Data-Over-Cable Service Interface Specifications
Flexible MAC Architecture

Flexible MAC Architecture (FMA) System Specification

CM-SP-FMA-SYS-I03-220126

ISSUED

Notice

This DOCSIS® specification is the result of a cooperative effort undertaken at the direction of Cable Television Laboratories, Inc. for the benefit of the cable industry and its customers. You may download, copy, distribute, and reference the documents herein only for the purpose of developing products or services in accordance with such documents, and educational use. Except as granted by CableLabs® in a separate written license agreement, no license is granted to modify the documents herein (except via the Engineering Change process), or to use, copy, modify or distribute the documents for any other purpose.

This document may contain references to other documents not owned or controlled by CableLabs. Use and understanding of this document may require access to such other documents. Designing, manufacturing, distributing, using, selling, or servicing products, or providing services, based on this document may require intellectual property licenses from third parties for technology referenced in this document. To the extent this document contains or refers to documents of third parties, you agree to abide by the terms of any licenses associated with such third-party documents, including open source licenses, if any.

© Cable Television Laboratories, Inc., 2019-2022

DISCLAIMER

This document is furnished on an "AS IS" basis and neither CableLabs nor its members provides any representation or warranty, express or implied, regarding the accuracy, completeness, noninfringement, or fitness for a particular purpose of this document, or any document referenced herein. Any use or reliance on the information or opinion in this document is at the risk of the user, and CableLabs shall not be liable for any damage or injury incurred by any person arising out of the completeness, accuracy, infringement, or utility of any information or opinion contained in the document.

CableLabs reserves the right to revise this document for any reason including, but not limited to, changes in laws, regulations, or standards promulgated by various entities, technology advances, or changes in equipment design, manufacturing techniques, or operating procedures described, or referred to, herein.

This document is not to be construed to suggest that any company modify or change any of its products or procedures, nor does this document represent a commitment by CableLabs or any of its members to purchase any product whether or not it meets the characteristics described in the document. Unless granted in a separate written agreement from CableLabs, nothing contained herein shall be construed to confer any license or right to any intellectual property. This document is not to be construed as an endorsement of any product or company or as the adoption or promulgation of any guidelines, standards, or recommendations.

Document Status Sheet

Document Control Number:	CM-SP-FMA-SYS-I03-220126			
Document Title:	Flexible MAC Architecture (FMA) System Specification			
Revision History:	W03 – Released 03/22/2019 D01 – Released 04/30/2019 D02 – Released 07/24/2019 D03 – Released 10/23/2019 D04 – Released 02/18/2020 D05 – Released 05/04/2020 D06 – Released 08/18/2020 D07 – Released 09/15/2020		I01 – Released 09/30/2020 I02 – Released 05/26/2021 I03 – Released 01/26/2022	
Date:	January 26, 2022			
Status:	~~Work in Progress~~	~~Draft~~	**Issued**	~~Closed~~
Distribution Restrictions:	~~Focus Team Only~~	~~CL/Member~~	~~CL/ Member/ Vendor~~	**Public**

Key to Document Status Codes

Work in Progress	An incomplete document, designed to guide discussion and generate feedback that may include several alternative requirements for consideration.
Draft	A document in specification format considered largely complete, but lacking review by Members and vendors. Drafts are susceptible to substantial change during the review process.
Issued	A generally public document that has undergone Member and Technology Supplier review, cross-vendor interoperability, and is for Certification testing if applicable. Issued Specifications are subject to the Engineering Change Process.
Closed	A static document, reviewed, tested, validated, and closed to further engineering change requests to the specification through CableLabs.

Trademarks

CableLabs® is a registered trademark of Cable Television Laboratories, Inc. Other CableLabs marks are listed at http://www.cablelabs.com/specs/certification/trademarks. All other marks are the property of their respective owners.

Contents

1 SCOPE 11

1.1 Introduction and Purpose 11
1.2 FMA Specification Documents 14
1.3 Requirements 14

2 REFERENCES 15

2.1 Normative References 15
2.2 Informative References 19
2.3 Reference Acquisition 20

3 TERMS AND DEFINITIONS 21

4 ABBREVIATIONS AND ACRONYMS 24

5 FMA SYSTEM OVERVIEW 32

5.1 MPEG Video 32
5.2 OOB Communications 32
5.3 IP-Based Residential Services 32
5.4 PacketCable 32
5.5 Lawful Intercept (LI) 32
5.6 Business Services over DOCSIS - L2VPN 32
5.7 Mobile Xhaul (MXh) Transport Service 33

6 TECHNICAL OVERVIEW 34

6.1 Introduction 34
6.2 System Diagram 35
6.2.1 Hub Access Network 36
6.2.2 Optical Access Network 36
6.2.3 Coax Access Network 36
6.2.4 Location of the MAC-NE and RF Requirements 36
6.2.5 Location of the MAC Network Element Functions 37
6.2.6 Location of the FMA Components 37
6.3 System Architecture 37
6.3.1 System Components 37
6.3.2 MAC-NE Device Architecture 41
6.3.3 MAC Manager Architecture 43
6.3.4 PacketCable Architecture 44
6.3.5 Lawful Intercept Architecture 53
6.4 Auxiliary Core System Operation 61
6.4.1 CCAP Core Types 61
6.4.2 Auxiliary CCAP Core RCP Fields 61
6.4.3 FMA R-DTI Operations 61
6.4.4 FMA GCP Operations 62
6.4.5 Traffic Engine Operations 63
6.4.6 L2TPv3, R-DEPI, and R-UEPI 64
6.4.7 L2TPv3 Requirements 64
6.4.8 R-DEPI MPEG Video Multicast Support 64
6.4.9 Video Core Architecture 65
6.4.10 OOB Core Architecture 66
6.4.11 OOB Support 66
6.4.12 R-PHY OOB Architecture Summary 67
6.4.13 FMA OOB Architecture 69
6.4.14 NDx Core Architecture 70
6.4.15 Pilot Tones 70

6.4.16 Leakage Detection Signal Generation 71
6.4.17 Timing Architecture 72
6.4.18 Remote DTI 73
6.5 Network Considerations 73
6.5.1 S-Leaf to MAC Network Element 73

7 FMA SYSTEM OPERATIONS 74

7.1 Reuse of MHAv2 Requirements 74
7.1.1 Out-of-Band Operations 74
7.1.2 DEPI Operations 74
7.1.3 UEPI Operations 74
7.2 DSG Operation 75
7.2.1 DSG IPv6-Only Operation 75
7.3 Networking Considerations 76
7.3.1 Network MTU 76
7.4 FMA Operation with Static Pseudowires 77
7.5 Multiple Core Operation 78
7.5.1 Introduction 78
7.5.2 MAC-NE Startup with Multiple Cores 78
7.5.3 Resource Sets and Auxiliary Resource Assignment 80
7.5.4 MAC-NE Read Operations 81
7.5.5 MAC-NE Write Operations 81
7.6 Resilient Software Upgrades 82

8 MAC MANAGER AND MAC-NE MANAGEMENT OPERATIONS 84

8.1 Management Interface Models 84
8.1.1 Southbound Interface Protocols 84
8.1.2 Southbound Interface Models 84
8.1.3 Northbound Interface Protocols 84
8.1.4 Northbound Interface Models 84
8.2 FCAPS Management Operations 84
8.2.1 Fault Management 84
8.2.2 Configuration Management 85
8.2.3 Accounting Management 85
8.2.4 Performance Management 85
8.2.5 Security Management 85
8.3 Subscriber Database 86
8.3.1 MAC Manager Requirements 86
8.3.2 MAC-NE Requirements 86
8.4 Interface Management 87
8.4.1 MAC Domain Configuration 87

9 MAC-NE TIMING AND SYNCHRONIZATION 88

9.1 Introduction 88
9.2 Timing Deployment Scenarios for the MAC-NE 88
9.2.1 No External Timing Requirements 88
9.2.2 External Frequency Synchronization 88
9.3 Precision Timing Synchronization Services 90
9.3.1 Market Forces Drive Need for Precise Timing 90
9.3.2 Mobile Xhaul (MXh) Timing for Small Cells 90

10 DATA FORWARDING 92

10.1 MAC-NE Data Forwarding Assumptions 92
10.2 MAC-NE Internal Data Forwarding Model 94
10.2.1 High-Level Model 94
10.2.2 NSI Encapsulation 95

10.2.3 Forwarding Functions 96
10.2.4 IP Host Interfaces 97
10.3 MAC-NE Data Forwarding 99
10.3.1 Data Forwarding Flowcharts 99
10.4 IP Services 101
10.4.1 MAC and IP Address Learning 101
10.4.2 Address Resolution 102
10.4.3 Cable Bundles 102
10.4.4 Unicast IP Services Forwarding 107
10.4.5 IP Services Forwarding Database (IP-FDB) 108
10.4.6 Multicast IP Services 111
10.5 L2VPN Services 112
10.6 Multiple Ethernet Ports 112
10.7 Quality of Service 112
10.7.1 IP Services QoS 113
10.7.2 IP Services QoS 113
10.7.3 L2VPN Services QoS 114
10.7.4 L2TPv3 Services QoS 114
10.7.5 Upstream Packet Scheduling 114
10.7.6 Downstream Packet Scheduling 115

11 SECURE PROVISIONING OF CM AND CPE 116

11.1 DHCP 116
11.2 Source Address Verification (SAV) 117

12 SECURITY AND CERTIFICATE OPERATIONS 119

12.1 Security Overview 119
12.2 MACsec Link Layer Security 119
12.3 Transport Layer Security (TLS) for the FMA System 119
12.3.1 Application of TLS to the Interfaces of the FMA System 119
12.3.2 TLS Requirements in the FMA System 120
12.3.3 General TLS Version Requirements 122
12.3.4 General TLS Cipher Suite Requirements 123
12.3.5 TLS Event Requirements of the MAC-NE 124
12.3.6 TLS Event Requirements of the MAC Manager 125
12.3.7 TLS Event Requirements of the PAG 125
12.3.8 General TLS Authentication And Version Requirements 126
12.3.9 TLS Certificate Basic Path Validation Requirements of the FMA System 126
12.3.10 Connection Persistence and Re-Use 127
12.3.11 Session Caching 127
12.4 Re-Use from MHAv2 R-PHY 127
12.4.1 X.509 Certificate Operations 127
12.4.2 Physical Protection of Keys 127
12.5 Secure Shell Requirements 127
12.5.1 MAC-NE SSH Requirements 128
12.5.2 MAC Manager SSH Server Requirements 130
12.5.3 MAC Manager SSH Client Requirements 132
12.6 Secure Boot 134
12.7 Tamper Resistance/Detection 134
12.8 Secure Software Download 134
12.8.1 Introduction 134
12.8.2 Overview 135
12.8.3 The MAC-NE Software Update Procedure 137
12.8.4 Software Code Upgrade Requirements 140

13 ONBOARDING/INITIALIZATION OF MAC-NE 143

13.1 Overview 143
13.2 Logging and Recording 144
13.3 Local Device Initialization 144
13.4 Wait for CIN Link 145
13.4.1 CIN Interface Timeout 147
13.5 Network Authentication 147
13.5.1 Problem Definition 147
13.5.2 802.1x Authentication 148
13.6 Link Layer Discovery 153
13.6.1 LLDP Support 153
13.6.2 Operation 153
13.7 Link Aggregation (LAG) 153
13.7.1 Link Aggregation Control Protocol (LACP) 154
13.7.2 Distributing Utilization Across Links and Frame Ordering Considerations 155
13.7.3 LAG and LACP Events 156
13.8 Address Assignment 156
13.8.1 DHCP Options 160
13.8.2 MAC Manager Suboptions 161
13.8.3 Failures 162
13.8.4 DNS Support 162
13.8.5 Security Implications 162
13.9 DNS SRV Discovery Method of Operation 162
13.10 SSH 163
13.11 Get ToD (Time of Day) 163
13.11.1 ToD Acquisition 163
13.11.2 ToD Conflicts and Problems 164
13.11.3 ToD Security Implications 165
13.12 Apply Saved Configuration 165
13.13 Listen for RESTCONF Connections 165
13.13.1 Start RESTCONF Listener 166
13.13.2 Connection Process 166
13.14 MAC-NE Announce to MAC Managers 167
13.14.1 Redirection 168
13.14.2 Connection Pending 168
13.14.3 Announce Timer 169
13.14.4 Software Upgrade 169
13.15 Synchronization 169
13.15.1 Synchronization Failures 170
13.16 Move to Operational 171
13.16.1 MAC Manager 171
13.16.2 Auxiliary Core 171
13.16.3 PacketCable Aggregator 171

14 PROACTIVE NETWORK MAINTENANCE 173

14.1 Scope 173
14.2 System Description 173
14.2.1 PNM Supported Functions 173
14.2.2 PNM Architectures 174
14.2.3 Common Collection Framework (CCF) 176

ANNEX A CORE TO MAC-NE MTU (NORMATIVE) 178

A.1 L2TPv3 Lower Layer Payload Size 178
A.2 Maximum Frame Size for DEPI 178
A.3 Path MTU Discovery 178

ANNEX B GCP USAGE (NORMATIVE) 180

B.1 R-PHY Control Protocol 181
B.1.1 RCP over GCP EDS Message 181
B.1.2 RCP over GCP EDS Response Messages 181
B.1.3 RCP over GCP Device Management Message 181
B.1.4 RCP over GCP Notify Message 181
B.1.5 RCP TLV Format, TLV Types and Nesting Rules 181
B.1.6 RCP Message Structure 181
B.1.7 RCP Messages Types 181
B.1.8 RCP Protocol Rules 181
B.1.9 Protocol Extensibility 181
B.1.10 Protocol Versioning 181
B.1.11 Information Model Extensibility 181
B.1.12 Vendor-Specific Extensions 181
B.1.13 Inclusion of DOCSIS Messages 181
B.1.14 Event Reporting 182
B.1.15 Error Handling 182
B.1.16 RCP Message Examples 182
B.2 RPD Initialization 182
B.2.1 GCP Connection Initialization Sequence 182
B.2.2 Initialization RCP Messages RPD and Cores 182
B.2.3 RPD Initialization States 184
B.2.4 Reconnect Messages 184
ANNEX C MPEG STREAM ANALYSIS (NORMATIVE) 185
ANNEX D CERTIFICATE HIERARCHY AND PROFILES (NORMATIVE) 186
ANNEX E RECEIVE POWER LEVEL MANAGEMENT (NORMATIVE) 187
ANNEX F DOCSIS 4.0 OFDM MODIFICATIONS FOR REMOTE PHY (NORMATIVE) 188
APPENDIX I ACKNOWLEDGEMENTS 189
APPENDIX II REVISION HISTORY 190

Figures

Figure 1 - Traditional I-CCAP System 11
Figure 2 - Aggregation to Disaggregation of the CCAP 12
Figure 3 - Target Architecture for Flexible MAC Architecture (FMA) 13
Figure 4 - FMA Support of Vendor Interoperability 14
Figure 5 - Double-Tagged Ethernet Service Frame 33
Figure 6 - MAC-NE Internal and External Connections 35
Figure 7 - FMA High-Level System Diagram with MAC-NE 35
Figure 8 - FMA Phase 1 Reference Architecture 38
Figure 9 - MAC Network Elements (MAC-NEs) 39
Figure 10 - Secure Leaf (S-Leaf) and Leaf 40
Figure 11 - MAC-NE Device Architecture 42
Figure 12 - FMA Management Overview 43
Figure 13 - MAC Manager High Level Architecture 44
Figure 14 - FMA with PacketCable Aggregator 45
Figure 15 - Example Message Flow with PacketCable Aggregator 47
Figure 16 - PacketCable Aggregator Redundancy 49
Figure 17 - Recommended FMA Lawful Intercept 54

Figure 18 - Recommendations for FMA LAES and HSD Intercept 55
Figure 19 - FMA Lawful Intercept Tap Provisioning 56
Figure 20 - DQoS-based Voice Call Intercept Set-up 57
Figure 21 - DQoS-based Voice Call Intercept Tear-down 58
Figure 22 - FMA Lawful Intercept Event and Content Delivery 59
Figure 23 - FMA Video - Aux Core OOB System Components 63
Figure 24 - FMA Video - OOB System Components 64
Figure 25 - FMA Video System Components 66
Figure 26 - NDF - NDR Signal Processing in the MAC-NE 67
Figure 27 - FMA Operation of SCTE 55-1 OOB 68
Figure 28 - FMA Operation of SCTE 55-2 OOB 69
Figure 29 - FMA OOB Architecture 70
Figure 30 - NDx Core Architecture 70
Figure 31 - MAC-NE Timing Architecture 72
Figure 32 - Timing Distribution to MAC-NE through Switches - Routers (T-BCs) 72
Figure 33 - DSG Operation 75
Figure 34 - DSG IPv6-Only Operation 76
Figure 35 - Mobile Xhaul (MXh) Timing for Small Cells 91
Figure 36 - Typical S-Leaf Connectivity 93
Figure 37 - High-Level Data Forwarding Model 94
Figure 38 - Detailed Data Forwarding Model 96
Figure 39 - IP Host Interfaces 98
Figure 40 - Downstream Data Passing Flowchart 100
Figure 41 - Upstream Data Passing Flowchart 101
Figure 42 - Cable Bundles 102
Figure 43 - Cable Sub-Bundle 103
Figure 44 - CM Subnet-based Sub-bundles 106
Figure 45 - TLS-Secured Logical Interfaces of the FMA System 120
Figure 46 - TLS version 1.2 Connection Sequence 121
Figure 47 - TLS version 1.3 Connection Sequence 122
Figure 48 - Typical Code Validation Hierarchy 137
Figure 49 - Secure Software Download FSM 139
Figure 50 - MAC-NE Initialization FSM 143
Figure 51 - Local Device Initialization FSM 145
Figure 52 - Wait for CIN Link FSM 146
Figure 53 - Remote MACPHY - Trusted Domain and Untrusted Domain 147
Figure 54 - Network Authentication Signaling 149
Figure 55 - MAC-NE Topologies for 802.1x Network Authentication 150
Figure 56 - Network Authentication FSM 151
Figure 57 - Start LLDP FSM 153
Figure 58 - LAG Groups 154
Figure 59 - IP Address Assignment FSM 157
Figure 60 - DHCP Network Diagram 159
Figure 61 - DHCPv4 Signaling 159

Figure 62 - DHCPv6 Signaling....160
Figure 63 - MAC Manager DHCP Suboption IPv4....162
Figure 64 - MAC Manager DHCP Suboption IPv6....162
Figure 65 - Start SSH FSM....163
Figure 66 - Get ToD FSM....164
Figure 67 - Apply Saved Configuration FSM....165
Figure 68 - Listen for RESTCONF FSM....166
Figure 69 - Announce to MAC Manager FSM....167
Figure 70 - Test Points in CMTS and CM Supporting PNM....173
Figure 71 - Integrated CMTS - CCAP PNM Architecture....174
Figure 72 - Integrated CMTS - CCAP Upstream Sweep Architecture....175
Figure 73 - FMA PNM Architecture....175
Figure 74 - FMA Upstream Sweep Alternatives....176
Figure 75 - FMA with CCF....177

Tables

Table 1 - List of FMA Specifications....14
Table 2 - PacketCable Security Interfaces (PacketCable Aggregator)....46
Table 3 - ConfiguredCoreTable Object....78
Table 4 - CcapCoreIdentification Table....78
Table 5 - ResourceSet Table....80
Table 6 - PermitAuxSelfConfiguration Table....81
Table 7 - Bundle Configuration Example....105
Table 8 - Sub-bundle Configuration Example....105
Table 9 - Sub-bundle Identification Rules Example....105
Table 10 - Default IP Services Upstream Priority Mapping....114
Table 11 - TLS Version 1.2 Cipher Suites....123
Table 12 - TLS Version 1.3 Cipher Suites....124
Table 13 - MAC-NE CIN Interface Timeout....147
Table 14 - 802.1x Network Authentication Entities....148
Table 15 - EAP Authentication Timing....152
Table 16 - Router Advertisement M Bit and O Bit Settings For SLAAC....158
Table 17 - FMA DHCP Options....160
Table 18 - MAC-NE PTP Timeout....170
Table 19 - FMA PNM Tests....174
Table 20 - MTU of DEPI (for PSP)....178

1 SCOPE

1.1 Introduction and Purpose

Flexible MAC Architecture (FMA) technology provides a modular headend, which supports an all-digital hybrid fiber-coax (HFC) plant. In an FMA-based digital HFC plant, the fiber portion utilizes a suitable baseband network transmission technology such as Ethernet, EPON (Ethernet passive optical network), GPON (gigabit passive optical network), or any Layer 2 technology that would support a fiber-based Layer 1. FMA specifies interoperable management, control and data plane interfaces on the digital fiber between the core components and a series of MAC-NEs (MAC-layer Network Elements) of which there are several types. Types of MAC-NE devices are typically identified by their location in the network. One of the most common locations for a MAC-NE is at an optical node device at the junction of the fiber and coaxial plants. An appliance containing such a component is typically referred to as a Remote MACPHY Node (RMN). The other typical location for an FMA network device is in the HFC hub or headend. An appliance containing such components is termed a Remote MACPHY Shelf (RMS). The terms RMN and RMS are interchangeable with the term RMD (the generic term for either an RMS or an RMN) or can be referred to as a MAC-NE, which is a class of device providing services over the coaxial plant.

FMA Remote MACPHY (R-MACPHY) is designed to enable software-defined networking (SDN) and headend/hub network function virtualization (NFV) in mind. R-MACPHY avoids creating unnecessary latency, does not require tight coupling, provides interoperability, and extensibility through its particular approach to the separation of the DOCSIS MAC, the upper layers, and the upstream scheduler function from the DOCSIS PHY. R-MACPHY has been shown to optimally support future services such as DOCSIS 4.0; support for necessary capabilities requires no additional standardization and interoperability testing to accommodate the base DOCSIS architecture. It is anticipated that R-MACPHY will also optimally support Low Latency DOCSIS (LLD) and Mobile Xhaul (MXh) over DOCSIS, given that these features involve changes to the DOCSIS upstream scheduler.

The FMA standard leverages advancements made during R-PHY standardization that enable video and out-of-band (OOB) support over digital optics between core components and access network elements.

Figure 1 represents the traditional I-CCAP (e.g., I-CCAP DOCSIS or PON OLT system) and the composition of component functions.

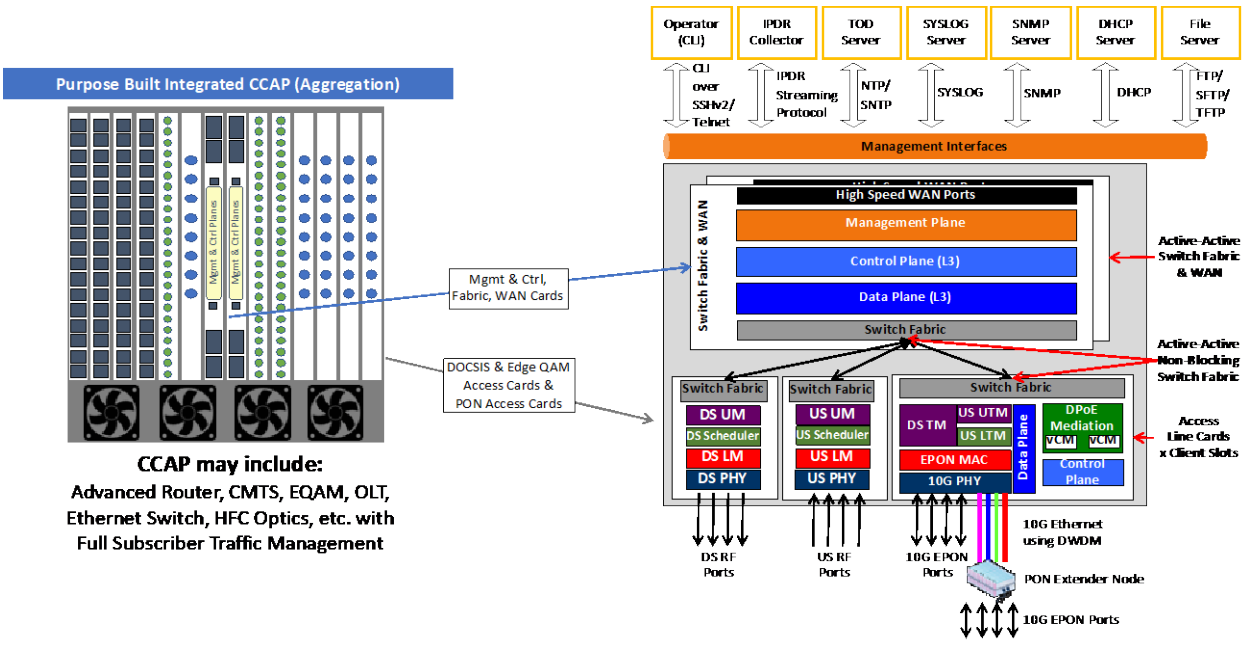

Figure 1 - Traditional I-CCAP System

Figure 2 represents the aggregation to disaggregation of the CCAP (CAA to DAA + SDN + NFV).

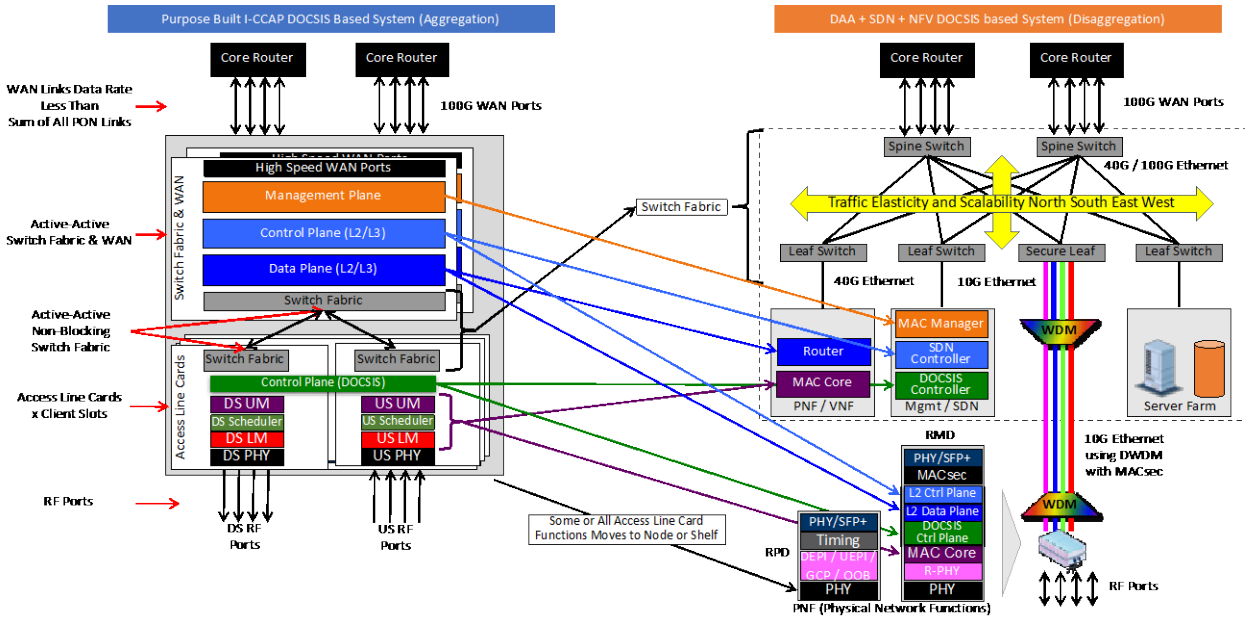

Figure 2 - Aggregation to Disaggregation of the CCAP

Figure 3 represents the target architecture for Flexible MAC Architecture (FMA).

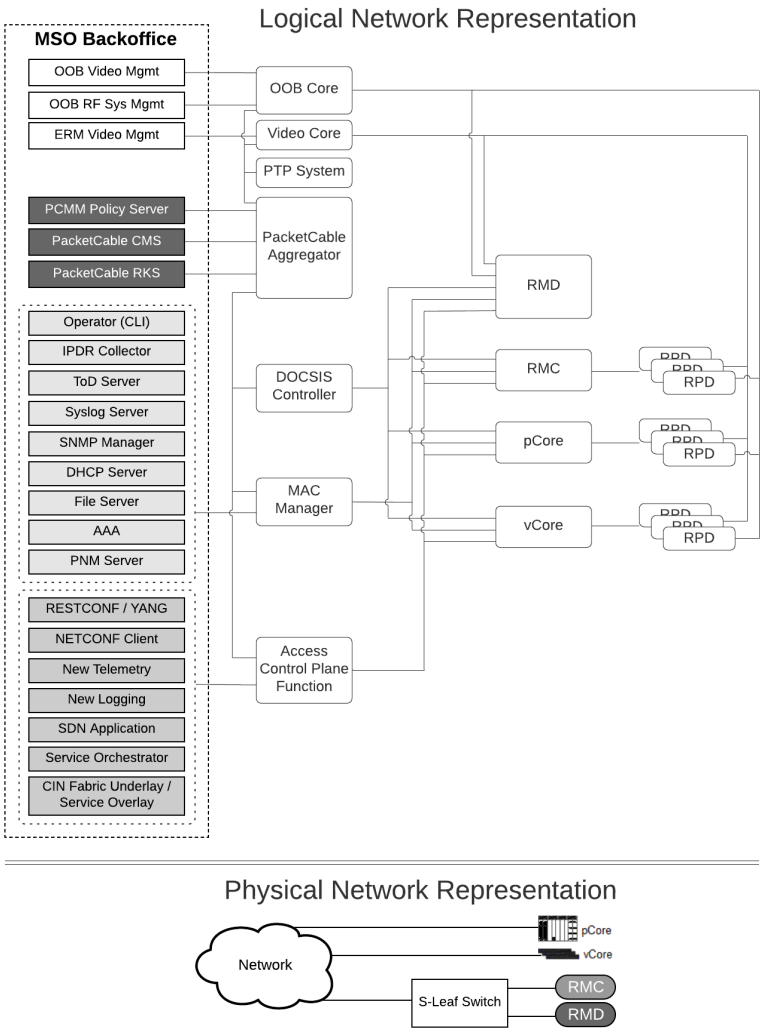

Figure 3 - Target Architecture for Flexible MAC Architecture (FMA)

The core objective of the cable operators is to define interfaces for a disaggregated CCAP such that the components may be interoperable between suppliers, as shown in Figure 4.

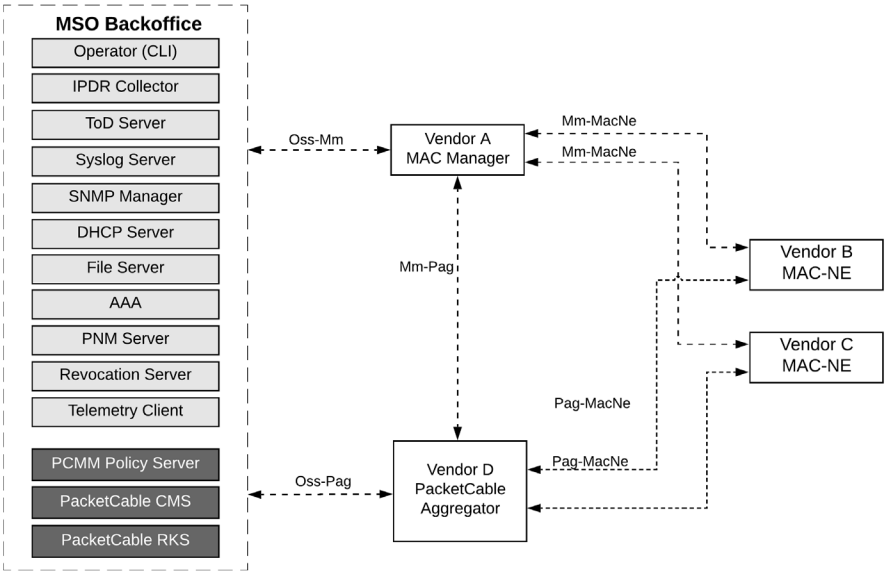

Figure 4 - FMA Support of Vendor Interoperability

1.2 FMA Specification Documents

A list of the documents in the FMA family of specifications is provided in Table 1, below. For updates, refer to http://www.cablelabs.com/specs/specification-search/.

Table 1 - List of FMA Specifications

Designation	Title
CM-SP-FMA-SYS (this spec)	Flexible MAC Architecture System Specification
CM-SP-FMA-OSSI	Flexible MAC Architecture Operations Support System Interface Specification
CM-SP-FMA-MMI	Flexible MAC Architecture MAC Manager Interface Specification
CM-SP-FMA-PAI	Flexible MAC Architecture PacketCable Aggregator Interface Specification
CM-SP-FMA-YANG	Flexible MAC Architecture YANG Object Models

1.3 Requirements

Throughout this document, the words that are used to define the significance of particular requirements are capitalized. These words are:

"MUST"	This word means that the item is an absolute requirement of this specification.
"MUST NOT"	This phrase means that the item is an absolute prohibition of this specification.
"SHOULD"	This word means that there may exist valid reasons in particular circumstances to ignore this item, but the full implications should be understood, and the case carefully weighed before choosing a different course.
"SHOULD NOT"	This phrase means that there may exist valid reasons in particular circumstances when the listed behavior is acceptable or even useful, but the full implications should be understood, and the case carefully weighed before implementing any behavior described with this label.
"MAY"	This word means that this item is truly optional. One vendor may choose to include the item because a particular marketplace requires it or because it enhances the product, for example; another vendor may omit the same item.

2 REFERENCES

2.1 Normative References

In order to claim compliance with this specification, it is necessary to conform to the following standards and other works as indicated, in addition to the other requirements of this specification. Notwithstanding, intellectual property rights may be required to use or implement such normative references.

All references are subject to revision, and parties to agreement based on this specification are encouraged to investigate the possibility of applying the most recent editions of the documents listed below.

[BBF TR-156]	Broadband Forum Technical Report BBF-TR-156, Using GPON Access in the context of TR-101, Issue 4, November 2017.
[CANN]	CableLabs' Assigned Names and Numbers, CL-SP-CANN-I20-200715, July 15, 2020, Cable Television Laboratories, Inc.
[CANN DHCP]	CableLabs' DHCP Options Registry, CL-SP-CANN-DHCP-Reg-I16-200715, July 15, 2020, Cable Television Laboratories, Inc.
[CBI]	Cable Broadband and Wi-Fi Intercept Specification, CM-SP-CBI2.0-I11-160602, June 2, 2016, Cable Television Laboratories, Inc.
[CCAP-OSSIv3.1]	DOCSIS 3.1 CCAP OSSI Specification, CM-SP-CCAP-OSSIv3.1-I22-211022, October 22, 2021, Cable Television Laboratories, Inc.
[CCAP-OSSIv4.0]	DOCSIS 4.0 CCAP OSSI Specification, CM-SP-CCAP-OSSIv4.0-I05-210927, September 27, 2021, Cable Television Laboratories, Inc.
[CM-OSSIv3.1]	DOCSIS 3.1 Cable Modem OSSI Specification, CM-SP-CM-OSSIv3.1-I21-211022, October 22, 2021, Cable Television Laboratories, Inc.
[CM-OSSIv4.0]	DOCSIS 4.0 Cable Modem OSSI Specification, CM-SP-CM-OSSIv4.0-I05-211027, September 27, 2021, Cable Television Laboratories, Inc.
[DOCS-SEC-MIB]	DOCSIS Security MIB, DOCS-SEC-MIB-2020-03-19, http://mibs.cablelabs.com/MIBs/DOCSIS/
[DRFI]	DOCSIS Downstream Radio Frequency Interface, CM-SP-DRFI-I16-170111, January 11, 2017, Cable Television Laboratories, Inc.
[DSG]	DOCSIS Set-top Gateway Interface Specification, CM-SP-DSG-I25-170906, September 6, 2017, Cable Television Laboratories, Inc.
[eDOCSIS]	eDOCSIS Specification, CM-SP-eDOCSIS-I30-190213, February 13, 2019, Cable Television Laboratories, Inc.
[FIPS-186-4]	NIST CSRC FIPS 186-4, Digital Signature Standard (DSS), July 2013.
[FMA-MMI]	Flexible MAC Architecture MAC Manager Interface Specification, CM-SP-FMA-MMI-I03-220126, January 26, 2022, Cable Television Laboratories, Inc.
[FMA-OSSI]	Flexible MAC Architecture Operations Support Systems Interface Specification, CM-SP-FMA-OSSI-I01-211101, November 1, 2021, Cable Television Laboratories, Inc.
[FMA-PAI]	Flexible MAC Architecture PacketCable Aggregator Interface Specification, CM-SP-FMA-PAI-I01-200930, September 30, 2020, Cable Television Laboratories, Inc.
[FMA-YANG]	Flexible MAC Architecture YANG Models, http://mibs.cablelabs.com/YANG/DOCSIS/FMA/
[GCP]	Generic Control Plane Specification, CM-SP-GCP-I05-200323, March 23, 2020, Cable Television Laboratories, Inc.
[IEEE 802.1ab]	IEEE Std 802.1ab-2016, IEEE Standard for Local and Metropolitan Area Networks—Station and Media Access Control Connectivity Discovery, March 2016.
[IEEE 802.1Qcy]	IEEE 802.1Qcy-2019 - IEEE Standard for Local and Metropolitan Area Networks--Bridges and Bridged Networks Amendment 32: Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) Extension to Support Network Virtualization Overlays Over Layer 3 (NVO3).

[IEEE 802.1ae]	IEEE Std 802.1ae-2018, IEEE Standard for Local and Metropolitan Area Networks—Media Access Control (MAC) Security, December 2018.
[IEEE 802.1ar]	IEEE Std 802.1ar-2018, IEEE Standard for Local and Metropolitan Area Networks—Secure Device Identity, June 2018.
[IEEE 802.1as]	IEEE Std 802.1as-2020, IEEE Standard for Local and Metropolitan Area Networks--Timing and Synchronization for Time-Sensitive Applications, January 2020.
[IEEE 802.1AX]	802.1AX-2020 - IEEE Standard for Local and metropolitan area networks -- Link Aggregation.
[IEEE 802.1Q]	802.1Q-2019 - IEEE Standard for Metropolitan Area Networks--Bridges and Bridged Networks Amendment 32: Virtual Station Interface (VSI) Discovery and Configuration Protocol (VDP) Extension to Support Network Virtualization Overlays Over Layer 3 (NVO3).
[IEEE 802.1x]	IEEE Std 802.1x-2010, IEEE Standard for Local and Metropolitan Area Networks—Port-Based Network Access Control, February 2010.
[IEEE 802.3]	IEEE Std 802.32018, IEEE Standard for IT—Telecommunications and Information Exchange Between Systems—LAN/MAN—Specific Requirements—Part 3: CSMA/CD Access Method and Physical Layer Specifications, June 2018.
[IEEE 1588]	IEEE Std 1588-2019, IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems, November 2019.
[IPDR/SSDG]	IPDR Service Specification Design Guide, Version 3.8, TM Forum, October 2009.
[ISO 13818.1]	ISO/IEC 13818.1:2019/AMD 1:2020, Information technology — Generic coding of moving pictures and associated audio information — Part 1: Systems — Amendment 1: Carriage of JPEG XS in MPEG-2 TS, January 2020.
[ITU-T G.781]	ITU-T Recommendation G.781, Synchronization layer functions for frequency synchronization based on the physical layer, April 2020.
[ITU-T G.984.1]	ITU-T Recommendation G.984.3, Amendment 1, Gigabit-capable passive optical networks (G-PON): Transmission convergence layer specification, March 2020.
[ITU-T G.987.3]	ITU-T Recommendation G.987.3, Amendment 3, 10-Gigabit-capable passive optical networks (XG-PON): Transmission convergence (TC) layer specification, January 2014.
[ITU-T G.988]	ITU-T Recommendation G.988, ONU management and control interface (OMCI) specification, November 2017.
[ITU-T G.989.3]	ITU-T Recommendation G.989, Amendment 3, 40-Gigabit-capable passive optical networks (NG-PON2): Definitions, abbreviations and acronyms, March 2020.
[ITU-T G.8013/Y.1731]	ITU-T Recommendation G.8013/Y.1731, Amendment 1, Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks, November 2018.
[ITU-T G.8032/Y.1344]	ITU-T Recommendation G.8032/Y.1344, Ethernet ring protection switching, March 2020.
[ITU-T G.8260]	ITU-T Recommendation G.8260, Definitions and terminology for synchronization in packet networks, August 2015.
[ITU-T G.8261.1]	ITU-T Recommendation G.8261/Y.1361, Amendment 1, Timing and synchronization aspects in packet networks, May 2014.
[ITU-T G.8262/Y.1362]	ITU-T Recommendation G.8262/Y.1362, Internet protocol aspects – Transport Timing characteristics of synchronous equipment slave clock, March 2020.
[ITU-T G.8264.1 /Y.1364]	ITU-T Recommendation G.8264.1/Y.1364, Amendment 1, Distribution of timing information through packet networks, March 2018.
[ITU-T G.8265.1]	ITU-T Recommendation G.8265.1/Y.1365.1, Amendment 1, Precision time protocol telecom profile for frequency synchronization, August 2019.
[ITU-T G.8273.1]	ITU-T Recommendation G.8273/Y.1368, Amendment 1, Framework of phase and time clocks, March 2020.
[ITU-T G.8275.1]	ITU-T Recommendation G.8275.1/Y.1369,1 Amendment 1, Precision time protocol telecom profile for phase/time synchronization with full timing support from the network, November 2020

[ITU-T G.8275.2]	ITU-T Recommendation G.8275.2/Y.1369.2, Amendment 1, Precision time protocol telecom profile for phase/time synchronization with partial timing support from the network, November 2020.
[ITU-T G.9807.1]	ITU-T Recommendation G.9807, Amendment 1, 10-Gigabit-capable symmetric passive optical network (XGS-PON), December 2017.
[ITU-T X.509]	ITU-T Recommendation X.509 (10/19): Information Technology—Open Systems Interconnection—The Directory: Public-Key and Attribute Certificate Frameworks, October 2019.
[L2VPN]	Business Services over DOCSIS® Layer 2 Virtual Private Networks, CM-SP-L2VPN-I15-150528, May 28, 2015, Cable Television Laboratories, Inc.
[MULPIv3.1]	DOCSIS 3.1 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv3.1-I22-211110, November 10, 2021, Cable Television Laboratories, Inc.
[MULPIv4.0]	DOCSIS 4.0 MAC and Upper Layer Protocols Interface Specification, CM-SP-MULPIv4.0-I04-210826, August 26, 2021, Cable Television Laboratories, Inc.
[PHYv3.1]	DOCSIS 3.1 Physical Layer Specification, CM-SP-PHYv3.1-I19-21110, November 10, 2021, Cable Television Laboratories, Inc.
[PHYv4.0]	DOCSIS 4.0 Physical Layer Specification, CM-SP-PHYv4.I04-210826, August 26, 2021, Cable Television Laboratories, Inc.
[PKT-DCI]	Delivery Function to Collection Function Interface Specification, PKT-SP-ES-DCI-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-DQOS1.5]	PacketCable 1.5 Dynamic Quality of Service Specification, PKT-SP-DQOS1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
[PKT-EM1.5]	PacketCable 1.5 Event Messages Specification, PKT-SP-EM1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
[PKT-ES1.5]	PacketCable 1.5 Electronic Surveillance Specification, PKT-SP-ESP1.5-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
[PKT-ES2.0]	PacketCable 2.0 Electronic Surveillance Intra-Network Functions Specification, PKT-SP-ES-INF-C01-140314, March 14, 2014, Cable Television Laboratories, Inc.
[PKT-MM]	PacketCable Multimedia Specification, PKT-SP-MM-C01-191120, November 20, 2019, Cable Television Laboratories, Inc.
[R-DEPI]	Remote Downstream External PHY Interface Specification, CM-SP-R-DEPI-I15-201207, December 7, 2020, Cable Television Laboratories, Inc.
[R-DTI]	Remote DOCSIS Timing Interface Specification, CM-SP-R-DTI-I08-200323, March 23, 2020, Cable Television Laboratories, Inc.
[R-OOB]	Remote Out-of-Band Specification, CM-SP-R-OOB-I12-200323, March 23, 2020, Cable Television Laboratories, Inc.
[R-OSSI]	Remote PHY OSSI Specification, CM-SP-R-OSSI-I16-210903, September 3, 2021, Cable Television Laboratories, Inc.
[R-PHY]	Remote PHY Specification, CM-SP-R-PHY-I16-210804, August 4, 2021, Cable Television Laboratories, Inc.
[R-UEPI]	Remote Upstream External PHY Interface Specification, CM-SP-R-UEPI-I13-201207, December 7, 2020, Cable Television Laboratories, Inc.
[RFC 868]	IETF RFC 868, Time Protocol, May 1983.
[RFC 1034]	IETF RFC 1034, Domain Names – Concepts and Facilities, November 1987.
[RFC 1035]	IETF RFC 1035, Domain Names – Implementation Specification, November 1987.
[RFC 1191]	IETF RFC 1191, Path MTU Discovery, November 1990.
[RFC 1256]	IETF RFC 1256, ICMP Router Discovery Messages, September 1991.
[RFC 1350]	IETF RFC 1350, The TFTP Protocol (Revision 2), July 1992.

[RFC 1886] IETF RFC 1886, DNS Extensions to support IP version 6, December 1995.

[RFC 1945] IETF RFC 1945, Hypertext Transfer Protocol – HTTP/1.0.

[RFC 1981] IETF RFC 1981, Path MTU Discovery for IP version 6, August 1996.

[RFC 2131] IETF RFC 2131, Dynamic Host Configuration Protocol, March 1997.

[RFC 2132] IETF RFC 2132, DHCP Options and BOOTP Vendor Extensions, March 1997.

[RFC 2315] IETF RFC 2315, PKCS #7: Cryptographic Message Syntax Version 1.5, March 1998.

[RFC 2348] IETF RFC 2348, TFTP Blocksize Option, May 1998.

[RFC 2578] IETF RFC 2578, Structure of Management Information Version 2 (SMIv2), April 1999.

[RFC 2616] IETF RFC 2616, Hypertext Transfer Protocol – HTTP/1.1, June 1999.

[RFC 2626] IETF RFC 2626, The Internet and the Millennium Problem (Year 2000), June 1999.

[RFC 2818] IETF RFC 2818, HTTP Over TLS, May 2000.

[RFC 2874] IETF RFC 2874, DNS Extensions to Support Ipv6 Address Aggregation and Renumbering, July 2000.

[RFC 3046] IETF RFC 3046, DHCP Relay Agent Information Option, January 2001.

[RFC 3164] IETF RFC 3164, The BSD Syslog Protocol, August 2001 (obsoleted by RFC 5424).

[RFC 3268] IETF RFC 3268, AES Ciphersuites for TLS, June 2002.

[RFC 3646] IETF RFC 3646, DNS Configuration options for Dynamic Host Configuration Protocol for IPv6 (DHCPv6), December 2003.

[RFC 3748] IETF RFC 3748, Extensible Authentication Protocol (EAP), June 2004.

[RFC 3924] IETF RFC 3924, Cisco Architecture for Lawful Intercept in IP Networks, October 2004.

[RFC 3931] IETF RFC 3931, Layer Two Tunneling Protocol—Version 3 (L2TPv3), March 2005.

[RFC 4250] IETF RFC 4250, The Secure Shell (SSH) Protocol Assigned Numbers, January 2006.

[RFC 4251] IETF RFC 4251, The Secure Shell (SSH) Protocol Architecture, January 2006.

[RFC 4252] IETF RFC 4252, The Secure Shell (SSH) Authentication Protocol, January 2006.

[RFC 4253] IETF RFC 4253, The Secure Shell (SSH) Transport Layer Protocol, January 2006.

[RFC 4254] IETF RFC 4254, The Secure Shell (SSH) Connection Protocol, January 2006.

[RFC 4291] IETF RFC 4291, IP Version 6 Addressing Architecture, February 2006.

[RFC 4388] IETF RFC 4388, Dynamic Host Configuration Protocol (DHCP) Leasequery.

[RFC 4604] IETF RFC 4604, Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast, August 2006.

[RFC 4649] IETF RFC 4649, Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Relay Agent Remote-ID Option.

[RFC 4704] IETF RFC 4704, The Dynamic Host Configuration Protocol for IPv6 (DHCPv6) Client Fully Qualified Domain Name (FQDN) Option, October 2006.

[RFC 4821] IETF RFC 4821, Packetization Layer Path MTU Discovery, March 2007.

[RFC 4861] IETF RFC 4861, Neighbor Discovery for IP version 6, September 2007.

[RFC 4862] IETF RFC 4862, IPv6 Stateless Address Autoconfiguration, September 2007.

[RFC 4941] IETF RFC 4941, Privacy Extensions for Stateless Address Autoconfiguration in IPv6, September 2007.

[RFC 4994] IETF RFC 4994, DHCPv6 Relay Agent Echo Request Option.

[RFC 5007] IETF RFC 5007, DHCPv6 Leasequery.

[RFC 5246] IETF RFC 5246, The TLS Protocol Version 1.2, August 2008.

[RFC 5656] IETF RFC 5656, Elliptic Curve Algorithm Integration in the Secure Shell Transport Layer, December 2009.

[RFC 5280] IETF RFC 5280, Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, May 2008.

[RFC 5424] IETF RFC 5424, The Syslog Protocol, March 2009.

[RFC 5460] IETF 5460, DHCPv6 Bulk Leasequery.

[RFC 6020] IETF RFC 6020, YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF), October 2010.

[RFC 6032] IETF RFC 6032, Cryptographic Message Syntax (CMS) Encrypted Key Package Content Type, December 2010.

[RFC 6052] IETF RFC 6052, IPv6 Addressing of Ipv4/Ipv6 Translators, October 2010.

[RFC 6144] IETF RFC 6144, Framework for IPv4/Ipv6 Translation, April 2011.

[RFC 6187] IETF RFC 6187, X.509v3 Certificates for Secure Shell Authentication, March 2011.

[RFC 6221] IETF RFC 6221, Lightweight DHCPv6 Relay Agent.

[RFC 6585] IETF RFC 6585, Additional HTTP Status Codes, April 2012.

[RFC 6818] IETF RFC 6818, Updates to the Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile, January 2013.

[RFC 6926] IETF RFC 6926, DHCPv4 Bulk Leasequery.

[RFC 6960] IETF RFC 6960, X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP, June 2013.

[RFC 7568] IETF RFC 7568, Deprecating Secure Sockets Layer Version 3.0, June 2015.

[RFC 8040] IETF RFC 8040, RESTCONF Protocol, January 2017.

[RFC 8415] IETF RFC 8415, Dynamic Host Configuration Protocol for IPv6 (DHCPv6), November 2018.

[RSA 3] RSA Laboratories, PKCS #3: Diffie-Hellman Key Agreement Standard, Version 1.4, RSA Security, Inc., November 1993.

[SCTE 55-1] ANSI/SCTE 55-1 2019, Digital Broadband Delivery System: Out of Band Transport Part 1: Mode A.

[SCTE 55-2] ANSI/SCTE 55-2 2019, Digital Broadband Delivery System: Out of Band Transport Part 2: Mode B.

[SECv3.1] DOCSIS 3.1 Security Specification, CM-SP-SECv3.1-I10-211110, November 10, 2021, Cable Television Laboratories, Inc.

[SECv4.0] DOCSIS 4.0 Security Specification, CM-SP-SECv4.0-I03-210826, August 26, 2021, Cable Television Laboratories, Inc.

[SYNC] Synchronization Techniques for DOCSIS® Technology Specification, CM-SP-SYNC-I02-210407, April 7, 2021, Cable Television Laboratories, Inc.

[TR 101] Migration to Ethernet-Based Broadband Integration, Technical Report 101, Issue 2, Broadband Forum, July 2011.

2.2 Informative References

This specification uses the following informative references.

[CCAP-TR] Converged Cable Access Platform Architecture Technical Report, CM-TR-CCAP-V03-120511, May 11, 2015, Cable Television Laboratories, Inc.

[RFC 8446]	IETF RFC 8446, The Transport Layer Security (TLS) Protocol Version 1.3, August 2018.
[UML Guidelines]	UML Modeling Guidelines, CM-GL-OSS-UML-V01-180627, June 27, 2018, Cable Television Laboratories, Inc.
[XCCF]	CableLabs PNM Combined Common Collection Framework Architecture Technical Report, CL-TR-XCCF-PNM-V01-180814, August 14, 2018, Cable Television Laboratories, Inc.

2.3 Reference Acquisition

- Cable Television Laboratories, Inc.: 858 Coal Creek Circle, Louisville, CO 80027; Phone: +1-303-661-9100; Fax: +1-303-661-9199; http://www.cablelabs.com
- Internet Engineering Task Force (IETF) Secretariat: 48377 Fremont Blvd., Suite 117, Fremont, California 94538, USA; Phone: +1-510-492-4080; Fax: +1-510-492-4001; http://www.ietf.org
- SCTE•ISBE—Society of Cable Telecommunications Engineers Inc.: 140 Philips Road, Exton, PA 19341; Phone: +1-610-363-6888 / 800-542-5040; Fax: +1-610-363-5898; http://www.scte.org/
- Institute of Electrical and Electronics Engineers (IEEE); Phone: +1-800-422-4633 (USA and Canada); http://www.ieee.org
- ITU Recommendations: Place des Nations, CH-1211, Geneva 20, Switzerland; Phone: +41-22-730-51-11; Fax +41-22-733-7256; http://www.itu.int
- MEF—Metro Ethernet Forum: 6033 W. Century Blvd, Suite 830, Los Angeles, CA 90045; Phone: +1-310-642-2800; Fax: +1-310-642-2808; http://metroethernetforum.org
- International Organization for Standardization (ISO) Central Secretariat: Chemin de Blandonnet 8, CP 401, 1214 Vernier, Geneva, Switzerland; www.iso.org

3 TERMS AND DEFINITIONS

This specification uses the following terms.

cable modem	modulator-demodulator at subscriber locations intended for use in conveying data communications on a cable television system
CCAP Core	CCAP device that uses MHAv2 protocols to interconnect to an RPD
Converged Interconnect Network (CIN)	network (generally gigabit Ethernet) that connects a CCAP Core to an RPD or MAC-NE
customer premises equipment (CPE)	equipment at the end user's premises; may be provided by the service provider
DELETE	[RFC 8040] command used in RESTCONF to remove a target resource
DIAMETER	An authentication, authorization, and accounting (AAA) protocol for computer networks. It evolved from the earlier RADIUS protocol and unlike that protocol is not an abbreviation. Diameter Applications extend the base protocol by adding new commands and/or attributes, such as those for use with the Extensible Authentication Protocol (EAP).
Downstream	(1) transmissions from CMTS to CM (2) RF spectrum used to transmit signals from a cable operator's headend or hub site to subscriber locations
Dynamic Host Configuration Protocol (DHCP)	network protocol enabling a server to automatically assign an IP address to a network element
Edge QAM Modulator	headend or hub device that receives packets of digital video or data—It re-packetizes the video or data into an MPEG transport stream and digitally modulates the digital transport stream onto a downstream RF carrier using quadrature amplitude modulation (QAM).
Flow	stream of packets in DEPI used to transport data of a certain priority from the CCAP Core to a particular QAM channel of an RPD or MAC-NE; in PSP operation, several flows can exist per QAM channel
FMA Core Access Device	The headend components of the FMA system including the MAC Manager(s), S-Leaf(s), Video Core(s), OOB Core(s) and NDx Core(s).
GET	[RFC 8040] command used in RESTCONF to retrieve data and metadata for a resource
HEAD	[RFC 8040] command used in RESTCONF to retrieve just the header fields (which contain the metadata for a resource) that would be returned for the comparable GET method, without the response message-body
hybrid fiber/coax (HFC) system	broadband bidirectional shared-media transmission system that uses optical fiber trunks between the headend and the fiber nodes and coaxial cable distribution from the fiber nodes to the customer locations
Institute of Electrical and Electronics Engineers (IEEE)	voluntary organization that, among other things, sponsors standards committees and is accredited by the American National Standards Institute (ANSI)
Internet Engineering Task Force (IETF)	standards body responsible for, among other things, developing standards used in the Internet
Internet Protocol (IP)	Internet network-layer protocol
I/Q Samples	I/Q sampling is a technique used in NDR/NDF Out Of Band (OOB) data processing that represents I/Q data using a cosine function. *I* is the current momentary amplitude of the signal (i.e. the Real signal). *Q* is the momentary amplitude of the signal phase shifted -90°.
L2SS	Layer 2 Specific Sublayer—DEPI is an L2SS of L2TPv3.
L2TP Access Concentrator (LAC)	If an L2TP Control Connection Endpoint (LCCE) is being used to cross-connect an L2TP Session directly to a data link, it is referred to as an L2TP Access Concentrator (LAC). Because an LCCE can act as both an LAC for some sessions and an LNS for others, these terms can only be used within the context of a given set of sessions unless the LCCE is, in fact, single purpose for a given topology.
L2TP Attribute Value Pair (AVP)	L2TP variable-length concatenation of a unique Attribute (represented by an integer), a length field, and a Value containing the actual value identified by the attribute
L2TP Control Connection	reliable control channel used to establish, maintain, and release individual L2TP sessions, as well as the control connection itself

L2TP Control Connection Endpoint (LCCE)	L2TP node that exists at either end of an L2TP Control Connection; can be referred to as an LAC or LNS depending on whether tunneled frames are processed at the data link (LAC) or network layer (LNS)
L2TP Control Connection ID	field that contains the identifier for the control connection, a 32-bit value—The Assigned Control Connection ID AVP, Attribute Type 61, contains the ID being assigned to this control connection by the sender. The Control Connection ID specified in the AVP is required to be included in the Control Connection ID field of all control packets sent to the peer for the lifetime of the control connection. Because a Control Connection ID value of 0 is used in this special manner, the zero value cannot be sent as an Assigned Control Connection ID value.
L2TP Control Message	L2TP message used by the control connection
L2TP Data Message	message used by the data channel
L2TP Endpoint	node that acts as one side of an L2TP tunnel
L2TP Network Server (LNS)	If a given L2TP session is terminated at the L2TP node and the encapsulated network layer (L3) packet is processed on a virtual interface, the L2TP node is referred to as an L2TP Network Server (LNS). Because an LCCE can act as both an LNS for some sessions and an LAC for others, these terms can only be used within the context of a given set of sessions unless the LCCE is, in fact, single purpose for a given topology.
L2TP Pseudowire	emulated circuit as it traverses a packet-switched network; there is one Pseudowire per L2TP Session
L2TP Pseudowire Type	payload type being carried within an L2TP session; examples include PPP, Ethernet, and Frame Relay
L2TP Session	the entity created between two LCCEs in order to exchange parameters for and maintain an emulated L2 connection; multiple sessions may be associated with a single control connection
L2TP Session ID	32-bit field containing a non-zero identifier for a session—L2TP sessions are named by identifiers that have local significance only. That is, the same logical session will be given different Session IDs by each end of the control connection for the life of the session. When the L2TP Control Connection is used for session establishment, Session IDs are selected and exchanged as Local Session ID AVPs during the creation of a session. The Session ID alone provides the necessary context for all further packet processing, including the presence, size, and value of the Cookie; the type of L2SS; and the type of payload being tunneled.
MAC Domain	grouping of Layer 2 devices that can communicate with each other without using bridging or routing—In DOCSIS, it is the group of CMs that are using upstream and downstream channels linked together through a MAC forwarding entity.
MAC Manager	(1) management entity that represents a single point of contact for operator back-office management systems to FMA access network elements within its scope of control (2) FMA functional entity that uses FMA protocols to manage a MAC-NE
MAC-NE Unique Identifier	The MAC-NE Unique Identifier serves as the principal identifier for the MAC-NE in the FMA management system. The MAC-NE Unique Identifier is unique within a management domain. The MAC-NE Unique Identifier is a string with a format similar to a IEEE 802 MAC address. However, this string does not need to be an actual MAC address. It is recommended that for appliance-based MAC-NE devices, a unique ID beginning with a vendor OUI and ending with a unique value are used to allow for deployment without the need for operator pre-deployment configuration.
Maximum Transmission Unit (MTU)	maximum size of the Layer 3 payload of a Layer 2 frame
Media Access Control (MAC)	refers to the Layer 2 element of the system, which would include DOCSIS framing and signaling
modulation error ratio (MER)	ratio of average signal constellation power to average constellation error power—that is, digital complex baseband signal-to-noise ratio—expressed in decibels (equal to $10\log_{10}(P_{OUT}/P_{IN})$
Multiple System Operator (MSO)	corporate entity that owns and/or operates more than one cable system
OPTIONS	[RFC 8040] command method used in RESTCONF to discover which methods are supported by the server for a given target resource (e.g., GET, POST, DELETE)
packet identifier (PID)	(system) unique integer value used to identify elementary streams of a program in a single or multi-program transport stream as described in [ISO 13818.1], section 2.4.3 (also published as ITU-T Rec. H.222.0)
pilot tones	required in the HFC network to ensure that amplifiers in the network are operating correctly—Amplifiers use these tones to adjust gain and keep signals at the appropriate output level.

PacketCable Aggregator	control plane entity that represents a single point of contact for operator back-office PacketCable and Policy Server systems to FMA access network elements within its scope of control
POST	[RFC 8040] command method used in RESTCONF creating a data resource or invoking an operation resource
Precision Time Protocol (PTP)	protocol used to synchronize clocks throughout a network
PUT	[RFC 8040] command method used in RESTCONF to create or replace the target data resource
Pseudowire	IP tunnel between two points in an IP network
QAM Channel	analog RF channel that uses quadrature amplitude modulation (QAM) to convey information
quadrature amplitude modulation	modulation technique in which an analog signal's amplitude and phase vary to convey information, such as digital data
radio frequency	in cable television systems, refers to electromagnetic signals in the range of 5–1000 megahertz
Radio Frequency Interface	term encompassing the downstream and upstream radio frequency interfaces
Request for Comments	technical policy document of the IETF; these documents are available online at http://www.rfc-editor.org/
Remote MAC Core (RMC)	device in the network that implements the FMA specifications to provide MHAv2 DOCSIS Core functionality closer to the RPD than is possible in the MHAv2 architectural model
Remote-MACPHY Device (RMD)	a device in the network that implements the FMA specifications to provide conversion from digital Ethernet transport to analog RF transport and is a kind of MAC-NE
Remote-PHY Device (RPD)	device in the network that implements the Remote-PHY specification ([R-PHY]) to provide conversion from digital Ethernet transport to analog RF transport
Session	L2TP data plane connection from the CCAP Core to the QAM Channel—There can only be one session per QAM Channel. There is one DEPI pseudowire type per session. There may be one MPT flow or one or more PSP flows per session. Multiple sessions may be bound to a single control connection.
SLAAC	method of IPv6 addressing whereby a client (RPD, in this case) addresses itself based on prefix information advertised onto the network by a router
Streaming Telemetry Client	The part of the system receiving network telemetry from a streaming telemetry server.
Streaming Telemetry Server	The part of the system providing network telemetry to the streaming telemetry client.
Trivial File Transfer Protocol (TFTP)	file transfer protocol generally used for automated transfer of configuration or boot files between machines
upstream	(1) transmissions from CM to CMTS (2) RF spectrum used to transmit signals from a subscriber location to a cable operator's headend or hub site
upstream channel descriptor	MAC Management Message used to communicate the characteristics of the upstream physical layer to the cable modems

4 ABBREVIATIONS AND ACRONYMS

This specification uses the following abbreviations.

