

OpenFlow™ Data Plane Abstraction (OF-DPA): Abstract Switch Specification

Version 2.0

October 6, 2014



Copyright © 2012, 2013, 2014 Broadcom Corp.

All Rights Reserved

Information contained within this document is the property of
Broadcom Corp.

No part of this document may be reproduced without
express written consent of Broadcom Corp.

All specifications and information are subject to change without notice.

Last Edit: September 4, 2015 11:22 AM

File Name: OF-DPA 2.0 - v1 - 6 Oct 2014.docx

Broadcom Corporation
3151 Zanker Road
San Jose, CA 95134

Broadcom®, the pulse logo, Connecting everything®, and the Connecting everything logo are among the registered trademarks of Broadcom Corporation and/or its subsidiaries in the United States, certain other countries, and/or the EU.

Any other trademarks or trade names mentioned are the property of their respective owners.

Confidential and Proprietary Information: This document is the property of Broadcom Corporation.

Revision History

Revision Number	Date	Change
2.0	6 Oct 2014	Initial release of OF-DPA 2.0.

This document represents the initial release of the complete specification for OF-DPA 2.0. It is intended for external distribution to solicit feedback and comments. As such, it is subject to change based on feedback received.

Table of Contents

1	Introduction	17
2	OF-DPA Components	18
3	The OF-DPA Abstract Switch	21
3.1	Abstract Switch Overview	21
3.2	Bridging and Routing	23
3.3	Data Center Overlay Tunnels	24
3.4	MPLS-TP Customer Edge Device	25
3.4.1	VPWS.....	26
3.4.2	MPLS-TP L2 VPN Groups	28
3.5	MPLS Label Edge Router (LER)	29
3.6	MPLS Label Switch Router (LSR)	31
3.7	MPLS-TP Protection Switching	32
3.8	Quality of Service (QoS)	33
3.8.1	QoS Concepts – Traffic Class and Color	33
3.8.2	Meter Table Operation	34
3.8.3	Bridging and Routing QoS	37
3.8.4	MPLS QoS.....	38
3.9	Operation, Administration, and Maintenance (OAM)	40
3.9.1	OAM Concepts – Maintenance Points	40
3.9.2	Network Protection Apps.....	42
3.9.3	MPLS-TP OAM.....	43
3.9.3.1	Ethernet OAM over MPLS-TP	45
3.9.3.2	G.8113.1 OAM for MPLS-TP	47
3.10	Protection Switching	50
4	OF-DPA Object Descriptions	52
4.1	Flow Tables	53
4.1.1	Ingress Port Flow Table.....	53
4.1.1.1	Flow Entry Types and Match Fields.....	53
4.1.1.2	Instruction Types.....	54
4.1.1.3	Actions.....	54
4.1.1.4	Counters and Flow Expiry.....	55
4.1.2	VLAN Flow Table	55
4.1.2.1	Flow Entry Types and Match Fields.....	55
4.1.2.2	Instruction Types.....	58
4.1.2.3	Actions.....	58
4.1.2.4	Counters and Flow Expiry.....	60
4.1.3	VLAN 1 Flow Table	60
4.1.3.1	Flow Entry Types and Match Fields.....	60
4.1.3.2	Instruction Types.....	61

4.1.3.3	Actions.....	62
4.1.3.4	Counters and Flow Expiry.....	63
4.1.4	MPLS L2 Port Flow Table.....	63
4.1.4.1	Flow Entry Types and Match Fields.....	64
4.1.4.2	Instruction Types.....	64
4.1.4.3	Counters and Flow Expiry.....	64
4.1.5	Termination MAC Flow Table	65
4.1.5.1	Flow Entry Types and Match Fields.....	65
4.1.5.2	Instruction Types.....	66
4.1.5.3	Counters and Flow Expiry.....	67
4.1.6	Bridging Flow Table.....	67
4.1.6.1	Flow Entry Types and Match Fields.....	69
4.1.6.2	Instruction Types.....	69
4.1.6.3	Action Set	70
4.1.6.4	Counters and Flow Expiration	70
4.1.7	Unicast Routing Flow Table	71
4.1.7.1	Flow Entry Types and Match Fields.....	71
4.1.7.2	Instruction Types.....	72
4.1.7.3	Action Set	73
4.1.7.4	Counters and Flow Expiration	73
4.1.8	Multicast Routing Flow Table	74
4.1.8.1	Flow Entry Types and Match Fields.....	74
4.1.8.2	Instruction Types.....	75
4.1.8.3	Action Set	76
4.1.8.4	Counters and Flow Expiration	76
4.1.9	MPLS Flow Tables	77
4.1.9.1	Flow Entry Types and Match Fields.....	77
4.1.9.2	Instruction Types.....	81
4.1.9.3	Actions.....	81
4.1.9.4	Counters and Flow Expiry.....	83
4.1.10	Maintenance Point Flow Table	84
4.1.10.1	Flow Entry Types and Match Fields	84
4.1.10.2	Instruction Types	84
4.1.10.3	Actions.....	85
4.1.10.4	Counters and Flow Expiry	85
4.1.11	Policy ACL Flow Table	86
4.1.11.1	Flow Entry Types and Match Fields	87
4.1.11.2	Instruction Types	90
4.1.11.3	Action List Actions	91
4.1.11.4	Action Set Actions.....	91
4.1.11.5	Counters and Flow Expiration.....	92
4.1.12	Color Based Actions Flow Table.....	93
4.1.12.1	Flow Entry Types and Match Fields	93
4.1.12.2	Instruction Types	94
4.1.12.3	Actions.....	94
4.1.12.4	Counters and Flow Expiry	95
4.2	Egress Flow Tables	95
4.2.1	Egress VLAN Flow Table.....	95
4.2.1.1	Flow Entry Types and Match Fields.....	95
4.2.1.2	Instruction Types.....	96
4.2.1.3	Actions.....	97
4.2.1.4	Counters and Flow Expiry.....	98
4.2.2	Egress VLAN 1 Flow Table.....	98

4.2.2.1	Flow Entry Types and Match Fields	98
4.2.2.2	Instruction Types	99
4.2.2.3	Actions	99
4.2.2.4	Counters and Flow Expiry	100
4.2.3	Egress Maintenance Point Flow Table	100
4.2.3.1	Flow Entry Types and Match Fields	100
4.2.3.2	Instruction Types	101
4.2.3.3	Actions	101
4.2.3.4	Counters and Flow Expiry	102
4.3	Group Table	102
4.3.1	OF-DPA L2 Interface Group Entries	103
4.3.1.1	Naming Convention	103
4.3.1.2	Action Buckets	104
4.3.1.3	Counters	104
4.3.2	OF-DPA L2 Unfiltered Interface Group Entries	105
4.3.2.1	Naming Convention	105
4.3.2.2	Action Buckets	105
4.3.2.3	Counters	106
4.3.3	OF-DPA L2 Rewrite Group Entries	106
4.3.3.1	Naming Convention	106
4.3.3.2	Action Buckets	107
4.3.3.3	Counters	107
4.3.4	OF-DPA L3 Unicast Group Entries	107
4.3.4.1	Naming Convention	108
4.3.4.2	Action Buckets	108
4.3.4.3	Counters	108
4.3.5	OF-DPA L2 Multicast Group Entries	109
4.3.5.1	Naming Convention	109
4.3.5.2	Action Buckets	109
4.3.5.3	Counters	110
4.3.6	OF-DPA L2 Flood Group Entries	110
4.3.6.1	Naming Convention	110
4.3.6.2	Action Buckets	111
4.3.6.3	Counters	111
4.3.7	OF-DPA L3 Interface Group Entries	111
4.3.7.1	Naming Convention	112
4.3.7.2	Action Buckets	112
4.3.7.3	Counters	112
4.3.8	OF-DPA L3 Multicast Group Entries	113
4.3.8.1	Naming Convention	113
4.3.8.2	Action Buckets	113
4.3.8.3	Counters	114
4.3.9	OF-DPA L3 ECMP Group Entries	114
4.3.9.1	Naming Convention	114
4.3.9.2	Action Buckets	115
4.3.9.3	Counters	115
4.3.10	OF-DPA L2 Overlay Group Entries	115
4.3.10.1	OF-DPA L2 Overlay Group Sub-Types	115
4.3.10.2	Naming Convention	117
4.3.10.3	Action Buckets	118
4.3.10.4	Counters	118
4.3.11	OF-DPA MPLS Interface Group Entry	118
4.3.11.1	Naming Convention	118

4.3.11.2	Action Buckets	119
4.3.11.3	Counters	119
4.3.12	OF-DPA MPLS Label Group Entries	120
4.3.12.1	Naming Convention	120
4.3.12.2	MPLS VPN Label Action Buckets	120
4.3.12.3	MPLS Tunnel Label 1 Action Buckets	122
4.3.12.4	MPLS Tunnel Label 2 Action Buckets	123
4.3.12.5	MPLS Swap Label Action Buckets	124
4.3.12.6	Counters	125
4.3.13	OF-DPA MPLS Fast Failover Group Entry	125
4.3.13.1	Naming Convention	125
4.3.13.2	Action Buckets	126
4.3.13.3	Counters	126
4.3.14	OF-DPA MPLS ECMP Group Entry	127
4.3.14.1	Naming Convention	127
4.3.14.2	Action Buckets	127
4.3.14.3	Counters	127
4.3.15	OF-DPA MPLS L2 Tag Group Entry	128
4.3.15.1	Naming Convention	128
4.3.15.2	Action Buckets	128
4.3.15.3	Counters	129
4.4	Meters	129
4.4.1	Meter Table Entries	130
4.4.2	Meter Bands	132
5	Configuration	133
5.1	Ports	133
5.1.1	Physical Ports	134
5.1.1.1	Features	134
5.1.1.2	Counters	137
5.1.2	Reserved Ports	138
5.1.3	Logical Ports	139
5.1.3.1	Overlay Tunnels	139
5.1.3.2	VXLAN Tunnel Logical Port Configuration	140
5.1.3.3	OAM Protection Liveness Logical Ports	144
5.2	Queues	144
5.2.1	Configuration	145
5.2.2	Counters	145
5.3	OAM Message Processing	145
5.3.1	MPLS-TP Ethernet OAM Configuration	146
5.3.2	MPLS-TP G.8113.1 OAM Configuration	147
5.4	Protection	149
5.4.1	MPLS-TP Linear Protection	149
6	Vendor Extension Features	151
6.1	Source MAC Learning	151
6.1.1	Controller Managed Learning	152
6.1.1.1	Configuration	152
6.2	Additional Group Properties	152

6.3	MTU Check.....	152
6.4	Table Numbering	153
6.5	Experimenter Features	154
6.5.1	OAM Data Plane Counter Table	154
6.5.2	Drop Status Action Table	155
6.5.3	MPLS Label Remark Action Tables	157
6.5.3.1	MPLS VPN Label Remark Action Table	157
6.5.3.2	MPLS Tunnel Label Remark Action Table	158
6.5.4	Actions	159
6.5.5	Match Fields.....	161
6.5.6	Color Set Meter Band.....	163
APPENDIX A :	References	164

List of Figures

Figure 1 OF-DPA Component Layering	18
Figure 2 Abstract Switch Objects Used for Bridging and Routing	23
Figure 3 Abstract Switch Objects Used for Overlay Tunnels.....	25
Figure 4 Abstract Switch Objects for MPLS-TP Initiation (VPWS)	26
Figure 5 Abstract Switch Objects for MPLS-TP Termination (VPWS).....	27
Figure 6 MPLS L2 VPN Groups	28
Figure 7 Abstract Switch Objects for MPLS L3 VPN Initiation.....	30
Figure 8 Abstract Switch Objects for MPLS L3 VPN Termination	31
Figure 9 Abstract Switch Objects Used for an MPLS LSR	31
Figure 10 MPLS LSR Groups.....	32
Figure 11 MPLS 1:1 Protection	33
Figure 12 Token Bucket Operation	34
Figure 13 TrTCM Meter Operation	35
Figure 14 SrTCM Meter Operation.....	36
Figure 15 Bridging and Routing Abstract Switch With QoS Objects.....	38
Figure 16 MPLS-TP VPWS Initiation with QoS Objects	38
Figure 17 MPLS-TP VPWS Tunnel Termination with QoS Objects.....	39
Figure 18 MPLS L3 VPN Initiation with QoS Objects	39
Figure 19 MPLS L3 VPN Termination with QoS Objects.....	40
Figure 20 MPLS LSR with QoS Objects.....	40
Figure 21 OAM MEP and MIP Examples	41
Figure 22 OAM and Protection Overview	42
Figure 23 MPLS-TP Service OAM Examples.....	44
Figure 24 Ethernet OAM Over MPLS-TP Packet Format.....	45
Figure 25 MPLS-TP Initiation - Ethernet over MPLS-TP OAM Data Frame	45
Figure 26 MPLS-TP Termination - Ethernet over MPLS-TP OAM Data Frame...	46
Figure 27 MPLS-TP Initiation - Ethernet over MPLS-TP OAM PDU	46
Figure 28 MPLS-TP Termination - Ethernet over MPLS-TP OAM PDU	47
Figure 29 OAM MPLS-TP G.8113.1 Packet Formats.....	48
Figure 30 MPLS-TP Initiation - G.8113.1 OAM Data Frame	48
Figure 31 MPLS-TP Termination - G.8113.1 OAM Data Frame	49
Figure 32 MPLS-TP - G.8113.1 OAM LSR Data Frame.....	49
Figure 33 MPLS-TP - G.8113.1 OAM PDU	49
Figure 34 Protection Switching Process.....	51
Figure 35 OF-DPA L2 Overlay Flood Over Unicast Tunnels	116
Figure 36 OF-DPA L2 Overlay Flood Over Multicast Tunnels.....	116
Figure 37 OF-DPA L2 Overlay Multicast Over Unicast Tunnels.....	117
Figure 38 OF-DPA L2 Overlay Multicast Over Multicast Tunnels.....	117
Figure 39 Meter Entry Example (TrTCM)	130
Figure 40 Port Properties Configuration.....	134
Figure 41 OpenFlow Feature Sub-Classes	135
Figure 42 Tunnel Logical Port Configuration	140
Figure 43 VXLAN Tunnel Configuration	142
Figure 44 OAM Ethernet Fault Management Configuration	146
Figure 45 OAM Ethernet Performance Monitoring Configuration	147

