Network Working Group
Request for Comments: 4139
Category: Informational

D. Papadimitriou
Alcatel
J. Drake
Boeing
J. Ash
ATT
A. Farrel
Old Dog Consulting
L. Ong
Ciena

July 2005

Requirements for Generalized MPLS (GMPLS) Signaling Usage and Extensions for Automatically Switched Optical Network (ASON)

Status of This Memo

Copyright Notice

Abstract

The Generalized Multi-Protocol Label Switching (GMPLS) suite of protocols has been defined to control different switching technologies and different applications. These include support for requesting Time Division Multiplexing (TDM) connections, including Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) and Optical Transport Networks (OTNs).

This document concentrates on the signaling aspects of the GMPLS suite of protocols. It identifies the features to be covered by the GMPLS signaling protocol to support the capabilities of an Automatically Switched Optical Network (ASON). This document provides a problem statement and additional requirements for the GMPLS signaling protocol to support the ASON functionality.

1. Introduction

The Generalized Multi-Protocol Label Switching (GMPLS) suite of protocol specifications provides support for controlling different switching technologies and different applications. These include support for requesting Time Division Multiplexing (TDM) connections, including Synchronous Optical Network (SONET)/Synchronous Digital Hierarchy (SDH) (see [ANSI-T1.105] and [ITU-T-G.707], respectively), and Optical Transport Networks (see [ITU-T-G.709]). In addition, there are certain capabilities needed to support Automatically Switched Optical Networks control planes (their architecture is defined in [ITU-T-G.8080]). These include generic capabilities such as call and connection separation, along with more specific capabilities such as support of soft permanent connections.

This document concentrates on requirements related to the signaling aspects of the GMPLS suite of protocols. It discusses the functional requirements required to support Automatically Switched Optical Networks that may lead to additional extensions to GMPLS signaling (see [RFC3471] and [RFC3473]) to support these capabilities. In addition to ASON signaling requirements, this document includes GMPLS signaling requirements that pertain to backward compatibility (Section 5). A terminology section is provided in the Appendix.

2. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC2119].

While [RFC2119] describes interpretations of these key words in terms of protocol specifications and implementations, they are used in this document to describe design requirements for protocol extensions.

3. Problem Statement

The Automatically Switched Optical Network (ASON) architecture describes the application of an automated control plane for supporting both call and connection management services (for a detailed description see [ITU-T-G.8080]). The ASON architecture describes a reference architecture, (i.e., it describes functional components, abstract interfaces, and interactions).

The ASON model distinguishes reference points (representing points of information exchange) defined (1) between a user (service requester) and a service provider control domain, a.k.a. user-network interface (UNI), (2) between control domains, a.k.a. external network-network interface (E-NNI), and, (3) within a control domain, a.k.a. internal

network-network interface (I-NNI). The I-NNI and E-NNI interfaces are between protocol controllers, and may or may not use transport plane (physical) links. It must not be assumed that there is a one-to-one relationship between control plane interfaces and transport plane (physical) links, control plane entities and transport plane entities, or control plane identifiers for transport plane resources.

This document describes requirements related to the use of GMPLS signaling (in particular, [RFC3471] and [RFC3473]) to provide call and connection management (see [ITU-T-G.7713]). The functionality to be supported includes:

(a) soft permanent connection capability
(b) call and connection separation
(c) call segments
(d) extended restart capabilities during control plane failures
(e) extended label association
(f) crankback capability
(g) additional error cases

4. Requirements for Extending Applicability of GMPLS to ASON

The following sections detail the signaling protocol requirements for GMPLS to support the ASON functions listed in Section 3. ASON defines a reference model and functions (information elements) to enable end-to-end call and connection support by a protocol across the respective interfaces, regardless of the particular choice of protocol(s) used in a network. ASON does not restrict the use of other protocols or the protocol-specific messages used to support the ASON functions. Therefore, the support of these ASON functions by a protocol shall not be restricted by (i.e., must be strictly independent of and agnostic to) any particular choice of UNI, I-NNI, or E-NNI used elsewhere in the network. To allow for interworking between different protocol implementations, [ITU-T-G.7713] recognizes that an interworking function may be needed.