1RU	1 Rack Unit
AAA	Network Authentication, Authorization, and Accounting
ACK	L2TPv3 Explicit Acknowledgement message
ADC	analog-to-digital converter
AGC	automatic gain control
AKE	authenticated key agreement
AM	Application Manager
API	application programming interface
ARP	Address Resolution Protocol (IPv4)
ARPD	Advanced Return Path Demodulator
ATM	Asynchronous Transfer Mode
AVP	L2TPv3 Attribute Value Pair
B2B	Business to Business
BBF	Broadband Forum
BC	Boundary Clock
BPI	Baseline Privacy Interface
BSoD	Business Services over DOCSIS
BSS	Business Support Systems
BW-REQ	Bandwidth Request message
C-TAG	customer tag
C-VLAN ID	customer VLAN identifier
CA	certificate authority
CAA	Centralized Access Architecture
CAK	Connectivity Association Key
CCAP	Converged Cable Access Platform
CE	Customer Edge
CIN	Converged Interconnect Network
CLI	command-line interface
CM	cable modem
CMS	Call Management Server
CMTS	Cable Modem Termination System
COPS	Common Open Policy Service
CPE	customer premises equipment
CPI	controller plane interface
CRC	cyclic redundancy check
CRL	Certificate Revocation List
CSMA	Carrier Sense Multiple Access
C-TAG	Customer VLAN Tag
CVC	Code Verification Certificate
CW	Continuous Wave
CWDM	Coarse Wavelength-Division Multiplexing
D-MPT	DOCSIS MPT Mode
DAA	Distributed Access Architecture

DAC	Digital-to-Analog Converter
DAD	Duplicate Address Detection
DCA	Distributed CCAP Architecture
DCI	DOCSIS controller interface
DEPI	Downstream External PHY Interface
DF	Don't Fragment (bit)
DHCP	Dynamic Host Configuration Protocol
DHCPv4	Dynamic Host Configuration Protocol version 4
DHCPv6	Dynamic Host Configuration Protocol version 6
DHCT	Digital Home Communications Terminal
DMM	Delay Measurement
DNCS	Digital Network Control System
DNS	Domain Name System
DOCSIS	Data-Over-Cable Service Interface Specifications
DP	Data Plane; Destination Port
DPD	Downstream Profile Descriptor
DQoS	Dynamic Quality of Service
DRFI	Downstream Radio Frequency Interface
DS	Downstream
DSA	Dynamic Service Add
DSCP	differentiated services code point
DSC	Dynamic Service Change
DSD	Dynamic Service Delete
DSG	DOCSIS Set-top Gateway
DSI	DOCSIS-side interface
DSP	Digital Signal Processor
DST	Destination
DSX	Dynamic Service "X" (DSA, DSC, DSD)
DTI	DOCSIS Timing Interface
DTP	DOCSIS Time Protocol
DUID	DHCP Unique Identifier
DUT	Device Under Test
DWDM	dense wavelength-division multiplexing
E-MTA	Embedded Multimedia Terminal Adapter
E-NNI	external network-to-network interface
EAP	Extensible Authentication Protocol
EAPoL	Extensible Authentication Protocol over LAN
EAS	Emergency Alert System
eCM	Embedded Cable Modem
EEC	Ethernet Equipment Clock
EPG	Electronic Program Guide
EPON	Ethernet passive optical network
EQAM	Edge QAM
ERM	Edge Resource Manager
ES	Electronic Surveillance
ESP	Encapsulating Security Payloads
ETSI	European Telecommunications Standards Institute

EVPL	Ethernet virtual private line
FCS	Frame Check Sequence
FDM	Frequency Division Multiplex
FDX	Full Duplex
FEC	Forward Error Correction
FMA	Flexible MAC Architecture
FQDN	Fully Qualified Domain Name
FSM	Finite State Machine
FTP	File Transfer Protocol
Gbps	gigabit per second
GCP	generic control plane
GE	Gigabit Ethernet (Gig E)
GM	grand master
GNSS	Global Network Satellite System
GPON	gigabit passive optical network
GPS	Global Positioning System
HDLC	High-Level Data Link Control
HERD	Headend Re-architected as a Datacenter
HFC	hybrid fiber-coax
HMAC	hash-based message authentication code
HP	homes passed
HRM	hypothetical reference model
HSD	high-speed data
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IAP	intercept access point
I-CCAP	integrated CCAP
I-CMTS	Integrated CMTS
I-NNI	internal network-to-network interface
ICMP	Internet Control Message Protocol
ID	Identifier
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
IGMP	Internet Group Management Protocol
IGP	Interior Gateway Protocol
IKE	Internet Key Exchange
IP	Internet Protocol
IRA	Identification and Resource Advertising
IPDR	Internet Protocol Detail Record
IPDR/SP	Internet Protocol Detail Record/Streaming Protocol
IPPV	Impulse Pay Per View
IPsec	Internet Protocol Security
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
ISO	International Organization for Standardization
ITU	International Telecommunication Union
ITU-T	Telecommunication Standardization Sector of the International Telecommunication Union

IWF	Internetworking Function
Km	kilometer
L2	Layer 2
L2PDU	Layer 2 Protocol Data Unit
L2SS	Layer 2 Specific Sublayer
L2TP	Layer 2 Tunneling Protocol
L2TPv3	Layer 2 Tunneling Protocol version 3
L2VPN	Layer 2 VPN
L3	Layer 3
L3VPN	Layer 3 VPN
LAC	L2TP Access Concentrator
LACP	Link Aggregation Control Protocol
LAG	link aggregation group
LAN	local area network
LCCE	L2TP Control Connection Endpoint
LI	Lawful Intercept
LLD	Low Latency DOCSIS
LLDP	Link Layer Discovery Protocol
LNS	L2TP Network Server
LQ	Leasequery
LTE TDD	Long Term Evolution Time Division Duplexing
MAC	Media Access Control
MAC-NE	MAC network element
MACPHY	Media Access Control and physical layers
MACsec	Media Access Control Security
MAP	Upstream Bandwidth Allocation Map
MSI	Main Software Image
MDU	Multiple Dwelling Unit
MEF	Metro Ethernet Forum
MER	Modulation Error Ratio
MHA	Modular Headend Architecture
MHAv2	Modular Headend Architecture version 2
MIB	Management Information Base
MKA	MACsec Key Agreement (protocol)
MLD	Multicast Listener Discovery
MMI	MAC Manager Interface
MPEG	Moving Picture Experts Group
MPEG-TS	Moving Picture Experts Group–Transport Stream
MPLS	Multiprotocol Label Switching
MPT	MPEG-TS mode of R-DEPI
MPTS	Multi Program Transport Stream
MSK	master secret key
MSO	Multiple System Operator
MTU	maximum transmission unit
MXh	Mobile Xhaul
NAD	Network Access Device; Network Authenticator Device
NAT	network address translation

NC	Network Controller
NCS	Network-based Call Signaling
ND	Neighbor Discovery
NDF	Narrowband Digital Forward
NDR	Narrowband Digital Return
NDx	refers to NDF, NDR, or both
NE	network element
NETCONF	Network Configuration Protocol
NFV	network function virtualization
NMS	Network Management System
NNI	network-to-network interface
NPR	Noise Power Ratio
NS	Neighbor Solicitation
Ns	nanosecond
NSI	network-side interface
NTP	network time protocol
OCC	optimistic concurrency control
OCD	OFDM Channel Descriptor
OFDM	orthogonal frequency-division multiplexing
OFDMA	orthogonal frequency-division multiple access
OLT	optical line terminal
ONU	optical network unit
OOB	out-of-band
OSP	outside plant
OSS	Operations Support System
OSSI	Operations Support System Interface
PAG	PacketCable Aggregator
PAE	Port Access Entity
PAI	PacketCable Aggregator Interface
PCEM	PacketCable Event Messages
PCES	PacketCable Electronic Surveillance
PCMM	PacketCable Multimedia
pCore	Physical Core
PDV	packet delay variation
PE	Provider Edge
PHB	Per Hop Behavior
PHB-ID	per hop behavior identifier
PHY	Physical Layer
PID	Packet Identifier
PIO	Prefix Information Option
PKI	public key infrastructure
PL	payload length
PLC	PHY Link Channel
PMA	Profile Management Application
PMTUD	Path MTU Discovery
PNF	Physical Network Function
PNM	Proactive Network Maintenance

PON	Passive Optical Network
PoP	Point of Presence
P-PoP	Primary Point of Presence
PoT	Point of Termination
PPP	Point-to-Point Protocol
PRC	Primary Reference Clock
PRTC	Primary Reference Time Clock
PS	Policy Server
PSP	Packet Streaming Protocol
PTP	Precision Time Protocol
PtP WDM	point-to-point wavelength-division multiplexing
PW	Pseudowire
QAM	quadrature amplitude modulation
QoS	quality of service
R-DEPI	Remote Downstream External PHY Interface
R-MACPHY	Remote MACPHY
R-OSSI	Remote OSSI
R-PHY	Remote PHY
R-UEPI	Remote Upstream External PHY Interface
RA	Router Advertisement
RADD	Remote Addressable DANIS/DLS
RADIUS	Remote Authentication Dial-In User Service
RCP	R-PHY Control Protocol
RDC	regional data center
RESTCONF	Representational State Transfer Configuration Protocol
RESTful	An API that uses HTTP requests to GET, PUT, POST and DELETE data
RF	radio frequency
RFI	Radio Frequency Interface
RFC	Request for Comments
RKS	Record Keeping Server
RMC	Remote MAC Core
RMD	Remote MACPHY Device
RMN	Remote MACPHY Node
RMS	Remote MACPHY Shelf
RNG-REQ	Ranging Request message
RPD	Remote PHY Device
RSA	Rivest-Shamir-Adleman (cryptosystem)
RX	Receive
RxMER	receive modulation error ratio
S-Leaf	Secure Leaf
S-PoP	Secondary Point of Presence
S-TAG	service VLAN tag
S-VLAN ID	service VLAN identifier
S/C-PCP	service/customer priority code point
SA	Security Association
SC-QAM	Single Carrier QAM
SCTE	Society of Cable Telecommunications Engineers

SDN	software-defined networking
SFD	Start Frame Delimiter
SFP+	enhanced small form-factor pluggable
SFTP	Secure File Transfer Protocol
SI	Source Interface
SID	Service Identifier
SIP	Session Initiation Protocol
SLAAC	Stateless Address Autoconfiguration
SLM	synthetic loss measurement
SNMP	Simple Network Management Protocol
SNTP	Simple Network Time Protocol
SP	Source Port
SpecMAN	Spectrum Management (pseudowire)
SPTS	Single Program Transport Stream
SRC	Source
SRMA	S-Leaf Router MAC Address
SSD	secure software download
SSH	Secure Shell
SSL	Secure Socket Layer
STB	set-top box
SyncE	Synchronous Ethernet
T-BC	Telecom Boundary Clock
T-GM	Telecom GrandMaster
T-TSC	Telecom Time Slave Clock
TBD	To Be Determined
TCP	Transmission Control Protocol
TDM	time division multiplex
TE	Time Error
TFTP	Trivial File Transfer Protocol
TLS	Transport Layer Security
TLV	type-length-value
ToD	Time of Day
ToR	Top of Rack
TPID	Tag Protocol Identifier
TR MB	Trigger Message Block
TWDM	time- and wavelength-division multiplexing
TX	Transmit
UA	User Agent
UCD	upstream channel descriptor
UDP	User Datagram Protocol
UEPI	Upstream External PHY Interface
UML	Unified Modeling Language
UNI	user-network interface
URL	Uniform Resource Locator
US	Upstream
USM	user-based security model
UTC	Coordinated Universal Time

UTSA	Upstream Triggered Spectrum Analysis
UTSC	Upstream Triggered Spectrum Capture
vCM	Virtual Cable Modem
vCore	Virtual Core
VoIP	Voice over IP
VLAN	virtual LAN
VNF	virtual network function
VPN	virtual private network
VSA	vector signal analyzer
WAN	Wide Area Network
WDM	Wave Division Multiplexing
XCCF	Combined Common Collection Framework
xPON	refers to EPON, GPON, or both
YANG	Yet Another Next Generation
ZBL	Zero Bit Loading

5 FMA SYSTEM OVERVIEW

5.1 MPEG Video

The support for traditional digital video using MPEG-TS will be supported in the Flexible MAC Architecture (FMA). The delivery of video services in FMA will share those requirements defined in the Remote PHY specification [R-PHY] to every degree possible.

5.2 OOB Communications

When MAC-NEs are deployed in certain regions (particularly North America), support for the downstream and upstream legacy set-top box (STB) radio frequency (RF) control channels is required. These control channels are often referred to as OOB, or out-of-band, because these narrow-band RF channels normally lie outside of normal quadrature amplitude modulation (QAM) video or DOCSIS data RF channels.

5.3 IP-Based Residential Services

The support for IP-based residential services such as high-speed data service or residential internet service, voice service, and IP video service are all required to be carried over the FMA.

5.4 PacketCable

Support for PacketCable will be provided in FMA. An FMA deployment may require PacketCable Dynamic Quality of Service (DQoS), PacketCable Multimedia (PCMM), PacketCable Event Messages (PCEM), any combination of the three, or no PacketCable support at all. FMA will support PacketCable functionality without impacting legacy Call Management Server (CMS), PCMM Policy Server, and Record Keeping Server (RKS) applications and interfaces that currently provide it.

5.5 Lawful Intercept (LI)

Lawful Intercept (LI), also commonly referred to as electronic surveillance (ES), is a process in which a service provider or network operator collects and provides authorized law enforcement agencies with intercepted communications of private individuals or organizations.

Detailed requirements for LI can differ from one jurisdiction to another, but the general requirements are the same. For example, the LI system provides transparent interception of specified traffic only, and the subject cannot be aware of the interception. The service provided to other users cannot be affected during interception to ensure service stability and ensure that the subject cannot determine if and when interception happens. In FMA, Lawful Intercept can be performed by the MAC-NE or an S-Leaf or can be done elsewhere in the Operator’s network.

5.6 Business Services over DOCSIS - L2VPN

In today's environment, a high demand for reliability, Quality of Service, and Quality of Experience associated with different applications - voice, data, video - means that DOCSIS L2VPN [L2VPN] is closely associated with carrier-grade Ethernet/Carrier Ethernet services. For carrier-grade Ethernet, the Metro Ethernet Forum (MEF) provides a common framework that many network operators leverage. The MEF defines attributes such as the user-network interface (UNI), which is typically the Ethernet port of a cable modem, and network-to-network interfaces (NNIs), of which the internal NNI (I-NNI) is nominally the MAC-NE Ethernet port (network-side interface) or MAC Core network-side interface.

Achieving this provider bridging/service provider model as per the MEF framework requires support for [IEEE 802.1Qcy] (or "Q-in-Q") tagging to allow end-user virtual circuits to be carried within service-provider virtual circuits. At the I-NNI, the inner [IEEE 802.1Q] C-TAG VLAN ID (C-VID) field (from the C-TAG) is used to address each individual access virtual circuit within a connectivity virtual circuit represented by the outer [IEEE 802.1Qcy] S-TAG VLAN ID (S-VID) (of the S-TAG. See Figure 5. Note that the same C-VID may appear through

different S-TAGs on a given I-NNI. In such cases, the C-VIDs always address different cable modems. The S/C-PCP field is used to communicate priority information both across the UNI/NNI boundaries and within the network provider's network.

Flexibility is provided regarding the allocation of VLAN tags:

- S-TAG and C-TAG values could be assigned by the network/service operator; and
- S-TAG values could be assigned by the network/service operator, and inner C-TAG values could be assigned/managed by the customer.

Each VLAN tag at the I-NNI represents an L2VPN Business Services over DOCSIS (BSoD) instance. Therefore, within the cable modem, configuration file TLV 43.5 is used to pass the L2VPN information. For example, TLV 43.5.1 contains the L2VPN encoding, and TLV 43.5.2.3 contains the double-tagged [IEEE 802.1Qcy] VLAN IDs.

For example, 0x09FA00AD (hex) indicates C-VID = 2554 (decimal) and S-VID = 173 (decimal).

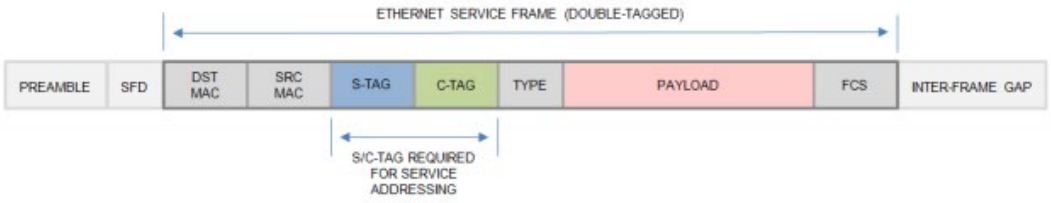

Figure 5 - Double-Tagged Ethernet Service Frame

5.7 Mobile Xhaul (MXh) Transport Service

Mobile wireless operators are looking to access providers for backhaul transport services than can support both legacy technologies and the newest radio technologies. Cable operators' hybrid fiber-coax (HFC) networks that use DOCSIS-based systems seem to be a perfect fit because they will improve reachability and increase capacity. The new radio technology may be microcell, picocell, or femtocell, but it will help mobile operators improve wireless services. As the volume of radio deployment increases, so will the number of network connections available for backhaul services.

One critical aspect of mobile base stations and broadcast transmitters is their need to synchronize the timing of the radio interfaces. Accurate frequency synchronization among base stations allows user handsets to seamlessly handover between base stations, reduces interference between cells, and optimizes radio bandwidth capacity. Another critical aspect of some base stations is the need for phase or time-of-day (ToD) synchronization. If cable operators are to be utilized for wireless backhaul services, the FMA elements needs to support precision time synchronization and latency requirements.

6 TECHNICAL OVERVIEW

6.1 Introduction

In a Flexible MAC Architecture, the classic integrated CCAP (I-CCAP) is separated into multiple distinct core and digital access network components.

The core components of the FMA are MAC Manager, PacketCable Aggregator, Video Cores, OOB Cores, Core Timing Devices, and Core Access Devices. The MAC Manager serves as the focal point between an operator's Operations Support System (OSS) and back-office management system and a potentially large number of MAC-NEs deployed in the access network. The PacketCable Aggregator serves as the focal point between an operator's PacketCable back-office system and a potentially large number of MAC-NEs deployed in the access network. The Video Core contains all video processing functions that an Edge QAM (EQAM) provides today, with the exception of video PHY and QAM-related processing in the MAC-NE. The OOB Core contains the functionality required to manage and control legacy set-top boxes (STBs) but in a way that is compatible with digital optical transport and MAC-NEs. OOB Cores have been defined that uniquely support requirements from [SCTE 55-1] or [SCTE 55-2]. There are also OOB Cores that support Narrowband Digital Forward (NDF) and Narrowband Digital Return (NDR), which can transparently pass OOB and RF signals over digital transport. These OOB Cores can be used as substitutes for SCTE 55-1 and SCTE 55-2 OOB Cores in suitable environments.

The access component of FMA is the MAC-NE.

The MAC-NE is an intelligent access device with the following functions:

- to provide DOCSIS MAC and the upper layer protocols, including all DOCSIS signaling functions, downstream and upstream bandwidth scheduling, and DOCSIS framing;
- to provide all PHY-related circuitry needed for HFC high-speed data and video services, such as downstream QAM modulators and digital-to-analog converters (DACs) and upstream analog-to-digital converters (ADCs) and QAM demodulators;
- to convert downstream MPEG video and OOB signals received from an Auxiliary Core over a digital medium such as Ethernet or a passive optical network (PON) into analog for transmission over RF or linear optics;
- to convert upstream OOB signals received from an analog medium such as RF or linear optics into digital for transmission over Ethernet or PON to an Auxiliary Core; and
- to provide pseudowire logic to connect to Auxiliary Cores.

The DOCSIS functionality of the MAC-NE is defined by [PHYv4.0], [MULPIv4.0], and [DRFI].

The MAC-NE also provides a subset of the following external interfaces:

CIN Facing

- One or more 10G, 40G Ethernet or PON ports

Access Network Facing

- One or more 10G Ethernet or higher bit rate ports or PON ports.
- One or more RF ports providing connectivity to the access network. RF ports may be unidirectional (for an external combiner) or bidirectional (for an internal combiner). RF port output may be over coaxial cable or over analog optics.

An example reference implementation based on Ethernet is shown in Figure 6.

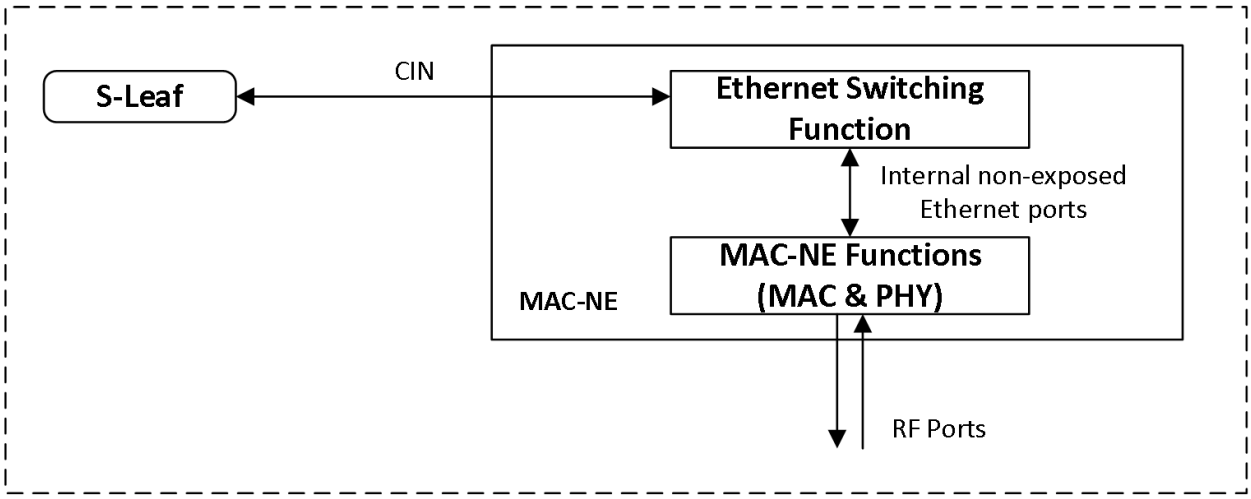

Figure 6 - MAC-NE Internal and External Connections

The FMA specifications describe how the core components and the MAC-NE interface with each other. Note that FMA functionality and signaling and DOCSIS functionality and signaling are completely separate. DOCSIS functionality and signaling remain the same for both I-CCAP and FMA solutions. FMA focuses on a simple deconstruction of the CCAP that moves the CCAP PHY and DOCSIS MAC elements into an external MAC-NE while keeping the DOCSIS CMTS-to-CM signaling untouched.

6.2 System Diagram

Figure 7 shows an abstracted view of a cable operator's network. Note that there are more aggregation points beyond a headend, such as a super headend or a regional data center (RDC). For the scope of this specification, the focus will be on the headend aggregation point.

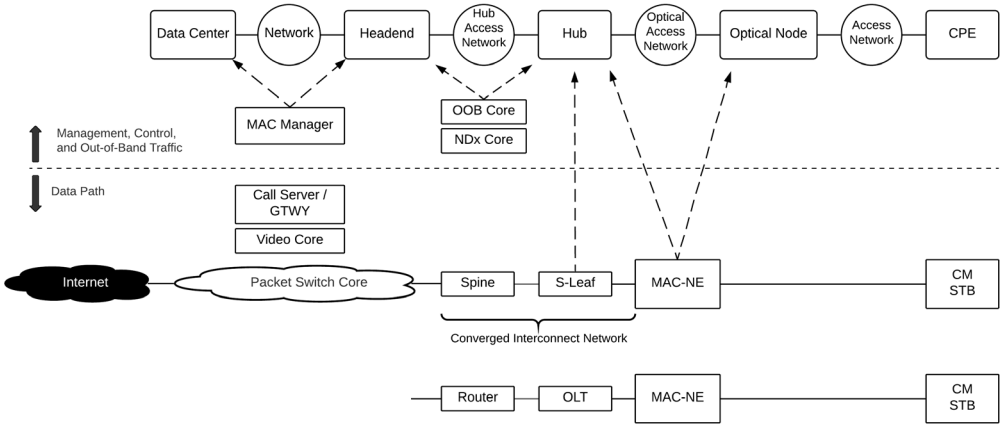

Figure 7 - FMA High-Level System Diagram with MAC-NE

The headend is where most of the equipment that does not require direct connectivity to the access network resides. Video channel line-ups are often created in the headend. The headend aggregates several hubs, with each hub containing equipment that requires direct connectivity to the HFC plant. One example of equipment in the hub is the I-CCAP. The hub aggregates several optical nodes, which are in the field and convert between a long point-to-point optical run and a local coax network. The optical node aggregates traffic from several subscriber endpoints such as DOCSIS cable modems (CMs) and video STBs.

6.2.1 Hub Access Network

The hub access network connects the headend and the hub. It can be either a switched L2 network or a routed L3 network. It typically is a multi-hop network, which means there can be multiple switches and/or routers between equipment in the headend and the hub.

6.2.2 Optical Access Network

The optical access network connects the hub and the optical node. It has a forward path and a reverse path.

6.2.2.1 Using Linear Optics

The classic HFC plant uses linear optics where the RF spectrum from the coax is modulated onto an optical wavelength. The only type of signal that can traverse this type of network is an RF modulated signal such as analog channels, QAM, or OFDM (orthogonal frequency-division multiplexing) signals.

6.2.2.2 Using Digital Optics Only

A variation of the classic HFC plant uses digital optics in the return path. The RF spectrum is sampled and digitized at the optical node, sent to the headend, and then reconstructed into an analog signal. From the viewpoint of this specification, this variation will be considered to be a subset of a linear optics HFC plant because its operation is transparent to the transmission path, which is still a modulated signal such as QAM or OFDMA (orthogonal frequency-division multiple access).

6.2.2.3 Using Digital Optics with IP

A new HFC plant architecture is available that can use any fiber-compatible baseband networking technology, such as Ethernet or PON, to drive the fiber portion of the HFC plant. However, the coax portion of the HFC plant remains the same. With digital optics, the optical access network could be directly connected from the core elements to the optical node. Because the hub is aggregating many optical nodes, the access network may contain one or more network elements, which could be L2 switches or L3 routers. Note that this model includes network elements that may be physically located at the hub but are connected between the core components and the optical node.

One of the goals of MHAv2 and FMA is to accommodate this new digital IP-based HFC plant architecture while minimizing impact on the CCAP definition and operation. In this manner, I-CCAP, MHAv2, and FMA implementations may be used as needed for different HFC plant architectures while maintaining a common CCAP feature set and software loads.

6.2.3 Coax Access Network

The coax portion of the network can be an FDD (frequency division duplexing) plant that carries RF-modulated signals with an upper frequency bound and a frequency range that is split between the upstream and downstream spectrums. The coax portion of the network can alternatively be a mixed FDD and FDX (full duplex) plant that has a lower frequency range dedicated to upstream spectrum, a middle frequency range dedicated to DOCSIS FDX, and an upper frequency range dedicated to downstream spectrum.

6.2.4 Location of the MAC-NE and RF Requirements

The MAC-NE contains network interfaces such as Ethernet or PON, the DOCSIS MAC functions, DOCSIS PHY functions, QAM video PHY functions, and RF functions contained in a network element. These functions can be split within the network element, but the MAC-NE contains all of these functions. The location of the MAC-NE is determined by the form factor it takes.

If the MAC-NE is placed in the outside plant (OSP), such as on a strand or in a vault, then the MAC-NE functions would be placed in node enclosure; this MAC-NE may be referred to as an RMD-Node or RMN. If the MAC-NE is placed in a headend, hub, or cabinet, then the device is referred to as an RMD-Shelf or RMS.

A MAC-NE MUST support RF requirements as described in [PHYv4.0]; [DRFI]; [DRFI], Annex D, "Additions and Modifications for Remote PHY Device for DOCSIS and EuroDOCSIS;" and [R-PHY], Annex E, "Receive Power Level Management," and Annex F, "DOCSIS 4.0 OFDM Modifications for Remote PHY."

6.2.5 Location of the MAC Network Element Functions

Other MAC Network Elements include Physical Cores (pCores), and Virtual Cores (vCores). The DOCSIS MAC functions and possibly other upper layer functions as defined in [MULPIv4.0] are contained in these MAC Network Elements, whereas the DOCSIS PHY functions, QAM video PHY functions, and RF functions are contained in the RPDs attached to them.

A pCore, or vCore could be located at a headend, an HFC hub, or an optical node. A pCore or vCore, on the other hand, is a shelf device in a trusted facility. Further details of these network elements as well as their locations are left for further studies.

An RPD attached to these network elements could be located at an HFC hub or a field location.

6.2.6 Location of the FMA Components

An I-CCAP and its associated OOB equipment is located at the headend or at the hub where the RF ports can have direct connectivity to the access network. This restriction does not apply in FMA because the RF ports are in the optical node devices. The functional entities of the core can be located at the hub or headend (or at another location beyond the headend, like the regional data center) and need not be collocated, but they will generally be referred to collectively in this document.

The CIN is between the core components and the MAC-NE. It encompasses a Secure Leaf (S-Leaf) or Leaf, which may provide L2 or L3 aggregation for the DAA network elements such as RPD or MAC-NE. Additionally, the CIN could also include a spine switch that interconnects many leaf switches for the purposes of aggregating access layer network components together.

6.3 System Architecture

6.3.1 System Components

The reference architecture for an FMA system is shown in Figure 8. This architecture contains both physical and logical components/functions and defines both physical and logical interfaces. This section briefly introduces each component and interface. Certain components/functions are referred to collectively as core components, but this is merely a convenience for the purposes of writing this specification. Core components may or may not be physically collocated, and no particular implementation is implied by this reference architecture.

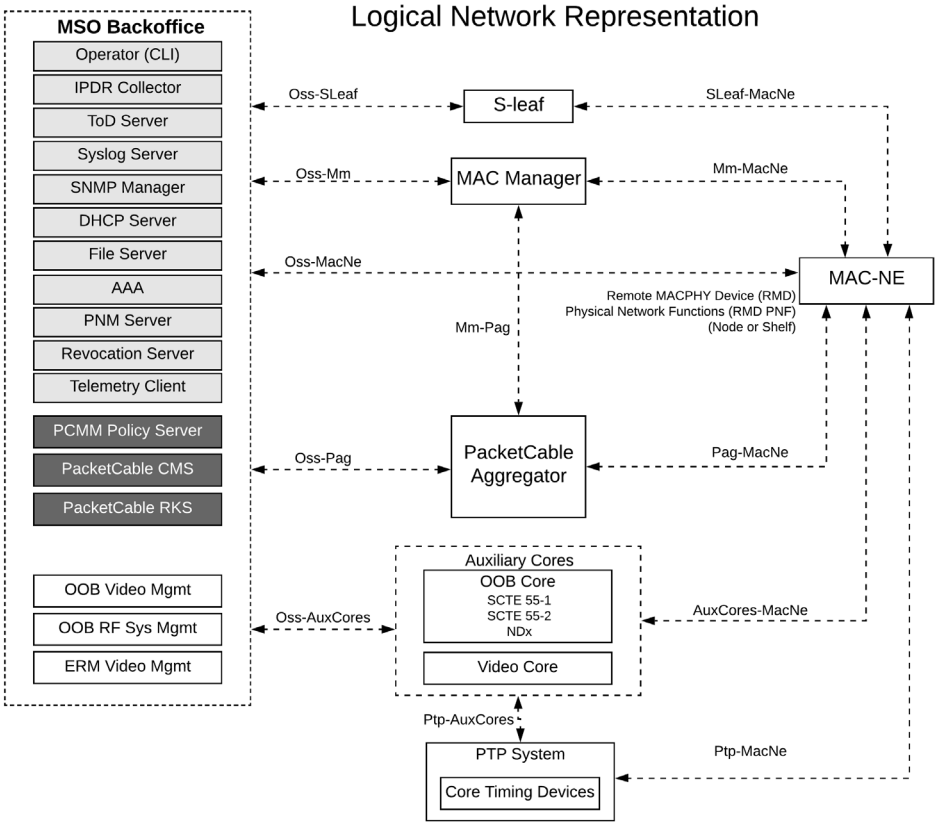

Figure 8 - FMA Phase 1 Reference Architecture

6.3.1.1 MAC Manager

MAC Manager is a functional entity that uses FMA management plane protocols to manage one or more MAC-NEs.

6.3.1.2 PacketCable Aggregator (PAG)

PacketCable Aggregator is an optional functional entity that supports PacketCable and the PacketCable Multimedia back-office element to MAC-NE scalability. In MHAv2, this scale is assumed to be handled by the CMTS/DOCSIS Core, which translates PacketCable and PacketCable Multimedia northbound interfaces to and from associated DOCSIS control plane signaling interfaces.

6.3.1.3 Auxiliary Cores

Auxiliary Cores are Video Core, OOB Core, or NDx Core. These core entities may be referred to collectively and in conjunction with the associated MAC Manager and PacketCable Aggregator as FMA core components.

Video Core contains all video processing functions that an EQAM provides today except for video PHY and RF-related processing, which are in the MAC-NE.

OOB Core supports legacy STBs. There is a unique OOB Core defined for SCTE 55-1 support and another for SCTE 55-2 support.

NDx Core supports NDF and NDR capabilities (1) by digitally sampling downstream RF traffic at the core and remodulating it into RF at the MAC-NE (NDF) and (2) by demodulating and digitally sampling upstream RF traffic at the MAC-NE and remodulating it into RF at the core (NDR).

Throughout this document, "OOB Core" is used to refer to the OOB Core, the NDx Core, or both. Only when the difference is significant will it be explicitly noted.

Core Access Devices collectively provide switched connectivity between the Core Network and the digital optical transport network. In practice, any routed or switched design is possible.

6.3.1.4 MAC Network Elements (MAC-NE)

MAC-NEs are network devices in one of two classes: physical network functions (PNFs) or virtual network functions (VNFs). In this case, the FMA MAC-NEs have been split into four device types, with three in the PNF class and one in the VNF class. The network elements in the PNF class include the RMD and the pCore. The network element in the VNF class is the vCore. The PNF network elements are purpose-built network platforms, and the VNF network elements are common off-the-shelf servers in most cases. Figure 9 illustrates the four FMA MAC-NE device types.

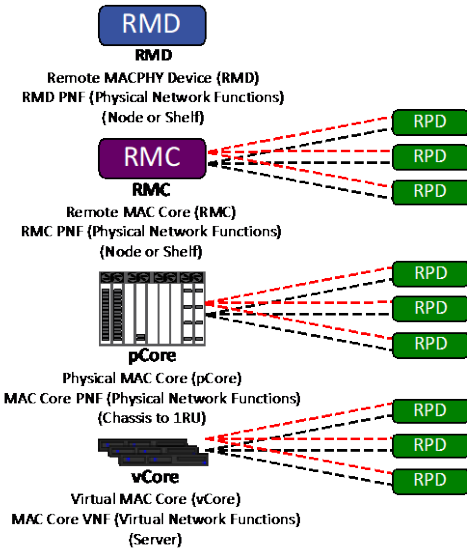

Figure 9 - MAC Network Elements (MAC-NEs)

A PNF is a device built to perform a specific network function, such as an RMD that performs DOCSIS MAC and PHY functions. Other examples of PNFs include a pCore. PNFs can also take the form of chassis-based hardware found in cable operator facilities, which could be network elements of 1 rack unit (1RU) or larger.

MAC-NEs could also take the form of VNFs, such as a vCore. The MAC functions at minimum could include higher layer functions and could also operate on off-the-shelf servers found in data centers.

The RMD has a digital optical network interface on one side and an RF interface on the other side (not shown in Figure 9). The RMD provides a PHY for all services and provides the DOCSIS MAC and upper layers for high-speed data, VoIP, and IP video services and for native transport and pseudowire capabilities for CIN transport between Core Access Devices and the RMD. The RMD also provides OOB translation to and from the network-side digital encapsulation and drop-side RF. The RMD RF output could be RF combined with other overlay services such as analog or digital video services.

6.3.1.5 Precision Time Protocol (PTP) System

The IEEE 1588 Precision Time Protocol system distributes timing synchronization from a source such as a Global Network Satellite System or Global Positioning System to MAC-NE applications.

Core Timing Devices support the distribution of timing and synchronization between FMA elements.

6.3.1.6 Secure Leaf (S-Leaf)

Secure Leaf (S-Leaf) and Leaf provide the access point to the digital optical transport to the MAC-NE. The term Secure Leaf signifies that this is not a traditional Leaf switch and will have added security, such as the use of MACsec for the connection with the MAC-NE. Figure 10 shows an architecture with Spine and S-Leaf switches as well as a placeholder for the Leaf to connect to management entities. In later phases, the pCore and vCore MAC-NEs will be defined and will connect to Leaf switches and, depending on functions, to Spine switches. Use of this structure will be the operator's choice. In practice, any routed or switched design is possible, creating an opportunity for the operator to evolve toward a transport network based on software-defined networking (SDN) and to maximize network function virtualization (NFV) possibilities. SDN is being designed into FMA from the start, but specific details are for further study.

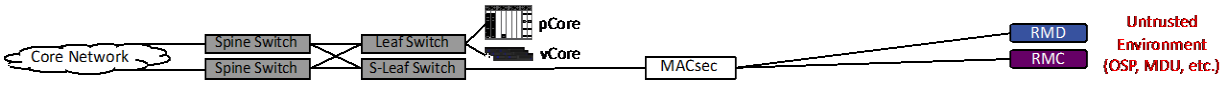

Figure 10 - Secure Leaf (S-Leaf) and Leaf

6.3.1.7 Operations Support System to S-Leaf Interface (Oss-SLeaf)

Oss-SLeaf is the interface between the operator OSS/back office and the S-Leaf. This interface is an area for further study.

6.3.1.8 S-Leaf to MAC-NE Interface (SLeaf-MacNe)

SLeaf-MacNe is the interface between S-Leaf and the MAC-NE. This interface is an area for further study.

6.3.1.9 Operations Support System to MAC Manager Interface (Oss-Mm)

Oss-Mm is the northbound interface between the MAC Manager and operator OSS/back office and represents DOCSIS OSSI and Remote OSSI (R-OSSI) plus extensions added to support FMA.

Oss-Mm carries management protocols such as SNMP, IPDR, DHCP, and TFTP. In some cases, this interface carries information between backoffice systems and the MAC Manager acting as a proxy for a MAC-NE. Refer to [FMA-OSSI] for details.

6.3.1.10 Operations Support System to MAC-NE Interface (Oss-MacNe)

O**ss**-MacNe represents a set of interfaces from operator back-office systems that pass directly to the MAC-NE. Refer to [FMA-OSSI] for details.

6.3.1.11 Policy Server to PacketCable Aggregator Interface (Ps-Pag)

Ps-Pag is the interface between a PacketCable Multimedia (PCMM) Policy Server in the operator backoffice system and a PacketCable Aggregator for policy protocols such as COPS and DQoS. Refer to [FMA-PAI] for details.

6.3.1.12 Call Management Server to PacketCable Aggregator Interface (Cms-Pag)

Cms-Pag is the interface between a PacketCable Call Management Server (CMS) in the operator backoffice system and a PacketCable Aggregator for call management signaling. Refer to [FMA-PAI] for details about this interface.

6.3.1.13 Records Keeping Server to PacketCable Aggregator Interface (Rks-Pag)

Rks-Pag is the interface between a PacketCable Records Keeping Server (RKS) in the operator backoffice system and a PacketCable Aggregator for records-keeping data and signaling. Refer to [FMA-PAI] for details about this interface.

6.3.1.14 PacketCable Aggregator to MAC-NE Interface (Pag-MacNe)

Pag-MacNe is the interface between the PacketCable Aggregator and the MAC-NE. Refer to [FMA-PAI] for details about this interface.

6.3.1.15 MAC Manager to MAC-NE Interface (Mm-MacNe)

Mm-MacNe is the southbound interface between the MAC Manager and MAC-NE Management Client and is specified as RESTCONF. Refer to [FMA-MMI] for definition of protocols and [FMA-YANG] for definition of object models.

6.3.1.16 Precision Time Protocol to Out-of-Band Core Interface (Ptp-OobCore)

Ptp-OobCore is the interface between the Precision Time Protocol (PTP) system and the OOB Core. Refer to Section 6.4.17, Timing Architecture for details.

6.3.1.17 Precision Time Protocol to Video Core Interface (Ptp-VidCore)

Ptp-VidCore is the interface between the PTP System and the Video Core. Refer to Section 6.4.17, Timing Architecture for details.

6.3.1.18 Precision Time Protocol to MAC-NE Interface (Ptp-MacNe)

Ptp-MacNe is the interface between the Precision Time Protocol system and the MAC-NE for timing protocols such as [IEEE 1588]. Refer to Section 6.4.17, Timing Architecture for details.

6.3.1.19 OOB Video Management System to Out-of-Band Core Interface (VidMgmt-OobCore)

VidMgmt-OobCore is the interface between the OOB video management system in the operator's backoffice and the OOB Core. Refer to Section 6.4, Auxiliary Core System Operation for details.

6.3.1.20 OOB RF System Management to Out-of-Band Core Interface (RfSysMgmt-OobCore)

RfSysMgmt-OobCore is the interface between the OOB RF System Management system in the operator's backoffice and the OOB Core. Refer to Section 6.4, Auxiliary Core System Operation for details.

6.3.1.21 Edge Resource Manager to Video Core Interface (ErmMgmt-VidCore)

ErmMgmt-VidCore is the interface between the Edge Resource Manager (ERM) and the Video Core. Refer to Section 6.4.9, Video Core Architecture for details.

6.3.1.22 Out-of-Band Core to MAC-NE Interface (OobCore-MacNe)

OobCore-MacNe is the interface between the Out-of-Band (OOB) core(s) and the MAC-NE. This interface carries STB control packets for video via R-DEPI and R-UEPI pseudowires. Refer to Section 6.4, Auxiliary Core System Operation for details.

6.3.1.23 Video Core to MAC-NE Interface (VidCore-MacNe)

VidCore-MacNe is the interface between the Video Core and the MAC-NE. This interface carries MPEG video packets via R-DEPI pseudowires. Refer to Section 6.4, Auxiliary Core System Operation for details.

6.3.2 MAC-NE Device Architecture

Figure 11 shows the architecture for a MAC-NE. This section is meant to be explanatory and is not meant to be all-inclusive. Specific implementations may differ.

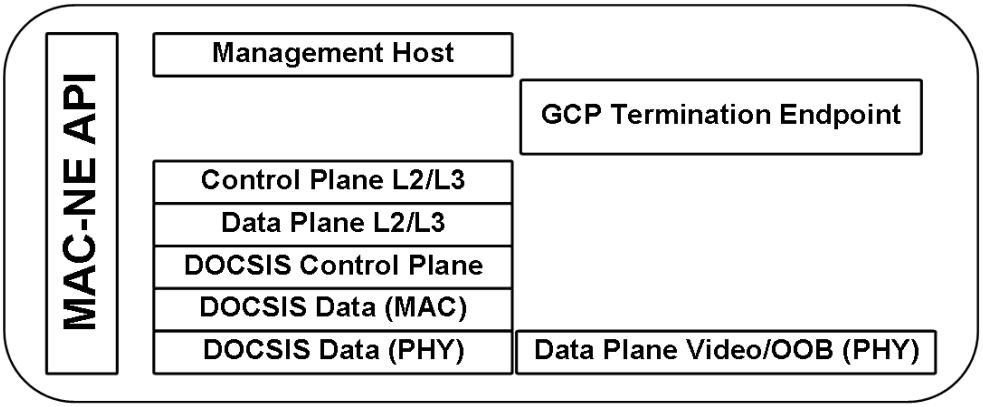

Figure 11 - MAC-NE Device Architecture

Management plane traffic from the MAC Manager is processed by the MAC-NE Management Host functional entity, which is assumed to include a logical RPD GCP Termination Endpoint for the handling of management transactions, which are common between a MAC-NE and an RPD. The MMI Southbound Interface connects the MAC Manager and MAC-NE Management Client.

L2/L3 control plane traffic is processed by the MAC-NE L2/L3 Control Plane functional entity. Processing depends upon whether the MAC-NE is operating with an L2 or an L3 data plane. It is expected that, in future phases of the FMA, the L2/L3 Control Plane functional entity will evolve to support SDN capabilities of the operator transport network.

L2/L3 data plane traffic is processed by the MAC-NE L2/L3 Data Plane functional entity. Processing depends upon whether the MAC-NE is operating with an L2 or an L3 data plane. All high-speed data, VoIP, and IP video traffic will be handled in the MAC-NE L2/L3 data plane layer, which is located above the DOCSIS MAC layer in the MAC-NE.

Control plane traffic from the PacketCable Aggregator is processed by the MAC-NE DOCSIS Control Plane functional entity. The PacketCable Aggregator Interface (PAI) connects the PacketCable Aggregator and the MAC-NE DOCSIS Control Plane entity. It is expected that, in future phases of FMA, the DOCSIS Control Plane functional entity will evolve to support NFV with a virtualized DOCSIS Controller.

DOCSIS MAC and upper layer functionality such as downstream forwarding and upstream scheduling resides in the MAC-NE DOCSIS MAC functional entity. Locating DOCSIS MAC and upper layer processing on the same platform as the DOCSIS PHY eliminates concerns of failover/redundancy, synchronization, latency/jitter, extensibility, and interoperability brought about by separation of MAC from PHY. It also means that high-speed data, VoIP, and IP video traffic can be transported over any standard L2 and/or L3 data plane on the CIN, enabling future possibilities for SDN.

DOCSIS PHY functionality resides in the MAC-NE DOCSIS PHY functional entity. There is also an RPD DOCSIS PHY entity, but the entity in the MAC-NE does not need DEPI decapsulation or UEPI encapsulation of DOCSIS MAC frames, nor is there a need in the MAC-NE to configure and manage DEPI or UEPI tunnels for DOCSIS data traffic.

Video and OOB traffic are processed by the MAC-NE Video/OOB PHY functional entity. In the downstream direction, DEPI framing is terminated and the payload is extracted, framed, modulated, and transmitted out the cable interface. In the upstream direction, the received signal on the coax is digitized, demodulated, placed into UEPI encapsulation, and transmitted to the Video/OOB Core components. Additional processing is required for OOB, particularly in the case of [SCTE 55-2]. Details can be found in [R-OOB].

Each MAC-NE maintains a unique identifier. The MAC-NE Unique Identifier serves as the principal identifier for the MAC-NE in the FMA management system. The MAC-NE Unique Identifier is embedded in the digital identity certificates assigned to the MAC-NE. The format of MAC-NE Unique Identifier is defined in [FMA-YANG].

The FMA specifications are written under the assumption that the MAC-NE Unique Identifier is permanently assigned to a MAC-NE and that all MAC-NEs within a management domain are assigned distinct MAC-NE Unique Identifiers. The method for changing of the MAC-NE Unique Identifier association to a particular MAC-NE is outside of the scope of FMA specifications.

6.3.3 MAC Manager Architecture

The primary role of the MAC Manager is to serve as an interface point for management plane functions of the MAC-NE as illustrated in Figure 12. The MAC Manger is an aggregation point for what will be a large volume of MAC-NEs in order to make the many appear as one to the back-office systems. A MAC Manager could have a 1:1 relationship to a MAC-NE or a 1:N relationship to many MAC-NE.

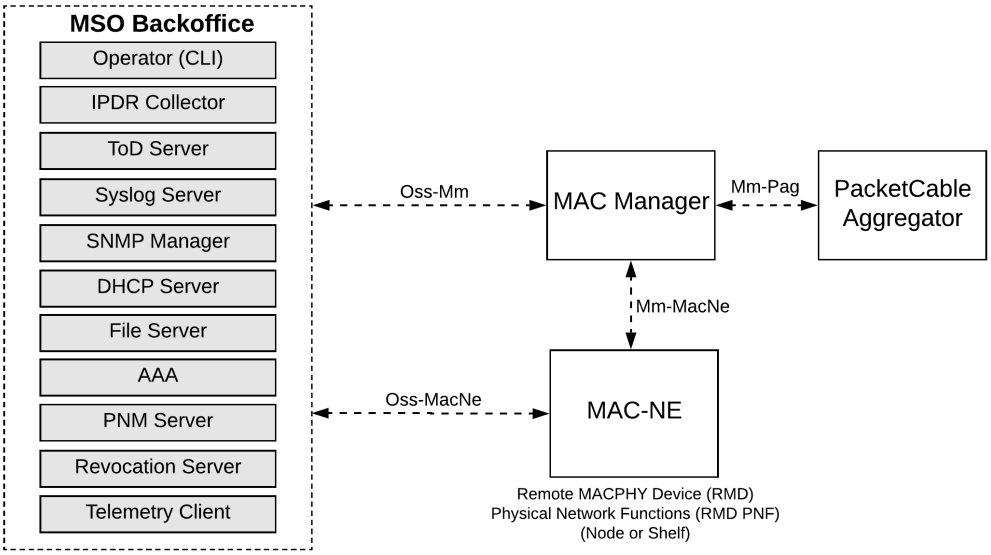

Figure 12 - FMA Management Overview

The MAC Manager Northbound Interface (FMA OSSI) implements all defined OSSI management interfaces including but not limited to SNMP, SNMP Notifications/Traps, Syslog, and IPDR while NETCONF is optionally specified for configuration management. Refer to Section 8 for additional details.

The MAC Manager Southbound Interface, referred to as MAC Management Interface (MMI), enables the MAC Manager to support [FMA-OSSI] interfaces such as: SNMP, SNMP Notifications/Traps, Syslog, and support for legacy and vendor CLI commands. See Figure 13 below and refer to [FMA-OSSI] for a detailed MAC Manager Management Architecture. The MMI is based on RESTCONF and YANG as specified in [FMA-MMI].

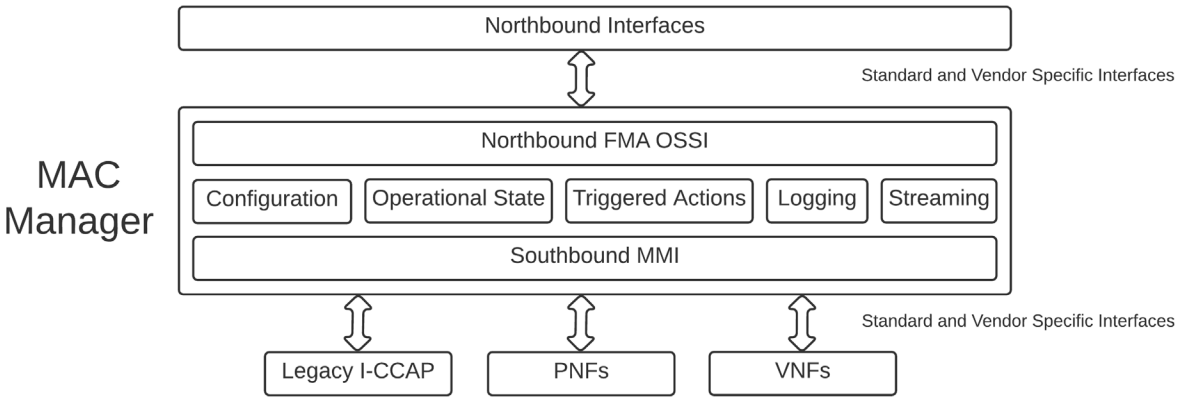

Figure 13 - MAC Manager High Level Architecture

6.3.3.1 FMA OSSI Overview

Refer to Section 5 of [FMA-OSSI] for an overview of FMA OSSI including FCAPS management models, management architecture, I-CCAP abstraction, use cases, key features, information model overview, and data mediation layer concepts.

6.3.4 PacketCable Architecture

FMA with PacketCable Aggregator (PAG) support is illustrated in Figure 14. A PAG makes the FMA appear as a single functional entity to the CMS, Policy Server, and RKS elements in an operator network. This setup is analogous to I-CCAP, CMTS, and MHAv2 architectures, in which a large number of subscribers can be dealt with through a single point of contact.

The PAG also enables each MAC-NE to have a single secure connection to the operator network rather than unique connections to every PacketCable Policy Control entity (i.e., CMS, Policy Server, RKS).

Figure 14 illustrates the first phase of FMA, where the MAC-NE connected with the PAG is a MAC-NE. A single PacketCable Application and Policy Server are shown, along with a single COPS interface from the Policy Server to the PAG. It is also possible to have multiple Application Managers per Policy Server as well as multiple Policy Servers per PAG, each having a unique COPS interface with the PAG. For simplicity, these alternatives are omitted from the figure.

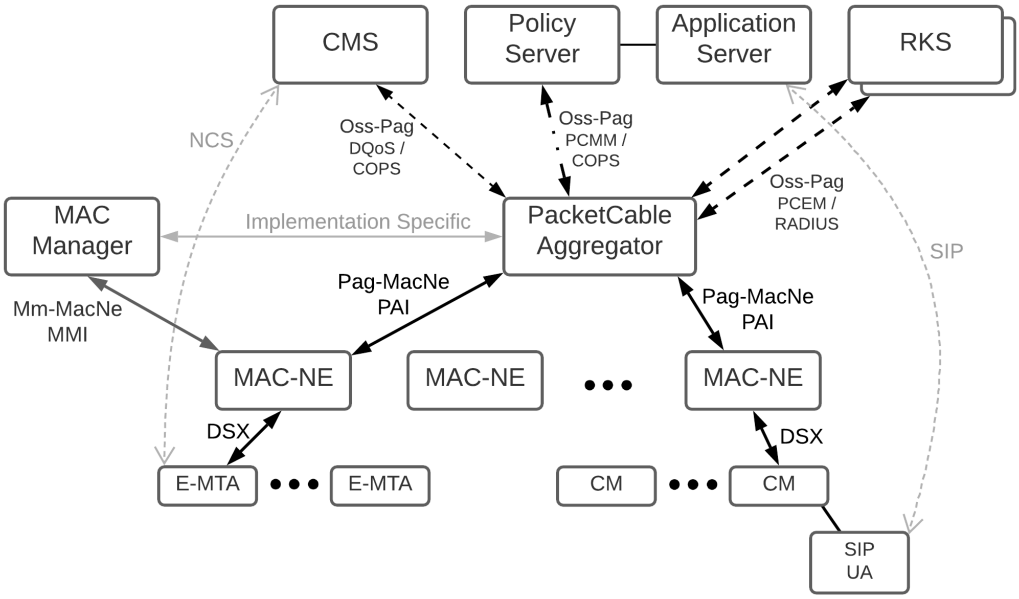

Figure 14 - FMA with PacketCable Aggregator

Interfaces north of the PAG have the following requirements: DQoS/COPS, [PKT-DQOS1.5]; PCMM/COPS, [PKT-MM]; PCEM/RADIUS, [PKT-EM1.5]. There will potentially be one DQoS interface, one or more PCMM interfaces, and a pair of PCEM interfaces from the PAG, depending on which functionality the network requires.

A PacketCable Aggregator Interface (PAI) is introduced between the PAG and the MAC-NE. The PAI is specified in [FMA-PAI]. There will be a unique PAI from the PAG to each MAC-NE in its domain of control.

The PacketCable configuration of an FMA system is performed on the MAC Manager, as specified in [FMA-OSSI].

The PacketCable configuration of MAC-NEs will be through the MAC Manager and the MMI. Associated requirements are specified in [FMA-YANG].

The PacketCable configuration of PAGs will be through the MAC Manager and is implementation-specific.

6.3.4.1 Scalability

The PacketCable Aggregator architecture can support any range of scale, from a few MAC-NE devices per PAG to many devices. This scalability is possible because the PAG makes all MAC-NEs and subscribers in its domain of control appear as a single entity to the CMS, Policy Server, and RKS. This structure is consistent with current practice and expectations for these elements.

The MAC-NE need support only a single secure PAI per PAG, rather than as many as four interfaces (one each to a CMS and a Policy Server and two to paired primary and secondary RKSs). Using only one interface can potentially reduce processing and product cost burden on the node and supports MAC-NEs of any size most effectively.

6.3.4.2 Security

Interfaces between the PAG and the CMS, the PAG and the Policy Server, and the PAG and the RKS are assumed to be within a trusted environment. As such, legacy IPsec-based security mechanisms specified for DQoS, PCMM, and PCEM are believed to be adequate. In fact, these mechanisms are often not even used because of the trusted nature of the environment in which the COPS and RADIUS interfaces are deployed. In any event, the operator may choose from legacy IPsec-based security mechanisms specified for DQoS, PCMM, and PCEM. The requirements for each of the pertinent legacy interfaces are listed in Table 2. Normative requirements exist in [PKT-DQOS1.5], [PKT-MM], and [PKT-EM1.5].

Table 2 - PacketCable Security Interfaces (PacketCable Aggregator)

Scope	PacketCable Interface	Security Method	Key Management
PacketCable Aggregator to CMS	pkt-s8	IPsec ESP	IKE w/ pre-shared keys
PacketCable Aggregator to Policy Server	pkt-mm-2	IPsec ESP	IKE w/ pre-shared keys
PacketCable Aggregator to RKS	pkt-mm-5	IPsec ESP	IKE w/ pre-shared keys

The PAI between the PAG and the MAC-NE(s) cannot be assumed to be completely within a trusted environment. FMA requires support for MACsec and 802.1x authentication, so the final link between the S-Leaf and MAC-NE can be assumed to be trusted. End-to-end security between the PAG and MAC-NE is provided via TLS.

6.3.4.3 Configuration

CMS and Policy Server elements are configured with enough information to be able to establish COPS connections with the PAG. Such information includes the PAG IP address and, if IPsec is in use, the pre-shared keys used for IKE and IPsec Security Association (SA) establishment. Similarly, the PAG is configured with enough CMS and Policy Server information for it to connect with each of those entities, such as pre-shared keys for IKE and IPsec SA establishment if IPsec is in use.

The CMS and Policy Server also need to be provided information that identifies subscribers served by a given PAG. This typically includes ranges of subscriber IP addresses (subnets) served by the PAG. This information will be made available to the CMS and Policy Server by the MAC Manager, as specified in [FMA-OSSI].

When IPsec is in use, RKS elements are configured with information needed to be able to establish IPsec SAs with the PAG. In this case, an RKS may receive IP address and pre-shared keys from the PAG. Similarly, the PAG would be configured with enough RKS information to establish IPsec SAs with its primary and secondary RKSs, that is, at least the IKE pre-shared keys.

IP addresses and ports for primary and secondary RKSs are provided in the DQoS and PCMM Gate Control messages sent by the CMS and Policy Server.

The configuration of the CMS, Policy Server, and RKS are unchanged by and out-of-scope for FMA.

The configuration of the MAC Manager/PAG and MAC-NE are in-scope for FMA.

CMS, Policy Server, and RKS connectivity information will be configured on the MAC Manager as specified in [FMA-OSSI]. The way in which the MAC Manager configures the PAG with this information is implementation specific.

PAG connectivity information is configured on the MAC-NE during onboarding by the MAC Manager as specified in [FMA-YANG].

PacketCable is configured on the MAC Manager as specified in [FMA-OSSI]. For example, PacketCable configuration includes COPS protocol timers, CMTS Gate ID range, and PacketCable Event Messaging controls. The MAC Manager configures this information on the PAG and the MAC-NEs it serves. The MAC Manager configuration of the PAG is implementation-specific. Configuration of MAC-NEs is specified in [FMA-YANG].

6.3.4.4 Connection Establishment

Subsequent to configuring the CMS, Policy Server, RKS, MAC Manager, and PAG with the appropriate information, IPsec SA (if used) and connectivity are established. Requirements for establishing IPsec SAs and connecting the CMS, Policy Server, and RKS to the PAG are as specified in [PKT-DQOS1.5], [PKT-MM], and [PKT-EM1.5], respectively, and are unchanged in FMA.

Requirements for securely connecting a MAC-NE to a PAG are specified in Section 13.16.3 as part of MAC-NE onboarding. Requirements for initializing and maintaining the PAI connection between the PAG and MAC-NE are specified in [FMA-PAI].

6.3.4.5 Operational Flow

An example operational flow for PCMM and PCEM is shown in Figure 15. It illustrates the PAG as an intermediary between a PCMM Policy Server and a MAC-NE as well as between the MAC-NE and an RKS. The PAG terminates the northbound COPS and RADIUS connections and communicates with the MAC-NE via the PAI.

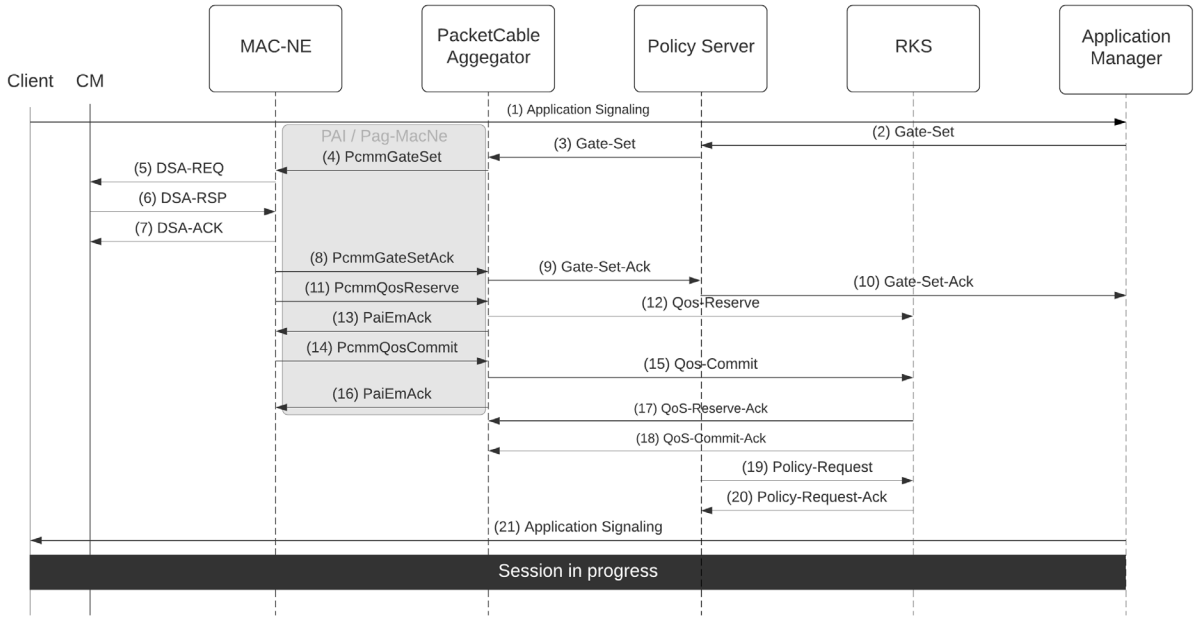

Figure 15 - Example Message Flow with PacketCable Aggregator

Gate Control from the Policy Server (i.e., Gate-Set) is relayed toward the MAC-NE by the PAG. Each MAC-NE is responsible for making its own Gate ID assignments and for returning the assigned Gate ID to the PAG, which in turn relays a unique Gate ID over the northbound COPS interface to the Policy Server in the associated Gate Control response (i.e., Gate-Set-Ack). The PAG applies an implementation-specific mapping algorithm to ensure the uniqueness of the Gate ID on the COPS interface.

PacketCable Event Messages (i.e., PcmmQosReserve, PcmmQosCommit) are transmitted by the MAC-NE on the PAI and relayed toward the RKS over the RADIUS interface by the PAG. Event Message sequence numbers are required to be monotonically increasing on the RADIUS interface. Since each MAC-NE generates its own sequence numbers for Event Messages on the PAI, the PAG is required to provide an implementation-specific mapping algorithm to ensure the proper order of sequence numbers for Event Messages it relays onto the RADIUS interface.

For PacketCable Event Messages, the PAG and RKS are required to maintain a common NTP timebase. The PAG is required to overwrite Event Message timestamps with its own view of NTP time. PacketCable provides special Event Messages to report internal time adjustments at the PAG to the RKS, and the PAG is required to generate this message when its view of time deviates by plus or minus 100ms from the current NTP time.

6.3.4.6 Subscriber Database

The subscriber database maintained by the MAC Manager is derived from the IP Services Forwarding Database. The FMA Subscriber Database keeps mappings of CM and CPE to MAC-NE, among numerous other things. It includes the MAC and IP addresses of registered CMs, the MAC and IP addresses of online CPE, and a mapping from CM and CPE to the MAC-NE on which they are resident. The Subscriber Database allows the PAG to represent all CMs and CPE in its domain of control collectively to the CMS and Policy Server. This database simplifies operational aspects for FMA in a manner analogous to CMTS, I-CCAP, and MHAv2 architectures.

The Subscriber Database is generally applicable to a number of FMA applications, so it is maintained by the MAC Manager. Requirements for the Subscriber Database are provided in Section 8.3.

6.3.4.7 Gate Control and Gate Database

The CMS, Policy Server, and PAG each maintain a Gate Database to track active DQoS/PCMM Gates within their scopes of control. The methods and protocol support for performing Gate Control and keeping the Gate Databases in sync between these entities are as specified in [PKT-DQOS1.5] and [PKT-MM].

The PAG and MAC-NE each maintain a Gate Database to track active DQoS/PCMM Gates within their scopes of control. The methods and protocol support for performing Gate Control and keeping the Gate Database in sync between these entities are specified here and in [FMA-PAI].

In both DQoS and PCMM, the Gate ID is assigned by the Policy Enforcement Point, which in legacy environments is a CMTS or I-CCAP that supports large numbers of subscribers. This is also true in MHAv2, where the CCAP Core represents many subscribers. In these cases, the CMTS, I-CCAP, or CCAP Core can be configured with a range of Gate IDs that it is allowed to assign, which minimizes the need for Gate ID collision resolution mechanisms in CMS and Policy Servers that communicate with more than one CMTS, I-CCAP, or CCAP Core. The CmtsGateIdValue in [CCAP-OSSIv3.1] is the lower 13 bits of the GateID, allowing the CMTS or CCAP to assign the remaining bits. In other words, it assigns the CMTS or CCAP a range of Gate IDs. FMA supports the same configurability as provided in [CCAP-OSSIv3.1], with the MAC Manager assigned a unique CmtsGateIdValue for the "CMTS" being represented by the PAG and its subtended MAC-NE population. This same CmtsGateIdValue will be configured on the PAG and on each MAC-NE in its scope of control by the MAC Manager. This CmtsGateIdValue is used by each MAC-NE when generating its Gate IDs.

In FMA, each MAC-NE is responsible for making its own Gate ID assignments. As a result, Gate ID conflicts can occur at the PAG. The PAG is responsible for preventing any such Gate ID assignment conflicts from traversing the COPS interfaces toward the Policy Server or CMS. The mechanism used by the PAG to do so is implementation specific.

The PAG is responsible for synchronizing its Gate Database with the Gate Databases in each MAC-NE. The majority of the synchronization is done by way of processing Gate Control messages: inserting MAC-NE-generated Gate IDs into and removing Gate IDs from the database based on Gate Control (e.g., Gate-Set, Gate-Delete) and associated transactional responses. The PAG, MAC-NE and PAI also support bulk Gate Database synchronization to recover from loss of PAI connectivity. This bulk synchronization is analogous to the bulk synchronization support provided in the DQoS and PCMM protocols on the COPS interfaces.

Although PAI protocol support will be provided for bulk synchronization, PAG processing related to recovery from out-of-synch is implementation specific. For example:

- If a MAC-NE gate entry is nonexistent in the PAG Gate Database, then the PAG adds it.
- If a PAG gate entry is nonexistent in the MAC-NE Gate Database, then the PAG deletes that entry from its own database, reports gate deletion to the Policy Decision Point and reports quality of service (QoS) resources released to the RKS.

6.3.4.8 Redundancy/Failover

Redundancy within the FMA PacketCable architecture is illustrated in the example shown in Figure 16.

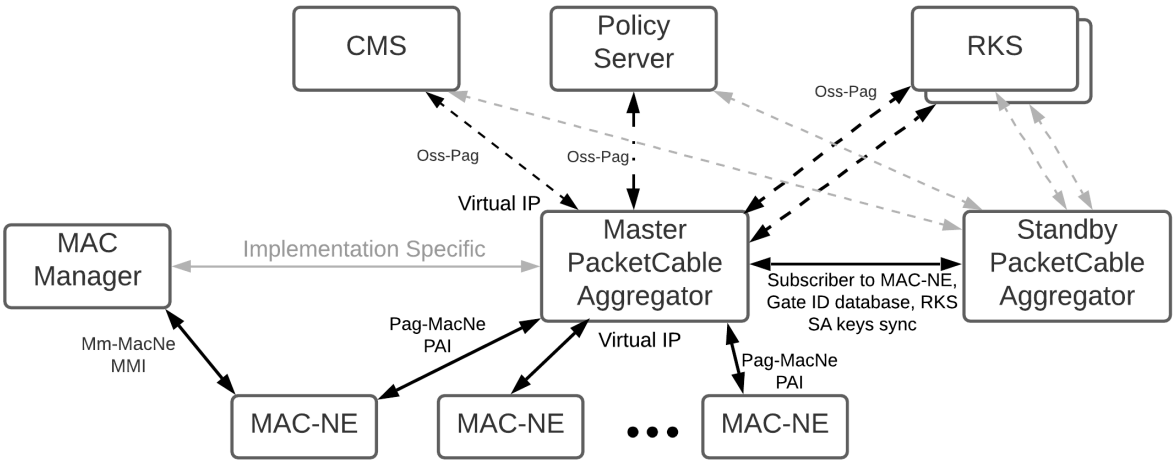

Figure 16 - PacketCable Aggregator Redundancy

Tools are available in DQoS, PCMM, and PCEM to support redundancy and failover implementations. For example, COPS Keep-Alive messages in DQoS and PCMM allow protocol endpoints to detect loss of connectivity more quickly than does the underlying TCP. PCMM sync procedures are provided to resync protocol endpoints after restoration of connections. RADIUS three-way handshakes verify communication with the primary RKS, and dual RKS connectivity supports near-real-time failover between the Primary and Secondary RKS.

CMS, Policy Server, and RKS redundancy/failover implementations can leverage the available protocol tools, but they are implementation-specific. PAG COPS and RADIUS interface redundancy/failover methods can leverage the available tools, but they also are implementation specific, just as are those for I-CCAP, CMTS, and MHAv2 CCAP Core.

Redundancy/failover support mechanisms for the PAI are covered in [FMA-PAI]. These include protocol support for loss of connectivity detection and failover/recovery.

6.3.4.9 Requirements for FMA with PacketCable Aggregator

The requirements in this section are applicable when the FMA supports Policy Control with a PacketCable Aggregator.

6.3.4.9.1 MAC Manager Requirements

The MAC Manager MUST accept PacketCable Event Message interface configuration for the PAG as specified in [FMA-OSSI]. The applicable parameters for PacketCable Event Message interface configuration are listed below:

- RADIUS Event Message retry timer. PacketCable specifies a range of 10ms to 10sec.
- RADIUS Event Message retry count. PacketCable specifies a range of 0-9.
- PAG Element_ID for use in RADIUS Event Message Headers. PacketCable specifies a space-padded, right-justified 5 digit ASCII character string unique within the "PacketCable Domain".

The MAC Manager MUST apply PacketCable Event Message interface configuration to the PAG, but it can do so in a way that is implementation-specific.

The MAC Manager MUST accept MAC-NE PAG onboarding, PAI initialization, and PAI connection management configuration for any MAC-NE subtended from the PAG as specified in [FMA-OSSI]. The applicable parameters for PAG onboarding, PAI initialization, and PAI connection management are listed below.

- PAG URL (see Section 13.16.3).

- PAG TLS session establishment retry count (see Section 13.16.3).
- PAG connection establishment timer (see Section 13.16.3).
- PAG connection establishment retry count (see Section 13.16.3).
- PAI initialization timer (see [FMA-PAI]).
- PAI initialization retry count (see [FMA-PAI]).
- PAI KeepAliveMaxInterval (see [FMA-PAI]).