Figure 46 OAM G.8113.1 Fault Management Configuration 148
Figure 47 OAM G.8113.1 Performance Monitoring Configuration..... 149
Figure 48 MPLS-TP Linear Protection Configuration 150

List of Tables

Table 1 OpenFlow 1.3.4 Features Required by OF-DPA v2.0	19
Table 2 TrTCM Color-Aware Operation.....	35
Table 3 SrTCM Color-Aware Operation	36
Table 4 Ingress Port Flow Table Entry Types	53
Table 5 Ingress Port Flow Table Match Fields	54
Table 6 Ingress Port Flow Table Instructions	54
Table 7 Ingress Port Flow Table Action List	55
Table 8 Ingress Port Flow Table Counters	55
Table 9 VLAN Flow Table Flow Entry Types.....	56
Table 10 VLAN Flow Table Match Fields	57
Table 11 VLAN Flow Table Instructions	58
Table 12 VLAN Flow Table Action List Actions	59
Table 13 VLAN Flow Table Action Set Actions	59
Table 14 VLAN Flow Table Counters	60
Table 15 VLAN 1 Flow Table Flow Entry Types.....	60
Table 16 VLAN 1 Flow Table Match Fields	61
Table 17 VLAN 1 Flow Table Instructions	61
Table 18 VLAN 1 Flow Table Action List Actions	62
Table 19 VLAN 1 Flow Table Action Set Actions	62
Table 20 VLAN 1 Flow Table Counters	63
Table 21 VLAN 1 Flow Table Expiry.....	63
Table 22 MPLS L2 Port Metadata Naming Convention.....	63
Table 23 MPLS L2 Port Flow Table Flow Entry Types	64
Table 24 MPLS L2 Port Flow Table Match Fields	64
Table 25 MPLS L2 Port Flow Table Instructions	64
Table 26 MPLS L2 Port Flow Table Counters	65
Table 27 Termination MAC Flow Table Entry Types	65
Table 28 Termination MAC Flow Table Match Fields.....	66
Table 29 Termination MAC Flow Table Instruction Set	67
Table 30 Termination MAC Flow Table Counters	67
Table 31 Tunnel Id Metadata Naming Convention	68
Table 32 Bridging Flow Table Flow Entry Types	68
Table 33 Bridging Flow Table Match Fields	69
Table 34 Bridging Flow Table Instructions	69
Table 35 Bridging Flow Table Action Set	70
Table 36 Bridging Flow Table Counters	70
Table 37 Bridging Flow Table Flow Entry Expiration	71
Table 38 Unicast Routing Flow Table Entry Types	71
Table 39 Unicast Routing Flow Table IPv4 Header Match Fields	72
Table 40 Unicast Routing Flow Table IPv6 Header Match Fields	72
Table 41 Unicast Routing Flow Table Instructions	73
Table 42 Unicast Routing Flow Table Action Set	73
Table 43 Unicast Routing Flow Table Counters	73
Table 44 Unicast Routing Flow Table Flow Entry Expiration.....	74
Table 45 Multicast Routing Flow Table Entry Types	74

Table 46 Multicast Routing Flow Table IPv4 Match Fields	74
Table 47 Multicast Routing Flow Table IPv6 Match Fields	75
Table 48 Multicast Routing Flow Table Instructions	75
Table 49 Multicast Routing Flow Table Action Set	76
Table 50 Multicast Routing Flow Table Counters	76
Table 51 Multicast Routing Flow Table Flow Entry Expiration	76
Table 52 MPLS Flow Table 0 Flow Entry Types	77
Table 53 MPLS Flow Table 1 and 2 Flow Table Entry Types	78
Table 54 MPLS Flow Table Match Fields	80
Table 55 MPLS Flow Table Instructions	81
Table 56 MPLS Flow Table Action List	81
Table 57 MPLS Flow Table Action Set Actions	83
Table 58 MPLS Flow Table Counters	83
Table 59 Maintenance Point Flow Table Entry Types	84
Table 60 Maintenance Point Flow Table Match Fields	84
Table 61 Maintenance Point Flow Table Instructions	84
Table 62 Maintenance Point Flow Table Actions	85
Table 63 Maintenance Point Flow Table Counters	85
Table 64 Policy ACL Flow Table Flow Entry Types	86
Table 65 Policy ACL Flow Table IPv4 Match Fields	87
Table 66 Policy ACL Flow Table IPv6 Match Fields	88
Table 67 Policy ACL Flow Table Instructions	90
Table 68 Policy ACL Flow Table Action List Actions	91
Table 69 Policy ACL Flow Table VLAN Flow Entry Action Set	91
Table 70 Policy ACL Flow Table Tunnel Flow Entry Action Set	92
Table 71 Policy ACL Flow Table Counters	92
Table 72 Policy ACL Flow Table Expiry	93
Table 73 Color Based Actions Flow Table Entry Types	93
Table 74 Color Based Actions Flow Table Match Fields	94
Table 75 Color Based Actions Flow Table Instructions	94
Table 76 Color Based Actions Flow Table Actions	94
Table 77 Color Based Actions Flow Table Counters	95
Table 79 Egress VLAN Flow Table Flow Entry Types	95
Table 80 Egress VLAN Flow Table Match Fields	96
Table 81 Egress VLAN Flow Table Instructions	97
Table 82 Egress VLAN Flow Table Action List	97
Table 83 Egress VLAN Flow Table Counters	98
Table 85 Egress VLAN 1 Flow Table Flow Entry Types	98
Table 86 Egress VLAN 1 Flow Table Match Fields	99
Table 87 Egress VLAN 1 Flow Table Instructions	99
Table 88 Egress VLAN 1 Flow Table Action List	99
Table 89 Egress VLAN 1 Flow Table Counters	100
Table 91 Egress Maintenance Point Flow Table Entry Types	101
Table 92 Egress Maintenance Point Flow Table Match Fields	101
Table 93 Egress Maintenance Point Flow Table Instructions	101
Table 94 Egress Maintenance Point Flow Table Actions	101
Table 95 Egress Maintenance Point Flow Table Counters	102
Table 97 OF-DPA Group Table Entry Identifier Naming Convention	102

Table 98 OF-DPA L2 Interface Group Entry Type Naming Convention	103
Table 99 OF-DPA L2 Interface Group Entry Bucket Actions	104
Table 100 OF-DPA L2 Interface Group Entry Counters	104
Table 101 OF-DPA L2 Unfiltered Interface Group Naming Convention	105
Table 102 OF-DPA L2 Unfiltered Interface Group Bucket Actions	105
Table 103 OF-DPA L2 Unfiltered Interface Group Entry Counters	106
Table 104 OF-DPA L2 Rewrite Group Entry Type Naming Convention	106
Table 105 OF-DPA L2 Rewrite Group Entry Bucket Actions	107
Table 106 OF-DPA L2 Rewrite Group Entry Counters	107
Table 107 OF-DPA L3 Unicast Group Entry Naming Conventioin	108
Table 108 OF-DPA L3 Unicast Bucket Actions	108
Table 109 OF-DPA L3 Unicast Group Entry Counters	109
Table 110 OF-DPA L2 Multicast Group Entry Type Naming Convention	109
Table 111 OF-DPA L2 Multicast Bucket Actions	109
Table 112 OF-DPA L2 Multicast Group Entry Counters	110
Table 113 OF-DPA L2 Flood Group Entry Naming Convention	110
Table 114 OF-DPA L2 Flood Bucket Actions	111
Table 115 OF-DPA L2 Flood Group Entry Counters	111
Table 116 OF-DPA L3 Interface Group Entry Type Naming Convention	112
Table 117 OF-DPA L3 Interface Group Entry Bucket Actions	112
Table 118 OF-DPA L3 Interface Group Entry Counters	112
Table 119 OF-DPA L3 Multicast Group Entry Naming Convention	113
Table 120 OF-DPA L3 Multicast Bucket Actions	113
Table 121 OF-DPA L3 Multicast Group Entry Counters	114
Table 122 OF-DPA L3 ECMP Group Entry Naming Convention	114
Table 123 OF-DPA L3 ECMP Group Entry Bucket Actions	115
Table 124 OF-DPA L3 ECMP Group Entry Counters	115
Table 125 OF-DPA L2 Overlay Group Entry Naming Convention	117
Table 126 OF-DPA L2 Overlay Group Sub-Type Entry Bucket Actions	118
Table 127 OF-DPA L2 Overlay Group Sub-Type Entry Counters	118
Table 128 OF-DPA MPLS Interface Group Entry Naming Convention	118
Table 129 OF-DPA MPLS Interface Group Entry Bucket Actions	119
Table 130 OF-DPA MPLS Interface Group Type Entry Counters	119
Table 131 OF-DPA MPLS Label Group Entry Naming Convention	120
Table 132 OF-DPA MPLS L2 VPN Label Group Bucket Actions	120
Table 133 OF-DPA MPLS Tunnel Label 1 Group Bucket Actions	122
Table 134 OF-DPA MPLS Tunnel Label 2 Actions	123
Table 135 OF-DPA MPLS Swap Label Actions	124
Table 136 OF-DPA MPLS Label Group Sub-Type Entry Counters	125
Table 137 OF-DPA MPLS Fast Failover Group Entry Naming Convention	125
Table 138 OF-DPA MPLS Fast Failover Group Entry Bucket Actions	126
Table 139 OF-DPA MPLS Fast Failover Tag Group Entry Counters	126
Table 140 OF-DPA MPLS ECMP Group Entry Naming Convention	127
Table 141 OF-DPA MPLS ECMP Group Entry Bucket Actions	127
Table 142 OF-DPA MPLS ECMP Group Entry Counters	128
Table 143 OF-DPA MPLS L2 Tag Group Entry Naming Convention	128
Table 144 OF-DPA L2 Tag Group Entry Bucket Actions	128
Table 145 OF-DPA MPLS L2 Tag Group Entry Counters	129

Table 146 OF-DPA Meter Table Entry Parameters.....	130
Table 147 OF-DPA Meter Entry Naming Convention.....	131
Table 148 Meter Entry Counters	131
Table 149 Meter Band Configuration Parameters.....	132
Table 150 Meter Band Counters	132
Table 151 Port Type Numbering Conventions	133
Table 152 OF-DPA Port Features	135
Table 153 Port Features Bitmap	136
Table 154 OF-DPA Physical Port Counters	137
Table 155 OF-DPA Reserved Ports	138
Table 156 VXLAN Tunnel Endpoint Logical Port Configuration Parameters.....	143
Table 157 VXLAN Access Logical Port Configuration Parameters	143
Table 158 VXLAN Tenant Configuration Parameters.....	144
Table 159 VXLAN Next Hop Configuration Parameters.....	144
Table 160 OF-DPA Queue Configuration Parameters	145
Table 161 OF-DPA Queue Counters	145
Table 162 Source MAC Learning Feature Configuration	152
Table 163 Flow Table Number Assignments	153
Table 164 OAM Data Plane Counter Table Entry	154
Table 165 Drop Status Table Entry	156
Table 166 MPLS Label Remark Table Entry Fields	157
Table 167 OF-DPA Experimenter Actions.....	159
Table 168 OF-DPA Experimenter Match Fields	161

Terminology

Term	Acronym	Description
Capabilities		Switch features as understood by controllers.
Flow		Sequence of packets with the same selection of header field values. Flows are unidirectional.
Flow Table		OpenFlow flow table as defined in the OpenFlow 1.3.3 specification
Flow Entry		Entry in an OpenFlow flow table with its match fields and instructions.
Hybrid Switch		Switch that has OpenFlow programmability in addition to legacy control plane features. There is no standard hybrid switch model although different approaches have been identified, ranging from integrated to completely independent pipelines.
Group Table		The OpenFlow group table, consisting of group table entries.
Meter Table		The OpenFlow meter table, consisting of meter table entries.
OpenFlow Logical Switch	OFLS	A set of OpenFlow resources that can be managed as a single entity, includes a datapath and a control channel.
Open Networking Foundation	ONF	Open Networking Foundation (ONF) is a user-driven organization dedicated to the promotion and adoption of SDN through open standards development, primarily OpenFlow.
Software Defined Networking	SDN	The principles of SDN as defined by ONF are: separation of control and forwarding functions; logically centralized intelligence; programmable; open-standards based; and vendor neutral programming interfaces [1].
Table Type Pattern	TTP	Formal description of an OpenFlow 1.3 abstract switch in terms of programmable pipeline objects.
Unit		A member switch within a chassis or switch stack

1 INTRODUCTION

This document specifies an OpenFlow abstract switch model – called OpenFlow Data Plane Abstraction (OF-DPA) for Broadcom® Ethernet switch devices. The primary goal of this open specification is to enable Broadcom-based devices to be programmable using the OpenFlow protocol as an OpenFlow Logical Switch, and in so doing foster further growth of the ecosystem of open source and commercial OpenFlow agents and controllers that can be utilized to enable wider OpenFlow-based network infrastructure deployments.

