or degradation on the 'working' entity and there is no administrative configuration or requests that cause traffic to

transmit over the 'protection' entity. Upon a fault condition (failure or degradation) or a specific administrative request, the traffic is switched over to the 'protection' entity.

Note that in the non-revertive behavior (see section 4.3.5), data traffic can be transmitted over the Protection entity also in normal conditions. This can happen after a failure condition on the Working entity (which caused a recovery action) is eliminated.

- In each transmission direction, the source of the protection domain bridges the traffic into the appropriate entity and the sink of the protected domain selects the traffic from the appropriate entity. The source and the sink need to be coordinated to ensure that the bridging and the selection are done to and from the same entity. For that sake a signaling coordination protocol is needed.

- In bidirectional protection switching, both ends of the protection domain switch to the 'protection' entity (even when the failure is unidirectional). This requires a protocol to synchronize the protection state between the two end points of the Protection Domain.

- When there is no failure, the resources of the 'idle' entity may be used for less priority traffic. When protection switching is performed, the less priority traffic may be pre-empted by the protected traffic.

1+1 linear protection:

- The data traffic is copied at fed at the source to both the 'working' and the protection' entities. The traffic on the 'working' and the 'protection' entities is transmitted simultaneously to the sink of the protected domain, where a selection between the 'working' and 'protection' entities is made (based on some predetermined criteria).

  In 1+1 unidirectional protection switching there is no need to coordinate the recovery state between the protection controllers at both ends of the protection domain. In 1+1 bidirectional protection switching, there is a need for a protocol to coordinate the protection state between the edges of the Protection Domain.

  In both protection schemes when the failure condition is eliminated, operation, when the failure condition is eliminated, the protected traffic may revert back into the Working entity. To verify that the network has stabilized, and to avoid frequent switching in case of intermittent failures, traffic is not switched back to the Working entity before the Wait-to-Restore (WTR) timer

has expired.

Revertive/non-revertive operations are provided as network operator options.

The protection switching may be performed when:

- A fault condition ('failed' or 'degraded') is declared on the active entity and is not declared on the standby entity. OAM CC&V (Continuity and Connectivity Verification) monitoring of both Working and Protection entities may be used to enable the fast detection of a fault condition. For protection switching, it is common to run a CC&V every 3.33ms. In the absence of three consecutive CC messages, a 'failed' condition is declared. In order to monitor the Working and the Protection entities, an OAM Maintenance Entity should be defined for each of the entities. OAM information should be provided as input to the protection switching controllers.

  Input from OAM performance monitoring indicating degradation in the Working entity may also be used as a trigger for protection switching. In the case of degradation, switching to the Protection entity is needed only if the Protection entity can guarantee better conditions.

  Note that in bidirectional protection switching, an attempt is made to coordinate the protection switching state between both end points of the Protection Domain when a unidirectional failure is detected or when an external administrative requests is received. A PSC (Protection State Coordination) protocol may be used for this purpose. This protocol is also used to detect mismatches between the provisioned protection switching configuration and the two ends of a Protection Domain.

  Note that in order to achieve 50ms protection switching it is recommended to use inband signaling protocol to coordinate the protection states.

- An indication is received from a lower layer server that there is a network failure.

- An external operator command is received (e.g., 'Forced Switch', 'Manual Switch'). For details see Section 5.1.2.

- A request to switch over is received from the far end (relevant in case of bidirectional 1:1 protection switching only).

Linear protection provides a clear indication of the protection

status.

1:n linear protection:

In 1:n linear protection, one Protection entity is used to protect n Working entities. The Protection entity might not have sufficient resources to simultaneously protect all of the Working entities that may be affected by fault conditions at the same time.

Revertive behavior is recommended when 1:n is supported.

P2MP linear protection:

Linear protection may apply to protect P2MP path using 1+1 protection architecture. The source/root LSR bridges the user traffic to both the Working and Protected entities. Each sink/leaf LSR selects the traffic from one entity based on some predetermined criteria. Note that when there is a fault condition on one of the branches of the P2MP path, some leaf LSRs may select the Working entity, while other leaf LSRs may select traffic from the Protection entity.

In a 1:1 P2MP protection scheme, the source/root LSR needs to identify the existence of a fault condition on any of the branches of the network. This requires the sink/leaf LSRs to notify the source/root LSR of any fault condition. This required also a return path from the sinks/leaves to the source/root LSR.

When protection switching is triggered, the source/root LSR selects the recovery transport path to transfer the traffic.