In support of the G.8080 end-to-end call model across different control domains, end-to-end signaling should be facilitated regardless of the administrative boundaries, protocols within the network, or the method of realization of connections within any part of the network. This implies the need for a clear mapping of ASON signaling requests between GMPLS control domains and non-GMPLS control domains. This document provides signaling requirements for G.8080 distributed call and connection management based on GMPLS, within a GMPLS based control domain (I-NNI), and between GMPLS based control domains (E-NNI). It does not restrict use of other (non GMPLS) protocols to be used within a control domain or as an E-NNI or UNI. Interworking aspects related to the use of non-GMPLS protocols,

such as UNI, E-NNI, or I-NNI -- including mapping of non-GMPLS protocol signaling requests to corresponding ASON signaling functionality and support of non-GMPLS address formats -- is not within the scope of the GMPLS signaling protocol. Interworking aspects are implementation-specific and strictly under the responsibility of the interworking function and, thus, outside the scope of this document.

By definition, any User-Network Interface (UNI) that is compliant with [RFC3473] (e.g., [GMPLS-OVERLAY] and [GMPLS-VPN]) is considered to be included within the GMPLS suite of protocols and MUST be supported by the ASON GMPLS signaling functionality.

Compatibility aspects of non-GMPLS systems (nodes) within a GMPLS control domain (i.e., the support of GMPLS systems and other systems that utilize other signaling protocols or some that may not support any signaling protocols) is described. For example, Section 4.5, 'Support for Extended Label Association', covers the requirements for when a non-GMPLS capable sub-network is introduced or when nodes do not support any signaling protocols.

4.1. Support for Soft Permanent Connection (SPC) Capability

A Soft Permanent Connection (SPC) is a combination of a permanent connection at the source user-to-network side, a permanent connection at the destination user-to-network side, and a switched connection within the network. An Element Management System (EMS) or a Network Management System (NMS) typically initiates the establishment of the switched connection by communicating with the node that initiates the switched connection (also known as the ingress node). The latter then sets the connection using the distributed GMPLS signaling protocol. For the SPC, the communication method between the EMS/NMS and the ingress node is beyond the scope of this document (as it is for any other function described in this document).

The end-to-end connection is thus created by associating the incoming interface of the ingress node with the switched connection within the network, along with the outgoing interface of the switched connection terminating network node (also referred to as egress node). An SPC connection is illustrated in the following figure. This shows the user's node A connected to a provider's node B via link #1, the user's node Z connected to a provider's node Y via link #3, and an abstract link #2 connecting the provider's node B and node Y. Nodes B and Y are referred to as the ingress and egress (respectively) of the network switched connection.

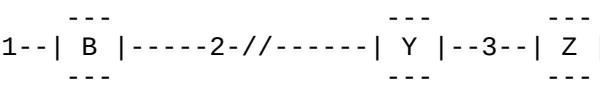

In this instance, the connection on link #1 and link #3 are both provisioned (permanent connections that may be simple links). In contrast, the connection over link #2 is set up using the distributed control plane. Thus, the SPC is composed of the stitching of link #1, #2, and #3.

Thus, to support the capability of requesting an SPC connection:

- The GMPLS signaling protocol MUST be capable of supporting the ability to indicate the outgoing link and label information used when setting up the destination provisioned connection.

- In addition, due to the inter-domain applicability of ASON networks, the GMPLS signaling protocol SHOULD also support indication of the service level requested for the SPC. In cases where an SPC spans multiple domains, indication of both source and destination endpoints controlling the SPC request MAY be needed. These MAY be done via the source and destination signaling controller addresses.

Note that the association at the ingress node, between the permanent connection and the switched connection, is an implementation matter that may be under the control of the EMS/NMS and is not within the scope of the signaling protocol. Therefore, it is outside the scope of this document.

4.2. Support for Call and Connection Separation

A call may be simply described as "An association between endpoints that supports an instance of a service" [ITU-T-G.8080]. Thus, it can be considered a service provided between two end-points, wherein several calls may exist between them. Multiple connections may be associated with each call. The call concept provides an abstract relationship between two users. This relationship describes (or verifies) the extent to which users are willing to offer (or accept) service to/from each other. Therefore, a call does not provide the actual connectivity for transmitting user traffic; it only builds a relationship by which subsequent connections may be made.

A call MAY be associated with zero, one, or multiple connections. For the same call, connections MAY be of different types and each connection MAY exist independently of other connections (i.e., each connection is setup and released with separate signaling messages).

The concept of the call allows for a better flexibility in how end-points set up connections and how networks offer services to users. For example, a call allows:

- An upgrade strategy for control plane operations, where a call control component (service provisioning) may be separate from the actual nodes hosting the connections (where the connection control component may reside).

- Identification of the call initiator (with both network call controller, as well as destination user) prior to connection, which may result in decreasing contention during resource reservation.

- General treatment of multiple connections, which may be associated for several purposes; for example, a pair of working and recovery connections may belong to the same call.

To support the introduction of the call concept, GMPLS signaling SHOULD include a call identification mechanism and SHOULD allow for end-to-end call capability exchange.