This model is based on the stable OpenFlow 1.3.4 specification [7] and utilizes its provisions to provide access to multiple tables implemented in Broadcom switch Application-Specific Integrated Circuits (ASICs). It is formally defined as a Table Type Pattern (TTP) using the notation specified in the ONF OpenFlow Table Type Patterns specification [8]. The intent is to facilitate general availability of production-quality OpenFlow switches from product vendors as well as provide a development platform for use in academic and industrial research networks.

This document represents the first feature-complete specification for OF-DPA version 2.0. OF-DPA 2.0 incorporates and adds functionality to OF-DPA version 1.0 [9]. It is published openly and meant to be used alongside the *Open Flow Data Plan Abstraction (OF-DPA) API Guide and Reference Manual*¹ for developing OpenFlow 1.3.4 agents and controllers. While the specification is deemed complete for features supported in OF-DPA 2.0, Broadcom solicits feedback and comments at all times to further improve the specification. As such, it may be subject to change based on feedback received from interested parties.

This document assumes familiarity with OpenFlow 1.3.4 and the goals of related Software Defined Networking (SDN) technologies.

¹ Available as an HTML document in the OF-DPA v2.0 software release package

2 OF-DPA COMPONENTS

OF-DPA is a software component that provides a hardware adaption layer between OpenFlow and Broadcom switch ASICs. It is layered above the Broadcom switch software development switch (SDK) that, in turn, provides the driver for configuring, programming, and controlling the Broadcom switch ASICs.

The OF-DPA API, as defined in the *Open Flow Data Plan Abstraction (OF-DPA) API Guide and Reference Manual*, presents a specialized hardware abstraction layer (HAL) that allows programming Broadcom ASICs using OpenFlow abstractions. However, it does not process OpenFlow protocol messages. To create a complete OpenFlow switch using OF-DPA, an OpenFlow agent is required. In addition, an OpenFlow Controller is required to field an OpenFlow network deployment using OF-DPA enabled switches. Figure 1 illustrates the relationship of OF-DPA with the other OpenFlow system components.

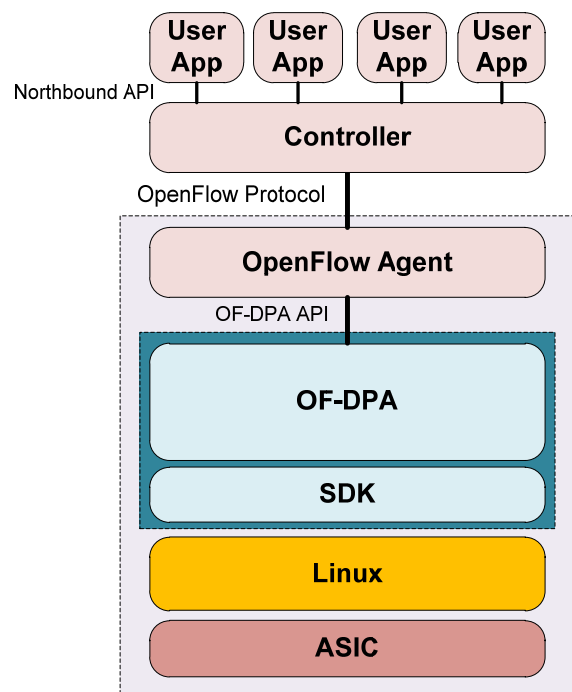


Figure 1 OF-DPA Component Layering

In Figure 1, user applications obtain services from an OpenFlow Controller via a Northbound API. The Northbound API enables applications to communicate with and control one or more OpenFlow switches. In addition, the Controller would likely provide advanced services such as discovery and enumeration of OpenFlow switches, along with a network-wide database of network resources including internal and external interfaces.

In Figure 1, the OpenFlow Controller communicates with an OpenFlow switch using OpenFlow as the Southbound API. The Controller maintains a (secure) channel with each OpenFlow switch over which it exchanges OpenFlow protocol messages. At the switches, OpenFlow agents maintain their end of the (secure) channel, processing received OpenFlow protocol messages and sending OpenFlow messages in response to local events.

Controllers are available from multiple sources. Any controller should be usable as long as it supports the OpenFlow 1.3.4 features defined by the Open Networking Foundation (ONF) specification and required by OF-DPA. These required features are listed in Table 1². In addition to commercial products, there are a number of readily available open source controllers that support OpenFlow 1.3 and later, including Ryu [18] and OpenDaylight [19].

Table 1 OpenFlow 1.3.4 Features Required by OF-DPA v2.0

Feature	Description
Pipeline Match Fields	Metadata fields that accompany the packet during pipeline processing but are not parsed from packet headers.
Experimenter Protocol Extensions	The Experimenter facility provides a standard way to extend the OpenFlow protocol to support additional functionality. OF-DPA 2.0 defines new Experimenter symmetric messages, multipart messages, flow table match fields, actions, and Meter bands.
Select and Fast Failover Group Types	These are optional types in OpenFlow 1.3.4 but required.
LOCAL Reserved Port	This is optional in OpenFlow 1.3.4 but required.

The OF-DPA API represents hardware objects to the agent in terms of objects such as flow tables, group table entries, queues, and ports that can be programmed using the protocol described in the OpenFlow 1.3.4 specification. The OF-DPA Abstract Switch provides instances of many of the OFLS objects defined in the OpenFlow 1.3.4 specification. As OF-DPA maintains the state that maps OpenFlow to the hardware, an agent is expected to do a relatively straightforward translation of OpenFlow messages into OF-DPA API calls and vice-versa, while maintaining a minimal amount of state.

² OF-DPA also has configuration APIs that require deployments to define a configuration approach.

In addition to standard objects, OF-DPA 2.0 defines new ancillary table objects needed to support use cases such as MPLS-TP, OAM, and QoS. These enhance the standard OFLS with features such as per-flow loss management counters and per-flow packet dropping. Analogous to the per-flow meter table introduced in OpenFlow 1.3, they are managed using new protocol message types, but invoked in action lists or action sets using new actions. To support certain OAM functions, OF-DPA 2.0 makes use of the Egress Tables planned for OpenFlow 1.5. Egress tables permit match action processing after groups are applied and the output port is decided. All OF-DPA 2.0 extensions are programmable using existing OpenFlow 1.3.4 protocol facilities such as experimenter message types³. Refer to Section 6 “Vendor Extension Features” on page 151 for detailed descriptions of these extensions.

The OF-DPA 2.0 Abstract Switch is formally defined using the Table Type Pattern (TTP) framework [8] developed by the ONF Forwarding Abstractions Working Group. A TTP is an OpenFlow abstract switch model that describes the forwarding behaviors that can be programmed by a controller. In particular, the TTP framework permits the application writer to express switch requirements, and a controller and switch to agree on supported features at run time.

Once the application gains access to an OF-DPA enabled network device, it can orchestrate and implement packet processing functions by adding flow entries to OpenFlow flow tables with action lists and/or action sets for packet editing and forwarding. Most packet forwarding uses OpenFlow group entries. An application can interrogate the status of OpenFlow ports and queues, and receive events such as port state changes or flow expiration through services of the Controller via the OpenFlow agent on the switch.

The next section of this document provides an overview of the OF-DPA Abstract Switch using diagrams to illustrate the objects and the relationships between them for particular packet flows. For a detailed description of the objects, flows, and relationships consult the TTP⁴. For details of the OF-DPA API, consult the [Open Flow Data Plan Abstraction \(OF-DPA\) API Guide and Reference Manual](#).

³ Some of these are forward-looking in that they anticipate features proposed for OpenFlow 1.5.

⁴ The TTP is included as a text document in the OF-DPA v2.0 software release package.

3 THE OF-DPA ABSTRACT SWITCH

The OF-DPA Abstract Switch is a specialization of the OpenFlow 1.3.4 OFLS with some extensions. This section describes the OF-DPA Abstract Switch in terms of OpenFlow abstract objects as visible to the OpenFlow controller.

The OF-DPA Abstract Switch objects can be thought of as programming points for Broadcom ASICs. These include flow tables with action sets, group table entries, logical and physical ports, and queues. The OF-DPA adaptation layer provides support for OpenFlow specific state, for example, statistics counters. It also maps OpenFlow objects to hardware and manages hardware resources.

Supporting OpenFlow in switch hardware involves some tradeoffs. As has been noted elsewhere, the generality promised by OpenFlow can come at a cost of latency, as well as cost and power inefficiencies. In addition, to effectively use this generality a specific multi-table pipeline needs to be designed and configured first. The OF-DPA Abstract Switch may be viewed as coming pre-configured and optimized to support single pass, full bandwidth packet processing performance that makes efficient use of the hardware and available table memory resources, trading off unrestricted generality in favor of latency, performance and cost, while enabling a logically centralized control plane with programming flexibility [1].

The OF-DPA Abstract Switch includes functionality to support: bridging; routing; data center tunnel gateways; MPLS provider edge and label switch routing; and QoS use case packet flows. Although all of these packet flows are available simultaneously in the same switch, different packet flows utilize different sequences of tables and group entries.

The OF-DPA Abstract Switch packet flows described in this section are those supported in OF-DPA 2.0. These flows expose a proper subset of the functionality available in Broadcom ASICs. Future versions of OF-DPA are expected to support additional features and packet flow use cases.

3.1 Abstract Switch Overview

OF-DPA flow tables accommodate specific types of flow entries with associated semantic rules, including constraints such as which match fields are available, which instructions and actions are supported, how priorities can be assigned to flow entries, which next table(s) flow entries can go to, and so forth.

The flow tables conform to the OpenFlow 1.3.4 specification [7]⁵. In addition to normal flows, two types of special flow entries are supported as follows:

- **Built-in.** Built-in flow entries come pre-installed in specific tables. They are visible to the controller but cannot be modified or deleted. Built-in entries have pre-assigned match fields, priority, and cookie values. They are typically used for default entries⁶.
- **Automatic.** Automatic flow entries are added by the switch as a side effect of the controller adding a flow entry.⁷ They are visible to the controller but cannot be directly modified or deleted except by modifying or deleting the rule that caused the automatic entry to be added. Match fields and priority are pre-determined, and the switch assigns the same cookie value as the initiating rule.

OF-DPA provides API calls to support interrogating tables for capabilities. These capabilities can include supported match fields, actions, instructions, etc. They also include status properties such as current resource usage.

In addition to flow tables, OF-DPA defines a set of group table entry types. The OpenFlow 1.3.4 specification defines four types of groups: indirect, all, select, and fast failover. OF-DPA further types group entries according to how they can be used in packet flows. This is done using specific naming conventions, properties, and supported action buckets. All OF-DPA group table entry types can be programmed using OpenFlow 1.3.4 as long as group mods respect the typing conventions.

One motivation for group typing is supporting fundamental differences in use case requirements. For example, in order to support “one-arm” routing using group table entries there needed to be a way to override OpenFlow’s default source removal and allow routing back to IN_PORT. This was accomplished by defining L3 group entry types with different properties from L2 groups. Group entry typing is also useful to enforce constraints on group entry chains and for VLAN configuration on physical ports.

Remember that OF-DPA tables are programming abstractions and do not necessary directly correspond one-to-one with hardware tables. However they are designed to faithfully capture both use case requirements and the hardware packet flow semantics, while being straightforward to program from standard controllers.

Users must program flow tables and group entries according to the allowed entry types. The OF-DPA API validates calls and returns errors if constraints and/or conventions are violated. This includes the requirement that objects must exist before they can be referenced from other objects. The OpenFlow agent that interfaces to OF-DPA may also do some argument validation and execute local iterative procedures.

Many forwarding and editing actions for encapsulation/push and field modify are programmed using one or more action buckets in group table entries. This not only proves to be a very efficient and modular

⁵ Some features however borrow from extensions proposed for OpenFlow 1.5. Experimenter fields, actions, and messages are defined so that an OpenFlow 1.3.4 controller can program these features.

⁶ OVS uses built-in rules, and they are described in the ONF TTP specification [8].