4.5.2. Ring Networks

Several Service Providers have expresses a high level of interest in operating MPLS-TP in ring topologies and require a high level of survivability function in these topologies.

Different criteria for optimization are considered in ring topologies, such as:

1. Simplification of the operation of the Ring in terms of the number of OAM Maintenance Entities that are needed to trigger the recovery actions, the number of elements of recovery, the number of management plane transactions during maintenance operations, etc.

2. Optimization of resource consumption around the ring, like the number of labels needed for the protection paths that cross the

network, the total bandwidth needed in the ring to ensure the protection of the paths, etc.

[MPLS-TP-REQ] introduces a list of requirements on ring protection that cover the recovery mechanisms need to protect traffic in a single ring and traffic that traverses more than one ring. Note that configuration and the operation of the recovery mechanisms in a ring must scale well with the number of transport paths, the number of nodes, and the number of ring interconnects.

The requirements for ring protection are fully compatible with the generic requirements for recovery.

The architecture and the mechanisms for ring protection are specified in separate documents. These mechanisms need to be evaluated against the requirements specified in [MPLS-TP-REQ]. The principles for the development of the mechanisms should be:

1. Reuse existing procedures and mechanisms for recovery in ring topologies as along as their performance is as good as new potential mechanisms.

2. Ensure complete interoperability with the mechanisms defined for arbitrary topologies to allow end-to-end protection.

4.5.3. Protection and Restoration Domains

Protection and restoration are performed in the context of a recovery domain. A recovery domain is defined between two recovery reference points which are located at the edges of the recovery domain and are responsible for performing recovery for a 'working' entity (which may be one of the elements of recovery defined above) when an appropriate trigger is received. These reference points function as recovery controllers.

As described in section 4.2 above, the recovery element may constitute a spam, a tandem connection (i.e. either an LSP segment or a PW segment), an end-to-end LSP, or an end-to-end PW.

The method used to monitor the health of the recovery element is unimportant, provided that the recovery controllers receive information on its condition. The condition of the recovery element may be OK, 'failed', or degraded.

When the recovery operation is launched by an OAM trigger, the recovery domain is equivalent to the OAM maintenance entity which is defined in [MPLS-TP-OAM], and the recovery reference points are defined at the same location as the OAM MEPs.

4.6. Recovery in Layered Networks

In multi-layer or multi-region networking, recovery may be performed at multiple layers or across cascaded recovery domains.

The MPLS-TP recovery mechanism must ensure that the timing of recovery is coordinated in order to avoid races, and to allow either the recovery mechanism of the server layer to fix the problem before recovery takes place at the MPLS-TP layer, or to allow an upstream recovery domain to perform recovery before a downstream domain. In inter-connected rings, for example, it may be preferable to allow the upstream ring to perform recovery before the downstream ring, in order to ensure that recovery takes place in the ring in which the failure occurred.

A hold-off timer is required to coordinate the timing of recovery at multiple layers or across cascaded recovery domains. Setting this configurable timer involves a trade-off between rapid recovery and the creation of a race condition where multiple layers respond to the same fault, potentially allocating resources in an inefficient manner. Thus, the detection of a failure condition in the MPLS-TP layer should not immediately trigger the recovery process if the hold-off timer is set to a value other than zero. The hold-off timer should be started and, on expiry, the recovery element should be checked to determine whether the failure condition still exists. If it does exist, the defect triggers the recovery operation.

In other configurations, where the lower layer does not have a restoration capability, or where it is not expected to provide protection, the lower layer needs to trigger the higher layer to immediately perform recovery.

Reference should be made to [RFC3386] that presents the near-term and practical requirements for network survivability and hierarchy in current service provider environments.

4.6.1. Inherited Link-Level Protection

TBD

4.6.2. Shared Risk Groups

TBD

4.6.3. Fault Correlation

TBD

5. Mechanisms for Providing Protection in MPLS-TP

This section describes the existing mechanisms available to provide protection within MPLS-TP networks and highlights areas where new work is required. It is expected that, as new protocol extensions and techniques are developed, this section will be updated to convert the statements of required work into references to those protocol extensions and techniques.

5.1. Management Plane

As described above, a fundamental requirement of MPLS-TP is that recovery mechanisms should be capable of functioning in the absence of a control plane. Recovery may be triggered by MPLS-TP OAM fault management functions or by external requests (e.g. an operator request for manual control of protection switching).

The management plane may be used to configure the recovery domain by setting the reference points (recovery controllers), the 'working' and 'protection' entities, and the recovery type (e.g. 1:1 bidirectional linear protection, ring protection, etc.). Additional parameters associated with the recovery process (such as a hold-off timer, revertive/non-revertive operation, etc.) may also be configured.