The MAC Manager MUST apply MAC-NE PAG onboarding, PAI initialization, and PAI connection management configuration to every MAC-NE in its scope as specified in [FMA-YANG].

The MAC Manager MUST accept PacketCable configuration for the PAG and subtended MAC-NE devices as specified in [FMA-OSSI].

The MAC Manager MUST apply PacketCable configuration to the PAG, but it can do so in a way that is implementation-specific. All PacketCable configuration parameters will be applied to the PAG. However, the PAG is not required to make use of all parameters during operation.

The MAC Manager MUST apply PacketCable configuration to every MAC-NE in its scope as specified in [FMA-MMI]. All PacketCable configuration parameters will be applied to the MAC-NE. However, the MAC-NE is not required to make use of all parameters during operation.

6.3.4.9.2 PacketCable Aggregator Requirements

The PAG MUST support IPsec ESP and IKE with pre-shared keys for the COPS interface to a PacketCable Call Management Servers.

The PAG MUST support IPsec ESP and IKE with pre-shared keys for the COPS interface to a PacketCable Multimedia Policy Server.

The PAG MUST support IPsec ESP and IKE with pre-shared keys for the RADIUS interface to a PacketCable Record Keeping Server.

The PAG MUST accept PacketCable Event Message interface configuration from the MAC Manager.

The application of PacketCable Event Message interface configuration to the PAG by the MAC Manager is implementation-specific.

The PAG MUST accept PacketCable configuration from the MAC Manager. All PacketCable configuration parameters will be applied to the PAG. However, the PAG is not required to make use of all parameters during operation.

The application of PacketCable configuration to the PAG by the MAC Manager is implementation-specific. The PAG is required to connect to the MAC-NE as specified in Section 13.16.3.

The PAG MUST obtain subscriber-to-MAC-NE mappings from the Subscriber Database. The way in which the PAG obtains subscriber-to-MAC-NE mappings from and stays synchronized with the Subscriber Database is implementation-specific.

The PAG MUST support all published PacketCable Multimedia specification versions.

The PAG MUST support all published PacketCable Dynamic Quality of Service specification versions.

The PAG MUST support PacketCable Dynamic Quality of Service CMTS requirements as specified in [PKT-DQOS1.5] and PacketCable Multimedia CMTS requirements as specified in [PKT-MM], except as noted below.

- The PAG MUST appear to the CMS and Policy Server to be the Policy Enforcement Point.
- The PAG MUST rely on the MAC-NE to be the actual Policy Enforcement Point.
- The PAG MUST serve as intermediary between the CMS/Policy Server and the MAC-NE by mapping COPS-based transaction messages to analogous PAI-based transaction messages specified in [FMA-PAI].

- The PAG MUST serve as intermediary between the MAC-NE and the CMS/Policy Server by mapping PAI-based transaction messages specified in [FMA-PAI] to analogous COPS-based transaction messages.

The PAG MUST log a COPS initialization success notice event (70090101) upon a successful first time COPS connection from a new CMS or Policy Server.

The PAG MUST log a COPS initialization failure critical event (70090100) upon failure to receive a COPS Client-Accept from a new CMS or Policy Server during first time COPS connection.

The PAG MUST log a COPS Keep-Alive-Ack missed warning event (70090200) upon failure to receive Keep-Alive-Ack from a CMS or Policy Server.

The PAG MUST log a COPS connection failure error event (70090201) upon loss of COPS connection with a CMS or Policy Server.

The PAG MUST log a COPS connection failure condition cleared notice event (70090202) upon successful reestablishment of COPS connection with a CMS or Policy Server.

The PAG MUST ensure that MAC-NE-assigned Gate IDs are unique on any one of its given northbound COPS interfaces. The mechanism that the PAG uses to ensure uniqueness of Gate IDs on a given COPS interface is implementation-specific.

When configured with a CmtsGateIdValue, the PAG MUST ensure that Gate IDs that traverse its northbound COPS interfaces preserve that CmtsGateIdValue in the lower 13 bits of the Gate ID.

The PAG MUST maintain a Gate Database containing Gate IDs for all active gates within the PAG's scope of control.

The PAG MUST synchronize its Gate Database to that of a given MAC-NE after loss of connectivity with that MAC-NE by using PAI tools specified in [FMA-PAI].

The PAG MAY originate COPS Gate-Delete messages as a result of Gate Database synchronization processing (e.g., when a MAC-NE is missing a Gate ID that is present in the PAG's Gate Database).

PCMM SharedResourceIDs are generated by a MAC-NE to identify that a particular Multicast Gate is associated with a specific Multicast Group Service Flow. A SharedResourceID essentially identifies a Multicast Group Service Flow. In FMA, Multicast Gates and their mapping to Multicast Group Service Flows are specific to a given MAC-NE and unique from the perspective of a PCMM Policy Server. However, it is possible that different MAC-NEs can generate the same SharedResourceID for different Multicast Gates and Multicast Group Service Flows. The PAG MUST ensure that SharedResourceIDs in Gate Control messages relayed from the PAI onto its northbound COPS interface are unique when coming from different MAC-NE devices. The method for doing so is implementation-specific.

The PAG MUST support PacketCable Event Messages CMTS requirements as specified in [PKT-EM1.5] and [PKT-MM], except as noted below.

- The PAG MUST serve as intermediary between the MAC-NE and the Primary RKS by mapping PAI-based Event Messages specified in [FMA-PAI] to analogous RADIUS-based Event Messages.
- The PAG MUST populate the Event Message Header Element_ID of all northbound RADIUS Event Messages with its configured PAG Element ID.

The PAG MUST immediately acknowledge Event Messages received on the PAI from a MAC-NE.

The PAG MUST ensure that Event Messages relayed from the PAI onto its northbound RADIUS interface contain monotonically increasing sequence numbers by inserting PAG-generated, monotonically increasing integer sequence number values. The PAG-generated sequence number replaces that provided in the MAC-NE generated PAI Event Message.

The PAG MUST ensure a common timebase for all Event Messages sent on the northbound RADIUS interface by providing its local NTP-based timestamp in the Event Message Header. The PAG timestamp replaces that provided in the MAC-NE generated PAI Event Message.

The PAG MUST acquire its time of day (e.g., via NTP) at least once every hour.

The PAG MUST transmit a Time-Change Event Message on its northbound RADIUS interface when the NTP time deviates by +/- 100ms from the PAG's local time of day.

The PAG MAY originate Event Messages related to QoS resource release as a result of Gate Database synchronization processing (e.g., when an MAC-NE is missing a Gate ID that is present in the PAG's Gate Database).

The PAG MUST support real time Event Message generation for all PCMM-initiated Gates created with a Gate Control message that specified Event Message generation is required.

The PAG MUST support real time Event Message generation for all DQoS-initiated Gates created with a Gate Control message that specified real time Event Message generation is required.

The PAG MUST support batch mode Event Message generation for all DQoS-initiated Gates created with a Gate Control message that specified batch mode Event Message generation is required. When DQoS Gate Control messages specify batch mode event generation for Gates, the PAG MUST buffer Event Messages received for those Gates and package multiple Event Messages into a single RADIUS message. RADIUS messages have a maximum size of 4096 bytes.

The PAG MUST log a RADIUS event message send failure Primary RKS warning event (70090203) upon failure to receive a RADIUS ACK for the event after exhausting the configured number of RADIUS event message retries on the Primary RKS.

The PAG MUST log a RADIUS event message send failure Secondary RKS error event (70090204) upon failure to receive a RADIUS ACK for the event after exhausting the configured number of RADIUS event message retries on the Secondary RKS after switching to the Secondary RKS from the Primary RKS.

The PAG SHOULD store a copy of all unacknowledged RADIUS event messages in a local event message error file.

6.3.4.9.3 MAC-NE Requirements

The MAC-NE MUST accept MAC-NE PAG onboarding, PAI initialization, and PAI connection management configuration from the MAC Manager as specified in [FMA-YANG].

The MAC-NE MUST accept PacketCable configuration from the MAC Manager as specified in [FMA-MMI]. All PacketCable configuration parameters will be applied to the MAC-NE. However, the MAC-NE is not required to make use of all parameters during operation.

The MAC-NE is required to connect to the PAG as specified in Section 13.16.3

The MAC-NE MUST support all published PacketCable Multimedia specification versions.

The MAC-NE MUST support all published PacketCable Dynamic Quality of Service specification versions.

The MAC-NE MUST support PacketCable Dynamic Quality of Service CMTS requirements as specified in [PKT-DQOS1.5] and PacketCable Multimedia CMTS requirements as specified in [PKT-MM], except as noted below.

- When working with a PAG, the MAC-NE MUST support PAI-based transactions, as specified in [FMA-PAI], rather than the analogous COPS-based transactions.

The MAC-NE MUST respond to synchronization requests from the PAG by providing Gate IDs for all of its active gates.

The MAC-NE MUST ensure that MAC-NE-assigned Gate IDs are unique on its PAI. The mechanism which the MAC-NE uses to ensure uniqueness of Gate IDs on a PAI is implementation-specific.

When configured with a CmtsGateIdValue, the MAC-NE MUST ensure that Gate IDs that traverse its northbound PAI preserve that CmtsGateIdValue in the lower 13 bits of the Gate ID.

The MAC-NE MUST support PacketCable Event Messages CMTS requirements as specified in [PKT-EM1.5] and [PKT-MM], except as noted below.

- The MAC-NE MUST support PAI-based transactions, as specified in [FMA-PAI], rather than the analogous RADIUS-based transactions.

- The MAC-NE MUST NOT implement the Event Message retry and associated RKS failover processing requirements of PacketCable on its PAI connection.
- The MAC-NE MUST NOT generate TimeChange Event Messages.
- The MAC-NE MAY buffer unsent Event Messages on loss of PAI connectivity for sending upon restoral of PAI connectivity.
- The MAC-NE MUST support real-time Event Message generation for all PCMM-initiated Gates created with a Gate Control message that specified real-time Event Message generation is required.
- The MAC-NE MUST support real-time Event Message generation for all DQoS-initiated Gates created with a Gate Control message that specified real-time Event Message generation is required.
- The MAC-NE MUST support real-time Event Message generation for all DQoS-initiated Gates created with a Gate Control message that specified batch mode Event Message generation is required. The MAC-NE MUST NOT implement batch mode Event Message generation for any Gate. Batching of Event Messages is controlled entirely by the PAG.

6.3.5 Lawful Intercept Architecture

The FMA Lawful Intercept (LI) architecture is illustrated in Figure 19, which shows tap provisioning aspects, and Figure 22, which details surveillance event and mirrored intercept content delivery. The LI architecture and these specific topics are described extensively in the sections that follow.

There are three types of Lawful Intercept that are applicable to the FMA Architecture. They include the interception of PacketCable 1.x voice calls (CALEA), the interception of PacketCable 2.0 voice and multimedia calls (LAES) and the interception of subscriber data traffic (also falls under LAES). Traditionally, the CMTS/CCAP provided the capability to perform all three types of packet-based intercepts. In FMA, Lawful Intercept is provided by the FMA System. The FMA System is defined as the disaggregated parts of the traditional integrated CMTS/CCAP that have been implemented in components of the FMA Architecture:

- The S-Leaf
- The MAC Manager
- The PacketCable Aggregator
- The MAC-NE

MSOs have traditionally used both the CMTS/CCAP and network-based devices to implement their Lawful Intercept services. The FMA Architecture continues that support.

6.3.5.1 FMA Lawful Intercept Functional Requirements

The FMA System SHOULD support Lawful Intercept of subscriber voice and data traffic.

The MAC-NE SHOULD support Lawful Intercept of PacketCable 1.x voice traffic based on CMTS requirements in the PacketCable 1.5 Electronic Surveillance Specification [PKT-ES1.5].

The FMA System SHOULD support Lawful Intercept of PacketCable 2.0 voice traffic. The MAC-NE MAY support Lawful Intercept of PacketCable 2.0 voice traffic. If the MAC-NE supports the Lawful Intercept of PacketCable 2.0 voice traffic, the MAC-NE MUST support the CMTS requirements in the PacketCable 2.0 Electronic Surveillance Intra-Network Specification [PKT-ES2.0] and the PacketCable 2.0 Electronic Surveillance Delivery Function to Collection Function Interface Specification [PKT-DCI]. If the MAC-NE supports the Lawful Intercept of PacketCable 2.0 voice traffic, the MAC-NE MUST support the configuration of the lawful intercept tap using the management objects in [FMA-OSSI].

The FMA System SHOULD support Lawful Intercept of subscriber data traffic. The MAC-NE MAY support Lawful Intercept of subscriber data traffic. If the MAC-NE supports Lawful Intercept of subscriber data traffic, the MAC-NE MUST support the CMTS requirements in the Cable Broadband and Wi-Fi Intercept Specification [CBI]. If the MAC-NE supports the Lawful Intercept of subscriber data traffic, the MAC-NE MUST support the configuration of the lawful intercept tap using the management objects in [FMA-OSSI].

Note that when the MAC-NE is the point of intercept for subscriber data traffic and the service rates for that traffic are in the gigabit or greater range, the MAC-NE's ability to process these streams could be limited due to resource constraints or egress bandwidth constraints. It is recommended that operators avoid these conditions by performing Lawful Intercept of data traffic to and from these subscribers in an alternate location (e.g., the S-Leaf) to maintain non-repudiation requirements in [CBI] and related specifications.

6.3.5.2 FMA Lawful Intercept Recommendations

This section details several recommendations for MSOs regarding the different Lawful Intercept types.

6.3.5.2.1 PacketCable 1.x Voice Call Interception Recommendations

For PacketCable 1.5 voice calls, the recommendation is to continue common practice of performing Lawful Intercept on the CMTS/CCAP by utilizing the MAC-NE as the point of intercept. The reasoning for this is based on the fact that for these voice calls the need to perform call content intercept is signaled in the DQoS messaging that is sent from the Call Management Server to the PAG and then to the MAC-NE in the FMA System. The bit rates and associated stream processing requirements for these flows are suitable for the MAC-NE to take on when signaled to do so. This recommendation is shown in the following figure.

It is noted that many MSOs have instead implemented Lawful Intercept for PacketCable 1.x voice calls in alternate locations like media gateways and have not enabled it on the traditional CMTS/CCAP. For these operators, FMA System support of PacketCable 1.x voice call Lawful Intercept would not be necessary. See Figure 17 for a depiction of the recommended FMA lawful intercept architecture.

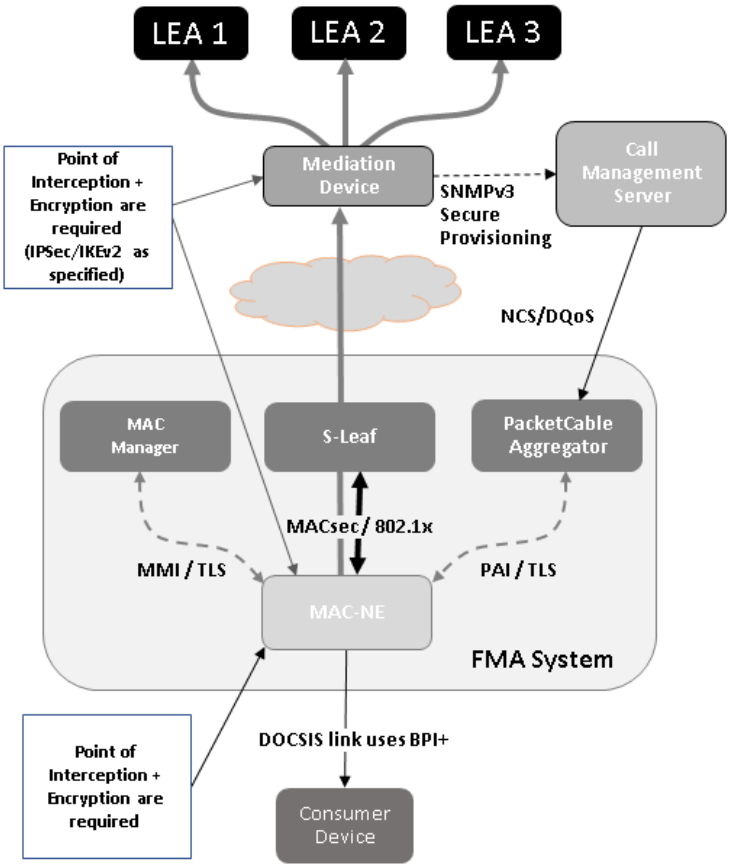

Figure 17 - Recommended FMA Lawful Intercept

6.3.5.2.2 PacketCable 2.0 Voice Call Interception Recommendations

For PacketCable 2.0 voice and multimedia calls, the recommendation is to perform intercepts in the MSO network rather than on the MAC-NE, if possible. IMS entities are used to perform the signaling and to control bearer path flows in the PacketCable 2.0 network. These IMS entities establish the IP endpoints for PacketCable 2.0 voice and multimedia calls. Intercepts are traditionally established using a secure SNMPv3 provisioning infrastructure involving the Mediation Device, the IMS entities, and the chosen points of bearer traffic intercept. When these intercepts would normally be conducted on the CMTS/CCAP, the FMA System will provide similar functionality utilizing the MAC Manager and the MAC-NE, and the point of intercept is then the MAC-NE. As with PacketCable 1.x voice calls, these calls tend to be lower in bandwidth and could be intercepted at the MAC-NE should the operator wish to utilize the MAC-NE as the point of intercept.

6.3.5.2.3 Subscriber High Speed Data Interception Recommendations

For subscriber data traffic, the recommendation is to perform intercepts in the MSO network. If the operator chooses to implement the interception of subscriber data traffic in the FMA System, the recommendation is to utilize the S-Leaf as the point of intercept. If the MAC-NE is the point of intercept, for subscriber traffic and the service rates for that traffic in the gigabit or greater range, the MAC-NE's ability to process these streams could be limited due to resource constraints or egress bandwidth constraints. It is recommended that operators avoid these conditions by performing Lawful Intercept of these subscribers in an alternate location (e.g. the S-Leaf) to maintain non-repudiation requirements in [CBI] and related specifications. The architecture for this recommendation is shown in Figure 18.

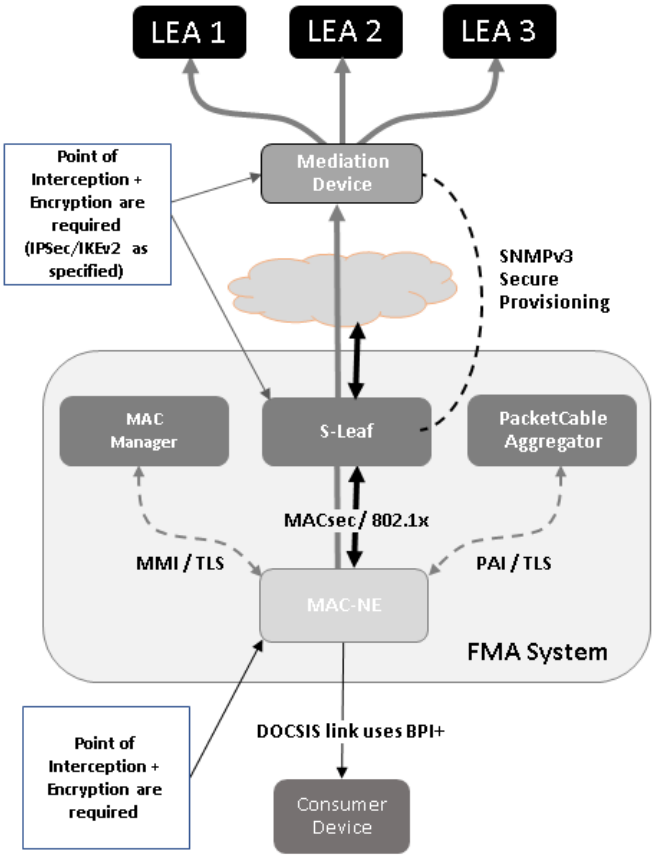

Figure 18 - Recommendations for FMA LAES and HSD Intercept

6.3.5.3 Intercept/Tap Provisioning

FMA Lawful Intercept tap provisioning is illustrated in Figure 19. There are two methods for provisioning taps in FMA. One is based on SNMPv3 with TAP-MIBs. The other is per voice call provisioning via DQoS.

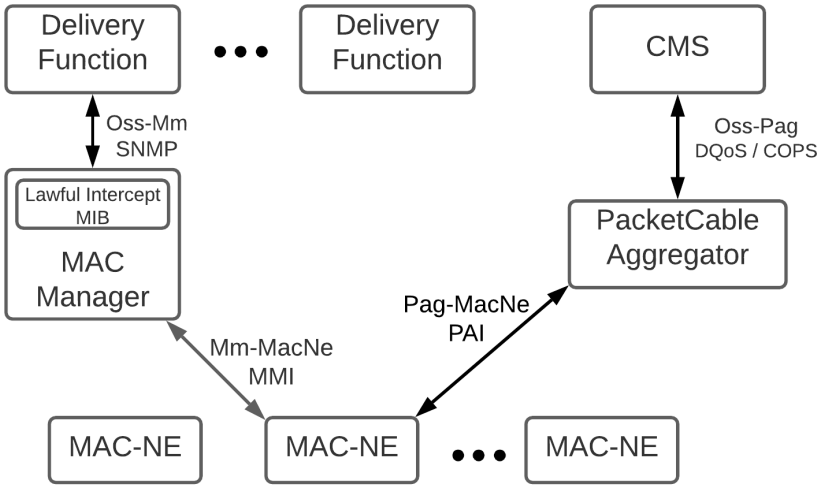

Figure 19 - FMA Lawful Intercept Tap Provisioning

6.3.5.3.1 SNMPv3-based Intercept/Tap Provisioning

SNMPv3 is used to provision intercepts from the Delivery Function/Mediation Device when the MAC-NE is the point of interception. The standards-based approach with SNMPv3 is specified in [PKT-ES2.0]. SNMPv3 is the only SNMP method supported for FMA tap provisioning.

The SNMPv3-based provisioning method includes participation from the MAC Manager. MAC Manager support for these MIBs is specified in [FMA-OSSI]. How the MAC Manager configures the MAC-NE with taps provisioned by way of these MIBs is specified in [FMA-YANG].

The MAC Manager configures the MAC-NE with tap filters to select intercept traffic to mirror, transport methods used to deliver mirrored traffic, and Delivery Function IP address and ports for delivering mirrored content. This will cause the MAC-NE to mirror intercept content toward the Delivery Function.

6.3.5.3.2 DQoS-based Intercept/Tap Provisioning

Provisioning for PacketCable 1.x voice calls (CALEA) is defined in [PKT-DQOS1.5] and [PKT-ES1.5] requirements. This method sets up voice call related event message and call content intercepts on a per call basis using PacketCable 1.5 Dynamic Quality of Service signaling from a PacketCable Call Management Server (CMS) to the PAG in the FMA System. The CMS is configured to intercept and report either call related event messages or call content, or both, to the Delivery Function. How this configuration is performed is outside the scope of this specification.

The DQoS-based provisioning method includes participation from the PAG, which processes DQoS Gate-Set messages from the CMS that have electronic surveillance parameters included in them. If a Gate-Set contains the PacketCable-Electronic-Surveillance-Parameters object and DUP-EVENT is set in the Flags field of that object, the PAG will replace the DF-IP-Address-for-CDC and DF-Port-for-CDC with its own PAI IP address and port for the respective MAC-NE before forwarding the Gate-Set on the PAI. This will cause the MAC-NE to duplicate PacketCable Electronic Surveillance (PCES) event messages for the session to the PAG over its PAI. The PAG can then forward the PCES events over RADIUS to the Delivery Function on behalf of the MAC-NE.

If a Gate-Set contains the PacketCable-Electronic-Surveillance-Parameters object and DUP-CONTENT is set in the Flags field of that object, the object will also contain the DF-IP-Address-for-CCC and DF-Port-for-CCC. This information will be forwarded by the PAG in the Gate-Set sent to the MAC-NE over the PAI, causing the MAC-NE to mirror traffic for the voice call session toward the Delivery Function.

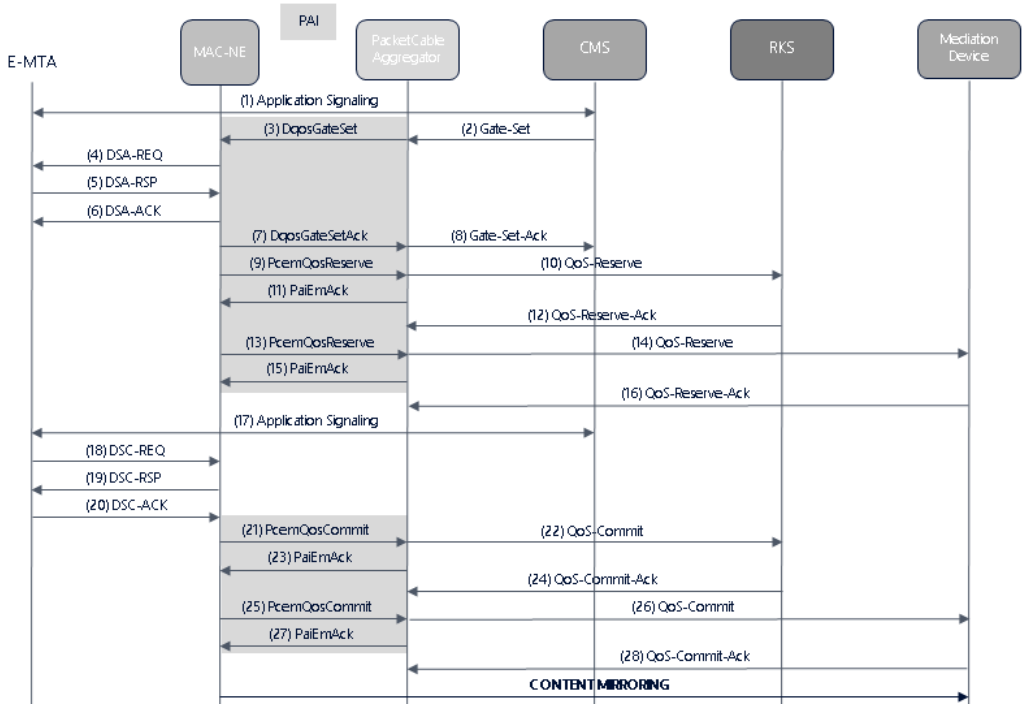

Figure 20 - DQoS-based Voice Call Intercept Set-up

A message flow and mirrored content delivery for voice calls intercepted by way of DQoS-based tap provisioning is illustrated in Figure 20. This figure shows an incoming call to an intercept subject and the associated event message reporting and call content mirroring. Details of how the CMS is provisioned are out-of-scope for FMA and are omitted. Details of the end-to-end NCS-based call signaling are abbreviated as "Application Signaling" and are also out-of-scope.

An example message flow for intercepted voice call release is shown in Figure 21. There are a number of ways the voice call and intercept can be torn down. This example shows the typical view of normal call release leading to E-MTA initiated DSD processing.

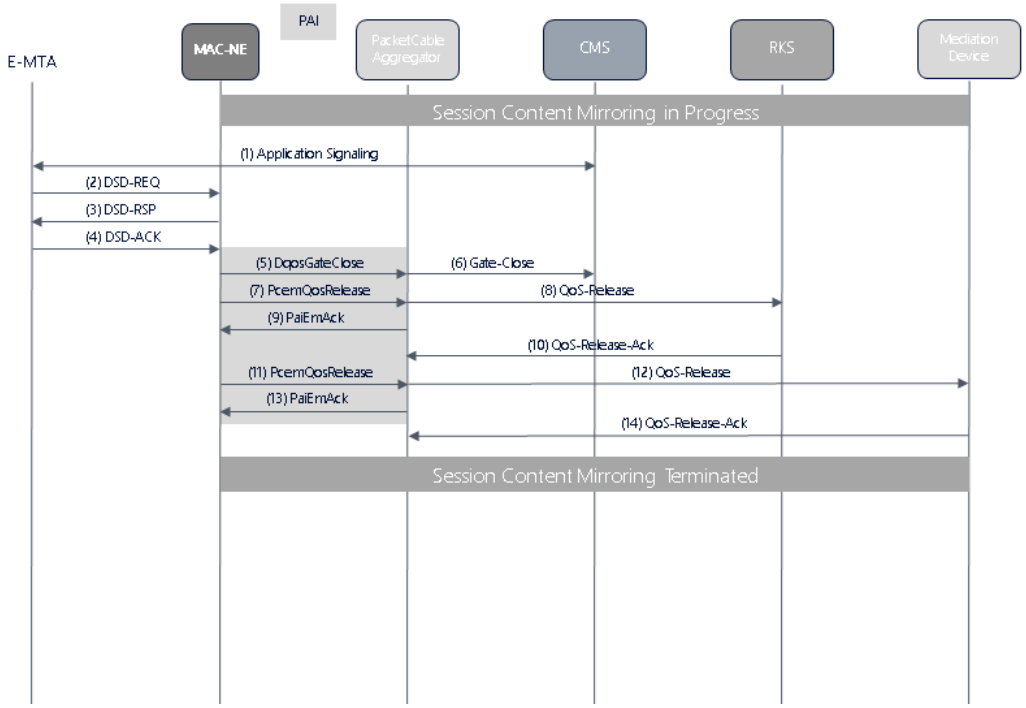

Figure 21 - DQoS-based Voice Call Intercept Tear-down

6.3.5.4 Content and Related Events Delivery When MAC-NE is Point of Interception

FMA Lawful Intercept events and mirrored content delivery is illustrated in Figure 22. Event message delivery is unique to DQoS initiated intercepts.

Mirrored intercept content is delivered by the MAC-NE directly to the Delivery Function. Only UDP is used for DQoS initiated intercepts. In some cases, the MAC-NE will insert a DSCP or ToS setting into the encapsulating IP header.

For DQoS-based intercepts, PCES event messages can be duplicated and delivered by the MAC-NE to the PAG over the PAI as specified in [PKT-DQOS1.5]. This depends on the electronic surveillance parameters present in the Gate-Set from the CMS.

The PAG receives duplicated PCES event messages for the intercepted session from the MAC-NE over its PAI and forwards these PCES event messages over RADIUS to the Delivery Function.

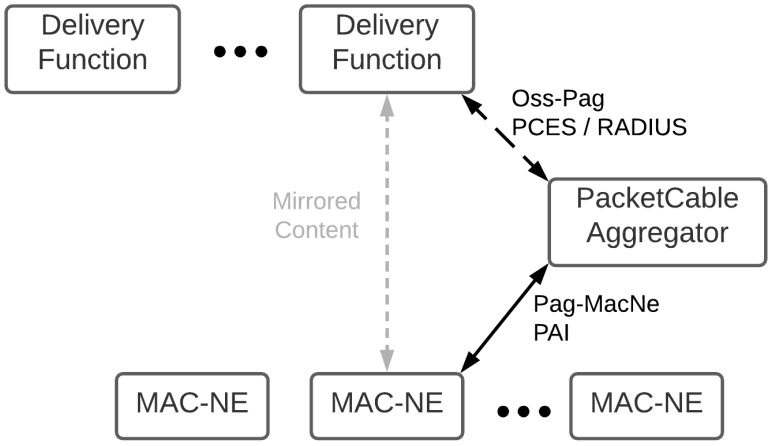

Figure 22 - FMA Lawful Intercept Event and Content Delivery

6.3.5.4.1 PacketCable Aggregator Requirements

The PAG will support DQoS-based intercept provisioning from the CMS, as specified in [PKT-DQOS1.5] for a CMTS.

If the PAG receives a DQoS Gate-Set that contains a PacketCable-Electronic-Surveillance-Parameters object and DUP-EVENT is set in the Flags field of that object, the PAG MUST replace the DF-IP-Address-for-CDC and DF-Port-for-CDC with its own PAI IP address and port for the respective MAC-NE before forwarding the Gate-Set as a DqosGateSet on the PAI to the MAC-NE. The PAG retains the DF-IP-Address-for-CDC and DF-Port-for-CDC for this Gate for later use in forwarding intercept-related PCES event messages to the Delivery Function.

The PAG will support forwarding PCES event messages to the Delivery Function over RADIUS. See [PKT-EM1.5], Appendix A, "PCES Support."

When the PAG receives a PcemQosReserve message on the PAI with the event_msg_header event_msg_type field set to 1 (PCES Event Message), the PAG MUST forward the message in PCEM/PCES Qos-Reserve format over RADIUS to the configured Delivery Function for this Gate.

When the PAG receives a PcemQosCommit message on the PAI with the event_msg_header event_msg_type field set to 1 (PCES Event Message), the PAG MUST forward the message in PCEM/PCES Qos-Commit format over RADIUS to the configured Delivery Function for this Gate.

When the PAG receives a PcemQosRelease message on the PAI with the event_msg_header event_msg_type field set to 1 (PCES Event Message), the PAG MUST forward the message in PCEM/PCES Qos_release format over RADIUS to the configured Delivery Function for this Gate.

6.3.5.4.2 MAC-NE Requirements

The requirements in this list apply when the MAC-NE is configured as the point of intercept for Lawful Intercept for PacketCable 1.x voice call content.

- The MAC-NE MUST support DQoS-based intercept provisioning from the CMS, as specified in [PKT-DQOS1.5] for a CMTS, except that the provisioning will arrive in DqosGateSet messages from the PAG over the PAI.
- The MAC-NE MUST support duplicating and forwarding PCES event messages to the Delivery Function by way of the PAG and PAI. See [PKT-EM1.5], Appendix A, "PCES Support."

- If the MAC-NE receives a PAI DqosGateSet that contains a DqosElecSurvParms object and dup_events is set in that object, the MAC-NE MUST duplicate PcemQosReserve, PcemQosCommit, and PcemQosRelease messages associated with that Gate toward the PAG over the PAI.
- When the MAC-NE duplicates PcemQosReserve messages and transmits them over the PAI to the PAG, it MUST set the event_msg_header event_object field to 1 (PCES Event Message).
- When the MAC-NE duplicates PcemQosCommit messages and transmits them over the PAI to the PAG, it MUST set the event_msg_header event_object field to 1 (PCES Event Message).
- When the MAC-NE duplicates PcemQosRelease messages and transmits them over the PAI to the PAG, it MUST set the event_msg_header event_object field to 1 (PCES Event Message).
- If the MAC-NE receives a PAI DqosGateSet that contains a DqosElecSurvParms object and dup_content is set in that object, the MAC-NE MUST duplicate upstream and downstream RTP streams associated with that Gate, encapsulate them as specified in [PKT-DQOS1.5], and transmit them to the DF_addr_for_CCC and DF_port_for_CCC.
- The MAC-NE MUST continue to duplicate PCES Event Messages and RTP streams associated with the Gate until the Gate is deleted.
- The MAC-NE will transmit tap-related status and event notifications to the MAC Manager (e.g., tap installation/start, tap termination/end). See [FMA-MMI] for normative requirements.

The requirements in this list apply when the MAC-NE is configured as the point of intercept for Lawful Intercept for PacketCable 2.0 voice calls and or subscriber high speed data content.

- The MAC-NE will accept MMI-based intercept provisioning from the MAC Manager, as specified in [FMA-MMI]. If the MAC-NE is configured to send intercepted content the MAC-NE MUST filter according to the provisioned filtering parameters, intercept the content, and transmit it toward the provisioned Delivery Function IP address and port using the provisioned transport method and with the provisioned DSCP or ToS settings.
- The MAC-NE MUST continue to mirror content toward the provisioned Delivery Function until the tap is deprovisioned by the MAC Manager or times out, whichever occurs first.

6.3.5.4.3 MAC Manager Requirements

The MAC Manager MUST support a method (API or SNMP MIBs) to provide information regarding the relationship of a cable modem, its associated CPE devices to a MAC-NE to aid on the provisioning of the Lawful Intercept session.

6.3.5.4.4 S-Leaf As An Intercept Point

The requirements for Lawful Intercept when the S-Leaf is the point of interception, the details of provisioning, and port/traffic mirroring functions are left to future study and definition.

6.3.5.5 Security Requirements

When the FMA System is configured for Lawful Intercept and the MAC-NE is configured as a point of intercept for PacketCable 2.0 and subscriber high speed data intercepts, the MAC Manager MUST support an SNMPv3-based method for configuring the taps as defined in the PacketCable 2.0 Electronic Surveillance Intra-Network Specification [PKT-ES2.0] and PacketCable 2.0 Electronic Surveillance Delivery Function to Collection Function Interface Specification [PKT-DCI]. The corresponding MIBs for configuring taps is documented in [FMA-OSSI]. The YANG based configuration for the MAC-NE is documented in [FMA-YANG].

When the FMA system is configured for Lawful Intercept and the S-Leaf is configured as a point of intercept for PacketCable 2.0 and subscriber high speed data intercepts, the intercept configuration is managed through an alternate interface, which is left for future study.

Duplicated PCES events sent from the MAC-NE are secured by way of transport over the PAI. PAI security is specified in [FMA-PAI].

When the MAC-NE is the point of intercept, the intercepted content sent from the MAC-NE travels in the data plane and is secured between the MAC-NE and S-Leaf by MACsec and/or IPSec.

6.4 Auxiliary Core System Operation

In order to make use of existing non-DOCSIS [R-PHY]-compliant Auxiliary Cores deployed in the operator's network, the MAC-NE needs to be able to support a subset of the MHAv2 specification's CCAP Core device definitions.

The MAC-NE MUST participate as a GCP client and initiate GCP connections as per [R-PHY] to any Auxiliary Cores configured by the MAC Manager. The MAC-NE MUST be capable of participating in mutual authentication negotiation using IPsec/IKEv2 secure communication with the Auxiliary Cores as defined in [R-PHY]. Note that this is a different mutual authentication mechanism than is specified in the FMA specifications between the MAC Manager and MAC-NE based on the TLS protocol.

The MAC-NE MUST accept GCP configuration from MHAv2 Auxiliary Cores, following the restricted rules for Auxiliary Cores defined in [R-PHY].

The MAC-NE MUST fail GCP connections and configuration from MHAv2 Principal Cores and log the event.

The MAC-NE MUST fail any GCP configuration related to DOCSIS services and log the event.

6.4.1 CCAP Core Types

In FMA, all CCAP Cores are defined as a type of Auxiliary Core with no support for Principal Cores. The MAC-NE MUST prevent a CCAP Core connected over GCP from attempting to become a Principal CCAP Core in the ConfiguredCoreTable. The MAC Manager, although it provides functionality similar to that of the MHAv2 Principal Core, is identified in the DHCPv4 and DHCPv6 sub-option encodings and identified in the MAC-NE's ConfiguredCoreTable, as a non-Principal Core.

A MAC-NE receiving a connection from one or more MAC Managers can also connect to multiple Auxiliary Cores. Each Auxiliary Core can utilize an independent subset of MAC-NE resources (e.g., one or more RF channel resources for delivering video and OOB). An Auxiliary Core can manage these resources on its own.

The MAC Manager is responsible for configuring the Auxiliary Cores that a MAC-NE will connect to over GCP. The MAC-NE MUST receive ConfiguredCoreTable entries for the Auxiliary Cores from the MAC Manager during RESTCONF configuration before allowing GCP configuration from those Auxiliary Cores to commence.

6.4.2 Auxiliary CCAP Core RCP Fields

The MAC-NE MUST provide write access to the subset of the RCP fields defined in [R-PHY], Annex B, "GCP Usage."

A future version of this specification will define requirements around which RCP trees are restricted.

6.4.3 FMA R-DTI Operations

MHAv2 [R-DTI] specifies timing and synchronization requirements on CCAP Cores and RPDs. The requirements apply to Video and OOB Cores and MAC-NEs, respectively.

The following statements apply to FMA.

- It is recommended that CCAP Cores are not used as the master timing clock over PTP or SynchE (for example, Grandmaster or Boundary Clock).
- For FMA, the Video Core MUST meet all EQAM-related CCAP Core requirements specified in [R-DTI].
- For FMA, the OOB Core MUST meet all OOB-related CCAP Core requirements specified in [R-DTI].

The following statements apply to MXh.

- For MXh operation with an MAC-NE, the MAC-NE MUST meet the timing requirements specified for the DAA CMTS in [SYNC].

- For MXh operation with an MAC-NE, the MAC-NE MUST meet the timing and performance requirements specified for the RPD in [SYNC].
- For MXh operation with a pCore/MAC-NE, the MAC-NE MUST meet the timing requirements specified for the RPD in [SYNC].
- For MXh operation with a vCore/MAC-NE, the MAC-NE MUST meet the timing requirements specified for the RPD in [SYNC].

The 1588 GM (Grandmaster) should be used to synchronize all FMA network elements requiring synchronization.

6.4.3.1 MAC-NE DTI Operations (non-MXh)

- The System Timing Clock of the MAC-NE MUST meet the MPEG-TS timing requirements in [R-DTI], Section 6.2, "MPEG-TS Video," (and all subsections).
- For protocol support, the MAC-NE MUST support the protocol requirements listed for the RPD in [R-DTI], Section 7.1, "Protocol Support."
- DOCSIS Timing Requirements
  - In Node_Slave mode as shown in Figure 32 and Figure 33, both the DOCSIS master clock frequency and the DOCSIS timestamp are driven by the PTP source and need to meet all applicable [DRFI] and [PHYv4.0] timing and synchronization requirements.
  - The MAC-NE MUST support the Node_Slave mode requirements listed for the RPD in [R-DTI], Section 7.2, "DOCSIS Timing Requirements."
  - The MAC-NE MUST meet the DOCSIS timing requirements specified in [R-DTI], Section 6.1, "DOCSIS Timing Requirements."

6.4.3.2 pCore DTI Operations (non-MXh)

- The pCore MUST meet all CMTS Core timing requirements specified in [R-DTI] for Node_Slave mode operation.
- Node_Master timing mode operation specified in [R-DTI] is not recommended for use with a pCore architecture as previously stated in this specification.

6.4.3.3 vCore DTI Operations

- The vCore MUST meet all CMTS Core timing requirements specified in [R-DTI] for Node_Slave mode operation.

6.4.4 FMA GCP Operations

MHAv2 [R-PHY] specifies GCP operation requirements on CCAP Cores and MAC-NEs.

- For FMA, the Auxiliary Core MUST meet all CCAP Core requirements that apply to Auxiliary Cores specified in [R-PHY] for GCP operation.

6.4.4.1 MAC-NE GCP Operations

The MAC-NE MUST meet all RPD requirements specified in [R-PHY] for GCP operation that supports communication with Auxiliary Cores.

6.4.4.2 pCore GCP Operations

The pCore MUST meet all CCAP Core requirements specified in [R-PHY] for GCP operation.

6.4.4.3 vCore GCP Operations

The vCore MUST meet all CCAP Core requirements specified in [R-PHY] for GCP operation.

Figure 23 shows the Auxiliary Core components of an FMA system and their interfaces to the MAC-NE. The following subsections explain the behavior and functionality of these components.

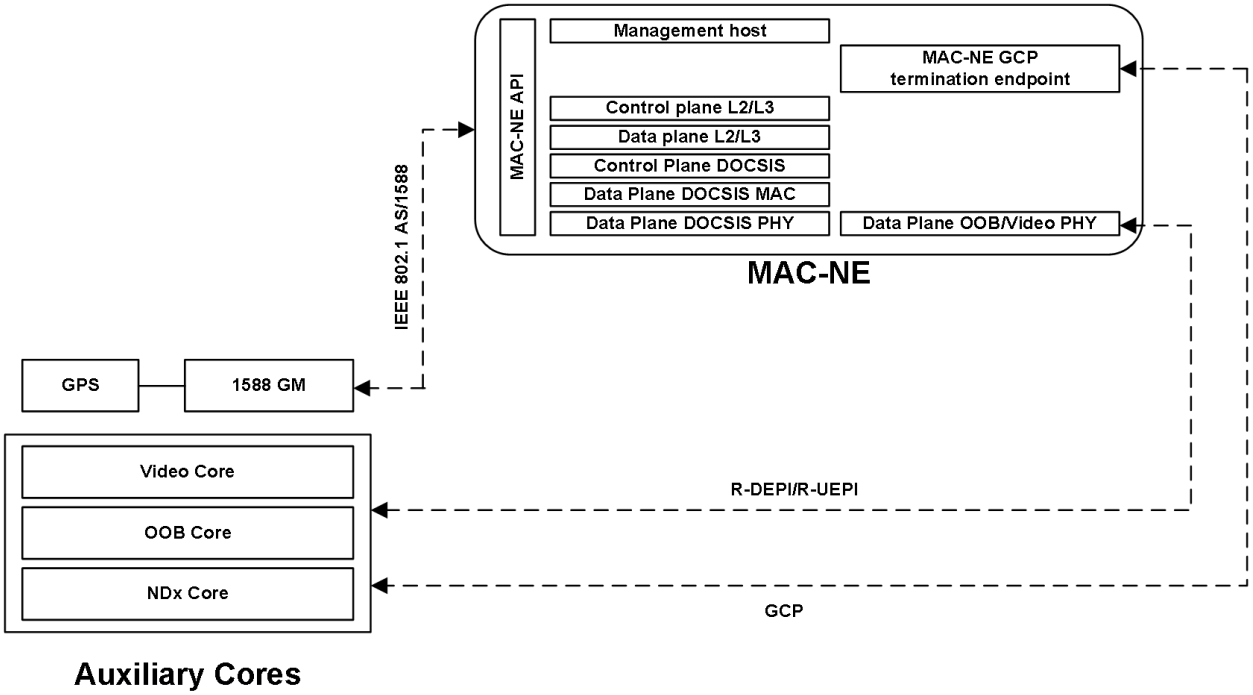

Figure 23 - FMA Video - Aux Core OOB System Components

6.4.5 Traffic Engine Operations

In lieu of an Auxiliary Core, an operator can deploy Traffic Engines for either or both video and OOB. Traffic Engines are not Auxiliary Cores and implement only a small subset of the [R-PHY] protocol. A Traffic Engine implements only the data plane functions of a CCAP Core that are limited to the reception or transmission of data on static L2TPv3 pseudowires. These devices generally do not support the L2TPv3 control plane layer of the protocol nor the GCP protocol required for the TE to perform configuration of the MAC-NE. Instead of supporting GCP, a Traffic Engine relies on the MAC Manager to configure the MAC-NE's resources (e.g., one or more RF channels) on its behalf. Specifically, the MAC Manager configures static pseudowires over which the MAC-NE receives or transmits data to/from the Traffic Engine. Certain types of parameters, such as downstream power, are common across resource sets. The MAC Manager is responsible for the configuration of these common parameters for the MAC-NE and for certain device management functions to maintain operation of the service.

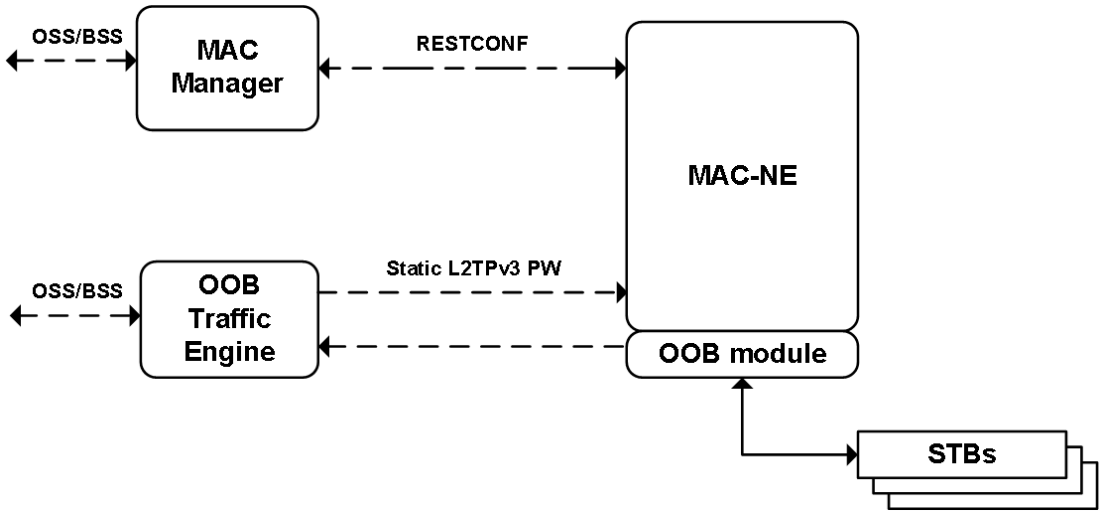

Figure 24 - FMA Video - OOB System Components

6.4.6 L2TPv3, R-DEPI, and R-UEPI

L2TPv3 Services refers to the MAC-NE forwarding of L2TPv3 pseudowires to and from head-end Auxiliary Cores and Traffic Engines developed to serve RPDs over the R-DEPI and R-UEPI protocols.

6.4.7 L2TPv3 Requirements

The MAC-NE MUST support dynamic DEPI signaling, as defined in [R-DEPI], to set up L2TPv3 pseudowires for MPEG video and OOB services.

The MAC-NE MUST support GCP setup of static PWs as specified in the [R-PHY] specifications and limited to L2TPv3 pseudowires specifically for MPEG video and OOB services.

The MAC-NE MUST support MAC Manager MMI setup, as specified in [FMA-OSSI], of L2TPv3 pseudowires for MPEG video and OOB services.

The MAC Manager MUST support MMI setup of L2TPv3 pseudowires for MPEG video and OOB services. This is required because Traffic Engines, unlike Auxiliary Cores, cannot configure the MAC-NE's channel resources on its own.

6.4.8 R-DEPI MPEG Video Multicast Support

Requirements in this section apply to L2TPv3 MPEG video multicast.

As defined in the [R-PHY] specifications, an MAC-NE MUST meet all RPD requirements specified in the [R-PHY] system specification and [R-DEPI] for support of dynamic and static MPEG video pseudowires. In MAC-NE deployments, the MAC-NE MUST meet all RPD requirements specified in [R-PHY] for synchronous and asynchronous video operating modes.

MAC-NE implementations MUST implement the [FMA-YANG] Static Pseudowire objects. The details of these objects are detailed in [FMA-YANG]. The MAC Manager MUST support the [FMA-YANG] Static Pseudowire objects to program an MAC-NE to join a video engine multicast output and direct the video content to a specific set of downstream channels.

R-DEPI is an IP-based pseudowire based on the L2TPv3 protocol that operates between (1) the EQAM MAC in the Video Core and the Video/OOB PHY in the MAC-NE and (2) the OOB and NDx Core and the Video/OOB PHY in the MAC-NE in the downstream direction. R-UEPI is an extension to R-DEPI that uses the same control plane structure and a unique set of encapsulations operating in the upstream direction. R-UEPI is used between the MAC-NE and OOB Core but not between the MAC-NE and a Video Core.

R-DEPI's role is to encapsulate MPEG-TS video packets generated by the Video Core (using MPT transport sub-type) and STB control packets generated by the OOB Core, transporting them through a Layer 2 or Layer 3 network, and delivering them to the MAC-NE for transmission to STBs. R-DEPI also encapsulates digital samples of downstream RF traffic for NDF and delivers them to the MAC-NE for transmission.

R-UEPI's role is to encapsulate OOB STB control packets that have been received and demodulated by the upstream PHY in the MAC-NE in order to transport them to the OOB Core for processing. Similar to the way in which R-DEPI operates, R-UEPI encapsulates digital samples of upstream RF traffic for NDR and delivers them to the OOB Core for processing.

The base protocol used for R-DEPI is the IETF Layer 2 Tunneling Protocol, version 3 (L2TPv3) specified in [RFC 3931]. L2TPv3 is a generic protocol for creating a pseudowire, mechanism for transparently transporting Layer 2 protocols over a Layer 3 network. Some examples of protocols supported by L2TPv3 include ATM, HDLC, Ethernet, Frame Relay, and PPP.

Each data packet contains a 32-bit session ID. In the MPT encapsulation, that session ID is associated with a single QAM Channel. When part of an L2TPv3 encapsulation, the UDP header is not used in the MHA protocols. The L2TPv3 session ID directly follows the IP header. It is worth noting that the L2TPv3 session ID lands in the same part of a packet as a classic UDP destination port/source port, allowing network equipment capable of classifying based on UDP headers to be leveraged in for classifying L2TPv3 headers.

L2TPv3 permits creation of a subheader whose definition is specific to the payload being carried. The control channel allows signaling messages to be sent between the core components and the MAC-NE. Typical control messages will first set up a "control connection" between the core components and the MAC-NE and then set up multiple data sessions (e.g., one for each downstream video QAM, OOB QAM, or NDF session, and one for each upstream OOB QAM or NDR session). Each session can be marked with different Differentiated Services Code Points (DSCPs) and can support different encapsulation protocols.

There are two main pseudowire techniques defined by R-DEPI, each supporting a variety of subtypes. The first technique, known as MPEG-TS mode, transports multiple 188-byte MPEG-TS video packets by placing them into the L2TPv3 payload with a unique subheader that contains a sequence number so packet drops can be detected. The encapsulation of MPTS (Multi Program Transport Stream) and SPTS (Single Program Transport Stream) video packets into MPEG-TS packets is performed in the Video Core. MPT is also used for SCTE 55-1 and SCTE 55-2 R-DEPI tunnels.

The second technique, known as the Packet Streaming Protocol (PSP), allows frames to be both concatenated, to increase network performance, and fragmented, in case the tunneled packets exceed the network maximum transmission unit (MTU) size. PSP is used for NDF R-DEPI tunnels and NDR R-UEPI tunnels. It is required in MHAv2 to carry DOCSIS 4.0 OFDM and OFDMA frames between the CCAP Core and RPD, but it is not required in R-MACPHY for this purpose because high-speed data, VoIP, and IP video traffic are carried natively over the operator digital optical transport network to the MAC-NE, where all DOCSIS processing resides.

6.4.9 Video Core Architecture

Figure 25 focuses on the functional entities and interfaces of the FMA Video Core architecture. The FMA Video Core architecture is identical to the Video Core architecture specified in MHAv2, with the exception of an FMA MAC Manager in place of the MHAv2 Principal Core in the overall MAC-NE configuration and management flow. The Edge Resource Manager (ERM) Video Management requirements, ERM interface, Video Core requirements, and video transport and control via DEPI and the GCP are all unchanged from MHAv2. RPD requirements apply as is to MAC-NEs, except as noted in the normative sections of this specification.

The MAC Manager identifies the Video Core to the MAC-NE so that the MAC-NE can connect to the Video Core using GCP, as is done in MHAv2.

Video management and control traffic from the ERM Video Management entities is terminated in the Video Core, which has a GCP-based configuration and management interface to the MAC-NE. GCP is used for initial MAC-NE configuration by the Video Core, as well as steady-state reconfiguration and management.

Video traffic from the Core Network is sent to the Video Core(s) for processing before being transported by way of the Core Access Devices (abbreviated in this diagram as "Access Switching") to the MAC-NE. Video traffic is encapsulated in DEPI tunnels between the Video Core and MAC-NE.

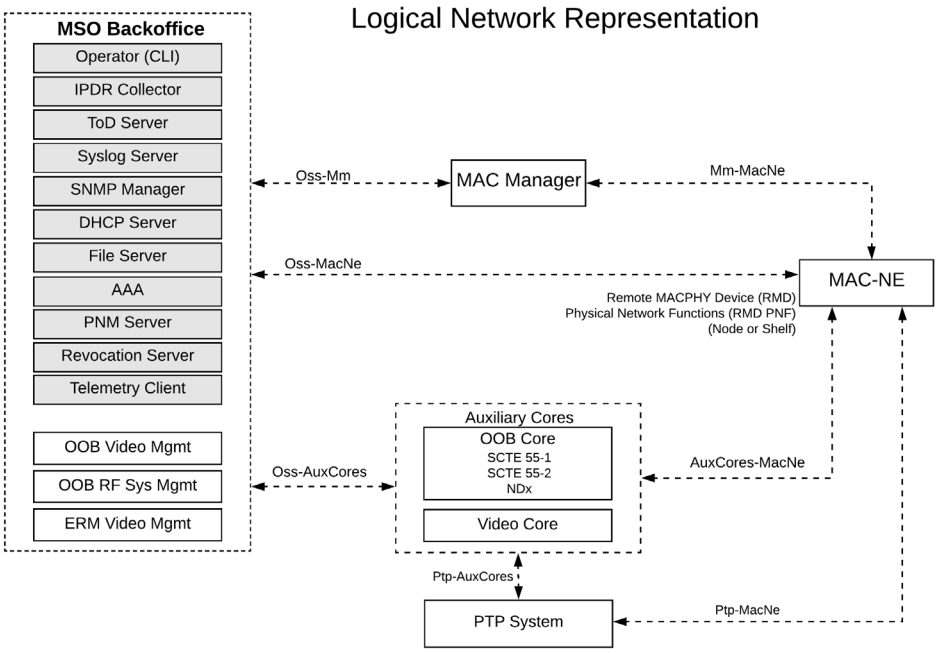

Figure 25 - FMA Video System Components

Several other video architectures have been considered in the Distributed CCAP Architecture (DCA). They are listed here for informational purposes only. These architectures have not been formally specified by CableLabs.

Remote EQAM is an architecture in which the entire EQAM is in the MAC-NE, effectively replicating a Video Core instance in each MAC-NE.

Video Engine is an architecture in which RF video from unmodified head-end video systems (e.g., I-CCAP, EQAM) is digitized before being transported to the MAC-NE. The MAC-NE processes the received digital samples and regenerates RF video downstream of the MAC-NE.

Discussions of video distribution (whether it be digital or analog video) over analog fiber are outside the scope of this specification.

6.4.10 OOB Core Architecture

The FMA OOB Core architecture is derived from the R-PHY OOB Core architecture with the following changes:

- All OOB data passing functions of the "CCAP Core" are performed by a separate "OOB Core" component.
- All configuration is performed by the MAC Manager.

The FMA OOB architecture is best understood by reviewing the R-PHY OOB architecture in [R-OOB].

6.4.11 OOB Support

The implementation for each of the OOB signals in an R-MACPHY system follows the same rules detailed in the R-PHY OOB Specification [R-OOB].

As shown in Figure 26, the signal processing of NDF and NDR in the MAC-NE also follows the same rules detailed in [R-OOB].

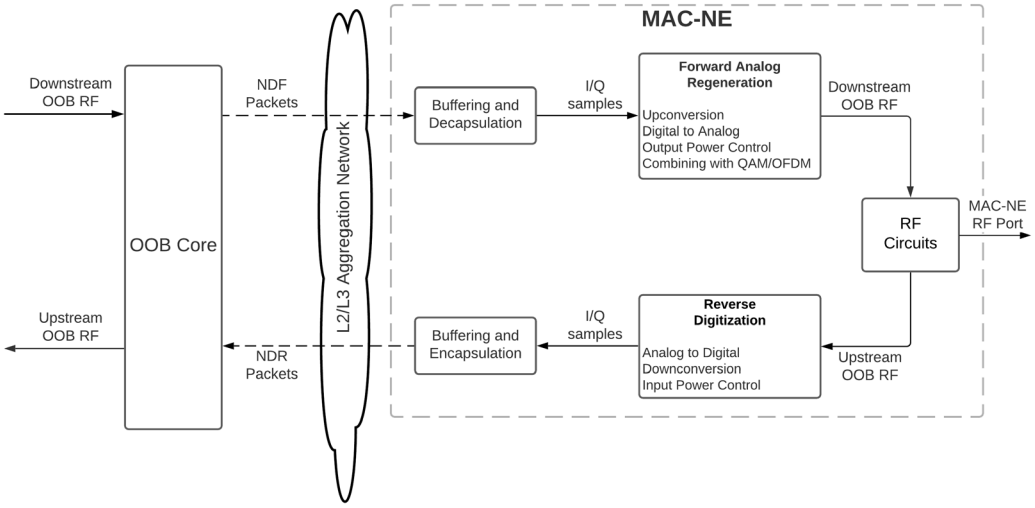

Figure 26 - NDF - NDR Signal Processing in the MAC-NE

6.4.12 R-PHY OOB Architecture Summary

Two legacy architectures for video OOB management with different "Headend RF" components are deployed to send and receive OOB RF signals to video STBs in a traditional HFC deployment:

- SCTE 55-1 features an "OM2000" OOB Modulator component for sending downstream OOB RF and an Advanced Return Path Demodulator (ARPD) for receiving upstream OOB RF.
- SCTE 55-2 features a "55-2 Controller" component that contains "55-2 Downstream Modules" for sending downstream OOB RF and "55-2 Upstream Modules".

The R-PHY OOB architecture calls for the CCAP Core to emulate the northbound interfaces of the Headend RF components to eliminate changes to other headend components. The R-PHY architecture splits the functions of the Headend RF components themselves between the CCAP Core and the RPD to emulate "virtual" Headend RF components.

Figure 27 depicts the R-PHY operation of SCTE 55-1 OOB.

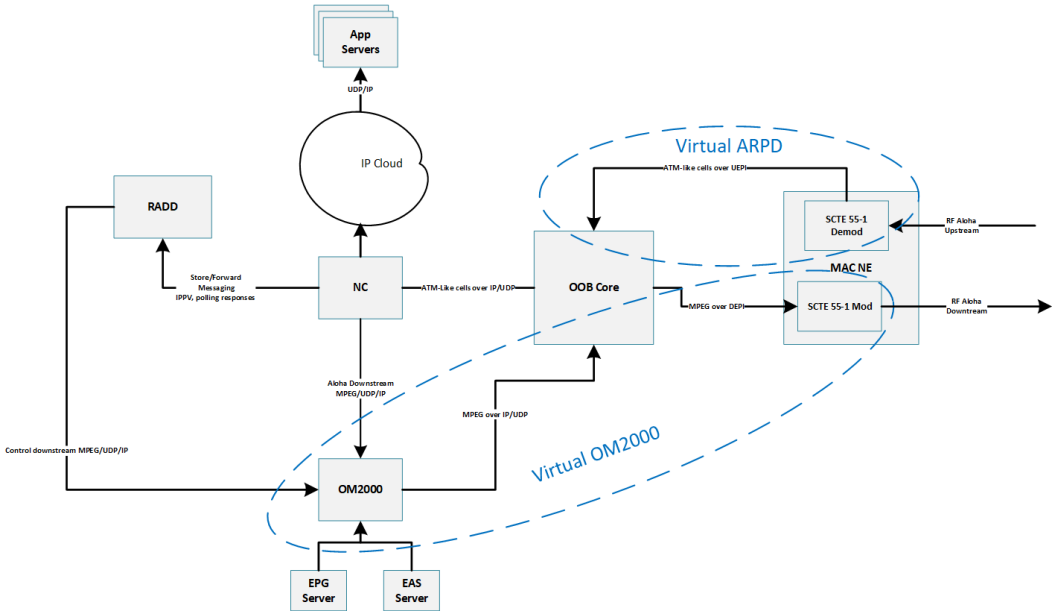

Figure 27 - FMA Operation of SCTE 55-1 OOB

Figure 28 depicts the operation of the SCTE 55-2 FMA architecture, in which the CCAP Core assumes the function of the "55-2 Controller".

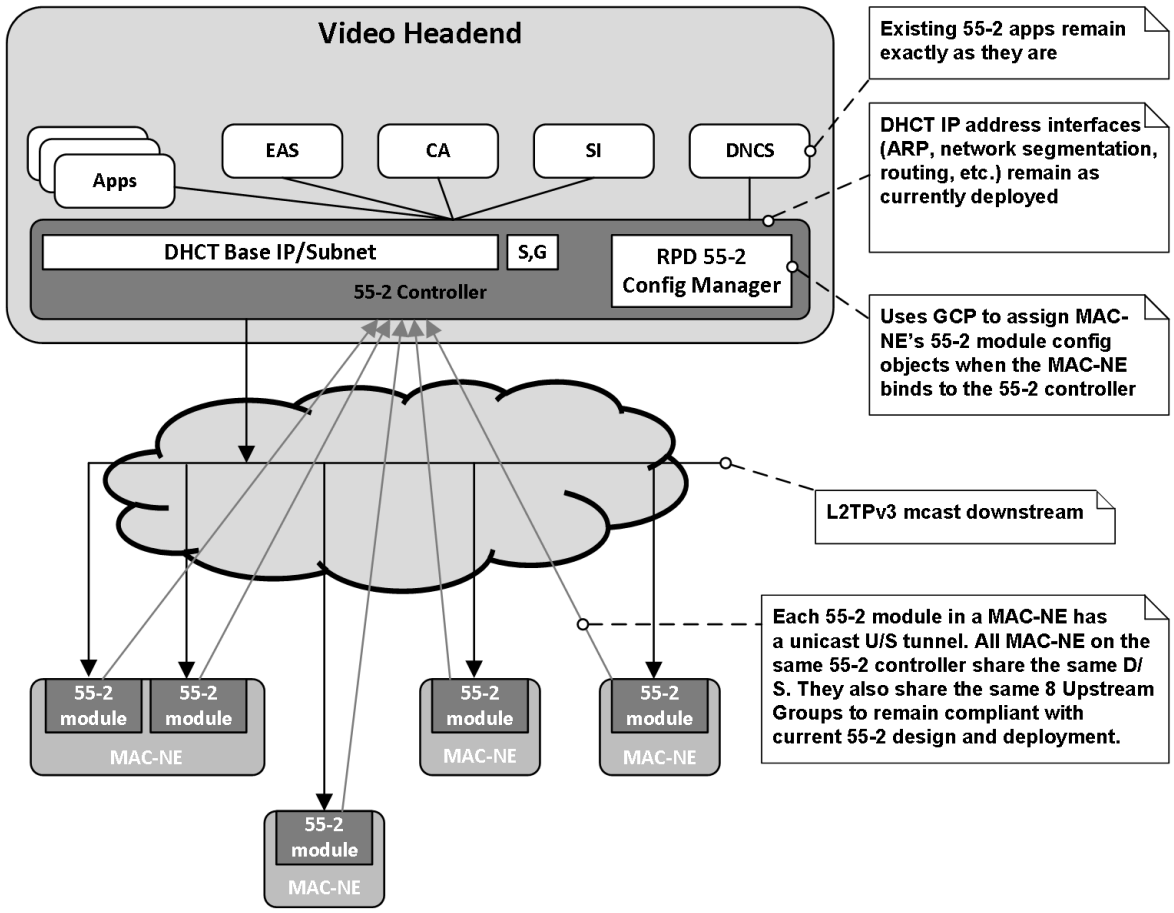

Figure 28 - FMA Operation of SCTE 55-2 OOB

In FMA, the Video Management component northbound of the OOB Core represents all functions northbound of the CCAP Core in the above diagrams.

The R-PHY OOB specification provides detailed requirements for the CCAP Core and RPD for each SCTE architecture. This specification uses the term "OOB Core Headend Functions" to refer to the functions defined in the R-PHY OOB specification for the CCAP Core for either SCTE 55-1 or SCTE 55-2. This specification uses the term "OOB RPD Functions" to refer to the functions defined in the R-PHY OOB specification for the RPD for either SCTE 55-1 or SCTE 55-2.

6.4.12.1 Static OOB L2TPv3 Pseudowires

Note that R-PHY specifications support either

- a static L2TPv3 pseudowire setup at RPDs with GCP, e.g., by a "Principal" CCAP core other than a "Data" CCAP Core; or
- a dynamic pseudowire setup with L2TPv3 control signaling from the "Data" CCAP Core.

The FMA OOB architecture supports only static OOB L2TPv3 pseudowires from the MAC-NE to an OOB core.

6.4.13 FMA OOB Architecture

The FMA OOB architecture is depicted in Figure 29.

