



Management Plane (OSS) Support for Control Plane Networks

OIF-CP-MGMT-01.0

Executive Summary

CP capabilities are instrumental in streamlined management of network resources and improvements in service delivery time. At the same time, they can provide scalable restoration combining the advantages of networks scalability (as opposed to rings) and the speed of network-based restoration (as opposed to centralized OS-driven restoration). However, to fully take advantage of these technologies, network operators are required to evolve their network management systems with added functionality.

This document attempts to document some of the network management functionality that is needed for the management of multi-domain intelligent optical networks. It provides some background information, assumptions, application scenarios, and discussions of functionality specific to the management of CP driven networks. Carriers have expressed interest in the promotion of operational support for multi-domain CP driven networks (to form a logically integrated CP framework) and in the rapid deployment of new CP technologies and capabilities without network service disruption. The CP has a critical role in realizing self-running and self-governing intelligent networks across multiple technologies and heterogeneous platforms from multiple vendors. Throughout this paper it is assumed that different carriers have different processes, different functional groupings of functional blocks. This document is intended to provide a generic set of guidelines to facilitate deployment of CP in carrier networks. The objective of the paper is to provide recommendations to make the control plane technology operational in the carrier's network.

Table of Contents

Table of Contents	2
List of Figures	3
List of Tables	4
1 Introduction and Project Objectives.....	5
2 Scope.....	6
3 Assumptions.....	7
4 Reference Models	8
5 Mediation Models	14
5.1 EMS Mediation	14
5.2 Management View of CP	15
6 Management Plane Guidelines for CP-driven Networks.....	17
6.1 Control Plane ID Management.....	17
6.2 Summary of CP IDs – CP Identifiers and Related Name Spaces	26
6.3 Management of Routing Controller and Routing Area Hierarchies	27
6.4 SCN and Control Channel (CC) Management, Configuration, and Security	29
6.5 MP in Capacity Planning.....	30
6.6 MP Management of Service Recovery.....	31
6.7 Security Considerations for OSS Management of CP	33
6.8 Service Management across (Heterogeneous) Domains: Administration of Cross OSS CP Connection: CP Connection Spanning Multiple MP/Servers	34
6.9 Transition Support.....	39
7 References.....	51
8 Appendix A: List of Contributors	53
9 Appendix B: About the OIF.....	54
10 Appendix C: Glossary.....	55
11 Appendix D: Acronym Definitions.....	57

List of Figures

Figure 1: TMF513 Model with NMS to NMS Interface.....	8
Figure 2: TMF513: Reference Model (No NMS Interface).....	9
Figure 3: G.7718 Model with NMS Managing Two EMSs and Intra-carrier Domains ...	10
Figure 4: G.7718 Scenario with an NMS Managing an EMS that Manages More than One Intra Carrier Domain	11
Figure 5: Scenario with NMS Managing Multiple EMSs with an EMS Handling Traditional Domain (Non CP)	12
Figure 6: Scenario with More than One Traditional Domain and Non ASON Inter- domain Connectivity to a Control Domain.....	13
Figure 7: EMS Mediation Model.....	15
Figure 8: NMS and EMS Model with CP Resources	16
Figure 9: SRLG: SRLG of CP Link A1-A2 = {12, 13, 14}	25
Figure 10: MP Role in CP ID Management.....	26
Figure 11: MP Role in CP ID Management – EMS Mediation Model	27
Figure 12: MP Service Spanning Multiple Domains.....	35
Figure 13: End-to-End Service with Three Different CPs and Management Plane Handoffs.....	36
Figure 14: End-to-end CP	37
Figure 15: Configuration of an End-to-End Call with Multiple Call Segments and Connection Segments within Call Segments	37
Figure 16: Path Retrieval Problem Definition	39
Figure 17: MP-to-CP Transition Use Cases.....	41
Figure 18: CP-to-MP Transition Use Cases.....	44
Figure 19: CP Mixed Network Routing Configuration	46
Figure 20: Transition Scenarios	47
Figure 21: Routing Configurations	48

List of Tables

Table 1: Control Plane Identifiers and Related Name Spaces	17
Table 2: Call and Connection Attributes	38

1 Introduction and Project Objectives

Control Plane (CP) capabilities are ubiquitously deployed in telecommunications networks at layers two and higher; the best-known examples are the Internet and the voice network. Recent years have witnessed the development and deployment of CP-based optical networks, concurrent with the development of CP standards for both intra- and inter-domain protocols.

CP capabilities are instrumental in streamlined management of network resources and improved service delivery time. They also can provide scalable restoration combining the advantages of networks scalability (versus rings) and the speed of network-based restoration (versus centralized OS-driven restoration). However, to realize the full advantage of these technologies, network operators are required to enhance their network management systems with additional functionality.

This document identifies some of the network management functionality needed for the management of multi-domain intelligent optical networks. It also provides background information, assumptions, application scenarios, and discussions of functionality specific to the management of CP-driven networks. Carriers have expressed interest in the promotion of operational support for multi-domain CP-driven networks to form a logically integrated CP framework and in the rapid deployment of new CP technologies and capabilities without network service disruption. The CP plays a crucial role in realizing self-running and self-governing intelligent networks across multiple technologies and heterogeneous platforms from multiple vendors.

Throughout this paper, it is assumed that individual carriers have different processes and groupings of functional blocks. It is intended to provide a generic set of guidelines to facilitate deployment of CP within these diverse carrier networks.

The text and diagrams are intended to provide additional information; e.g. the architecture reference models in Section 4 form the basis for understanding the scope of the paper.

We also cover Management Plane (MP) implications for CP support of Soft Permanent Connections (SPC) and Switched Connections (SC) for OTN, SONET/SDH, and Ethernet bearer services.

The term *multi-domain* is used to indicate that the CP is partitioned; the CP partitions are called “control domains”, or “domains” throughout this document. A *control domain (CD)* is an abstraction entity that allows the details of the control implementation within the domain to be represented by a single architectural component [G.8080]. This allows independent control implementations within each domain, yet hides the internal control details. A control domain defines the scope of a CP function on a network. For example, a routing control domain is a control domain for a routing controller.

2 Scope

This document is intended to guide work on MP specifications and architecture for CP-based multi-domain networks. It also documents network management architecture alternatives for managing multi-domain networks consisting of CP-driven and legacy domains. CP functions are instances of existing network management applications (e.g., calculating a path for a connection); therefore, CP management includes applications and interactions with non-CP network management applications.

We assume consistency with relevant ITU-T and TMF documents [19], [5], and [12] and the OIF Carrier WG Guideline Document [18].

To provide timely guidance, no attempt has been made to provide comprehensive or detailed requirements at this time; instead, the focus is on topics of particular interest or concern to the carrier community. We assume that the reader of this document is willing and able to derive the specific requirements necessary to define a product that meets its intent.

This document provides additional guidance through the documentation of architecture alternatives and MP implications regarding functionality not already provided by the TMF specifications. Further, it focuses on the changing role of management systems, as specific functions (such as path computation) are added to network elements. It addresses the role played by management systems in heterogeneous networks with CP- and MP-driven domains. The following general topics are addressed:

1. CP ID Management
2. Management of routing controller and routing area.
3. SCN and Control Channel management, configuration and security
4. MP in Capacity Planning
5. MP management of service recovery
6. Security considerations for OSS management of CP
7. Service Management across heterogeneous domains
8. Transition Support

3 Assumptions

This document is based on the following assumptions:

- ITU-T [5] architecture
- ITU-T [19] and [6] transport architecture
- MEF and ITU-T defined Ethernet services
- TMF MTNM requirements, architecture, and interfaces (TMF 814)
- A multi-vendor, multi-domain heterogeneous environment (i.e., domains may use different signaling and routing protocols, or architectures)
- Domains agnostic to each other's internal signaling or routing protocols

4 Reference Models

This section provides reference models to be assumed when considering the guidelines in Section 6 and other sections of this whitepaper. It also presents the NMS-EMS models and provides a context for the whitepaper.

An OSS is considered an instance of the NMS and EMS implementations, where both functions are embedded in the same suite of systems. The models specified in this section do not include an OSS and are consistent with G.7718.

This section introduces several models for managing a multi-domain network. Each model has implications on the visibility of network resources by management systems and the placement of functional components in various systems.

Figure 1 depicts (at a very high level) the layering and interactions of network management layers SMS, NMS, and EMS.

NOTE: The scenarios with E-NNI Proxy are not within the current scope of the OSS CP Management whitepaper and could be considered for discussion beyond its scope.




Figure 1: TMF513 Model with NMS to NMS Interface

The model depicted in Figure 1 shows a Service Management System interfacing with multiple NMSs, with each NMS controlling the respective subnetwork through the EMS instance and each EMS instance managing its own subnetwork (i.e., EMS A manages the ring network while EMS B manages the point-to-point network). The model also shows an interface between the NMSs.

TMF513 Reference Model

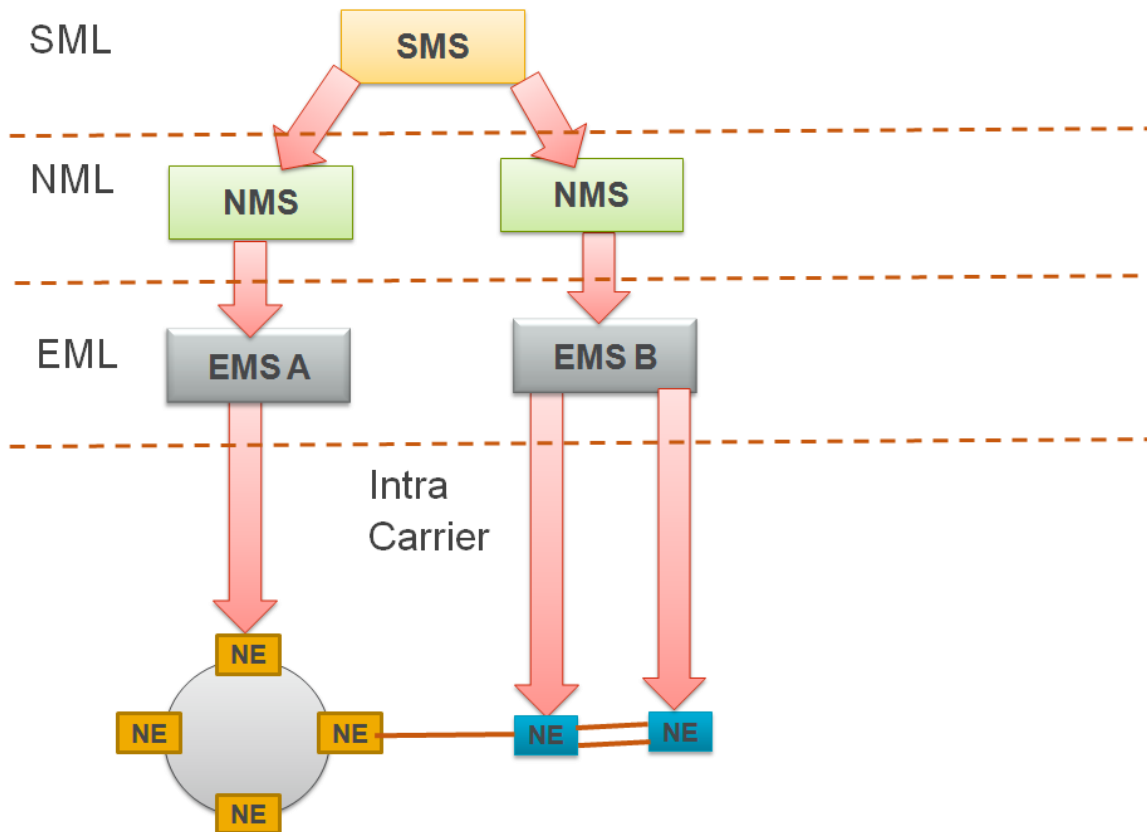


Figure 2: TMF513: Reference Model (No NMS Interface)

The model shown in Figure 2 illustrates a Service Management System interfacing with multiple NMSs and each NMS controlling the respective subnetwork through the EMS instance with each EMS instance managing its own subnetwork (i.e., EMS A manages the ring network while EMS B manages the point-to-point network). This differs from the previous model in that there is no interface between the NMSs.




Figure 3: G.7718 Model with NMS Managing Two EMSs and Intra-carrier Domains

The model shown in Figure 3 depicts an architecture where each EMS has management responsibility and visibility for one individual control domain. The NMS interacts with an individual EMS to perform its functions, with each EMS managing an individual control domain.