For instance, a feasible structure for the call identifier (to guarantee global uniqueness) MAY concatenate a globally unique fixed ID (e.g., may be composed of country code or carrier code) with an operator specific ID (where the operator specific ID may be composed of a unique access point code - such as source node address - and a local identifier). Other formats SHALL also be possible, depending on the call identification conventions between the parties involved in the call setup process.

4.3. Support for Call Segments

As described in [ITU-T-G.8080], call segmentation MAY be applied when a call crosses several control domains. As such, when the call traverses multiple control domains, an end-to-end call MAY consist of multiple call segments. For a given end-to-end call, each call segment MAY have one or more associated connections, and the number of connections associated with each call segment MAY be different.

The initiating caller interacts with the called party by means of one or more intermediate network call controllers, located at control domain boundaries (i.e., at inter-domain reference points, UNI or E-NNI). Call segment capabilities are defined by the policies associated at these reference points.

This capability allows for independent (policy based) choices of signaling, concatenation, data plane protection, and control plane driven recovery paradigms in different control domains.

4.4. Support for Extended Restart Capabilities

Various types of failures may occur, affecting the ASON control plane. Requirements placed on control plane failure recovery by [ITU-T-G.8080] include:

- Any control plane failure (i.e., single or multiple control channel and/or controller failure and any combination thereof) MUST NOT result in releasing established calls and connections (including the corresponding transport plane connections).

- Upon recovery from a control plane failure, the recovered control entity MUST have the ability to recover the status of the calls and the connections established before failure occurrence.

- Upon recovery from a control plane failure, the recovered control entity MUST have the ability to recover the connectivity information of its neighbors.

- Upon recovery from a control plane failure, the recovered control entity MUST have the ability to recover the association between the call and its associated connections.

- Upon recovery from a control plane failure, calls and connections in the process of being established (i.e., pending call/connection setup requests) SHOULD be released or continued (with setup).

- Upon recovery from a control plane failure, calls and connections in the process of being released MUST be released.

4.5. Support for Extended Label Association

It is an ASON requirement to enable support for G.805 [ITU-T-G.805] serial compound links. The text below provides an illustrative example of such a scenario, and the associated requirements.

Labels are defined in GMPLS (see [RFC3471]) to provide information on the resources used on a link local basis for a particular connection. The labels may range from specifying a particular timeslot, indicating a particular wavelength, or to identifying a particular port/fiber. In the ASON context, the value of a label may not be consistent across a link. For example, the figure below illustrates

the case where two GMPLS capable nodes (A and Z) are interconnected across two non-GMPLS capable nodes (B and C), where all of these nodes are SONET/SDH nodes, providing, for example, a VC-4 service.

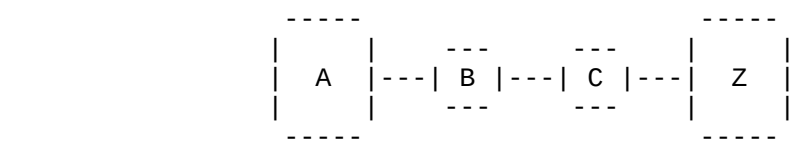

Labels have an associated implicit imposed structure based on [GMPLS-SONET] and [GMPLS-OTN]. Thus, once the local label is exchanged with its neighboring control plane node, the structure of the local label may not be significant to the neighbor node, as the association between the local and the remote label may not necessarily be the same. This issue does not present a problem in simple point-to-point connections between two control plane-enabled nodes in which the timeslots are mapped 1:1 across the interface. However, if a non-GMPLS capable sub-network is introduced between these nodes (as in the above figure, where the sub-network provides re-arrangement capability for the timeslots), label scoping may become an issue.

In this context, there is an implicit assumption that the data plane connections between the GMPLS capable edges already exist prior to any connection request. For instance, node A's outgoing VC-4's timeslot #1 (with SUKLM label=[1,0,0,0,0]), as defined in [GMPLS-SONET]), may be mapped onto node B's outgoing VC-4's timeslot #6 (label=[6,0,0,0,0]), or may be mapped onto node C's outgoing VC-4's timeslot #4 (label=[4,0,0,0,0]). Thus, by the time node Z receives the request from node A with label=[1,0,0,0,0], node Z's local label and timeslot no longer correspond to the received label and timeslot information.