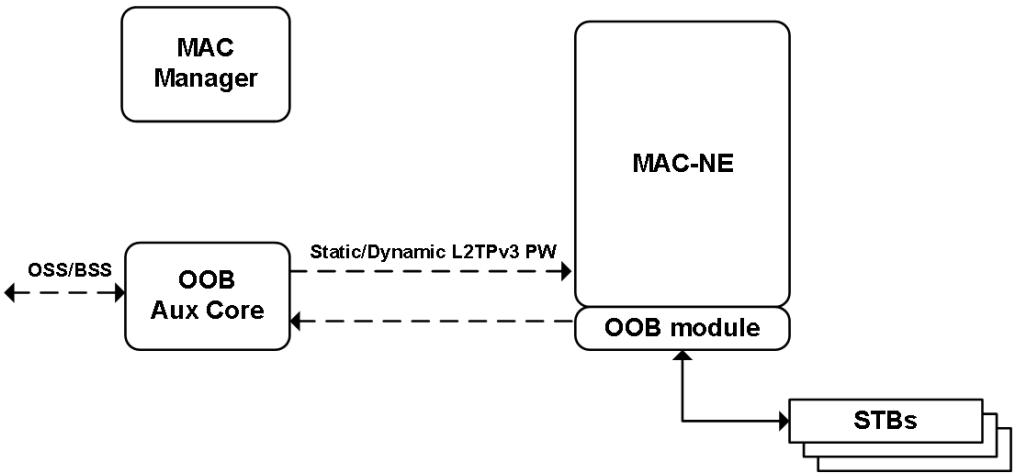

Figure 29 - FMA OOB Architecture

In general, the MAC-NE implements all OOB RPD functions by using RESTCONF rather than GCP for configuration, monitoring, and control of static OOB pseudowires.

The MAC-NE MUST meet all OOB RPD Function requirements not mentioning GCP, replacing the term "MAC-NE" for "RPD".

The MAC Manager MUST implement RESTCONF configuration of static L2TPv3 pseudowires for OOB data to the MAC-NE. The MAC-NE MUST implement all MAC Manager RESTCONF configuration of OOB static L2TPv3 pseudowires.

MAC Manager configuration via RESTCONF of communication between the OOB Core and Video Management applications northbound of it is for further study.

6.4.14 NDx Core Architecture

The NDx Core Architecture is shown in Figure 30. NDF/NDR channel requirements described in [R-OOB] will be adapted for FMA where NDF in MAC-NE processes receives a DEPI packet in downstream from the NDx Core and converts it to an analog stream. For NDR, a MAC-NE digitizes a received analog stream, converts it to a UEPI packet, and sends it upstream to the NDx Core. One or more NDF/NDR channels can reside in the NDx Core or MAC-NE.

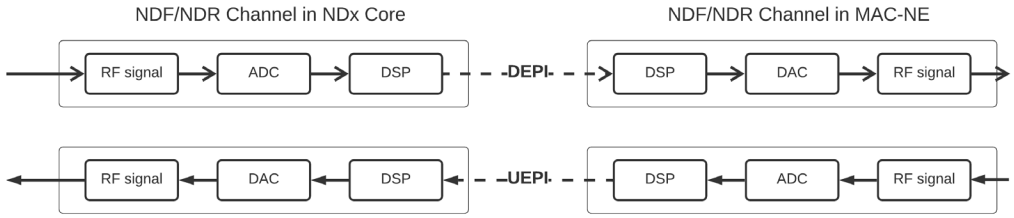

Figure 30 - NDx Core Architecture

6.4.15 Pilot Tones

Changes in the ambient temperature of outside plant causes changes to the attenuation characteristics of the coaxial cable, resulting in varying RF power or signal levels in the plant during various times of the day and across changing seasons. It is for this reason that pilot tones are often used in the HFC network to ensure that amplifiers in the network can correctly adjust their gain to compensate for changes in attenuation, thus keeping signals at the target

output level. Amplifiers accomplish this by using pilot tones in concert with the mechanism known as automatic gain control. The AGC circuits monitor the signal level of one or more RF signals called pilots and adjust the active device's RF output levels accordingly. Pilots can be analog or digital channels or might be implemented as dedicated continuous wave carriers.

To support pilot tone generation in FMA systems, the MAC-NE is required to be capable of generating these AGC tones and placing them in the appropriate portion of the downstream spectrum under control of the MAC Manager. Because R-PHY systems have the same need to support pilot tone generation, pertinent sections in the R-PHY specifications can be used for the FMA requirements.

6.4.15.1 Pilot Tone Generation

The requirements in [R-OOB], Section 8.1, "Pilot Tone Generation," apply to MAC-NEs with the following change:

- " MAC-NE" replaces "RPD" in all instances.

6.4.15.2 Pilot Tone Power Level and Frequency, Range, and Performance

The requirements in [R-OOB], Section 8.1.1, "Pilot Tone Power Level and Frequency, Range and Performance," apply to MAC-NEs with the following change:

- "MAC-NE" replaces "RPD" in all instances.

6.4.15.3 Alignment Tone Power Level and Frequency, Range, and Performance

The requirements in [R-OOB], Section 8.1.2, "Alignment Tone Power Level and Frequency, Range and Performance," apply to MAC-NEs with the following change:

- "MAC-NE" replaces "RPD" in all instances.

6.4.15.4 Pilot and Alignment Tones and Other OOB Signals Fidelity

The requirements in [R-OOB], Section 8.1.3, "Pilot and Alignment Tones and Other OOB Signals Fidelity," apply to MAC-NEs with the following change:

- "MAC-NE" replaces "RPD" in all instances.

6.4.15.5 Operation Outline

The requirements in [R-OOB], Section 8.1.3.1, "Pilot and Alignment Tones and Other OOB Signals Fidelity," apply to MAC-NEs and the MAC Manager with the following changes:

- "MAC-NE" replaces "RPD" in all instances,
- "MAC Manager" replaces "Principal CCAP Core", and
- "MAC Manager" replaces "CCAP Core".

6.4.16 Leakage Detection Signal Generation

To support RF leakage detection in FMA systems, the MAC-NE needs to generate the signals that are detected in the field by mobile and handheld leakage detection systems. Since Remote PHY systems had the same need to support leakage detection signals, the requirements for FMA can be specified by referencing the pertinent sections in Remote PHY specifications.

For informative information, see [R-OOB], Section 9.1, "Digital Leakage Detection Signals," and Section 9.2, "Legacy Analog Leakage Detection Support in Remote PHY." Appendix IV of [R-OOB], "Leakage Detection System Support in Remote PHY," describes how these systems work and regulations that operators face.

6.4.16.1 FMA Leakage Detection Signal Support Requirements

The requirements in [R-OOB], Section 9.3, "Remote PHY Leakage Detection Signal Support Requirements," and Section 9.3.1, "Frequency Accuracy," apply to MAC-NEs with the following change:

- "MAC-NE" replaces "RPD" in all instances.

6.4.17 Timing Architecture

Precision network timing is normally derived from a GNSS (Global Network Satellite System) or GPS (Global Positioning System) at a few centralized locations in operator networks. Timing from these locations is distributed to MAC-NEs via [IEEE 1588] Precision Time Protocol (PTP) over L2 Ethernet networks (via CIN).

Synchronization for most MAC-NE applications will be via the packet-based IEEE 1588 PTP. A simplified diagram showing a high-level view of this synchronization is shown in Figure 31.

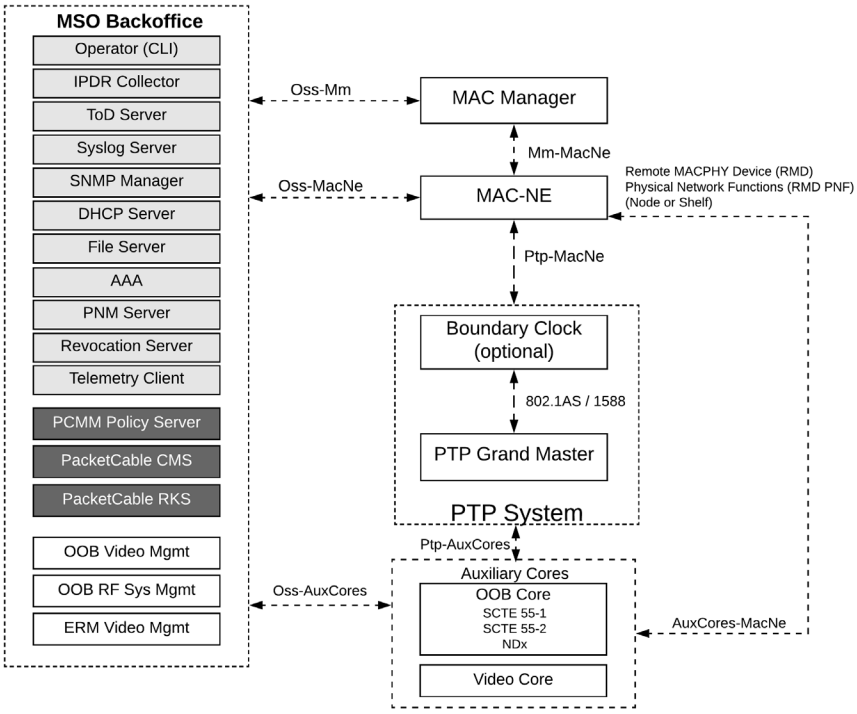

Figure 31 - MAC-NE Timing Architecture

Although a direct connection between the MAC-NE and the 1588 grand master (GM) is possible, a more general application may include a number of intermediate switches or routers between the 1588 GM and the MAC-NE, as shown in Figure 32.

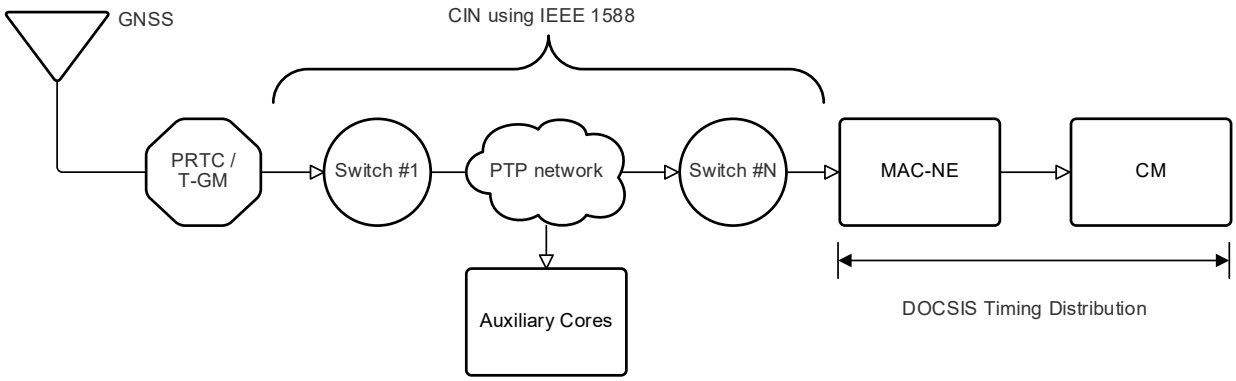

Figure 32 - Timing Distribution to MAC-NE through Switches - Routers (T-BCs)

The switches/routers shown in the above figure might or might not be 1588-aware. In some use cases, partial network timing support with the PTP profile defined in [ITU-T G.8275.2] will be a solution. Other use cases, like Mobile Xhaul (MXh), might require full network timing support with the PTP profile defined in [ITU-T G.8275.1][ITU-T G.8275.1][ITU-T G.8275.1]. In this second case, the switches/routers shown in the above figure include 1588-aware Telecom Boundary Clock (T-BC) functionality and provide significantly less timing degradation between the 1588 GM and the MAC-NE/RPD than the first case with 1588-unaware switches/routers. The [IEEE 1588] PTP assumes L2 links with symmetric upstream/downstream latency, which is not the case for DOCSIS. Thus, for precision timing applications such as MXh, the DOCSIS Time Protocol (DTP) would be used to bridge the timing link between the MAC-NE/RPD and the CM.

Additional information on timing distribution requirements for MAC-NEs (including MXh applications) is included in Section 9, MAC-NE Timing and Synchronization.

6.4.18 Remote DTI

Remote DTI (see [R-DTI]) provides timing synchronization between the Core Timing Devices and MAC-NEs based on the IEEE 1588v2 standard [IEEE 1588]. The protocol supports basic synchronization between the core components and MAC-NE for video and OOB services and precision time synchronization for services such as MXh. Additional information on timing distribution requirements for MAC-NEs (including the MXh application) is included in Section 9, MAC-NE Timing and Synchronization.

6.5 Network Considerations

6.5.1 S-Leaf to MAC Network Element

A "Leaf" is a switch that operates in the access layer of the network and typically does not support MACsec encryption. A "Secure Leaf" or "S-Leaf" switch is likely to support the MACsec encryption protocol and acts as the authenticator for the purposes of 802.1x network authentication of the MAC-NE.

The link between the S-Leaf and a MAC-NE generally use single-mode fiber and can be up to 80 km between endpoints without an amplifier using suitable optical modules. The link is expected to be at least 10 Gbps in each direction, which implies use of SFP+ optical modules. The MAC-NE or HFC FMA device are required to be capable of operating with either CWDM or DWDM optics. However, the wavelength plan is not specified.

At the MAC layer, the MAC-NE MUST use standard Ethernet II/DIX 2.0 framing in its link with the S-Leaf. Synchronous Ethernet ([ITU-T G.8261.1], [ITU-T G.8262/Y.1362], [ITU-T G.8264.1 /Y.1364]) capability may be required to provide timing accuracy for some use cases. See Section 9, MAC-NE Timing and Synchronization, for further detail.

The MAC-NE SHOULD implement a MACsec per [IEEE 802.1ae].

If the S-Leaf switch requires configuration to support VLANs; it is recommended that the switch offer an SDN-compatible control-plane interface so that it may be configured programmatically without manual intervention.

7 FMA SYSTEM OPERATIONS

7.1 Reuse of MHAv2 Requirements

7.1.1 Out-of-Band Operations

MHAv2 [R-OOB] specifies out-of-band processing requirements on CCAP Cores and RPDs. The requirements apply to OOB Cores and MAC-NEs, respectively.

- For FMA, the OOB Core MUST meet all CCAP Core requirements specified in [R-OOB].
- For FMA, the MAC-NE MUST meet all RPD requirements specified in [R-OOB].

7.1.2 DEPI Operations

MHAv2 [R-DEPI] specifies DEPI tunneling extensions and requirements on CCAP Cores and RPDs. The requirements apply to Video and OOB Cores and MAC-NEs, respectively.

- For FMA, the Video Core MUST meet all EQAM-related CCAP Core requirements specified in [R-DEPI], except as noted in this section. For FMA, the OOB Core MUST meet all OOB-related CCAP Core requirements specified in [R-DEPI], except as noted in this section.
    1. R-DEPI requirements that pertain only to transmission of downstream DOCSIS traffic do not need to be implemented in FMA core components.
    2. DEPI Latency Measurement does not need to be implemented in FMA core components. DEPI Latency Measurement is only for DOCSIS channels.
- For FMA, the MAC-NE MUST meet all RPD requirements specified in [R-DEPI], except as noted in this section.
    1. The MAC-NE MUST implement PSP Mode requirements applicable to NDF downstream traffic. The MAC-NE MAY implement PSP Mode requirements applicable to non-NDF downstream traffic. R-DEPI PSP is only specified to carry DOCSIS, NDF, and "future services" traffic at this time.
    2. R-DEPI requirements that pertain only to reception of downstream DOCSIS traffic do not need to be implemented at the MAC-NE.
    3. DEPI Latency Measurement does not need to be implemented at the MAC-NE. DEPI Latency Measurement is for DOCSIS channels.
    4. In MAC-NE deployments, the MAC-NE MUST meet all RPD requirements specified in [R-PHY] and [R-DEPI] for synchronous and asynchronous video operating modes.

7.1.3 UEPI Operations

MHAv2 [R-UEPI] specifies UEPI tunneling extensions and requirements on CCAP Cores and RPDs. The requirements apply to OOB Cores and MAC-NEs, respectively.

- For FMA, the OOB Core MUST meet all CCAP Core requirements specified in [R-UEPI], except as noted in this section.
    1. R-UEPI requirements that pertain only to reception of upstream DOCSIS traffic do not need to be implemented in FMA core components.
    2. OOB Cores are not required to implement upstream latency related requirements specified in R-UEPI. FMA does not contain DOCSIS MAC and upper layer processing in core components, including the upstream DOCSIS scheduler, and therefore does not experience MAP advance prediction issues related to CIN latency and jitter and RPD MAC preprocessing delay.
- The MAC-NE MUST meet all RPD requirements specified in [R-UEPI], except as noted in this section.

1. R-UEPI requirements that pertain only to transmission of upstream DOCSIS traffic do not need to be implemented at the MAC-NE.
2. The MAC-NE is not required to implement upstream latency and skew related requirements specified in R-UEPI. FMA MAC-NEs contain all DOCSIS MAC and upper layer processing, including the upstream DOCSIS scheduler, which makes requirements associated with CIN latency and jitter, MAC preprocessing delay, and skew on latency through an RPD unnecessary.

7.2 DSG Operation

The DOCSIS Set-top Gateway [DSG] specification defines an interface and associated protocol that allows OOB messaging between set-top devices and headend set-top controllers over a DOCSIS transport.

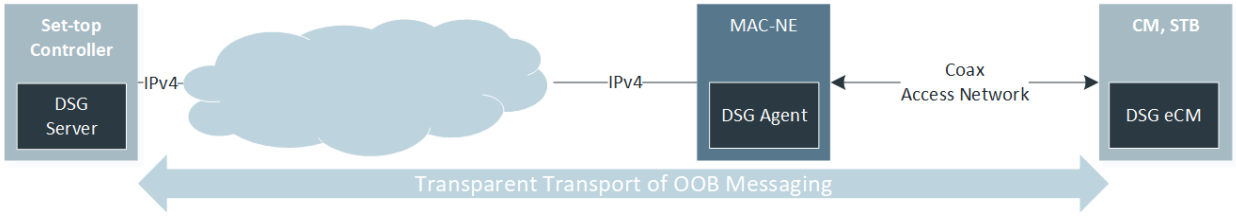

Figure 33 - DSG Operation

[DSG] specifies that a CMTS contains a DSG Agent, which translates IPv4 Multicast packet streams into DOCSIS flows and vice-versa. As per [DSG], DSG is intended to be an extension to IPv4 Multicast; the addressing of the IPv4 Multicast packet and the DSG Tunnel are the same. [DSG], as specified, cannot operate over IPv6-only networks because it utilizes IPv4 addressing in control plane communication, DOCSIS messages, TLV definitions, MIB definition, DOCSIS classifiers, and other IPv4-based assumptions. In FMA deployments over IPv4-capable networks, the MAC-NE is expected to implement a DSG Agent and comply with the requirements outlined in [DSG]. In FMA deployments without IPv4 (IPv6-only deployments), the FMA MAC-NE is still expected to be an IPv4 DSG Agent but will operate behind an IP-version network address translation (NAT) gateway, converting IPv6 Multicast groups to and from IPv4 Multicast.

The MAC-NE MUST implement a DSG Advanced Mode DSG Agent per [DSG].

In IPv4-capable networks, where the MAC-NE has received an IPv4 address, the MAC-NE MUST utilize DSG Agent functionality only over the IPv4 network even in cases where the MAC-NE also received an IPv6 address.

The MAC-NE MAY implement a DSG Basic Mode DSG Agent per [DSG].

If DSG Basic Mode is supported and the MAC-NE received an IPv4 address, the MAC-NE MUST utilize DSG Agent functionality only over the IPv4 network even in cases where the MAC-NE also received an IPv6 address.

[DSG] contains provisions for tunneling DSG IPv4 Multicast over IPv4 Unicast for legacy deployments that do not support IPv4 Multicast forwarding. This specification uses this framework to provide an IPv4 Multicast-over-IPv6 Multicast forwarding model. This specification proposes a set of requirements against a DSG NAT46 function, which could be implemented at the DSG Server or, more likely, as a separate network function.

7.2.1 DSG IPv6-Only Operation

DSG operation in IPv6-only networks (Figure 34) uses NAT64 [RFC 6144] and stateless address translation [RFC 6032]. This section contains a set of requirements against an IPv6-only MAC-NE and a set of expectations on a network function called the DSG NAT46 function. The DSG NAT46 function needs to be provided by an external device (DSG Server or a virtualized/physical device on the network) and is not an FMA-defined device or device under test. This solution does not need DNS, DNS46, or DNS64 functionality. The CIN between the DSG NAT46 function and the MAC-NE needs to support Multicast Listener Discovery (MLD) and be capable of carrying IPv6 Multicast traffic.

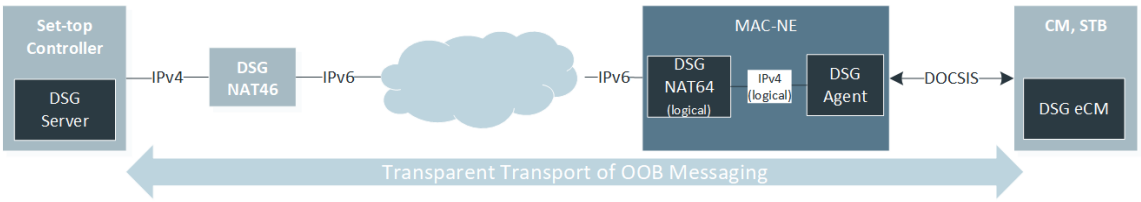

Figure 34 - DSG IPv6-Only Operation

The DSG Server will provide IPv4 Multicast traffic to a NAT46 function, which will convert the IPv4 Multicast packet headers to IPv6 Multicast packet headers. These IPv6 Multicast packets transit the CIN and reach the MAC-NE. The MAC-NE embeds a (logical) NAT64 gateway, which converts the IPv6 Multicast packet headers back into IPv4 Multicast packet headers before providing the IPv4 Multicast packets to the DSG Agent in the MAC-NE. Because the IPv4 Multicast address is an important part of the DSG protocol, the original IPv4 address needs to be preserved across this translation. The NAT64 function could be a separate network device, a virtualized network function, or logically embedded in the set-top controller and is left as an implementation detail. The logical NAT64 functionality exists within the MAC-NE device and is specified herein.

The DSG NAT46 network component will provide an IPv6 Multicast packet stream containing the DSG packets. The DSG NAT46 network component will replace the Multicast IPv4 header with a Multicast IPv6 header. The DAG NAT46 network component will populate the destination address with an "IPv4-Embedded IPv6 Address", in which the lower 32 bits of the Multicast address will contain the original DSG IPv4 Multicast address. The DSG NAT46 network component will provide an IPv6 Multicast packet stream in the FF05::/8 space.

The MAC-NE MUST implement a (logical) NAT64 function as per [RFC 6144], converting a subset of IPv6 Multicast packets into IPv4 packets to be (logically) provided to a DSG Agent within the MAC-NE. This communication does not need to be visible externally or available via a test point; therefore, depending on implementation, it can only exist logically. The MAC Manager MUST provide the list of IPv6 Multicast addresses to join through the [FMA-YANG] interface.

The MAC-NE MUST support "IPv4-Embedded IPv6 Address Prefix and Format" ([RFC 6052], section 2) Stateless NAT addressing in its NAT64 function, with the exception of subsection "Well-Known Prefix," which will be the Multicast prefix "FF05::/8". The MAC-NE MUST accept a PL value of 96, which places the IPv4 address in the bottom 32 bits of the IPv6 Multicast address. Other PL values can be supported. When the PL value is 96, the MAC-NE DSG NAT64 function MUST extract the lower 32 bits of the IPv6 Multicast address and use that address for all [DSG] requirements that require the DSG IPv4 Multicast address. If other PL values are supported, the MAC-NE MUST extract the original IPv4 address from the appropriate bits. Depending on vendor implementation, the MAC-NE NAT64 function MAY replace the IPv6 header with an IPv4 header before providing the packets to the DSG Agent.

The MAC-NE MUST support the MAC Manager providing a list of IPv6 Multicast Addresses over [FMA-YANG]. The MAC-NE MUST join the provided IPv6 Multicast addresses via MLDv2 [RFC 4604]. The MAC-NE MUST accept these addresses in the IPv4-embedded notation of "x:x:x:x:x:x:d.d.d.d". The MAC-NE MUST also be able to accept these addresses in the standard IPv6 notation. The MAC-NE SHOULD display the DSG Multicast addresses in IPv4-Embedded notation of "x:x:x:x:x:x:d.d.d.d".

7.3 Networking Considerations

It is important to distinguish between the terms "PHB-ID" and "DSCP" as used in MHAv2 and FMA specifications:

- a "PHB-ID" is a 6-bit value appearing in an L2TPv3 Attribute Value Pair (AVP), and
- a "DSCP" is a 6-bit value appearing in an IP packet header.

7.3.1 Network MTU

This section details the requirements for data packets that are forwarded across the CIN, which is the NSI of the MAC-NE.

7.3.1.1 Data Plane Packet Size Requirements

Data plane packets include any packets that are forwarded through the MAC-NE traffic and all control and management traffic terminated on the MAC-NE. The requirements related to packet size for DOCSIS MAC Management packets are specified in [MULPIv4.0].

The MAC-NE MUST support receiving an Ethernet frame size large enough to support the DOCSIS Ethernet frame size of 2000 bytes per [MULPIv4.0]. This will also need to account for DOCSIS headers and NSI encapsulation. The MAC-NE MAY support receiving Ethernet frame sizes larger than 2000 bytes.

The MAC-NE MUST adhere to the DOCSIS MAC management packet size requirements as specified in [MULPIv4.0].

7.3.1.2 Control Plane Packet Size Requirements

This section details the requirements related to packet size for control plane packets that are being transmitted to or received from the CIN. Control plane packets include any packets that are not data plane packets. They typically include traffic sourced from or terminated at the MAC-NE, including management and control traffic but does not include any DOCSIS MAC Management traffic. The requirements related to packet size for DOCSIS MAC Management packets are specified in [MULPIv4.0].

The IP MTU is defined as the maximum size of the Ethernet payload that a MAC-NE may transmit in the control plane. For IP packets, the IP MTU is the length of the IP datagram including the IP header. The IP MTU does not apply to data plane packets, as there are alternate methods for controlling data plane MTU.

A MAC-NE MUST default its IP MTU to 1500 bytes. The MAC-NE MAY support non-volatile configuration to modify the IP MTU.

A MAC-NE that uses an IP MTU of 1500 allows for up to 500 bytes of lower layer headers and trailers. Every Ethernet packet will include 14 bytes of standard Ethernet header and 4 bytes of Ethernet frame check sequence (FCS), leaving 482 bytes available for additional prefixes and suffixes required by other encapsulation protocols such as MACsec, Multiprotocol Label Switching (MPLS), 802.1Q, etc. See [IEEE 802.3] for further details on Ethernet MAC frame formatting.

To ensure that control plane services are not compromised because of path MTU incompatibilities, Path MTU Discovery (PMTUD) can be used on the MAC-NE for control plane traffic.

A MAC-NE SHOULD implement Path MTU Discovery per [RFC 1191] and [RFC 1981]. If Path MTU Discovery is supported, the MAC-NE MUST provide non-volatile configuration to enable/disable it.

A MAC-NE MAY implement Packetization Layer Path MTU Discovery per [RFC 4821]. If Packetization Layer Path MTU Discovery is supported, the MAC-NE MUST provide non-volatile configuration to enable/disable it.

7.4 FMA Operation with Static Pseudowires

See [R-PHY], Section 12.5, "Operation with Static Pseudowires," for requirements for CCAP Cores, Principal Cores, and RPDs.

It is intended that MAC-NEs in FMA will utilize the static pseudowire infrastructure, messaging, and behaviors specified for [R-PHY] RPDs. It is also expected that Auxiliary Cores in FMA utilizing static pseudowires will comply with the requirements related to that feature in [R-PHY].

In MAC-NE deployments, the MAC-NE MUST meet all RPD requirements specified in [R-PHY] for static pseudowires.

In MAC-NE deployments, the MAC-NE MUST implement the [FMA-YANG] static pseudowire objects. The details of these objects are in [FMA-YANG]. The MAC Manager MAY use the [FMA-YANG] static pseudowire objects to program a MAC-NE to join a video engine multicast output and direct the video/OOB content to a specific QAM resource.

7.5 Multiple Core Operation

7.5.1 Introduction

FMA permits MAC-NEs to be managed by one or more MAC Managers and zero or more Auxiliary Cores. The MAC Manager is a control entity for a MAC-NE. The MAC Manager controls the allocation of MAC-NE resources, software upgrade, and reset or reboot operations for the MAC-NE.

The MAC-NE MUST support at least two MAC Managers that are included in the unified MAC Manager list. The details of the operation of more than one MAC Manager in the system are defined exclusively in [FMA-MMI]. Please see [FMA-MMI] for definition of the Optimistic Concurrency Control mechanism.

An Auxiliary Core manages a subset of MAC-NE resources (e.g., particular channels). Each MAC-NE establishes its own GCP session with the Auxiliary Core, and then the Auxiliary Core establishes L2TPv3 control sessions with the MAC-NE. See Section 6.4, Auxiliary Core System Operation for further details. In the case of a Traffic Engine (TE), GCP is not used, rather the MAC Manager is responsible for configuring channel resources on the MAC-NE on behalf of the TE using one or more static pseudowires (PW). See Section 6.4.5, Traffic Engine Operations for further details.

7.5.2 MAC-NE Startup with Multiple Cores

7.5.2.1 Configured Core Table

An MAC-NE MUST implement a CcapCoreIdentification table object. The MAC-NE MUST perform a RESTCONF POST to the CcapCoreIdentification table.

The table has an entry for each MAC Manager or Auxiliary Core. The MAC Manager and Auxiliary Cores allocates and populates the entries during the MAC-NE initialization.

The MAC Manager MUST support the ability to modify the CcapCoreIdentification table.

Table 3 - ConfiguredCoreTable Object

Attribute Name	Type	Access	Type Constraints	Units	Default
ConfiguredCoreTable		RW			
Index	UnsignedByte	RW			
ConfiguredCoreIp	IpAddress	RW	Key		

A starting MAC-NE initially populates its ConfiguredCoreTable with the *CCAP Core-IP-Address* list learned from the MAC Manager configuration.

A MAC-NE MUST attempt to contact each core in the ConfiguredCoreTable in the order in which they appear in the table.

7.5.2.2 CCAP Core Identification Table

A MAC-NE MUST implement a CcapCoreIdentification table object. The MAC-NE MUST support the RESTCONF equivalent of a GCP AllocateWrite operation on the CcapCoreIdentification table.

The table has an entry for each MAC Manager or Auxiliary Core. The MAC Manager and Auxiliary Cores allocate and populate the entries during the MAC-NE initialization.

The MAC Manager MAY modify the CcapCoreIdentification table.

Table 4 - CcapCoreIdentification Table

Attribute Name	Type	Access	Type Constraints	Units	Default
CcapCoreIdentification		RW			
Index	UnsignedByte	RW	Key		

Attribute Name	Type	Access	Type Constraints	Units	Default
CoreId	HexBinary	RW	NULL if not allocated		
CoreIpAddress	IpAddress	RW			
IsPrincipal	Boolean	RW			
CoreName	String	RW			
VendorId	UnsignedShort	RW			
CoreMode	unsignedByte	RW			
InitialConfigurationComplete	Boolean	RW			
MoveToOperational	Boolean	RW			
CoreFunction	unsignedShort	RW			
ResourceSetIndex	unsignedByte	RW			
ProtocolSupport	unsignedShort	RW			

Descriptions for the CcapCoreIdentification attributes follow.

CcapCoreIdentification	No type defined
Index	index to the table
CoreId	a hex-binary string providing a unique identifier for the MAC Manager or Auxiliary Core; for example, a MAC address This field is used to mark the ownership of the entry.
CoreIpAddress	the IP address of the MAC Manager or Auxiliary Core
IsPrincipal	This field is set to "true" if this is the current active MAC Manager.
CoreName	user-friendly name for the MAC Manager or Auxiliary Core [related requirement follows this list]
VendorId	identifies the manufacturer of the MAC Manager or Auxiliary Core
CoreMode	indicates mode in which the core is currently acting: ContactPending, Active, Backup, NotActing, DecisionPending, OutOfService, Reconnecting
InitialConfigurationComplete	Boolean set to indicate that initial configuration of MAC-NE is complete
MoveToOperational	—
CoreFunction	a bitmap describing the functionality that the core will provide This field is valid for a core with CoreMode = Active or Reconnecting. This field is informational and is intended for use in problem resolution/debugging. Bit 0 Principal; Bit 1 DOCSIS Core; Bit 2 Broadcast video; Bit 3 Narrowcast video; Bit 4 SCTE 55-1 OOB; Bit 5 SCTE 55-2 OOB; Bit 6 NDF; Bit 7 NDR; Bit 8–15 reserved
ResourceSetIndex	an index into the resource table indicating the resources the core is using This field is valid for a core with CoreMode = Active or Reconnecting.
ProtocolSupport	a bitmap describing the protocols that the core can support Bit 0 GCP control plane; Bit 1 L2TPv3 control plane; Bit 2; Bit 3 L2TPv3 downstream data plane; Bit 4 L2TPv3 upstream data plane; Bit 5–15 reserved

The MAC-NE MUST set CoreMode to OutOfService if contact with an Active Core is lost.

After initial configuration from the MAC Manager, a MAC-NE MUST attempt to establish EAP-TLS authentication and a GCP TCP session (one per core) with each of the non-Principal Core ConfiguredCoreTable IP addresses, according to Section 13.5.2.5, 802.1x Mutual Authentication.

A CCAP Core that is contacted by the MAC-NE MUST populate its entry in the CcapCoreIdentification table as described in [R-PHY], Section "Initialization RCP Messages RPD and Cores," and Section "Reconnect Messages."

The MAC-NE MUST attempt to maintain an IPsec-authorized GCP session to each active non-Principal Core IP address in its CcapCoreIdentification table.

The MAC-NE MUST attempt to maintain an IPsec-authorized GCP session to each non-Principal Backup Core IP address in its CcapCoreIdentification table.

7.5.3 Resource Sets and Auxiliary Resource Assignment

A MAC-NE MUST implement the ResourceSet table, which identifies which Auxiliary Cores may control which MAC-NE object. The MAC-NE MUST support the AllocateWrite operation by Auxiliary Cores on the ResourceSet table. The MAC-NE MUST support a RESTCONF POST to the ResourceSet table.

Table 5 - ResourceSet Table

Attribute Name	Type	Access	Type Constraints	Units	Default
ResourceSet					
ResourceSetIndex	UnsignedByte	RW	Key		
CcapCoreOwner	HexBinary	RW	NULL if not allocated		
DsRfPortStart	Integer	RW	-1 if unused		
DsRfPortEnd	Integer	RW	-1 if unused		
DsChanGroup	NA	RW			
DsChanGroupIndex	Integer	RW	-1 if unused		
DsChanType	UnsignedByte	RW			
DsChanIndexStart	Integer	RW	-1 if unused		
DsChanIndexEnd	Integer	RW	-1 if unused		
UsRfPortStart	Integer	RW	-1 if unused		
UsRfPortEnd	Integer	RW	-1 if unused		
UsChanGroup	NA	RW			
UsChanGroupIndex	Integer	RW	-1 if unused		
UsChanType	UnsignedByte	RW			
UsChanIndexStart	Integer	RW	-1 if unused		
UsChanIndexEnd	Integer	RW	-1 if unused		

A resource set consists of a range of channels from start to end (inclusive) on particular RF ports from start to end (inclusive). The channels will not necessarily be in a single contiguous block, so multiple channel ranges can be assigned within a single ResourceSet entry.

The MAC-NE MUST reject a request to assign a resource to more than one ResourceSet. Thus, no two entries in the ResourceSet table can overlap, i.e., include the same channel on the same RF port.

An active MAC Manager MUST use a RESTCONF POST operation to allocate an entry in the ResourceSet table. The MAC Manager uses an entry in the ResourceSet table to define the resource sets that an Auxiliary Core is assigned to use. An active Auxiliary Core MUST use the GCP AllocateWrite operation to allocate an entry in the ResourceSet table. An Auxiliary Core uses this entry to define the resource set it intends to use. The AllocateWrite operation and resource release methods are described in [R-PHY], Section "AllocateWrite." The MAC-NE MUST NOT allow a resource set to be assigned to more than one active core at any time. The MAC-NE MUST send an error response to the MAC Manager if the MAC Manager attempts to assign a resource set to more than one active core at any time. The MAC-NE MUST send an error response to the Auxiliary Core if the Auxiliary Core attempts to assign to itself a resource set that is already assigned to another active core.

The MAC-NE MUST allow the MAC Manager to modify any entries in the ResourceSet and CcapCoreIdentification tables. The MAC-NE MUST implement sufficient entries in the AuxResourceAssignment table to assign each supported Auxiliary Core to one resource set.

The CcapCoreOwner field identifies the owner of the entry in the table. The CcapCoreOwner field is a hex-binary string providing unique identification of the Auxiliary Core, such as the MAC address of the core.

When resource allocation checking is enabled, the MAC-NE MUST use CcapCoreOwner to ensure that an Auxiliary Core that is not the owner of the table entry is not allowed to modify the entry.

The MAC-NE MUST NOT permit more than one active core to be configured to the same ResourceSetIndex. The MAC-NE MUST send an error response to the MAC Manager if an attempt is made to configure more than one active core to the same ResourceSetIndex.

An MAC-NE MUST implement the PermitAuxSelfConfiguration object to control whether Auxiliary Cores are permitted to configure their own resource sets.

Table 6 - PermitAuxSelfConfiguration Table

Attribute Name	Type	Access	Type Constraints	Units	Default
PermitAuxSelfConfiguration	Boolean	RW	Writeable by Principal Core only		False

An MAC-NE MUST enforce the policy that only the MAC Manager can write to PermitAuxSelfConfiguration.

When PermitAuxSelfConfiguration is "true", the MAC-NE MUST permit an Auxiliary Core to

- write to the ResourceSet table in order to claim a set of resources and
- assign the resources to itself by writing the ResourceSetIndex to the CcapCoreIdentification table entry for the Auxiliary Core.

Even when PermitAuxSelfConfiguration is "true", the MAC-NE MUST enforce the policy that an Auxiliary Core writes only its own ResourceSetIndex into the CcapCoreIdentification table.

When PermitAuxSelfConfiguration is "false", the MAC Manager is solely responsible for writing to the ResourceSetTable and writing the ResourceSetIndex to the CcapCoreIdentification table.

- When PermitAuxSelfConfiguration is "false", the MAC Manager MUST create an entry in the ResourceSetTable for each active Auxiliary Core.
- When PermitAuxSelfConfiguration is "false", the MAC Manager MUST populate the entry in the CcapCoreIdentification table for each Auxiliary Core including the ResourceSetIndex to identify the resources assigned to the auxiliary.
- When PermitAuxSelfConfiguration is "false", the MAC Manager MUST set the CoreMode in the CcapCoreIdentification table entry for each Auxiliary Core to ContactPending.
- When PermitAuxSelfConfiguration is "false", an Auxiliary Core MUST set the CoreMode in the CcapCoreIdentification table entry during configuration, as described in [R-PHY], Section B.3.2, "Initialization RCP Messages RPD and Cores."
- When PermitAuxSelfConfiguration is "false", the MAC-NE MUST NOT allow an Auxiliary Core to write to the ResourceSetTable.

7.5.4 MAC-NE Read Operations

The MAC-NE MUST support concurrent reads of any objects by any connected MAC Manager or Auxiliary Core.

7.5.5 MAC-NE Write Operations

Writes to objects specified in the ResourceSetTable are controlled by the ResourceAllocationCheck and PermitAuxSelfConfiguration control attributes.

ResourceAllocationCheck controls whether a MAC-NE performs resource ownership checks when a MAC Manager or Auxiliary Core writes to a MAC-NE variable.

The MAC-NE MUST allow the active MAC Manager to modify ResourceAllocationCheck.

The MAC-NE MUST allow the active MAC Manager to modify PermitAuxSelfConfiguration.

The MAC-NE MUST NOT allow a MAC Manager that is not the active MAC Manager to modify ResourceAllocationCheck.

The MAC-NE MUST NOT allow an Auxiliary Core to modify ResourceAllocationCheck.

The MAC-NE MUST NOT allow a MAC Manager that is not the active MAC Manager to modify PermitAuxSelfConfiguration.

The MAC-NE MUST NOT allow an Auxiliary Core to modify PermitAuxSelfConfiguration.

If ResourceAllocationCheck is True and PermitAuxSelfConfiguration is True, the MAC-NE MUST reject an attempt by an Auxiliary Core to write to an object that is specifically authorized to another core by a combination of ResourceSetTable and CcapCoreIdentification table entries (e.g., a specific port + channel).

If ResourceAllocationCheck is True and PermitAuxSelfConfiguration is True, the MAC-NE MUST reject an attempt by the MAC Manager to write to an object that is specifically authorized to another Auxiliary Core by a combination of ResourceSetTable and CcapCoreIdentification table entries.

If ResourceAllocationCheck is True and PermitAuxSelfConfiguration is False, the MAC-NE MUST reject an attempt by an Auxiliary Core to write to an object that is specifically authorized to another Auxiliary Core by a combination of ResourceSetTable and CcapCoreIdentification table entries.

If ResourceAllocationCheck is True and PermitAuxSelfConfiguration is False, the MAC-NE MUST allow the MAC Manager to write to an object that is specifically authorized to another Auxiliary Core by a combination of ResourceSetTable and CcapCoreIdentification table entries.

If ResourceAllocationCheck is False and PermitAuxSelfConfiguration is True, the MAC-NE MUST NOT perform any resource checks based on the ResourceSetTable and the CcapCoreIdentification table entries.

If ResourceAllocationCheck is False and PermitAuxSelfConfiguration is False, the MAC-NE MUST NOT perform any resource checks based on the ResourceSetTable and the CcapCoreIdentification table entries.

7.6 Resilient Software Upgrades

MAC-NEs support software image management features that help to eliminate failures in unsuccessful software upgrades and unrecoverable software errors. For example, an MAC-NE could "brick" after a software upgrade, becoming inoperable or losing the ability to connect to the network. Such failures are highly undesirable because they require costly service interventions. Software management features, such as the support for multiple, redundant software images combined with the capability to fall back to another image in case of repeated boot or system failures reduce issues resulting from SW upgrades and lower the overall FMA network maintenance costs. This section is intended to provide guidance on the support for multiple software images in the MAC-NE.

A MAC-NE MUST support a minimum of two software images.

The software images are identified by "software-image-index" attribute of the [FMA-YANG]. The MAC-NE advertises its support for multiple SW images via a set of capabilities defined in the YANG grouping "software-image-capabilities". These capabilities define how many software images the MAC-NE supports and which images can be upgraded via SSD.

Under normal circumstances, the MAC-NE executes the Main Software Image (MSI). The MSI is the most recent SW image identified with "software-image-index" value of 0. The MAC-NE downloads the MSI, successfully validates and boots for operation under normal circumstances. The method by which the MAC-NE validates the downloaded software image after SSD procedure is left to vendor implementation.

When the MAC-NE successfully completes the SSD procedure of the MSI, the MAC-NE MAY relocate the current MSI to another software-image-index value. This optional functionality simplifies SW image management flow

because it permits the operator to always upgrade the MSI, and to maintain previously used SW images as the backup to the MSI.

The MAC-NE MAY decide to boot an image different then MSI. This can happen, for example, in case of repeated failures of the MSI image. The criteria and the procedure(s) by which the MAC-NE decides to boot image other than the MSI is left to vendor implementation.

The MAC-NE MUST allow the MAC Manager to activate any software image which was previously downloaded and validated by the MAC-NE. The MAC-NE MAY reject an attempt to activate an image it does not qualify as fit for the activation. For example, the MAC-NE can reject the activation of an image that has previously repeatedly failed to boot or encountered SW exceptions. Once activated, the MAC-NE uses this image for subsequent boots until another image is downloaded or activated.

The status of support for the MAC-NE's ability to perform partial software upgrades is for further study.

8 MAC MANAGER AND MAC-NE MANAGEMENT OPERATIONS

The MAC Manager serves as an aggregation point for one or more MAC-NE devices for management purposes. Northbound, it exposes all defined OSSI management interfaces [FMA-OSSI] including SNMP, SNMP Notifications/Traps, Syslog, and IPDR. Southbound, it connects to MAC-NE devices using the MAC Manager Interface (MMI) defined in [FMA-MMI].

The MAC-NE supports [FMA-MMI] by acting as a server for the defined RESTCONF interface and YANG modules. The MAC Manager supports [FMA-MMI] by acting as a client that reads and writes the YANG modules defined in [FMA-YANG].

Refer to [FMA-OSSI] for a description of the FCAPS Network Management Model and the FMA management architectural overview.

8.1 Management Interface Models

8.1.1 Southbound Interface Protocols

RESTCONF [RFC 8040] with a fully-defined YANG model [FMA-YANG] is the defined API between the MAC Manager and the MAC-NE.

8.1.2 Southbound Interface Models

Southbound interfaces between the MAC Manager and the MAC-NE are defined using YANG data models. To manage those data models, RESTCONF has been chosen as the protocol for creating, reading, updating, and deleting instances of YANG objects between the MAC Manager and the MAC-NE. This interface is specified in [FMA-YANG].

8.1.3 Northbound Interface Protocols

Northbound interface protocols between the MAC Manager and back office are defined in [FMA-OSSI].

8.1.4 Northbound Interface Models

Northbound interface models between the MAC Manager and back office are defined in [FMA-OSSI].

8.2 FCAPS Management Operations

8.2.1 Fault Management

8.2.1.1 Event Logging and Reporting

This section defines requirements for remote monitoring/detection, diagnosis, reporting, and correction of problems for the MAC Manager and the MAC-NE. Refer to [FMA-OSSI] for additional details.

Asynchronous notifications are supported over [FMA-MMI] through the use of RESTCONF push notifications. The notifications are defined in [FMA-YANG]. Refer to [FMA-MMI] for additional details.

8.2.1.2 MAC Manager Syslog Support

Refer to [FMA-OSSI] for additional details.

8.2.1.3 MAC-NE Syslog Support

Refer to [FMA-OSSI] for additional details.

8.2.2 Configuration Management

8.2.2.1 Southbound Configuration Management Operations

The [FMA-MMI] and [FMA-YANG] standards define the configuration interface for MAC Manager to MAC-NE configuration.

8.2.2.2 Northbound Configuration Management Operations

Refer to [FMA-OSSI] for additional details.

8.2.3 Accounting Management

8.2.3.1 Southbound Accounting Management Operations

This is not specified at the MMI.

8.2.3.2 Northbound Accounting Management Operations

Refer to [CCAP-OSSIv3.1] for additional details on the Subscriber Accounting Management Interface Specification.

8.2.4 Performance Management

8.2.4.1 Southbound Performance Management Operations

The [FMA-MMI] and [FMA-YANG] standards define the status information for MAC Manager to MAC-NE status. This also includes PNM information.

8.2.4.2 Northbound Performance Management Operations

Refer to [FMA-OSSI] for additional details.

8.2.4.3 Streaming Telemetry

Refer to [FMA-OSSI] for further details.

8.2.4.4 IPDR/SP

This section is for further study.

8.2.5 Security Management

8.2.5.1 Security Considerations

Because of the disaggregated nature of distributed architecture systems, the subcomponents in FMA may not be collocated in a secure facility. In most cases, it is likely that some of the components will be placed in insecure locations. It is imperative that the links between system components are secured to ensure authenticity and integrity of data throughout the system.

The general approaches to network security to be considered include

- use of [ITU-T X.509] certificates for authentication of system interfaces and
- handshake protocols such as TLS, providing methods for mutual authentication and negotiation of cryptographic methods to provide for encrypted communication between elements.

Other considerations and good practices, such as those listed below, should also be followed, but they are outside the scope of standardization:

- secured storage of secrets or private keys associated with certificates and

- application of good design principles and security measures to remove or otherwise minimize the impact of attacks such as denial of service attacks, backdoor entries, operating system and application layer attacks, improper authorizations, password attacks, or other common methods of cyber-attacks.

8.2.5.2 Southbound Security Management Operations

Refer to [FMA-MMI] for additional details.

8.2.5.3 Northbound Security Management Operations

Refer to [FMA-OSSI] for additional details.

8.3 Subscriber Database

The FMA Subscriber Database holds mappings of the CM and CPE to MAC-NE, among other things. It includes the MAC and IP addresses of registered CMs, the MAC and IP addresses of online CPE, and a mapping from the CM and CPE to the MAC-NE on which they are registered.

Because of its general applicability in FMA, the Subscriber Database is maintained by the MAC Manager. Operational requirements for Subscriber Database maintenance on the MAC Manager and MAC-NE are provided in the sections below.

The way in which the PAG and other FMA applications access and synchronize to the Subscriber Database is implementation specific.

8.3.1 MAC Manager Requirements

The MAC Manager MUST maintain a Subscriber Database that contains MAC and IP addresses for all online CMs and CPE within the MAC Manager's scope of control, as well as a mapping of CM and CPE to MAC-NE IP address.

The MAC Manager MUST update the Subscriber Database when a CM comes online (i.e., is registered) or goes offline.

The MAC Manager MUST update the Subscriber Database when a CPE comes online or goes offline (i.e., obtains an IP address).

The MAC Manager MUST provide real-time access to the containing MAC-NE IP address using any one of the CM MAC, CM IP, CPE MAC or CPE IP addresses as the key into the Subscriber Database.

The MAC Manager MUST provide the PAG with an interface to determine the mapping of online CM and CPE to MAC-NE IP Address. The MAC Manager to PAG interface is implementation-specific.

The MAC Manager MUST synchronize the Subscriber Database after a loss of connectivity with an MAC-NE (e.g., due to RMD reboot).

8.3.2 MAC-NE Requirements

The MAC-NE MUST notify the MAC Manager when a CM comes online (i.e., is registered). The MAC-NE MUST indicate the CM's MAC and IP addresses when notifying the MAC Manager of a CM coming online.

The MAC-NE MUST notify the MAC Manager when a CM goes offline. The MAC-NE MUST indicate the CM's MAC address when notifying the MAC Manager of a CM going offline.

The MAC-NE MUST notify the MAC Manager when a subscriber CPE comes online (i.e., obtains an IP address). The MAC-NE MUST indicate the CPE's MAC and IP addresses when notifying the MAC Manager of a CPE coming online.

The MAC-NE MUST notify the MAC Manager when a subscriber CPE goes offline. The MAC-NE MUST indicate the CPE's MAC address when notifying the MAC Manager of a CPE going offline.

The MAC-NE MUST respond to synchronization requests from the MAC Manager by providing MAC and IP addresses for all of its online CMs and CPE.

8.4 Interface Management

The MAC Manager is responsible for configuration of interfaces supported by the MAC-NE. The following section provides informative guidance and normative requirements for MAC-NE interface management.

8.4.1 MAC Domain Configuration

The MAC-NE MUST support the ability to create DOCSIS MAC domains via MMI. The MAC-NE MUST support configuration of mac-domain list elements with mac-domain-name in the form of an arbitrary string.

9 MAC-NE TIMING AND SYNCHRONIZATION

9.1 Introduction

The requirements for timing and synchronization of the MAC-NE come from the following areas:

- existing DOCSIS specification requirements,
- Modular Headend v2 requirements,
- precision timing for applications like MXh.

These requirements place definitions and constraints on the use of the DOCSIS master clock and the DOCSIS timestamp. They are the minimum reasons for doing timing synchronization. There could be other requirements that also exist.

There could be different use cases of MAC-NE deployments that would affect the set of timing requirements that the MAC-NE should follow. There are 3 main use cases, listed below.

- The MAC-NE provides services that do not require any external timing synchronization. These services can include DOCSIS services, Video (Async mode), or OOB (55-1 and/or 55-2 Async mode). In these use cases, there are no external timing synchronization requirements imposed on the MAC-NE.
- The MAC-NE provides services that require frequency synchronization. These services can include Video (Sync mode), OOB (55-2 Sync mode), NDF/NDR, some forms of Leakage Detection Signal Generation, or MXh frequency delivery (LTE FDD).
- The MAC-NE provides services that require phase and frequency synchronization like MXh (LTE TDD and 5G). In this case, the MAC-NE needs to support the strict timing requirements of MXh defined in [SYNC].

9.2 Timing Deployment Scenarios for the MAC-NE

9.2.1 No External Timing Requirements

Similar to the I-CCAP case, when the MAC-NE is not required to support services that require timing synchronization (neither frequency nor phase), it can work in a "free-run" mode, in which it relies on its internal oscillator to drive the DOCSIS clock.

When the MAC-NE is supporting services that do not require external timing synchronization, the MAC-NE MUST support running in a "free-run" mode.

The MAC-NE MUST meet the DOCSIS timing requirements as specified in [DRFI] and [PHYv4.0].

9.2.2 External Frequency Synchronization

The following services require the MAC-NE clock to be synchronized to an external timing source:

- Video (Sync mode),
- OOB (55-2 Sync mode),
- NDF/NDR,
- MXh (frequency delivery),
- signal leakage detection.

As defined in [ITU-T G.8261.1], there are two main classes of synchronization methods in the Ethernet network:

- The MAC-NE is deployed with plesiochronous and network synchronous methods (i.e., a reference timing signal distributed over the synchronous physical layer). Synchronous Ethernet network is defined for this case.

- The MAC-NE is deployed with packet-based methods for frequency synchronization. (PTP is one of the methods using [ITU-T G.8265.1].)

The methods supported by the MAC-NE node depend on the level of accuracy required by various applications.

It is advised that the MAC-NE not attempt to synchronize to phase in these cases, as the phase adjustments needed might influence the frequency accuracy and, in some cases, may conflict with DOCSIS clock requirements.

When the MAC-NE provides services that require frequency synchronization, the MAC-NE MUST meet the frequency accuracy requirements defined in [R-DTI]..

9.2.2.1.1 Synchronous Ethernet

A Layer 1 frequency synchronization technique provides a higher level of frequency accuracy and stability, and as a result, a higher level of time accuracy becomes possible as well. Synchronous Ethernet is designed to deliver a physical layer clock through the network. A reference timing signal traceable to a Primary Reference Clock (PRC) is injected into the Ethernet switch using an external clock port. This signal is extracted and processed via a synchronization function before injecting timing onto the Ethernet bit stream. The synchronization function provides filtering and may require holdover.

First defined in the 2006 version of, then complemented by the 2008 version of [ITU-T G.8261.1], [ITU-T G.8262/Y.1362], [ITU-T G.8264.1 /Y.1364], and a new release of [ITU-T G.781], Synchronous Ethernet specifies not only the method and requirements for frequency recovery and transmission, but also standardizes the advertisement of clock quality through the network. Like all Layer 1 frequency synchronization techniques, all network elements between network segments need to be capable of recovering and passing the frequency downstream. Therefore, changing a path from Ethernet to Synchronous Ethernet requires all nodes in-line to be changed to use Ethernet Equipment Clock (EEC).

Synchronous Ethernet provides frequency traceability but not time information.When the MAC-NE is providing services that require frequency synchronization, the MAC-NE MAY support Synchronous Ethernet in order to lock to an external frequency.

9.2.2.1.2 Precision Time Protocol

PTP (see [IEEE 1588]) can be used to achieve frequency synchronization. [ITU-T G.8261.1] specifies the hypothetical reference model (HRM) and the packet delay variation (PDV) network limits applicable when frequency synchronization is carried via packets and is recovered according to the adaptive clock recovery method as defined in [ITU-T G.8261.1] and [ITU-T G.8260].

The [ITU-T G.8265.1] profile is dedicated for frequency synchronization. Because frequency information does not require measurement of the round-trip delay between the slave and its master, this profile allows both one-way communication (only Announce and Sync/follow-up messages are used) and two-way communication (Delay request/response messages are used as well). The one-way communication reduces the amount of PTP packets flowing towards the slave and eliminates any PTP packets sent from the slave.

The [ITU-T G.8275.2] profile is aimed at both frequency and phase synchronization. Only two-way communication is supported in this profile. Because [R-DTI] adopted [ITU-T G.8275.2] as the profile to be used for R-PHY, it may be used for the MAC-NE as well, assuming only frequency information is used.

When the MAC-NE is supporting services that require frequency synchronization, the MAC-NE MUST support being a 1588 OC slave and lock to its 1588 master frequency.

When the MAC-NE is supporting services that require frequency synchronization, the MAC-NE MUST support PTP messages over the following encapsulations defined in [IEEE 1588]:

- PTP over UDP over IPv4
- PTP over UDP over IPv6.

The MAC-NE MUST support the unicast model defined in [IEEE 1588].

When the MAC-NE is supporting services that require frequency synchronization, the MAC-NE MUST support either the [ITU-T G.8265.1] or [ITU-T G.8275.2] profile.

When the MAC-NE is supporting services that require frequency synchronization, the MAC-NE MAY attempt to lock to its 1588 master phase.

When the MAC-NE is supporting services that require frequency synchronization, the MAC-NE MUST meet the frequency accuracy requirement as defined in [R-DTI].

9.3 Precision Timing Synchronization Services

9.3.1 Market Forces Drive Need for Precise Timing

In recent years, synchronization in access networks has become important because of the evolution from time division multiplex (TDM)-based networks to packet-based networks. In particular, mobile wireless operators are struggling to increase their backhaul capacity to meet the requirements of the newest radio technologies. They are pushing themselves in order to provide greater bandwidth and improved services, which are important to their customers.

Although the introduction of smaller capacity base stations (namely microcell, picocell, and femtocell) permits mobile and broadcast operators to improve wireless service by providing better coverage, it also requires an increase in the number of network connections. Such trends lead to optimization of the MXh infrastructure. Various transmission options for the aggregation networks are being considered and/or utilized. As one of the last-mile access transport technologies, the cable industry's HFC networks are being considered, especially as the bandwidth of DOCSIS-based systems continues to increase.

One critical aspect of mobile base stations and broadcast transmitters is their need to synchronize their radio interfaces. Accurate frequency synchronization among base stations allows user handsets to seamlessly handover between base stations, reduces interference between cells, and optimizes radio bandwidth capacity. Another critical aspect of some base stations is the need for phase or ToD synchronization. In order to be utilized in the wireless market space, the MAC-NE needs to support precision time synchronization.

9.3.2 Mobile Xhaul (MXh) Timing for Small Cells

One of the key drivers for precision time and phase transport over a DOCSIS network is to support operator small-cell deployments. A key initial objective is to support LTE TDD and 5G small cells, which not only requires transporting user and control traffic to the small cell (often with low-latency requirements), but also requires precision phase synchronization of the LTE TDD or 5G small cell RF air interface. This topic has been studied thoroughly by the CableLabs MXh Sync over DOCSIS working group, which has produced an architecture and performance specification [SYNC].

The synchronization path from a precision time source (GPS/UTC) to the MAC-NE will most likely not take the same path as the main traffic paths to the MAC-NE. Figure 35 shows an example synchronization path.

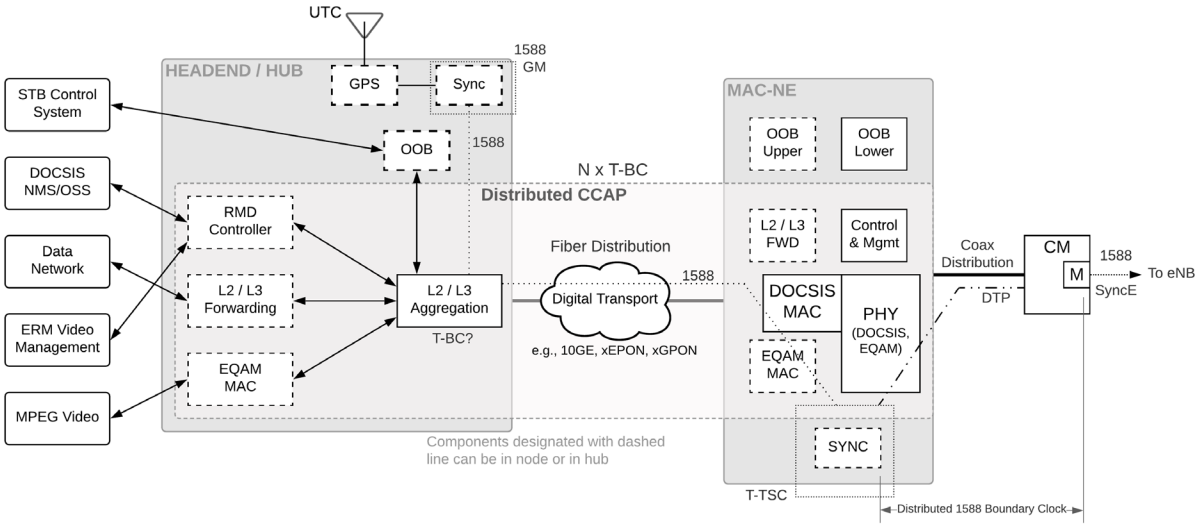

Figure 35 - Mobile Xhaul (MXh) Timing for Small Cells

The protocols most often used to transport precision time and frequency synchronization over L2 networks are [IEEE 1588] and Synchronous Ethernet (Sync-E). Sync-E is good at providing physical-layer frequency synchronization but lacks the capability of transporting precision time/phase, which is required for LTE TDD mode and 5G. Thus, supplementing Sync-E with IEEE 1588 allows for the transport of high-level frequency stability (Sync-E) as well as precision time/phase (IEEE 1588). However, IEEE 1588 requires downstream (DS)/upstream (US) latency symmetry to properly function. Because DOCSIS 4.0 networks are highly asymmetric in terms of DS/US latency, a combination of the DOCSIS 4.0 OFDM carrier stability and DOCSIS Time Protocol (DTP) over the DOCSIS 4.0 link needs to be used to transport high-frequency stability and precision time/phase. The combination of the MAC-NE and the CM form a distributed IEEE 1588 Boundary Clock (BC).

The DTP defined in [MULPIv4.0] is a set of techniques coupled with extensions to the DOCSIS signaling messages. DTP allows the DOCSIS timing and frequency system to be interfaced to external timing protocols such as PTP with high accuracy.

The MAC-NE MUST support DTP for the applications that require precision synchronization.

However, supporting DTP alone is not enough. In order to be placed in a network chain that delivers precise timing, the MAC-NE needs to comply with requirements in [SYNC], which defines the requirements that each DOCSIS element (e.g., the MAC-NE) supports based on the deployment use case. When supporting [SYNC] for small cells, the MAC-NE needs the following:

- DTP in [MULPIv4.0],
- [IEEE 1588] slave functionality,
- Sync-E, and
- low Time Error (TE) between the 1588/Sync-E input and the DOCSIS 4.0/DTP output.

This is not a complete list of requirements; all requirements can be found in [SYNC].

When the MAC-NE is supporting MXh Sync for Small Cell timing services, the MAC-NE MUST meet the requirements for the MAC-NE defined in [SYNC].

10 DATA FORWARDING

10.1 MAC-NE Data Forwarding Assumptions

From a data forwarding perspective, the system composed from MAC-NE and S-Leaf acts like an I-CCAP or a CCAP Core, both of which are Layer 3 forwarders.

The key assumption is that the MAC-NE is only a Layer 2 forwarder, whereas the S-Leaf performs the Layer 3 data forwarding functions. For the MAC-NE, this means that downstream Layer 3 forwarding decisions are based mainly on the destination MAC address of the Protocol Data Unit (PDU), called a frame in this case, and that the source and destination MAC addresses of PDUs that pass through the MAC-NE will remain unmodified. The default gateway for CMs, CPEs, and the MAC-NE IP host is the S-Leaf and not the MAC-NE itself.

Even though it is primarily a Layer 2 forwarder, the MAC-NE maintains some Layer 3 awareness similar to an I-CCAP or CCAP Core device. Thus, functions that require DOCSIS network knowledge, like Source Address Verification (SAV), Subscriber Management Filters, Layer 2 DHCP Relay, etc., are performed by the MAC-NE. The MAC-NE implements all DOCSIS functions of the I-CCAP MAC Domain, such as the RF Interface PHY, DOCSIS lower MAC, DOCSIS scheduler, and the traffic management functions of the DOCSIS upper MAC. The MAC-NE forwards Ethernet MAC Domain frames to and from an S-Leaf port via a Network System Interface (NSI), optionally with an NSI encapsulation such as a VLAN tag.

The S-Leaf is a combined Layer 2 and Layer 3 forwarder. Its main function is to provide a Layer 3 interface that acts as a default gateway for CMs, CPEs, and MAC-NE hosts. The S-Leaf performs Layer 3 forwarding for IP Services traffic to and from the MAC-NE.

Figure 36 shows a typical connectivity of a MAC-NE-based headend:

- Multiple MAC-NEs are connected to the same S-Leaf
- Usually the entire MAC-NE population connected to the same S-Leaf supports the same set of services, thus:
    - S-Leaf has a separate VLAN configured for each service (Management, IP Services, Video, etc.).
    - Each MAC-NE port (Layer 2 interface) on the S-Leaf is member of all the service VLANs required on the MAC-NE. Typically, Management VLAN is untagged while all the other service VLANs are tagged.
    - The S-Leaf has a virtual Layer 3 interface configured for each VLAN that acts as a default gateway for the devices in that particular VLAN.

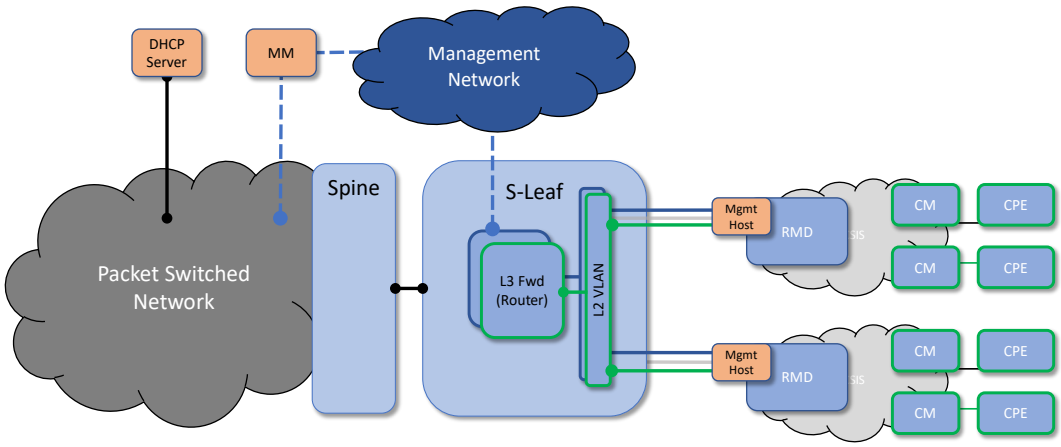

Figure 36 - Typical S-Leaf Connectivity

S-Leaf Typical Layer 2 Forwarding:

- The S-Leaf learns the location of the hosts and builds a MAC address table for each VLAN.
- Frames are forwarded based on destination MAC address only:
  - Unicast frames with a known destination MAC address (learned previously) are forwarded only to the respective port.
  - Unicast frames with an unknown destination MAC address, broadcast and multicast frames are flooded to all VLAN member ports, except the port the frames were received on.

Some S-Leaf switches support additional features that might be deployed for additional security or for optimized forwarding, like:

- IGMP/MLD snooping for optimized multicast forwarding
- DHCP snooping/learning for IP address - MAC address binding verification
- Private VLANs
- Layer 2 forwarding restrictions

Please note that a typical S-Leaf might forward traffic at Layer 2 between all the ports members of a VLAN, unless some of the additional features mentioned are deployed. As a result, devices (CMs, CPEs) within a VLAN could communicate to each other at Layer 2 without being routed through the S-Leaf Layer 3 virtual interface. Either the MAC-NE or S-Leaf could implement filtering or Layer 2 port isolation to prevent Layer 2 communication between devices.