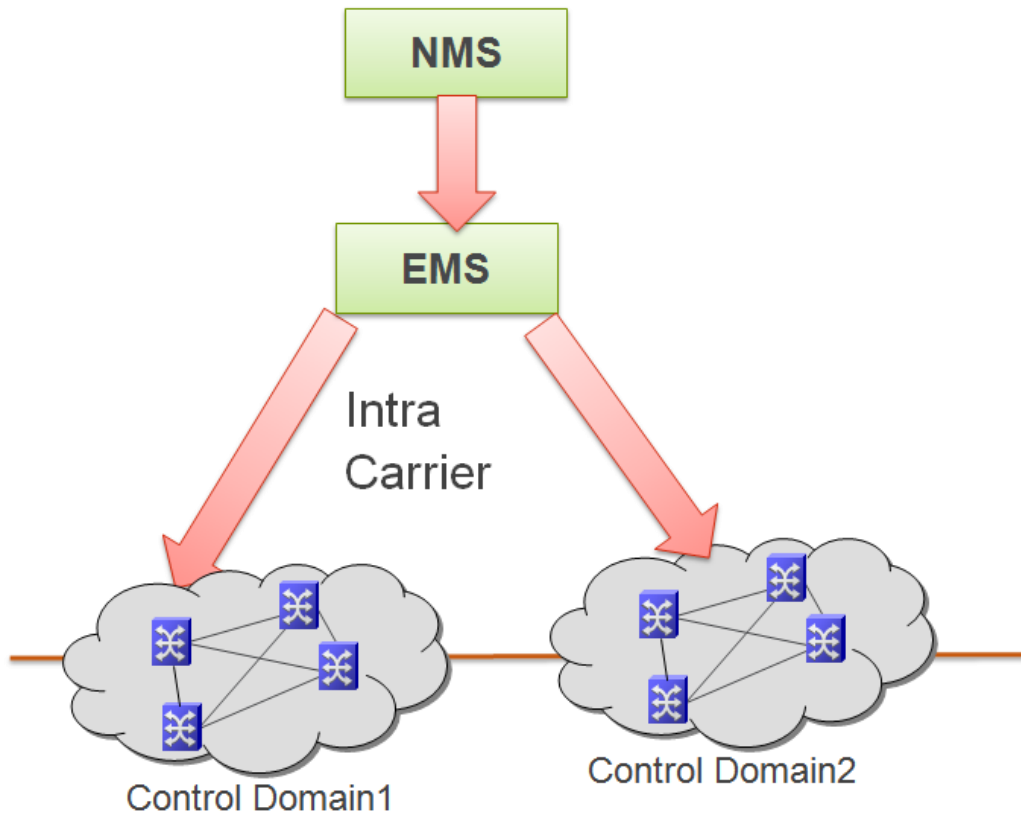


Figure 4: G.7718 Scenario with an NMS Managing an EMS that Manages More than One Intra Carrier Domain

Figure 4 represents a single NMS/EMS instance managing more than one intra-carrier domain, i.e., the control domains 1 and 2 are under a single NMS/EMS instance.

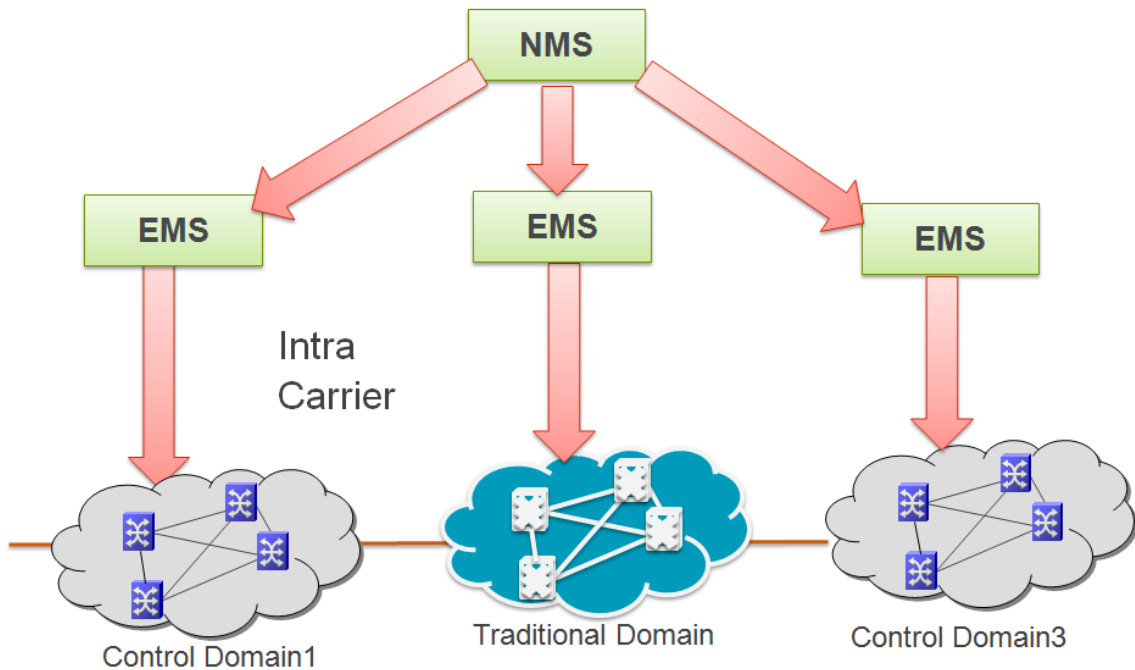


Figure 5: Scenario with NMS Managing Multiple EMSs with an EMS Handling Traditional Domain (Non CP)

Figure 5 shows an intra-carrier scenario, partly controlled by traditional management and partly controlled via CP. In this scenario, multiple EMS instances manage the respective traditional and control domains under the same NMS instance. This arrangement does not convey any E-NNI implications, i.e., the link between the control domain1 and the traditional domain is a traditional, non-CP link. The link between the traditional domain and Control Domain3 is also a traditional link.

Service management for this scenario poses some challenges since it needs to provide a unified view of the transport services. Section 6.8.1 discusses the alternative in greater detail, i.e., where all three domains are controlled under different NMS/EMS instances. In this case, the traditional domain switches at the same layer as the two control domains.




Figure 6: Scenario with More than One Traditional Domain and Non ASON Inter-domain Connectivity to a Control Domain

The model shown in Figure 6 illustrates a hybrid intra-carrier network management scenario for two traditionally managed domains interworking across the ASON domain, with each domain (traditional or control) under different EMS instances, but under the same NMS instance.

5 Mediation Models

This section describes the two mediation models adopted for this whitepaper, i.e., with the OSS and with the NMS/EMS. All mediation models illustrated in this paper show references with the NMS/EMS since OSS is an implementation of NMS/EMS, thereby including the OSS model, i.e., the OSS could comprise of NMS and/or EMS instances. As such, OSS refers to an architecture of EMS, NMS implementing various functions. There could also be other OSS applications (HP OV) that collect alarms in SNMP directly from the NEs.

Network Management of transport networks consist of functions, interfaces, and processes. These can be combined (or packaged) into many different systems and instantiated in centralized and/or distributed platforms.

Functions have traditionally been generically described in the FCAPS taxonomy of Fault, Configuration, Accounting, Performance, and Security [7]. These general functions are combined in various network management applications. For example, an application to create a connection may use Configuration, Accounting, and Security type functions.

While there are many interfaces in network management systems, the two main mediation models adopted for the OIF OSS project are the OSS-to-NE, and the NMS-EMS-NE model from TMF [1].

Processes in network management are developed to operate a network and are driven by business needs. The Business Process Framework from TMF “is an industry consensus on the Service Provider processes” and covers planning and lifecycle management, operational management, and corporate (or business-support) management. It is also published in ITU-T [9]. The eTOM provides an alternative taxonomy to FCAPS that is more business-focused (this is often referred to as FAB¹). In this paper, it is assumed that an OSS contains applications used in most operational processes, such as order handling, billing, service configuration, and resource provisioning.

5.1 EMS Mediation

An Element Management Systems (EMS) is a collection of applications (often on a single platform) performing various FCAPS or FAB functions on a set of resources in the network. EMSs interact with a Network Management System (NMS) using the TMF814 (or similar) interface. An EMS-to-Network Element (NE) interface (which may be either open or proprietary) is used for information exchange between the EMS and the NE. Management functions on the NE are scoped to the NE and interact with the EMS.

Functionally, the NE-EMS, EMS-NMS, and NE-OSS interfaces are similar in their one-to-many relationship, where a higher level management entity has a set of lower level management entities within its scope. This is true for the EMS-NE relationship since nodal level network management resides on the NEs. NE/EMS/NMS/OSS entities can be viewed as locations for applications containing various FCAPS (or FAB) functions.

¹ Emphasising a key focal point in eTOM of Fulfilment, Assurance and Billing

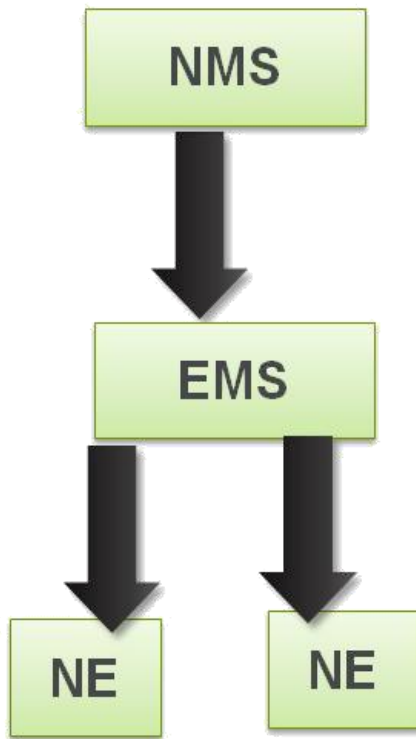


Figure 7: EMS Mediation Model

As illustrated in Figure 7, the network can be managed by the OSS applications, the NMS/EMS, or a combination of both.

5.2 Management View of CP

From an architecture perspective, the CP is essentially the automation of MP applications that existed prior to the CP; these applications are service and connection management.

The management of CP entities is similar to the management of other network management applications and controls the following:

- CP components. These are instances of routing controllers, call controllers, connection controllers, link resource managers, discovery agents, and termination and adaptation performers. They are defined in the ASON architecture [5]. Management of CP components includes relevant FCAPS/FAB of protocols, protocol communications (adjacencies over SCN), addresses of the components, scope of the protocols, and management commands to these components.
- CP resource addressing. The CP applications use different name spaces than other network management applications for resources. Configuration (and reconfiguration) of the routing topology (SNPP name space) and service end points (UNI Transport Resource Identifiers or OIF TNAs) is required prior to activation of CP components.

- Mappings between the CP resource addressing and other network management name spaces, e.g., building/floor/shelf/card name spaces is needed so that users can view CP-created connections in terms of name spaces used by other management applications.

The MP view of a CP includes the routing topology view, the distribution of RC/SC functions, and the reference points, as shown in Figure 8.

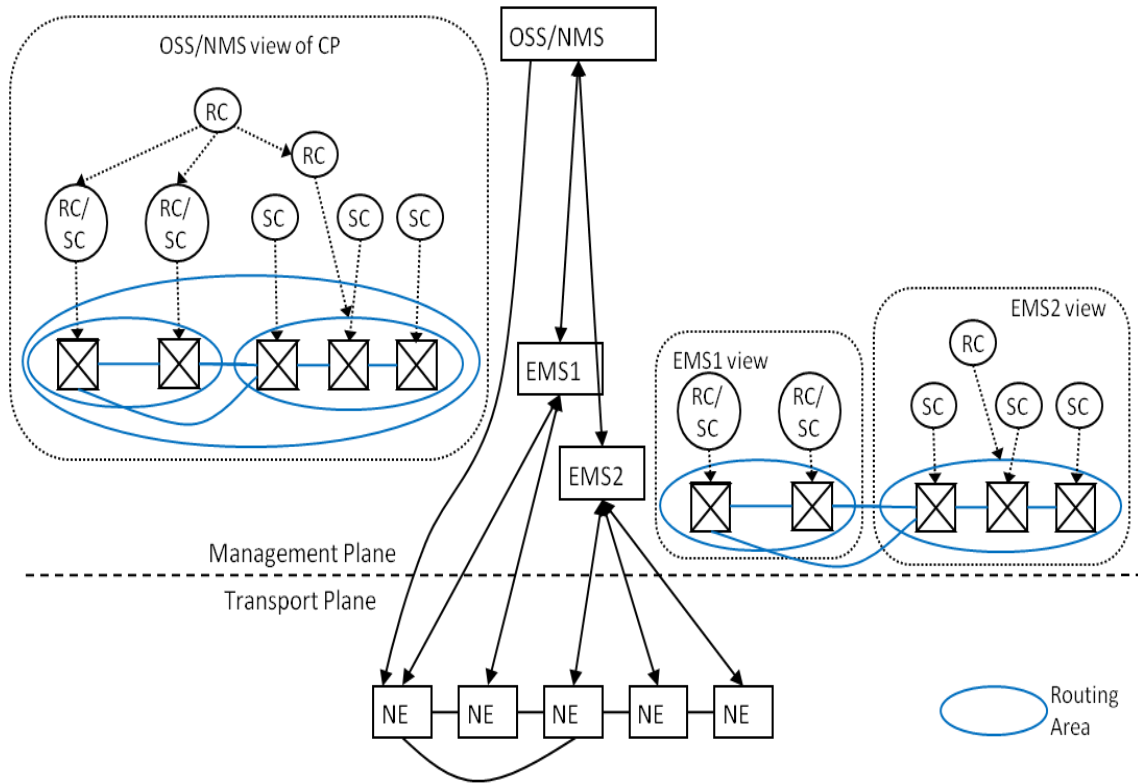


Figure 8: NMS and EMS Model with CP Resources

6 Management Plane Guidelines for CP-driven Networks

6.1 Control Plane ID Management

Table 1 describes the relationships between name spaces used by TMF MTNM, ASON, and specific protocols for the CP categories described in Section 5.

The first column contains elements from MTNM v.3.5, the reference information model for the OSS CP management project. The second column contains elements from ASON, the reference architecture for the CP. Each row represents a class of identifiers; the main description of the MTNM is in the first column. Equivalent identifiers from each different model/implementation are shown in columns of the same row. Network management applications for the CP need to map to the MTNM model which then should be mapped to CP implementations.