In addition, the management plane may initiate manual control of the protection switching function. Either the fault condition or the operator request should be prioritized.

Since provisioning the recovery domain involves the selection of a number of options, mismatches may occur at the different reference points. The MPLS-TP OAM Automatic Protection Switching (APS) protocol may be used as an in-band (i.e., data plane-based) control protocol to align both ends of the protected domain.

It should also be possible for the management plane to monitor the recovery status.

5.1.1. Configuration of Protection Operation

In order to implement the protection switching mechanism, the following entities and information should be provisioned:

- The protection controllers (reference points)

- The protection group consisting of a 'working' entity (which may be one of the recovery elements defined above) and a 'protection' entity. To guarantee protection, the paths of the 'working' and the

'protection' entities should have complete physical diversity.

- The protection type that should be applied

- Revertive/non-revertive behavior

5.1.2. External Manual Commands

The following external, manual commands may be applied to a protection group; they are listed in descending order of priority:

- Blocked protection action - a manual command to prevent data traffic from switching to the 'protection' entity. This command actually disables the protection group.

- Force protection action - a manual command that forces a switch of normal data traffic to the 'protection' entity.

- Manual protection action - a manual command that forces a switch of data traffic to the 'protection' entity when there is no failure in the 'working' or the 'protection' entity

5.2. Fault Detection

TBD

5.3. Fault Isolation

TBD

5.4. OAM Signaling

TBD

5.4.1. Fault Detection

TBD

5.4.2. Fault Isolation

TBD

5.4.3. Fault Reporting

TBD

5.4.4. Coordination of Recovery Actions

TBD

5.5. Control Plane

The GMPLS control plane has been proposed as the control plane for MPLS-TP [RFC5317]. Since GMPLS was designed for use in transport networks, and has been implemented and deployed in many networks, it is not surprising that it contains many features to support a high level of survivability function.

The signaling elements of the GMPLS control plane utilize extensions to the Resource Reservation Protocol (RSVP) as documented in a series of documents commencing with [RFC3471] and [RFC3473], but based on [RFC3209] and [RFC2205]. The architecture for GMPLS is provided in [RFC3945], and [RFC4426] gives a functional description of the protocol extensions needed to support GMPLS-based recovery (i.e., protection and restoration).

A further control plane protocol called the Link Management Protocol (LMP) [RFC4204] is part of the GMPLS protocol family and can be used to coordinate fault isolation and reporting.

Clearly, the control plane techniques described here only apply where an MPLS-TP control plane is deployed and operated. All mandatory survivability features must be enabled even in the absence of the control plane, but where the control plane is present it may provide alternative mechanisms that may be desirable by virtue of their ease of automation or richer feature-set.

5.5.1. Fault Detection

The control plane is not able to detect data plane faults. However, it does provide mechanisms to detect control plane faults and these can be can be used to deduce data plane faults where it is known that the control and data planes are fate sharing. Although [MPLS-TP-REQ] specifies that MPLS-TP must support an out-of-band control channel, it does not insist that this is used exclusively. That means that there may be deployments where an in-band (or at least in-fiber) control channel is used. In this case, the failure of the control channel can be used to infer a failure of the data channel or at least to trigger an investigation of the health of the data channel.

Both RSVP and LMP provide a control channel "keep-alive" mechanism (called the Hello message in both cases). Failure to receive a message in the configured/negotiated time period indicates a control plane failure. GMPLS routing protocols ([RFC4203] and [RFC5307]) also

include keepalive mechanisms designed to detect routing adjacency failures and, although these keep-alive mechanisms tend to operate at a relatively low frequency (order of seconds) it is still possible that the first indication of a control plane fault will be through the routing protocol.

Note, however, care must be taken that the failure is not caused by a problem with the control plane software or processor component at the far end of a link.

Because of the various issues involved, it is not recommended that the control plane be relied upon as the primary mechanism for fault detection in an MPLS-TP network.

5.5.2. Testing for Faults

The control plane may be used to initiate and coordinate testing of links, LSP segments, or whole LSPs. This is important in some technologies where it is necessary to halt data transmission while testing, but may also be useful where testing needs to be specifically enabled or configured.

LMP provides a control plane mechanism to test the continuity and connectivity (and naming) of individual links. A single management operation is required to initiate the test at one end of the link, and LMP handles the coordination with the other end of the link. The test mechanism for an MPLS packet link relies on the LMP Test message inserted into the data stream at one end of the link and extracted at the other end of the link. This mechanism need not be disruptive to data flowing on the link.