⁷ OpenFlow 1.5 introduces “flow template” actions to automatically add rules based on a rule match. This is traffic and not controller driven and should not be confused with automatic rules.

programming approach, in that the controller can optimize hardware resources better than the switch, but the controller intrinsically has more CPU power and memory than the control processor on a typical switch for this task. The controller also understands what the application is trying to do, especially when programming requires updating multiple tables. However, when compared with OpenFlow 1.0 programming, it may require more messages between the controller and switches, since more objects need to be programmed. It also potentially requires the controller to keep track of more switch state, although this state can be interrogated as needed.

Some functionality must be configured using logical ports. In general, this is to handle packet processing functions that OpenFlow 1.3 has historically not equipped to handle, such as adding and removing encapsulation headers. VXLAN data center overlay tunnels are handled by specialized configuration of logical ports rather than by directly programming flow and group entries.

Note: OF-DPA does not support reassembly. The switch drops fragments by default.

The following sections are organized by use case packet flow.

3.2 Bridging and Routing

The OF-DPA Abstract Switch objects that can be programmed for bridging and routing are shown in Figure 2. This packet flow is essentially the same as in OF-DPA 1.0 [9].

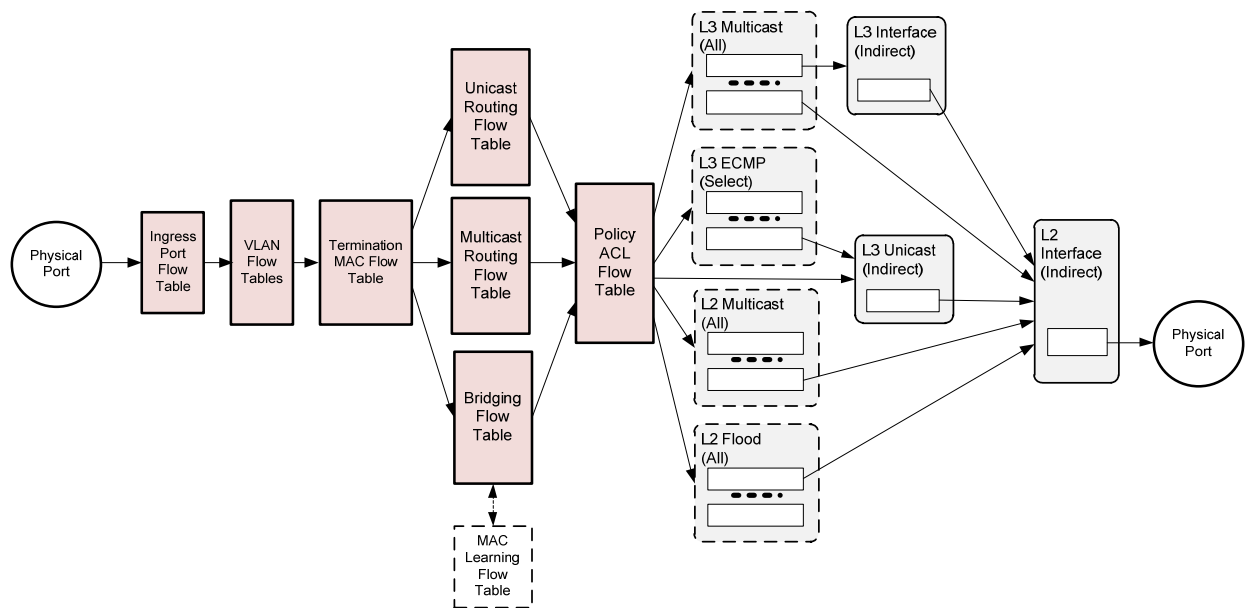


Figure 2 Abstract Switch Objects Used for Bridging and Routing

Packets enter and exit the pipeline on physical ports local to the switch. The Ingress Port Flow Table (table 0) is always the first table to process a packet. Flow entries in this table can distinguish traffic from different types of input ports by matching associated Tunnel Id metadata. Normal bridging and routing packets from physical ports have a Tunnel Id value of 0. To simplify programming this table provides a default rule that passes through packets with Tunnel Id 0 that do not match any higher priority rules.

All packets in the Bridging and Routing flow must have a VLAN. The VLAN Flow Table can do VLAN filtering for tagged packets and VLAN assignment for untagged packets. If the packet has more than one VLAN tag, the outermost VLAN Id is the one used for forwarding.

The Termination MAC Flow Table matches destination MAC addresses to determine whether to bridge or route the packet and, if routing, whether it is unicast or multicast. MAC learning is supported using a “virtual” flow table that is logically synchronized with the Bridging Flow Table.

When MAC learning is enabled, OF-DPA does a lookup in the Bridging Flow Table using the source MAC, outermost VLAN Id, and IN_PORT. A miss is reported to the controller using a Packet In message. Logically this occurs before the Termination MAC Flow Table lookup. The MAC Learning Flow Table cannot be directly read or written by the controller.⁸

The ACL Policy Flow Table can perform multi-field wildcard matches, analogous to the function of an ACL in a conventional switch.

OF-DPA makes extensive use of OpenFlow Group entries, and most forwarding and packet edit actions are applied based on OpenFlow group entry buckets. Groups support capabilities that are awkward or inefficient to program in OpenFlow 1.0, such as multi-path and multicast forwarding, while taking advantage of functionality built into the hardware.

3.3 Data Center Overlay Tunnels

OF-DPA objects for Data Center Overlay tunnels are shown in Figure 3. These are unchanged from OF-DPA 1.0.

Data Center Overlay Tunnel processing forwards traffic for different tenants in isolated forwarding domains. Packets are forwarded based on a non-zero Tunnel-id value that identifies a particular tenant forwarding domain. A rule must be programmed in the Ingress Port Flow table to admit Data Center Overlay tunnel packets for a particular tenant. In addition, flow entries in the Bridging Flow Table and ACL Policy Flow Table match tunnel traffic by Tunnel Id and not VLAN Id.

OF-DPA uses a naming convention for Tunnel-id metadata where the high order 16 bits identify the tunnel type. This permits flow entries to distinguish among different types of tunnel traffic.

⁸ The MAC Learning Flow Table has a “virtual” table number which is reported to the Controller in a table miss Packet-In message. It does not appear as part of the pipeline since its table number assignment would violate the OpenFlow requirement for packets to traverse tables in monotonically increasing order.

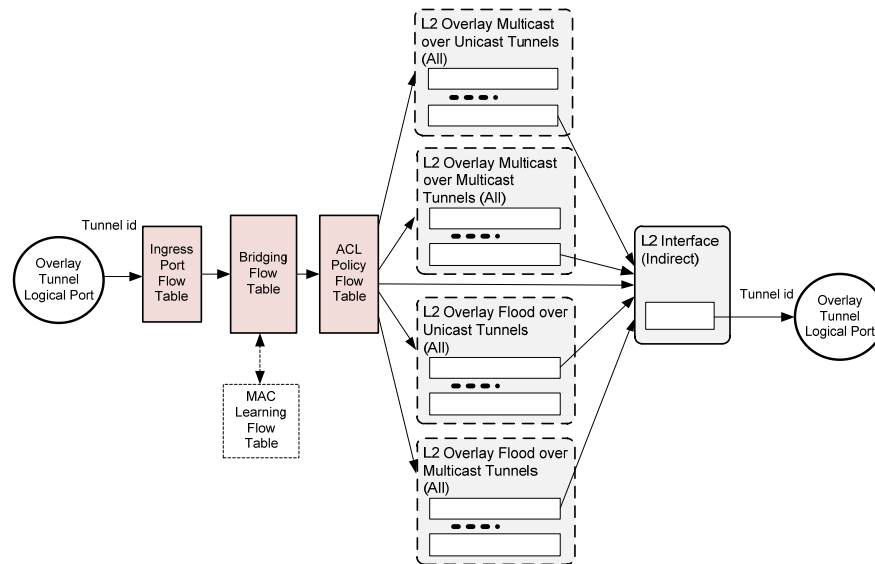


Figure 3 Abstract Switch Objects Used for Overlay Tunnels

In OF-DPA 1.0 Data Center Overlay tunneling is implemented using a combination of configuration APIs, logical ports, and flow tables. In particular the Abstract Switch pipeline receives inner packets from logical ports along with Tunnel Id metadata after encapsulation headers have been removed, and forwards packets to logical ports for encapsulation with Tunnel Id metadata. The tunnel endpoints themselves are handled outside of OpenFlow.

OF-DPA 1.0 supplied an information model for tunnel configuration intended to be used as an extension to OF-Config. The configuration differed from that in OF-Config 1.2 [10] in that it is both able to relate Tunnel Id to VxLAN VNI and allow the pipeline to make forwarding decisions taking into account different VxLAN tunnel endpoints (VTEP). It also configures the forwarding behavior for overlay frames instead of relying on some unspecified routing function. The OF-DPA 1.0 information model was described used the same approach as was used for OF-Config to facilitate incorporation into its YANG model, but could but it also be used with ovsdb [11].

Configuring tunnels this way depends on having a suitable way to do configuration. In general this would require using a configuration protocol and associated configuration agent in addition to requiring an OpenFlow agent.

3.4 MPLS-TP Customer Edge Device

This section describes the OF-DPA objects for MPLS-TP L2 Customer Edge (CE). This functionality is new in OF-DPA 2.0.

The OpenFlow 1.3.4 OFLS does not intrinsically provide sufficient functionality to support MPLS-TP. As a result OF-DPA 2.0 must introduce some extensions, most of which are in the form of additional pipeline metadata match fields and related actions. As opposed to the approach used for Data Center Overlay tunneling, OF-DPA implements MPLS-TP tunneling directly in the pipeline. To support this OF-DPA 2.0

introduces new L2 header encapsulation and de-encapsulation actions, since OpenFlow 1.3.4 does not support the required actions to push and pop entire Ethernet headers.

Note: These extensions are modeled after extensions proposed for OpenFlow 1.5. The experimenter programming code points are described in Section 6⁹

3.4.1 VPWS

Figure 4 shows just the OF-DPA Abstract Switch objects used in VPWS initiation packet flows. The MPLS-TP L2 VPN groups that are enclosed within the dotted lines are used for 1:1 linear and ring protection switching and for label processing, and are discussed in more detail in Section 3.4.2.

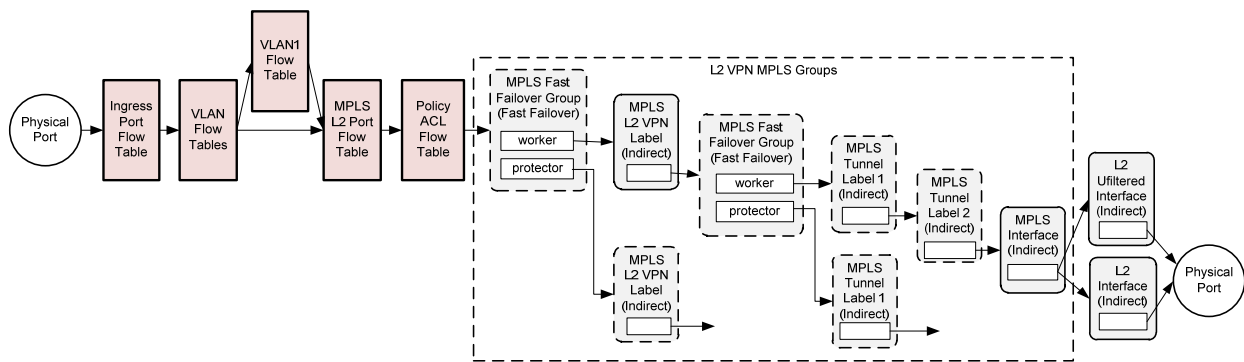


Figure 4 Abstract Switch Objects for MPLS-TP Initiation (VPWS)

VPWS is a point-to-point service. In forwarding terms the pseudo-wire directly determines the packet destination without the need for a bridging lookup. VPWS does not require learning, flooding, or multicast support.

For VPWS initiation, packets are classified to a particular customer pseudo-wire based on some combination of ingress port and packet VLAN header fields. To accommodate IEEE 802.1Q VLAN stacking user packets may need to be classified based on both outer and inner VLAN tags. However, OpenFlow 1.3.4 only supports matching the outermost VLAN tag. OF-DPA uses the VLAN 1 Flow Table to match two tags. A flow entry in the VLAN Flow table is first used to match the outermost VLAN, with an action list that has a Set-Field action to set a new pipeline match field (OVID) to the value matched and then do a Pop-VLAN action to make the inner tag the new outermost tag. The rule Goto-Table instruction specifies the VLAN 1 Flow Table where a flow entry can then match both VLAN values.¹⁰

To facilitate features that will be introduced in later OF-DPA releases, such as QoS classification or VPLS, another pipeline match field, MPLS L2 Port, is defined to represent a logical ingress interface for the pseudo-wire. An MPLS L2 Port can represent either a local attachment (UNI) or network (NNI) logical

⁹ At this writing OpenFlow 1.5 is a work in progress. As a result, there may be changes to OF-DPA to align OpenFlow 1.5 as features become better defined. If history is any indication, OpenFlow 1.5 support in agents and controllers may not be available for some time.