The COMMON LANGUAGE® Identifiers column refers to several identifier spaces used by several major carriers.

Specific functions of the MTNM model are not included in the table. These are typically associated with managed objects and are APIs. Descriptions of CP related functions can be found in SD1-45 of MTNM v3.5.

Table 1: Control Plane Identifiers and Related Name Spaces

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE® Identifiers
<p>Signaling Control Network. No capabilities to configure and manage in v3.5. It is assumed to be present. Missing capabilities are requested for future versions of the information model.</p>	<p>Signaling Communication Channel (SCN). Addresses for communicating over the SCN.</p>	<p>IP network. Addressing is separate from other CP and network management name spaces</p> <p>OIF UNI 2.0: UNI-C/UNI-N SC PC SCN IP address</p> <p>OIF ENNI 2.0 Signaling: eNNI-D/eNNI-U SCN IP address</p> <p>OIF ENNI 2.0 Routing:</p>	<p>No impact identified at this point.</p>	<p>No impact identified at this point.</p>

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE ® Identifiers
		Source/Dest IP Addresses (for RC PC SCN address)		
<p>Routing Controller. No capabilities to configure and manage in v3.5. It is assumed to be present. Missing capabilities are requested for future versions of the information model.</p>	<p>CP component: Routing Controller: Controller ID and SCN attachments</p> <p>Protocol parameters (e.g., timers)</p> <p>Routing Adjacencies</p> <p>Routing Area in scope</p> <p>Routing Information Database</p> <p>Remote Route Query identifier and SCN address</p>	<p>OIF ENNI 2.0 Routing</p> <p>Router Address TLV (advertising RC)</p> <p>Router ID in OSPF packet Header and Advertising Router ID (RC PC ID)</p>	<p>No impact identified at this point.</p>	<p>No impact identified at this point.</p>
<p>Signaling Controller. No capabilities to configure in v3.5. It is assumed to be present. Missing capabilities are requested for future versions of the information model.</p> <p>Can establish and release calls, assign/de-assign TNAs</p> <p>TNAName</p> <p>TNA</p> <p>CallEnd</p>	<p>CP component: Call Controller: Controller ID and SCN attachments</p> <p>Signaling Adjacencies</p> <p>Service endpoints (TNAs)</p> <p>Call Identifiers (dynamically created)</p> <p>Call Parameters (diversity, restoration, etc.)</p> <p>Call Detail record(s)</p>	<p>OIF UNI 2.0 and ENNI 2.0 Signaling.</p> <p>call_ID</p> <p>SESSION_ATTRIBUTE</p> <p>DESTINATION_TNA</p> <p>SOURCE_TNA</p> <p>DIVERSITY</p> <p>SERVICE_LEVEL</p> <p>ENNI 2.0 Routing: Node IPv4 Local</p>	<p>TNA Address (TNAID)</p>	<p>CLFI (Bundled TE Link)</p> <p>Service Identifiers (ASR as an example & other service orders as appropriate)</p>

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE ® Identifiers
<p><i>TNAName</i> Group <i>TNAName</i></p> <p>CallSNC call ID (<i>CP assigned</i>) A-End <i>CallEnd</i> Z-End <i>CallEnd</i> CallState CallParameterProfile (classOfService, classOfServiceparameters, degradedThreshold, severelyDegradedThreshold) CallDiversity (nodeSRGType, linkSRGType, nodeDiversityLevelOfEffort, linkDiversityLevelOfEffort, coroutingLevelOfEffort) NodeDiversityViolations (<i>List <DiversityInfo></i>) LinkDiversityViolations (<i>List <DiversityInfo></i>)</p> <p>MLSNPP Link A-End <i>TNAName</i> (implies same value across all layers) Z-End <i>TNAName</i> (implies same value across all</p>	<p>Connection(s) associated with call</p>	<p>Address sub-TLV and Node IPv6 Local Address Sub-TLV (TNAs in Node Attribute TLV) NSAP TNA Sub-TLV (TNAs in Node Attribute TLV)</p>		

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE ® Identifiers
<p>layers)</p> <p>A-End Group <i>TNAName</i> (implies same value across all layers)</p> <p>Z-End Group <i>TNAName</i> (implies same value across all layers)</p> <p>SNPP <i>TNAName</i> Group <i>TNAName</i></p> <p>SNP <i>TNAName</i> TerminationPoint Name</p>				
<p>Can establish and release connections. Can add a connection to a call.</p> <p>TopLevelConnection (SNC)</p> <p>Connection Segment/RA Connection(SNC)</p> <p>RN Connection (XC)</p> <p>RouteGroupInfo</p> <p>SNP TNA Name</p>	<p>CP component: Connection Controller:</p> <p>Controller ID and SCN attachments</p> <p>Signaling Adjacencies</p> <p>Connection Identifiers (dynamically created)</p>	<p>OIF UNI 2.0:</p> <p>UNI_IPv4_SESSION</p> <p>EGRESS_LABEL/SPC_LABEL</p> <p>SONET/SDH_SENDER_TSPEC</p> <p>G.709_SENDER_TSPEC</p> <p>ETHERNET_SENDER_TSPEC</p> <p>UNI_IPv4_SESSION</p> <p>LSP_TUNNEL_IPv4_FILTER_SPE</p> <p>IPv4_IF_ID_RSVP_HOP</p>	<p>Logical Port ID</p>	<p>CLCI name space for connection identification</p> <p>Service code modifiers/services identifiers for bandwidth identification.</p>

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE ® Identifiers
(containing LP ID and Index/channel) TerminationPoint Name (mapping into MLSN namespace)		SONET/SDH_F LOWSPEC G.709_FLOWSPEC ETHERNET_FLOWSPEC ENNI 2.0 Signaling.		
Three levels of routing hierarchy are supported (Top-Level, Intermediate and RN). MultiLayerRoutingArea (MLSN) MLSN Name (“TPRA” for Top-Level MLRA, CP-assigned RAId for Intermediate RA and NodeID for RN) – This is discovered from CP	CP resources (SNPP name space): Routing Area (parent RA, child RAs, level, Shared Risk Group etc.)	ENNI 2.0 Routing.	No impacts identified at this point.	No impacts identified at this point.
No configuration of SNPP links. Needed in future. MLSNPPLink A-END MLRA Name Z-END MLRA Name Direction (<i>UNI, BI</i>) InterfaceType (<i>UNI, I-NNI,</i>	CP resources (SNPP name space): Node identifiers Link end identifiers. Link attributes (Shared Risk Group, direction, layer ID, capacity, etc.)	ENNI 2.0 Signaling ENNI 2.0 Routing: (TE) Link TLV (link attributes) Local and Remote TE Router ID sub-TLV (node IDs)	Node ID (Node AID) Shared Risk Link Group	CLLI name space is part of AID name. Please note that Node ID to CLLI mapping is a many to one (i.e., many Node IDs in the CP name space can map to one and the same CLLI).

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE ® Identifiers
<p><i>Internal-E-NNI, External-E-NNI, UNSPECIFIED)</i></p> <p>SignalingProtocol (CRLDP, RSVP, PNNI)</p> <p>LinkSRG (<i>List <RiskType, Values></i>)</p> <p><i>List <LayeredSNPPLink></i></p> <p>CallEnd</p> <p>SNPP ID (<i>CP-assigned</i>)</p> <p>SNP ID (<i>CP-assigned</i>)</p> <p>TerminationPoint Name</p> <p>LayeredSNPPLink</p> <p>LayerRate</p> <p><i>List <SNPPLink></i></p> <p>SNPPLink</p> <p>SNPPLink ID (<i>CP assigned</i>)</p> <p>A-END <i>SNPP</i></p> <p>Z-END <i>SNPP</i></p> <p>SNPP</p> <p>SNPP ID (CP assigned)</p> <p>RA ID</p> <p><i>List <SNP></i></p>				<p>CLLI identifies the location of a given Node ID.</p> <p>CLFI name space is for links</p> <p>CLCI (Connections for a given bandwidth within a carrier's network)</p>

TMF MTNM v3.5 Information Model	ASON Architectural Model	Protocol Realization	Other Network Management Application	COMMON LANGUAGE® Identifiers
SNP SNP ID (<i>CP assigned</i>) TerminationPoint Name				

NOTE: In rows identified with “No Impacts identified at this time” the expectation is that the ID can be managed with its definition in the name spaces specified in the other columns.

This section addresses the MP assignment and/or administration of CP Identifiers. It specifically addresses the requirements on the naming authority for CP identifiers and the tools that need to be provided to administer the identifiers, including support for bidirectional mapping between MP IDs and CP IDs (where applicable). The section also includes the source of the identifiers (Network versus MP) and the guidelines for the bidirectional mapping of the application identifiers in the MP to the CP identifiers.

CP-enabled networks must be defined with unique fields of Network ID associated with a Multilayer Sub network (MLSN) per the TMF 814 definition.

NOTE: E-NNI 2.0 Signaling allows logical port ids to exist at an E-NNI interface, i.e., not within the scope of a TNA. Signaling uses the logical port ID since it does not distinguish component link id/te-link id in the same manner as routing. Also, E-NNI 2.0 signaling appears to use logical port ID and interface id interchangeably.

6.1.1 Node ID

The Node ID format supported is currently IPv4. Although standards allow more than just the IPv4 address, the scope of this paper is confined to the agreed format in OIF (IPv4):

- Node ID (each layer [i.e., technology SONET, OTN etc.] has a separate Node ID, i.e., the OTN has a different Node ID than SONET). The reference to layer here also includes different payloads within the same technology, such as ODU1, ODU2, etc., within OTN. CP IDs may correspond to one or more technology layers or routing levels. Links can have one or more interface ID per technology layer or routing level.
- Interface ID pairs as attributes, i.e., the CP edge port is layer specific.

6.1.2 Logical Port ID

The Logical Port ID corresponding to the CP edge port AID (Access Identifier) is, in general, CP generated. Logical Port IDs are always used within the scope of a TNA. Note that this definition differs slightly from the E-NNI signaling definition [3]. If the logical port ID is not assigned by the CP, the MP generates the logical port ID and

maintains the Logical Port ID as an attribute of the CP edge port. Logical Port IDs are stored against both the ingress and egress CP ports. The MP service order process uses the MP constructs for order creation, updates, etc.; however, the MP stores the equivalent mapping of MP Identifiers and CP Identifiers. The Bidirectional ID mapping and the IDs in both MP and CP provides the optional capability for MP operations to view and use IDs in both the MP and the CP. However, this does not imply any real-time mapping between the LP IDs and the ports.

6.1.3 TNA

The TNA is used by CP to identify a single client (a connection end point) or a group of clients served by a transport node; when a TNA is assigned to a group of clients, “Logical Port ID (LPI)” must be specified to identify an individual client (i.e., a port or a data link) when the TNA is applicable for I-NNI and E-NNI routing.

The TNA can be mapped to multiple client ports, or each client port can have many TNAs. Logical Port IDs are unique within a TNA, i.e., we can reuse logical port IDs per Node.

NOTE: G.8080 discusses TNA uniqueness.

6.1.3.1 Valid Formats of TNA Formats: IPv4, IPv6, and NSAP

CP uses TNA (and may use the Logical Port ID with the TNA) in signaling, which the MP stores as an attribute of the CP Port. Every CP-enabled client facility has an associated TNA. If the TNA is assigned to a group of facilities, each CP facility has a Logical Port ID unique per Node ID (where the TNA and Logical Port ID are unique per CP edge port). MP understands the MP name space separate from the CP name space. The structure of the name space in the MP cannot be assumed to be the same. MP needs to store the CP IDs as associated attributes of a CP resource (facility) and/or the CP edge port to facilitate the bidirectional mapping of the IDs. MP needs to perform bidirectional mapping between CP IDs and the MP IDs (such as the TID [Managed Element ID], AID [Facility ID] nomenclature) depending on usage. Every AID (CP edge facility) of a CP facility edge port in the MP has an associated TNA of the appropriate format (IPv4, IPv6, or NSAP address).

6.1.4 Component Link ID (Bundled to a TE Link ID)

The Component Link ID may or may not be associated with the CP edge port.

6.1.5 TE Link ID

All components of a bundled link are incorporated in the MP, i.e., the bundled links are managed as associated (or related) links with the appropriate identifiers for the TE Links. MP system messages (MP) have link bundling modeled and represented in the interface as an explicit entity. The TE Link identifier of the bundle is stored as an attribute of component links and ports that are part of the bundle. Each link terminates on a physical port. A physical port is the termination (or a facility) designated as a CP resource, where each port is associated with a link (i.e., each link has two physical ports associated with it denoting the two end points or terminations).

6.1.5.1 Interface ID (IFINDEX)

IFINDEX is retrieved against each CP-enabled edge port facility and stored as an attribute of the CP edge port facility (TE Link ID). Logical Port ID + TNA mapping to an IFINDEX is maintained in the MP. The IFINDEX for a port is based upon the Port type (I-NNI, ENNI) and AID if the port has an AID (Facility Identifier).

6.1.6 SRLG

SRLG (Shared Risk Link Group) is a CP link attribute that is the union of the SRG identifiers of all resources used to construct the CP link. Since the CP expects to be provided with SRLG information from the MP, it is essential that this information be carefully provided in the MP if it does not already exist in some form.

SRLG (Shared Risk Link Group) attributes are used by the CP to create/manage disjoint or co-routed paths. It may be created as the union of the SRG identifiers of all resources associated with a link. CP expects to be provided with SRLG information from the MP; thus, this information is provided by the MP or the operational staff.