An S-Leaf is expected to prevent forwarding at Layer 2 of IP Services traffic between ports of an IP Services VLAN. This includes non-unicast IP address resolution (ARP, IPv6 Neighbor Discovery) as well as unicast IP Services traffic between devices (CMs and CPEs). In case the S-Leaf does not provide such functionality, the MAC-NE has optional requirements to do so.

The S-Leaf might also perform other tasks traditionally implemented on Layer 3 forwarders (a.k.a., routers), as follows:

- Routing protocols (unicast and multicast),
- ARP/NDP,
- IGMP/MLD,
- Layer 3 DHCP relay, and
- interface ACLs.

The full S-Leaf configuration (VLANs, Addressing, etc.) for I-CCAP operation is outside the scope of FMA standardization, but certain required features (e.g., local proxy ARP) are mentioned in this specification.

10.2 MAC-NE Internal Data Forwarding Model

This section subdivides MAC-NE data packet processing into a number of logical blocks for purposes of specifying forwarding behavior. It does not necessarily reflect the operation of any particular implementation.

10.2.1 High-Level Model

In the downstream direction, the model (Figure 37) calls for a Layer 2 Dispatch block to demultiplex downstream traffic to a number of Forwarding Function blocks based on packet contents that include the packet's VLAN tag (if any), Layer 2 destination MAC address, Layer 3 destination IP address and protocol, and for some functions, its L2TPv3 session ID.

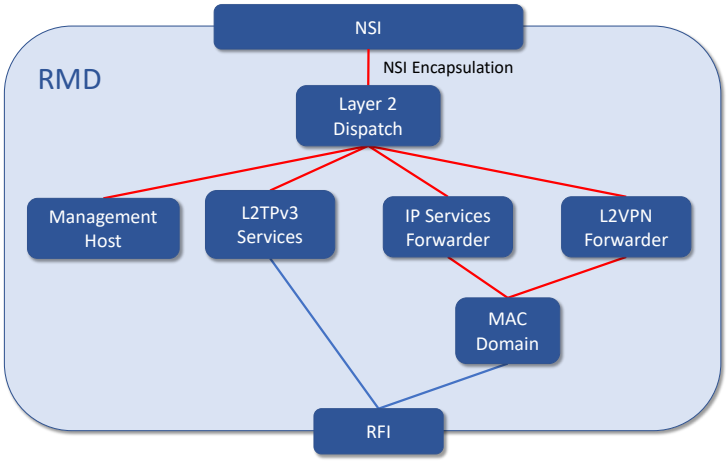

Figure 37 - High-Level Data Forwarding Model

NSI (Network Side Interface): the network-side interface between the MAC-NE and the S-Leaf.

MAC Domain: a MAC-NE subcomponent that provides data forwarding services to a set of DOCSIS downstream and upstream channels.

Layer 2 Dispatch: an internal component that dispatches downstream Layer 2 NSI traffic to the proper internal forwarding function, stripping any NSI encapsulation and forwarding with relevant information based on mapping that encapsulation (e.g., a logical L2VPN or IP Address Space identifier). In the upstream direction, this component conceptually adds the Layer 2 NSI encapsulation based on conceptual information provided with a packet from the forwarding function.

Management Host: an internal MAC-NE subcomponent that terminates management protocol sessions used to manage and configure the MAC-NE itself. The Management Host exposes management protocols (GCP, SSH, TELNET, SNMP, REST, etc.) via its own IP stack.

IP Services Forwarder: an internal MAC-NE subcomponent that performs forwarding and processing of data packets between the NSI and MAC Domain for IP protocol packets.

L2VPN Forwarder: an internal MAC-NE subcomponent that performs forwarding and processing of data packets between the NSI and MAC Domain for only L2VPN/BSOD services.

L2TPV3 Services: the block that terminates static or dynamic L2TPV3 sessions and forwards to/from RFI channels. Its functionality is to translate between the DEPI/UEPI protocols and the RF interface for MPEG-TS video, OOB, and NDF/NDR as defined in [R-DEPI], [R-UEPI], [R-OOB], and [R-PHY].

10.2.2 NSI Encapsulation

The term Network Side Interface (NSI) refers to an interface that connects directly to the S-Leaf and forwards Layer 2 Ethernet frames. In Phase 1 of FMA, an NSI is either a single physical Ethernet port or a link aggregation group (LAG) of multiple Ethernet ports. Operation with more than one NSI is vendor-specific and beyond the scope of this specification.

An NSI "subinterface" is a subset of Ethernet traffic on an NSI that implements a virtual LAN identified by an explicitly tagged encapsulation consisting of;

- a single [IEEE 802.1Q] VLAN tag or;
- a dual [IEEE 802.1Q] VLAN tag.

In FMA Phase 1, only [IEEE 802.1Q] single and dual VLAN tags are supported for explicit NSI encapsulation of subinterfaces.

The NSI encapsulation of an Ethernet frame determines how it is processed by an MAC-NE forwarding function. Operators can implement VLANs on the S-Leaf infrastructure for traffic isolation and forwarding. Each NSI encapsulation on the MAC-NE usually corresponds to an S-Leaf VLAN.

There may be no NSI encapsulation at all, or there may be services that share the same NSI encapsulation, in which case the L2 Dispatch function recognizes the service and forwards based on the Ethernet frame's L2/L3/L4 fields only.

In FMA Phase 1, the NSI encapsulation may be a VLAN tag or a tag pair that distinguishes the service. See the description of a service for how the encapsulation is configured in the MAC-NE and for particular downstream Ethernet frame field matching requirements.

The L2 Dispatch adds NSI encapsulation (if any) in the upstream direction. Upstream packets from a forwarding service use the same NSI encapsulation as for the downstream service.

For future consideration are NSI encapsulations using other Layer 2 Ethernet tunneling approaches such as MPLS and VXLAN. For brevity, this specification uses the term "encapsulation" to mean either inserting a VLAN tag or tag pair or encapsulating in a tunneled encapsulation. One of the choices for NSI encapsulation is "untagged."

The MAC-NE MUST support untagged NSI encapsulation.

The MAC-NE MUST support single [IEEE 802.1Q] VLAN tag NSI encapsulation.

The MAC-NE MUST support dual [IEEE 802.1Q] VLAN tag NSI encapsulation.

The MAC-NE MUST support standard single VLAN tag with TPID 0x8100 NSI encapsulation.

The MAC-NE MUST support double tagging operations for the purposes of NSI encapsulation using the following:

- Both inner and outer VLAN tags using TPID 0x8100
- The standard double tagging as per [IEEE 802.1Q] with an inner C-TAG with TPID 0x8100 and an outer S-TAG with TPID 0x88a8
- The non-standard but commonly used double tagging with an inner C-TAG with TPID 0x8100 and an outer S-TAG with TPID 0x9100
- The non-standard but commonly used double tagging with an inner C-TAG with TPID 0x8100 and an outer S-TAG with TPID 0x9200.

The MAC-NE MUST allow configuration of a VLAN for a dedicated control path to the MAC-NE (e.g., dedicate a particular VLAN to manage traffic to the MAC-NE). The MAC-NE MUST use a VLAN that does not require pre-configuration prior to installation. The MAC-NE MUST learn the dedicated management VLAN during onboarding (i.e., from LLDP or through an equivalent mechanism).

With the exception of L2VPN service, NSI encapsulation for the various services is configured by the MAC Manager.

The MAC-NE MUST encapsulate upstream traffic from an RFI-side host on the NSI with the same encapsulation with which it receives downstream traffic for that host from the NSI.

For any service, the MAC-NE MAY implement a vendor-specific encapsulation with configuration in a vendor-specific manner. Vendor-specific encapsulation for a service is intended for use only when the MAC Manager omits configuring the encapsulation. A MAC-NE MUST reject committing to operating with a conflicting vendor-specific configuration and MAC Manager configuration of the NSI encapsulation for any service. The MAC-NE would do so by rejecting the most recent vendor-specific configuration or MAC Manager configuration for a service encapsulation that conflicts with the existing configuration.

10.2.3 Forwarding Functions

Figure 38 elaborates on the high-level model by identifying separate forwarding functions. The requirements for each forwarding function are provided in existing DOCSIS specifications.

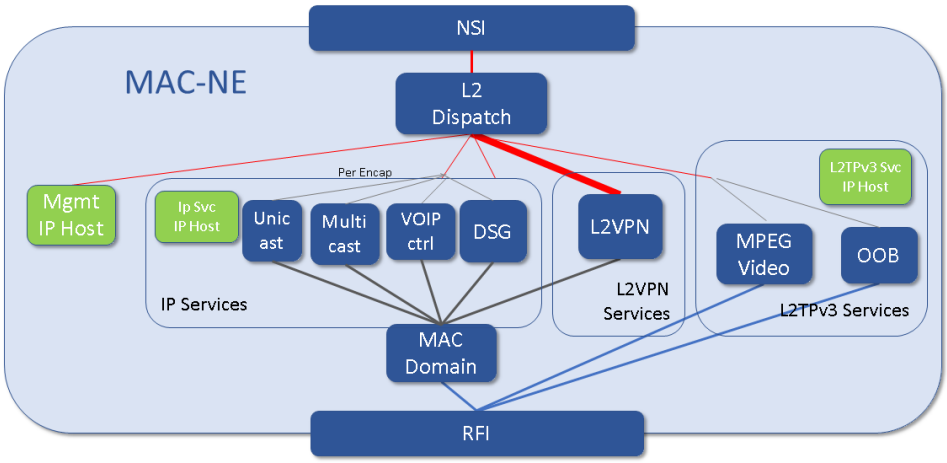

Figure 38 - Detailed Data Forwarding Model

Downstream traffic from L2 Dispatch to forwarding functions is considered to be internally marked with an internal service identifier mapped from the NSI encapsulation.

IP Services has multiple forwarding functions.

- Unicast Forwarding refers to the filtering and forwarding of unicast IP packets to CMs and CPEs. It can require the IP Services Host to implement a TFTP proxy on behalf of CM/CPE hosts or to exchange DHCPLEASEQUERY messages with the external DHCP server. Unicast Forwarding is also considered to include ARP and IPv6 Neighbor Discovery (ND) non-unicast traffic required for the resolution of unicast IP addresses to unicast MAC addresses. Overall forwarding requirements for IP services are specified in Section 10.4.
- Multicast Forwarding refers to forwarding downstream and upstream IP multicast packets. It can require the IP Services Host to serve as an IP Multicast Proxy. Forwarding requirements are specified in Section 10.4.6

- The VOIP function of IP services refers to control plane interaction with external PacketCable protocol entities, using the IP Services Host. VOIP control functions are specified in [PKT-DQOS1.5], [PKT-MM], and [FMA-PAI].
- The DSG service of IP services refers to the conversion of downstream DSG traffic (usually IPv4 or IPv6 multicasts) to DSG tunnels. MAC-NE forwarding requirements are specified in [DSG], with IPv6 translation as specified in Section 7.2.

The interaction between IP Services and MAC Domain is not specified in this document but is assumed to be sufficient to meet the requirements of [MULPIv4.0].

An "IP Services NSI encapsulation" refers to the encapsulation of a packet for it to be considered eligible for IP Services forwarding downstream or the encapsulation to be applied by the MAC-NE when sending packets upstream from an IP Services function. It is the encapsulation of the NSI interface or subinterface configured for a "cable bundle" or "cable sub-bundle" for a DOCSIS MAC Domain, as described in Section 10.4.3.

The L2TPv3 Services function refers to the termination of L2TPv3 pseudowires on an NSI interface and the forwarding of traffic between RFI channels and those pseudowires. The L2TPv3 Services function is considered to have a separate function for each pseudowire type.

- MPEG Video services refers to the termination and forwarding of MPEG packets on statically or dynamically configured L2TPv3 pseudowires, as specified in [R-DEPI] and [R-PHY].
- OOB services refers to termination and forwarding of static or dynamic L2TPv3 pseudowires for all protocols in the R-OOB specification, i.e., SCTE 55-1, SCTE 55-2, NDR, and NDF, as specified in [R-OOB] and [R-UEPI]. It includes the function of Tone Generation.

10.2.4 IP Host Interfaces

The MAC-NE sends and receives IP over Ethernet traffic as an IP host from an "IP Host interface". An IP Host interface has the following attributes:

- an administratively assigned "ip-interface-name" string to uniquely identify the IP interface on the MAC-NE;
- an association to a single NSI, i.e., Ethernet port or LAG group, identified with an "eth-interface-name";
- an optional NSI encapsulation (e.g., VLAN tag); and

one IPv4 and/or Ipv6 address, together with the prefix length of the IP subnet of each IP host address. An "eth-interface-name" uniquely identifies an interface (physical or logical) capable of sending Ethernet frames. The eth-interface-name is determined by the MAC-NE vendor. For example, it can be the concatenation of an Ethernet interface type and a zero-based index (e.g., "eth0" or "bond0").

An NSI may have multiple IP host interfaces, and those IP Host interfaces can have the same or different NSI encapsulations. IP host interfaces may be automatically provisioned with DHCP, statically provisioned with non-volatile vendor-specific configuration, or dynamically provisioned with MMI. The MAC-NE MUST permit the creation of an IP host Interface with an IP prefix that overlaps the IP prefix of an existing IP host interface when the existing IP host interface has a different NSI encapsulation. The MAC-NE MUST reject the creation of an IP Host interface with an IP prefix that overlaps the IP prefix of an existing IP host interface when the existing IP host interface has the same NSI encapsulation.

An IP host Interface does not have a separate "admin-status" attribute. An IP host Interface is considered to be operationally enabled when its associated Ethernet interface has an admin-status of "up".

IP Host Interfaces are designated for use by one or more data forwarding services.

Figure 39 depicts potential IP hosts in the MAC-NE.

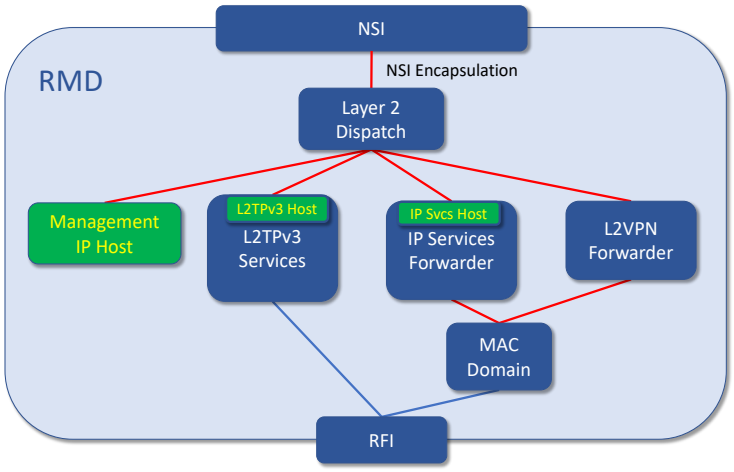

Figure 39 - IP Host Interfaces

A MAC-NE MUST support an IP host Interface designated as the "Management IP host". The Management IP host is automatically provisioned during onboarding initialization with the S-Leaf, with an IP address learned from DHCP or SLAAC and an optional NSI encapsulation learned from 802.1x, LLDP, or DHCP (in that order).

One IP host Interface can be configured by the MAC Manager as the "Default L2TPv3 Service host" to terminate L2TPv3 sessions as specified in [R-PHY], [R-DEPI], and [R-OOB]. These functions include the following:

- receiving from L2TPV3 pseudowires for MPEG Video;
- joining IP multicast sessions for MPEG Video multicasts; and
- receiving from and sending on L2TPV3 pseudowires for OOB pseudowire types, which include SCTE 55-1, SCTE 55-2, NDF, or NDR as described in [R-OOB].

A MAC-NE MAY implement additional (non-default) IP host Interfaces for one or more L2TPv3 pseudowires. The MAC Manager can configure which IP host terminates which pseudowires. The MAC Manager also configures the MAC-NE with all IP addresses in the RPD GCP configuration of L2TPv3 static pseudowires.

If no IP host interface is designated as the "Default L2TPv3 Services Host" interface, the MAC-NE performs these functions on the "Management IP Host" interface.

The MAC Manager may configure a "Default IP Services host" to implement certain IP host functions of a CMTS:

- sending/receiving traffic as a "TFTP proxy",
- sending/receiving traffic as an "IP Multicast Proxy",
- sending/receiving PacketCable Aggregator Interface (PAI) traffic,
- sending/receiving DHCPLEASEQUERY messages
- joining a DSG IPv4 or IPv6 multicast session

A MAC-NE MUST support configuration of at least three (3) IP host interfaces. This requirement is to support use of NSI encapsulation to isolate traffic for

- Management IP Host,
- Default L2TPv3 Services Host, and
- Default IP Services Host.

If the "Default IP Services Host" is not configured for a function, the MAC-NE MUST perform that function using the Management IP Host. If the "Default L2TPv3 Services Host" is not configured for a function, the MAC-NE MUST perform that function using the Management IP Host.

If a MAC-NE implements the Lawful Intercept function, the MAC-NE MUST support configuration of an additional "Lawful Intercept IP Host" for the LI function. If a MAC-NE implements the Lawful Intercept function, and the LI IP Host is not configured, the MAC-NE MUST use the Management IP Host for the LI function. Note that LI IP Host can be configured with the same NSI interface and IP address as any other IP Host, e.g., the Default IP Services Host. The MAC-NE does not advertise any capabilities for IP host Interfaces. The number of IP host interfaces supported by an MAC-NE and any restrictions on their use are assumed to be known *a priori* by the operator when configuring the MAC Manager.

The assignment of FQDNs to address network elements in the FMA system allows for easier management and can speed familiarization of the network for new personnel. Naming conventions for individual network elements, IP host(s) and the determination to use sub-domains or a single domain name across all internal infrastructure is determined by the operator and follows existing practices and/or the best practices promulgated by the industry.

10.3 MAC-NE Data Forwarding

The word "accept" in this section is used to describe frames that can be dispatched by the L2 Dispatch block to one of the four forwarding blocks in the high-level model. As a general guideline, the Layer 2 dispatching of a frame can be performed based on the NSI encapsulation and destination MAC address.

A known NSI encapsulation for a MAC-NE is one that is configured on the MAC-NE or learned from CM configuration for a L2VPN/BSOD service. Untagged NSI encapsulation is, by default, known without being configured because initialization communication can occur over the untagged NSI encapsulation, which does not require specific configuration.

The MAC-NE will maintain a list of known MAC addresses relevant to each service or known NSI encapsulation. Depending on the services configured on an MAC-NE, the known MAC address list could include the MAC addresses of CMs, CPEs, and any of MAC-NE hosts as well as the addresses of the multicast groups of which the device is a member. For instance, if an MAC-NE is configured with a specific NSI encapsulation for the HSD service, then the known MAC address list for that HSD service/NSI encapsulation would consist of the MAC addresses of all the CMs registered for the HSD service and the CPEs behind them.

The MAC-NE MUST accept any downstream NSI frames, whether untagged or tagged with a known NSI encapsulation, that include a known MAC address in the destination MAC address field for that NSI encapsulation.

The MAC-NE MUST reject any downstream NSI frames tagged with an unknown NSI encapsulation.

The MAC-NE MUST reject any downstream NSI frames tagged with a known NSI encapsulation but destined for an unknown MAC address for that service or NSI encapsulation.

The MAC-NE MUST remove the NSI encapsulation for all downstream frames before forwarding them to the RFI.

The MAC-NE MUST add the appropriate NSI encapsulation for upstream frames before sending them to the NSI.

For Management traffic, the MAC-NE MUST forward traffic between the NSI and the local MAC-NE IP Hosts.

For L2TPv3 Services traffic, the MAC-NE MUST forward traffic between the NSI and the RFI per [R-PHY], [R-DEPI].

The MAC-NE MUST forward all IP Services traffic between the NSI and RFI as specified in [MULPIv4.0].

The MAC-NE MUST forward all L2VPN traffic between the NSI and the RFI as specified in [L2VPN].

10.3.1 Data Forwarding Flowcharts

The following flowcharts (Figure 40 and Figure 41) depict the overall operation of data forwarding at a high level.

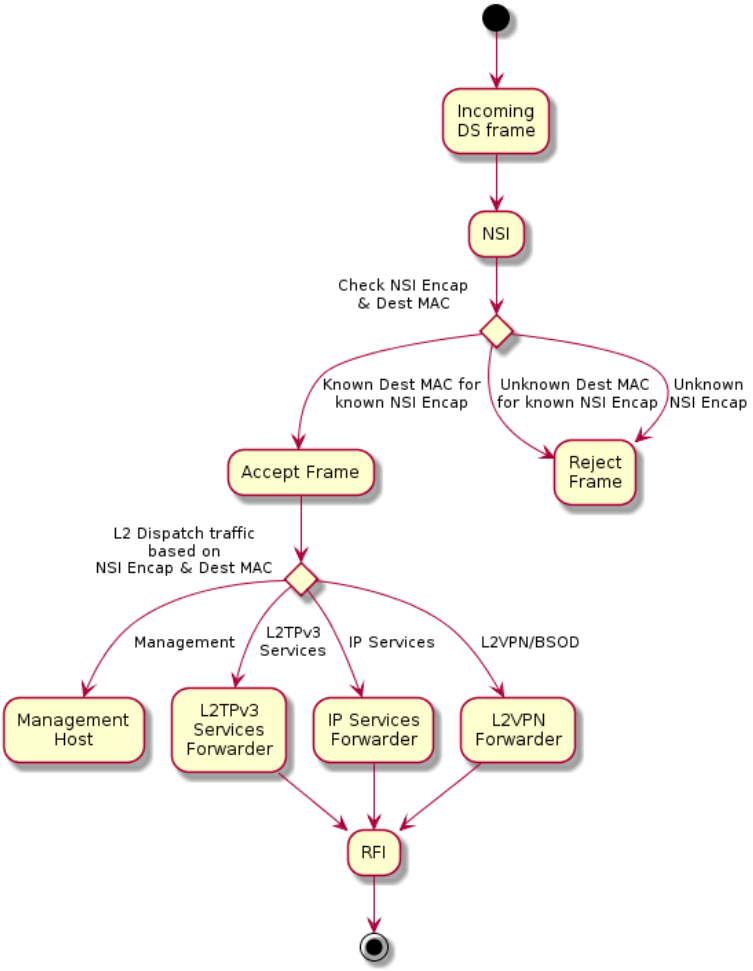

Figure 40 - Downstream Data Passing Flowchart

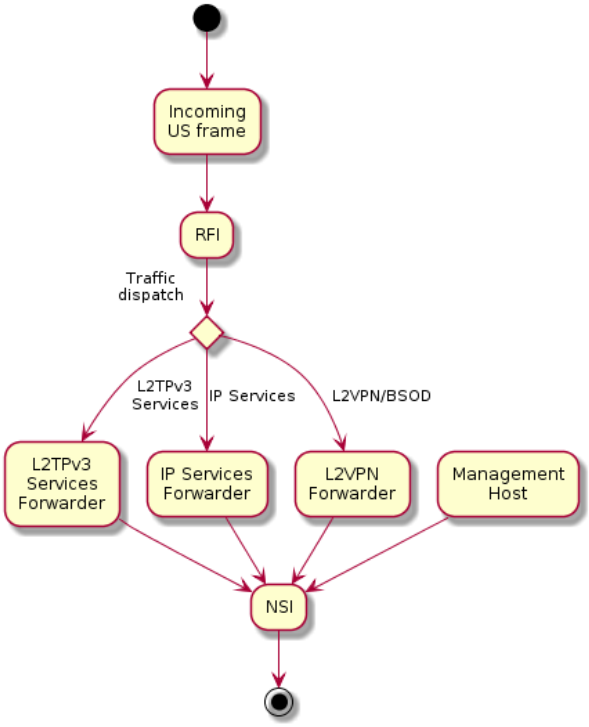

Figure 41 - Upstream Data Passing Flowchart

10.4 IP Services

IP Services refers to the forwarding of IP traffic to subscriber CMs and CPEs based on Layer 3 IP routing. The S-Leaf and MAC-NE share specific IP Services forwarding responsibilities.

For downstream IP Services traffic, the S-Leaf routes a packet's destination IP address to an L3 interface VLAN, resolves the destination IP to a destination MAC address, builds an L2 Ethernet frame, and forwards that frame to one or more MAC-NEs on an NSI with an NSI encapsulation.

At the MAC-NE, a downstream NSI encapsulation that matches a configured "cable bundle" of MAC Domains (Section 10.4.3) or "cable sub-bundle" (Section 10.4.3.1) enables IP Services forwarding of the packet, although other services may use the same NSI encapsulation.

The MAC-NE learns MAC addresses and authorized IP address assignments of the CMs and CPEs on its DOCSIS MAC Domains (Section 10.4.1) and builds an IP Services "Forwarding Database" (IP-FDB). When the S-Leaf issues an ARP or ND query for a learned CM/CPE IP address in the IP-FDB, the MAC-NE responds directly to the S-Leaf with the IP-FDB's MAC address for the CM/CPE rather than forwarding the query downstream (Section 10.4.2).

In the upstream direction, all DOCSIS MAC Domain traffic not otherwise forwarded for L2VPN services is considered to be forwarded by IP Services. The MAC-NE enforces Source Address Verification (SAV) (Section 11.2) based on its IP-FDB.

10.4.1 MAC and IP Address Learning

A MAC-NE learns the association of MAC addresses to the IP addresses of CMs and CPEs through DHCP snooping and/or ARP/ND snooping of CPE IP addresses authorized by CM configuration file TLVs. The MAC-NE uses these associations for purposes such as forwarding, SAV, Subscriber Management filters, and cable sub-bundle classification criteria.

The MAC-NE MUST maintain an IP Services Forwarding Database (IP-FDB) table of known (learned) CM and CPE MAC addresses and corresponding IP address mappings. The forwarding table includes the NSI encapsulation of the bundle or sub-bundle containing the IP address.

IP services CM and CPE forwarding does not require the MAC-NE to learn MAC addresses or associations of MAC addresses with IP addresses on the NSI ports. The CM and CPE use normal means (like ARP and ND) to learn MAC addresses of the IP peers reachable through NSI ports.

10.4.2 Address Resolution

The requirements in this section are intended to prevent an unauthorized host from receiving downstream traffic simply by that host responding with its MAC address to a downstream ARP broadcast for IPv4 addresses or IPv6 multicast ND Neighbor Solicitation multicast query.

The MAC-NE MUST intercept downstream ARP and ND non-unicasts with IP Services NSI encapsulation for possibly responding in proxy or forwarding downstream.

When the MAC-NE intercepts a downstream ARP or ND non-unicast query and an IP-FDB entry exists for the query's NSI encapsulation and target IP address and the IP-FDB entry's CM is online, the MAC-NE MUST respond in proxy for the target with the IP-FDB entry's MAC address.

When the MAC-NE intercepts a downstream ARP or ND non-unicast query and an IP-FDB entry does not exist for the query's NSI encapsulation and target IP address for an online CM, the MAC-NE MUST forward the downstream ARP or ND non-unicast query downstream to reach all CMs on the MAC Domains in a cable bundle configured for the query's NSI encapsulation.

10.4.3 Cable Bundles

A cable bundle is a set of MAC Domains with CMs and CPEs that share the same set of IP subnets, as depicted in Figure 42.

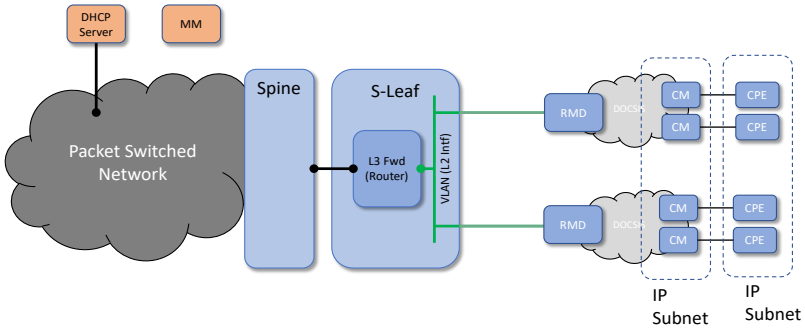

Figure 42 - Cable Bundles

For routing IP packets through MAC-NEs, the S-Leaf implements a Layer 2 VLAN shared across multiple S-Leaf Ethernet ports, each connected to a different MAC-NE. Each MAC-NE implements separate MAC Domains consisting of downstream and upstream RF channels. CMs and CPEs operate in a single MAC Domain on each MAC-NE at a time. The S-Leaf implements a Layer 3 router interface on the shared VLAN that defines the gateway IP addresses for those IP subnets. Usually, a cable operator DHCP server assigns IP addresses to CMs in one set of IP subnets and to CPEs in a different set of IP subnets. Each MAC Domain is configured to belong to a single cable bundle.

The MAC Manager is expected to be configured in a vendor-specific manner. Each MAC Domain of a MAC-NE can be a member of exactly one such cable bundle. Multiple MAC Domains in an MAC-NE can be members of the same or different cable bundles. Because a CM registers on a single MAC Domain, a CM belongs to a single cable bundle. A CPE is considered to belong to the same MAC Domain as the CM through which it attaches.

A cable bundle configuration references a particular NSI interface or subinterface. The NSI encapsulation of the cable bundle is thus configured by its interface or subinterface. Multiple cable bundles can reference the same NSI interface or subinterface. Downstream traffic on that interface or subinterface is forwarded only to the MAC Domains belonging to the cable bundle.

The MAC-NE MUST encapsulate upstream IP Services traffic with the NSI encapsulation for the packet's identified bundle or sub-bundle.

10.4.3.1 Cable Sub-Bundles

A cable sub-bundle (Figure 43) represents a set of traffic within a cable bundle. It is considered to be a specialization of an existing cable bundle that inherits and optionally overrides the non-key attributes of its parent cable bundle.

All CMs and CPE hosts within a cable bundle are characterized by the same attributes (IP subnets, NSI encapsulation, DHCP policy, etc.), unless sub-bundles are defined to override them. Sub-bundles are usually characterized by a unique set of attributes that are different from the parent cable bundle.

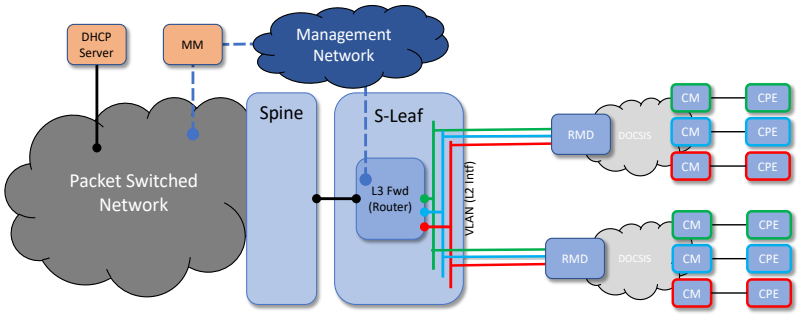

Figure 43 - Cable Sub-Bundle

A cable sub-bundle is configured with a set of criteria for selecting the IP traffic considered to belong to the sub-bundle; for example, traffic from IP addresses for an eMTA device type could be selected for a cable sub-bundle. Any IP Services traffic not selected for a sub-bundle is considered as selected for the top-level cable bundle.

Each sub-bundle is configured with a particular NSI encapsulation (e.g., single or double tag value) that is unique on the MAC-NE. It is expected that all S-Leaf ports and MAC-NEs will use the same NSI encapsulation for a sub-bundle on different MAC-NEs. Although it is physically possible to represent the same cable sub-bundle on different S-Leaf ports with different NSI encapsulations, such complexity is discouraged and there is no requirement for MAC Managers to support vendor-specific configuration to do this.

The MAC-NE encapsulates upstream traffic identified for the sub-bundle with the NSI encapsulation for that sub-bundle. In many cases, the MAC-NE can identify the sub-bundle of an upstream DHCP request and forwards the upstream DHCP with the NSI encapsulation of the sub-bundle. For example, the MAC-NE might match all upstream traffic from a particular host type, e.g., CMs, external CPEs, or an embedded CPE type such as eMTAs.

In the downstream direction, the NSI encapsulation of a packet can identify a particular sub-bundle for IP Services. For downstream IP Services traffic, a destination MAC address can be learned or configured on one or more sub-bundles. The MAC-NE MUST reject downstream unicast IP services traffic that has a sub-bundle NSI encapsulation when that NSI encapsulation does not match one of the sub-bundles on which the destination MAC address was learned or configured.

A valid MAC-NE configuration can have, at most, only one sub-bundle with the untagged NSI encapsulation. A valid MAC-NE configuration has each cable bundle and sub-bundle with a unique NSI encapsulation.

The Cable Helper list for a cable bundle or sub-bundle is the set of DHCP Server IP addresses authorized to assign CM/CPE host IP addresses to hosts in the cable bundle. The MAC-NE uses this list to authorize the source IP address of downstream DHCP traffic. The MAC-NE uses this list as the destination IP address of DHCP Leasequery requests for upstream IP Source Address Validation.

Cable bundle and sub-bundle helper addresses and NSI encapsulations are expected to be configured the same on the S-Leaf and MAC Manager, but there is no provision in the current FMA standard to check or enforce this compatibility. FMA operation is not defined with incompatible cable bundle or sub-bundle configuration on the S-Leaf and MAC Manager.

The MAC-NE is expected to isolate downstream non-unicast IP traffic to different sub-bundles with different Downstream Service IDs (DSIDs) [MULPIv4.0].

10.4.3.2 Cable Sub-bundle Identification

A cable sub-bundle is configured with a set of identification criteria organized as "terms" joined by "OR", i.e., the sub-bundle is selected when any of the terms are matched. Each of those terms consists of one or more "rules" joined by "AND", i.e., the term is matched only when all of its rules are matched. When more than one "OR" term applies to an upstream packet, the particular sub-bundle for forwarding it is resolved by the term with the highest "priority". When two terms with the same priority are matched, the particular term (and hence sub-bundle) selected by the MAC-NE is not specified.

Sub-bundle identification criteria are matched at the time a CM or CPE IP address becomes known to the MAC-NE, i.e., when snooping a downstream DHCP assignment or with static configuration.

The cable bundle identification rules are as follows.

- *DOCSIS Device Type*: assigns an IP host to a sub-bundle by matching the DOCSIS device type (if any) signaled in the IPv4 option 43 (Vendor Specific) suboption 2 (Device Type) or in the IPv6 option 17 (OPTION_VENDOR_OPTS) suboption 2 (CL_OPTION_DEVICE_TYPE). The DOCSIS Device Type is an enumeration selected from one of the following [eDOCSIS].
    - cpe: matching an external CPE device that had no DOCSIS Device Type option
    - docsis: matching any device that signaled a DOCSIS Device Type option
    - ecm: matching a DOCSIS embedded cable modem
    - emta: matching a DOCSIS embedded Multimedia Terminal Adapter
    - estb: matching a DOCSIS embedded set-top box
    - eps: matching a DOCSIS embedded portal services device
    - edva: matching a DOCSIS embedded digital voice adapter
    - etea: matching a DOCSIS TDM emulation agent
    - esg: matching a DOCSIS SMA (Security/Monitoring/Automation) Gateway
    - sesg: matching a EuroDOCSIS Standalone Enterprise SIP Gateway
- *IP Subnet*: assigns the sub-bundle for a CM or CPE based on its IP address matching a particular IPv4 or IPv6 prefix.
- *DHCP Field List*: selects an IP host for a sub-bundle based on its DHCP assignment matching an arbitrary set of fields in DHCPv4 or DHCPv6 options or suboptions. An example would be the value of a Client Identifier option. All fields in the list are required to match in order to match the "DHCP Field List" rule. Each DHCP field is identified by IP protocol (IPv4 or IPv6), option number, suboption number, field offset, field octet length, and either ASCII text or hexadecimal binary field value. Note that the DHCP Field List rule can be encoded to match any current or future standardized DOCSIS Device Type field.
- *CM Subnet-Based*: selects a sub-bundle for a CPE based on the CM that reaches it having an IP address within an IP subnet specified in the rule. This is intended for an operator to support multiple subscriber Internet Service Providers (ISPs) in a manner similar to existing I-CCAP cable sub-bundles. Section 10.4.3.3 describes an example of this rule criteria.

10.4.3.3 Example of CM Subnet-based Rule Matching

This section provides an example of using CM Subnet-Based rules for identifying the cable sub-bundle of a CPE in the deployment with multiple subscriber ISPs. An example with three ISPs "Green", "Blue", and "Red" is depicted

in Figure 44, where IP subnets and VLANs for the different ISPs are identified with their corresponding colors. The S-leaf implements an L3 interface host address on each subscriber VLAN which serves as the default gateway for each ISP's subscribers and the giaddr for forwarding subscriber CPE DHCP to an DHCP server specific for each ISP. The DHCP server for each ISP assigns IP addresses to the CPEs of its subscribers. The border of a CM or CPE indicates the bundle or sub-bundle of its IP address.

The operator assigns the IP addresses of all CMs. The peach color border indicates IP addresses that are in the default top-level bundle. In this example, the top-level cable bundle is assigned to a LAG sub-interface, so its traffic is also tagged with a VLAN tag for that sub-interface. Table 7 summarizes the configuration of the MAC-NE bundle and sub-bundles for this example:

Table 7 - Bundle Configuration Example

Cable Bundle	Designation	Bundle references NSI	Referenced NSI sub-interface VLAN Tag	Operator "Peach" DHCP server assigns CM IP addresses within ranges
1	"Peach"	Subinterface lag0.2	2	10.10.10.0/24 Peach "A" 10.10.20.0/24 Peach "B" 10.10.30.0/24 Peach "C"

Table 8 shows the configuration of sub-bundles in the example.

Table 8 - Sub-bundle Configuration Example

Cable Sub-Bundle	Designation	Cable sub-bundle	Sub-Bundle References NSI	NSI Sub-interface VLAN Tag	ISP DHCP Server assigns CPE IP addresses in range
1	"Green"	11	lag0.11	1011	24.24.11.0/24
	"Blue"	12	lag0.12	1012	24.24.12.0/24
	"Red"	13	lag0.13	1013	24.24.13.0/24

Table 9 shows how the cable bundle is configured with the following sub-bundle identification rules:

Table 9 - Sub-bundle Identification Rules Example

Rule Index	Rule Type	Rule	Selected Sub-Bundle
1	"CM Subnet"	CM IP within 10.10.10.0/24	11
2	"CM Subnet"	CM IP within 10.10.20.0/24	12
3	"CM Subnet"	CM IP within 10.10.30.0/24	13

Operation of CM subnet-based sub-bundling proceeds as follows. Please refer to Figure 44.

1. When the MAC-NE receives an initial upstream DHCP request from a CM, it attempts to find the highest priority sub-bundle identification rule in the mac-domain's cable bundle that matches the packet. In the case of a CM DHCP request, there are no such matching rules, so the MAC-NE assigns the CM's MAC address to the top-level "peach" cable bundle and forwards its DHCP request on the cable bundle's sub-interface lag0.2 on VLAN 2.
2. The cable operator's DHCP server has static reservations for each ISP subscriber's CM and assigns the IP address of a CM for an ISP's subscriber into a particular IP subnet for that ISP. That is, the DHCP server assigns Green ISP subscriber CMs to IPv4 subnet 10.10.10.0/24 (peach "A") Blue ISP subscriber CMs to IPv4 subnet 10.10.20.0/24 (peach "B"), and Red subscriber CMs to IPv4 subnet 10.10.30.0/24 (peach "C"). All CM IP addresses are assigned to one of these three peach subnets.

3. When the MAC-NE receives an initial upstream DHCP request from a CPE, it finds that the highest priority sub-bundle identification rule for forwarding the DHCP packet. In this case, it will be one of the "CM Subnet" based rules.
4. The MAC-NE assigns the CPE MAC address to the identified cable sub-bundle. The MAC-NE will forward all subsequent upstream traffic from the CPE's MAC address on the NSI sub-interface for the sub-bundle, including the initial upstream DHCP packets themselves. For example, an upstream DHCP Discover from a CPE attached to a CM with IP address 10.10.30.101 (in "peach C" subnet) would match the "CM Subnet-based" rule for sub-bundle 13 ("Red"), so the MAC-NE forwards the upstream DHCP from the CPE with VLAN tag 1013.
5. The S-Leaf relays the CPE's DHCP packet on VLAN 1013 to the "Red" ISP's DHCP server, e.g., by configuring it as the DHCP "helper" for a S-leaf router interface on VLAN 1013. The Red ISP DHCP server assigns a publicly reachable IP address for the CPE (e.g., 24.24.10.17) and responds downstream toward the S-leaf on the same VLAN 1013. The Red DHCP server provides a next-hop gateway to the Internet (e.g., 24.24.10.1) through a "Red" internet router (not depicted) that is also attached to VLAN 1013.
6. The S-Leaf forwards the downstream DHCP assignment on VLAN 1013 to the MAC-NE.
7. The MAC-NE recognizes that the "Red" DHCP server is authorized as a "cable-helper" for sub-bundle 13 and forwards the downstream DHCP traffic to the CPE. If SAV is enabled, the MAC-NE notes the CPE IP address and marks it as an SAV authorized source IP address for cable bundle 13.
8. The MAC-NE considers all subsequent upstream IP services traffic from the CPE source MAC address as from sub-bundle 13 and verifies the source IP address of the traffic against the authorized SAV entries for only that bundle.
9. The S-Leaf forwards all subsequent internet traffic to and from the "Red" subscriber through that ISP's internet router on VLAN 1013.

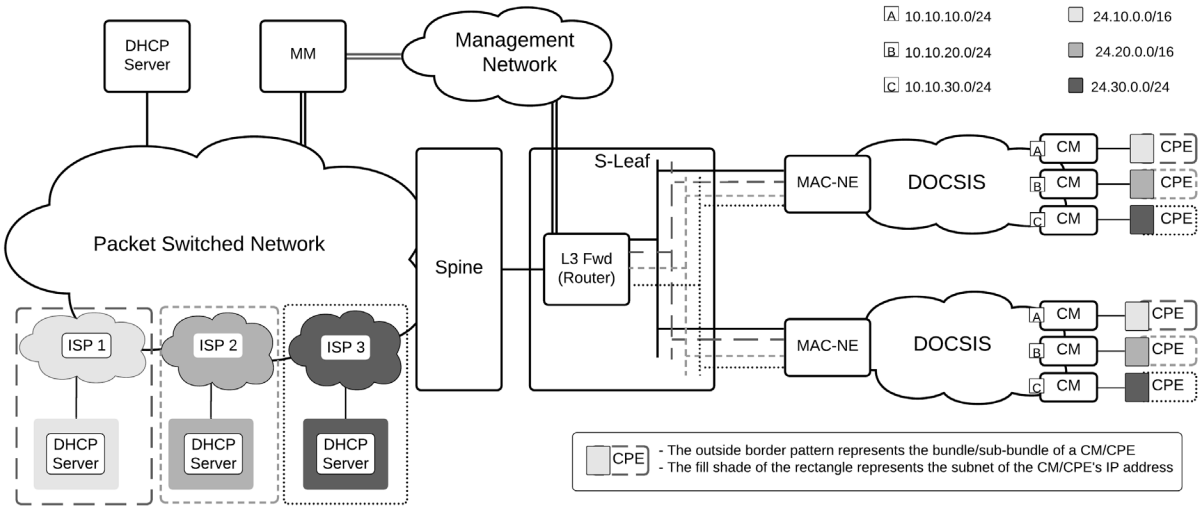

Figure 44 - CM Subnet-based Sub-bundles

10.4.3.4 Cable sub-bundles and Upstream DHCP packets

The following applies to rules for upstream DHCPv4 Discover or DHCPv6 Solicit packets.

- A cable sub-bundle DOCSIS Device Type rule can always be matched.
- A cable sub-bundle IP Subnet rule can never be matched.

- Whether a DHCP Field List rule can be matched depends on the option code(s) in the list.
- A CM Sub-Bundle rule can always be matched.

When an upstream DHCPv4 Discover or DHCPv6 Solicit packet matches any of the cable sub-bundle criteria, the MAC-NE MUST encapsulate it with the NSI encapsulation for that cable sub-bundle.

When an upstream DHCPv4 Discover or DHCPv6 Solicit packet does not match any of the cable sub-bundle criteria, the MAC-NE MUST encapsulate it with the NSI encapsulation of the top-level cable bundle. If necessary, it is the responsibility of the S-Leaf to forward upstream DHCP packets with top-level cable bundle encapsulation to a DHCP server capable of assigning sub-bundle IP addresses.

Downstream DHCP responses are expected to be encapsulated by the S-Leaf with the same encapsulation as the upstream DHCP request.

10.4.4 Unicast IP Services Forwarding

Unicast IP Services refers to the forwarding of unicast IP traffic for CMs and CPEs. The MAC-NE forwards downstream unicast IP traffic based on the destination Layer 2 MAC address in a "Non-Routing Mode" as described in [CCAP-TR]. All Layer 3 routing of unicast IP traffic is performed by a Layer 3 forwarder (i.e., router) in the S-Leaf.

MAC-NE Unicast IP service includes special handling of non-unicast ARP, ND, and DHCP traffic associated with the Layer 3 IP address of CMs and CPEs.

The MAC-NE forwards IP Services packets in the downstream direction based on the NSI encapsulation and destination MAC address. The MAC-NE does not verify the destination IP address of downstream traffic. The MAC-NE MUST discard downstream packets destined to an unknown MAC address.

The MAC-NE forwards all upstream IP Services traffic received on a DOCSIS interface to the S-leaf. The MAC-NE MUST forward upstream traffic destined for any of its IP Hosts to the S-Leaf, which is responsible for hair-pinning the traffic back to the RMD. The MAC-NE MUST NOT forward packets received on its DOCSIS interface(s) from any CM, eSAFE, or CPE device to its downstream DOCSIS interface(s).

The S-Leaf performs "local proxy-arp" for IPv4 and "local proxy-ND" for IPv6, meaning it will respond to ARP requests and ND Solicits on behalf of all devices on the same subnet.

10.4.4.1 IP Services Layer 2 Port Isolation

As mentioned in Section 10.1, the S-Leaf is expected to isolate the layer 2 forwarding of IP Services VLAN traffic on each MAC-NE access port to be point-to-point with an internal layer 3 router interface. The S-Leaf is expected to never forward or flood Ethernet layer 2 packets tagged with an IP Services NSI Encapsulation directly from one MAC-NE access port to another. This expectation for "IP Services Layer 2 Port Isolation" can be implemented in several ways on modern S-Leaf products using product features such as "port isolation" or "private VLANs". It can even be implemented with a careful set of access control list (ACL) rules. This section describes an *optional* set of requirements for the MAC-NE to enforce IP Services Layer 2 Port Isolation in the event the S-Leaf does not do so. An "IP Service frame" is defined in the upstream as originating from a SID identified with IP Services or from an MAC-NE internal IP Host. An IP Service frame in the downstream is one tagged with an NSI encapsulation configured for IP Services. For brevity, this section uses the acronym "SRMA" for "S-Leaf Router MAC Address", the ethernet MAC address of an S-Leaf's layer 3 routing interface.

A MAC-NE MAY support the Port Isolation requirements of this section. If an operator does not implement Port Isolation on the S-Leaf, the operator is expected to deploy a MAC-NE that does support these requirements.

MAC-NE-enforced IP Services Layer 2 Port Isolation is specified with the following assumptions:

- IP Services is forwarded only on an interface or sub-interface with an explicit (tagged) NSI encapsulation;
- The S-Leaf uses the same source MAC address for downstream DHCP relay and IP L3 interface router forwarding;

- The S-Leaf uses the same source MAC address of each IP L3 interface on all MAC-NE access ports;
- The S-Leaf enables IPv4 ICMP Internet Router Discovery Protocol [RFC 1256] and IPv6 Router Advertisement for each IP L3 interface.

10.4.4.1.1 SRMA MAC-NE Requirements for IP Services Layer 2 Port Isolation

If a MAC-NE supports IP Services Layer 2 Port Isolation, it MUST meet the following:

- Until the SRMA for an NSI encapsulation is determined, the MAC-NE does not filter frames in that encapsulation for IP Services Layer 2 Port Isolation.
- The MAC-NE snoops the source MAC address of the first downstream IPv4 [RFC 1256] or IPv6 Router Advertisement since a sub-interface became operational with an IP Service NSI Encapsulation and designate it as the "S-Leaf Router MAC Addresses" (SRMA) for that NSI encapsulation;
- The MAC-NE supports associating the same SRMA with more than one IP Services NSI encapsulation, up to and including all such encapsulations;
- The MAC-NE supports MMI configuration of the SRMA for each supported IP Services NSI Encapsulation;
- The MAC-NE prevents snooping to override an MMI configured SRMA;
- The MAC-NE supports MMI configuration to override or remove a snooped SRMA;
- The MAC-NE reports status for each SRMA, including whether it was snooped or configured.

Additionally, if a MAC-NE supports IP Services Layer 2 Port Isolation, it MAY limit the number of unique SRMAs supported to a vendor-defined maximum of at least a value of one (1).

10.4.4.1.2 General MAC-NE Requirements for IP Services L2 Port Isolation

If a MAC-NE supports IP Services Layer 2 Port Isolation for an IP Services NSI encapsulation with a designated SRMA:

- The MAC-NE MUST discard downstream traffic in that encapsulation when the Ethernet source MAC address differs from the SRMA;
- The MAC-NE MUST discard upstream packets in that encapsulation with a source MAC address that matches the SRMA.

10.4.5 IP Services Forwarding Database (IP-FDB)

The IP Services Forwarding Database (IP-FDB) of a MAC-NE describes forwarding of unicast IP addresses with IP Services.

Configurable attributes include:

- "IP-FDB Offline Timeout" which is the number of seconds after a CM goes offline and remains offline before removing the IP-FDB entry;
- "Maximum IP-FDB Notification Delay", which is the number of milliseconds the MAC-NE defers sending an IP-FDB notification involving a single CM in order to aggregate the updates of multiple CMs into the same notification message.

The IP-FDB reports as status how the MAC-NE forwards packets for IP Services.

At the top level, the IP-FDB status is a list of "CM Forwarding Entry" nodes. Each CM Forwarding Entry contains:

- MAC Domain Name<key>:
  Name of MAC domain associated with the CM. When a MAC domain is deleted or admin-downed, the MAC-NE MUST remove all CM Forwarding Entry nodes that reference that MAC domain.

- CM Mac Address <key>:
  MAC address of a CM reaching a set of CM/CPE IP addresses.
- Update Timestamp
  Date and time at which any field of the CM Forwarding Entry was last updated.
- CM Online {Boolean}:
  Whether the CM is currently online, i.e., reports a CM registration status of 'OPERATIONAL'.
- IP Forwarding Lists:
  Separate lists of IP Forwarding Entries for IPv4 and IPv6, each of which describes authorized MAC-NE forwarding of IP traffic to or from L2 MAC addresses reached through the CM.
  Each IP Forwarding Entry contains the following elements:
    - NSI Ref Name<key>
      Vendor-selected name for a reference to an NSI interface or sub-interface on which IP packets are forwarded. When an NSI interface or sub-interface referenced by an IP Forwarding Entry is deleted or admin-downed, the MAC-NE MUST delete the IP Forwarding Entry referencing it.
    - Forwarding MAC Address{key}:
      The Ethernet MAC address resolved to reach the Forwarding IP Prefix of this IP Forwarding Entry.
    - Forwarding IP Prefix {key}
      The IP address and prefix length reached at the Forwarding MAC Address for IP Services traffic forwarded across the referenced NSI. For the host IP address of a CM or CPE itself, this is the full-length host IP prefix (i.e. with prefix length "/32" or "/128"). For an IPv6 prefix delegated to a CM, this is the delegated prefix. For vendor-specific reporting of IPv4 or IPv6 routes advertised by a router in the CM or CPE, this is the IP prefix of the advertised route.
    - LearnedBy:
      Enumeration that describes how the IP Forwarding Entry was learned:
        - other(0) -- learned by another mechanism, e.g., vendor-specific operation;
        - dhcp(1) -- learned by snooping fully-completed DHCP, e.g.; IPv4 DHCP-ACK or IPv6 DHCPv6-REPLY;
        - cmConfig(2) – learned by CM configuration file of either MAC address or IP address and subsequently resolved;
        - slaac(3) – learned via SLAAC.
    - DOCSIS Device Type
      Enumeration identifying the DOCSIS device type as learned from:
        - The eDOCSIS DHCP device type option, e.g., "ECM" for an eDOCSIS embedded CM or various other choices for embedded eSAFE CPEs;
        - "CM" for a CM not reporting an eDOCSIS device type with DHCP;
        - "CPE" for an external CPE.
    - DHCP Lease Expiration
      If applicable, indicates the expiration date and time of a DHCP Lease.

A MAC-NE MUST accept on-change telemetry subscriptions to an individual CM Forwarding Entry. A MAC-NE MUST accept on-change telemetry subscriptions to the IP Forwarding Database. When an individual CM Forwarding Entry node is reported with RESTCONF or telemetry, the RMD MUST report its IPv4 and IPv6 lists of IP Forwarding Entry nodes as an atomic snapshot at a point in time. When reporting multiple CM Forwarding Entry nodes, a MAC-NE MAY report each CM Forwarding Entry node with snapshots at different points in time.

The IP-FDB is not required to report the IP prefixes advertised by an eRouter or CPE operating as a router. A MAC-NE MAY report them as a vendor-specific enhancement with a LearnedBy field of other(0).

Note that because IP Forwarding Entries require IP/MAC address resolution, unresolved assignments of either IP addresses or MAC addresses via configuration are not included in the IP-FDB. Because the IP-FDB describes only forwarding across an NSI, it does not include link-local IPv6 addresses.

The MAC-NE MUST implement a scalar "IP-FDB Offline Timeout" with a default of 0 seconds, which means IP Forwarding entries never time out.

When a CM goes offline, the MAC-NE MUST set the expiration timestamp for all IP Forwarding entries reached through the CM to "IP-FDB Offline Timeout" seconds in the future. An RMD MUST maintain all IP Forwarding entries for an offline CM until they expire. A MAC-NE MUST remove all IP-FDB entries that have expired. When an IP-FDB Offline Timeout configuration change has occurred, the RMD is not required to modify the expiration time of existing IP-FDB entries until they expire.

The MAC-NE MUST implement an MMI action to immediately expire all IP-FDB CM Forwarding entries for offline CMs.

The MAC-NE MUST implement an MMI action to remove a single IP Forwarding entry regardless of CM online/offline status, e.g., for clearing a stale entry when a CPE moves from one CM to another.

The MAC-NE MUST implement an MMI action to remove all IP Forwarding entries for an offline CM, e.g., for clearing stale entries when a CM moves to a different MAC-NE.

The particular algorithms for adding and deleting entries in the IP-FDB are vendor-specific, subject to the following requirements:

- The MAC-NE MUST report an IP-FDB entry for each known CM (online or offline).
- The MAC-NE MUST report an IP-FDB entry for each SAV-authorized CPE.
- The MAC-NE MUST report IP_FDB entries for each known delegated IPv6 Prefix.

The MAC-NE implements the requirements of DOCSIS subscriber management, e.g., limiting the maximum number of CPEs reached through a CM or filtering traffic from those CPEs.

The MAC-NE reports subsets of IP-FDB information in other YANG objects. Only the IP-FDB reports the CPE MAC addresses associated with a CPE IP address as a "Forwarding MAC Address".

10.4.5.1 Notifications

When a MAC Manager connects or reconnects to a MAC-NE, it is expected to resynchronize its knowledge of the MAC-NE's IP-FDB by initially reading the entire contents "in bulk" and thereafter updating that knowledge with IP-FDB notifications from the MAC-NE. Because the initial bulk read and notifications can occur concurrently, the update-timestamp or remove-timestamp of a notification permit the MAC Manager to identify the most recent change by the MAC-NE. A MAC Manager SHOULD subscribe to IP-FDB notifications before the initial bulk read.

The MAC-NE implements the following YANG notifications regarding the IP Forwarding Database, which forward as RESTCONF notifications to the MAC Manager:

- An "IP-FDB Update" notification that indicates the MAC-NE is adding or replacing the entire contents of "CM Forwarding Entry" nodes;
- An "IP-FDB Remove CMs" notification that indicates the MAC-NE is removing the entire contents of one or more "CM Forwarding Entry" nodes; and
- An "IP-FDB Remove Mac Domain" notification that indicates that the MAC-NE is removing all CM Forwarding Entry nodes on a Mac Domain.

The MAC Manager MUST subscribe to a particular IP-FDB notification on a MAC-NE in order to receive it. A MAC-NE MUST send an IP-FDB notification to a MAC Manager that has subscribed to it.

The MAC-NE MUST send an "IP-FDB Update" notification to a MM when any leaf or list entry of an IP-FDB "CM Forwarding Entry" (except DHCP Lease Expiration) is added, changed, or deleted and the CM Forwarding Entry continues to exist. The MAC-NE is permitted to defer sending individual updates in order to aggregate multiple IP-

FDB updates into a single notification message. The MAC-NE MUST defer "IP-FDB Update" notifications no longer than the configured "Maximum IP-FDB Update Delay".

When one or more CM Forwarding Entry nodes are deleted but a Mac Domain continues to exist as administratively enabled, the MAC-NE MUST send an "IP-FDB Remove CMs" notification. The MAC-NE is permitted to defer sending notifications for a single CM in order to aggregate multiple CMs into a single notification message. The MAC-NE MUST defer "IP-FDB Remove CMs" notifications no longer than the configured "Maximum IP-FDB Update Delay".

When an entire MAC Domain is deleted or administratively disabled, the MAC-NE MUST send an "IP-FDB Remove Mac Domain" notification in lieu of sending a notification for each removed CM Forwarding Entry on that Mac Domain.

10.4.6 Multicast IP Services

10.4.6.1 Introduction

Multicast can provide significant bandwidth savings in a network. Multicast is especially attractive in the cable network because of the broadcast nature of the cable downstream. In addition to providing end to end bandwidth savings, the cable RF network can be used effectively to distribute multicast streams to multiple downstream devices.

Broadcast video services using multicast allows the system to scale by allowing a single AUX video core element to generate and serve streams to all the MAC-NEs that are configured to receive the same broadcast lineup. Because broadcast Serving Groups can be quite large (~100,000 or more subscribers), using multicast to deliver the same copy to hundreds of MAC-NE devices provides significant cost savings for operators.

IP multicast for supporting unicast IP forwarding (IPv6 ND) is discussed in Section 10.4.

10.4.6.2 Subscriber Multicast Forwarding

For forwarding of downstream IP multicast sessions to subscriber CPEs, the MAC-NE acts as a CMTS, not as a transparent bridge.

The MAC-NE MUST implement CMTS requirements specified in [MULPIv4.0], Section 9.2, "Multicast Forwarding."

The MAC-NE MUST implement CMTS requirements specified in [MULPIv4.0], Section G.3, "Multicast Support for Interaction with Pre-3.0 DOCSIS Devices."

The cable operator can configure the cable modem to join IP multicast sessions during registration. Such multicast sessions are called Static Multicast Sessions. If configured, the MAC-NE MUST set up Static Multicast Sessions upon CM registration, as specified in [MULPIv4.0].

10.4.6.3 Explicit Tracking of CPEs Joined to a Multicast Group

As specified in [MULPIv4.0], in DOCSIS 3.0, all multicast control planes have been centralized and moved to MAC entity (MAC-NE) where each multicast joined from a CPE is handled and assigned a DSID flow.

The MAC-NE implementing this functionality will internally keep a list of joined CPEs; there will be a defined [FMA-YANG] object for getting the list of the joined CPEs and corresponding multicast groups.

10.4.6.4 Multicast and Sub-bundles

By using appropriate cable sub-bundle assignment criteria, a DOCSIS MAC domain can contain subsets of host IP addresses (CMs and CPEs) distinct from each other in different sub-bundles. Note that multiple sub-bundles can be configured with the same NSI encapsulation.

DOCSIS already requires the CMTS to mark all downstream IP multicast packets with a per-session Multicast DSID (Downstream Service ID) to prevent duplicated delivery on different Receive Channel Sets. In order to permit isolation of IP multicasts in the same session to hosts in different sub-bundles, the MAC-NE supports assigning separate Multicast DSIDs for each *combination* of sub-bundle NSI encapsulation and multicast session.

MAC-NE Requirements:

- The MAC-NE MUST support a per-sub-bundle Boolean configuration for "IP Multicast Isolation Enabled" with a default of 'false'.
- When IP Multicast Isolation is enabled for a sub-bundle, the MAC-NE MUST associate each joined IP multicast session with the NSI encapsulation of the sub-bundle to which the host joining that session is assigned.
- The MAC-NE MUST assign a separate DSID for each combination of mac-domain, session and NSI encapsulation for multicast sessions joined by MD hosts in a sub-bundle enabling IP Multicast Isolation.
- When a downstream IP multicast packet is received with an NSI encapsulation of a sub-bundle enabling IP Multicast Isolation, the MAC-NE MUST forward the packet downstream tagged with the Multicast DSID assigned for the combination of mac-domain, session, and NSI encapsulation.
    1. Example: A single mac-domain MD1 is assigned to top-level bundle 1 which has sub-bundles 1.2 and 1.3, both of which are enabled for IP Multicast Isolation. Sub-bundle 1.2 configures an NSI encapsulation with VLAN ID 102 and sub-bundle 1.3 configures an NSI encapsulation with VLAN ID 103.
    2. A host H2 in sub-bundle 1.2 joins session 224.10.10.10, for which the RMD assigns (if not assigned already) a Multicast DSID 1002 for the combination of (MD1, 224.10.10.10, VLAN ID 102)
    3. A host H3 in sub-bundle 1.3 joins the same session 224.10.10.10, so the RMD assigns a different Multicast DSID 1003 for the combination of (224.10.10.10, VLAN ID 103).
    4. A downstream multicast packet received on VLAN ID 102 is forwarded with DSID 1002, and so reaches only H2.
    5. A downstream multicast packet received on VLAN ID 103 is forwarded with DSID 1003, and so reaches only H3.

10.5 L2VPN Services

L2VPN Services relates to the use of the Layer 2 Forwarder and thus the use of the L2VPN/BSOD capability as defined by the L2VPN specification [L2VPN].

The MAC-NE MUST support L2VPN/BSOD functionality as defined by the L2VPN specification [L2VPN], but only limited to IEEE802.1Q NSI encapsulations (single and double tagged). Any other NSI encapsulations described in [L2VPN] are beyond the scope of this document.

Single and double VLAN tags will be deployed for both IP Services (configured statically via YANG) and BSOD/L2VPN services (configured dynamically via CM TLVs). Thus, in order to avoid accidental VLAN ID conflicts, the MAC-NE MUST support configuration option to reserve L2VPN service VLAN ID ranges via YANG.

The MAC-NE MUST reject the configuration of Ethernet sub-interface if the configured inner and/or outer VLAN IDs is within the L2VPN service reserved range. The registration of a L2VPN CM MUST be rejected if the provisioned inner or outer VLAN IDs are not in the reserved L2VPN VLAN ID ranges, if configured.

10.6 Multiple Ethernet Ports

A MAC-NE can be implemented with one or more Ethernet ports. The MAC-NE supports the aggregation of multiple Ethernet ports into one or more Link Aggregation Groups (LAGs) per the [IEEE 802.1AX] standard.

10.7 Quality of Service

Each packet forwarded by the MAC-NE between an NSI and RFI is considered to be forwarded by the IP Service, L2VPN Service, or L2TPv3 Service. How the MAC-NE schedules egress downstream or upstream traffic from different services is vendor-specific.

This section specifies how QoS is provided between flows of the same service when forwarding between the NSI and RFI.

This section specifies the rules for determining a 3-bit "Upstream NSI User Priority" for a particular upstream packet from each service. How that priority is encoded in an NSI encapsulation depends on the encapsulation. In the absence of vendor-specific configuration, the MAC-NE encodes the "Upstream NSI User Priority" in an 802.1Q single tag in the PCP bits, and in a double 802.1Q tag in the outer tag PCP bits.

10.7.1 IP Services QoS

IP Services forward packets from an NSI downstream to a DOCSIS Mac Domain, and upstream from a DOCSIS MAC Domain to an NSI.

10.7.1.1 Downstream IP Services QoS

IP Services forwarding determines the set of Mac Domains eligible to receive downstream IP unicast or multicast traffic base on its NSI encapsulation. For IP unicast downstream traffic, the destination MAC address determines a CM or CPE host that belongs to a single MAC Domain. For IP multicast downstream traffic, [MULPIv3.1] determines whether each eligible MAC domain actually forwards the packet. A MAC-NE conceptually replicates the downstream IP multicast packet to each MAC Domain that forwards it.

[MULPIv3.1] specifies how downstream Ethernet packets forwarded on a MAC Domain are classified to a service flow. Downstream QoS between service flows on the same MAC Domain is specified in [MULPIv3.1] with most details left as vendor-specific.

10.7.1.2 Upstream QoS

Each packet forwarded by the RMD between an NSI and RFI is considered to be forwarded by the IP Service, L2VPN Service, or L2TPv3 Service. How the MAC-NE schedules egress downstream or upstream traffic from different services is vendor-specific.

This section specifies how QoS is provided between flows of the same service when forwarding between the NSI and RFI.

This section specifies the rules for determining a 3-bit "Upstream NSI User Priority" for a particular upstream packet from each service. How that priority is encoded in an NSI encapsulation depends on the encapsulation. In the absence of vendor-specific configuration, the MAC-NE encodes the "Upstream NSI User Priority" in an 802.1Q single tag in the PCP bits, and in a double 802.1Q tag in the outer tag PCP bits.

10.7.2 IP Services QoS

IP Services forward packets from an NSI downstream to a DOCSIS Mac Domain, and upstream from a DOCSIS Mac Domain to an NSI.

10.7.2.1 Downstream IP Services QoS

IP Services forwarding determines the set of Mac Domains eligible to receive downstream IP unicast or multicast traffic base on its NSI encapsulation. For IP unicast downstream traffic, the destination MAC address determines a CM or CPE host that belongs to a single MAC Domain. For IP multicast downstream traffic, [MULPIv3.1] determines whether each eligible MAC domain actually forwards the packet. A MAC-NE conceptually replicates the downstream IP multicast packet to each MAC Domain that forwards it.

[MULPIv3.1] specifies how downstream Ethernet packets forwarded on a MAC Domain are classified to a service flow. Downstream QoS between service flows on the same MAC Domain is specified in [MULPIv3.1] with most details left as vendor-specific.

10.7.2.2 Upstream IP Services QoS

IP Service upstream traffic is classified to DOCSIS service flows by a CM and scheduled on upstream RF channels according to [MULPIv3.1], although most details are again vendor-specific. Conceptually, upstream IP service traffic is delivered for forwarding to an NSI port with its service flow known.

All DOCSIS upstream service flows are considered to have a 3-bit "Traffic Priority" that is one of the QoS Parameter Set attributes of a service flow.

All upstream NSI Ethernet packets are considered to have a 3-bit "User Priority". In the case of an [IEEE 802.1Q] encapsulation, the 3-bit NSI User Priority is directly encoded in the outer tag. Other encapsulations (e.g., VXLAN) can also support the representation of a 3-bit "User Priority" for Ethernet-layer QoS.

The RMD MUST support an MMI-configurable feature called "IP Services Upstream Priority Mapping" that is disabled by default. When the IP Services Upstream Priority Mapping feature enabled, the MAC-NE MUST map the 3-bit DOCSIS Traffic Priority of an IP Services upstream service flow to a 3-bit NSI User Priority.

Table 10 provides a 1-to-1 mapping of priorities.

Table 10 - Default IP Services Upstream Priority Mapping

Upstream DOCSIS Traffic Priority	Upstream NSI User Priority
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7

When IP Services Upstream Priority Mapping is disabled, the upstream NSI User Priority for IP Services packets is vendor specific.