¹⁰ This approach was used rather than defining new VLAN header match fields to avoid issues with enabling OpenFlow 1.3.4 to differentiate between tags by TPID.

ingress interface. OF-DPA uses a naming convention for MPLS L2 Port identifiers that partitions into UNI and NNI ranges in order to distinguish the source type in flow entries.

In addition a Tunnel Id metadata value must be assigned to the flow. Analogous to the case with Data Center Overlay Tunnel logical ports, the MPLS forwarding pipeline implements an isolated forwarding domain for a particular customer pseudo-wire, with packets identified by Tunnel-id. A separate tunnel-id range is defined in order for flow entries to distinguish MPLS-TP packets from data center overlay tunnel packets.

Both MPLS L2 Port and Tunnel-id must be assigned to MPLS-TP flows. Together these are used to represent packets and direction for a particular customer flow.

Figure 5 shows just the OF-DPA Abstract Switch objects that are used for VPWS termination packet flows. For simplicity, the two VLAN flow tables are shown “stacked”. For VPWS termination flows typically require no more than one VLAN tag to be matched.




Figure 5 Abstract Switch Objects for MPLS-TP Termination (VPWS)

The Termination MAC Flow Table identifies MPLS frames that require MPLS tunnel termination processing using flow entries that match destination MAC address and Ethertype and that have a Goto-Table instruction that specifies the MPLS Flow Tables.

For the same reason that two VLAN tables were used to enable matching two VLAN tags, multiple MPLS Flow Tables are used to enable matching up to three MPLS labels. Each table is used to match an outermost MPLS shim header. The first table, MPLS 0, can be used to match and pop an outermost LSP label¹¹. MPLS Flow Tables 1 and 2 can be used to match another LSP label or a pseudo-wire bottom of stack label. In the latter case OF-DPA provides additional new match fields and actions to identify and pop a control word if one is expected, and to remove the outermost L2 header¹². The pseudo-wire label also assigns the Tunnel-id and references a group entry for forwarding the packet.

MPLS Flow Table 0 only supports a subset of the actions supported by MPLS Flow Tables 1 and 2. All flow entries in MPLS Flow Table 0 should also be installed in MPLS Flow Tables 1 and 2. MPLS Flow Tables 1

¹¹ The dotted lines indicate that MPLS 0 is optional and may not be supported on all platforms.

¹² OpenFlow expects the MPLS pop action to specify the packet Ethertype, but in this case the result is Ethernet in Ethernet. For consistency it could be programmed to Transparent Ethernet Bridging (0x6558) [25].

and 2 are synchronized and contain the same rules, so updating MPLS Flow Table 1 also updates MPLS Flow Table 2.

Note: The OF-DPA 2.0 API does not permit MPLS Flow Table 2 to be modified directly.

Flexible policies can be applied to the resulting packet using the Policy ACL Flow Table.

If the frame has a service delimiting tag [30] an MPLS L2 Tag group entry can optionally be used to process the tag (the fact that this group entry is optional is indicated by its dotted outline). The resulting frame is forwarding to an egress interface using either an L2 Interface group entry or an L2 Unfiltered Interface group entry. As in the Bridging and Routing flows, these groups are used to represent the tagging properties of the egress port.

Note: In diagrams overlaid L2 Interface group entries are shown as in Figure 5. This should be interpreted as representing either an L2 Interface or an L2 Unfiltered group. L2 Unfiltered group entries are described in Section 4.3.2.

3.4.2 MPLS-TP L2 VPN Groups

Figure 6 shows the L2 VPN MPLS groups referenced from the packet flow diagrams. This is essentially the same as was included in Figure 4. The group entries in dotted lines are optional for pseudo-wire initiation.

OF-DPA 2.0 supports MPLS-TP 1:1 linear and ring protection. MPLS Fast Failover group entries for two level 1:1 protection are shown in Figure 6.

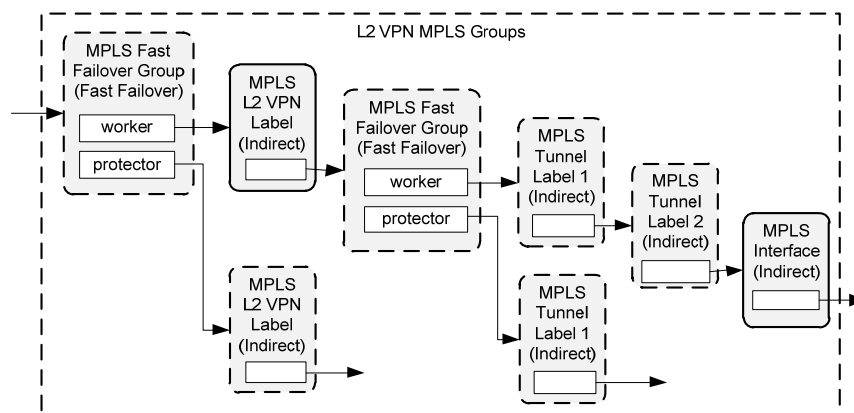


Figure 6 MPLS L2 VPN Groups

MPLS Fast Failover group entries have two buckets. The primary (worker) bucket is for the normal working path and the second (protector) bucket is for the protection path. An external network protection switching process decides which to use and when to switch over from one to the other. The operation of the network protection switching process and its relation to OAM objects is described in Section 3.9.

Note: The number of levels of protection available is platform dependent.

The MPLS L2 VPN Label group entry bucket actions can include:

- Push L2 Header (initial values are zero; to be populated by actions in the MPLS Interface group)
- Push CW (new action to push control word)
- Push MPLS shim header (for PW label; must set the Ethertype in the L2 header to 0x8847)
- Set-Field MPLS Label
- Set-Field BOS (bottom of stack)
- Set-Field TC (EXP) – explicit value or from table based on Traffic Class and Color. Optional.
- Set-Field TTL
- Set-Field PCP (in outermost L2 header) – explicit value or from table. Optional.
- Group (MPLS Tunnel Label)

The MPLS Tunnel Label 1 group entry bucket actions can include:

- Push MPLS shim header (for LSP label; must set the Ethertype to 0x8847)
- Set-Field MPLS Label
- Set-Field TC (EXP) – explicit value, copy out, or set from table. Optional
- Set-Field TTL (value or copy out)
- Set-Field PCP (in L2 header) – explicit value or from table (Optional)
- Group (MPLS Tunnel Label 2 or MPLS Interface)

The MPLS Tunnel Label 2 Group entry may be used to push a second outer tunnel label and has the same actions as MPLS Tunnel Label 1 except the chained group can only be MPLS Interface.

The MPLS Interface group entry sets the outer MAC-DST, MAC-SRC, and VLAN Id. The outer L2 header is assumed to have a VLAN field.

3.5 MPLS Label Edge Router (LER)

The OF-DPA objects for MPLS L3 VPN are shown in Figure 7 for initiation and Figure 8 for termination.

The MPLS Label Edge Router packet flow supports routing into and out of MPLS L3 VPN tunnels. An LER is both an IP router and an MPLS tunnel endpoint. The LER can support multiple VPNs for different customers.

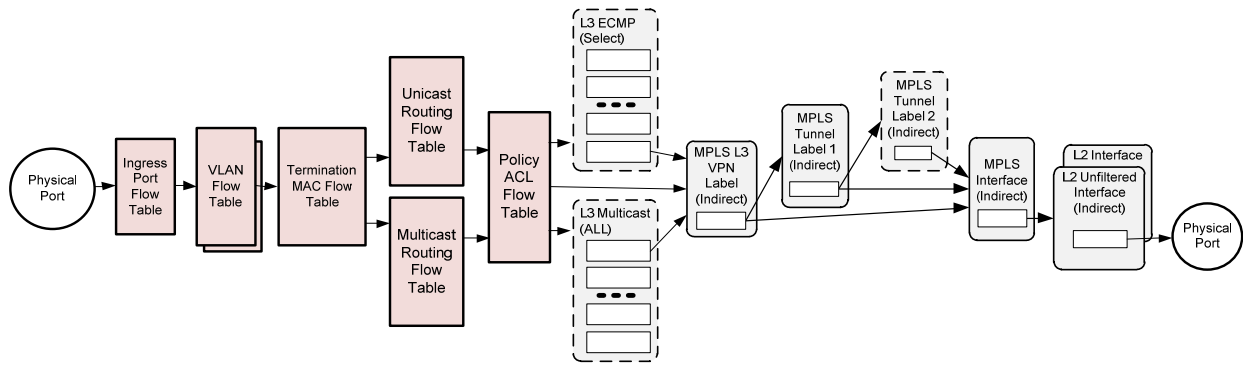


Figure 7 Abstract Switch Objects for MPLS L3 VPN Initiation

For tunnel initiation IP packets are routed to an MPLS L3 VPN. For isolation multiple virtual routing tables (VRF) are supported, selected using VRF pipeline metadata. OF-DPA defines a pipeline match field for VRF.

Customer ingress traffic is differentiated based on VLAN, and VLAN Flow Table rules can optionally assign a VRF for a customer's traffic. IP packets can be forwarded either directly or through an L3 Multicast or L3 ECMP group entry to MPLS Label Group Entries that push MPLS labels and update fields in the Ethernet header. In these groups the buckets must reference appropriate L3 group entry types.

The MPLS L3 VPN Label group entry bucket actions include:

- Push MPLS shim header (for VPN label; this sets the Ethertype to 0x8847)
- Set-Field MPLS Label
- Set-Field BOS (bottom of stack)
- Set-Field MPLS TC (EXP, optional)
- Decrement and Check TTL (depends on label)
- Set-Field TTL (value or copy out)
- Set-Field PCP (in L2 header, optional)
- Group (MPLS Tunnel Label)

The MPLS Tunnel Label group entry bucket actions are the same as for the L2 VPN case.

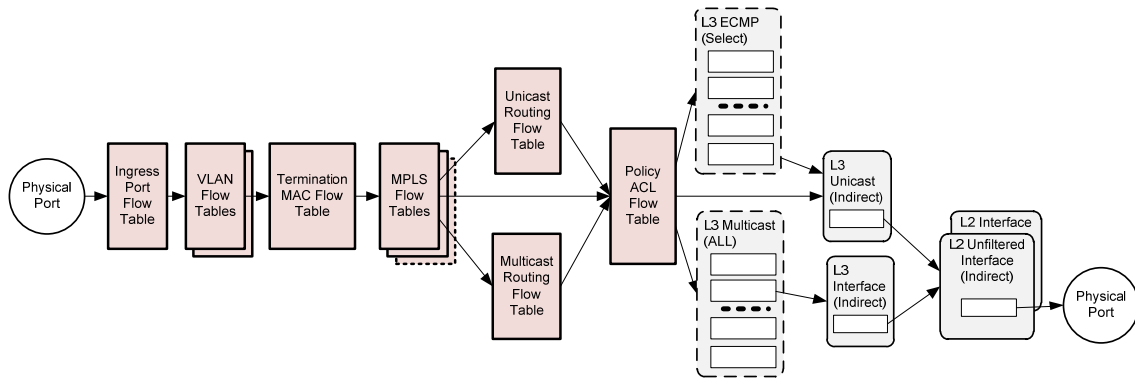


Figure 8 Abstract Switch Objects for MPLS L3 VPN Termination

For tunnel termination, MPLS frames are identified by destination MAC and VLAN in the Termination MAC Flow Table. MPLS shim headers are processed as indicated by the MPLS Flow Tables.

For MPLS L3 VPN termination, popping the shim headers turns these frames into IP packets that can be routed using the Routing Flow Tables or directly forwarded to L3 Unicast, L3 Multicast, or L3 ECMP Group Entries. Packets to be routed that have multicast group destinations are forwarded using the Multicast Routing Flow Table rather than the Unicast Routing Flow Table.

The MPLS bottom of stack label can assign VRF pipeline metadata for the VPN. VPN traffic is isolated by being forwarded to the VLAN associated with that VPN.

3.6 MPLS Label Switch Router (LSR)

The OF-DPA objects for an MPLS Label Switch Router (MPLS-TP P node) are shown in Figure 9. The flow uses many of the same objects as the LER case.

An LSR forwards MPLS frames by optionally popping one or more labels and then swapping a label. For OF-DPA these are programmed as actions in the MPLS Flow Tables. Swap can apply either to a tunnel (LSP) label that is not bottom of stack, and also to a PW label that is bottom of stack for PW stitching scenarios.




Figure 9 Abstract Switch Objects Used for an MPLS LSR

Figure 9 shows optional forwarding to an MPLS ECMP group to support multipath load balancing. This is only used for MPLS L3 VPN and not MPLS-TP. The path selection function is implemented in the hardware platform.

The LSR group entries are shown in Figure 10. OF-DPA 2.0 only supports Fast Failover groups for an MPLS-TP P-node and not for MPLS L3 VPN.