MP may be provided with the SRLG attributes manually by network operator personnel or the MP may include a system (or systems) responsible for creating, assigning, and tracking these attributes. However, it should be noted that each network operator may have his/her own guidance on managing these attributes.

Some guidance for SRLG use is provided in an OIF Carrier Working Group Guideline [18].

In a multi-domain environment, it is desirable to ensure label uniqueness for the SRLG attributes associated with the inter-domain links that support constrained based path computation for multi-domain connections. Therefore, the same SRLG value (or list of values) must be configured by the two EMSs that control the terminations of the E-NNI link. Within each specific E-NNI routing area, the SRLG must be unique (if it is in an independent SRG) regarding the other E-NNIs. If abstraction is used for one I-NNI domain or for one nested domain (with the abstract link model), care must be taken in the creation of the exposed information to avoid reusing the same value as the one used for an existing E-NNI link (if it is in an independent SRG).

Figure 9 depicts an example of an association of multiple SRG labels to an SRLG.




Figure 9: SRLG: SRLG of CP Link A1-A2 = {12, 13, 14}

6.2 Summary of CP IDs – CP Identifiers and Related Name Spaces

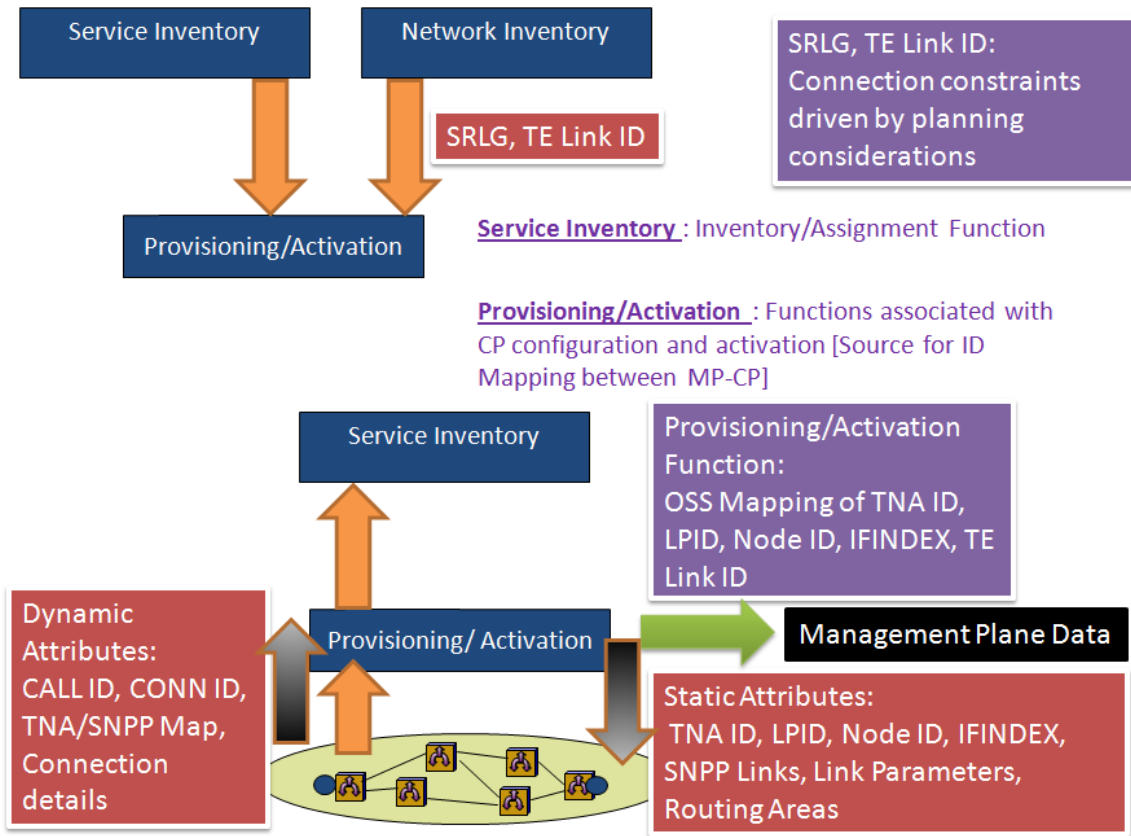


Figure 10: MP Role in CP ID Management

Figure 10 shows the relationships between multiple inventory and provisioning functions managing various attributes and identifiers.

References to the Inventory and the Activation functions in the previous figures present the view of MP-driven attributes versus Network-driven attributes. The MP-driven attributes are connection constraints that require a planning/engineering functional component to define methods and procedures for generating these attributes and feeding data to other systems within the MP. Figure 10 represents the functional components more in the context of OSS building blocks, while Figure 11 represents the MP role in CP ID management from an NMS/EMS perspective. The previous and following figures represent the role of MP in the management of CP IDs represented with a view of the systems, and the functional components.

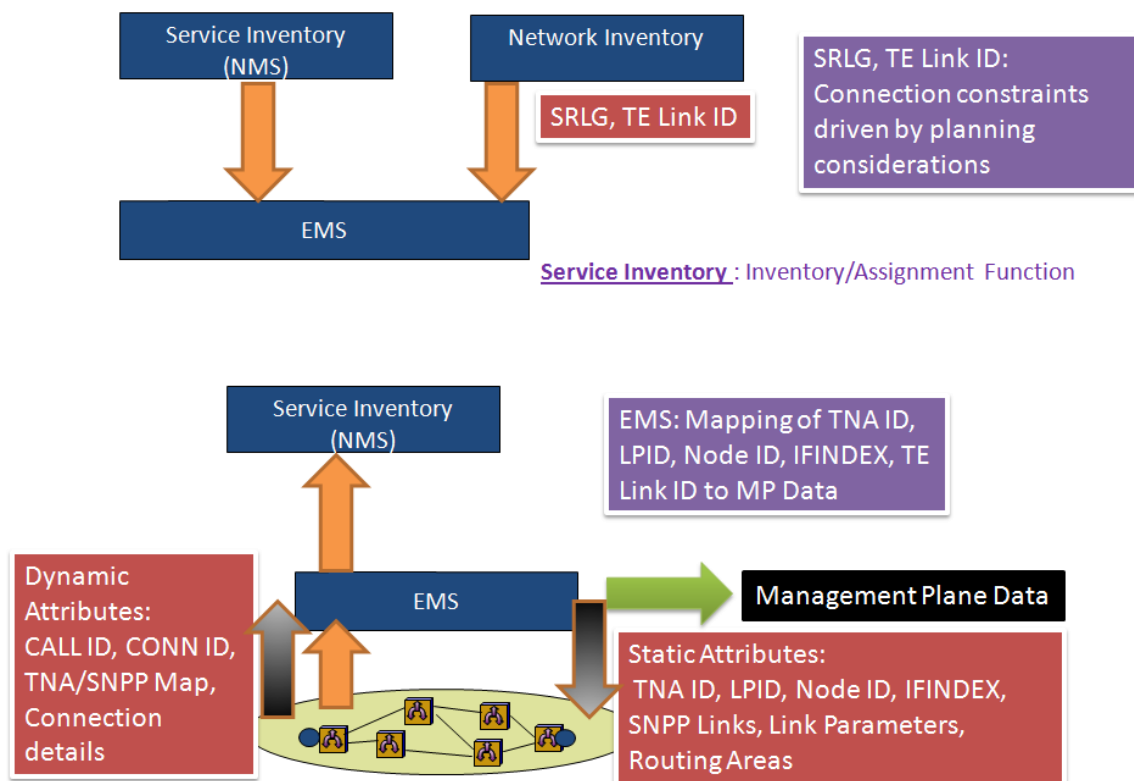


Figure 11: MP Role in CP ID Management – EMS Mediation Model

An MP component must assign, de-assign, and track assignments for TNA. However, once assigned, other systems within the MP may treat the TNA as an NE-assigned attribute. If the LPID (Logical Port ID) is not discovered from the network (i.e., it is CP-generated), the LPID (Logical Port ID) is automatically created with an MP algorithm to ensure unique values for each port of a CP-enabled network. Figure 11 also shows the distinction between static CP attributes provisioned from the MP versus the dynamic CP attributes discovered from the network and subsequently registered in the MP. Management of static and dynamic attributes may not occur within the same time frame. While static CP attributes are established during the planning/engineering phase, dynamic attributes are set/managed during the run phase of the life cycle.

6.3 Management of Routing Controller and Routing Area Hierarchies

After the routing area and the routing hierarchy configuration, the MP interface supports the assignment of CP components to hierarchical routing levels.

Next, MP must support the route initialization function. The operational style of routing (such as distributed/centralized path computation) is a prerequisite for activating the routing controller. Commissioning or initializing the CP includes configuring the routing component, such as the routing controller ID, routing controller IP address, routing area hierarchy, and the routing area IDs. The procedures of local initialization include CP component activation and CP component binding. The MP is responsible for commissioning activities; these could be performed from an EMS or a network element craft interface. CP component activation refers to instantiation of the routing controller

protocol controller by creating a routing protocol object. The routing area is created as an MLRA subnetwork with the routing area identifier. The next step in the commissioning process for which the MP is responsible is configuring the component binding; here, the CP components (RC PC IDs) are bound to each other, i.e., RC adjacencies configuration. The MP interface supports assignment of CP components to routing areas and the routing function. This controls routing protocol neighbor adjacencies and addressing, including routing area hierarchy and subnetwork identifiers. In the next step of the routing initialization process, MP configures the association between routing areas. This is accomplished by providing the appropriate routing area identifiers during the creation of the routing area objects.

Adjacencies between routing controllers and routing controller protocol controllers (RC PC) within a routing area are established to support advertising and maintenance of the available SNPP link topology. In the case of hierarchical routing, adjacencies between lower level RCs and upper level RCs in the corresponding parent routing area are established when the "parent routing area" and "child routing area" parameters are configured on the "routing area". It is expected that carriers manage the assignment of the routing area ID to ensure the routing controller ID is unique within the routing area. Routing control domains are described in relation to components associated with a layer network.

There are four typical reconfiguration operations on routing areas: 1) splitting a routing area into two or more separate routing areas, 2) merging two or more routing areas into one routing area, 3) inserting a routing area into the routing hierarchy, and 4) deleting a routing area from the routing hierarchy. All these reconfigurations involve reconfiguration of area Identifiers and routing controllers. Routing area reconfiguration may require further work at the protocol level.

The MP must be aware of the routing control domain. A routing control domain may, for example, contain components that represent one or more routing areas at one or more levels of aggregation, depending upon the routing method/protocol used throughout the domain. When a routing area is split into two or more routing areas, a group of links within that routing area must be split into multiple routing areas, assuming the control domain does not change.

Routing Control Domain (RCD) is a type of control domain where the criterion for membership is a common routing capability containing zero or more rerouting domains. Routing control domains are derived from routing controller components, and a rerouting domain is derived from a set of connection controller and network call controller components that share responsibility for the rerouting/restoration of connections/calls traversing that domain. In both examples, the operation that occurs (either routing or rerouting) is entirely contained within the routing control domain. In essence, when a routing area is split into multiple routing areas, it needs to be performed with the knowledge of the routing control domain from an MP perspective. A routing control domain can have multiple routing areas.

A rerouting (recovery) domain is defined as a type of routing control domain whose control components at the edge of the domain coordinate rerouting operations for all calls/connections traversing the rerouting domain. In the OIF, rerouting domain is referred to as the recovery domain. Recovery includes restoration and protection.

Therefore, network resources associated with a rerouting domain must be contained entirely within a routing area. So, if a routing area split occurs, the two or more routing areas from the original routing area each constitute a rerouting domain. Further, MP must update the recovery domains based on the routing area reconfiguration.

If two or more routing areas are merged into a single routing area, this represents expansion of multiple routing domains into a single routing domain/rerouting domain. If a routing area is added to the routing hierarchy, this becomes another rerouting domain from an MP perspective.

6.4 SCN and Control Channel (CC) Management, Configuration, and Security

This section discusses MP support for the CP SCN network.

In-Fiber SCN configurations include the configuration of SCN node address and identifier, SCN link configuration (e.g., DCC, Address configuration) and control channel, SCN Protocol stack configuration, and static routes [resiliency of control channels, if any]. An Out-of-Fiber SCN involves an External IP-SCN or LAN configuration. If in-fiber, the overhead bytes (control channels - DCC, GCC) or OSC are used. There is technology-dependent provisioning of protocol stack, encapsulation, and security. If using in-band techniques, such as DCC or GCC, these bytes may be terminated and need to be remapped at domain boundaries; this may require provisioning from the MP. If using out-of-band techniques, an address resolution (or gateway function) may need to be configured to propagate SCN traffic across the boundaries. In a multi-layer CP network, if an intermediate server layer domain and its SCN need to support a gateway function (unique to OIF/ASON multi-layer model) for forwarding client layer CP messages, this may require provisioning. SCN configuration includes:

- Provisioning these addresses on the SCN per node. This would include provisioning SCN addresses for SC PC and RC PC.
- Configuring associations between SCN addresses and protocol controllers reachable via that address (SCN adjacency)
- Configuring protocol stack, encapsulation, and security mechanisms for SCN traffic (SCN technology-dependent)
- Monitoring SCN for traffic congestion
- Requiring SCN to detect and report congestion conditions
- Monitoring SCN for failures/degradations and switchover to backup resources (if any).