The IP Services Upstream Priority Mapping feature does not apply to upstream packet forwarded with L2VPN Services or L2TPv3 Services.

10.7.3 L2VPN Services QoS

The RMD meets the mandatory QoS requirements in [L2VPN] for point-to-point L2VPN services applicable to the MAC-NE's supported NSI encapsulations.

10.7.4 L2TPv3 Services QoS

10.7.4.1 L2TPv3 Services Upstream QoS

All upstream L2TPV3 Services traffic is originated from an IP Host which is configured with an explicit VLAN Tag and 3-bit NSI User Priority. IP Hosts are described in Section 10.2.4.

10.7.4.2 L2TPv3 Services Downstream QoS

All QoS considerations for downstream L2TPv3 are in [R-DEPI] and [R-OOB].

10.7.5 Upstream Packet Scheduling

Upstream (outbound) NSI traffic from the MAC-NE to the S-Leaf is specified to have a 3-bit NSI User Priority value to be encoded in the NSI encapsulation, but this document does not specify any requirements on the internal upstream egress scheduling of packets with different User Priority values, e.g., with strict priority or weighted-fair-queueing discipline. Nor does this document specify any requirements regarding the upstream scheduling of packets with the same or different NSI User Priority settings from different services (e.g., IP services vs. L2VPN services). Upstream egress scheduling for all packets from all data forwarding services is vendor-specific.

10.7.6 Downstream Packet Scheduling

Downstream (inbound) NSI packets are generally marked for QoS in their NSI encapsulation, e.g., a 3-bit 802.1Q PCP field. Egress scheduling by an S-Leaf of packets with different NSI encapsulation priority is beyond the scope of this specification.

Once downstream packets are received by an NSI interface, this document specifies no requirements for relative quality of service between ingress packets with different NSI encapsulation beyond what is specified for a particular service. For example, ingress NSI encapsulation priority can affect the service flow to which downstream packets are classified for L2VPN service. This document specifies no requirements for relative quality of service for packets with the same or different NSI encapsulation priority intended for different services.

11 SECURE PROVISIONING OF CM AND CPE

The MAC-NE supports several features intended to secure the download of CM configuration files:

- a capability to prevent the disclosure of the IP address of the configuration file server (see [SECv4.0], Section 9.4.2.1, "TFTP Proxy");
- a capability to enforce that a CM downloads the correct configuration file according to the DHCP configurations offered to the CM (see [SECv4.0], Section 9.4.2.3, "Configuration File Name Authorization"); and
- a capability to verify that a CM registers with settings that match those in the downloaded configuration file (see [SECv4.0], Section 9.4.2.4, "Configuration File Learning").

11.1 DHCP

Because the MAC-NE is a Layer 2 forwarder, it cannot act as a typical Layer 3 DHCP Relay Agent. The DHCP functions performed by an I-CCAP or CCAP Core device will now be split between the MAC-NE and the S-Leaf.

The MAC-NE acts as a Layer 2 DHCP Relay Agent, which consists of adding various DHCP Relay Options to upstream DHCP packets and possibly removing those options in the downstream direction. The MAC-NE does not alter DHCP packets other than inserting or removing DHCP Relay options.

The S-Leaf performs regular Layer 3 DHCP Relay functionality. For example, the S-Leaf converts broadcast packets to unicast packets in the upstream direction and forwarding responses to clients in the downstream direction (honoring "Broadcast Bit" functionality).

The MAC-NE MUST support Layer 2 DHCP Relay functionality as described in section 3.8.1 ("Layer 2 DHCP Relay Agent") of [TR 101] and [RFC 3046]. The MAC-NE MUST insert DHCP Relay options (Option 82) to DHCP messages intercepted in the upstream direction before forwarding them to the S-Leaf, as follows:

- Option 82.1 - Circuit ID (with MAC-NE's MacNeUniqueId);
- Option 82.2 - Remote ID (CM MAC, according to [MULPIv4.0]); and
- Option 82.9 - Vendor-Specific Info (according to [MULPIv4.0] and [CANN DHCP]).

The MAC-NE treats the RF interfaces as "Untrusted" (see [TR 101], section 3.8.1.1 "Basic Operation"), thus the MAC-NE MUST discard all upstream DHCP messages that include Option 82. The MAC-NE MUST remove Option 82 information from all DHCP reply messages received from the S-Leaf before forwarding them to the RF interface.

The MAC-NE uses Option 82 information to forward the DHCP message to the appropriate CM.

The MAC-NE MUST support Lightweight DHCPv6 Relay Agent as specified in [RFC 6221]. The MAC-NE MUST insert DHCPv6 Relay options in DHCPv6 messages intercepted in the upstream direction before forwarding them to the S-Leaf, as follows:

- Interface ID Option ([RFC 8415] with MacNeUniqueId);
- Remote ID Option ([RFC 4649] with CM's MAC);
- CMTS Capabilities Option (as described in [MULPIv4.0] and [CANN DHCP]; and
- CM MAC Address Option (as described in [MULPIv4.0] and [CANN DHCP]).

When performing DHCP Relay, the MAC-NE MUST discard all RELAY-FORW messages received from RF interfaces in the upstream direction. For DHCP Relay, the MAC-NE MUST discard all RELAY-REPL messages in the downstream direction before transmission on the RF interface.

In order to perform Layer 2 DHCP Relay, the MAC-NE snoops/intercepts DHCP packets both in the upstream and downstream direction. The information gathered during snooping will be used for the IP Forwarding Database].

The MAC-NE MUST support DHCPv4 Leasequery as described in [RFC 4388].

The MAC-NE MUST support DHCPv4 Bulk Leasequery as described in [RFC 6926].

The MAC-NE MUST support DHCPv6 Leasequery as described in [RFC 5007].

The MAC-NE MUST support DHCPv6 Bulk Leasequery as described in [RFC 5460].

The MAC-NE makes all Leasequery requests from the IP Services host (see Section 10.2.4).

The MAC-NE SHOULD send a Bulk Leasequery after reinitialization of the device.

11.2 Source Address Verification (SAV)

Source Address Verification is a feature described in [SECv4.0] for verifying that the source IP address of upstream IP services traffic is "assigned" by the operator to the CM originating the traffic. The MAC-NE MUST implement all CMTS requirements of Source Address Verification as specified in [SECv3.1].

The MAC-NE does not verify the destination IP address of unicast downstream traffic.

A MAC-NE MUST implement the following MMI configuration objects:

- a per-mac-domain attribute that configures whether SAV is enforced on the MAC Domain;
- a per-MAC-NE attribute that configures whether the RMD enforces per-CM SAV Authorization Encodings (CM TLV 43.7);
- a per-MAC-NE object that configures a list of IP prefixes associated with an SAV Group Name;
- a per-cable-bundle/sub-bundle object that is a list of "helper" addresses that the MAC-NE considers to be a list of authorized DHCP servers;
- a per-cable-bundle/sub-bundle container for configuring DHCP Leasequery throttling consisting of a count of requests and a count of the number of seconds in an interval during which the MAC-NE throttles DHCP Leasequery requests from each CM; and
- a per-MAC-NE object that configures RMD behavior when a conflict occurs between static SAV assignment with CM configuration and dynamic DHCP learned SAV assignments. The object is an enumeration with the following choices:
  - on a conflict, prefer the static CM configuration (default); or
  - on a conflict, prefer the dynamic DHCP assignment.

The following requirements apply to the list of cable bundle/sub-bundle "helper" addresses. When the \helper list is non-empty, the MAC-NE MUST discard downstream DHCP packets selected for IP Services when the source IP address is not in the cable bundle helper list. When the helper list is empty, the MAC-NE does not check the authorization of the source IP address of downstream IP Services DHCP packets. Operation is not defined when the MAC-NE's authorized DHCP servers do not match the S-Leaf configuration of DHCP relay agent servers. The MAC-NE MUST support configuration of at least four (4) authorized DHCP servers, with an arbitrary combination of IPv4 and IPv6 addresses of the cable helpers.

Note that when SAV Authorization Encoding is enabled for a MAC Domain, the MAC-NE supports both the SAV Group Name Subtype (CM TLV 43.7.1) and Static SAV Prefix Rule Subtype (CM TLV 43.7.2) TLVs, where only one such subtype is supported for any given CM in that MAC Domain.

The MAC-NE rejects upstream IP services traffic when its source IP address is assigned to a different CM.

Per [SECv3.1], the MAC-NE IP Services function considers a source IP address to be "assigned" in any of the following cases.

1. It matches the Subscriber Management CPE IPv4 List (CM TLV 36) or Subscriber Management CPE IPv6 List (CM TLV 61) signaled for that CM in a REG-REQ-MP.
2. It matches a prefix in the CMTS prefix list referenced by the SAV Group Name (43.7.1) Encoding signaled for that CM in a REG-REQ-MP.

3. It matches a Static SAV Prefix Encoding (CM TLV 43.7.2) signaled for that CM in a REG-REQ-MP.
4. It has learned from snooping a DHCP assignment by an authorized DHCP server; examples include the following:
    a. the MAC-NE extracts assigned IP address information from a DHCPv4 DHCPACK message sent in response to an upstream DHCPv4 DHCPREQUEST or from a DHCPv6 Reply message sent in response to an upstream DHCPv6 Request received from a CM's SID;
    b. the MAC-NE receives a reply to DHCPv4 DHCPLEASEQUERY or DHCPv6 Leasequery response that verifies the assigned IP address for a host source MAC address; or
    c. the MAC-NE extracts an assigned IPv6 prefix from an IPv6 prefix delegation option.

Note that the MAC-NE does not permit RESTCONF configuration of static SAV assignments.

Although [SECv3.1] did not require a CMTS to implement DHCP Leasequery, this specification does require a MAC-NE to implement that protocol in order to avoid requiring all CPEs to repeat DHCP after a MAC-NE restart. When its IP Services function enables SAV and receives an upstream IP packet with an unassigned source IP address, a MAC-NE MUST initiate a DHCP LEASEQUERY request to its set of authorized DHCP servers, subject to throttling. The MAC-NE MUST limit upstream Leasequery transmission on a per-modem basis to no more than a configured number of requests per configured interval. IPv4 Leasequery operation is described in [RFC 4388]; the DHCPv6 Leasequery protocol and related standards are specified in [RFC 5007] and [RFC 4994].

A static SAV assignment can conflict with a dynamic DHCP assignment by assigning the same IP address as reached through a different CM. When a CM registration request contains a static SAV assignment that conflicts with a Forwarding Database assignment for the previously offline CM:

- if the configured SAV conflict policy is to prefer the static assignment, the MAC-NE MUST replace the dynamic SAV assignment with the static assignment, or
- in all other cases, the MAC-NE MUST reject the CM registration due to the conflicting SAV assignment TLV.

When a dynamic snooped DHCP assignment conflicts with the static SAV assignment of an online CM,

- If the configured SAV conflict policy is to prefer dynamic assignment, the MAC-NE MUST replace the static SAV assignment with the dynamic snooped DHCP assignment, or
- in all other cases, the MAC-NE MUST ignore the DHCP SAV assignment and keep the static SAV assignment.

Upon detecting a conflict between static and dynamic SAV assignment, the MAC-NE MUST log an ERROR severity event.

When RESTCONF changes the SAV conflict policy, the MAC-NE does not change any existing SAV assignments.

12 SECURITY AND CERTIFICATE OPERATIONS

12.1 Security Overview

The FMA standard introduces new challenges in securing the interfaces that interconnect the distributed components of the system together. These interfaces are described throughout this section as *security associations.* Each security protocol is described in the sections below.

12.2 MACsec Link Layer Security

The IEEE 802.1ab standard defines an encryption mechanism that is applied to the link between the switch (such as the S-Leaf) and an endpoint (such as the MAC-NE). This encryption can be enabled, if supported, by the MAC-NE during IEEE 802.1x network authentication if the method is set to EAP-MKA or if manually configured. For further detail, see [IEEE 802.1x].The MAC-NE SHOULD support MACsec. MACSec provides an encrypted connection between the MAC-NE and the S-Leaf per [IEEE 802.1ab]. Further information about MACsec is found in Section 13.5.2.1.

The MAC-NE MUST support the enabling and disabling of MACsec if supported. Depending on the network policy of the operator, MACsec on a particular MAC-NE can be enabled or disabled as appropriate.

12.3 Transport Layer Security (TLS) for the FMA System

TLS (Transport Layer Security) provides interface security to the FMA system specifications. The use of TLS provides a secure and reliable TCP based connection for the FMA system's management and control interfaces for both onboarding and status reporting purposes. The TLS protocol was selected because it is implemented as an open standard supporting mutual authentication with adequate defense against eavesdropping of the connection. The current specification does not address secure connectivity for unreliable connections implemented using the UDP transport protocol.

TLS supports the negotiation and use of compression methods. These methods are not specified within TLS and thus are not to be used over the MMI or PAI.

12.3.1 Application of TLS to the Interfaces of the FMA System

The TLS protocol is the default mechanism to secure the configuration and management interfaces of the FMA system. These interfaces are described from a security perspective as *security associations* in the system architecture. The TLS security associations include the following:

- The MMI Interface (Mm-MacNe) uses **RESTCONF over TLS** between the MAC Manager and the MAC-NE
- The PAI Interface (Pag-MacNe) uses **Protocol Buffers over TLS** between the PAG and the MAC-NE
- The Streaming Telemetry Interfaces (Oss-MacNe) which provide Subscription and Publication services using **gNMI** and **gRPC over TLS** exists between the MAC-NE and Streaming Telemetry client

As shown in Figure 45, the FMA system leverages TLS across each of its defined interfaces to protect the integrity of the control and management communication’s paths between elements of the system.

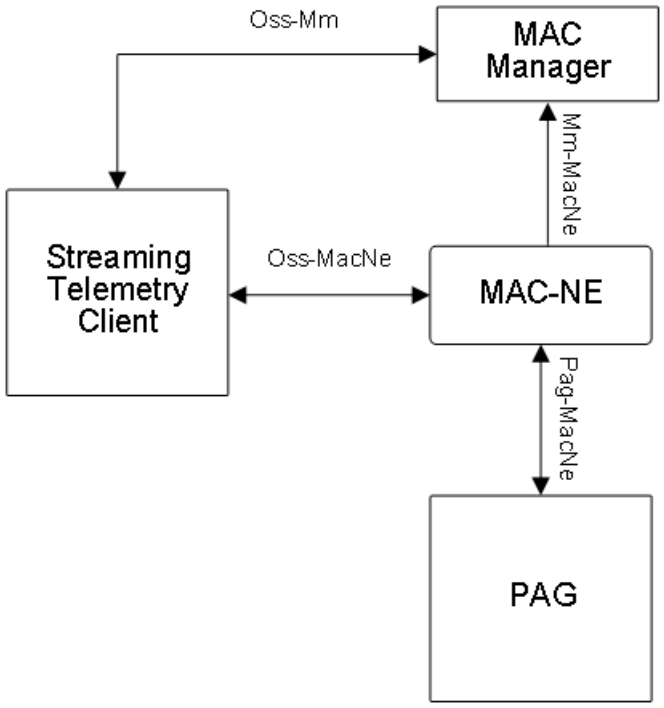

Figure 45 - TLS-Secured Logical Interfaces of the FMA System

All communications within FMA will use a secure channel (that is, a channel that is authenticated, authorized, encrypted, and attested). As the default secure channel for communications across the FMA system, TLS is leveraged for each interface for which another mechanism has not been explicitly defined. The exceptions are the use of SSH for some programmatic interfaces (such as NETCONF) and to support CLI access to some systems and the use of IPSec for GCP to support legacy video and lawful intercept, where applicable. For more detail on the Secure Shell protocol, see the SSH section of this document. For more detail on the Generic Control Protocol, review the [R-PHY] protocol specifications.

TLS will be implemented to ensure best current practices are exercised in securing the network and satisfying the following objectives:

- Perfect forward secrecy
- Mutual authentication
- Authentication based on PKI certificates

PKI certificate management requirements are defined in Section 12.4.1, X.509 Certificate Operations.

12.3.2 TLS Requirements in the FMA System

Each device in the FMA System is required to meet specific TLS requirements to ensure that it functions consistent with the RFCs and the desired behavior described further in this document.

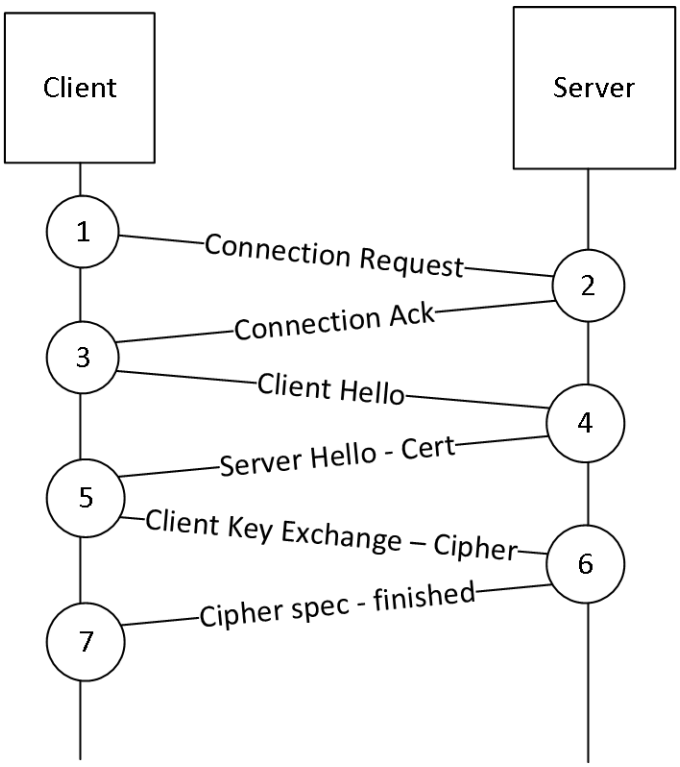

Figure 46 - TLS version 1.2 Connection Sequence

Compare Figure 46 to Figure 47 to better understand the steps that were collapsed in version 1.3 of the TLS protocol to speed the connection process in addition to other changes and improvements made to the protocol.

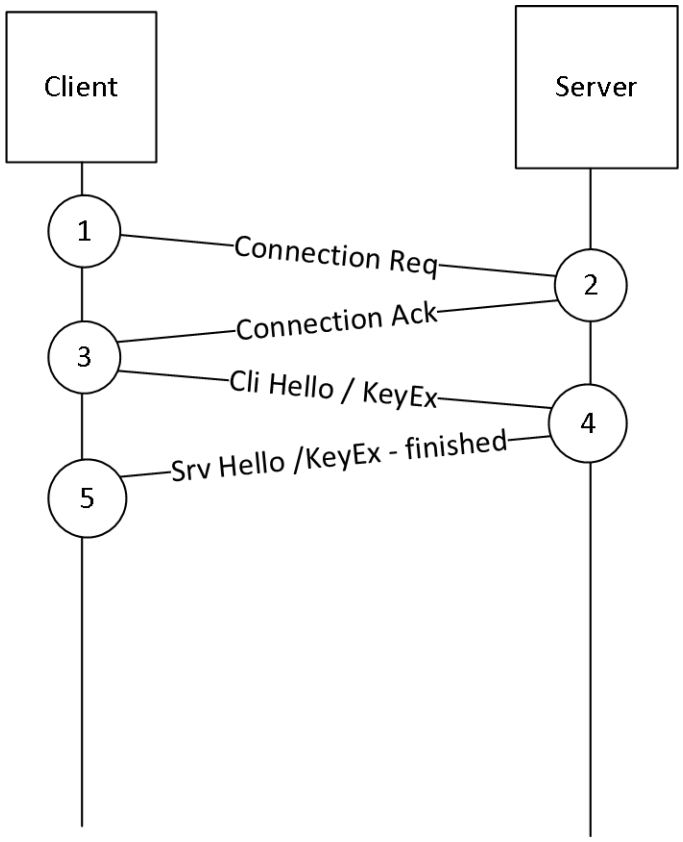

Figure 47 - TLS version 1.3 Connection Sequence

12.3.3 General TLS Version Requirements

The following sub-sections detail the TLS version requirements and behavior mandated for the MAC-NE, MAC Manager and PAG.

The MAC-NE MUST support TLS version 1.2. The MAC Manager MUST support TLS version 1.2. The PAG MUST support TLS version 1.2. The MAC-NE MAY support higher versions of the TLS protocol than 1.2. The MAC Manager MAY support higher versions of the TLS protocol than 1.2. The PAG MAY support higher versions of the TLS protocol than 1.2. The MAC-NE MUST disconnect a connection that attempts to use a TLS protocol version lower than 1.2. The MAC Manager MUST disconnect a connection that attempts to use a TLS protocol version lower than 1.2. The PAG MUST disconnect a connection that attempts to use a TLS protocol version lower than 1.2. The MAC-NE MUST prevent attempts to downgrade to a lower version of the TLS protocol after an end-point has initially connected using TLS version 1.2 or higher. The MAC Manager MUST prevent attempts to downgrade to a lower version of the TLS protocol after an end-point has initially connected using TLS version 1.2 or higher. The PAG MUST prevent attempts to downgrade to a lower version of the TLS protocol after an end-point has initially connected using TLS version 1.2 or higher. If an attempt to downgrade the TLS protocol from version 1.2 to a lower version occurs, the MAC-NE MUST disconnect the connection. If an attempt to downgrade the TLS protocol from version 1.2 to a lower version occurs, the MAC Manager MUST disconnect the connection. If an attempt to downgrade the TLS protocol from version 1.2 to a lower version occurs, the PAG MUST disconnect the connection. The Streaming Telemetry Client is expected to implement the TLS version and authentication behavior found in this section for compatibility with the FMA System.

TLS termination and load balancing that is external from the main business logic is common in virtualized or "cloud native" architectures and generally considered a best practice in system design. It's very possible that these functions could be provided by either a hardware appliance or by a 3rd-party service provider and therefore not be provided by the MAC Manager (or the PAG) itself. In such architectures, it would be difficult for the 3rd party or off-the-shelf

appliance to provide DOCSIS-formatted log messages. For this reason, the requirements for the MAC Manager (or the PAG) to generate such events as could result from a variety of TLS outcomes were defined in this specification as a series of recommendations rather than as mandatory requirements against the MAC Manager (or the PAG).

12.3.4 General TLS Cipher Suite Requirements

This section describes the general TLS requirements of the devices which comprise the FMA System. The FMA system devices include the MAC-NE, the MAC Manager, and the PAG.

If the MAC-NE supports ECC cryptography, it MUST support NIST curves per [FIPS-186-4]. If the MAC Manager supports ECC cryptography, it MUST support NIST curves per [FIPS-186-4]. If the PAG supports ECC cryptography, it MUST support NIST curves per [FIPS-186-4]. If the MAC-NE supports ECC cryptography, it MAY support additional curves beyond those defined in NIST with similar or stronger security. If the MAC Manager supports ECC cryptography, it MAY support additional curves beyond those defined in NIST with similar or stronger security. If the PAG supports ECC cryptography, it MAY support additional curves beyond those defined in NIST with similar or stronger security. If the MAC-NE supports RSA cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the MAC Manager supports RSA cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the PAG supports RSA cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the MAC-NE supports RSA cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the MAC Manager supports RSA cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the PAG supports RSA cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the MAC-NE supports RSA cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session. If the MAC Manager supports RSA cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session. If the PAG supports RSA cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session. If the MAC-NE supports ECC cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the MAC Manager supports ECC cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the PAG supports ECC cryptography, it MUST support up to 4096-bit keys for authentication and ephemeral key negotiation for the session. If the MAC-NE supports ECC cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the MAC Manager supports ECC cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the PAG supports ECC cryptography, it MUST support device certificates and private keys of 2048-bit length for authentication and ephemeral key negotiation for the session. If the MAC-NE supports ECC cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session. If the MAC Manager supports ECC cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session. If the PAG supports ECC cryptography, it MAY support 3072-bit and/or 4096-bit keys for its device certificate and private keys for authentication and ephemeral key negotiation for the session.

A MAC-NE MUST support the following cipher suites and algorithms for use with TLS v1.2 session encryption per Table 11 - TLS Version 1.2 Cipher Suites. A MAC Manager MUST support the following cipher suites and algorithms for use with TLS v1.2 session encryption per Table 11 - TLS Version 1.2 Cipher Suites. A PAG MUST support the following cipher suites and algorithms for use with TLS v1.2 session encryption per Table 11 - TLS Version 1.2 Cipher Suites. The Streaming Telemetry Client is expected to implement the cipher suites, algorithms and behavior mandated in this section in order to provide compatibility with FMA System requirements.

Table 11 - TLS Version 1.2 Cipher Suites

Pref Order	Cipher Suite	References
(1)	TLS_ECDHE_ECDSA-AES256-GCM-SHA384 TLS_ECDHE_RSA_AES_256_GCM_SHA384	RFC 5289

(2)	TLS_ECDHE_ECDSA-ChaCha20-poly1305-SHA256	RFC 7905
(3)	TLS_ECDHE_ECDSA-AES128-GCM-SHA256 TLS_ECDHE_RSA_AES_128_CBC_SHA256	RFC 5289

A MAC-NE implementing TLS v1.3 MUST support the following cipher suites and algorithms for use with TLS v1.3 session encryption per Table 12 - TLS Version 1.3 Cipher Suites. A MAC Manager implementing TLS v1.3 MUST support the following cipher suites and algorithms for use with TLS v1.3 session encryption per Table 12 - TLS Version 1.3 Cipher Suites. A PAG implementing TLS v1.3 MUST support the following cipher suites and algorithms for use with TLS v1.3 session encryption per Table 12 - TLS Version 1.3 Cipher Suites.

Table 12 - TLS Version 1.3 Cipher Suites

Preference Order	Cipher Suite	References
(1)	TLS_ECDHE_ECDSA-AES256-GCM-SHA384	RFC 5288
(2)	TLS_ECDHE_ECDSA-ChaCha20-poly1305-sha256	RFC 7905
(3)	TLS_ECDHE_ECDSA-AES128-GCM-SHA256	RFC 5289

12.3.5 TLS Event Requirements of the MAC-NE

The MAC-NE MUST log event ID 70070103 when it successfully completes TLS mutual authentication with the MAC Manager.

The MAC-NE MUST log event ID 70070115 when it fails to complete TLS mutual authentication with the MAC Manager due to timeout.

The MAC-NE MUST log event ID 70070116 when it fails to complete TLS mutual authentication with the MAC Manager due to cipher or algorithm mismatch.

The MAC-NE MUST log event ID 70070104 when it successfully completes TLS mutual authentication with the PAG.

The MAC-NE MUST log event ID 70070118 when it fails to complete TLS mutual authentication with the PAG due to timeout.

The MAC-NE MUST log event ID 70070119 when it fails to complete TLS mutual authentication with the PAG due to cipher or algorithm mismatch.

The MAC-NE MUST log event ID 89020100 when it successfully completes TLS mutual authentication with the Streaming Telemetry Client.

The MAC-NE MUST log event ID 89020101 when it fails to complete TLS mutual authentication with the Streaming Telemetry Client due to timeout.

The MAC-NE MUST log event ID 89020102 when it fails to complete TLS mutual authentication with the Streaming Telemetry Client due to mismatch in the cipher or algorithm.

The MAC-NE MUST log event ID 70070118 when the remote end attempts to downgrade the TLS version below version 1.2.

The MAC-NE MUST log event ID 70070119 when the remote end attempts to connect via a TLS version lower than version 1.2. The MAC-NE MUST log event ID 70070120 when the MAC Manager's certificate cannot be validated resulting in a rejection of authentication and a closure of the connection.

The MAC-NE MUST log event ID 70070121 when the MAC Manager's certificate has been revoked, has expired or revocation status is unknown after disconnecting the connection.

The MAC-NE MUST log event ID 70070126 when the PAG's device certificate cannot be validated resulting in a rejection of authentication and a closure of the connection.

The MAC-NE MUST log event ID 70070127 when the PAG's device certificate has been revoked, has expired or revocation status is unknown after disconnecting the connection.

The MAC-NE MUST log event ID 70070132 when the MDT Collector's device certificate cannot be validated resulting in a rejection of authentication and a closure of the connection.

The MAC-NE MUST log event ID 70070133 when the MDT Collector's device certificate has been revoked, has expired or revocation status is unknown after disconnecting the connection.

The MAC-NE MUST implement the Basic path Validation requirements per Section 12.3.9.1, General TLS Basic Path Validation Requirements.

12.3.6 TLS Event Requirements of the MAC Manager

The MAC Manager SHOULD log event ID 70080000 when it successfully completes TLS mutual authentication.

The MAC Manager SHOULD log event ID 70080001 when it fails to complete TLS mutual authentication due to a timeout.

The MAC Manager SHOULD log event ID 70080002 when it fails to complete TLS mutual authentication due to a mismatch in the cipher or algorithm.The MAC Manager SHOULD log event ID 70080003 when the remote end attempts to downgrade the TLS version below version 1.2.

The MAC Manager SHOULD log event ID 70080004 when the remote end attempts to connect via a TLS version lower than version 1.2.

The MAC Manager SHOULD log event ID 70080005 when the MAC-NE's device certificate cannot be validated resulting in a rejection of authentication and a closure of the connection.

The MAC Manager SHOULD log event ID 70080006 when the MAC-NE's device certificate has been revoked, has expired or revocation status is unknown after disconnecting the connection.

The PAG MUST implement the Basic path Validation requirements per Section 12.3.9.1, General TLS Basic Path Validation Requirements.

12.3.7 TLS Event Requirements of the PAG

The PAG MUST implement the cipher suites, algorithms and behavior mandated in Section 12.3.4, General TLS Cipher Suite Requirements. The PAG MAY implement optional functionality and ciphers found in Section 12.3.4, General TLS Cipher Suite Requirements.

The PAG MUST log event ID 70090000 when it successfully completes TLS mutual authentication with the MAC-NE.

The PAG MUST log event ID 70090001 when it fails to complete TLS mutual authentication with the MAC-NE due to timeout.

The PAG MUST log event ID 70090002 when it fails to complete TLS mutual authentication with the MAC-NE due to mismatch in the cipher or algorithm.

The PAG MUST log event ID 70090003 when the remote end attempts to downgrade the TLS version below version 1.2.

The PAG MUST log event ID 70090004 when the remote end attempts to connect via a TLS version lower than version 1.2.

The PAG MUST implement the Basic path Validation requirements per Section 12.3.9.1, General TLS Basic Path Validation Requirements.

The PAG MUST log event ID 70090005 when the MAC-NE's device certificate cannot be validated resulting in a rejection of authentication and a closure of the connection.

The PAG MUST log event ID 70090006 when the MAC-NE's device certificate has been revoked, has expired or revocation status is unknown after disconnecting the connection.

12.3.8 General TLS Authentication And Version Requirements

FMA System devices that rely on TLS to secure their control and management-plane communication paths are required to meet specific version requirements and perform authentication per the RFCs cited in the References section.

12.3.9 TLS Certificate Basic Path Validation Requirements of the FMA System

The sub-sections below detail the Certificate Basic Path Validation directives applied to the MAC-NE, MAC Manager, and PAG.

12.3.9.1 General TLS Basic Path Validation Requirements

The MAC-NE MUST have an ECC and/or RSA root CA certificate installed as a trust anchor for validating received certificates for each supported cryptographic scheme and PKI issuance authority. The MAC Manager MUST have an ECC and/or RSA root CA certificate installed as a trust anchor for validating received certificates for each supported cryptographic scheme and PKI issuance authority. The PAG MUST have an ECC and/or RSA root CA certificate installed as a trust anchor for validating received certificates for each supported cryptographic scheme and PKI issuance authority. Further detail of the certificate management for the FMA system is found in Section 12.4.1, X.509 Certificate Operations.

During TLS authentication, the MAC-NE is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange per [RFC 5246]. During TLS authentication, the MAC Manager is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange per [RFC 5246]. During TLS authentication, the PAG is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange per [RFC 5246]. See Figure 46 - TLS version 1.2 Connection Sequence.

If the MAC-NE implements TLS version 1.3, it is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange during authentication per [RFC 8446]. If the MAC Manager implements TLS version 1.3, it is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange during authentication per [RFC 8446]. If the PAG implements TLS version 1.3, it is required to include the issuing CA certificate with its own certificate in the TLS certificate message exchange during authentication per [RFC 8446]. See Figure 47 - TLS version 1.3 Connection Sequence.

FMA_SYS-REQ-998 The MAC-NE MUST validate certificates it receives using Basic Path Validation procedures as defined in [RFC 5280] and updated by [RFC 6818]. The MAC Manager MUST validate certificates it receives using Basic Path Validation procedures as defined in [RFC 5280] and updated by [RFC 6818]. The PAG MUST validate certificates it receives using Basic Path Validation procedures as defined in [RFC 5280] and updated by [RFC 6818]. The Streaming Telemetry Client is expected to implement the Basic path Validation requirements per [RFC 5280] and [RFC 6818] for compatibility with the FMA System.

The application of [RFC 5280] and [RFC 6818] to the FMA System demands the following for certificate handling behavior:

- If a certificate cannot be validated, authentication is rejected and the connection closed.
- If a certificate has been revoked, has expired or its revocation status is unknown, authentication will be rejected, the connection closed and the failure paths defined by the resiliency requirements sub-section of the X.509 Certificate Management section in this document will be followed.
- All OCSP responses and/or CRL responses for authenticity and freshness are validated before they can be used to check the revocation status of a certificate. See Section 12.4.1, X.509 Certificate Operations for guidance regarding acceptability of certificates and the freshness of the OCSP or CRL list maintained by the server.
- TLS stapling [RFC 8446] is also specified for the MAC-NE, providing OCSP revocation status of certificates via a somewhat more efficient means.

12.3.10 Connection Persistence and Re-Use

Because TCP connection and TLS session establishment (which relies on TCP) can be quite costly both in terms of performance and network latency, TLS sessions should be kept persistent as much as possible.

12.3.11 Session Caching

In TLS, it is possible to resume a previous session if it has been cached on both the TLS client and server. Resuming sessions drastically speeds up the session establishment, as fewer messages are exchanged, and authentication is based on previously generated symmetric keys.

It is recommended that TLS session caching be supported.When initiating an MMI TLS session, a MAC-NE MUST attempt to resume a cached session if it has retained a session for the MAC Manager.When initiating a PAI TLS session, a MAC-NE MUST attempt to resume a cached session if it has retained a session for the PAG. The duration for which a TLS client or server retains a cached session is a local policy and is implementation specific.

12.4 Re-Use from MHAv2 R-PHY

Network Authentication/Privacy/Integrity

- [IEEE 802.1x] Mutual Authentication using CableLabs public key infrastructure (PKI) Certificate Hierarchy
- MACsec (AES-GCM-128) with MACsec Key Agreement (MKA) protocol
- Optional at discretion of the operator

12.4.1 X.509 Certificate Operations

MHAv2 [R-PHY] specifies X.509 Certificate Management requirements on CCAP Cores and RPDs.

- For FMA, the MAC Manager MUST meet all CCAP Core requirements specified in [R-PHY], Section 10, "X.509 Certificate Management."
- For FMA, the Auxiliary Core MUST meet all CCAP Core requirements specified in [R-PHY], Section 10, "X.509 Certificate Management."
- For FMA, the MAC-NE MUST meet all RPD requirements specified in [R-PHY], Section 10, "X.509 Certificate Management."

12.4.2 Physical Protection of Keys

MHAv2 [R-PHY] specifies Physical Protection of Keys requirements on RPDs.

- For FMA, the MAC-NE MUST meet all RPD requirements specified in [R-PHY], Section 11, "Physical Protection of Keys in the RPD."

12.5 Secure Shell Requirements

During normal operation, the MAC-NE is managed through the MAC Manager. However, there are situations where the operator requires secure remote access to the MAC-NE for activities such as set-up prior to deployment, performing device maintenance or engaging in detailed troubleshooting. The MAC-NE provides a Secure Shell (SSH) server allowing remote access to the MAC-NE via a vendor-specific command line interface. Like the MAC-NE, there are similar scenarios that apply to the MAC Manager as well. Use cases for both the MAC Manager and the MAC-NE are detailed in Secure Shell section of [FMA-OSSI].

The FMA system is an application of the zero-trust security model. It is for this reason that whenever SSH is utilized, mutual authentication between endpoints and other best practices are assumed to have been followed in order to authenticate and secure the SSH interactions.

The FMA system can use X.509 certificates for a number of purposes including but not limited to SSH. For further detail, see X.509 Certificate Operations section of this document.

12.5.1 MAC-NE SSH Requirements

The requirements below establish the requirements against the MAC-NE for accepting SSH connections over its Converged Interconnect Network (CIN) interface.

The MAC-NE MUST NOT implement as SSH client.

The MAC-NE SHOULD provide an SSHv2 server.

If the MAC-NE implements an SSH server, it MUST provide a configurable object to enable or disable the SSH server.

If the MAC-NE implements an SSH server, it MUST be compliant with [RFC 4250], [RFC 4251], [RFC 4252], [RFC 4253], and [RFC 4254].

If the MAC-NE implements an SSH server, it MUST enable SSH server functionality by default.

If the MAC-NE implements an SSH server, it MUST NOT implement SSH port forwarding functionality.

If the MAC-NE implements an SSH server and the SSH server is disabled, the MAC-NE MUST log event ID 70070046.

If the MAC-NE implements SSH and the server is enabled, the MAC-NE MUST log event ID 70070047.

If the MAC-NE implements SSH and is unable to successfully start an SSHv2 server, the MAC-NE MUST log event ID 70070000.

If the MAC-NE implements SSH and the SSH server is enabled, then it MUST re-initialize per the onboarding process described in this specification if the MAC-NE is unable to successfully start an SSHv2 server.

When a MAC-NE that implements SSH successfully starts an SSHv2 server, the MAC-NE MUST log event ID 70070048.

A MAC-NE that implements SSH MUST allow ongoing SSH connections to remain open, subject to the SSH INACTIVITY_TIMER.

When the inactivity timeout is reached by the SSH client, the MAC-NE MUST close the connection to this SSH client and log event ID 70070101.

If the MAC-NE implements SSH, it MUST support the ecdh-sha2-* method for SSHv2 Key Exchange as specified in [RFC 5656]. (Note that in RFC 5656, the asterisks is used to denote the elliptic curve domain parameter identified as defined in the RFC sections 6.1 and 10).

A MAC-NE that implements SSH MUST support "required curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC-NE that implements SSH MAY support "recommended curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC-NE that implements SSH MUST support the following algorithms for SSHv2 Authentication:

- ECDSA
- RSA

A MAC-NE that implements an SSH Server MUST authenticate SSH client connection requests using the algorithm appropriate for the X.509 certificate and public key type provided in the connection request.

A MAC-NE that implements an SSH Server MUST support the following MAC algorithms:

- hmac-sha2-256
- hmac-sha2-512

A MAC-NE that implements an SSH Server MUST support the following algorithms for SSHv2 Encryption:

- aes128-gcm
- aes256-gcm (default)
- chacha20-poly1305

A MAC-NE that implements an SSH Server MUST refuse the connection with SSH clients if the SSH client attempts to use incompatible algorithms or a lower SSH version than the MAC-NE's SSH server and log event ID 70070102.

A MAC-NE that implements an SSH Server SHOULD NOT support password, host key, or public key authentication for SSH client authentication.

A MAC-NE that implements an SSH Server MUST log event ID 70070100 when an SSH client has successfully authenticated.

A MAC-NE that implements an SSH Server MUST log event ID 70070319 when an SSH client closed the SSH connection normally.

12.5.1.1 MAC-NE Support for Multi-Factor Authentication

If the MAC-NE implements SSH, it MUST support Multi-Factor Authentication (MFA) between an SSH client and the MAC-NE's SSH server.

When the MAC-NE detects an SSH authentication error from a client attempting to access the MAC-NE via multi factor authentication, the MAC-NE MUST log event ID 70070106.

12.5.1.2 MAC-NE Support for Certificate-Based Authentication

If the MAC-NE implements SSH and if NETCONF is used, the MAC-NE MUST support X.509 Certificates and validation procedures as defined in [RFC 5280] and [RFC 6187].

If the MAC-NE implements certificate-based SSH authentication, it MUST support mutual authentication between an SSH client and the MAC-NE SSH server using X.509 certificates.

If the MAC-NE implements certificate-based SSH authentication, it MUST provide a configurable object to enable or disable certificate revocation status checking.

If the MAC-NE implements certificate-based SSH authentication, it MUST support CRL validation for revocation checking as defined in [RFC 5280].

If the MAC-NE implements certificate-based SSH authentication, it MUST support CRL revocation checking to verify revocation status of the SSH Client Certificate before establishing the SSH connection.

If the MAC-NE implements certificate-based SSH authentication, it MUST support encoding of a URL for the CRL authoritative responder in either the certificate itself or as previously configured.

If the MAC-NE is unable to access the CRL revocation information in either the SSH client's certificate or as previously configured due to a missing URL, it MUST log event ID 70070104.

If the MAC-NE is unable to access the CRL revocation information due to misconfiguration or a lack of network availability in order to verify the SSH client's certificate, it MUST retry before logging event ID 70070104.

If OCSP revocation checking is configured and X.509 certificate authentication is used, the MAC-NE MUST verify OCSP status of the SSH Client Certificate before establishing the SSH connection.

The MAC-NE MUST support encoding of a URL for the OCSP authoritative responder in either the SSH client's certificate or as previously configured.

The MAC-NE MUST NOT proceed with the authentication process unless the SSH client's certificate has a value of "good" for the OCSP response in CertStatus field.

If the MAC-NE is unable to access the OCSP revocation information in either the SSH client's certificate or as previously configured due to a missing URL, it MUST log event ID 70070104.

The MAC-NE MUST provide a configurable retry counter for the purposes of determining the number of retries that OCSP revocation checking has attempted before timing out.

If the MAC-NE is unable to access the OCSP responder due to misconfiguration or a lack of network availability in order to verify the SSH client's certificate, it MUST retry before logging event ID 70070104.

In case the MAC-NE is not able to fetch the revocation information, the MAC-NE MAY allow the SSH connection to continue and log event ID 70070105.

When the MAC-NE detects an SSH authentication error from a client attempting to access the MAC-NE via X.509 certificate authentication, the MAC-NE MUST log event ID 70070104.

12.5.2 MAC Manager SSH Server Requirements

The requirements below establish the requirements against the MAC Manager for accepting SSH connections over its CIN interface.

The MAC Manager SHOULD provide an SSHv2 server.

If the MAC Manager implements an SSH server, it MUST provide a configurable object to enable or disable the SSH server.

If the MAC Manager implements an SSH server, it MUST be compliant with [RFC 4250], [RFC 4251], [RFC 4252], [RFC 4253], and [RFC 4254].

If the MAC Manager implements an SSH server, it MUST enable SSH server functionality by default.

If the MAC Manager implements an SSH server, it MUST NOT implement SSH port forwarding functionality.

If the MAC Manager implements an SSH server and the SSH server is disabled, the MAC Manager MUST log event ID 70080101.

If the MAC Manager implements SSH and the server is enabled, the MAC Manager MUST log event ID 70080102.

If the MAC Manager implements SSH and is unable to successfully start an SSHv2 server, the MAC Manager MUST log event ID 70080100.

When a MAC Manager that implements SSH successfully starts an SSHv2 server, the MAC Manager MUST log event ID 70080103.

A MAC Manager that implements SSH MUST allow ongoing SSH connections to remain open, subject to the SSH INACTIVITY_TIMER.

When the inactivity timeout is reached by the SSH client, the MAC Manager MUST close the connection to this SSH client and log event ID 70080008.

If the MAC Manager implements SSH, it MUST support the ecdh-sha2-* method for SSHv2 Key Exchange as specified in [RFC 5656]. (Note that in RFC 5656, the asterisks is used to denote the elliptic curve domain parameter identified as defined in the RFC sections 6.1 and 10).

A MAC Manager that implements SSH MUST support "required curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC Manager that implements SSH MAY support "recommended curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC Manager that implements SSH MUST support the following algorithms for SSHv2 Authentication:

- ECDSA
- RSA

A MAC Manager that implements an SSH Server MUST authenticate SSH client connection requests using the algorithm appropriate for the X.509 certificate and public key type provided in the connection request.

A MAC Manager that implements an SSH Server MUST support the following MAC algorithms:

- hmac-sha2-256
- hmac-sha2-512

A MAC Manager that implements an SSH Server MUST support the following algorithms for SSHv2 Encryption:

- aes128-gcm
- aes256-gcm (default)
- chacha20-poly1305

A MAC Manager that implements an SSH Server MUST refuse the connection with SSH clients if the SSH client attempts to use incompatible algorithms or a lower SSH version than the MAC Manager's SSH server and log event ID 70080009.

A MAC Manager that implements an SSH Server SHOULD NOT support password, host key, or public key authentication for SSH client authentication.

A MAC Manager that implements an SSH Server MUST log event ID 70080007 when an SSH client has successfully authenticated.

A MAC Manager that implements an SSH Server MUST log event ID 70080200 when an SSH client closed the SSH connection normally.

12.5.2.1 MAC Manager Support for Multi-Factor Authentication

If the MAC Manager implements SSH, it MUST support Multi-Factor Authentication (MFA) between an SSH client and the MAC Manager's SSH server.

When the MAC Manager detects an SSH authentication error from a client attempting to access the MAC Manager via multi factor authentication, the MAC Manager MUST log event ID 70080013.

12.5.2.2 MAC Manager Support for Certificate-Based Authentication

If the MAC Manager implements SSH and if NETCONF is used, the MAC Manager MUST support X.509 Certificates and validation procedures as defined in [RFC 5280] and [RFC 6187].

If the MAC Manager implements certificate-based SSH authentication, it MUST support mutual authentication between an SSH client and the MAC Manager SSH server using X.509 certificates.

If the MAC Manager implements certificate-based SSH authentication, it MUST provide a configurable object to enable or disable certificate revocation status checking.

If the MAC Manager implements certificate-based SSH authentication, it MUST support CRL validation for revocation checking as defined in [RFC 5280].

If the MAC Manager implements certificate-based SSH authentication, it MUST support CRL revocation checking to verify revocation status of the SSH Client Certificate before establishing the SSH connection.

If the MAC Manager implements certificate-based SSH authentication, it MUST support encoding of a URL for the CRL authoritative responder in either the certificate itself or as previously configured.

If the MAC Manager implements certificate-based SSH authentication, it MUST log event ID 70080018 if it is unable to access the CRL revocation information due to a missing URL or inability to reach the CRL authoritative responder.

If the MAC Manager is unable to access the CRL revocation information in either the SSH client's certificate or as previously configured due to a missing URL, it MUST log event ID 70080017.

If MAC Manager is unable to access the CRL revocation information due to misconfiguration or a lack of network availability in order to verify the SSH client's certificate, it MUST retry before logging event ID 70080017.

The MAC Manager MAY support OCSP for revocation checking as defined in [RFC 6960].

If OCSP revocation checking is configured and X.509 certificate authentication is used, the MAC Manager MUST verify OCSP status of the SSH Client Certificate before establishing the SSH connection.

The MAC Manager MUST support encoding of a URL for the OCSP authoritative responder in either the SSH client's certificate or as previously configured.

The MAC Manager MUST NOT proceed with the authentication process unless the SSH client's certificate has a value of "good" for the OCSP response in CertStatus field.

If the MAC Manager is unable to access the OCSP revocation information in either the SSH client's certificate or as previously configured due to a missing URL, it MUST log event ID 70080017.

The MAC Manager MUST provide a configurable retry counter for the purposes of determining the number of retries that OCSP revocation checking has attempted before timing out.

If the MAC Manager is unable to access the OCSP responder due to misconfiguration or a lack of network availability in order to verify the SSH client's certificate, it MUST retry before logging event ID 70080017.

In case the MAC Manager is not able to fetch the revocation information, the MAC Manager MAY allow the SSH connection to continue and log event ID 70080018.

When the MAC Manager detects an SSH authentication error from a client attempting to access the MAC Manager via X.509 certificate authentication, the MAC Manager MUST log event ID 70080017.

12.5.3 MAC Manager SSH Client Requirements

The requirements below establish the requirements against the MAC Manager for accepting SSH connections over its Convergent Interconnect Network (CIN) interface.

The MAC Manager SHOULD provide an SSHv2 client.

If the MAC Manager implements an SSH client, it MUST be compliant with [RFC 4250], [RFC 4251], [RFC 4252], [RFC 4253], and [RFC 4254].

If the MAC Manager implements SSH and is unable to initiate connection to the server, the MAC Manager MUST log event ID 70080104.

When a MAC Manager that implements SSH successfully initiates connection to the server, the MAC Manager MUST log event ID 70080105.

A MAC Manager that implements SSH MUST allow ongoing SSH connections to remain open, subject to the SSH INACTIVITY_TIMER.

When the inactivity timeout is reached by the SSH client, the MAC Manager MUST close the connection to this SSH server and log event ID 70080015.

If the MAC Manager implements SSH, it MUST support the ecdh-sha2-* method for SSHv2 Key Exchange as specified in [RFC 5656]. (Note that in RFC 5656, the asterisks is used to denote the elliptic curve domain parameter identified as defined in the RFC sections 6.1 and 10).

A MAC Manager that implements SSH MUST support "required curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC Manager that implements SSH MAY support "recommended curves" as specified in [RFC 5656] for ecdh-sha2-* method.

A MAC Manager that implements SSH MUST support the following algorithms for SSHv2 Authentication:

- ECDSA
- RSA

A MAC Manager that implements an SSH client MUST authenticate SSH server connection requests using the algorithm appropriate for the X.509 certificate and public key type provided in the connection request.

A MAC Manager that implements an SSH client MUST support the following MAC algorithms:

- hmac-sha2-256
- hmac-sha2-512

A MAC Manager that implements an SSH client MUST support the following algorithms for SSHv2 Encryption:

- aes128-gcm
- aes256-gcm (default)
- chacha20-poly1305

A MAC Manager that implements an SSH client MUST refuse the connection with SSH servers if the SSH server attempts to use incompatible algorithms or a lower SSH version than the MAC Manager's SSH client and log event ID 70080016.

A MAC Manager that implements an SSH client SHOULD NOT support password, host key, or public key authentication for SSH server authentication.

A MAC Manager that implements an SSH client MUST log event ID 70080014 when successfully authenticated by the SSH servers.

12.5.3.1 MAC Manager Support for Multi-Factor Authentication

If the MAC Manager implements SSH, it MUST support Multi-Factor Authentication (MFA) between an SSH server and the MAC Manager's SSH client.

When the MAC Manager fails to access the SSH server via multi factor authentication, the MAC Manager MUST log event ID 70080020.

12.5.3.2 MAC Manager Support for Certificate-Based Authentication

If the MAC Manager implements SSH and if NETCONF is used, the MAC Manager MUST support X.509 Certificates and validation procedures as defined in [RFC 5280] and [RFC 6187].

If the MAC Manager implements certificate-based SSH authentication, it MUST support mutual authentication between an SSH server and the MAC Manager SSH client using X.509 certificates.

If the MAC Manager implements certificate-based SSH authentication, it MUST provide a configurable object to enable or disable certificate revocation status checking.

If the MAC Manager implements certificate-based SSH authentication, it MUST support CRL validation for revocation checking as defined in [RFC 5280].

If the MAC Manager implements certificate-based SSH authentication, it MUST support CRL revocation checking to verify revocation status of the SSH Server Certificate before establishing the SSH connection.

If the MAC Manager implements certificate-based SSH authentication, it MUST support encoding of a URL for the CRL authoritative responder in either the certificate itself or as previously configured.

If the MAC Manager implements certificate-based SSH authentication, it MUST log event ID 70080020 if it is unable to access the CRL revocation information due to a missing URL or inability to reach the CRL authoritative responder.

If the MAC Manager is unable to access the CRL revocation information in either the SSH server's certificate or as previously configured due to a missing URL, it MUST log event ID 70080017.

If MAC Manager is unable to access the CRL revocation information due to misconfiguration or a lack of network availability in order to verify the SSH server's certificate, it MUST retry before logging event ID 70080017.

The MAC Manager MAY support OCSP for revocation checking as defined in [RFC 6960].

If OCSP revocation checking is configured and X.509 certificate authentication is used, the MAC Manager MUST verify OCSP status of the SSH Server Certificate before establishing the SSH connection.

The MAC Manager MUST support encoding of a URL for the OCSP authoritative responder in either the SSH server's certificate or as previously configured.

The MAC Manager MUST NOT proceed with the authentication process unless the SSH server's certificate has a value of "good" for the OCSP response in CertStatus field.

If the MAC Manager is unable to access the OCSP revocation information in either the SSH server's certificate or as previously configured due to a missing URL, it MUST log event ID 70080017.

The MAC Manager MUST provide a configurable retry counter for the purposes of determining the number of retries that OCSP revocation checking has attempted before timing out.

If the MAC Manager is unable to access the OCSP responder due to misconfiguration or a lack of network availability in order to verify the SSH server's certificate, it MUST retry before logging event ID 70080017.

When the MAC Manager detects an SSH authentication error from a server attempting to access the MAC Manager via X.509 certificate authentication, the MAC Manager MUST log event ID 70080017.

12.6 Secure Boot

Secure Boot offers the prospect of a hardware-verified, malware-free bootstrap process that can improve the security of many system deployments. Secure Boot technology ensures that the first code executed on a hardware platform is authentic and unmodified, establishing a root of trust in the system to build upon. This sets the foundation for establishing a full chain of trust to validate all levels of software executing on the hardware platform, including multiple phases of the boot process and the final executable image. The Secure Boot chain of trust is often attached to an immutable hardware anchor that provides secure device identity. For example, [IEEE 802.1ar] describes a standard for a secure device identifier, which is a cryptographic identity bound to a device and used for assertion of the device's identity. [IEEE 802.1ar]-compliant devices can rely on CableLabs issued RPD and MAC-NE device certificates and the executable signing procedures of Secure Software Download.

The MAC-NE SHOULD support Secure Boot technology.

The FMA will follow the Secure Boot requirements for the RPD as defined in the latest release of the CableLabs RPD specification.

12.7 Tamper Resistance/Detection

Section for future information, if needed.

12.8 Secure Software Download

12.8.1 Introduction

The FMA architecture supports Secure Software Download (SSD) to the MAC-NE. Authenticating the source and verifying the integrity of downloaded code is vital to the overall operation and security of the FMA architecture. The methods for Secure Software Download as well as the relevant specification text have been adapted from the DOCSIS 4.0 Security Specification [SECv4.0]. The Secure Software Download functionality is generally applicable to MAC-NEs installed in insecure locations.

In FMA, MAC Manager facilitates downloading software code for MAC-NEs. With respect to Secure Software Downloads, the MAC-NE assumes the function of a DOCSIS cable modem. It is envisioned that such an approach will allow operators to reuse much of the OSS infrastructure already deployed for CM software and security certificate management to perform equivalent functions for MAC-NEs. However, there are important differences to the upgrade procedure defined in this specification. These changes are summarized below and explained further within this section.

- The MAC-NE upgrade process relies on CVC certificates from the CableLabs DOCSIS 4.0 PKI. Legacy certificates are not supported.

- MAC-NE initialization involves connecting to and obtaining configuration information from a MAC Manager via the RESTCONF protocol. The software upgrade attributes received via the RESTCONF effectively replace equivalent TLVs received by a CM in a configuration file.
- Unlike a CM SSD process which needs to be enabled by inclusion of Code Verification Certificates (CVC), the MAC-NE is implicitly enabled for SSD at all times. The MAC Manager maintains control over this feature because it has a RESTCONF Client located at the south bound interface (see [FMA-MMI]) which has controls over RESTCONF configuration by using PUT methods to invoke Remote (RPC) operations that defined in FMA YANG models.
- The software code for the MAC-NE is signed with the CableLabs DOCSIS 4.0 PKI CVC certificates defined in [SECv4.0] and then validated by the MAC-NE.

12.8.2 Overview

The requirements defined in this section address these security objectives for the code download process:

- The MAC-NEs need to have a means to authenticate that the originator of any download code is a known and trusted source.
- The MAC-NEs need to have a means to verify that the downloaded code has not been altered from the original form in which it was provided by the trusted source.
- The process needs to simplify the operator's code file-handling requirements and provide mechanisms for the operator to upgrade or downgrade the code version of the MAC-NE on their network.
- The process allows operators to dictate and control their policies with respect to 1) which code files will be accepted by MAC-NE within their network; and 2) security controls that establish the security of the process on their network.
- The MAC-NEs are able to move freely among systems controlled by different operators.
- Support updating the Root CA Certificate in the MAC-NE (optional).
- Support updating the Device CA Certificate in the MAC-NE (optional).

The concerns of individual operators or the MAC-NE manufacturers may result in additional security related to the distribution or installation of code into the MAC-NE. This specification does not restrict the use of further protections, as long as they do not conflict with the requirements of this specification.

Multiple levels of protection are required to protect and verify the code download:

- The manufacturer of the MAC-NE code always applies a digital signature to the code file. The signature is verified with a certificate chain that extends up to the Root CA before accepting a code file. The manufacturer signature affirms the source and integrity of the code file to the MAC-NE.
- Though the manufacturer always signs its code file, an operator may later apply its code signature in addition to the manufacturer signature. If a second signature is present, the MAC-NE verifies both signatures with a certificate chain that extends up to the Root CA before accepting a code file.
- OSS mechanisms for the provisioning and control of the MAC-NE are critical to the proper execution of this process. SSDs are initiated by the MAC Manager during the initial MAC-NE configuration process, or during normal operation. The operator controls this process indirectly by connecting to the SSH server or SNMP Agent on the MAC Manager Northbound Interface. The MAC-NE SSD can be also initiated by operators via SSH directly on the MAC-NE.

The MAC-NE code file is built using a [RFC 2315]-compliant structure that is defined below, which is identical to the code structure used to upgrade CM software. Included in this structure are the following:

- the upgrade code image;
- the Code Verification Signature (CVS), i.e., the digital signature over the code image and any other authenticated attributes as defined in the structure; and

- the Code Verification Certificate (CVC), i.e., an [ITU-T X.509]-compliant certificate that is used to deliver and validate the public code verification key that will verify the signature over the code image. The DOCSIS Certificate Authority (CA), a trusted party whose public key is already stored in the MAC-NE, signs this certificate.

Figure 48 shows the basic steps required for the signing of a code image when the code file is signed only by the MAC-NE manufacturer, and when the code file is signed by the MAC-NE manufacturer and co-signed by an operator.

In FMA, the Root CA certificate is installed in each MAC-NE as a trust anchor. The code manufacturer builds the code file by signing the code image using a DOCSIS digital signature structure (see [SECv3.1]) with a Manufacturer CVC certificate and the issuing CVC CA certificate. The code file is then sent to the operator. The operator verifies that the code file is from a trusted DOCSIS manufacturer and has not been modified. At this point, the operator has the option of loading the code file on the Software Download server as-is, or of adding its signature and operator CVC and issuing CVC CA certificate to the code file. During the code upgrade process, the MAC-NE retrieves the code file from the Software Download server and verifies the new code image using the Root CA Certificate trust anchor before installing it. See Annex D for CVC chain details. See Figure 48 for a depiction of the code validation hierarchy.

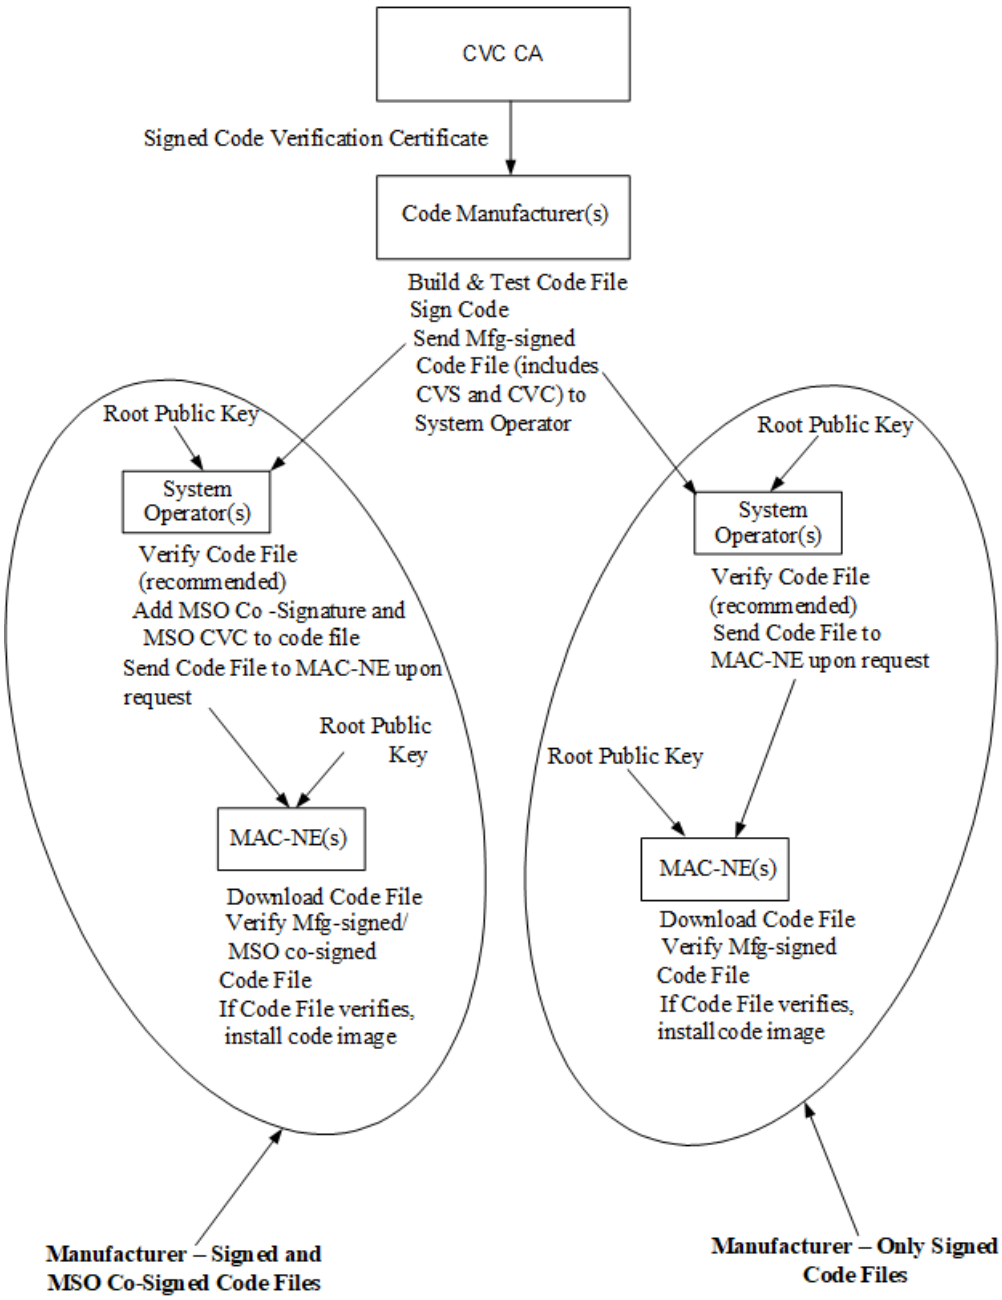

Figure 48 - Typical Code Validation Hierarchy

12.8.3 The MAC-NE Software Update Procedure

The secure software download (SSD) can be initiated in the following situations:

- During MAC-NE onboarding with the MAC Manager
    - The SSD process can be ordered by the MAC Manager after the establishment of the RESTCONF connection. The Software upgrade can occur immediately after the MAC Manager connects to the MAC-NE in order to initiate critical SW upgrades to resolve incompatibility preventing successful pairing of the MAC-NE and the MAC Manager.

  - It's also possible that the upgrade during onboarding of the MAC-NE is a result of previously issued instruction by the operator that prompts the MAC Manager to perform software upgrade of the MAC-NE.
  - After the MAC Manager has already configured the MAC-NE and it has reached operational state when the operator instructs the MAC Manager to initiate SSD for the MAC-NE. This can happen at any time when the device is operational. However, upgrades are typically performed during low service activity periods, i.e. within designated maintenance windows in the early morning hours.
- SSD is ordered by the operator using SSH. This may occur before or after the MAC-NE establishes a RESTCONF connection with MAC Manager to perform configuration. In this case, the operator provides necessary SSD parameters and command MAC-NE to update its software.

In the first and the second use case the SSD is initiated by MAC Manager via RESTCONF RPC. The YANG model for the SSD RPC can be found in [FMA-YANG].

The MAC-NE software upgrade procedure is presented in Figure 49.

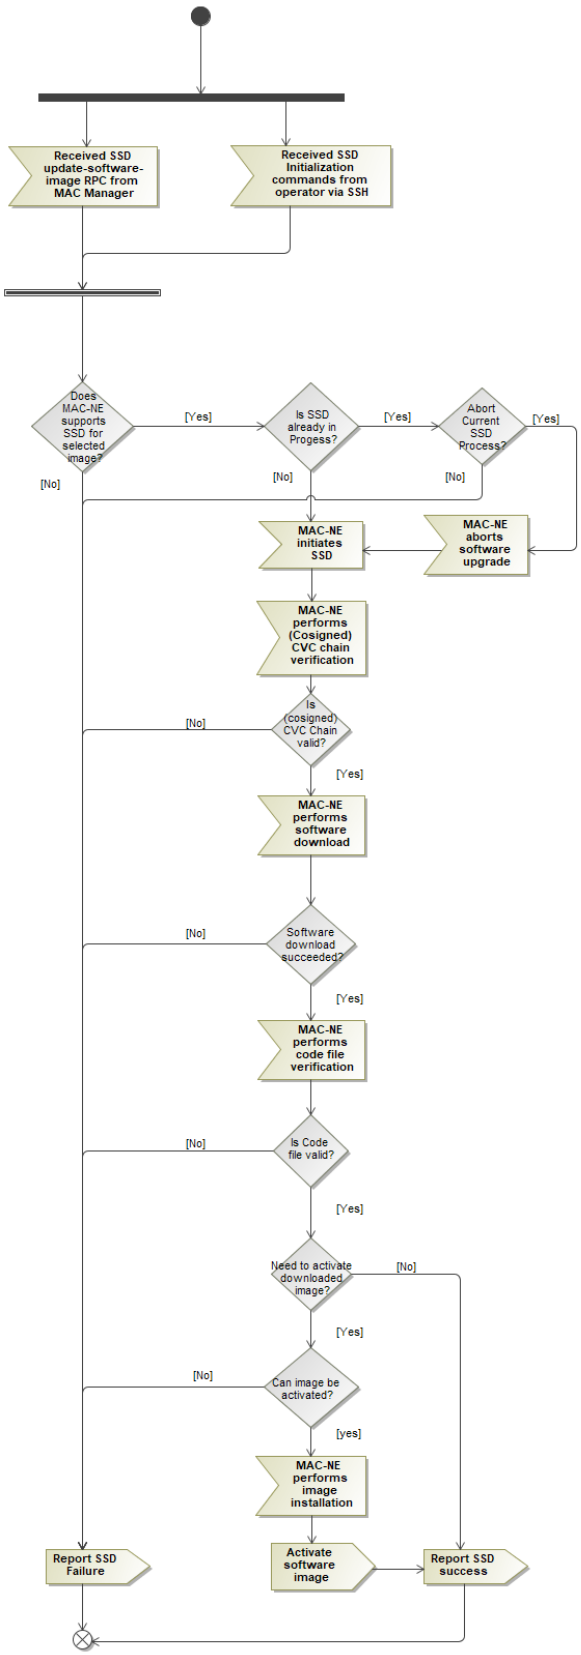

Figure 49 - Secure Software Download FSM

The MAC Manager is then ready for either receiving the initialization command from the operator via SSH or invoking the secure software download for a given MAC-NE via RESTCONF RFC as defined in [FMA-YANG].After connecting to the MAC-NE, if there is a need to upgrade software, the MAC Manager needs to select the Main Software Image (MSI) or any other image that it wants to update and provides the MAC-NEs with a set of SSD parameters as defined in [FMA-YANG]. These parameters include:

- Software Filename
- IP address of software download server (IPv4 or IPv6)
- Code Manufacturer's Code Verification Certificate.
- Co-signer's Code Verification Certificate.
- Software image index
- Transport protocol (TFTP, HTTP or HTTPS)

If the selected software image is upgradeable, the target MAC-NE then starts to verify the received CVC chain.