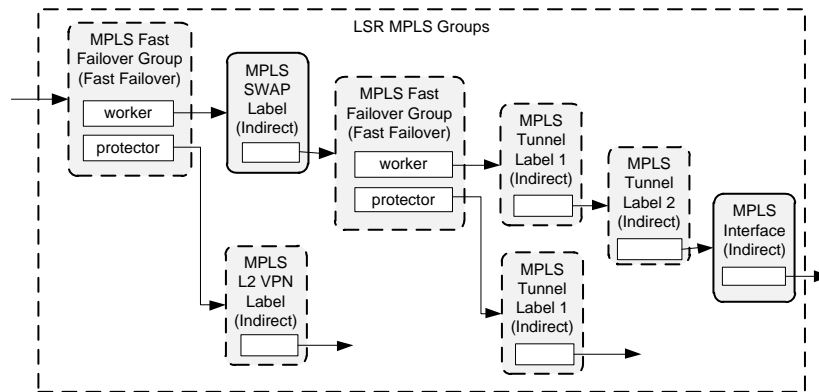


Figure 10 MPLS LSR Groups

The MPLS SWAP Label group entry bucket actions for P include:

- Push MPLS Label (if LSP label popped)
- Set-Field MPLS Label (set or swap label)
- Set-Field TC (EXP, optional)
- Decrement and check TTL
- Set-Field TTL (optional)
- Set-Field PCP (in outermost L2 header) – explicit value or from table. Optional.
- Group (MPLS Tunnel Label)

The MPLS Tunnel Label group entry bucket actions are the same as for the earlier cases.

The MPLS Interface group entry sets the outer MAC-DST, MAC-SRC, and VLAN Id for forwarding. The outer L2 header is assumed to have a VLAN field.

Figure 9 also shows optional protection switching using one (or more) MPLS Fast Failover Group entries.

3.7 MPLS-TP Protection Switching

OF-DPA 2.0 supports MPLS-TP 1:1 linear and ring protection. The MPLS Fast Failover group entry for 1:1 protection is shown in Figure 11.




Figure 11 MPLS 1:1 Protection

MPLS Fast Failover group entries have two buckets. The primary bucket is for the normal working path and the second bucket is for the protection path. An external network protection switching process decides which to use and when to switch over from one to the other.

MPLS Fast Failover buckets must be programmed with a watch_port that specifies “watching” an OAM Protection Liveness logical port. The Protection Switching process can change the state of an OAM Protection Liveness logical port in order to switchover to the protection path. The OAM Protection Liveness logical port must be included in the MPLS Linear Protection configuration as shown in Figure 48. The operation of the network protection switching process and its relation to OAM objects is described in Section 3.9.

3.8 Quality of Service (QoS)

This section describes the OF-DPA Abstract Switch objects supporting QoS for Bridging and Routing and for MPLS. Note that OpenFlow 1.3.4 does not fully support the metering and marking capabilities required. As a result, some concepts need to be introduced before describing the abstract switch packet flows.

3.8.1 QoS Concepts – Traffic Class and Color

QoS packet processing following the DiffServ model [22] comprises the following:

- Classification - assigning a Traffic Class and Color, or drop precedence. Trust policies are associated with ingress ports and can determine whether or not to assign Traffic Class and Color based on incoming packet header fields;
- Metering – policing, which can change the Color based on flow properties such as packet and byte rate;
- Marking - setting QoS fields in the packet headers based on Traffic Class and Color; and
- Shaping – queuing or dropping the packet based on Traffic Class and Color, respectively. Queues are serviced based on the scheduling discipline. Typically queues implement a color-based admission scheme such that they stop accepting packets with higher drop precedence after the number of entries in the queue exceeds some threshold¹³.

OF-DPA 2.0 defines new pipeline match fields for Traffic Class and Color. In general higher Traffic Class values correspond to higher priorities; Color can be one of Green (0), Yellow (1), or Red (2).

¹³ Shaping properties and configuration are platform dependent.

The Policy ACL Flow Table is used for QoS classification and metering. It provides an optional Meter instruction to reference the Meter Table entry to use to police a flow. It also can add a Set-Queue action to the action set. The Policy ACL Color Based Actions Flow Table can be used for marking actions.

3.8.2 Meter Table Operation

OF-DPA uses a type of Meter Table entry with two meter bands. Each meter band can indicate actions to apply immediately to the packet. In OpenFlow 1.3.4, the meter instruction must be evaluated before other instructions.¹⁴

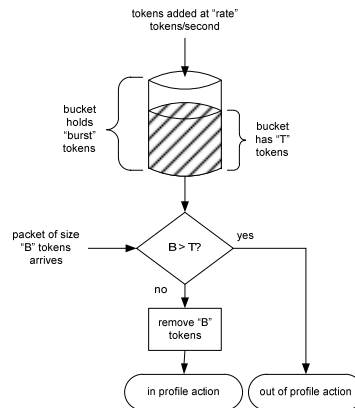


Figure 12 Token Bucket Operation

As opposed to OpenFlow 1.3.4 Meters, OF-DPA Meters are implemented using the standard token bucket algorithm and hence the configuration implies not only the rate and burst limits but also how they are measured. For the most part, rate and burst are configured in terms of standard OpenFlow Meter Band parameters. Figure 12 illustrates the operation of the token bucket algorithm. OpenFlow Meter Band rate and burst parameters are used to configure the rate at which tokens are added and the bucket size in tokens, as well as whether the tokens represent kilobits or packets.¹⁵ However, OF-DPA meters also support two modes (two rate three color or single rate three color) and two forms of coloring (whether color aware or color blind).

Operation of the two rate three color mode (TrTCM), illustrated in Figure 13, is as specified in RFC 2698 [24].

¹⁴ In the future OpenFlow may deprecate the meter instruction in favor of a meter action. This would not materially change the operation of OF-DPA 2.0 meters but would change how meters were programmed from the controller.

¹⁵ OpenFlow meter bands specify packet rates and burst sizes independent of how measured.

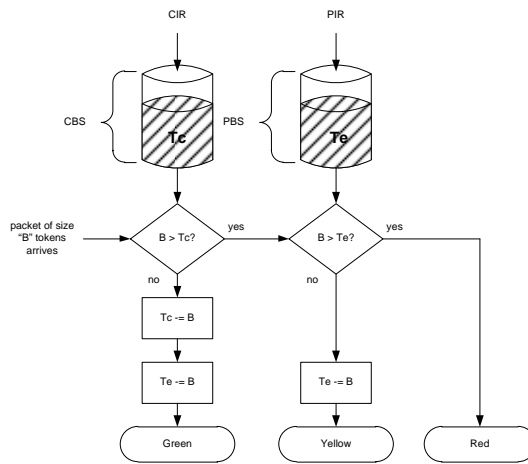


Figure 13 TrTCM Meter Operation

TrTCM meters can be color-aware or color-blind. TrTCM color-aware operation in terms of actions and incoming packet color is as shown in Table 2. In color-blind mode the incoming packet is always Green.

Table 2 TrTCM Color-Aware Operation

Incoming Color	$B \leq T_c?$	$B \leq T_e?$	Result Color	T_c	T_e
Green	Yes	Yes	Green	$T_c -= B$	$T_e -= B$
	Yes	No	Red	-	-
	No	Yes	Yellow	-	$T_e -= B$
	No	No	Red	-	-
Yellow	Yes	Yes	Yellow	-	$T_e -= B$
	Yes	No	Red	-	-
	No	Yes	Yellow	-	$T_e -= B$
	No	No	Red	-	-
Red	Yes	Yes	Red	-	-
	Yes	No	Red	-	-

Incoming Color	$B \leq Tc?$	$B \leq Te?$	Result Color	Tc	Te
	No	Yes	Red	-	-
	No	No	Red	-	-

The single rate three color mode (SrTCM), illustrated in Figure 14, operates as specified in RFC 2697 [23].

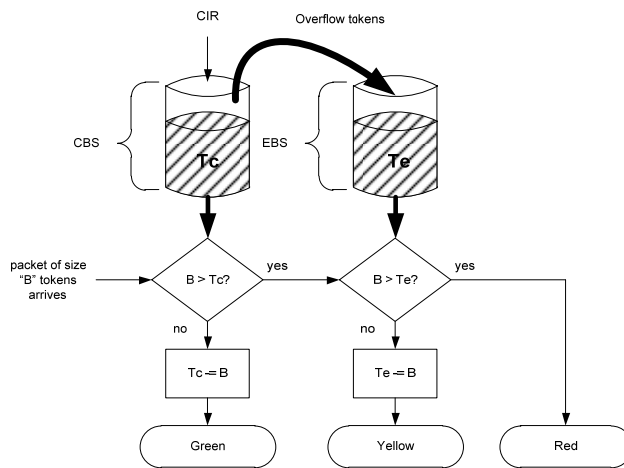


Figure 14 SrTCM Meter Operation

As with TrTCM, SrTCM meters can be color-aware or color-blind. SrTCM color-aware operation is as shown in Table 3. In color-blind mode the incoming packet is always Green.

Table 3 SrTCM Color-Aware Operation

Incoming Color	$B \leq Tc?$	$B \leq Te?$	Result Color	Tc	Te
Green	Yes	Yes	Green	$Tc -= B$	
	Yes	No	Green	$Tc -= B$	-
	No	Yes	Yellow	-	$Te -= B$
	No	No	Red	-	-
Yellow	Yes	Yes	Yellow	-	$Te -= B$

Incoming Color	$B \leq T_c?$	$B \leq T_e?$	Result Color	T_c	T_e
	Yes	No	Red	-	-
	No	Yes	Yellow	-	$T_e - B$
	No	No	Red	-	-
Red	Yes	Yes	Red	-	-
	Yes	No	Red	-	-
	No	Yes	Red	-	-
	No	No	Red	-	-

In OF-DPA 2.0 the packet Traffic Class value is used to directly assign the queue. Queues can be configured in terms of minimum and maximum rates using the OF-DPA queue configuration APIs. As mentioned, the packet Color determines the drop precedence. Queues reserve a guaranteed size for green packets. Yellow packets are dropped if the queue exceeds this size, otherwise they are queued. Red packets are always dropped.

Note: The Traffic Class values are the same as the Queue values for this version of OF-DPA.

Note: Queues are associated with output ports and serviced by default in strict priority order in this version of OF-DPA.

Meter Table entries are configured using OF-DPA meter table APIs. An OpenFlow controller would program OF-DPA Meter Table entries using the experimenter message fields described in Section 6.

3.8.3 Bridging and Routing QoS

The OF-DPA 2.0 Bridging and Routing pipeline including additional objects for QoS is shown in Figure 15. Basically QoS adds Meter table instructions to the Policy ACL Flow Table and re-mark actions for IP and Ethernet using the Color Based Actions Flow Table.

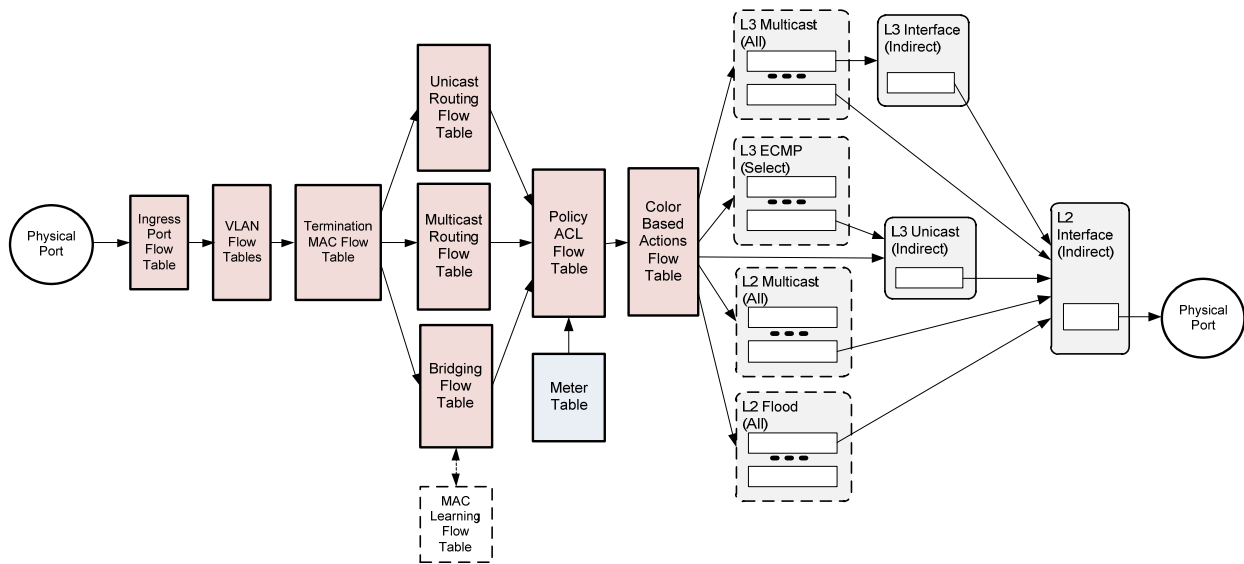


Figure 15 Bridging and Routing Abstract Switch With QoS Objects

Metering is used to police traffic rate according to bandwidth profiles. OF-DPA policing effectively implements three profiles: red, yellow, and green. The meter can change the value of the Color pipeline match field based on the profile.