The MP must support the ability to receive a notification for a failure/degradation in SCN. It must also support the ability to receive a notification of SCN congestion. SCN is similar to the Transport network and failures/degradations in the Transport network. MP may support the ability to correlate the failed SCN (degraded). The methodology for the MP to perform this correlation requires further analysis. MP needs to process the Signaling adjacency failures/degradations from a fault management perspective. Signaling and routing adjacency establishments and failures will be reported back to the

MP. Health check related notifications are supported on signaling and routing adjacency failures/degradations.

6.5 MP in Capacity Planning

For traditional networks without CP capabilities, the MP has traditionally tracked capacity of all resources (e.g., links, nodes, and specific topologies) established when the infrastructure was created. The MP also tracked the network utilization for all MP services, (i.e., using spare capacity is an inherent part of service assignment). In addition, the MP tracked the percent utilization of all links in a network, and also managed a threshold bandwidth per link for all MP networks. For CP-driven networks, this functionality must be coordinated between the MP and the CP.

The MP continues to support functions associated with running the business. For legacy domains, it controls all assignments; however, for CP-driven domains, since the MP may not necessarily have up-to-date information, it may need to receive (or request) information (as appropriate) from the CP. For example, this may be performed via queries (scheduled or on-demand) or via autonomous notifications from the CP. Alternately, the MP could monitor routing controller advertisements for capacity planning purposes.

The MP is responsible for generating reports. For example, these reports may contain information obtained from the CP and may be:

- Per Link
- Per Node (Only if it is a blocking Node)
- Per Module (if blocking)
- Per Domain
- Per Connection
- Per Client

Further, for networks with resource control split between MP and CP, the MP is responsible for providing reports to operator personnel containing up-to-date status and details of resource utilization. These reports are optional and network provider driven.

The CP is aware of the resources used by MP and takes this into account during bandwidth accounting (and advertised by OSPF-TE for instance). Carriers may prefer to retrieve this data directly from the CP.

In addition, if a bandwidth threshold per link is specified at network creation time, the MP also provides reports with data on links that have used bandwidth in excess of the thresholds. Such thresholds may be used to help assess necessary network augments.

Capacity planning for multi-domain networks may be integrated under one system, receiving (or owning) pertinent data from each of the domains; alternately, separate systems may be responsible for subsets of such domains. In the latter case, inter-domain capacity may be coordinated according to prevailing operator policy.

The need for capacity and utilization monitoring is also driven by the need to account for facilities required to mitigate potential network failures (including capacity needed to support diversity and protected connections and restoration).

Hence, capacity planning systems need to incorporate functionality to support restoration needs, while considering diversity criteria such as Shared Risk Link Groups (SRLG) regarding links and nodes. However, it is expected that individual operators would establish specific criteria to be used within their networks.

6.6 MP Management of Service Recovery²

With the introduction of CP-based recovery mechanisms (and potentially, various grades of services for CP enabled services) the MP must incorporate the functionality necessary to support these capabilities. This section discusses some of this functionality.

One essential function is the capability to provision recoverable services (e.g., establishing protected and/or CP-restorable services). This includes support for the following:

- a) Capability to provide (or request from the CP) support for optional routing constraints for restoration of the service (constraints for calculation of restoration route, e.g., via use of SRLGs):
 - i. Inclusion support (Strict or Loose)
 - ii. Exclusion support
 - iii. Diversity
 - iv. Latency
- b) Capability to provision (or request from the CP) support for a restoration priority level for the service (e.g., it can be chosen based on the SLA for the service, and used by the CP to prioritize/reconcile restoration demands).
- c) Capability to optionally provision pre-emptor ability for the service restoration (i.e., to optionally allow the service to be a pre-emptor on restoration)³.
- d) Capability to retrieve and make changes (or request the CP to support) the above restoration parameters for the service.
- e) Capability to configure and reconfigure the Service Level of a service/call (i.e., capability to configure/reconfigure restorability and/or protection support for a service/call)⁴.

Also important is the capability to provision services that employ a combination of protection and restoration mechanisms (e.g., protected services which, in turn, leverage secondary restoration for the working and/or protect legs)⁵. This includes support for the following:

² i.e., MP management of service level (including restoration and/or protection).

³ This is dependent on whether pre-emptible services are applicable in the service provider's network and service portfolio. If applicable, then the capability to optionally configure pre-emptibility of a service should also be supported.

⁴ e.g., capability to reconfigure a non-restorable service to be restorable, capability to reconfigure a restorable service to be non-restorable, capability to add/remove protection support for a service, etc., as applicable to the Service Level configured for the service/call.

⁵ Based on the Service Level provisioned for the call.

- a) Capability to provision the precedence of the mechanisms (e.g., to provision the precedence of 1:N [shared] protection versus restoration). For example, this can include the capability to provision hold-off timers.
- b) Capability to retrieve and make changes to the above precedence
- c) Capability to retrieve and change the service levels as examples, based on the reconfigured service level, this can:
 - Add or remove protection from the service
 - Add or remove restoration from the service.
- d) Capability to provision Dual Homing (single service with working and protect diverse paths with separate CP calls).

In addition to retrieval of the parameters/attributes specified above for the support of restorable service, the following need to be retrievable from the CP:

- a) An indication of whether the service is restorable or not
- b) The actual route in use by a restorable service (whether working route or restoration route)
- c) The working (preferred) route of the service
- d) The status of restoration (e.g., on-working, on-restoration, in-progress/pending)
 - i. Route query (“what-if”) for a service, e.g., regardless of the current route in use by the service (on working or on restoration), query to see if diverse restoration (diverse alternate route) is possible (e.g., route query for a diverse alternate route/restoration route [based on current network status and based on imposed restoration route constraints, if any]).

The MP also needs to support the capability to provision reversion parameters for recoverable services. This includes support for the following:

- a) Provision/retrieve/change - whether or not auto-reversion is used
- b) If revertive, provision/retrieve/change WTR Time
- c) Retrieve the amount of time remaining in the (active) WTR timer (following a successful restoration action or protection action).

To support data consistency between the CP and the MP with regard to the routes of the connections for CP-managed restorable service, it may be desirable that the CP provides the MP with information related to CP-initiated notifications/events associated with restoration actions. Subject to performance needs and individual operations processes, this information may be provided by the CP to the MP as autonomous notifications, either on a scheduled basis or on demand. The information includes support for notifications/events for the following CP restoration actions:

- a) Service Restoration in Progress (Restoration Initiated/Invoked)
 - i. Service Restored to Restoration Route (Restoration Successfully Completed) (optionally, to include current route/resources)
- b) Service Unable to Restore (Restoration Failed [No Restoration Route Found])

c) Service Reverted to Working (Preferred) Route.

The MP plays an important role in providing historical data beyond that supplied by the NEs. As such, the MP may be required to maintain a history (log) of CP restoration actions; it may also (potentially) be required to maintain historical data on the route changes (including specific timelines) for individual connections.

The MP also supports operational processes by providing various reporting capabilities. For restoration-related events, the reports should reflect the CP restoration status and restoration events, plus the impact of the individual connections. This should include support for reports by Service Levels⁶ and the following (this does not imply query):

- a) Services (Restorable Services) on a Restoration Route (optionally, provide indication of whether the service is on the preferred restoration route [if a preferred route was provisioned] or on a ‘best-effort’ route [e.g., preferred route may have been diverse from another service route, or may have been for a low-latency route]).

b) Services (Restorable Services) with Restoration Attempt In Progress

Services (Restorable Services) Unable to Restore (optionally, including indication of the root-cause for the service being unable to restore, e.g., 1) Route Not Found (with No Constraints Applied), 2) Route Not Found Meeting Constraints (Constraints Can’t Be Met), 3) Explicit (Strict) Route Not Available, and 4) CP Not Available).

c) Services (Restorable and Non-Restorable Services) on Working/Preferred Route

Services on Restoration Route with WTR Timer running (revertive services). The MP shall support the capability to provide performance reports for restorable services. This shall include support for Service Availability (service “up-time”). Optionally, correlation of availability/performance of a service to the route/resources is used by the service at the time.

6.7 Security Considerations for OSS Management of CP

OIF provides security IAs that define security capabilities for the CP [15], MP and logging [16], and auditing [17]. All these security mechanisms are optional. This section identifies the security functions that may need to be addressed in the MP depending on the security capabilities deployed.

CP interfaces at the UNI and E-NNI implementing the Security Extension may require MP support for:

- a) Equipping each SC PC ID (Signaling Controller Protocol Controller Identifier) and RC PC ID (Routing Controller Protocol Controller Identifier) with signature keys and certificates (or with pair wise pre-shared secrets) for each such entity. These certificates (or secrets) are used by the IKE (Internet Key Exchange) v2 protocol to establish IP sec security associations. SC PC ID and RC PC ID are names used to identify the entity across the secured signaling or routing interface, respectively.

⁶ For example, Service Level may correspond (directly) to service provider SLA or restoration priority level, or may be chosen (indirectly) based on requested SLA.

- b) Providing resources to check the certificate with the Online Certificate Status Protocol (OCSP). When certificates are used, this capability allows them to be checked for validity.
- c) Supporting initial security policy configuration and updates: IP sec enforces policy on the traffic protected and how it is protected. The MP may provide tools for this. For example, users may have a blanket security policy for all CP interfaces or specific, fine grained policies that depend on addresses, services, times of day, or similar variables.

CP-MP interfaces protected with any of the methods specified in the reference document [16] may also need security configuration. The specific need depends on the choices made for securing these interfaces. Details are specified in the MP security IA [16].

The OIF has an implementation agreement for CP logging and auditing [17] that is used in the various signaling, routing, and security implementation agreements. Implementations using OIFs logging and auditing [17] may need MP support. This may require that the MP provide a secure channel for receiving log messages from the CP and a secure repository for these messages. Additional MP functionality for log file verification and analysis may also be needed.

6.8 Service Management across (Heterogeneous) Domains: Administration of Cross OSS CP Connection: CP Connection Spanning Multiple MP/Servers

The objective of this section is to address operational challenges associated with CP services spanning multiple administrative NMS/EMS domains where control domains are under the management of different NMS/EMS instances. It depicts various scenarios showing service support across multiple control domains and discusses MP support for these scenarios.

Network operators provide services that span multiple business organizations within the operator's network. While provisioning, protection, and restoration requirements are critical within a CP domain, these requirements cannot be compromised in a multi-domain network. Service requirements that span multiple business organizations include:

- All advantages of a CP can be realized across MP for these converged networks
- End-to-End CP can be supported
- Restoration for end-to-end CP is supported given that an E-NNI can help support this convergence
- A common strategy for Carriers is to use separate MP for these networks
- Special MP strategies must be used to realize the CP value in these cases
- MPs support these converged networks with minimum impact

6.8.1 Services in the Management Plane Spanning Multiple Administrative Domains (Provisioning and Path Restoration)

6.8.1.1 Services in the Management Plane




Figure 12: MP Service Spanning Multiple Domains

Figure 12 shows a traditional network without CP NNI as a Network to Network traditional Handoff (NNTH) with no CP capabilities at the handoff between the networks. A traditional service (Management Plane) across MP business organizations shown in this figure is accomplished with three separate orders within each domain (Circuits 1, 2, and 3 in the diagram), along with 4 and 5 (the handoff points between the intra-carrier business organizations).




Figure 13: End-to-End Service with Three Different CPs and Management Plane Handoffs

Figure 13 illustrates three control domains interconnected by MP handoffs. NNTH is a Network-to-Network traditional Handoff (referred to as NNI in the figure) with no CP interfaces at the handoff between the networks. The end-to-end service is accomplished with three separate CPs (one within each business organization) with the handoffs being traditional MP connections. This is an end-to-end service, but not an end-to-end CP and presents the view where each boundary of MP manages the CP call within its administrative domain. In this case, no data is required to be shared between MP other than the facility within each domain for the inter-domain handoffs. This scenario consists of three independent CP calls within each domain and could be executed without any sequence. The interconnection between domains is an NNI- Network-to-Network Interface (traditional MP handoff) and is NOT a CP connection. These are concatenated calls with multiple independent signaling sequences, i.e., the objective is to not have an CP call span MP boundaries, but rather to confine the call within each administrative domain. This approach for intra-carrier CP services spanning business organizations conforms to the regulatory and customer billing issues associated with these services.