As such, to support this capability, a label association mechanism SHOULD be used by the control plane node to map the received (remote) label into a locally significant label. The information necessary to allow mapping from a received label value to a locally significant label value can be derived in several ways including:

- Manual provisioning of the label association

- Discovery of the label association

Either method MAY be used. In case of dynamic association, the discovery mechanism operates at the timeslot/label level before the connection request is processed at the ingress node. Note that in the case where two nodes are directly connected, no association is

required. In particular, for directly connected TDM interfaces, no mapping function (at all) is required due to the implicit label structure (see [GMPLS-SONET] and [GMPLS-OTN]). In these instances, the label association function provides a one-to-one mapping of the received to local label values.

4.6. Support for Crankback

Crankback has been identified as an important requirement for ASON networks. Upon a setup failure, it allows a connection setup request to be retried on an alternate path that detours around a blocked link or node (e.g., because a link or a node along the selected path has insufficient resources).

Crankback mechanisms MAY also be applied during connection recovery by indicating the location of the failed link or node. This would significantly improve the successful recovery ratio for failed connections, especially in situations where a large number of setup requests are simultaneously triggered.

The following mechanisms are assumed during crankback signaling:

- The blocking resource (link or node) MUST be identified and returned in the error response message to the repair node (that may or may not be the ingress node); it is also assumed that this process will occur within a limited period of time.

- The computation (from the repair node) of an alternate path around the blocking link or node that satisfies the initial connection constraints.

- The re-initiation of the connection setup request from the repair node (i.e., the node that has intercepted and processed the error response message).

The following properties are expected for crankback signaling:

- Error information persistence: the entity that computes the alternate (re-routing) path SHOULD store the identifiers of the blocking resources, as indicated in the error message, until the connection is successfully established or until the node abandons rerouting attempts. Since crankback may happen more than once while establishing a specific connection, the history of all experienced blockages for this connection SHOULD be maintained (at least until the routing protocol updates the state of this information) to perform an accurate path computation that will avoid all blockages.

- Rerouting attempts limitation: to prevent an endless repetition of connection setup attempts (using crankback information), the number of retries SHOULD be strictly limited. The maximum number of crankback rerouting attempts allowed MAY be limited per connection or per node:

  - When the number of retries at a particular node is exceeded, the node that is currently handling the failure reports the error message upstream to the next repair node, where further rerouting attempts MAY be performed. It is important that the crankback information provided indicate that re-routing through this node will not succeed.

  - When the maximum number of retries for a specific connection has been exceeded, the repair node that is handling the current failure SHOULD send an error message upstream to indicate the "Maximum number of re-routings exceeded". This error message will be sent back to the ingress node with no further rerouting attempts. Then, the ingress node MAY choose to retry the connection setup according to local policy, using its original path, or computing a path that avoids the blocking resources.

  Note: After several retries, a given repair point MAY be unable to compute a path to the destination node that avoids all of the blockages. In this case, it MUST pass the error message upstream to the next repair point.

4.7. Support for Additional Error Cases

To support the ASON network, the following additional category of error cases are defined:

- Errors associated with basic call and soft permanent connection support. For example, these MAY include incorrect assignment of IDs for the Call or an invalid interface ID for the soft permanent connection.

- Errors associated with policy failure during processing of the new call and soft permanent connection capabilities. These MAY include unauthorized requests for the particular capability.

- Errors associated with incorrect specification of the service level.

5. Backward Compatibility

As noted above, in support of GMPLS protocol requirements, any extensions to the GMPLS signaling protocol, in support of the requirements described in this document, MUST be backward compatible.

Backward compatibility means that in a network of nodes, where some support GMPLS signaling extensions to facilitate the functions described in this document, and some do not, it MUST be possible to set up conventional connections (as described by [RFC3473]) between any arbitrary pair of nodes and to traverse any arbitrary set of nodes. Further, the use of any GMPLS signaling extensions to set up calls or connections that support the functions described in this document MUST not perturb existing conventional connections.

Additionally, when transit nodes that do not need to participate in the new functions described in this document lie on the path of a call or connection, the GMPLS signaling extensions MUST be such that those transit nodes are able to participate in the establishment of a call or connection by passing the setup information onwards, unmodified.

Lastly, when a transit or egress node is called upon to support a function described in this document, but does not support the function, the GMPLS signaling extensions MUST be such that they can be rejected by pre-existing GMPLS signaling mechanisms in a way that is not detrimental to the network as a whole.

6. Security Considerations

Per [ITU-T-G.8080], it is not possible to establish a connection in advance of call setup completion. Also, policy and authentication procedures are applied prior to the establishment of the call (and can then also be restricted to connection establishment in the context of this call).

This document introduces no new security requirements to GMPLS signaling (see [RFC3471]).