If the target MAC-NE fails to verify the CVC verification, the MAC-NE reports this failure through a RESTCONF notification defined in [FMA-YANG].

Once the target MAC-NE verifies the CVC chain is valid, the target MAC-NE then uses the received parameters to connect to the file server and retrieve the software update file as defined in [FMA-OSSI]. If the target MAC-NE fails to download the image, the MAC Manager needs to report SSD failure to the Notification Receiver at operator's system. The SSD will then terminate.

If the selected software image is not upgradeable or does the provided index value is invalid, the target MAC-NE will report the error via RESTCONF. The MAC-NE will then terminate SSD process.

Notice that in some cases, the target MAC-NE may already perform the SSD for selected software image index via other methods (i.e., initiate SSD by providing MAC-NE SSD parameters via SSH).

When receiving the update request from the MAC Manager with a set of SSD parameters, the target MAC-NE, while still in process of performing SSD, will send back the error message via RESTCONF as defined in [FMA-OSSI]. The MAC Manager can choose to abort the current SSD process via RESTCONF RPC Operation as defined in [FMA-YANG] so that the target MAC-NE can proceed its request or wait for its completion.

After successfully downloading the software image, MAC-NEs need to verify the code file by the digital signature and CableLabs DOCSIS 4.0 PKI as defined in Section 12.8.4 before accepting the downloaded code file as valid. The MAC-NE MAY perform additional identity and integrity checks that are vendor-specific.

If the received code file is valid, the MAC-NE will automatically upgrade its software and reboot if the code file is for Main Software Image (MSI) or current running software. The MAC-NE then repeats the entire initialization process including authentication. Image validation uses the same method involving digital signatures and the PKI certificate.

If the received code file is invalid or if the software download process fails, the MAC-NE needs to report an SSD failure as defined in Section 12.8.

Alternatively, after successfully acquiring the selected software image and installing it, the MAC Manager can choose to activate the selected software image on the target MAC-NE via corresponding RESTCONF RPC operation as defined in [FMA-YANG].

If MAC-NE considers that the selected image cannot be activated, the MAC-NE needs to reject the activation commands from the MAC Manager and reports SSD failure via corresponding RESTCONF RPC operation as defined in [FMA-YANG].

12.8.4 Software Code Upgrade Requirements

The following sections define the requirements of the MAC-NE's software code upgrade and verification process. The MAC-NE code upgrade process is prepared and verified as described below.

12.8.4.1 Secure Software Download General Requirements

FMA_SYS-REQ-822 If the MAC Manager indicates that a software upgrade is required, then the MAC-NE MUST perform the upgrade. The mechanisms to initiate a software upgrade are detailed in [FMA-YANG], but the MAC Manager logic to determine when to request the MAC-NE upgrade are left as an implementation detail.

FMA_SYS-REQ-360 The MAC-NE MUST support a software download initiated by the MAC Manager via RESTCONF.

FMA_SYS-REQ-361 The MAC-NE MAY support a software download initiated by the operator via SSH.

FMA_SYS-REQ-362 The MAC-NE MUST log a DOCSIS event ID 70070601 when the SSD is initialized via RESTCONF.

FMA_SYS-REQ-363 The MAC-NE MUST log a DOCSIS event ID 70070600 when the SSD is initialized via other methods such as CLI (i.e., SSH).

FMA_SYS-REQ-364 The MAC-NE MUST support TFTP compliant with [RFC 1350] for software file downloads as specified in [R-PHY].

- FMA_SYS-REQ-365 The MAC-NE MUST include the TFTP block size option [RFC 2348] when requesting the software image file via TFTP.
- FMA_SYS-REQ-366 The MAC-NE MUST request a block size of 1448 octets if using TFTP over IPv4.
- FMA_SYS-REQ-367 The MAC-NE MUST request a block size of 1428 octets if using TFTP over IPv6.

FMA_SYS-REQ-368 The MAC-NE MUST support HTTP compliant with [RFC 1945] or [RFC 2616] for software file downloads.

FMA_SYS-REQ-369 The MAC-NE SHOULD support HTTPS compliant with [RFC 2818] for software file downloads.

FMA_SYS-REQ-823 If HTTPS is used for software file downloads, the MAC-NE MUST support at minimum TLSv1.2 as defined in [RFC 5246].

FMA_SYS-REQ-372 The MAC-NE MUST use the cipher suites specified in DOCSIS 4.0 for secure software downloads.

FMA_SYS-REQ-373 If the specified filename provided via RESTCONF does not match the current filename of the selected software image of the MAC-NE, the MAC-NE MUST request this specified file from the software file server.

FMA_SYS-REQ-374 If the specified filename provided via RESTCONF matches the current filename of the selected software image of the MAC-NE, the MAC-NE MUST ignore the software upgrade request. There is no DOCSIS event generated in this case.

FMA_SYS-REQ-375 If the MAC-NE communicates with MAC Manager via IPv4 and receives an IPv4 software file server address via RESTCONF, the MAC-NE MUST use the specified IPv4 software file server address.

FMA_SYS-REQ-376 The MAC-NE MUST ignore the SSD request with IPv6 address when it communicates MAC Manager using IPv4.

FMA_SYS-REQ-377 If the MAC-NE communicates with MAC Manager via IPv6 and receives an IPv6 software file server address via RESTCONF, MAC-NEs MUST use the specified IPv6 software file server address.

FMA_SYS-REQ-378 The MAC-NE MUST ignore the SSD request with IPv4 address when it communicates MAC Manager using IPv6.

FMA_SYS-REQ-379 The MAC-NE MUST log a DOCSIS event ID 70070603 if MAC-NEs cannot reach the configured server.

FMA_SYS-REQ-380 The MAC-NE MUST log a DOCSIS event ID 70070604 if the server does not contain configured file.

FMA_SYS-REQ-381 If the download fails to start, MAC-NEs MUST retry the download for a minimum of 3 times.

FMA_SYS-REQ-382 The MAC-NE MUST log a DOCSIS event ID 70070605 when the number of retries is exhausted.

FMA_SYS-REQ-383 If the download fails after it is started, MAC-NEs MUST retry the download for a minimum of 3 times.

FMA_SYS-REQ-384 The MAC-NE MUST log a DOCSIS event ID 70070602 when the number of retries is exhausted.

FMA_SYS-REQ-385 If the MAC-NE cannot complete the file transfer for any reason, MAC-NEs MUST remain capable of accepting new software download requests (without operator or user interaction), even if power or connectivity is interrupted between attempts.

FMA_SYS-REQ-386 The MAC-NE MUST log a DOCSIS event ID 70070608 if the software download is interrupted by a power failure.

FMA_SYS-REQ-387 The MAC-NE MAY report the failure asynchronously to the network manager.

FMA_SYS-REQ-388 During the period that the software upgrade is in progress, MAC-NEs MUST defer rebooting due to the expiration of any timer, except those related to the SSD, in order to avoid disrupting the upgrade.

FMA_SYS-REQ-389 The MAC-NE MUST continue to operate with the existing software if an upgrade cannot be performed.

FMA_SYS-REQ-390 The MAC-NE MUST verify that the downloaded image is valid for itself as defined in [R-PHY].

FMA_SYS-REQ-391 If the image is valid, the MAC-NE MUST write the new software image to non-volatile storage.

FMA_SYS-REQ-392 If the software upgrade was for the Main Software Image (MSI), the MAC-NE MUST restart with the new code image with a MAC-NE Initialization Reason of SW_UPGRADE_REBOOT.

FMA_SYS-REQ-393 If the software upgrade was for the currently running image, the MAC-NE MUST restart itself with the new code image with the MAC-NE Initialization Reason of SW_UPGRADE_REBOOT.

FMA_SYS-REQ-394 If the software upgrade was an image other than the Main Software Image (MSI) or the currently running image, the MAC-NE MUST continue operating.

FMA_SYS-REQ-395 The MAC-NE MUST log a DOCSIS event ID 70070610 when it successfully completes SSD initiated via RESTCONF.

FMA_SYS-REQ-396 The MAC-NE MUST log a DOCSIS event ID 70070609 when it successfully completes SSD initiated via CLI (i.e., SSH).

The MAC-NE MAY abort the existing SSD process.

If there is already a software upgrade in process for the selected software image and the MAC-NE chooses to abort that existing process, the MAC-NE MUST log a DOCSIS event ID 70070621.

If there is already a software upgrade in process for the selected software image and the MAC-NE chooses NOT to abort that existing process, the MAC-NE MUST log a DOCSIS event ID 70070619 and terminates the SSD process.

The MAC-NE MUST reject the activation command and log a DOCSIS event ID 70070620 if the MAC-NE considers that the selected image cannot be activated.

13 ONBOARDING/INITIALIZATION OF MAC-NE

13.1 Overview

When the MAC-NE device powers up, it goes through a series of steps before becoming operational. These steps are shown in Figure 50 and explained in this section. Note that Figure 50 is the highest level in a set of nested state machines; detailed processing is shown in lower level diagrams and described later in this section. As can be seen from Figure 50, failures reported back to the main level result in a MAC-NE reboot. Actions specific to the particular error and state are shown in the more detailed finite state machine (FSM) figures.

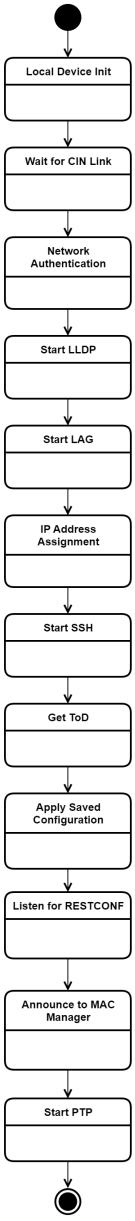

Figure 50 - MAC-NE Initialization FSM

13.2 Logging and Recording

During the initialization process, the MAC-NE MUST maintain a record of the current state of all FSMs and a log of significant events for all of its FSMs.

The MAC-NE MUST be able to report the FSM states and other device logs to the MAC Manager or other management system on request. The API for requesting this log is detailed in [FMA-MMI] and the object models in [FMA-YANG].

States that occur prior to the MAC-NE being granted network access are defined for completeness, but they cannot be reported synchronously.The MAC-NE MUST log these state transitions in the local log in order to allow the MAC Manager to parse them after initialization is complete through the [FMA-MMI] interface.

If the DHCP options presented to the MAC-NE contain one or more syslog server addresses, the MAC-NE MUST forward its retained FSM boot logs and any significant events accumulated prior to it having access to the remote syslog server.

Refer to [FMA-OSSI] for a list of standard events during initialization.

13.3 Local Device Initialization

Following a power cycle or reboot, the MAC-NE may need to run device-specific diagnostic and initialization code. This specification does not attempt to standardize this process other than to enable logging and error reporting in a standard manner. Refer to Figure 51 for further details of local MAC-NE initialization when there is saved configuration information.

The MAC-NE MUST log event ID 70070002 as soon in the boot process as possible.

If configuration parameters have been saved over a reboot, the MAC-NE MUST use them. If no saved parameters are available, the MAC-NE MUST use default parameters.

If it has appropriate stored configuration data, the MAC-NE MUST establish RF power and pilot tones as part of the local initialization before Announcing to the MAC Manager. Such actions are beneficial to the network.

If the stored configuration data cannot be applied, the MAC-NE SHOULD log event ID 70070005.

If a failure occurs during local initialization, the MAC-NE SHOULD log event ID 70070004 if possible.

If further operation is not possible following a failure, the MAC-NE SHOULD log event ID 70070002 with P1 = hardReset and then reboot.

Other actions following a failure are vendor-specific.

If local initialization completes successfully, the MAC-NE SHOULD log event ID 70070001 with P1 set to a vendor-specific event code or text string if possible.

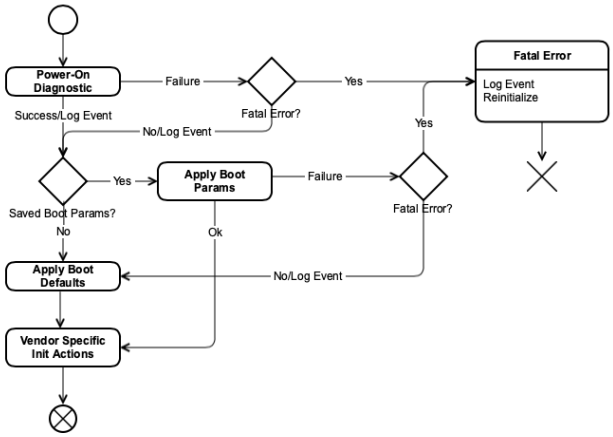

Figure 51 - Local Device Initialization FSM

13.4 Wait for CIN Link

The MAC-NE MUST attempt to establish a network connection with the CIN network on all enabled ports. This might include detecting and initializing SFP+ optics, PON networking components, or some other networking technology. This specification does not restrict the networking technologies used for providing connectivity. The diagram below assumes that Ethernet over optics using SFP+ modules connects the MAC-NE to the CIN network. See Figure 52 for further details.

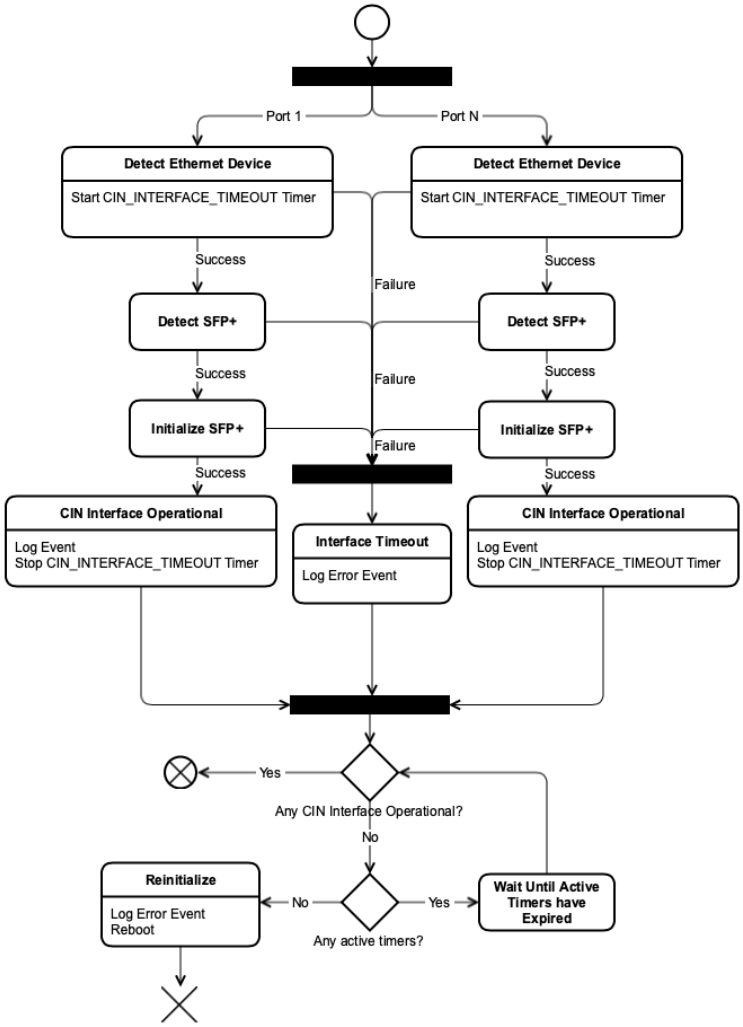

Figure 52 - Wait for CIN Link FSM

The MAC-NE MUST support an Ethernet-based Link Layer (L2) on its CIN interface(s). The MAC-NE MUST support an IP-based Network Layer (L3) on its CIN interface(s). The MAC-NE MUST support TCP-based Transport Layer (L4) on its CIN interface(s).

The MAC-NE MUST start the CIN_INTERFACE_TIMEOUT timer on each network interface port when local initialization completes.

The MAC-NE MUST stop the CIN_INTERFACE_TIMEOUT timer on each network interface port when the port becomes operational.

The MAC-NE MUST support saving the enabled/disabled state of each network port that connects the MAC-NE to the CIN interface across resets.

The MAC-NE MUST use the stored network port enabled/disabled state information when determining which ports will be enabled as part of initializing the CIN interface. Details pertaining to the CIN interface port configuration and how this information is preserved across reset is detailed in [FMA-YANG].

If failure occurs during initialization of the MAC-NE's CIN-facing network ports, the MAC-NE MUST wait until all enabled ports reach the CIN_INTERFACE_TIMEOUT before logging the event. If all of the MAC-NE's CIN-facing network ports reach the default or configured value of the CIN_INTERFACE_TIMEOUT and none of the ports

have become operational, the MAC-NE MUST log event ID 70070303 and re-initialize. Assigning Network Layer (L3) settings to the MAC-NE's ports is detailed later in this specification.

When a CIN interface becomes operational, the MAC-NE MUST log event ID 70070302.

When the first port that connects the MAC-NE to the CIN network interface becomes operational, the MAC-NE MUST proceed to perform authentication with the network.

13.4.1 CIN Interface Timeout

This attribute configures the interval (in seconds) that the MAC-NE waits for each of its network ports to become operational on the CIN interface after it has completed local initialization. If even a single port of the multi-port MAC-NE becomes operational with the CIN, the MAC-NE will proceed with initialization, logging an event for each of the ports that failed to come up.

The MAC-NE MUST provide a mechanism to configure CIN_INTERFACE_TIMEOUT during a staging process.

The MAC-NE MUST retain the value for CIN_INTERFACE_TIMEOUT in non-volatile storage.

Table 13 - MAC-NE CIN Interface Timeout

Constant	Valid Range	Default Value
CIN_INTERFACE_TIMEOUT	1...600 seconds	120 seconds

13.5 Network Authentication

13.5.1 Problem Definition

The FMA security architecture applies the concept of a zero trust. A MAC-NE in a zero trust environment is shown in Figure 53.

To access other devices in the FMA system, MAC-NEs are required to mutually authenticate with the network. This is accomplished by using [IEEE 802.1x] to perform this function using digital certificate credentials issued by a trusted public key infrastructure (PKI).

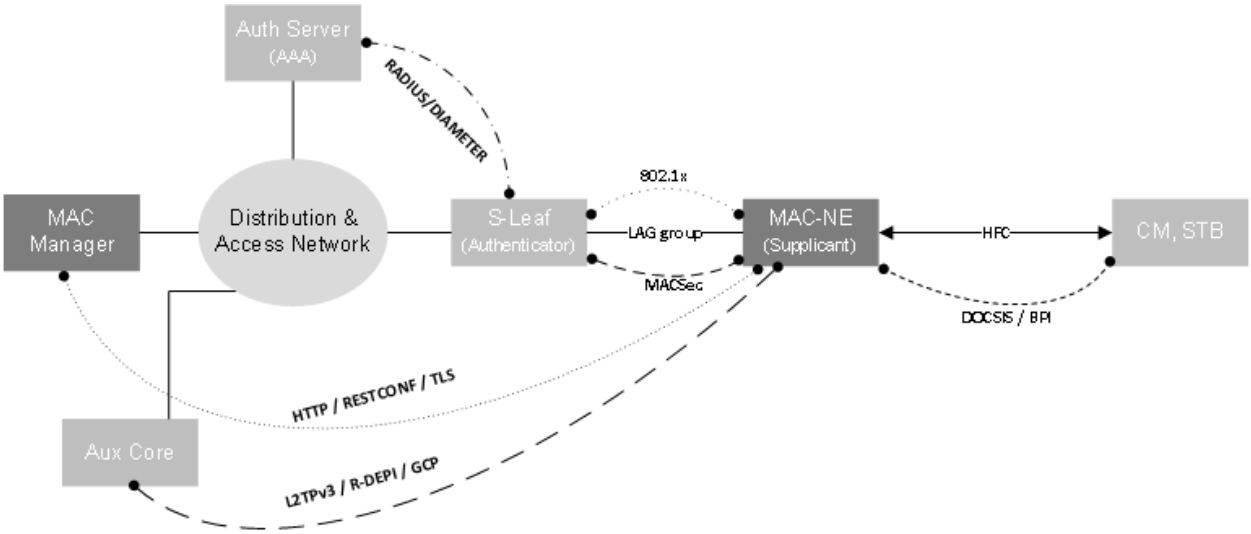

Figure 53 - Remote MACPHY - Trusted Domain and Untrusted Domain

The FMA architecture operates on a model of zero trust. In this model, all communications between devices are mutually authenticated. Whether network authentication is enabled for a given MAC-NE is determined by the configuration of the network rather than by the MAC-NE itself. To support "out of the box" operation, the MAC-NE is expected to attempt to authenticate with the network. If no response to its attempts to authenticate are received, its default mode of probationary allows the MAC-NE to assume network authentication is not supported by the network and attempt to operate without it until prompted to authenticate at some later point.

A MAC-NE MUST support [IEEE 802.1x]-based authentication, which requires the MAC-NE to act as an 802.1x supplicant during Link Layer network initialization. The MAC-NE MUST log event ID 70070129 when it has started EAPOL on any of its ports. The MAC-NE MUST log event ID 70070130 when the EAP-REQ is received on any of its ports. The MAC-NE MUST log event ID 70070131 when it has achieved negotiation of EAP on any of its ports. The MAC-NE MUST log event ID 70070133 when it has reached the maximum number of retries permissible in its attempts to reach the Authentication Server. The MAC-NE MUST log event ID 70070132 if it successfully authenticates with the network. The MAC-NE MUST log event ID 70070134 if it fails to authenticate with the network. Event ID 70070134 would be expected to appear in the logs immediately after event ID 70070133 to make it clear the cause of the failure was caused by unreachability of the Authentication Server.

A fundamental objective aiding in the successful deployment of MAC-NE devices is to avoid requiring prior configuration as much as possible. To achieve "out of the box" operation, the MAC-NE MUST be able to determine if network authentication is enabled, without vendor-specific pre-configuration by reading the [IEEE 802.1x] L2PDUs from its network ports.

When operating in an [IEEE 802.1x] authenticated network, the MAC-NE MUST also attempt to establish a MACsec encrypted link.

13.5.2 802.1x Authentication

Authentication is performed based on the [IEEE 802.1x] and MACsec [IEEE 802.1ae] standards.

802.1x is a Layer 2 protocol that uses EAP (Extensible Authentication Protocol) to provide authentication services.

For the MAC-NE, EAP-TLS is used to convey layer 2 messages that permit basic path validation of digital certificate credentials issued from the DOCSIS PKI.

The [IEEE 802.1x] standard defines three entities involved in the network authentication process, listed in Table 14 below.

Table 14 - 802.1x Network Authentication Entities

Entity	Description
Supplicant	The MAC-NE that requires authentication
Authenticator/NAD	Network element that allows or prevents the MAC-NE from gaining network access until authentication is achieved; also known as a Network Access Device (NAD)
Authentication Server	An Accounting, Authentication and Authorization (AAA) server that orchestrates network authentication that is able to communicate with the supplicant through the authenticator and able to contact the Certificate Authority to check credentials

FMA uses a standard version of the [IEEE 802.1x] protocol with EAP-TLS. This method is referred to as network authentication. The mutual authentication process occurs between the MAC-NE and the Authentication Server without the involvement of the MAC Manager.

Figure 54 shows how the EAP messages between the Authentication Server and the Authenticator are carried over the RADIUS or DIAMETER while EAP messages from the Authentication Server to the MAC-NE are carried over RADIUS/DIAMETER and 802.1x (EAPoL).

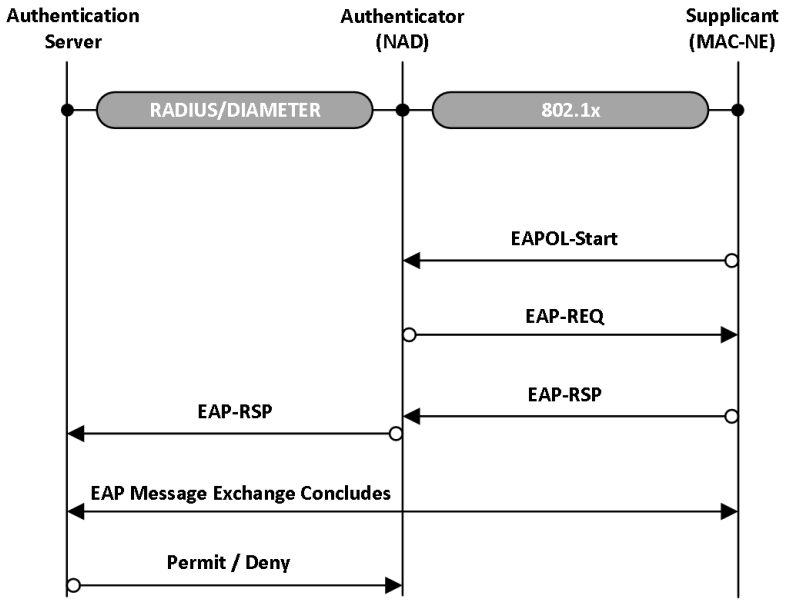

Figure 54 - Network Authentication Signaling

The Authenticator will transmit Layer 2 broadcast EAP Request messages periodically or in response to an EAPoL start message sent from a supplicant. A MAC-NE will respond with a Layer 2 unicast EAP Response. The Authenticator will forward the EAP Response to the Authentication Server using a RADIUS or DIAMETER protocol. The Authentication Server and the MAC-NE then communicate directly, using the Authenticator as a relay agent. When the Authentication Server has made a decision based on MAC-NE authentication and authorization status, it communicates that decision to the Authenticator. The Authenticator will then provide access or deny network access to the MAC-NE depending on its configuration. The MAC-NE also authenticates and validates the Authentication Server before continuing with connection setup.

13.5.2.1 MACsec

MACsec (see [IEEE 802.1ae]) is a Link Layer encryption mechanism that can be initiated at the conclusion of the [IEEE 802.1x] network authentication process between the MAC-NE and the S-Leaf, which doubles as the Authenticator in the network authentication scheme.

Security associations are created between the Authenticator and each MAC-NE based on keying material created during the EAP exchanges. These associations are used for encrypting data between the Authenticator and each MAC-NE, providing a higher level of security than basic 802.1x. With 802.1x, after authentication of a MAC-NE, the Authenticator port is open to receiving any messages from the authenticated MAC-NE's MAC address. This creates the possibility of a device spoofing the MAC-NE MAC address to gain access to the network. With MACsec, only devices in possession of legitimate security keys can send or receive traffic on the network. Therefore, it is recommended that MACsec be used with 802.1x authentication when supported.

The use of MACsec provides the following advantages.

- It enables secure access for multiple devices per port.
- It provides protection against potential man-in-the-middle attacks in both single and multiple devices per port use cases.

MACsec is relatively new and is not yet supported by all network switches. The following requirements apply if MACsec is supported.

- If the MAC-NE supports MACsec, the MAC-NE MUST support the MACsec Key Agreement (MKA) protocol for key exchange and management.

- If the MAC-NE supports MACsec, the MAC-NE MUST derive the Connectivity Association Key (CAK) from the EAP-MSK as defined by EAP-TLS and [IEEE 802.1x].
- If the MAC-NE supports MACsec, the MAC-NE MUST NOT use pre-shared CAKs.

13.5.2.2 MAC-NE Topology Support for 802.1x Network Authentication

Figure 55 shows a number of potential topologies for MAC-NE deployment and connectivity to the NAD. The various topologies are discussed in the following subsections.

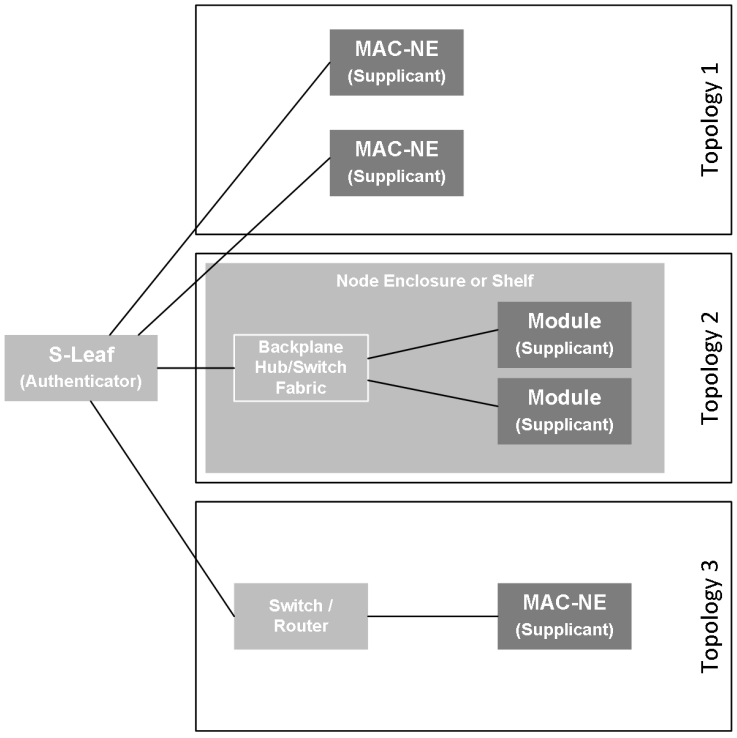

Figure 55 - MAC-NE Topologies for 802.1x Network Authentication

13.5.2.2.1 Topology 1: Single Host per Port

In its most basic form, 802.1x supports access by a single host per Ethernet switch port. This is defined in [IEEE 802.1x] and is widely supported by existing switches.

The MAC-NE MUST support the single host per port topology.

The MAC-NE MUST support [IEEE 802.1x] in a single host per port topology configuration.

The MAC-NE SHOULD support MACsec in a single host per port topology configuration.

13.5.2.2.2 Topology 2: Multiple MAC-NEs in Single Device

In this topology, a single NAD port is connected to an integrated device such as a node enclosure or shelf with multiple MAC-NE modules connecting to the S-Leaf via an internal hub or switching fabric. Thus, the NAD sees multiple devices (and multiple MAC source addresses) on the port. In this topology, MACsec can be used to establish independent security associations with each MAC-NE, based on the MAC-NE MAC address. This is defined in [IEEE 802.1ae] but is not widely supported in Ethernet switches at present.

The MAC-NE can support the topology in which multiple MAC-NEs are in a single device as follows.

- If supported, then each MAC-NE module MUST have a unique MAC address per Ethernet port.
- If supported, the internal hub/switching fabric of the MAC-NE SHOULD propagate the [IEEE 802.1x] EAP-REQ multicast messages to its individual modules.
- If supported, the MAC-NE module SHOULD support [IEEE 802.1x] in this multi-module integrated device configuration.
- If supported, the MAC-NE SHOULD support MACsec in this multi-module integrated device configuration.

13.5.2.2.3 Topology 3: Intermediate External Switch/Router

In this topology, one or more MAC-NEs are connected to the NAD through an intermediate switch or router.

Operators deploying this topology might not be able to utilize 802.1x or MACsec because of the forwarding restrictions on PAE multicast in Ethernet bridges and switches unless measures are taken to tunnel this traffic (as with MPLS over Ethernet). . The definition of the external switch behavior that would be required to support this topology is outside the scope of this specification.

13.5.2.3 Authenticator Location

The Authenticator (NAD) is hosted in the device that provides connectivity to the MAC-NEs. It can be a Layer 2 switch, a Layer 3 router, or even an FMA Core Access Device.

There can be zero or more Layer 2 switches and/or Layer 3 routers between the Authenticator and the FMA Core Access Device. There can be zero or more Layer 2 switches and/or Layer 3 routers between the Authenticator and the Authentication Server.

13.5.2.4 Operation

After powering up (and prior to starting link-layer discovery or obtaining an IP address), the MAC-NE MUST attempt to authenticate itself with the network using [IEEE 802.1x]. See Figure 56 - Network Authentication FSM for more details.

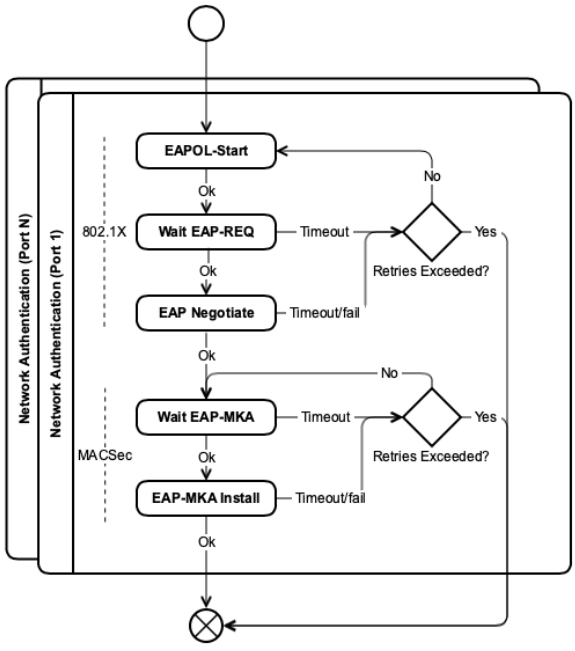

Figure 56 - Network Authentication FSM

The MAC-NE MUST send an EAPOL-START message to the Authenticator and wait for an EAP-REQ. If there is no EAP-REQ in response to the EOPOL-START within "EAP-REQ-TIMEOUT", then the MAC-NE MUST resend the EAPOL-START message and return to wait mode. If no EAP-REQ is received after EAPOL-START-RETRIES have been exhausted, then the MAC-NE MUST assume that the network is not authenticated, operate in a non-authenticated mode, and proceed with the LLDP phase (if supported) or IP address assignment phase of the initialization sequence (this is standard operating procedure for an 802.1x device).

If an EAP-REQ is received, the MAC-NE MUST proceed with [IEEE 802.1x] authentication. If the MAC-NE authentication is rejected, the Authentication Server authentication is rejected, or the authentication process fails after an EAP-REQ has been received, then the MAC-NE MUST wait for the defined [IEEE 802.1x] wait period before trying to reauthenticate.

Once authentication is completed successfully, the MAC-NE MUST attempt to establish a MACsec connection by waiting to for the NAD to broadcast an EAP-MKA. If the MAC-NE does not receive an EAP-MKA within the timeout period, then the MAC-NE MUST proceed to the next initialization phase without establishing MACsec.

13.5.2.5 802.1x Mutual Authentication

802.1x with EAP-TLS provides mutual authentication of the MAC-NE and the Authentication Server. The MAC-NE MUST use EAP-TLS per [RFC 5280] with certificates issued from the DOCSIS PKI managed by CableLabs (see [SECv4.0]). The CableLabs Root CA certificate is installed in the Authentication Server and MAC-NE as a trust anchor for validating received certificates. The MAC-NE Certificate and its private key, along with the issuing intermediate Device CA certificate, are installed in the MAC-NE. The Authentication Server Certificate and its private key, along with the issuing intermediate Service Provider CA certificate, are installed on the Authentication Server. During the EAP-TLS message exchange, the MAC-NE and Authentication Server will send their device/server certificates and the issuing intermediate CA certificate to each other in order to validate them against the root CA trust anchor certificate.

The MAC-NE MUST use the "Basic Path Validation" procedure defined in [RFC 5280] for validating received certificates.

The MAC Manager MUST use the "Basic Path Validation" procedure defined in [RFC 5280] for validating received certificates.

If the MAC-NE has not acquired or maintained the time of day, then it MUST NOT check the validity period expiration of the Authentication Server's certificate chain. If the MAC-NE skipped validity period checking of the Authentication Server's certificate due to not having the time of day, then the MAC-NE MUST check the validity period expiration after it acquires the time of day and periodic re-authentication occurs.

13.5.2.6 Authentication Failures

If an EAP-REQ message has been received by the MAC-NE indicating that authentication is in effect for the network, then any subsequent failures during the mutual authentication process MUST be handled per the [IEEE 802.1x] specification.If an EAP-REQ message has been received by the MAC Manager indicating that authentication is in effect for the network, then any subsequent failures during the mutual authentication process MUST be handled per the [IEEE 802.1x] specification. The wait period timer (which defines the time a device waits after a failed authentication attempt before another attempt is permitted) SHOULD NOT be reduced below the 60 second default time.

Retransmission behavior for EAP messages sent or received by the MAC-NE (which are forwarded from the authenticator to the authentication server) MUST follow [RFC 3748]. Retransmission behavior for EAP messages sent or received by the MAC Manager (which are forwarded from the authenticator to the authentication server) MUST follow [RFC 3748].

Table 15 - EAP Authentication Timing

Constant	Value
EAP-REQ-TIMEOUT	10 sec
EAPOL-START-RETRIES	3

13.6 Link Layer Discovery

13.6.1 LLDP Support

The MAC-NE SHOULD support LLDP [IEEE 802.1ab].

If implemented, the MAC-NE MUST include TLVs 0-8 in the LLDP PDU per [IEEE 802.1ab].

13.6.2 Operation

Upon completing [IEEE 802.1x] network authentication, or upon determining that the network is not authenticated, the MAC-NE SHOULD begin transmitting and listening for LLDP PDUs on its CIN-facing interface(s). See Figure 57 for a depiction of the LLDP finite state machine.

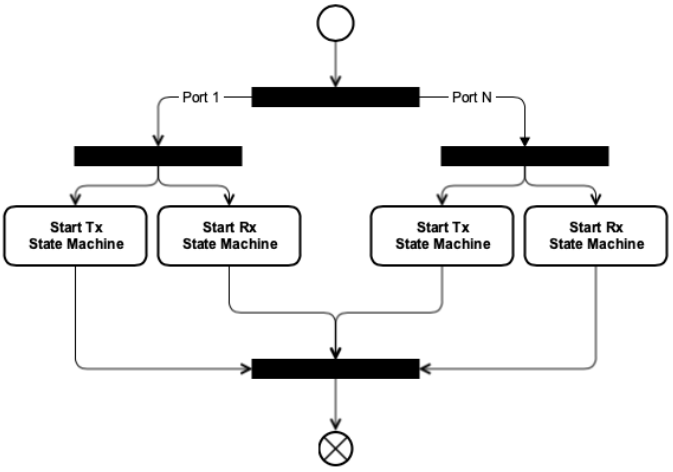

Figure 57 - Start LLDP FSM

13.7 Link Aggregation (LAG)

Link Aggregation is specified in [IEEE 802.1AX] to consider a group of Ethernet ports on an MAC-NE as a "Link Aggregation Group" (LAG) and operate as a single Network System Interface (NSI) of the MAC-NE.

The MAC-NE MUST support all "aggregated system" requirements of [IEEE 802.1AX].

The MAC-NE MUST support static Link Aggregation per [IEEE 802.1AX].

If it implements multiple Ethernet ports connected to an S-Leaf, the MAC-NE MUST support [IEEE 802.1AX] link aggregation of those ports and consider the aggregation to be the network-side interface (NSI) of the MAC-NE. The MAC-NE attempts to establish link aggregation and obtain an IP address for the LAG interface before attempting to obtain an IP address for any individual port of the LAG.

[IEEE 802.1AX] requires dispatching traffic for the same "conversation" to the same port, where the conversation is determined by vendor-specific hashing algorithms. Because the MAC-NE operates as a Layer 2 bridge, the MAC-NE MUST implement a hashing function that includes fields above Layer 2, e.g., source and destination IP addresses. Note that each MAC or IP address is incorporated in a hash key which a device uses to select the outbound interfaces over which the device will forward. This mechanism provides assurances that frames belonging to the same L2/L3 conversation are forwarded on the same physical link and are distributed among multiple links. The specifics of the hashing and distribution algorithm are left to vendor implementation detail.

The MAC-NE MUST ignore the LAG TLV (127) of LLDP.

Figure 58 depicts the operation of two MAC-NEs with LAG groups.

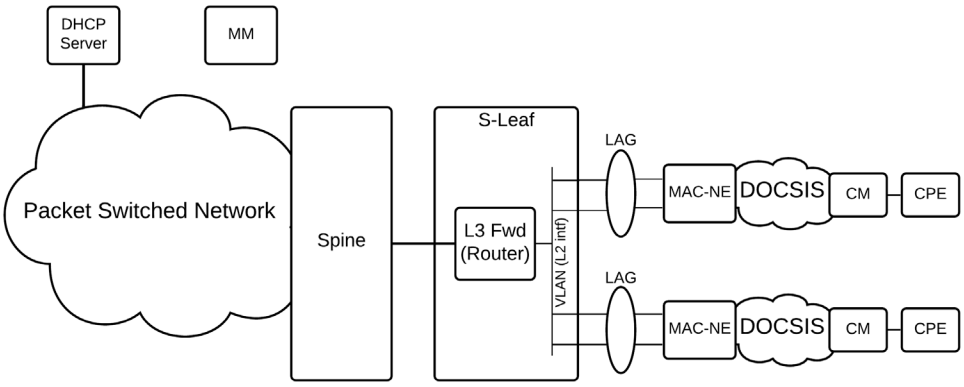

Figure 58 - LAG Groups

13.7.1 Link Aggregation Control Protocol (LACP)

The Link Aggregation Control Protocol provides a dynamic means of controlling the aggregation of several physical Ethernet ports of a device (i.e., network element) into one or more Ethernet links.

The MAC-NE MUST support dynamic Link Aggregation using the Link Aggregation Control Protocol (LACP) [IEEE 802.1AX].

LACP allows directly connected network peers to negotiate the automatic bundling of available links by sending control packets to one another periodically. LACP allows peers to respond to the presence of new links and react to link faults automatically, with the actions taken by directly connected peers dependent upon policies controlled by the mode of operation and the distribution algorithm configured for a given LAG.

LACP operates by sending frames (LACPDUs) between directly connected devices across all links on which LACP is enabled. If it finds a device on the other end of the link that responds to the LACPDUs, both devices will detect that multiple links exist between them and will combine these links into a single logical LAG interface.

LACP LAGs can be configured in one of two modes: active or passive mode. In active mode, the device will always send frames across links where LACP is enabled. In passive mode, the device will not send LACPDUs unless it receives them from the peer device first. This mode can be used as a method for controlling accidental loops when the other device on the link is operating in active mode.

The MAC-NE MUST support LACP active and passive modes of operation per [IEEE 802.1AX].

The MAC-NE MUST use the active LACP mode per [IEEE 802.1AX] by default for all LAGs.

Basic concepts in LACP include:

- Link aggregation provides a configurable maximum count for the number of bundled ports that the peer supports via the LACP message exchanges. Valid port counts of between 1 and 8 ports are supported. If the peers have different supported maximums (say one device supports 4 ports and the other 8 ports) the lowest common denominator between them is used in building their link aggregations.
- The multicast L2 address 01:80:c2:00:00:02 (01-80-c2-00-00-02) is used for LACP packets between peers. LACP supporting devices are required to join this multicast group during interface startup in order to participate in LACP exchanges.
- The LACP detection phase follows the following process:
  - LACP packets are transmitted once every second between peers until a response is received and steady-state is achieved.

  - The keep-alive timer is started and keep-alive messages are sent to a link member periodically after reaching steady-state. By default, these keep-alive messages are sent once every 30 seconds, though the range allows for a period down to one second.
  - Link-ID: an integer identifying the member link for the purposes of load balancing. The range is from 1 to 8;
  - The distribution mode can be configured based on traffic models or device capabilities.
  - A device configured for Active LACP enables LACP unconditionally on a LAG.
  - A device configured for Passive LACP enables LACP on the LAG when LACP messages are detected.
- The LACP "system priority" parameter determines the sequence in which devices at both ends of a link aggregation select active interfaces to join a LAG. In order for a LAG to be established, both devices are required to select the same interfaces as active interfaces. To achieve this, the device with the higher priority is responsible for selecting the active interfaces. The other device with a lower priority then selects the same active interfaces. In priority comparisons, numerically lower values receive higher priority.
- The LACP per-link "interface priority" parameter affects which interfaces of a link aggregation are selected as active interfaces. A smaller numerical value represents a higher priority. The interfaces with the highest LACP interface priority become active interfaces.

Unlike in static LAG configuration, failover can occur automatically in LACP-controlled LAGs. When a link that is a member of a LACP-controlled LAG fails, the peer device will continue sending and receiving traffic without error by bypassing the failed link. With static link aggregation, the peer is often unaware of the link fault and continues to send traffic over the failed link which results in potentially significant data flow disruption.

Another advantage provided by LACP-controlled LAGs over static configured LAGs is that the device can negotiate with the peer at the other end of the connection to confirm that it can perform link aggregation and configure the LAG with the appropriate parameters. With static configuration, errors made in configuring modes, policies or even cabling problems between peers could go undetected, resulting in undesirable network behavior.

13.7.2 Distributing Utilization Across Links and Frame Ordering Considerations

Similar load distribution and high utilization of all ports combined in the link aggregation are goals that can prove difficult to achieve in real-world use cases, though certain distribution schemes provide greater benefits than others. In worst-case scenarios, the distribution method chosen can only accomplish prevention of out-of-order transmission while providing little or no distribution of traffic load among members of the link aggregation.

When distributing traffic, the reordering of Ethernet frames in the same "conversation" between host protocol entities is avoided to prevent impacts on higher layer transport protocols and applications. TCP and other connection-oriented protocols can suffer significant overhead when contending with out-of-order packet transmissions. Link aggregation relies on hashing fields of an Ethernet frame to identify a conversation. All frames associated with a particular hash value are sent across the same LAG link, preventing out-of-order packet delivery across the LAG for the conversation. Common hash algorithms of link aggregation use only L2 or L2 and L3 tuples of source/destination pairs to create their hash keys (e.g., hashes based on MAC address and IP address information). The goal of the hash key is to ensure that a given conversation is always sent via the same physical link while providing sufficient dispersion of conversations across different links to provide similar loading on each link.

Unfortunately, the importance of avoiding out-of-order delivery negatively impacts the similar loading of packets across the links in the aggregation group. Achieving similar loading of traffic becomes especially problematic when only a single or a very small number of IP hosts are communicating with one another across the aggregation. This can occur when a large fraction of traffic is embedded within IP tunnels. A lack of distribution in the hashes used by the sender causes the frames to be transmitted on the same link even when additional links in the aggregation are available and underutilized. This can dramatically reduce the effective bandwidth available in the link aggregation.

More advanced LAG implementations employ a hash that contains L4 tuples that further aid in the prevention of the aforementioned link aggregation asymmetries by adding port numbers (e.g., source and destination TCP/UDP port

numbers) to the L2 and L3 tuples. This effectively increases variation in hash keys that results in a more similar loading of traffic across the links in the aggregation.

Note that L2TPv3 traffic to a MAC-NE is encapsulated in IP tunnels with the same L2 and L3 addresses. In the absence of L4 fields included in the distribution hashing algorithm, all such traffic will traverse the same Ethernet link. It is desirable for an S-Leaf to include the L2TPv3 session ID (an L4 field) in its downstream LAG hashing algorithm.

An S-Leaf is expected to hash on at least the same combination. Configuration of MAC-NE LAG hashing algorithms beyond this is vendor-specific. A MAC-NE is not required to implement L4 fields in its upstream LAG hash algorithm.

A MAC-NE is not required to implement LACP version 2 or optional TLVs of LACP [IEEE 802.1AX].

General recommendations:

- It is highly recommended that all members of a LAG bundle be of the same physical distance between peers. This is to prevent asymmetric latency of bidirectional flows across the bundle.

13.7.3 LAG and LACP Events

This section describes the events that the MAC-NE logs whenever the configuration is changed, or an operational event has occurred.

The MAC-NE MUST log event ID 70070307 when its LACP Activity mode is set to active on a given port.

The MAC-NE MUST log event ID 70070308 when its LACP Activity mode is set to passive on a given port.

The MAC-NE MUST log event ID 70070309 when the LACP Transmission Rate is set to Fast on a given port.

The MAC-NE MUST log event ID 700703010 when the LACP Transmission Rate is set to Slow on a given port.

The MAC-NE MUST log event ID 700703011 when LACP Need To Transmit is forced by configuration change on a given port.

The MAC-NE MUST log event ID 70070312 when the remote end does not respond to LACP frames within the short (3 sec) timeout.

The MAC-NE MUST log event ID 70070313 when the remote end does not respond to LACP frames within the long (90 sec) timeout.

The MAC-NE MUST log event ID 70070314 when it successfully constructs a LAG using LACP-enabled ports.

The MAC-NE MUST log event ID 70070315 when a LAG group is activated and in an ‘Up’ state.

The MAC-NE MUST log event ID 70070316 when a LAG fault occurs, and a link is removed.

The MAC-NE MUST log event ID 70070317 when link membership in the LAG changes.

The MAC-NE MUST log event ID 70070318 when a link in the LAG has remained idle for longer than vendor-specific timeout value.

13.8 Address Assignment

Once the MAC-NE is successfully authenticated to the network, the MAC-NE MUST obtain a Network Layer IP address. The MAC-NE MUST support IPv4 and IPv6 addressing of its IP host. See Figure 59 for a depiction of the IP address finite state machine.

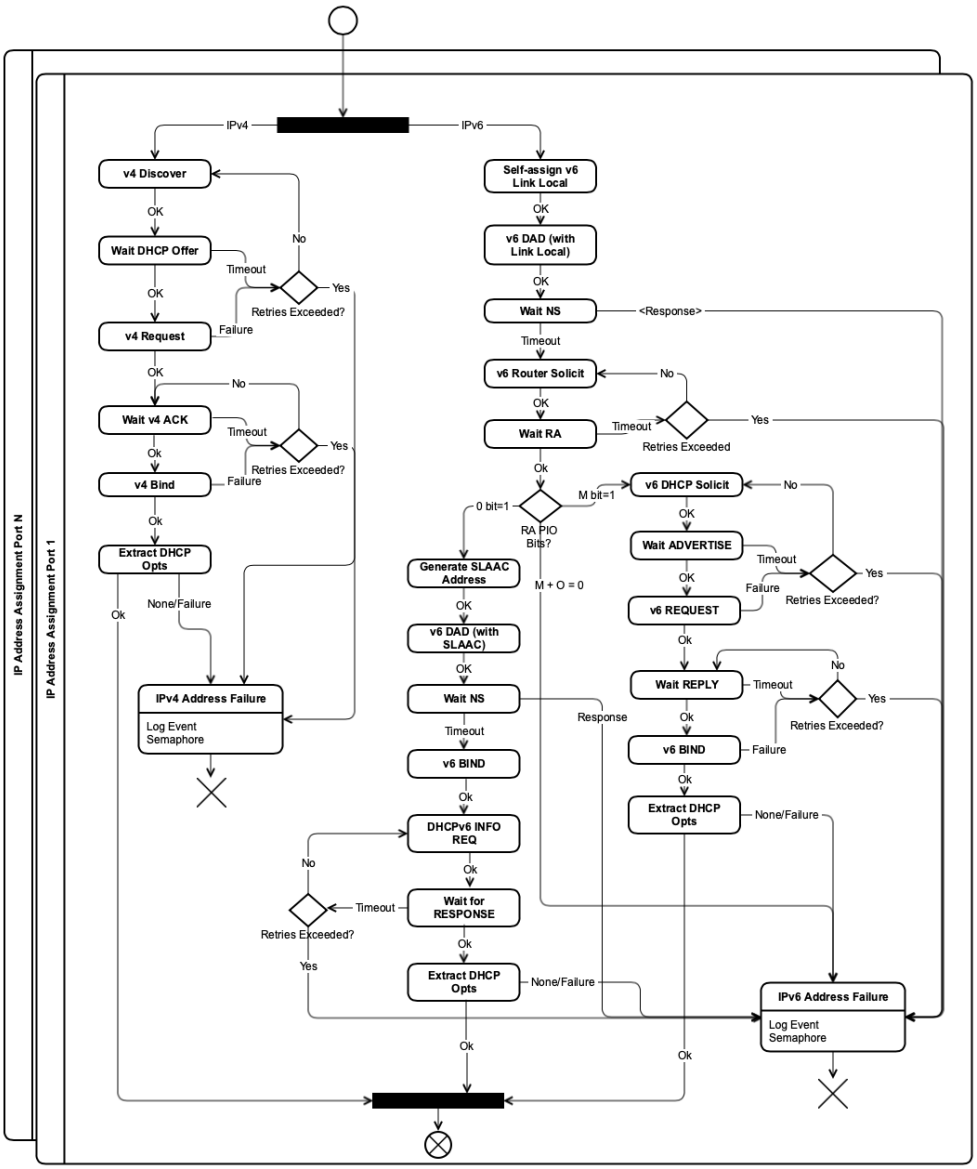

Figure 59 - IP Address Assignment FSM

The MAC-NE MUST be able to operate in single stack IPv6 mode.

The MAC-NE MUST be able to operate in single stack IPv4 mode.

The MAC-NE MUST be able to operate in dual stack (IPv4 + IPv6) mode.

The MAC-NE MAY support static configuration of IP addresses. The network cannot participate in static address assignment, and the mechanisms for configuring or storing static IP addresses are left as an implementation detail.

The MAC-NE MUST support DHCPv4 per [RFC 2131].

The MAC-NE MUST support DHCPv6 per [RFC 8415] and [RFC 4861].

The MAC-NE MUST support IPv6 address assignment via SLAAC [RFC 4862], with options provided via Stateless DHCPv6 [RFC 8415].

The MAC-NE MUST NOT utilize Privacy Extensions for Stateless Address Autoconfiguration in IPv6 [RFC 4941].

The MAC-NE mode of operation (v6, v4, or dual stack) is determined by the responses it receives from the network (such as v6 router advertisements) and DHCP servers (refer to [RFC 8415] and [RFC 4861] for details) and is thus controlled by the operator.

IPv6 address assignment of the MAC-NE is governed by the configuration bits set in the ICMPv6 Router Advertisement (RA) message and the presence of a valid prefix in the Prefix Information Option (PIO).

The MAC-NE MUST set the M bit and O bit properly to use either method for IPv6 addressing (DHCPv6 or SLAAC).

For IPv6, the MAC-NE MUST wait for an RA message containing the required information to proceed based on values set for the M bit and O bit in the RA in Table 16 - Router Advertisement M Bit and O Bit Settings For SLAAC.

Table 16 - Router Advertisement M Bit and O Bit Settings For SLAAC

	M Bit=0	M Bit=1
O Bit=0	This combination of bits is unsupported. The MAC-NE will ignore this combination of bits.	O-bit ignored. MAC-NE uses Stateful DHCPv6 for addressing and options.
O Bit=1	MAC-NE uses SLAAC for addressing and Stateless DHCPv6 for options.	O-bit ignored. MAC-NE uses Stateful DHCPv6 for addressing and options.

The MAC-NE is directed by the RA to use either SLAAC [RFC 4862] with Stateless DHCPv6 INFORMATION-REQUEST messages [RFC 8415] or Stateful DHCPv6 [RFC 8415].

The length of time the MAC-NE waits for a properly formed RA before declaring a failure to acquire an IPv6 address and the actions that the MAC-NE needs to take when such a failure occurs is relegated to implementation detail. The method used by the MAC-NE to select a prefix if several prefixes are presented in the PIO of the RA message is also left to implementation detail.

If the MAC-NE is instructed via RA to utilize SLAAC for address assignment, the IPv6 address MUST be constructed using the Modified-EUI method as described in [RFC 4291]. For example, if the PIO advertised in the RA was 2001:db8:a000:1234::/64 and the MAC-NE MAC address was 10:00:00:12:34:56, then the Modified EUI-64 Address would be 2001:db8:a000:1234:1200:00ff:fe12:3456.

If the MAC-NE obtains no IPv4 or IPv6 address from the network, then the MAC-NE MUST retry DHCP per [RFC 2131], [RFC 8415], and [RFC 4861].

If the MAC-NE obtains an IPv6 address but not an IPv4 address, then the MAC-NE MUST attempt to contact the MAC Managers specified in the CCAP Core option encoding learned from the DHCPv6 server using IPv6.

If the MAC-NE obtains an IPv6 address but not an IPv4 address, then the MAC-NE MUST retry DHCPv4 per [RFC 2131] unless instructed by the MAC Manager to use only IPv6 for all future communications.

If the MAC-NE obtains an IPv4 address but not an IPv6 address and IPv6 is enabled on the network, then the MAC-NE MUST attempt to contact the MAC Managers specified in the CCAP Core option encoding learned from the DHCPv4 server using IPv4.

If the MAC-NE obtains an IPv4 address but not an IPv6 address and IPv6 is enabled on the network, then the MAC-NE MUST retry DHCPv6 per [RFC 4861] unless instructed by the MAC Manager to use only IPv4.

If the MAC-NE obtains both an IPv4 address and an IPv6 address, then the MAC-NE MUST attempt to contact any MAC Managers specified in the DHCP suboption encoding learned from the DHCPv4 and DHCPv6 servers starting with IPv6 servers first. If the MAC-NE obtains both an IPv4 address and an IPv6 address, then the MAC-NE MUST operate in dual stack mode and accept IP packets to either address until instructed otherwise by the MAC Manager.

The entire address assignment of the MAC-NE can take place without the involvement of any FMA core component. An FMA core component such as a MAC Manager or an FMA Core Access Device MAY run a DHCP Relay agent but is not required to if relay is provided by another network component. Figure 60 shows a simplified access network containing a MAC-NE, FMA core components, and a DHCP Server.

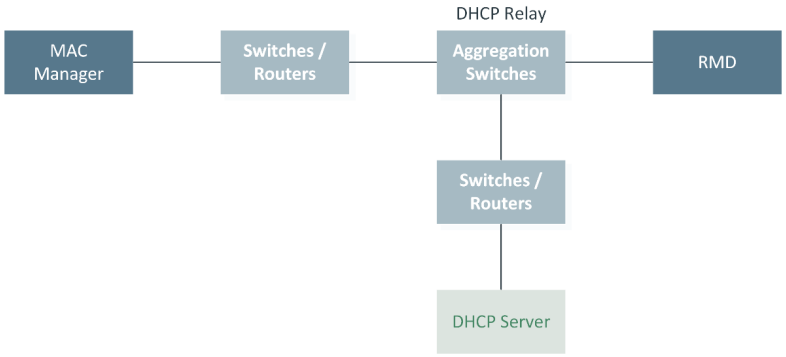

Figure 60 - DHCP Network Diagram

There may be zero or more Layer 2 switches between the MAC-NE and the DHCP Relay. The DHCP Relay may be hosted by a Layer 2 switch, a Layer 3 router, or an FMA core component. There may be zero or more Layer 2 switches and/or Layer 3 routers between the network element that is hosting the DHCP Relay and the FMA Core Access Devices. There may be zero or more Layer 2 switches and/or Layer 3 routers between the DHCP Relay and the DHCP Server.

Figure 61 shows the DHCPv4 signaling protocol, while Figure 62 shows the DHCPv6 signaling protocol. The MAC-NE issues a broadcast DHCPv4 discovery message when it needs to obtain an IP address. The DHCP agent responds with a unicast DHCP offer that contains the IPv4 address of one or more DHCP servers. The MAC-NE picks one of the DHCP servers and sends a DHCP request to it. The DHCP server sends a DHCP acknowledgement with an IP address for the MAC-NE device.

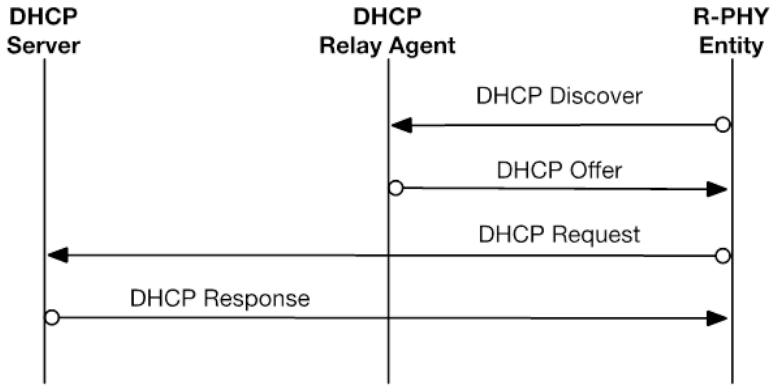

Figure 61 - DHCPv4 Signaling

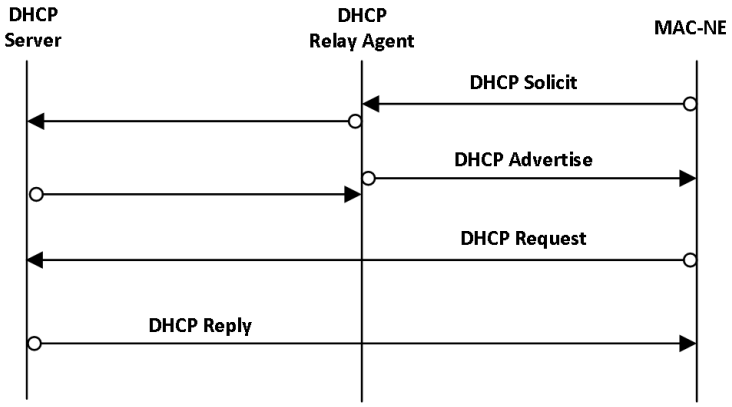

Figure 62 - DHCPv6 Signaling

Unlike the DOCSIS DHCP process, in which the CCAP is always the DHCP relay agent and can always snoop and append to the DHCP messages, the MAC Manager generally does not have any direct access to the DHCP message exchange and thus will not necessarily be aware of the IP address assignments given to MAC-NEs. The following mechanisms can be used to create an association between the MAC Manager and the MAC-NE.

- The DHCP Server is configured with the IP address, MAC address, and/or DUID of the MAC-NE. The DHCP Server then provides the IP address of the MAC Manager to the MAC-NE. Two additional DHCP options <MAC Manager IP> and <MAC Manager DN> (refer to Section 13.8.2) are added in FMA to support this mechanism.
- The MAC Manager can be configured with the IP address of the MAC-NE if the operator wishes. The MAC Manager can then directly contact the MAC-NE without the need of the DHCP options and without the MAC-NE announcing its presence to the MAC Manager.
- The MAC Manager MUST either accept the connection from any MAC-NE, deny the connection, or redirect the MAC-NE to another MAC Manager.

13.8.1 DHCP Options

Refer to [CANN] for details on specific options.

When the MAC-NE issues a request to the DHCPv4 server, it MUST populate *Option 60: Class Identifier* with "RMD".

When the MAC-NE issues a request to the DHCPv6 server, it MUST populate *Option 16: Vendor Class* with "RMD".

When the MAC-NE issues a request to the DHCPv4 server, it MUST populate the CableLabs suboptions from Table 17 - FMA DHCP Options under DHCPv4 *Option 43*.

When the MAC-NE issues a request to the DHCPv6 server, it MUST populate the CableLabs suboptions from Table 17 - FMA DHCP Options under DHCPv6 *Option 17*.

Table 17 - FMA DHCP Options

Suboption	Value	Use
2	<Device Type>	Set to "RMD"
3	<ECM: eSAFE>	Not used
4	<serial number>	Refer to [CANN]
5	<hw version>	Refer to [CANN]
6	<sw version>	Refer to [CANN]
7	<Boot ROM version>	Refer to [CANN]

Suboption	Value	Use
8	<OUI>	Refer to [CANN]
9	<Model Number>	Refer to [CANN]
10	<Vendor Name>	Refer to [CANN]

If the MAC-NE issued a DHCPv4 request, it expects to receive a DHCPv4 response message containing DHCPv4 options defined in [RFC 2132] and [CANN]; the MAC-NE reads the options to determine certain configuration settings. If the MAC-NE issued a DHCPv6 request, it expects to receive a DHCPv6 response message containing DHCPv6 options defined in [RFC 8415] and [CANN]; the MAC-NE reads the options to determine certain configuration settings.

- The MAC-NE MUST support use of DHCPv4 *Option 2: Time Offset* when setting its local clock.
- The MAC-NE MUST support use of DHCPv4 *Option 4: Time Server* to receive time of day servers during that initialization step.
- The MAC-NE MUST support DHCPv4 *Option 7: Log Server* to send log messages over the Syslog protocol to a remote server.
- The MAC-NE MUST support use of DHCPv4 *Option 6: Domain Server* when resolving domain names to IP addresses.
- The MAC-NE MUST support use of DHCPv4 *Option 15: Domain Name* to determine the domain name when resolving hostnames.
- The MAC-NE MUST support DHCPv4 *Option 43, Suboption 61:CCAP Cores* list of IP addresses, which will contain a list of 0 to N IP addresses, in priority order, of MAC Managers to contact.
- The MAC-NE MUST support DHCPv4 *Option 43, Suboption 62:MAC Manager Domain Names*, which will contain a list of 0 to N domain names, in priority order, of MAC Managers to contact.
- If the MAC-NE issued a DHCPv6 request, it expects to receive a DHCPv6 response message containing DHCPv6 options defined in [RFC 8415], [RFC 3646], [RFC 4704], and [CANN]; the MAC-NE reads the options to determine certain configuration settings.
- The MAC-NE MUST support DHCPv6 *Option 17, Suboption 34: Syslog Servers* list of IP addresses, which will contain a list of 0 to N IP addresses, in priority order, of syslog remote servers.
- The MAC-NE MUST support use of DHCPv6 *Option 17, Suboption 38: Time Offset* when settings its local clock.
- The MAC-NE MUST support use of DHCPv6 *Option 17, Suboption 37: RFC868 Servers* list of IP addresses, which will contain a list of 0 to N IP addresses, in priority order, of RFC868 Time of Day Servers to use during that initialization step.
- The MAC-NE MUST support DHCPv6 *Option 23: DNS Servers*, which will contain a list of 0 to N IP addresses, in priority order, of DNS Servers to use for domain name resolution.
- The MAC-NE MUST support DHCPv6 *Option 39: Client FQDN* to determine the domain name when resolving hostnames.
- The MAC-NE MUST support DHCPv6 *Option 17, Suboption 61: CCAP Cores* list of IP addresses, which will contain a list of 0 to N IP addresses, in priority order, of MAC Managers to contact.
- The MAC-NE MUST support DHCPv6 *Option 17, Suboption 62: MAC Manager Domain Names*, which will contain a list of 0 to N domain names, in priority order, of MAC Managers to contact.

13.8.2 MAC Manager Suboptions

The MAC Manager IPs suboption describes either IPv4 or IPv6 addresses, as shown in Figure 63 and Figure 64. The MAC Manager IPs suboption can also be used with DHCPv6.

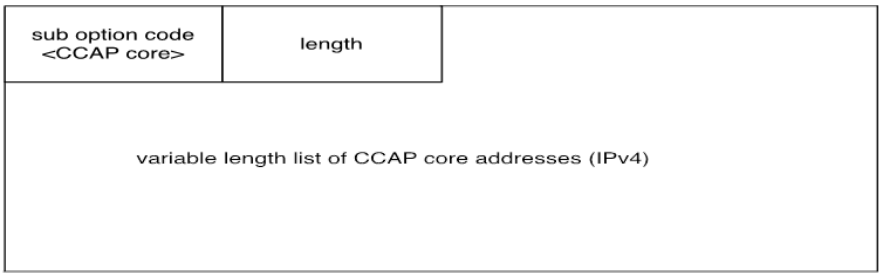

Figure 63 - MAC Manager DHCP Suboption IPv4

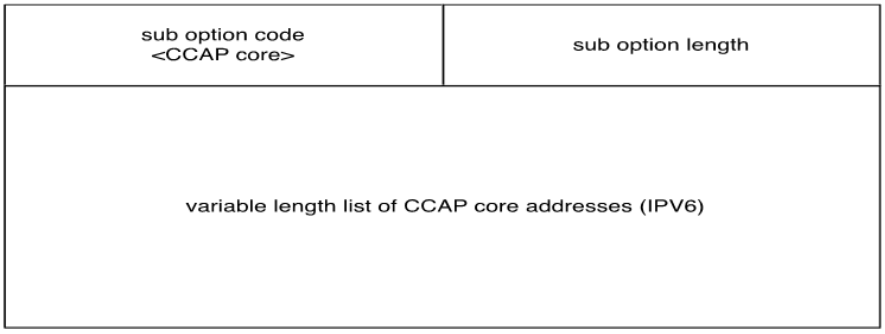

Figure 64 - MAC Manager DHCP Suboption IPv6

The MAC-NE MUST attempt to connect to all MAC Managers in the DHCP options list.