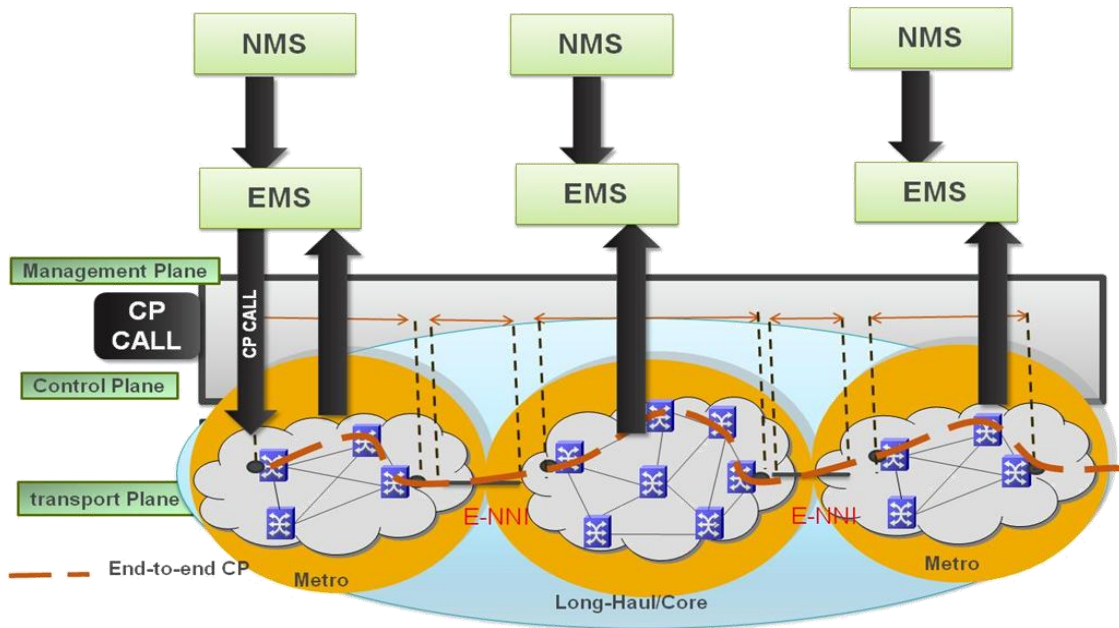


Figure 14: End-to-end CP

Figure 14 illustrates an end-to-end CP service with one CP across MP business organizations. The data contained in the CP, i.e., data advertised over the E-NNI, is retrieved from the network by the EMS managing that domain. Such data does not need to be exchanged over the MP interfaces. Data that needs to be exchanged between MP domains (such as different business organizations not exchanged across E-NNI) may be exchanged through various means, such as system interfaces or business process. These may be carrier-specific methods and procedures to share service level data for an end-to-end service across the different administrative domains per the established service ordering and billing process.

NOTE: Support for services spanning multiple domains is not a new issue introduced by the CP.

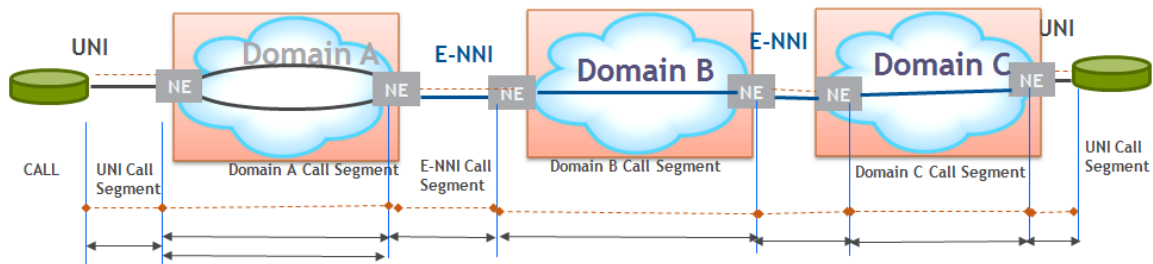


Figure 15: Configuration of an End-to-End Call with Multiple Call Segments and Connection Segments within Call Segments

The CP distributes TNA reachability using routing protocols. Hence, data such as the destination TNA may be sourced by the MP from the local CP. Processing of an end-to-end CP call is triggered in the MP by a service order. This triggers the record of an end-to-end connection in the inventory system for which the path needs to be updated. Separate MPs must have knowledge of all the established connections, including restoration paths made by the CP. In the case of a service spanning multiple domains and MP boundaries, this would mean that both MPs should be aware of the end-to-end connections and path restorations.

6.8.2 Other Data Attributes between NMSs in an End-to-End CP

6.8.2.1 SPC Call/connection Attributes from UNI 2.0 [2]

For an SPC, the MP needs to support the attributes listed in Table 2 to initiate the CP call.

Table 2: Call and Connection Attributes⁷

Attributes
Source Logical Port Identifier (M)
Source Generalized Label (O)
Destination TNA Name (M)
Destination Logical Port Identifier (M)
Destination Generalized Label (O)
Contract ID (O)
Source TNA Name (M)
Switching Type (M)
SONET/SDH, OTN or Ethernet traffic Parameters (M)
Directionality (O)
Generalized Payload Identifier (O)
Service Level (M-1, O-2)
Diversity (M)
Bandwidth Modification Support(M)

For example, service level attributes are tied to the SLA and the restoration/recovery strategy. The MP is responsible for supporting the business aspect and translating into the appropriate network operational strategy, such as recovery. Support for the attributes needs to be addressed by the MP including appropriate inventory, assignment, and activation capabilities.

Figure 16 shows the various components involved in path retrieval.

⁷ M-Mandatory, O-Optional

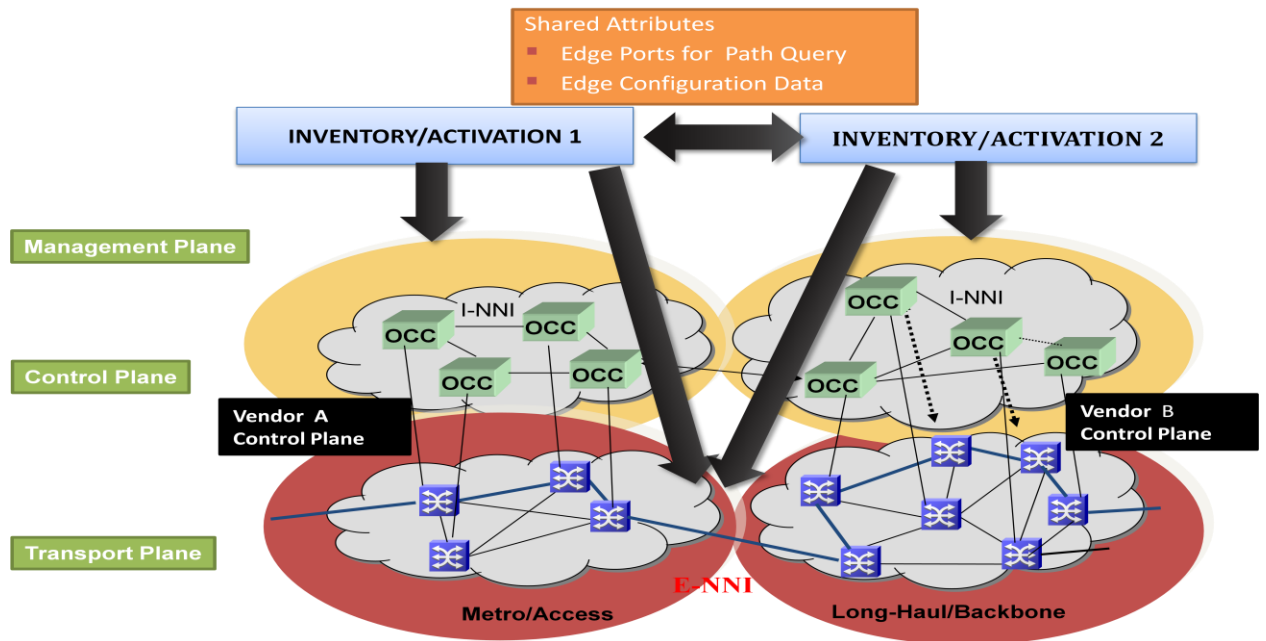


Figure 16: Path Retrieval Problem Definition

Path routing retrieval capability refers to the ability to retrieve the path end to end across MP domains. The path query from the source end provides the end-to-end path, assuming there are no SCN implications. An end-to-end CP call in this scenario can be achieved by multiple CP Call segments and with call segments divided into connection segments within each domain. The end-to-end path route information retrieval can be accomplished with a thin layer above the MP suite with discoveries performed per administrative domain as an alternative and the end-to-end path maintained and stored in a thin layer (i.e., it could be a system that stores the end-to-end path) with the overall view of the service route.

6.8.3 E2E (End-to-End OAM) for Services Spanning Multiple Domains

MP will interact with the data plane and/or CP to enforce operator commands at the end points of each domain of an end-to-end service crossing multiple domains. The critical requirement for the carriers on these services spanning multiple administrative domains is the need for end-to-end OAM, which will be enforced by operator commands initiated by the NMS/EMS/OSS across the boundaries of each domain.

6.9 Transition Support

6.9.1 Management Plane to CP Transition

It cannot be guaranteed that CP is always deployed in a green field environment. There may be situations where a service provider has a deployed network managed by traditional management systems and intends to upgrade this network to support a CP. For brown field applications such as this, it is necessary for the network elements and management systems to support the ability to transition a network from MP to CP.

It is also possible that the service provider has deployed a CP, but there remain portions of the network under MP control. The service provider may later decide to upgrade more

of his network to CP. This scenario would require the ability to expand the extent of a CP subnetwork.

Services are typically migrated to CP as part of the process of converting the network to CP. The service provider may choose to convert all existing services to CP or perhaps convert only a subset of the services. It is also possible that the service provider may choose not to transition any services and leave them under MP control.

Some CP implementations require setting a link attribute to indicate which entity is in control of services over that link. Possible settings include:


- MP-only. Only MP services can exist on the link.
- CP-only. The link only supports CP services. During transition from MP to CP, all services on this link must be transitioned to CP.
- Hybrid-MP-CP. The link supports both MP and CP services. (Note that under the hybrid MP/CP control, there may be a race condition where the MP and CP attempt to assign the same resource.)

The transition from MP to CP must not affect traffic. The conversion of the network should retain all existing services and not change the transport resources supporting these services. The transition should only be a logical change in the entity (i.e., MP or CP) responsible for the control of these resources.


Following transition, new services are typically established by CP over the transitioned subnetwork. However, it is also possible that the subnetwork can establish both CP and traditional MP services (e.g., if the subnetwork links are designated as hybrid MP/CP links).

There are several MP-to-CP transition use cases as shown in Figure 17:


- Case 1: The transition of an MP subnetwork to CP
- Case 2: The extension of an existing CP subnetwork
- Case 3: The complete service transition of a subnetwork to CP where all services within the subnetwork are transferred to the CP
- Case 4: The partial service transition of a subnetwork to CP where a subset (perhaps an empty subset) of the services are transferred to the CP. In this scenario, both MP and CP services can coexist.



Case 1: MP-to-CP Conversion of a Subnetwork




Case 2: Extension of an Existing CP Subnetwork



Case 3: Full Service Transition to CP

Links contain both MP and CP services



Case 4: Partial transition to CP Subnetwork

Figure 17: MP-to-CP Transition Use Cases

There are several different ways a network may be transitioned to CP; the following procedure is one example. Specific procedures would need to be developed to take into account vendor-specific NMS, EMS, and CP implementations.

The transition from MP to CP can be divided into three phases: NMS preparation, network transition, and NMS discovery and finalization.

Phase 1: NMS Preparation

The first step is that the service provider identifies the subnetwork being transitioned to CP. This subnetwork is defined by a set of links. Note that on a network element, some links may be marked for transition, while other links may be unaffected and remain under MP control. For implementations that identify the controlling entity for links, the set of link defining the subnetwork should be configured as either CP-only or hybrid CP/MP.

The service provider would mark this subnetwork within the NMS application. With the help of the NMS, the service provider may want to ensure that there are no pending service orders that would use resources within the subnetwork. The transition should not be started until all pending orders are completed.

Once all pending orders have been completed, the service provider would put the subnetwork into the MP-to-CP transition state. The service provider would also indicate if this is to be a full service transition or a partial service transition.

Once in the MP-to-CP transition state, the NMS may deny new service orders that would use resources within the subnetwork being transitioned.

Note that other order processing methods may be supported. It may be possible that pending orders are put on hold during the transition and resumed once the transition completes. It is also possible that orders may continue to be processed while the network is being transitioned, especially under the partial service transition use case.

Phase 2: Network Transition

The next phase is to transition the network to CP. One of the goals of the network transition phase is to minimize the amount of downtime for the subnetwork. Given this objective, it would be desirable to perform as much of the activities in this phase in parallel with (or prior to) locking the subnetwork in Phase 1.

The service provider must configure the CP infrastructure on all network elements affected by the transition. This includes assigning CP addresses and identifiers, provisioning protocol controllers, configuring parameters, establishing the SCN, and enabling CP. Links are also configured (if necessary) to indicate if they are controlled by CP or by both the control and MPs (i.e., hybrid mode). The EMS may also be configured to support the CP subnetwork.

The service provider would also configure security for the CP. This may include configuring CP protocol security, establishing protocol security between the management systems and CP, and establishing the CP logging and auditing infrastructure. Different security requirements may be applicable for the case of a full conversion to CP versus a partial transition.

Phase 3 Service Transition

The service provider then initiates the MP-to-CP conversion of the subnetwork within a management system. Note that the network transition may be supported by a different management application (e.g., initiated from an EMS, craft, or supplemental tool). The service provider would indicate if this is to be a full or a partial service transition. In the case of partial service transition, the service provider would also indicate which, if any, services are to be transitioned. Note that a given service can be a mix of MP and CP, i.e., a circuit end to end can be routed by the MP or CP (per segment).

Transition of services from the MP to the CP is performed on a per circuit basis. The management system initiates the MP-to-CP conversion and instructs the CP to transition each individual circuit.

The CP transitions the service within the bounds of the subnetwork. Noting that the service may extend beyond the boundaries of the subnetwork after the transition, the end-to-end service may be under combined MP/CP control.

The management system provides a success indication on a per service basis and at the completion of all services. The management system also provides an error indication if any of the services failed to be transitioned.

Phase 4: NMS Discovery and Finalization

The fourth phase is to finalize the transition in the NMS. The service provider transitions the MP-to-CP operation to be completed within the NMS. If the MP-to-CP network transition was initiated via the NMS, then the completion notification may be automated.

The NMS then discovers the CP infrastructure within the subnetwork. The NMS may also discover the services that have been transitioned to CP, if necessary.