3.8.4 MPLS QoS

The OF-DPA objects for MPLS-TP VPWS tunnel initiation with QoS are shown in Figure 16. In addition to policing, this adds an MPLS Label Re-Mark Table that can be accessed from the MPLS Label groups.

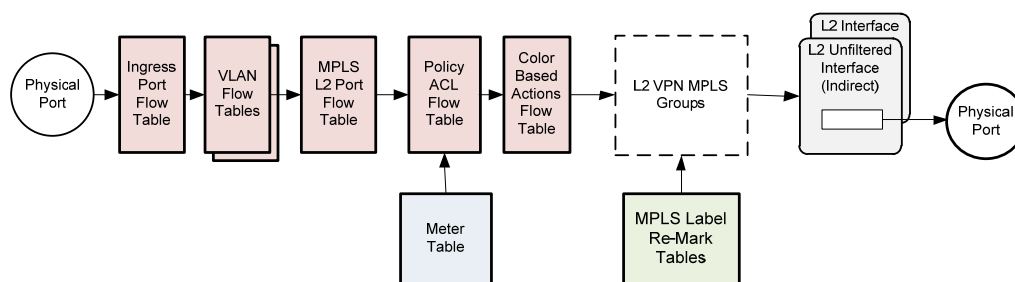


Figure 16 MPLS-TP VPWS Initiation with QoS Objects

For tunnel origination, OF-DPA 2.0 supports different options for setting the MPLS EXP and outermost VLAN PCP and DEI fields to accommodate different use cases, as follows:

- Fixed values programmed into the MPLS label group entries;
- Copy out from inner label, if one exists;
- Map using a MPLS Label Re-Mark Table with an index supplied in the label group;
- Use the existing value unchanged.

There are two types of MPLS Re-Mark tables. The MPLS VPN Label Remark Tables are referenced by actions from OF-DPA MPLS VPN Label Group entries when originating a tunnel, and from OF-DPA MPLS Swap Label Group entries when swapping a label. The MPLS Tunnel Label Remark Tables are referenced by actions from OF-DPA MPLS Tunnel Label 1 or OF-DPA MPLS Tunnel Label 2 Group Entries. Either type can specify an MPLS_TC value as well as new 802.1p PCP and DEI values.

OF-DPA 2.0 defines new ancillary table objects for MPLS Re-Mark tables. OpenFlow 1.3.4 does not provide a way to do such a table lookup in the context of a group action, or even to do a match action table lookup after evaluating a group entry.¹⁶

The OF-DPA objects for MPLS-TP VPWS tunnel termination with QoS are shown in Figure 17. This adds the Meter Table and Color Based Actions Flow Table described above.

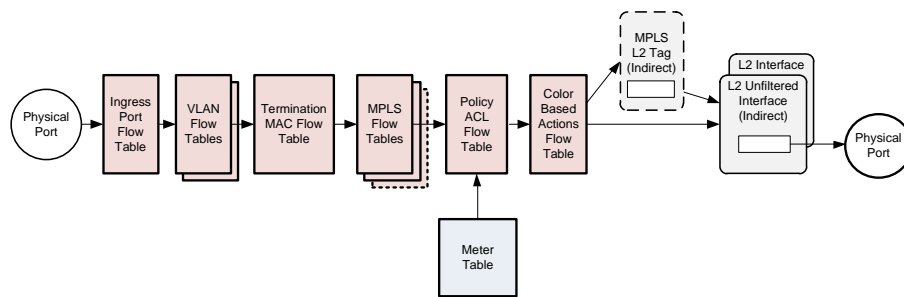


Figure 17 MPLS-TP VPWS Tunnel Termination with QoS Objects

The OF-DPA MPLS L3 VPN Initiation abstract objects with QoS are as shown in Figure 18. The corresponding MPLS L3 VPN Termination Abstract Switch objects including QoS objects for label re-mark is as shown in Figure 19.

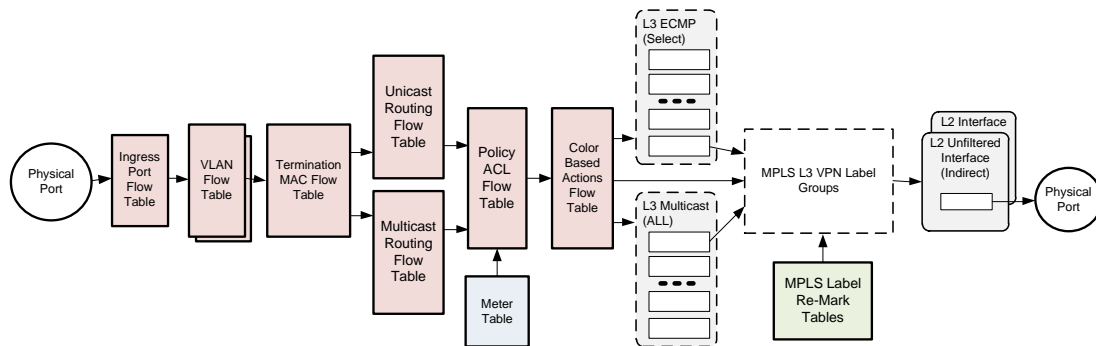


Figure 18 MPLS L3 VPN Initiation with QoS Objects

¹⁶ Although OF-DPA does make use of Egress Tables, they are only applied at the output port immediately before packet egress and after all groups.

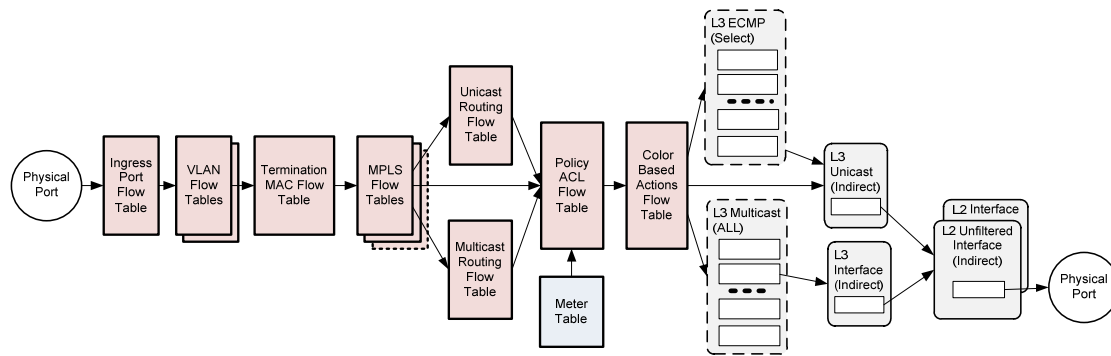


Figure 19 MPLS L3 VPN Termination with QoS Objects

For completeness, the MPLS LSR abstract objects including QoS objects is shown in Figure 20. The MPLS Label Re-Mark tables are as described above.

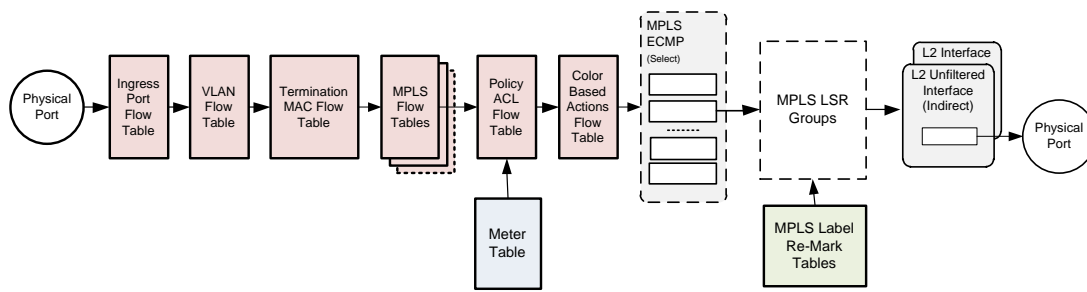


Figure 20 MPLS LSR with QoS Objects

Note that flow policing can change packet color and potentially affect queue drop precedence as well as re-marking.

3.9 Operation, Administration, and Maintenance (OAM)

OAM functions are used to detect and localize loss of connectivity or degradation of performance in order to maintain service level guarantees. In particular, identifying the defect location in the network enables rapid repair based on pre-provisioned redundant (protection) paths. This section provides some background context on OAM but is not intended to be a comprehensive tutorial.

Different networking scenarios use different types of OAM, and sometimes more than one can be used at the same time. OF-DPA includes support for ITU G.8113.1 and for Ethernet OAM over MPLS-TP. This section describes the OF-DPA objects used for OAM processing.

3.9.1 OAM Concepts – Maintenance Points

OAM functions are located at Maintenance Endpoints (MEP) and Maintenance Intermediate Points (MIP) belonging to a Maintenance Association (MA) for a fully connected network. MEPs are often, but not always, located at nodes that provide external access to the network. MIPs are located at strategic points

within the network. The MEPs transmit and receive OAM messages in order to detect defects on the protected path(s) between them.

OAM messages are used for fault management and performance monitoring. Depending on the type and purpose, messages might be sent periodically (proactive OAM) or on-demand. MIPs located along the path between MEPs respond to messages received from MEPs belonging to the same MA.

On a network element, MEPs and MIPs can be provisioned so that they face outward, towards the network, or inwards, towards the node. Network facing MEPs are referred to as Down MEPs and are used to test the path up to and out of an interface on the node. Inward facing MEPs are referred to as Up MEPs and are used to test the forwarding path through the node. An MA can consist of either Down MEPs or Up MEPs, but not both. Note that a MIP can have Up and Down orientations as well¹⁷.

Maintenance Domain (MD) is an Ethernet OAM concept. An MD is a collection of nested MAs identified by Maintenance Domain Level (MDL). Nested MDs are used to verify connectivity of a path subset in order to localize faults. All Ethernet OAM messages include an MDL field.

These concepts are illustrated in Figure 21, taken from the IEEE 802.1ag specification [27].

IS THIS COPYRIGHTED?

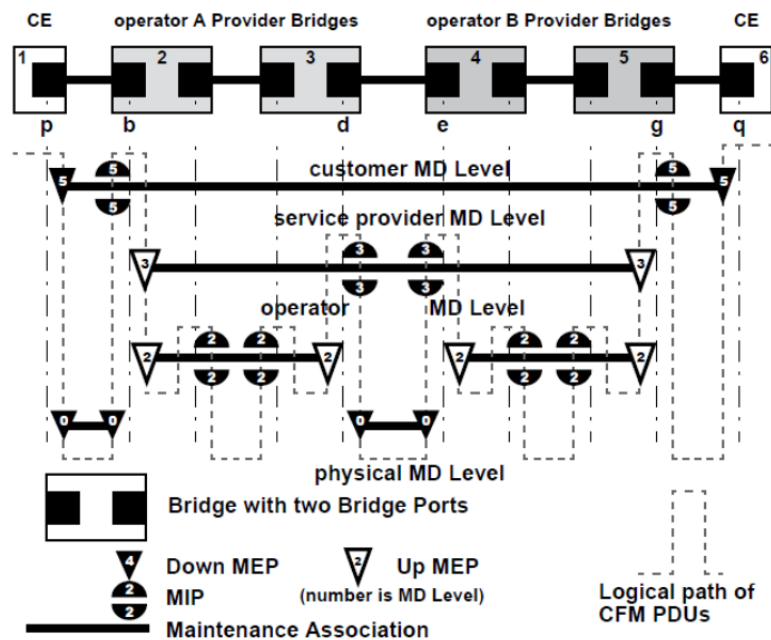


Figure 21 OAM MEP and MIP Examples

Note: The term Maintenance Association is defined in IEEE 802.1ag for Ethernet OAM. The equivalent term defined in ITU G.8113/Y.1731 is Maintenance Entity Group (MEG).

¹⁷Actually, MIPs are both, with an Up and Down MHFs (MIP Half Functions), represented with semi-circles.

OF-DPA OAM provides operations for provisioning and supporting MEPs and MIPs. These include using OpenFlow match action tables to identify frames that require processing at a MEP or MIP, and then performing OAM actions on these frames. On control frames (OAM messages), the actions can be to do OAM message processing. On data frames, the actions can include performance monitoring instrumentation, for example, updating transmit and receive counters for loss measurement (LM). The particular Maintenance Point in the MA (alternately, Maintenance Entity in the Maintenance Entity Group) is identified using a unique LMEP (local MEP) id¹⁸. OF-DPA defines an LMEP id pipeline metadata match field for this purpose.

3.9.2 Network Protection Apps

OF-DPA provides actions to send OAM frames to the controller for processing by a Network Protection App. However, since this approach cannot meet all OAM proactive fault management requirements, OF-DPA also supports offloading processing to one or more Local Engines hosting local Network Protection Apps. There are multiple implementation options for a Local Engine, such as: software on the local CPU; external hardware or software device such as an FPGA or sidecar processor; internal processing engine dedicated to OAM¹⁹. In all cases, the local Network Protection App can perform time-sensitive OAM functions such as transmitting and receiving continuity check (CCM) frames.