Note that a link in LMP may in fact be an LSP tunnel used to form a link in the MPLS-TP network.

GMPLS signaling (RSVP) offers two mechanisms that may also assist with testing for faults. First, [RFC3473] defines the Admin_Status object that allows an LSP to be set into "testing mode". The interpretation of this mode is implementation specific and could be documented more precisely for MPLS-TP. The mode sets the whole LSP into a state where it can be tested; this need not be disruptive to data traffic.

The second mechanism provided by GMPLS to support testing is provided in [GMPLS-OAM]. This protocol extension supports the configuration (including enabling and disabling) of OAM mechanisms for a specific LSP.

5.5.3. Fault Isolation

Fault isolation is the process of determining exactly where a fault has occurred. It is often the case the fault detection only takes place at key points in the network (such as at LSP end points, or MEPs). This means that the fault may be located anywhere within a segment of the LSP concerned.

If segment or end-to-end protection are in use, this level of information is often sufficient to repair the LSP. However, if a finer granularity of information is needed (either to implement optimal recovery actions or to diagnose the fault), it is necessary to isolate the fault more closely.

LMP provides a cascaded test-and-propagate mechanism specifically designed for this purpose.

5.5.4. Fault Reporting

GMPLS signaling uses the Notify message to report faults. The Notify message can apply to a single LSP or can carry fault information for a set of LSPs to improve the scalability of fault notification.

Since the Notify message is targeted at a specific node it can be delivered rapidly without requiring hop-by-hop processing. It can be targeted at LSP end-points, or at segment end-points (such as MEPs). The target points for Notify messages can be manually configured within the network or may be signaled as the LSP is set up. This allows the process to be made consistent with segment protection and the concept of Maintenance Entities.

GMPLS signaling also provides a slower, hop-by-hop mechanism for reporting individual LSP faults on a hop-by-hop basis using the PathErr and ResvErr messages.

[RFC4783] provides a mechanism to coordinate alarms and other event or fault information through GMPLS signaling. This mechanism is useful to understand the status of the resources used by an LSP and to help understand why an LSP is not functioning, but it is not intended to replace other fault reporting mechanisms.

GMPLS routing protocols ([RFC4203] and [RFC5307]) are used to advertise link availability and capabilities within a GMPLS-enabled network. Thus, the routing protocols can also provide indirect information about network faults. That is, the protocol may stop advertising or withdraw the advertisement for a failed link, or may advertise that the link is about to be shut down gracefully. This mechanisms is, however, not normally considered to be fast enough to

be used as a trigger for protection switching.

5.5.5. Coordination of Recovery Actions

Fault coordination is an important feature for certain protection mechanisms (such as bidirectional 1:1 protection). The use of the GMPLS Notify message for this purpose is described in [RFC4426], however, specific message field values remain to be defined or this operation.

A further piece of work in [GMPLS-REV] allows control and configuration of reversion behavior for end-to-end and segment protection.

5.5.6. Establishment of Protection and Restoration LSPs

It should not be forgotten that protection and recovery depend on the establishment of suitable LSPs. The management plane may be used to set up these LSPs, but the control plane may be used if it is present.

Several protocol extensions exist to make this process more simple:

- [RFC4872] provides features in support of end-to-end protection switching.

- [RFC4873] describes how to establish a single, segment protected LSP.

- [RFC4874] allows one LSP to be signaled with a request that its path excludes specified resources (links, nodes, SRLGs). This allows a disjoint protection path to be requested, or a recovery path to be set up avoiding failed resources.

Lastly, it should be noted that [RFC5298] provides an overview of the GMPLS techniques available to achieve protection in multi-domain environments.

6. Pseudowire Protection Considerations

The main application for the MPLS-TP network is currently identified as the pseudowire. Pseudowires provide end-to-end connectivity over the MPLS-TP network and may be comprised of a single pseudowire segment, or multiple segments "stitched" together to provide end-to-end connectivity.

The pseudowire service may, itself, require a level of protection as part of its SLA. This protection could be provided by the MPLS-TP

LSPs that support the pseudowire, or could be a feature of the pseudowire layer itself.

6.1. Utilizing Underlying MPLS-TP Protection

TBD

6.2. Protection in the Pseudowire Layer

TBD

7. Manageability Considerations

TBD

8. Security Considerations

TBD

9. IANA Considerations

This informational document makes no requests for IANA action.