13.8.3 Failures

The MAC-NE MUST respond to any errors during the DHCP process as per [RFC 2131].

Note: This results in the MAC-NE entering a time out and retry loop with a randomized exponential back off.

13.8.4 DNS Support

The MAC-NE MUST support DHCP provisioned DNS. The MAC-NE uses DNS for resolving domain names embedded in the configuration attributes received from the MAC Manager, among other uses. The MAC Manager MUST support a DNS client compliant with [RFC 1034], [RFC 1035], [RFC 1886], and [RFC 2874].

13.8.5 Security Implications

The MAC-NE MUST attempt to contact the DHCP server via the CIN interface. If 802.1x and MACsec are in place, this will provide secure access to the trusted network.

13.9 DNS SRV Discovery Method of Operation

As an alternative to DHCP-provided MAC Manager entities, the MAC-NE can use DNS SRV Discovery. This is intended as an optional discovery mechanism when DHCP directed discovery is not used by an operator. DNS SRV records are used with service discovery for multiple protocols and services. SRV records are symbolic names in the zone records that include transport protocol and port, used in combination with the domain name associated with the MAC-NE.

The MAC-NE MAY use DNS SRV records to discover MAC Manager entities only if no DHCP-directed (see MAC Manager Suboptions section of this document) addresses are present. If the MAC-NE uses DNS SRV, the MAC-NE MUST use DNS SRV record priority and weight when establishing a preferred MAC Manager connection request.

The DNS SRV resolved hostnames MUST be placed into the unified MAC Manager list in priority order for use later in the initialization.

When using DNS SRV, the MAC-NE MUST use the *service* "macmanager" and *protocol* "tcp".

Format of SRV Resource Records:

_service._protocol.domainName. TTL class SRV priority weight port target

Example SRV record for MAC Manager discovery:

_macmanager._tcp.example.comn. 3600 IN SRV 10 60 443 manager.example.com.

_macmanager._tcp.example.comn. 3600 IN SRV 10 10 443 manager1.example.com.

In the above example, both SRV records have an equal priority of 10 and a different weight value 60 and 10, respectively. A higher weight value within the same priority will be preferred over lower weight within the same priority group.

13.10 SSH

The MAC-NE starts an SSHv2 server listening on the IP management host after receiving IP address configuration. The MAC-NE will automatically re-initialize if it is unable to start an SSHv2 server. Details on SSH configuration and operation are in Section 12.5 of this document. See Figure 65 for a depiction of the SSH finite state machine.

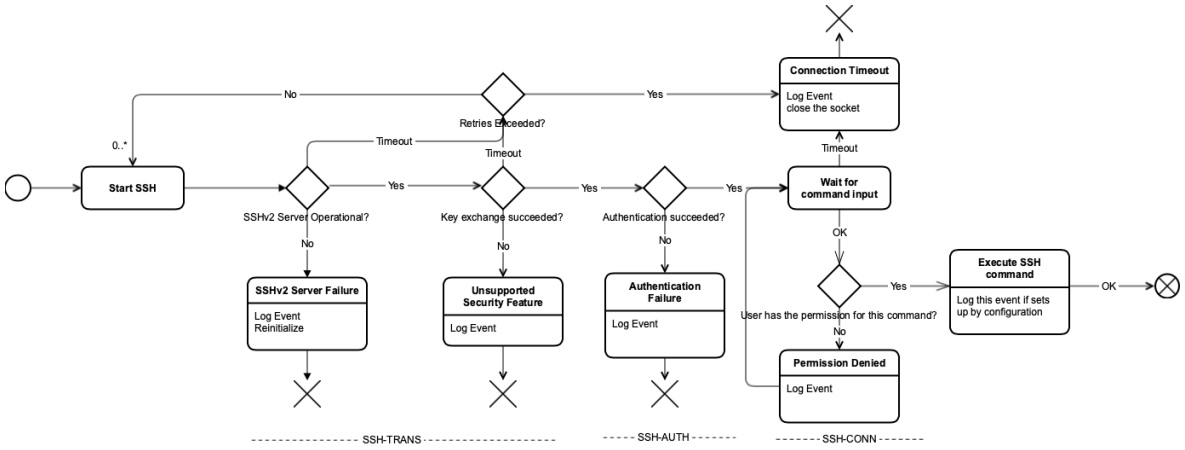

Figure 65 - Start SSH FSM

13.11 Get ToD (Time of Day)

The MAC-NE acquires the ToD (Time of Day) for the purpose of timestamping warnings, error logs and messages, validation of the CVC during a software upgrade, and validation of the MAC Manager and Auxiliary Core certificates during mutual authentication.

13.11.1 ToD Acquisition

The MAC-NE MUST attempt to obtain the current date and time by using the Time Protocol (see [RFC 868]) from one of the servers listed in the Time Server Option DHCP field. If this field is missing or invalid, the MAC-NE MUST initialize the current time to Jan 1, 1970, 0h00. If the time is initialized (reset), the MAC-NE MUST ignore the value, if any, of the Time Offset DHCP option.

The MAC-NE MUST use its SLAAC-acquired or DHCP-provided IP address for exchange of messages with the Time Protocol server. The MAC-NE MUST transmit the Time Protocol request using UDP. The MAC-NE MUST listen for the Time Protocol response on the same UDP port as is used to transmit the request. The MAC-NE MUST combine the time retrieved from the server (which is UTC) with the time offset received from the DHCP server to create a valid "local" time. See Figure 66 for a depiction of the Time-Of-Day finite state machine.

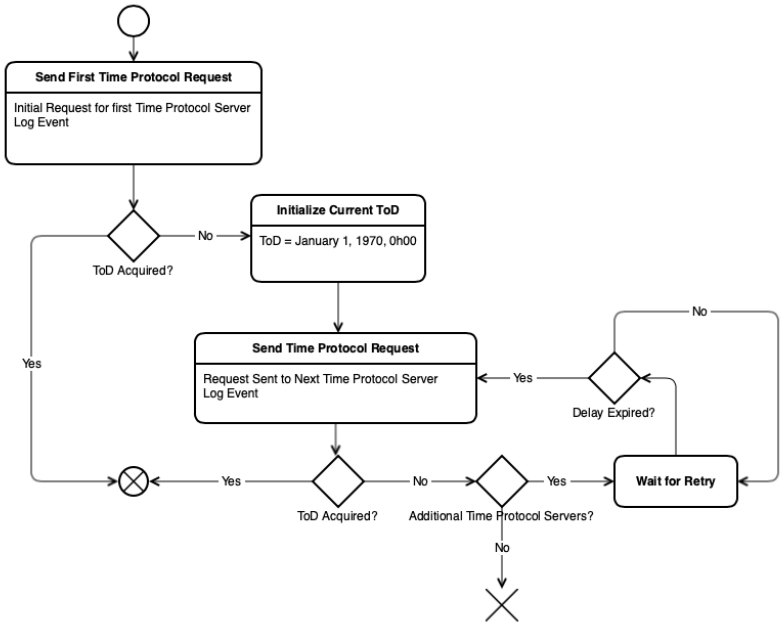

Figure 66 - Get ToD FSM

Once the MAC-NE acquires ToD, the MAC-NE MUST stop requesting, unless any of its ToD-related parameters (such as time offset or server address) are modified. If the MAC-NE's ToD-related parameters are modified, the MAC-NE MAY re-request ToD from the Time Protocol server(s).

When the MAC-NE successfully establishes its initial ToD, the MAC-NE SHOULD log event ID 70070402 in the local log.

13.11.2 ToD Conflicts and Problems

The DHCP server can return multiple IP addresses from multiple Time Protocol servers. The MAC-NE MUST attempt to obtain ToD from all the servers listed until it receives a valid response from any of the servers. The MAC-NE MUST contact the servers in batches of tries with each batch consisting of one try per server and each successive try within a batch at most one second later than the previous try and in the order listed by the DHCP message. If the MAC-NE fails to acquire ToD after any batch of tries, the MAC-NE MUST retry a similar batch using a truncated randomized binary exponential backoff with an initial backoff of 1 second and a maximum backoff of 256 seconds.

For each Time Protocol query with no response, the MAC-NE MUST log event ID 70070403 in the local log. For each Time Protocol response with an invalid data format, the MAC-NE MUST log event ID 70070404 in the local log.

If the MAC-NE does not obtain ToD in the initial request against the first server, the MAC-NE MUST initialize the current time to Jan 1, 1970, 0h00, and then subsequently initialize its current time once it receives a response from a Time Server.

If the MAC-NE fails to establish ToD, then it will not be able to validate the MAC Manager certificate. The failure may be due to Time Server failure or to an error in the DHCP option list.

The MAC-NE will follow the same mechanism as for a DHCP failure and restart the DHCP process as described in Section 13.8.

13.11.3 ToD Security Implications

The MAC-NE MUST attempt to contact the Time Server via the CIN-facing port on which the DHCP response was received. If 802.1x and MACsec are in place, this will provide secure access to the trusted network.

13.12 Apply Saved Configuration

The MAC-NE can have saved configuration stored across reboots in non-volatile memory. The MAC-NE MUST apply any saved configuration from non-volatile memory before the MAC-NE announces itself to the MAC Manager. If the MAC-NE does not have a saved configuration or if an error is detected when applying the saved configuration, the default configuration actions are applied. Additionally, vendor-specific configuration actions are also applied after the saved configuration and/or default configuration are applied.

The MAC-NE MUST log event ID 70070005 if the MAC-NE is unable to apply non-volatile configuration settings during local device initialization.

The MAC-NE MUST log event ID 70070006 if the MAC-NE is unable to apply default configuration settings during local device initialization.

The MAC-NE MUST log event ID 70070007 if the MAC-NE is unable to apply vendor-specific configuration settings during local device initialization.

See Figure 67 for a depiction of the finite state machine used during the application of saved configuration in the MAC-NE.

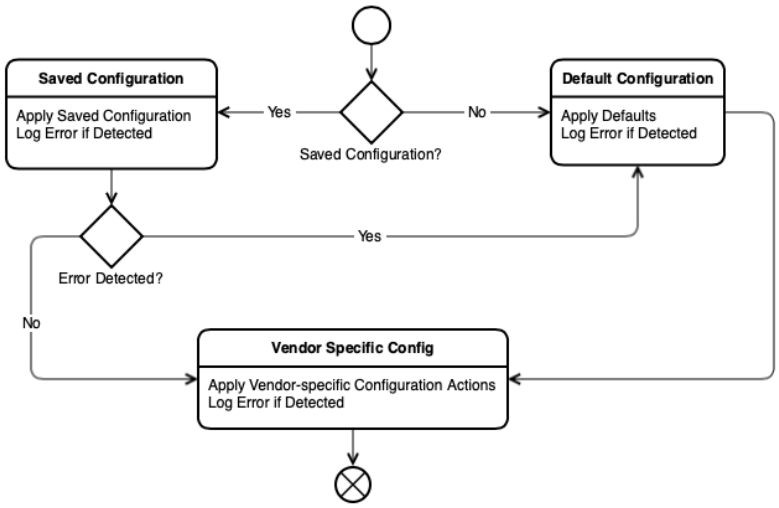

Figure 67 - Apply Saved Configuration FSM

13.13 Listen for RESTCONF Connections

After the MAC-NE has completed initialization and applied any locally saved configuration, it needs to listen for unsolicited RESTCONF connections. The MAC-NE MUST start a RESTCONF listener before it announces itself to a MAC Manager. See Figure 68 for a depiction of the finite state machine used by the MAC-NE during RESTCONF startup.

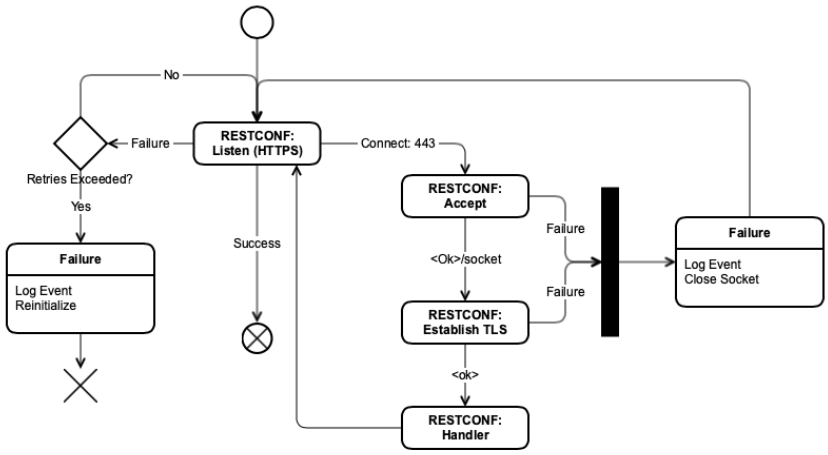

Figure 68 - Listen for RESTCONF FSM

13.13.1 Start RESTCONF Listener

The MAC-NE MUST start a RESTCONF server before announcing itself to the MAC Manager. The details of the RESTCONF configuration and operational models are documented in [FMA-MMI].

The MAC-NE MUST listen on TCP port 443 (HTTPS port) on any Layer 3-enabled networking ports.

The MAC-NE MAY listen on TCP port 80 (HTTP port). If the MAC-NE listens on TCP port 80, it MUST NOT service RESTCONF requests over that port. If the MAC-NE listens on TCP port 80, it MUST attempt to redirect hosts to port 443 utilizing the HTTP redirect code of "301 Moved Permanently".

The MAC-NE MUST accept TCP connections from any remote host connecting to port 443, then immediately attempt to establish a TLS session with the remote host. The MAC-NE MUST terminate communication to the remote host if it fails to participate in the TLS session.

13.13.2 Connection Process

The connection process between an MAC-NE and each MAC Manager component consists of the following phases:

- establishing a TCP connection between the MAC-NE and MAC Manager;
- performing TLS session establishment, which includes Mutual Authentication and Symmetric Cryptography; and
- MAC-NE configuration and operational monitoring using RESTCONF (see [FMA-MMI]).

The connection process between the MAC-NE and any[R-PHY] Auxiliary Cores are defined in the MHAv2 specifications, and the MAC-NE is expected to comply with that process.

13.13.2.1 Mutual Authentication and Connection Security

The mutual authentication and connection security between the MAC-NE and MAC Manager is independent from the authentication used for trusted network access described in Section 13.5, Network Authentication. The trusted network access in Section 13.5, Network Authentication, details connection trust across a single L2 pair-wise connection, but connection between the MAC-NE and MAC Manager is expected to be a multiple-hop L3 connection.

The connection between the MAC-NE and MAC Manager needs to be authenticated, private, and reliable across an arbitrary L2/L3 network. TLS sessions support these three goals through certificate validation, symmetric cryptography, and message integrity checks. See Section 12.3, Transport Layer Security (TLS) for the FMA System for more details.

13.14 MAC-NE Announce to MAC Managers

The MAC-NE is required to announce itself to the MAC Manager(s) during startup. This is depicted in Figure 69.

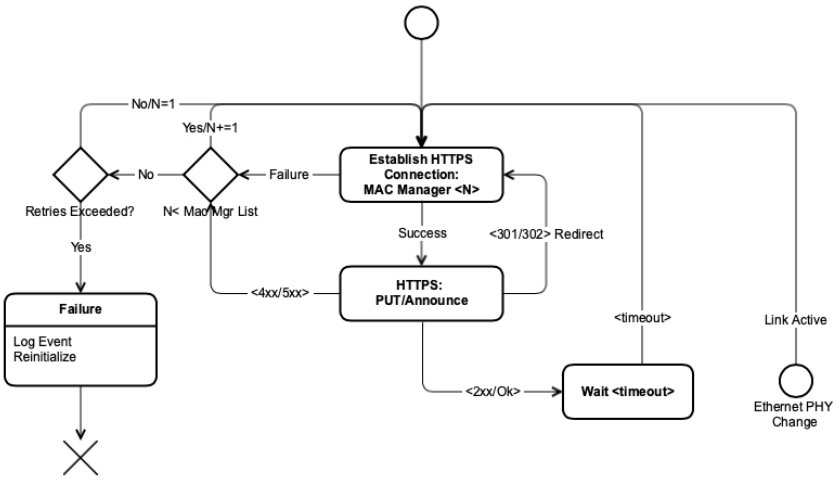

Figure 69 - Announce to MAC Manager FSM

After the MAC-NE starts the RESTCONF listening server, the MAC-NE proceeds to announce itself to one of the MAC Managers in the unified MAC Manager list. The Announcement function represents a steady-state, long-term operation and does not exit or complete. After the first successful Announce, the MAC-NE transitions to a wait state, in which the function will logically sleep until a timeout is completed. Once the timeout has occurred, the Announce function will re-announce following the same procedure as before. The Ethernet PHY link can also trigger a re-announce procedure.

The MAC-NE MUST first formulate a unified, ordered, MAC Manager list from the DHCP options, favoring the MAC Managers received from DHCPv6 over DHCPv4 and favoring IP addresses over domain names. The unified MAC Manager list will be used for all subsequent operations. The MAC-NE MUST reorder this list when it establishes a successful Announce by moving the successful IP address or domain name to the first entry of the list while keeping the rest of the list in the existing order. The MAC-NE MAY further update or reorder the MAC Manager list during run-time operation through methods left to implementation details.

The MAC-NE MUST iterate through the unified MAC Manager list, in order, attempting to successfully establish connections to each MAC Manager. With each MAC Manager in the unified list, the MAC-NE MUST attempt to establish a TCP connection to port 443 and attempt to establish a TLS session. For further detail, review the Transport Layer Security (TLS) for the FMA System section of this document.

Following TLS establishment, the MAC-NE MUST issue an HTTP PUT to the well-known URL "/FMA/v1/ANNOUNCE" with the metadata JSON payload defined in [FMA-YANG].

If the MAC Manager responds with a 4xx or 5xx HTTP error, then the MAC-NE MUST close the connection, record the failure by logging event 70070012, and move onto the next MAC Manager in the unified list.

If the MAC Manager responds with a 301 or 302 HTTP redirect, then the MAC-NE MUST replace the current entry in the unified list with the new MAC Manager information, log event 70070010, close the connection, record the redirection, and begin the connection process with the newly received MAC Manager information.

If the MAC Manager responds with a 201 HTTP success code, then the MAC-NE MUST close the connection, record the success by logging event 70070011, and continue the announce process to any further MAC Managers from the unified list.

Communication and coordination among MAC Managers are outside the scope of this specification.

As detailed in [FMA-MMI], the MAC Manager will populate the "X-CL-HW-IDENTIFIER" HTTP header of the Announce response with a unique ID of the MAC Manager. This UUIDv4 is included as a HTTP header on all subsequent RESTCONF operations from the same MAC Manager.

If a MAC Manager responds to an Announce with a 201 HTTP, the MAC-NE MUST store the value of response header "X-CL-HW-IDENTIFIER" for future authorization of RESTCONF operations from this MAC Manager. The MAC-NE's stored values of the "X-CL-HW-IDENTIFIER" responses are not preserved across reboots.

The MAC-NE MUST reject any future RESTCONF operations which contain a different "X-CL-HW-IDENTIFIER" HTTP header than that of the stored values or the header is absent.

If the MAC-NE is rejecting a RESTCONF operation due to an invalid or missing "X-CL-HW-IDENTIFIER", the MAC-NE MUST respond with a 403 HTTP and record the failure by logging event 70070049.

A 201 HTTP code from the MAC Manager indicates that the initialization process has completed for the MAC-NE. The MAC-NE MUST start or restart a connection_pending timer upon receiving the 201 success code from a MAC Manager in the unified list.

If the MAC Manager responds with any HTTP code that is not 4xx, 5xx, 301, 302, or 201, then the MAC-NE MUST close the connection, record a failure by logging event 70070012, and move onto the next MAC Manager in the unified list.

If the MAC-NE does not receive a 201 HTTP success code from any MAC Manager in the unified list, then it waits a capped random exponential backoff time between 1 and 60 seconds before restarting the Announce procedure from the first MAC Manager in the unified list. The MAC-NE MUST retry the MAC Manager list 4 times (which includes 3 retries) before recording an initialization failure and restarting the boot process.

13.14.1 Redirection

If a MAC Manager does not have configuration data for a MAC-NE or is not aware of the MAC-NE, then it MUST either reject the connection with a 404 HTTP error code or use an HTTP 301/302 code to redirect the MAC-NE to another MAC Manager. The MAC Manager can use the RMD metadata provided in the HTTP PUT to determine which behavior to use; this is left as an implementation detail.

If the MAC Manager decides to redirect the MAC-NE, then it MUST use an HTTP 301/302 code to redirect the MAC-NE.

13.14.2 Connection Pending

Once the MAC-NE has successfully announced itself to a MAC Manager and has a RESTCONF listener server running on TCP port 443, it is ready for configuration to be transferred from the MAC Manager.

If the MAC-NE does not currently have a connection_pending timer, then the MAC-NE MUST start a timer after receiving an HTTP 201 code from a MAC Manager during the Announcement. The connection_pending timer will begin at 5 minutes and count down to 0. If the MAC-NE does not get a successful RESTCONF connection from any valid remote host within the time limit, then the MAC-NE MUST log event 70070013 and restart the boot process.

After receiving the Announce message from the MAC-NE and responding with an HTTP code 201, the MAC Manager MUST issue a valid RESTCONF request to the MAC-NE within the connection_pending timer value of 300 seconds (5 minutes). The MAC Manager does not need to use the same FQDN or IP address when contacting the RMD-hosted RESTCONF server. The MAC Manager MUST use the same "X-CL-HW-IDENTIFIER" in the HTTP headers.

If the MAC-NE has the connection_pending timer counting down and receives a successful RESTCONF request, it MUST terminate the countdown and continue operating.

If a RESTCONF operation is rejected and the connection_pending timer is operational, the MAC-NE MUST keep the connection_pending timer operating.

If the operator desires for the MAC-NE to continue logging events after initialization is complete, no further action is required. If the operator desires the MAC-NE to use a different syslog server IP address, then the onboarding process will replace the DHCP acquired address(es) with the IP address(es) obtained via RESTCONF RPC. When

this occurs, the MAC-NE will cease to use the DHCP acquired syslog server IP address(es) and use only the newly obtained syslog server IP address(es).

In addition to support for syslog facilities to transmit events to a syslog server, the MAC-NE also has a local log for storing events. This local log stores events in situations where no syslog server IP address has been configured. The MAC-NE MUST continue to send events to the syslog server IP address(es) acquired via DHCP unless otherwise substituted via RESTCONF RPC configuration from the MAC Manager.

The MAC-NE MUST cease to transmit events to the syslog server IP address(es) if they were replaced by new IP addresses set during the onboarding process.

Details of the local log retention policy and the minimum number of events it is required to store are found in [FMA-OSSI].

The MAC-NE expects that configuration will be applied by the MAC Manager at this time. The details of the configuration and monitoring object models are defined in [FMA-YANG].

13.14.3 Announce Timer

The MAC-NE Announce function represents a background task that continues for the entire life cycle of the MAC-NE operation. Once a successful Announce has occurred, the MAC-NE MUST wait a random timeout of between 20 and 40 seconds before doing the Announce procedure again. The MAC-NE MUST allow the random timeout bounds to be configurable by the MAC Manager.

The MAC-NE MUST NOT reset the timer if the MAC-NE is still waiting for a connection_pending timer during the connection pending step (Section 13.14.2, "Connection Pending").

The MAC-NE continues to re-announce itself during the entire lifetime of the MAC-NE. The Announce procedure causes the MAC-NE to place the last successful MAC Manager as the first entry in the MAC Manager list, meaning the re-Announce procedure will contact that MAC Manager first.

The MAC-NE MUST immediately re-trigger the Announce process when the link status changes to link-present on the primary Ethernet PHY.

13.14.4 Software Upgrade

During the onboarding of the MAC-NE, if the MAC Manager indicates that software upgrade is required, it will instruct the MAC-NE to perform the upgrade. Details of this process and the associated requirements are found in the Secure Software Download (SSD) section of this document.

13.15 Synchronization

After the MAC-NE has successfully been configured by the MAC Manager, it may begin clock synchronization. The MAC-NE MUST support internal synchronization in which the MAC-NE acts as a clock master to itself. The MAC-NE MUST support external synchronization in which the MAC-NE acts as a clock slave to an external synchronization source over PTP [IEEE 1588]. The MAC-NE MAY support external synchronization in which the MAC-NE acts as a frequency slave to an external synchronization source over SyncE [ITU-T G.8261.1] and [ITU-T G.8262/Y.1362].

In both an FMA and R-PHY system, the downstream and upstream PHY timing are always aligned because the downstream and upstream PHY are co-located in the node. Specifically, in R-MACPHY, unlike in R-PHY, there is no need to be able to share a timestamp value between a CCAP Core and the node for DOCSIS upstream scheduling; this is a major benefit of R-MACPHY. There is also no need to consider moving the upstream scheduler to the node to fix latency, jitter, and timing issues in an R-MACPHY because the scheduler already resides in the MAC-NE.

Although precision timing is not required in FMA for DOCSIS upstream scheduling, there are applications for which it is necessary. The protocol specified in [R-DTI], PTP as defined by [IEEE 1588], will be used between the FMA Core Timing Devices and the MAC-NE in the FMA. PTP is used because it is a standard protocol whose accuracy can be enhanced when the CIN is built with [IEEE 1588]-compliant equipment. Note that it is not necessary for the network to be compliant to [IEEE 1588].

Encryption or authentication of PTP messages between the master clock and the MAC-NE (e.g., by using IPsec) would result in some loss of accuracy because intermediate nodes could not update the timing data. If security of PTP messages is required, then MACsec encryption of the untrusted portion of the network can be used.

The synchronization requirements can be summarized as follows.

- All specific operational requirements are stated in [R-DTI].
- The MAC-NE MUST be able to support PTP messages received over a MACsec-secured (see [IEEE 802.1ae]) link from the CIN.
- If the MAC Manager has specified a PTP Master source through [FMA-YANG] configuration, then the MAC-NE MUST use it.

During the initialization process, the MAC-NE MUST use a local clock source for frequency information. During the initialization process, the MAC-NE MUST use the ToD server for initial time information and use a local clock source to increment that time information.

The MAC Manager can provide PTP configuration as per [FMA-YANG]; if it does, then the MAC-NE MUST attempt to externally synchronize both time and frequency to the configured source.

The MAC Manager can provide SyncE configuration as per [FMA-YANG]; if it does, then the MAC-NE MAY attempt to externally synchronize its frequency to the configured source. If the MAC-NE does not support SyncE configuration, then it MUST report an error during the configuration process, as per [FMA-MMI].

The MAC Manager MUST configure PTP before configuring [R-PHY] Auxiliary Cores. If the MAC Manager provides [R-PHY] Auxiliary Cores, then the MAC-NE MUST comply with the Auxiliary Core PTP operation as specified in [R-PHY].

13.15.1 Synchronization Failures

13.15.1.1 MAC-NE Operating as a Timing Slave

If the MAC-NE does not receive a sync message within PTP_SYNC_TIMEOUT, then it MUST log a LOSS_OF_SYNC error and operate in a non-synchronized local clocking mode. It will continue to attempt to synchronize with a PTP clock master. When synchronization is re-established, the MAC-NE MUST log a SYNC_ESTABLISHED message.

13.15.1.2 MAC-NE Failure to Synchronize

If it does not achieve synchronization within RPD_PTP_OK_TIMEOUT, then the MAC-NE MUST send a PTP Timeout Notify to the MAC Manager.

Table 18 - MAC-NE PTP Timeout

Constant	Value
RPD_PTP_OK_TIMEOUT	60 seconds

13.15.1.3 Abort from Core Timing Devices

At any time during the synchronization process, the Core Timing Devices may abort the PTP process on a port, which will be reported by the MAC Manager via mechanisms defined in [FMA-MMI].

The MAC-NE MUST be prepared to receive and process an aborted PTP process at any time.

- the MAC-NE MUST stop active PTP processing on the port, after the MAC Manager sets the PTP AdminState for a port to down.
- the MAC-NE MUST continue operating any services which do not need to be synchronized, after the MAC Manager sets the PTP AdminState for a port to down.
- the MAC-NE MUST wait for further direction relating to PTP processing from the MAC Manager, after the MAC Manager sets the PTP AdminState for a port to down.

- the MAC-NE MUST continue PTP operation on other ports, after the MAC Manager sets the PTP AdminState for a port to down.

13.16 Move to Operational

This section describes the states change at the conclusion of the initialization process.

13.16.1 MAC Manager

When the MAC Manager deems that the MAC-NE initial configuration is complete, the logic of which is left as an implementation detail, the MAC Manager MUST configure the MAC-NE as Operational using the mechanisms defined in [FMA-YANG]. The MAC Manager will likely use optional PTP synchronization and [R-PHY] Auxiliary Core configurations as input to the move to Operational logic.

If the MAC-NE does not receive an indication from the MAC Manager to move to Operational state within 60 seconds, then the MAC-NE MUST send a notification to the MAC Manager (detailed in [FMA-YANG]), log the failure locally, reset the timer, and continue to wait.

13.16.2 Auxiliary Core

The MAC Manager configures [R-PHY] Auxiliary Cores on the MAC-NE through mechanisms defined in [FMA-YANG]. The MAC Manager MUST configure resources to be utilized by the Auxiliary Cores when it is deems the MAC-NE will likely succeed in connecting to those elements. The MAC-NE MUST comply with the Auxiliary Core Move to Operational requirements specified in [R-PHY].

13.16.3 PacketCable Aggregator

In an FMA implementation without PacketCable, the MAC Manager will receive PacketCable configuration that has both PacketCable DQoS and PCMM disabled. In this case there will be no PacketCable Aggregator (PAG) function necessary.

In an FMA implementation with PacketCable, the MAC Manager will receive PacketCable configuration that has PacketCable configured. In this case there will be a PAG function implemented. The MAC Manager configures the MAC-NE with a PAG URL, a PAI connection retry count, and a PAI TLS session establishment retry count through mechanisms defined in [FMA-YANG]. If PacketCable has been configured but the MAC Manager has not configured the MAC-NE with a PAG URL, a PAI connection retry count, and a PAG TLS session establishment retry count, the MAC-NE MUST abort the PAG connection attempt and log a PAG configuration failure critical event (70070026).

When configured to do so, the MAC-NE MUST initiate PAI connection to the PAG by using DNS to resolve the PAG URL to an IP address and TCP port. If the PAG URL could not be resolved to an IP address and TCP port, the MAC-NE MUST abort the PAG connection attempt and log a PAG URL resolution failure critical event (70070032).

After the MAC-NE obtains a PAG URL, the connection process between a MAC-NE and PAG consists of the following phases:

- establishing a TCP connection between the MAC-NE and PAG;
- performing TLS session establishment, which includes Mutual Authentication and Symmetric Cryptography; and
- performing PacketCable Aggregator Interface (PAI) initialization (see [FMA-PAI]).

The PAG MUST accept TCP connections from any remote host connecting to the PAG TCP port, then expect an immediate attempt to establish a TLS session by the remote host. The PAG MUST terminate the TCP connection to the remote host if the remote host fails to establish a TLS session.

The MAC-NE MUST attempt TCP connection to the PAG IP address at the PAG TCP port. If TCP connection with the PAG fails, the MAC-NE MUST retry the connection up to the PAI connection retry count or until a successful TCP connection, whichever comes first. If, after exhausting the number of PAI connection retries the MAC-NE has

not successfully established a TCP connection with the PAG, the MAC-NE MUST abort the connection attempt and log a PAI initialization failure no TCP connection critical event (70070033).

After establishing a TLS session with the PAG, the MAC-NE initializes its PAI as specified in [FMA-PAI].

13.16.3.1 Mutual Authentication and Connection Security

The connection between the MAC-NE and PAG needs to be authenticated, private, and reliable across an arbitrary L2/L3 network. TLS sessions support these three goals through certificate validation, symmetric cryptography, and message integrity checks.

Please see Section 12.3, Transport Layer Security (TLS) for the FMA System.

14 PROACTIVE NETWORK MAINTENANCE

14.1 Scope

This section defines the architecture and requirements for supporting Proactive Network Maintenance (PNM) in FMA. DOCSIS CMTS and CM features and capabilities can be leveraged to enable measurement and reporting of network conditions so undesired impacts such as plant equipment and cable faults, interference from other systems, and ingress can be detected and measured. With this information, cable network operations personnel can make modifications necessary to improve conditions and monitor network trends to detect when network improvements are needed.

PNM is considered a part of the Performance Management subset of FCAPS Management Operations defined for FMA in Section 8.2. FMA supports all PNM features and capabilities defined in [PHYv3.1], [MULPIv4.0], [CCAP-OSSIv4.0], and [CM-OSSIv4.0]. Requirements written against a CMTS or CCAP in those specifications are applied to the functional entities that comprise a CMTS or CCAP in the FMA. Requirements written against a CM in those specifications are unchanged by FMA but are an important part of the suite of PNM tests, so are also discussed in the FMA specifications.

Generally speaking, PNM tests are configured by the PNM Server on the MAC Manager and the CM using SNMP. PNM test results are reported to the PNM Server by the MAC Manager and CM using SNMP and TFTP. This section provides an overview these processes as well as discussion of use of the Common Collection Framework (CCF) for PNM.

[FMA-OSSI] contains detailed descriptions of the supported PNM tests and normative requirements on the FMA functional entities involved in them.

14.2 System Description

14.2.1 PNM Supported Functions

As shown in Figure 70, the CMTS and CM contain test points that include essential functions of a spectrum analyzer, vector signal analyzer (VSA), and network analyzer, while the cable plant is considered the device under test (DUT). The goal of PNM is to rapidly and accurately characterize, maintain, and troubleshoot the upstream and downstream cable plant to guarantee the highest throughput and reliability of service. The CMTS and CM make the specified measurements and report the results to the PNM Server.

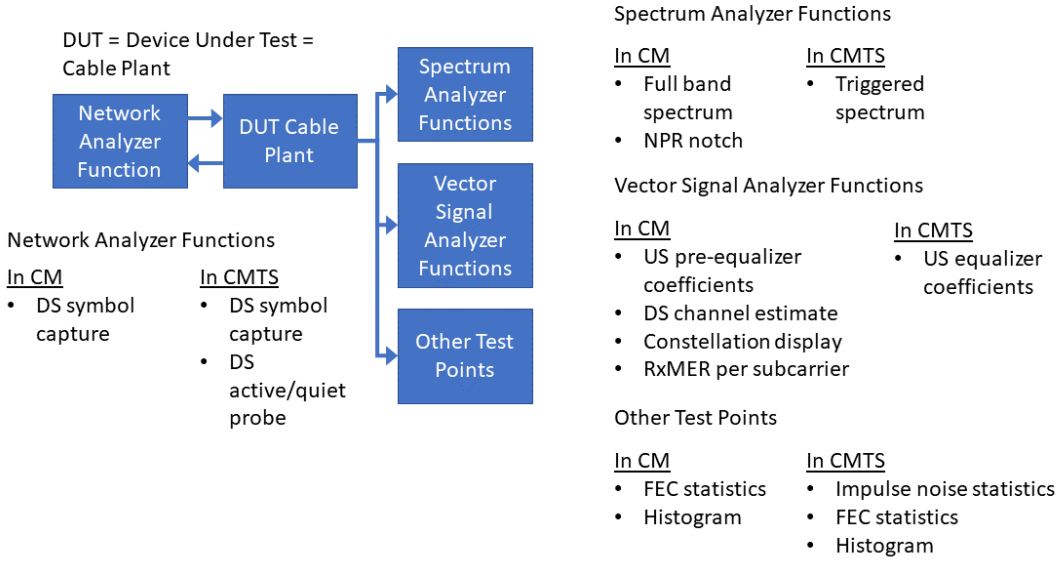

Figure 70 - Test Points in CMTS and CM Supporting PNM

PNM tests supported in FMA are included in Table 19, which lists the network elements participating in each test. CMs are listed in this table for completeness even though FMA places no new requirements on the CM for PNM. This is because CM participation in the listed tests is significant. Tests which only involve the CM but neither the MAC Manager nor MAC-NE are not included in the table or the FMA specifications. Please see [FMA-OSSI] for detailed test descriptions and normative requirements.

Table 19 - FMA PNM Tests

Test Name	MAC Mgr	MAC NE	CM
Downstream Symbol Capture	X	X	X
Downstream Noise Power Ratio (NPR) Measurement	X	X	X
Downstream OFDM Profile Test (OPT)	X	X	X
Upstream Capture for Active and Quiet Probe	X	X	X
Upstream Triggered Spectrum Capture	X	X	X
Upstream Impulse Noise Statistics	X	X	
Upstream Equalizer Coefficients	X	X	X
Upstream FEC Statistics	X	X	
Upstream Histogram	X	X	
Upstream OFDMA Rx Power	X	X	X
Upstream Receive Modulation Error Ratio (RxMER)	X	X	X
Upstream Sweep	X	X	

Upstream sweep is listed as a PNM test but is actually an application that utilizes upstream triggered spectrum capture.

14.2.2 PNM Architectures

14.2.2.1 Integrated CMTS/CCAP

The system architecture of PNM for an integrated CMTS/CCAP system is shown in Figure 71. In this architecture, the PNM Servers initiate PNM tests via SNMP, collect the associated test results from the CMTS/CCAP and CM via SNMP and/or TFTP, and process these results. The PNM Server is assumed in most cases to include a TFTP file server where bulk data files containing PNM test results are stored by the CMTS/CCAP, CM, or both. The TFTP server can but need not be on the same device as the PNM Server and can even be a shared resource.

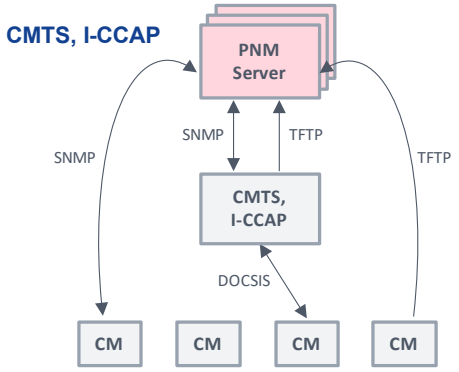

Figure 71 - Integrated CMTS - CCAP PNM Architecture

In the integrated CMTS/CCAP system, there are also proprietary PNM tests, such as Upstream Sweep, which bypass the CMTS/CCAP and make use of the RF combining network to connect the PNM Server to a Field Device. These

applications do not require interaction with the CMTS/CCAP. All control and data plane communication between the PNM Server and Field Device are in the RF domain, as shown in Figure 72.

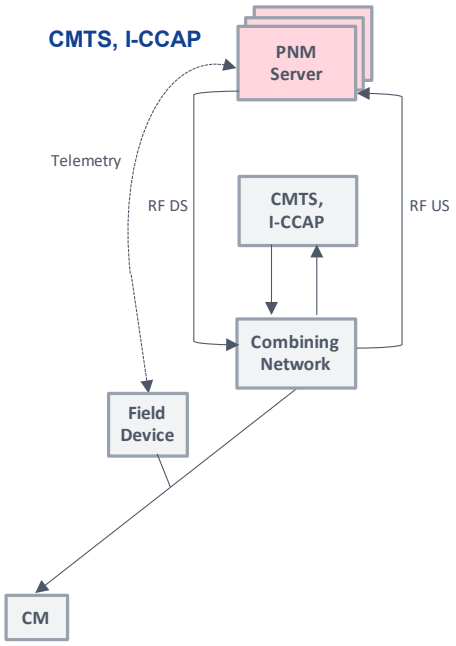

Figure 72 - Integrated CMTS - CCAP Upstream Sweep Architecture

14.2.2.2 Distributed CMTS/CCAP

In a Distributed CCAP Architecture (DCA) the CMTS/CCAP function is performed by disparate entities. The first phase of DCA is Remote PHY. The second phase of DCA is FMA, in which the MAC Manager performs management functions and the MAC-NE performs all DOCSIS control and data plane functions. FMA supports PNM as shown in Figure 73.

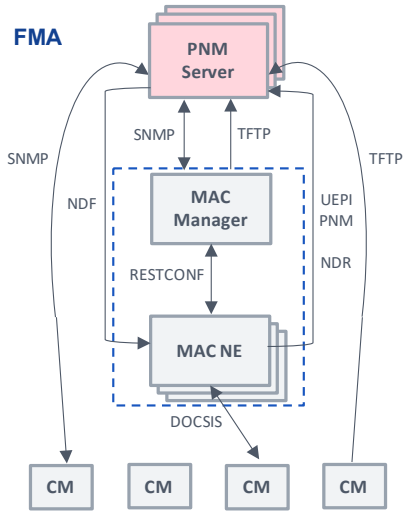

Figure 73 - FMA PNM Architecture

FMA supports the same data models and interfaces specified for integrated CMTS/CCAP to interface the PNM Server to distributed CMTS/CCAP functional entities. SNMP is used by the PNM Server to configure PNM tests and to receive test status and basic results. TFTP is used by the PNM Server to retrieve bulk test results, in which

large amounts of data are involved. The PNM Server is assumed in most cases to include a TFTP file server where bulk data files containing PNM test results are stored as sent by the MAC Manager, CM, or both.

FMA uses RESTCONF and YANG between MAC Manager and MAC-NE to provision PNM tests on the MAC-NE and to receive test results from it, as specified in [FMA-MMI].

Because the FMA does not include an RF combining network, an additional set of interfaces and requirements have been specified to support Upstream Sweep. A UEPI PNM pseudowire can be used to stream upstream spectrum captures from the MAC-NE directly to the PNM Server. Vendor-specific delivery of spectrum captures to the PNM Server can also be supported. NDF and NDR pseudowires between the MAC-NE and PNM Server communicate telemetry associated with the test between the PNM Server and a Field Device connected in the outside plant. Once the spectrum capture and telemetry transport mechanisms have been established from the MAC-NE to the PNM Server, the Upstream Triggered Spectrum Capture PNM test is configured by the PNM Server via SNMP on the MAC Manager. Upstream Triggered Spectrum Capture is set to run in free run mode. The MAC-NE will deliver the spectrum captures to the PNM Server for the duration of the test. The PNM Server can request a long duration and can manually stop the test at the point it completes.

A common architecture for Upstream Sweep in FMA is shown in Figure 74. In this architecture the UEPI PNM and NDR/NDF telemetry pseudowire processing is all asynchronous, with no timing synchronization or frequency lock required at the ends of the pseudowires. The PNM Server processes streamed spectrum captures from the UEPI PNM pseudowire as they are received, having no need for synchronization with the sender. The PNM Server also processes NDR data as it is received without any need to do digital to analog conversion as would be the case in many other NDR applications, and hence no need for frequency synchronization. It is understood that the MAC-NE will likely experience FIFO overrun and/or underrun in its NDF digital to analog processing buffers. It is expected that the MAC-NE will be able to recover from these conditions, but that the RF generated by the node will have occasional glitches. It is the responsibility of the vendor-proprietary Upstream Sweep telemetry protocols to be robust against these glitches.

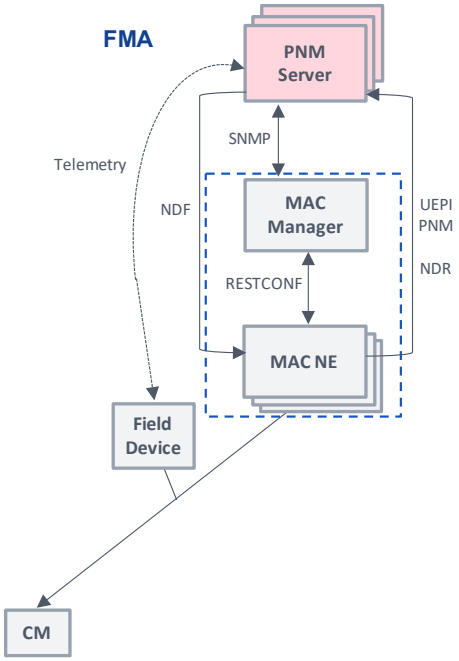

Figure 74 - FMA Upstream Sweep Alternatives

14.2.3 Common Collection Framework (CCF)

CableLabs has proposed the CCF to optimize data collection functions across multiple applications. This framework is documented in the CableLabs [XCCF] technical report.

CCF has a number of potential benefits.

- collects data once, uses many times
- supports multiple apps and functions—PNM, PMA, Certificate Authority, and more, with more to be determined
- provides a single point of contact for multiple apps to network elements (CMTS, CCAP) and CMs
- provides a single point of contact for network elements and CMs to have data collected for multiple apps
- supports a modern API for applications at the back office application layer which can support future network element interfaces such as RESTCONF or NETCONF, as an alternative to legacy SNMP and TFTP which exists today
- supports existing interfaces (i.e., SNMP, TFTP) for network elements (CMTS, CCAP) and CMs.

Although there are no normative requirements associated with CCF, the concepts are admirable, and new interfaces to CCF based on RESTCONF or NETCONF could easily be supported given the FMA use of RESTCONF between the MAC Manager and MAC-NE. In this way, the MAC Manager could merely serve as a simple intermediary for the RESTful API between the CCF framework interface and its MAC-NE MMI. Likewise, the MAC Manager could host the CCF function. The CCF architecture is shown for FMA in Figure 75.

CCF currently does not include support for reception of streamed test results, as can be used in the Upstream Sweep capability by FMA. CCF also does not include support for NDF and NDR pseudowires for telemetry, as has been architected for Upstream Sweep. But because of its architecture, CCF could support collection and presentation of streaming data, as applications require.

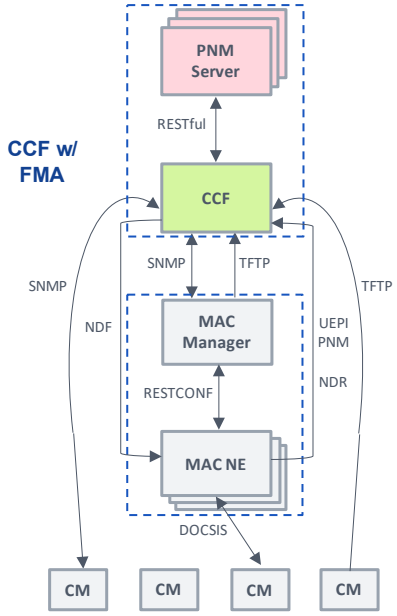

Figure 75 - FMA with CCF

Annex A Core to MAC-NE MTU (Normative)

A.1 L2TPv3 Lower Layer Payload Size

Typically, an interface calculates its default maximum payload size by asking the interface below it in the interface channel for its maximum payload size and considering its own encapsulation. For example, by default, Ethernet has a frame size of 1518 (without VLANs). The Ethernet encapsulation is 18 bytes, leaving 1500 bytes of payload (MTU) for its upper layer. IP then subtracts the IP header size (typically 20 bytes) to arrive at 1480 bytes available to its upper layer. For D-MPT, the remainder becomes 1472 bytes because the Session Field and the L2TPv3 Data Session Header comprise 8 bytes. For PSP, the PSP header including the maximum PSP segment table size needs to be taken into account.

The FMA core components (Video Core, OOB Core, NDx Core) needs to support expanded Ethernet Frame sizes, up to 2000 bytes long, in compliance with [MULPIv4.0]. For details on the data layer frames size requirements, see Section 7.3.1.1, Data Plane Packet Size Requirements.

A.2 Maximum Frame Size for DEPI

This section documents the maximum frame size of the DEPI when a PSP pseudowire is used without fragmenting or concatenation.

Table 20 - MTU of DEPI (for PSP)

Field			Size
DEPI Frame	Ethernet Header		14 bytes
	802.1Q Header		4 bytes
	DEPI MTU	IPv4 Header	20 bytes
		IPv6 Header	40 bytes
		L2TPv3 Header	8 bytes
		DEPI-PSP Header*	6 bytes
		DOCSIS Frame: DOCSIS Header**	6–246 bytes
		DOCSIS Frame: Ethernet Header	14 bytes
		DOCSIS Frame: 802.1Q Header	4 bytes
		DOCSIS Frame: Ethernet PDU	1500 or 2000 bytes
		DOCSIS Frame: Ethernet CRC	4 bytes
	Ethernet CRC		4 bytes
Total with PSP, no UDP, IPv4, no VLAN			**1570 to 1862 (or 2362)**

* A PSP header is 4 bytes plus 2 bytes for each segment. Only one segment is shown. (A D-MPT header is 4 bytes.)

** A typical DOCSIS header with BPI and no other extended headers is 11 bytes.

For simplicity, only one PSP segment is included in the above calculations. Additional segments are needed when PSP is concatenating or fragmenting. Note that a 2000 byte payload in a PSP frame could contain as many as 26 uncompressed TCP ACKs (64 byte Ethernet packets plus 6 to 11 bytes of DOCSIS overhead), which could create as many as 22 segments (first and last packets are fragmented), which would create a segment table size of 44 bytes, in addition to the standard 4 byte PSP header. For other payload types, such as VoIP packets with high codec compression and with PHS disabled or with larger MTUs, the number of segments could be even higher.

A.3 Path MTU Discovery

Path MTU Discovery relies on the fact that the network elements between the FMA core components (Video Core, OOB Core, NDx Core) and the MAC-NE all support this functionality [RFC 1191]. If these network elements do not support Path MTU Discovery, then this mechanism cannot be used, and the static configuration option should be used instead.

Path MTU Discovery (PMTUD) works when (1) the IP path MTU between the FMA core component and the MAC-NE is less than the total IP datagram size generated when using the payload size negotiated during L2TPv3 session establishment and (2) the Don't Fragment (DF) bit is set in the IP header. If the FMA core component (Video Core, OOB Core, NDx Core) sends packets larger than the network can support, then network elements between the FMA core component and the MAC-NE MAY generate an ICMP Destination Unreachable message with the code "Fragmentation needed and DF set" (ICMP Type 3 Code 4, also referred to as "Datagram Too Big" message) toward the source of the tunneled packet, if ICMP unreachables are allowed.

This ICMP error message includes at least the IP header and the next 8 bytes of the IP data (corresponding to the UDP header when using L2TPv3 over UDP or to the Session ID and first 4 bytes of the L2SS when using L2TPv3 over IP) from the offending packet. The FMA core component (Video Core, OOB Core, NDx Core) and the MAC-NE SHOULD have a way to map the source and destination IP address contained in the IP header embedded in the ICMP data to an L2TP Control Connection. As defined in [RFC 1191], a "PMTU is associated with a path, which is a particular combination of IP source and destination address and perhaps a Type-of-Service (ToS)."

Upon successfully processing the ICMP Destination Unreachable message, the FMA core component (Video Core, OOB Core, NDx Core) and MAC-NE SHOULD reduce the Max Payload of all the sessions associated with the control connection mapped from the ICMP Destination Unreachable message to the size requested in the Next-Hop MTU field of the message. Both the Max Payload and the size contained in the Next-Hop MTU field express a Layer 3 payload of a Layer 2 frame, including the IP header and IP data.

The Max Payload cannot be increased when an FMA core component (Video Core, OOB Core, NDx Core) or an MAC-NE receive an ICMP Destination Unreachable message. The FMA core component (Video Core, OOB Core, NDx Core) and MAC-NE MAY periodically attempt to increase the Max Payload of the session to its negotiated maximum and restart this process in case the path through the network has changed and larger MTUs are allowed. This technique is described in [RFC 1191]. The Max Payload size learned through this process will never be greater than the negotiated maximum learned during session establishment. The Path MTU Discovery procedures for IPv6 are described in [RFC 1981].

Annex B GCP Usage (Normative)

Below are notes on leveraging [R-PHY] Annex B requirements and definitions in FMA.

1. [R-PHY] Annex B consists largely of TLVs that define RCP extensions to GCP and associated processing requirements at the CCAP Core and RPD.
2. Many parts of [R-PHY] Annex B are reused, given that FMA leverages the R-PHY initialization and configuration procedures for OOB and Video Cores, which heavily utilize GCP and RCP.
3. Some of [R-PHY] Annex B does not apply to FMA. For example, TLVs, codepoints, and requirements related to tunneling and processing of DOCSIS MAC management messages (MAP, RNG-REQ, BW-REQ, OCD, DPD, UCD, etc.) and tunneling DOCSIS data traffic are applicable only to R-PHY. Special configuration of SIDs and snooping of MAC MAP messages for SID information so that RPD upstream burst receivers can be dynamically configured for segment header on and segment header off bursts is entirely unnecessary in FMA. In this particular example, all TLVs, codepoints, and requirements specified in [R-PHY] Sections B.5.4.20–27 are not relevant for FMA. As another example, all TLVs, codepoints, and requirements related to snooping BW-REQs from CMs, buffering them, and getting them onto special UEPI pseudowires in a timely fashion are out of scope for FMA.
4. Below is a list of rules for mapping requirements and definitions from R-PHY Annex B to FMA.
    a. The GCP usage requirements on the Principal Core and/or CCAP Core in [R-PHY] Annex B do not apply directly to the MAC Manager because the MAC Manager-to-MAC-NE interface does not use GCP. The MAC Manager Interface specification ([FMA-MMI]) contains the normative requirements on the MAC Manager-to-MAC-NE interface and on the data exchanged over this interface. In many cases, requirements in [FMA-YANG] will be based on requirements in [R-PHY] Annex B.
    b. FMA Auxiliary Cores MUST implement all [R-PHY] Annex B Auxiliary Core requirements.
    c. FMA Auxiliary Cores MUST implement all [R-PHY] Annex B CCAP Core requirements that apply to Auxiliary Cores. One such example is in [R-PHY] Section B.2.3.15: "If the CCAP Core is prepared to act as an Active Auxiliary and configure the RPD, the Core MUST generate a single IRA message which includes a write request for the CcapCoreIdentification attributes."
    d. MAC-NEs MUST implement all [R-PHY] Annex B RPD requirements that support communication with Auxiliary Cores.
    e. MAC-NEs SHOULD NOT implement any [R-PHY] Annex B RPD requirements that pertain only to communication with Principal Cores.
5. The approach taken in this annex is to assume the above requirements mappings for [R-PHY] Annex B into this annex and to only annotate the respective FMA-SYS Annex B section where exceptions or the need for clarifications exist.
6. [R-PHY] Annex B section numbering and names are preserved in this annex for traceability. Where there is nothing to say about a given section or its subsections (i.e., there is no difference of significance in using the section for FMA versus R-PHY), nothing is said. This does not rule out future additions to those sections, especially as new methods for configuring and managing the MAC-NE arise that differ from the methods described in [R-PHY].
7. Numerous TLVs from [R-PHY] Annex B do not apply to FMA. Certain code points within universally applicable TLVs do not apply to FMA. Numerous functional requirements do not apply to FMA. The assumption is that such TLVs, code points, and requirements will simply not be used in FMA implementations. As new methods for configuring and managing the MAC-NE are introduced in FMA, and as [R-PHY] Annex B continues to be modified, it is likely that additional unused TLVs, code points, and requirements could be identified. Compiling and maintaining such a list would be difficult, if not misleading, so this annex does not call out this type of information.

B.1 R-PHY Control Protocol

B.1.1 RCP over GCP EDS Message

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.2 RCP over GCP EDS Response Messages

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.3 RCP over GCP Device Management Message

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.4 RCP over GCP Notify Message

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.5 RCP TLV Format, TLV Types and Nesting Rules

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.6 RCP Message Structure

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.7 RCP Messages Types

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.8 RCP Protocol Rules

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.9 Protocol Extensibility

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.10 Protocol Versioning

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.11 Information Model Extensibility

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.12 Vendor-Specific Extensions

This [R-PHY] Annex B section specifies Principal Core and RPD requirements to support vendor-specific extensions, in particular, vendor-specific preconfiguration of the RPD. In FMA, the MAC Manager is responsible for vendor-specific preconfiguration of the MAC-NE. The [FMA-MMI] specification contains the normative requirements for vendor-specific extensions, including vendor-specific preconfiguration, on the MAC Manager and MAC-NE.

GCP will not be used for this capability.

B.1.13 Inclusion of DOCSIS Messages

This [R-PHY] Annex B section specifies CCAP Core and RPD requirements to support upstream channel and downstream OFDM profile change procedures in an R-PHY system in which tunneling of DOCSIS MAC Management messages in RCP is required for UCD, OCD, and DPD. Also specified are complex interactions of UCD processing with other MAC-related functions done on the RPD (i.e., MAP snooping and processing), as well

as synchronization rules related to applying OFDM downstream profile changes at an RPD in a coordinated fashion to all affected CMs.

These requirements do not apply to an R-MACPHY system.

B.1.14 Event Reporting

This [R-PHY] Annex B section specifies Principal Core and RPD requirements to support event logging at the RPD and event reporting from the RPD to the Principal Core. In FMA, the Principal Core is replaced by the MAC Manager, and the interface used for event reporting between the MAC Manager and MAC-NE is as defined in [FMA-MMI]. The interface used for event reporting between the MAC Manager and the back office is defined in [FMA-OSSI]. Consequently, all normative requirements related to event logging and reporting on the MAC Manager and MAC-NE are contained in those specifications.

GCP will not be used for this capability.

B.1.15 Error Handling

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.1.16 RCP Message Examples

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.2 RPD Initialization

This [R-PHY] Annex B section specifies CCAP Core and RPD requirements for GCP-based initialization and connection to CCAP Cores. In FMA, there is no Principal Core, and the MAC Manager is responsible for MAC-NE configuration during initialization, including initial MAC-NE Auxiliary Core configuration. MAC Manager control of MAC-NE initialization is as specified in Section 13 of this document. [FMA-MMI] also contains normative requirements for initialization on the MAC Manager and MAC-NE.

GCP is not used between the MAC Manager and MAC-NE. When functionality that is provided by GCP between a Principal CCAP Core and an RPD should also be considered for use between the MAC Manager and MAC-NE, it will be specified in [FMA-YANG]. The subsections that follow note specific instances of [R-PHY] Annex B GCP usage requirements in FMA that should be considered for specification in [FMA-YANG].

B.2.1 GCP Connection Initialization Sequence

The MAC Manager will configure Auxiliary Cores on the MAC-NE during initialization. This configuration replaces the DHCP-based configuration of Auxiliary Cores that is used in [R-PHY].

The GCP Core Redirect capability of R-PHY is not being carried forward in R-MACPHY.

B.2.2 Initialization RCP Messages RPD and Cores

B.2.2.1 IRA vs REX Usage

The MAC-NE will not send a StartUp NOTIFY to the MAC Manager, and the MAC Manager will not generate an IRA response, nor will it issue REX commands for subsequent configuration of the MAC-NE.

B.2.2.2 Start Up Notify

The MAC-NE will not send a StartUp NOTIFY to the MAC Manager.

B.2.2.3 RedirectResult Notify

The GCP Core Redirect capability of [R-PHY] is not being carried forward into FMA.

B.2.2.4 Ptp Notify

The MAC-NE will not generate a NOTIFY message to the MAC Manager when local PTP synchronization is achieved or a failure is detected.

B.2.2.5 Auxiliary Core Result Notify

The MAC-NE will not generate a NOTIFY message to the MAC Manager when connection to an Auxiliary Core is successful or a failure is detected.

B.2.2.6 Time Out Notify

Following a time out event, the MAC-NE will not generate a Time Out NOTIFY message to the MAC Manager.

B.2.2.7 IRA – Core Is Active Principal and Will Configure RPD

The MAC Manager is considered to act as an Active Principal and configures the MAC-NE. However, it will not use GCP/RCP to do so.

B.2.2.8 IRA – Core Is Backup (Not Active) Principal for RPD

Redundancy of MAC Managers in FMA is via loose coupling such that the connection from the MAC-NE to the MAC Manager is not required to be persistent. As a result, it is anticipated that the MAC-NE can operate autonomously from time to time and will not need to track active versus backup MAC Managers, as is the case in [R-PHY], in which RPDs actively participate in redundancy with the Principal Core/DOCSIS Core.

B.2.2.9 IRA – Core Is Not Active Principal or Backup Principal

B.2.2.10 IRA – Redirect

The GCP Core Redirect capability of [R-PHY] is not being carried forward in R-MACPHY.

B.2.2.11 IRA – Decision Pending

The GCP Core Redirect capability of [R-PHY] is not being carried forward in R-MACPHY.

B.2.2.12 IRA – S/W Upgrade

Software upgrade in FMA will be controlled by the MAC Manager, as specified in [FMA-YANG].

B.2.2.13 REX – Configuration by Active Principal Core

The MAC Manager is considered to act as an Active Principal and configures the MAC-NE. However, it will not use GCP/RCP to do so.

B.2.2.14 Software Upgrade before Redirection

The GCP Core Redirect capability of [R-PHY] is not being carried forward in R-MACPHY.

B.2.2.15 IRA – Core Is Active Auxiliary and Will Configure RPD

B.2.2.16 IRA – Core Will Act as Backup Auxiliary for RPD

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.2.2.17 IRA – Core Is Not Active and Not Backup Auxiliary for RPD

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.2.2.18 REX – Configuration by Active Auxiliary Core

No requirements, exceptions, or clarifications are necessary against the corresponding section of [R-PHY] Annex B.

B.2.3 RPD Initialization States

MAC Manager control of MAC-NE initialization is as specified in Section 13 of this document. [FMA-MMI] also contains normative requirements for initialization on the MAC Manager and MAC-NE, including MAC-NE initialization state reporting requirements. Below are requirements that are ported from the corresponding subsections of [R-PHY] Annex B and applied to FMA.

B.2.3.1 High Level RPD State

The MAC-NE MUST maintain a single state variable for the current high-level MAC-NE state.

The MAC-NE MUST report this state as specified in [FMA-YANG].

B.2.3.2 Network Authentication

Because Network Authentication is on a per port basis, the MAC-NE MUST maintain a state variable per port.

The RPD MUST be able to report Network Authentication state on a per port basis as specified in [FMA-YANG].

The MAC Manager MUST request Network Authentication state using the authentication port index as the key, as specified in [FMA-YANG].

B.2.3.3 Connect to Active Principal Core Sub State

The connection between the MAC Manager and MAC-NE will have an association operational state.

The connection status and reporting of MAC Manager connection state from the perspective of the MAC-NE will be as specified in [FMA-YANG].

B.2.3.4 Auxiliary Core State

Because connections to Auxiliary Cores are on a per-core basis, the MAC-NE MUST maintain a state variable per core.

The RPD MUST be able to report core state on a per Auxiliary Core basis as specified in [FMA-YANG].

The MAC Manager MUST request the Auxiliary Core state as specified in [FMA-YANG].

B.2.4 Reconnect Messages

B.2.4.1 ReconnectNotify

B.2.4.2 HandoverNotify

B.2.4.3 AuxCoreGcpStatusNotify

The MAC-NE will not generate a NOTIFY message to the MAC Manager when the status of a GCP connection to an Auxiliary Core changes.

Annex C MPEG Stream Analysis (Normative)

The MAC-NE MUST support RPD MPEG Stream Analysis requirements as specified in [R-PHY] Annex C, MPEG Stream Analysis (Normative).

Annex D Certificate Hierarchy and Profiles (Normative)

The MAC-NE MUST support RPD Certificate Hierarchy and Profiles requirements as specified in [R-PHY] Annex D, Certificate Hierarchy and Profiles (Normative).

The MAC Manager MUST support CCAP Core Certificate Hierarchy and Profiles requirements as specified in [R-PHY] Annex D, Certificate Hierarchy and Profiles (Normative).

Annex E Receive Power Level Management (Normative)

The MAC-NE MUST support RPD Receive Power Level Management requirements as specified in [R-PHY] Annex E, Receive Power Level Management (Normative).

Annex F DOCSIS 4.0 OFDM Modifications for Remote PHY (Normative)

The MAC-NE MUST support RPD DOCSIS 4.0 OFDM Modifications for Remote PHY requirements as specified in [R-PHY] Annex F.

The MAC-NE MUST support CMTS DOCSIS 4.0 OFDM Modifications for Remote PHY requirements as specified in [R-PHY] Annex F.

Appendix I Acknowledgements

FMA Leadership Participants

At the time this specification was submitted to CableLabs for approval, the Flexible MAC Architecture (FMA) Working Group (WG) and Task Forces (TF) under the WG were organized with the following leadership roles and affiliations.

Flexible MAC Architecture (FMA) Working Group (WG)

FMA WG Co-Chair	Michael "Mike" Emmendorfer, CommScope Inc.
FMA WG Co-Chair	Jon Schnoor, Cable Television Laboratories, Inc.
FMA WG Co-Chair	Jeff Finkelstein, Cox Communications

Task Force 1: MAC Manager Interface

Task Force 1 Co-Chair	Douglas Johnson, Vecima Networks Inc.
Task Force 1 Co-Chair	Stephen Kraiman, CommScope Inc.

Task Force 2: PacketCable and Lawful Intercept

Task Force 2 Co-Chair	Rex Coldren, Vecima Networks Inc.
Task Force 2 Co-Chair	Dan Torbet, CommScope Inc.

Task Force 3: FMA OSS Interface Specification

Task Force 3 Co-Chair	Steve Burroughs, Cable Television Laboratories, Inc.
Task Force 3 Co-Chair	Dwain Frieh, CommScope Inc.

Task Force 4: Data Forwarding

Task Force 4 Co-Chair	Mike Patrick, Harmonic Inc.
Task Force 4 Co-Chair	Mircea Orban, CommScope Inc.

Task Force 5: FMA-Proactive Network Maintenance

Task Force 5 Co-Chair	Jason Rupe, Cable Television Laboratories, Inc.
Task Force 5 Co-Chair	Santhana Chari, CommScope Inc.

We wish to thank the following participants contributing directly to this document.

Contributor	Company Affiliation
Amit Aggarwal	Broadcom
Margo Dolas	Broadcom
Victor Hou	Broadcom
Niki Pantelias	Broadcom
Steve Burroughs	CableLabs
Kirk Erichsen	CableLabs
Steve Goeringer	CableLabs
Brian Hedstrom	CableLabs
Jason Rupe	CableLabs
Jon Schnoor	CableLabs
Yuan Tian	CableLabs
Philip Anderson	Charter
Pawel Sowinski	Cisco

Contributor	Company Affiliation
Mike Emmendorfer	CommScope
Steve Foley	CommScope
Dwain Frieh	CommScope
Stephen Kraiman	CommScope
Mircea Orban	CommScope
Dan Torbet	CommScope
Jeff Finklestein	Cox
Tadeusz Ciesielski	Falcon V Systems
Michael Patrick	Harmonic
Rex Coldren	Vecima
Douglas Johnson	Vecima
Utku Yilmaz	Vecima

Appendix II Revision History

The following Engineering Change was incorporated into CM-SP-FMA-SYS-I02-210526.

ECN Identifier	Accepted Date	Title of EC	Author
FMA-SYS-N-21.2161-2	4/28/2021	FMA-SYS I02 Candidate	Schnoor

The following Engineering Change was incorporated into CM-SP-FMA-SYS-I03-220126.

ECN Identifier	Accepted Date	Title of EC	Author
FMA-SYS-N-21.2216-2	1/6/2022	FMA System Spec I03 candidate	Schnoor

* * *