7. Acknowledgements

The authors would like to thank Nic Larkin, Osama Aboul-Magd, and Dimitrios Pendarakis for their contribution to the previous version of this document, Zhi-Wei Lin for his contribution to this document, Deborah Brungard for her input and guidance in our understanding of the ASON model, and Gert Grammel for his decryption effort during the reduction of some parts of this document.

8. References

8.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

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

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

8.2. Informative References

[ANSI-T1.105] ANSI, "Synchronous Optical Network (SONET): Basic Description Including Multiplex Structure, Rates, and Formats", T1.105, October 2000.

[GMPLS-OTN] Papadimitriou, D., Ed., "Generalized MPLS (GMPLS) Signaling Extensions for G.709 Optical Transport Networks Control", Work in Progress, January 2005.

[GMPLS-OVERLAY] Swallow, G., Drake, J., Ishimatsu, H., and Y. Rekhter, "Generalize Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model", Work in Progress, October 2004.

[GMPLS-SONET] Mannie, E. and D. Papadimitriou, "Generalized Multi-Protocol Label Switching (GMPLS) Extensions for Synchronous Optical Network (SONET) and Synchronous Digital Hierarchy (SDH) Control", RFC 3946, October 2004.

[GMPLS-VPN] Ould-Brahim, H. and Y. Rekhter, Eds., "GVPN Services: Generalized VPN Services using BGP and GMPLS Toolkit", Work in Progress, May 2004.

[ITU-T-G.707] ITU-T, "Network Node Interface for the Synchronous Digital Hierarchy (SDH)", Recommendation G.707, October 2000.

[ITU-T-G.709] ITU-T, "Interface for the Optical Transport Network (OTN)", Recommendation G.709 (and Amendment 1), February 2001 (October 2001). http://www.itu.int

[ITU-T-G.7713] ITU-T "Distributed Call and Connection Management", Recommendation G.7713/Y.1304, November 2001. http://www.itu.int

[ITU-T-G.805] ITU-T, "Generic functional architecture of transport networks)", Recommendation G.805, March 2000. http://www.itu.int

[ITU-T-G.8080] ITU-T "Architecture for the Automatically Switched Optical Network (ASON)", Recommendation G.8080/Y.1304, November 2001 (and Revision, January 2003). http://www.itu.int

Appendix - Terminology

This document makes use of the following terms:

Administrative domain: See Recommendation G.805 [ITU-T-G.805].

Call: Association between endpoints that supports an instance of a service.

Connection: Concatenation of link connections and sub-network connections that allows the transport of user information between the ingress and egress points of a sub-network.

Control Plane: Performs the call control and connection control functions. The control plane sets up and releases connections through signaling, and may restore a connection in case of a failure.

(Control) Domain: Represents a collection of entities that are grouped for a particular purpose. G.8080 applies this G.805 recommendation concept (that defines two particular forms: the administrative domain and the management domain) to the control plane in the form of a control domain. Entities grouped in a control domain are components of the control plane.

External NNI (E-NNI): Interfaces are located between protocol controllers that are situated between control domains.

Internal NNI (I-NNI): Interfaces are located between protocol controllers within control domains.

Link: See Recommendation G.805 [ITU-T-G.805].

Management Plane: Performs management functions for the Transport Plane, the control plane, and the system as a whole. It also provides coordination between all the planes. The following management functional areas are performed in the management plane: performance, fault, configuration, accounting, and security management.

Management Domain: See Recommendation G.805 [ITU-T-G.805].

Transport Plane: Provides bi-directional or unidirectional transfer of user information, from one location to another. It can also provide transfer of some control and network management information. The Transport Plane is layered and is equivalent to the Transport Network defined in G.805 [ITU-T-G.805].

User Network Interface (UNI): Interfaces are located between protocol controllers, between a user and a control domain.

Authors' Addresses

Dimitri Papadimitriou
Alcatel
Francis Wellesplein 1,
B-2018 Antwerpen, Belgium

Phone: +32 3 2408491
EMail: dimitri.papadimitriou@alcatel.be

John Drake
Boeing Satellite Systems
2300 East Imperial Highway
El Segundo, CA 90245

EMail: John.E.Drake2@boeing.com

Adrian Farrel
Old Dog Consulting

Phone: +44 (0) 1978 860944
EMail: adrian@olddog.co.uk

Gerald R. Ash
ATT
AT&T Labs, Room MT D5-2A01
200 Laurel Avenue
Middletown, NJ 07748, USA

EMail: gash@att.com

Lyndon Ong
Ciena
PO Box 308
Cupertino, CA 95015, USA

EMail: lyong@ciena.com

Full Copyright Statement

Intellectual Property

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.