10. Acknowledgments

TBD

11. References

11.1. Normative References

[RFC2205] Braden, R. (Ed.), Zhang, L., Berson, S., Herzog, S. and S. Jamin, "Resource ReserVation Protocol -- Version 1 Functional Specification", RFC 2205, September 1997.

[RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V. and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.

[RFC3471] Berger, L., Editor, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Functional Description", RFC 3471, January 2003.

[RFC3473] Berger, L. (Ed.), "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions", RFC 3473, January 2003.

[RFC3945] Mannie, E., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.

[RFC4203] Kompella, K, and Rekhter, Y., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4203, October 2005.

[RFC4204] Lang, J., Ed., "The Link Management Protocol (LMP)", RFC 4204, September 2005.

[RFC4427] Mannie, E., and Papadimitriou, D., "Recovery (Protection and Restoration) Terminology for Generalized Multi-Protocol Label Switching (GMPLS)", RFC 4427, March 2006.

[RFC4428] Papadimitriou D. and E.Mannie, Editors, "Analysis of Generalized Multi-Protocol Label Switching (GMPLS)-based Recovery Mechanisms (including Protection and Restoration)", RFC 4428, March 2006.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and Farrel, A., " GMPLS Segment Recovery", RFC 4873, May 2007.

[RFC5307] Kompella, K, and Rekhter, Y., "IS-IS Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)", RFC 5307, October 2008.

[RFC5317] Bryant, S., and Andersson, L. "Joint Working Team (JWT) Report on MPLS Architectural Considerations for a Transport Profile", RFC 5317, February 2009.

[G.808.1] ITU-T, "Generic Protection Switching - Linear trail and subnetwork protection,", Recommendation G.808.1, December 2003.

[G.841] ITU-T, "Types and Characteristics of SDH Network Protection Architectures," Recommendation G.841, October 1998.

[MPLS-TP-REQ] B. Niven-Jenkins, et al., "Requirements for MPLS-TP", draft-ietf-mpls-tp-requirements, work in progress.

[MPLS-TP-OAM] Vigoureux, M., Betts, M., and Ward, D., "MPLS TP OAM Requirements (MPLS)", work in progress.

11.2. Informative References

[RFC3386] Lai, W. and D. McDysan, "Network Hierarchy and Multilayer Survivability", RFC 3386, November 2002.

[RFC3469] Sharma, V., and Hellstrand, F., "Framework for Multi-Protocol Label Switching (MPLS)-based Recovery", RFC 3469, February 2003.

[RFC4426] Lang, J., Rajagopalan B., and D. Papadimitriou, Editors, "Generalized Multiprotocol Label Switching (GMPLS) Recovery Functional Specification", RFC 4426, March 2006.

[RFC4783] Berger, L., "GMPLS - Communication of Alarm Information", RFC 4783, December 2006.

[RFC4872] Lang, J., Rekhter, Y., and Papadimitriou, D., "RSVP-TE Extensions in Support of End-to-End Generalized Multi-Protocol Label Switching (GMPLS) Recovery", RFC 4872, May 2007.

[RFC4873] Berger, L., Bryskin, I., Papadimitriou, D., and Farrel, A., "GMPLS Segment Recovery", RFC 4873, May 2007.

[RFC4874] Lee, CY., Farrel, A., and De Cnodder, S., "Exclude Routes - Extension to Resource ReserVation Protocol-Traffic Engineering (RSVP-TE)", RFC 4874, April 2007.

[RFC5298] Takeda, T., Farrel, A., Ikejiri, Y., and Vasseur, JP., "Analysis of Inter-Domain Label Switched Path (LSP) Recovery", RFC 5298, August 2008.

[GMPLS-OAM] Takacs, A., Fedyk, D., and Jia, H., "OAM Configuration Framework and Requirements for GMPLS RSVP-TE", draft-ietf-ccamp-oam-configuration-fwk, work in progress.

[GMPLS-REV] Takacs, A., Fondelli, F., Tremblay, B., "GMPLS RSVP-TE recovery extension for data plane initiated reversion", draft-takacs-ccamp-revertive-ps, work in progress.

12. Editors' Addresses

Nurit Sprecher
Nokia Siemens Networks
3 Hanagar St. Neve Ne'eman B
45241 Hod Hasharon, Israel
Tel. +972 9 7751229
Email: nurit.sprecher@nsn.com

Adrian Farrel
Old Dog Consulting
Email: adrian@olddog.co.uk

13. Author's Address

Himanshu Shah
Ciena
Email: hshah@ciena.com

14. Intellectual Property Statement

Disclaimer of Validity

Full Copyright Statement