Once completed, the NMS allows new service orders to proceed. If the transition was a full service transition, then typically all future service orders are established by CP within the subnetwork. However, it is possible for the subnetwork to support both CP and MP services, especially in the case of a partial service transition.

Under the partial service transition case, or if the subnetwork supports hybrid MP and CP links, a service provider may need to transition additional services from MP to CP in the future. The network elements and management systems should support this ability to transition services outside the subnetwork transition procedure.

6.9.1.1 CP to Management Plane Transition

There may also be a need to support the transition of a network from CP to MP. For instance, a service provider may want to revert back to an MP network if issues were discovered during the MP-to-CP transition.

The transition to MP could be for the entire CP subnetwork (i.e., the entire network is reverted back to MP) or only a portion of the CP subnetwork. In the later case, this would be equivalent to the reduction of the extent of the CP subnetwork.

Services would also be migrated to MP. All services over the affected region must be transitioned to MP since CP services are not supported over an MP network.

As in the previous section, the transition of services to MP must not affect traffic. All services would be retained and service routes remain unchanged. The transition would only reassign control of the transport resources to MP.

Following transition, new services must be established by MP across the MP network.

There are two MP transition use cases as shown in Figure 18.

- Case 5: Transition of a CP subnetwork to MP
- Case 6: Reduction of the extent of an existing CP subnetwork

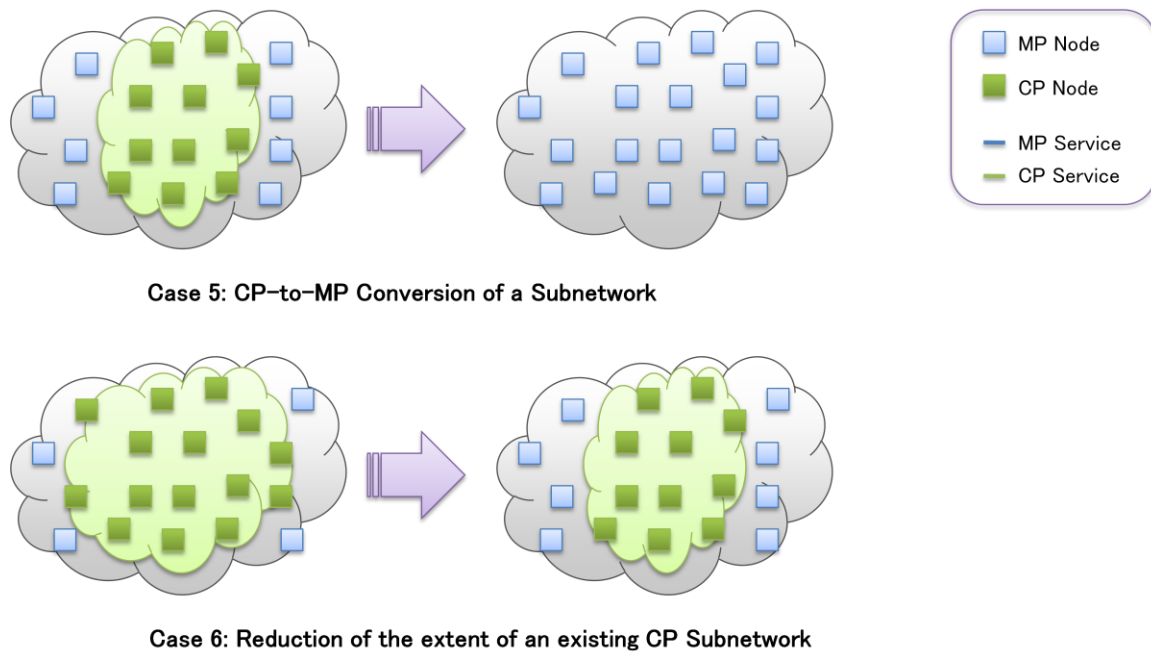


Figure 18: CP-to-MP Transition Use Cases

The following procedure is an example of transitioning a subnetwork from CP to MP. As in the MP-to-CP case, specific procedures would need to be developed to take into account vendor-specific OSS, NMS, EMS, and CP implementations.

The transition from CP to MP is divided into the same four phases as the MP-to-CP transition, i.e., NMS preparation, network transition, service transition, and NMS discovery and finalization.

Phase 1: NMS Preparation

The service provider must identify the subnetwork to be transitioned to MP. This subnetwork is defined by a set of links. For implementations that identify the controlling entity for links, the set of links defining the subnetwork would be set to MP-only at the end of this procedure. The service provider would mark this subnetwork within the NMS application.

The service provider should also ensure that there are no pending orders that could be routed over the subnetwork. The service provider should also put the CP in a state where no new services would be allowed.

The procedure should not continue until all pending orders have completed. This may be a manual operation or automated by the NMS.

Once all pending orders have been resolved, the service provider puts the subnetwork into the CP-to-MP transition state. For this operation, all services must be transitioned to MP. The NMS denies any new service orders over this subnetwork while in the CP-to-MP transition state.

As in the MP-to-CP case, other order processing methods may be supported. It may be possible that pending orders are put on hold during the transition and resumed once the transition completes. It is also possible that orders may continue to be processed while the network is being transitioned, especially if the orders were being completely supported by MP.

Phase 2: Network Transition

The next phase is to transition the network to MP. The service provider initiates the CP-to-MP transition of the subnetwork within a management system (e.g., from an NMS, EMS, craft, or supplemental tool). For each service that runs over the subnetwork, the management system instructs the CP to transition the service to MP.

Phase 3: Services Transition

Services are transitioned to MP during this phase.

The CP transitions the service to MP within the bounds of the subnetwork that is transitioning. Note that the transition to MP may not be end-to-end (i.e., it may be a reduction of the CP subnetwork as described in use case 6). In this case, the service would be a hybrid MP/CP service after the transition has completed.

The management system provides a success indication to the service provider once all services have been transitioned. Management systems also provide an error indication if any service transitions failed.

During this phase, links in the subnetwork should be configured to indicate that they are now controlled by MP.

The service provider can now deprovision the CP infrastructure within the subnetwork.

Phase 4: NMS Finalization

The next phase is to finalize the transition in the NMS. The service provider transitions the CP-to-MP operation to “complete” within the NMS. If the network transition was initiated via the NMS, the completion notification may be automated.

The NMS may discover the MP infrastructure and services, if necessary.

Once discovery is complete, the NMS allows new service orders to proceed. All new services using resources within the subnetwork use the MP to provision those resources.

6.9.1.2 Hybrid Services

Services are routed through a mix of MP and CP enabled networks, i.e., services need not originate and terminate in CP enabled networks.




Figure 19: CP Mixed Network Routing Configuration

The routing configuration in Figure 19 illustrates an end-to-end CP service that originates on a traditional network, rides a CP enabled network, and then terminates on an MP network with traditional links between the different networks. This routing configuration is a representative example of an end-to-end service across MP and CP networks.

In this configuration, the network supports a mix of MP and CP services. If there is any spare bandwidth in a network supporting MP services (traditional services), that remaining bandwidth could be used for CP services. In essence, there could be a mixture of MP and CP services in a given network. As such, the MP process for transition is triggered on a per-service basis. However, if existing services need to continue as MP, and new services using the spare bandwidth must be managed by the CP, the generic MLSN network model adopted by TMF supports such a capability. This is not viewed as a per-service transition; however, this is the case of supporting mixed services (MP, CP) within the same network.

The MP will provide the user with the capability to define the boundaries of the management-plane subnetwork to be transitioned to CP. This subnetwork is defined as the “MP-CP transition subnetwork”. The MP-CP transition subnetwork must be enabled to support CP prior to the execution of the transition operation, including the assignment of CP identifiers and provisioning of CP components. The MP shall identify the MP circuits (or portion of an MP circuit) to be transitioned to CP within the boundaries of the transition sub-network. For each MP circuit, the MP identifies the endpoints within the MP-CP transition subnetwork and initiates the transition operation at one of the endpoint nodes as shown in Figure 19. The capability to transition from an MP Network to a CP network (or vice versa) must be supported (per the use cases specified earlier in the section). An end-to-end service over MP networks could be routed over a CP network if the MP networks are transitioned to a CP-enabled network. Alternately, a service routed over CP networks can be routed over MP networks if the CP networks are transitioned to MP networks.

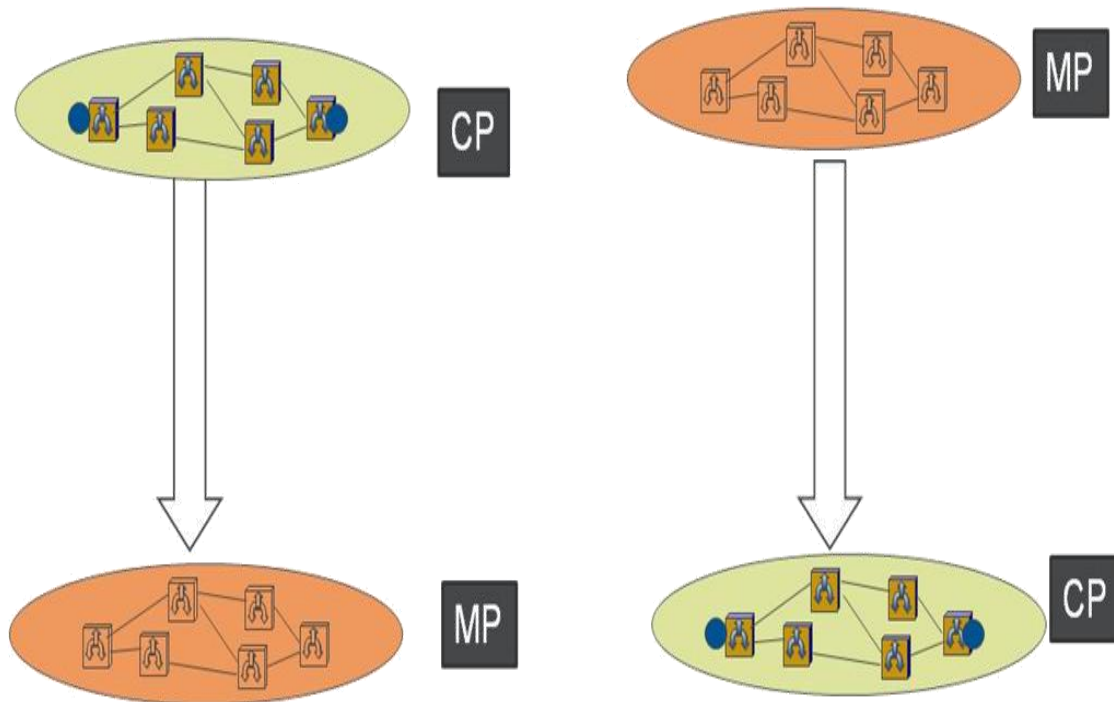


Figure 20: Transition Scenarios

6.9.1.3 Multilayer Support

While the intent is to document the procedures associated with the NE to perform this network transition, MP functionality in supporting this transition is crucial to a successful transition. There is more complexity in the multi-layer scenario. This multilayer scenario is outside the scope of this whitepaper.

6.9.1.4 Routing Configurations and Other Transition Considerations

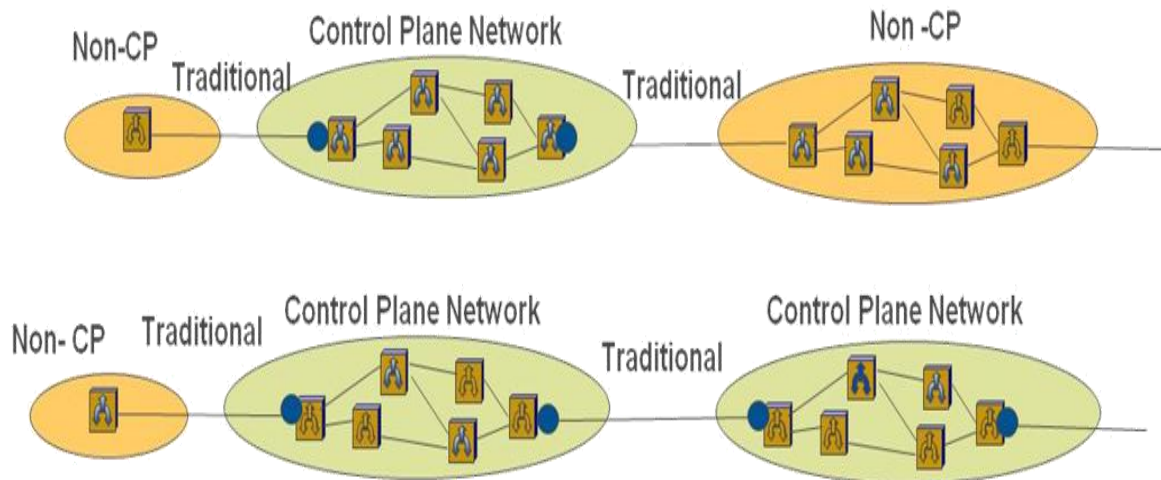


Figure 21: Routing Configurations

Figure 21 shows a typical network/service scenario with a mix of MP and CP. The MP contains the complete inventory of CP networks, facilities, and resources. This may include the optional view of the MP as a database of record (alternately, the network could be the database of record) with an updated real-time view of all CP resources, including available bandwidth and assigned bandwidth. The MP has the capability to provide the CP with an explicit manual route based on planning criteria driven by user input (for example, node or link levels diversity constraints or any other path optimization criteria). Note that this could mean integration with any vendor tools, the output of which would drive the MP-driven CP route. The MP supports the ability to support network queries and provides reports on specific data to end users on CP resources (for example, CP edge port bandwidth availability if the edge port has been used for traditional MP services or for CP services, assuming the MP keeps track of the bandwidth in the network).

There is not an absolute limit on the number of messages per unit time from the CP. However, the volume of messages does cause MP performance to deteriorate.

The following steps are involved in the transition of an MP Network to a CP Network:

- Steps to trigger the NE for such a transition from MP to CP (IETF RFC addresses this transition). NOTE: Signaling extensions may be required.
- For all Equipment and Cross-connections in the end-to-end path of an MP circuit, the CP supports the capability to query the call and the connections after the transition (MP→CP).
- The MP database is updated to reflect all MP connections to a CP connection at the source node. This involves the connections and the provisioning records at the pass-through nodes made in the MP to be erased and replaced, reflected as a CP record based on the responses from the CP to the path query.
- Upon successful completion of the transition in the network, all MP resources become controlled by the CP (refer to the transition process described in RFC 5493).

- It is expected that the query response returns data associated with the MP, such as the original connection before the transition.
- The MP updates its database and reflects the circuit provisioning-related data as a CP route. The MP also updates the traditional networks on which the original circuit was provisioned as a CP Network by assigning an MLRA ID to the network. This includes all networks in the end-to-end route of the CP.
- The above transition steps are performed In-effect; hence, the only network interaction (in addition to processing the autonomous messages) is associated with the network queries triggered by autonomous network notifications from the NE. A typical MP process is to synchronize its database with the Network Element database. However, an MP to CP transition should trigger an autonomous message (Event) with the identification of the head-end node Edge port data (which the NE tracks as the source of the CP call origination).
- The event or the autonomous message indicative of an MP-CP transition having taken place needs to be specific (i.e., unique from any other event notifications indicative of other conditions). NOTE: Events may not be specific.
- The transition process in the MP is performed once the network has been transitioned, i.e., when the network is CP-ready. The transition process in the MP is triggered on a per-service basis, i.e., updates to provisioning records in MP databases are for all services riding on the transitioned network on a per-circuit basis until all circuits are transitioned. Once the transition process has been completed, it is expected that all new services on the network will be provisioned and managed by the CP.

The following steps are involved in the transition of a CP Network to a Management Plane Network:

- Per Network Transition of connections implies that a network transition would initiate the MP to trigger all CALLs (each call serially) associated with the CP network transitioning to an MP connection (i.e., upon successful transition, all services on that network are managed by the MP).
- Assumption: end points are within the same routing area and MP Admin domain.
- For CP connections transitioned to an MP, CP connections at the head-end node (call origination) are transitioned to explicit connections end-to-end across the networks, including pass-through nodes.
- The MP database is updated to reflect all MP connections. This involves connections and provisioning records at the pass-through nodes made in the MP based on responses to the path query from the network element.
- It is expected that the query response returns data associated with the MP (after the transition), such as the original connection before the transition and/or an indication of this being a transition from a CP circuit to the MP.
- The MP updates its database and reflects the circuit provisioning-related data as an MP circuit. The MP also updates the networks on which the original circuit was provisioned by the CP. This includes all the networks in the end-to-end route, which is updated with a traditional network ID from an MLRA ID.

- The above transition steps are performed “In-effect”; hence, the only network interaction is associated with the network queries. A typical MP process is to synchronize its database with the Network Element database. However, a CP to MP transition should trigger an autonomous message (Event) with the identification of the original CP path.
- The event or the autonomous message indicative that a CP-MP transition has taken place needs to be specific (i.e., unique from any other event notifications indicative of other conditions). Upon query response, the MP creates provisioning records in its database end-to-end per node, and the inventory is updated accordingly. The updated view of the inventory is reflected in the MP for all further MP processing, such as rearrangement of a circuit or any timeslot/bandwidth-related processing. This scenario for transition from the MP to the CP assumes that the required infrastructure for routing services (MP or CP) is created through the MP so that the MP has an inventory of the infrastructure networks in its database, either with an MLRA ID (if CP enabled network) or Network ID (if MP network). In general, networks can have a mixture of CP and Non-CP services. However, based on carrier needs, networks can be dedicated for CP services or MP services. In general, a network transition may be more likely associated with the latter case (requiring a network to be dedicated for MP or CP). A network that can support a mixture of CP and MP services is inventoried as an MLSN when the infrastructure is created. However, such a network may have additional requirements on resource partitioning between the CP and MP. NOTE: In the transition scenario between MP-CP (and vice versa) it is assumed that the Nodes and the Links are absolutely the same.

7 References

1. TMF513 MTNM Business Agreement V3.2
2. OIF Implementation Agreement, “User Network Interface (UNI) 2.0 Signaling Specification: Common Part,” February 2008
3. OIF Implementation Agreement, “OIF E-NNI Signaling Specification (OIF E-NNI 2.0)”, April 2009
4. ITU-T Recommendation G.874, Version 4, July 2010
5. ITU-T Recommendation G.8080/Y.1304 (2006), “Architecture for the Automatically Switched Optical Network (ASON)”
6. ITU-T Recommendation G.8081/Y.1353 (2010), “Terms and definitions for Automatically Switched Optical Networks (ASON)”
7. ITU-T Recommendation, M.3400, Series M: Tmn And Network Maintenance:
8. International Transmission Systems, Telephone Circuits, Telegraphy, Facsimile and Leased Circuits, Telecommunications Management Network, April 1997. ITU-T Recommendation G.7718/Y.1709 (Version 3, July 2010), “Framework for ASON Management Routing Extensions in Support of GMPLS”, RFC4202, October 2005.
9. ITU-T Recommendation, M.3050, Series M: Tmn And Network Maintenance: International Transmission Systems, Telephone Circuits, Telegraphy, Facsimile and Leased Circuits, Telecommunications Management Network, May 2004.
10. ITU-T Recommendation E.101 (2009), Definitions of terms used for identifiers (names, numbers, addresses and other identifiers) for public telecommunication services and networks in the E-series Recommendations
11. ITU-T Recommendation E.191 (2000), B-ISDN Addressing
12. TeleManagement Forum Multi Technology Network Management (MTNM) Version 3.5
13. ITU-T Recommendation G.7718.1/Y.1709.1, “Protocol-neutral management information model for the CP view”, 2006
14. TeleManagement Forum Multi Technology Network Management (MTNM) Version 3.5. Supporting document SD1-24 “Multi-Technology Network Management support for a Naming Convention” version 3.3
15. CP security: Optical Internetworking Forum Implementation Agreement OIF-SEP-03.1, “Security Extension for UNI and NNI version 2.0,” October 2011.
16. Securing management interfaces to NEs: Optical Internetworking Forum Implementation Agreement OIF-SMI-03.0, “Security for Management Interfaces to Network Elements 2.0,” October 2011.
17. Logging and auditing: Optical Internetworking Forum Implementation Agreement OIF-SLG-01.2, “OIF CP Logging and Auditing with Syslog Version 1.1,” October 2011.

18. OIF-CWG-CPR-01.0 - OIF Carrier Working Group Guideline Document: CP Requirements for Multi-Domain Optical Transport Networks (July 2010).
19. ITU-T (G.805)- Series G: Transmission Systems and Media Digital Systems and Networks, March 2000

8 Appendix A: List of Contributors

TECHNICAL EDITOR

Sriram Subramanian

Telcordia

Address: RRC 6J337

Address 2: #1 Telcordia Drive

Phone: 1.732.699.8818

Email: ssubrama@telcordia.com

WORKING GROUP VICE CHAIR

Evelyne Roch

Ciena

3500 Carling Ave

2 Ottawa, Canada K2H 8E9

1.613.670.3193

Email: eroch@ciena.com

Other contributors:

Doug Zuckerman

Fred Gruman

George Frank

Hans-Martin Foisel

Jim Jones

Jonathan Sadler

Lyndon Ong

Martha Fratt

Monica Lazer

Robert Bryttingard

Stephen Liu

Stephen Shew

Thierry Marcot

Vishnu Shukla

9 Appendix B: About the OIF

The OIF is an international nonprofit organization with over 90 member companies, including the world's leading carriers and vendors. Being an industry group uniting representatives of the data and optical worlds, OIF's purpose is to accelerate the deployment of interoperable, cost-effective and robust optical internetworks and their associated technologies. Optical internetworks are data networks composed of routers and data switches interconnected by optical networking elements.

With the goal of promoting worldwide compatibility of optical internetworking products, the OIF actively supports and extends the work of national and international standards bodies. Working relationships or formal liaisons have been established with IEEE 802.1, IEEE 802.3ba, IETF, IP-MPLS Forum, IPv6 Forum, ITU-T SG13, ITU-T SG15, MEF, ATIS-OPTXS, ATIS-TMOC, TMF and the XFP MSA Group.

For additional information contact:

The Optical Internetworking Forum, 48377 Fremont Blvd.,

Suite 117, Fremont, CA 94538

510-492-4040 ☎ info@oiforum.com

www.oiforum.com

10 Appendix C: Glossary

Table 1: Definitions and Terminology

ASON (Automatically Switched Optical Network)	A switched transport network (connection oriented circuit or packet transport) as defined in ITU-T [6] whose connection management is implemented by means of a CP. ASON architecture is described in ITU-T [5].
Control Domain	Reference: [6]
Edge Port (CP Edge Port)	The edge port is the port on the Ingress or Egress Node where the CP call is provisioned. The CP call has source and destination TNA attributes.
Fault Management	The function responsible for processing alarms and events in the network, including root cause analysis, alarm correlation to links, etc.
FCAPS (Fault, Configuration, Accounting, Performance, Security)	The ISO Telecommunications Management Network model and framework for network management. These are the management categories into which the ISO model defines network management tasks. In non-billing organizations, accounting is sometimes replaced with administration.
Multi Layer Routing Area (MLRA)	A multi layer routing area managed by the EMS system. In the CP context, an MLRA consists of multiple congruent routing areas (RAs) at different layers. The CP RA consists of a set of subordinate routing areas, the SNPP links that interconnect them and the SNPPS representing the ends of the SNPP links exiting that RA. The way that RAs and SNPP links are grouped in RAs is entirely determined by the EMS. MLRAs are organized in a hierarchical structure based upon "containment" (in which the lower level MLRAs are completely contained within a single higher level MLRA).
Multi Layer Sub Network (MLSN)	The abstraction offered by the EMS to the NMS to represent a Subnetwork that is managed by the EMS system. It represents a logical grouping or partitioning of the resources of managed elements in a way that is entirely determined by the EMS or CP. In this document, Subnetwork and MultiLayer Subnetwork are used interchangeably. The MultiLayer Subnetwork may also represent a Multi Layer Routing Area (MLRA) in the CP.
Network Inventory	The Resource Inventory function stores and maintains a reference of all technical resources (physical and logical resources) used to provide a service. The resource inventory describes network and IT resources used during the service delivery process (mainly through the Resource Assignment function) and also during Assurance processes. This function includes the sub-function "Reconciliation" (resource inventory data should be automatically aligned to field data; it must also help and support manual data

	correction).
Operations Support Systems (OSS)	An implementation of the network management architecture and may incorporate functionality as described in the SMS, NMS, and EMS layers. OSS supports business processes enabling management, monitoring, and control of the telecommunications networks and services, including billing, directory services, maintenance, inventory management, activation/provisioning, and fault management. In general, a collection of systems that comprise the MP is often, informally referred to as OSS by operators and systems providers.
Provisioning/Activation	Function that supports the CP configuration and activation, including the mapping of CP and MP identifiers.
PTP (Physical Termination Points)	A Termination Point (TP) that includes a physical layer rate. It represents the signal termination and assurance functions supported by a port., not the physical connector itself. Examples of PTPs include T1 ports, T3 ports, OC-N optical ports, etc. PTPs have a containment relationship with Connection Termination Points (CTP). PTPs forming a UPSR pair are related and contain related CTPs. If the PTP is potentially able to support a CTP on a particular layer rate, then that CTP is represented at the Interface.
Rerouting Domain	Reference: [5]
Routing Control Domain	Reference: [5]
Service Inventory	A function that stores and maintains two kinds of data: Technical Nomenclature that defines the list of resource types that support service types described in the catalogue. Installed Technical Services (also called Technical Installed base) that describes the set of resources used to support the services of a dedicated customer. The Technical Installed base links resources with the Installed Services.
Shared Risk Link Group (SRLG)	A CP link attribute that is the union of the SRG identifiers of all resources used to construct the CP link.
Soft Permanent Connection (SPC)	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.
Transport Network Assigned (TNA) Name	UNI connection endpoints are identified by Transport Network Assigned (TNA) names. Each TNA is a globally unique name assigned by the transport network to one or more data-bearing links connecting a client and TNE and scoped by a UNI-C and UNI-N pair. The TNA is an instantiation of the [5] UNI Transport ResourceIdentifier.

11 Appendix D: Acronym Definitions

AID	Access Identifier
CP	Control Plane
EML	Element Management Layer
EMS	Element Management System
E-NNI	External Network to Network Interface
IGP	Interior Gateway Protocol
I-NNI	Internal Network to Network Interface
LSP	Label Switched Paths
MP	Management Plane
NMS	Network Management System
NNI	Network to Network Interface
NNTH	Network to Network Traditional Handoff
SML	Service Management Layer
SMS	Service Management System
SRG	Shared Resource Group
UNI	User Network Interface