OF-DPA accommodates both types of Network Protection Apps in a uniform way using OpenFlow Reserved Ports. As shown in Figure 22, controller Network Protection Apps receive and transmit OAM frames using Packet_In and Packet_Out messages.

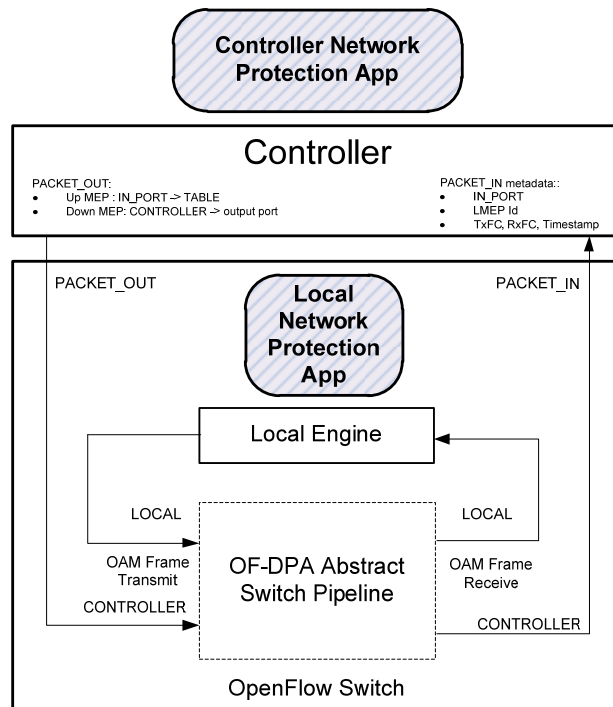


Figure 22 OAM and Protection Overview

¹⁸ LMEP is used generically and applies to both MEPs and MIPs

¹⁹ Available on certain Broadcom platforms

OF-DPA Abstract Switch objects send message frames to Network Protection Apps on the controller using an output action to the CONTROLLER Reserved Port. These frames are accompanied by OAM pipeline metadata fields. The agent forwards these frames to the controller in Packet In messages. Conversely, the controller Network Protection App can inject OAM frames by sending them to the switch in Packet_Out messages. Packet_Out messages associate source and destination ports with the packet to indicate where it should be injected relative to the OF-DPA Abstract Switch pipeline. For an Up MEP, the source is a physical or logical port to be used as the value of IN_PORT, and the destination is the TABLE Reserved Port, which indicates processing by the first table in the pipeline. For a Down MEP, the source is CONTROLLER and the destination is an output port. In this case the Network Protection App must output a packet that is already encapsulated as necessary and ready to be sent on the wire.

Similarly, OF-DPA can use the LOCAL Reserved Port to interface with a local Network Protection App. OF-DPA defines a LOCAL reserved port such that the local Network Protection App can inject OAM frames through this it by specifying source and destination ports as in the Packet Out case. For an Up MEP, the source is a physical or logical port to be used as the value of IN_PORT, and the destination is the TABLE Reserved Port, which indicates processing by the first table in the pipeline. For Down MEPs, the source is LOCAL and the destination is an output port.

This provides a uniform modeling approach for OAM message processing independent of whether done at the Controller or using a local engine. OF-DPA flow entries execute an output action to a Reserved Port and provide metadata fields to a Network Protection App. These pipeline match fields include the packet IN_PORT, LMEP Id, transmit and receive counters, and packet arrival timestamp. For frames to the Controller OAM App, this metadata is carried in the Packet_In message. For local OAM Apps, it is supplied using a local mechanism.

Note: The exact mechanism for associating metadata with the LOCAL Reserved Port is implementation dependent and outside of the scope of this specification.

Note: OF-DPA sends the complete frame as received, including all headers, to the Network Protection App²⁰. In addition the packet is accompanied by available pipeline metadata, including the LMEP id.

Note: The level of accuracy in loss measurements is affected by the latency between reading counters and inserting counts into injected OAM frames. In general, a Local Network Protection App should be used to provide better accuracy. It is further recommended that a local interface be provided for accessing counter values, although such a facility is outside of the scope of this specification.

3.9.3 MPLS-TP OAM

MPLS-TP support uses MPLS Fast Failover Groups to implement protection for network paths. Fast Failover Groups rely on OAM path fault detection for liveness. When OAM detects a fault on the worker path, the OAM connectivity verification process state machine will notify the network protection state machine, which in turn will invalidate the liveness property being monitored by the worker bucket. This then causes the Fast Failover Group to switch over to the protector path. The MPLS Fast Failover Group

²⁰ This varies from the OpenFlow specification, which outputs the packet in its currently edited condition, i.e., following application of actions such as pop, push, set-field, etc.

bucket “watch_port” parameter is configured to watch the operational status of an OAM Network Protection Logical Port defined for this purpose.²¹

The requirements for MPLS-TP OAM are described in RFC 5860 [26]. Examples are illustrated in Figure 23.

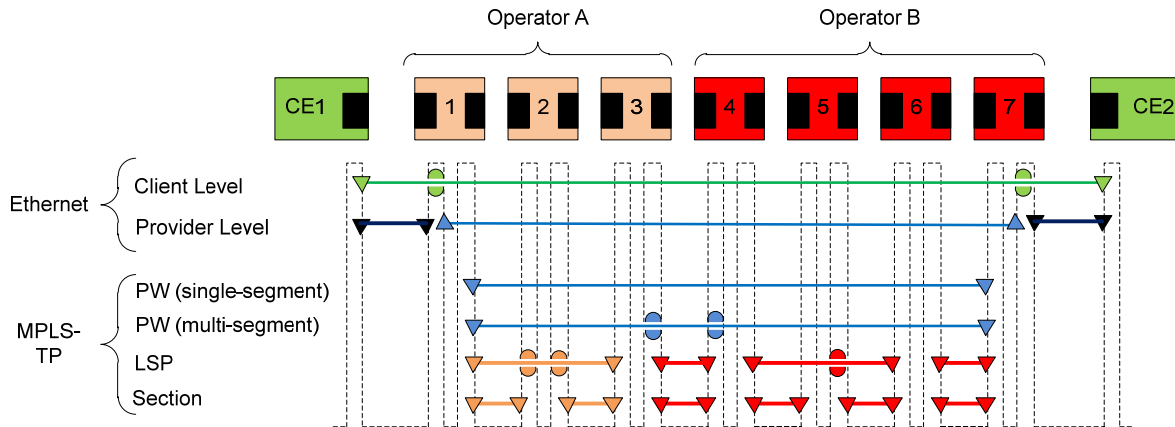


Figure 23 MPLS-TP Service OAM Examples

As shown in Figure 23, MPLS-TP requires support for the following different types of OAM:

- Client Service OAM is end-end between client points of service (UNI-C). Since MPLS-TP provides an Ethernet service, client service OAM exchanges Ethernet OAM frames through the MPLS tunnel between Up MEPs on customer equipment (CE).
- Provider Service OAM is end-end between service provider points of service (UNI-N). This also exchanged Ethernet OAM frames through the MPLS tunnel between Up MEPs in LERs.
- MPLS-TP Channel OAM is at the pseudo-wire (PW) level and may be end-end (single-segment, SS-PW), or stitched across multiple operator domains (multi-segment, MS-PW). Two operator domains are shown in Figure 23. Multi-domain end-end at the PW level is also shown. Channel OAM is only between Down MEPs. MS-PW may also involve one or more MIPs on the path.
- MPLS-TP Path OAM is at the LSP level and is used to protect the path between two network elements. Path OAM MEs are always within a single operator domain and only between Down MEPs but may also involve MIPs on the path where warranted. While one LSP is shown above, there may be segments with nested LSPs.
- MPLS-TP Section OAM is on the single hop between two network elements. Section OAM is between Down MEPs on adjacent nodes.

Each requires identifying one or more packet formats. Packet formats are described in the following sections.

²¹ In OpenFlow 1.3.4, Fast Failover Group buckets are configured to “watch” either port or group liveness, with group liveness determined (at run time) by port liveness. OF-DPA MPLS Fast Failover Group buckets “watch” a logical port, relying on the Network Protection App to control bucket liveness by changing the state of the watched logical port.

The following sections provide details of OF-DPA MPLS-TP OAM processing for Ethernet and G.8113.1. Configuration for message processing and liveness determination can be found in Section 5.3 on page 145.

3.9.3.1 Ethernet OAM over MPLS-TP

Ethernet OAM methods are described in ITU-T Y.1731 [17] and IEEE 802.1ag [27]. Ethernet OAM frames are encapsulated within MPLS. These frames are identified by having ETH-TYPE=0x8902.

The packet format is shown in Figure 24.




Figure 24 Ethernet OAM Over MPLS-TP Packet Format

For tunnel initiation at the Maintenance Point, the OF-DPA data path requirement is to count transmitted user data frames if LM is enabled. User frames are recognized in the VLAN Flow Table as for the general initiation case. A new OAM Data Plane Counter Table is defined along with a new action to increment based on LMEP Id and Traffic Class for the affected flow. In addition, an action to query the Drop Status table is used to drop frames during a Lock condition. The flow tables involved are shown in the pipeline fragment in Figure 25.

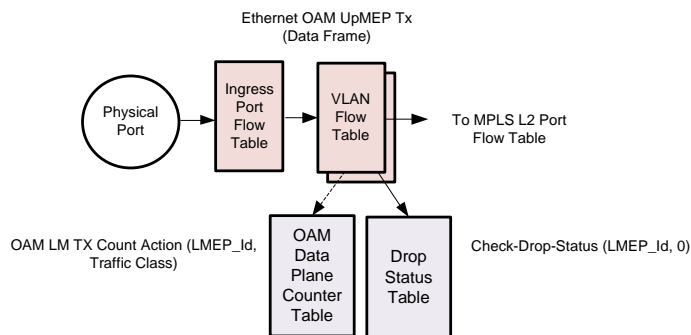


Figure 25 MPLS-TP Initiation - Ethernet over MPLS-TP OAM Data Frame

The connection to the OAM Data Plane Counter table is shown as a dotted line to highlight the fact that the count action is performed using the action set. Counters are not actually incremented until the action set is applied after the last match action table in the pipeline. This accommodates the fact that the packet may be dropped before that point due to a policing action.

Figure 26 shows the corresponding pipeline fragment for user data frame tunnel termination at a Maintenance Point. Note that this requires the use of egress tables. Egress tables permit match action processing in the context of an egress port on the payload frame after encapsulation headers are removed.

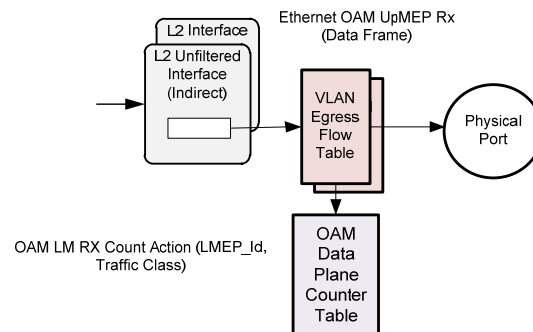


Figure 26 MPLS-TP Termination - Ethernet over MPLS-TP OAM Data Frame

For a client level tunnel initiation MIP, OAM messages are classified to a flow using the same match criteria as customer data frames. However these are identified by ETH-TYPE=0x8902 and MAC-DST, rendering them candidates for OAM processing. Figure 27 shows the pipeline fragment for this. The VLAN Flow Tables identify the MP and set the LMEP Id. The Maintenance Point Flow Table matches on the opcode, Maintenance Level from the frame, and LMEP-id to decide whether and where to output the frame for OAM message processing; to count and forward higher MDL frames as data; to both count and forward in the case of a MIP; to drop for lower MDL frames, or to drop because of an administrative lock condition.

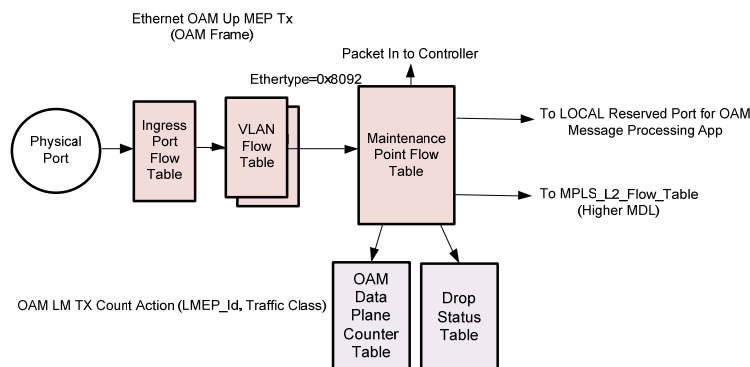


Figure 27 MPLS-TP Initiation - Ethernet over MPLS-TP OAM PDU

Note: An arrival timestamp value is associated with OAM PDUs. This timestamp is a pipeline metadata value. Its maintenance is a local matter.

A new Set-Counter-Fields action is defined for setting the counter pipeline metadata fields by reading current values in the OAM Data Plane Counter Table for a specified LMEP and Traffic Class, and for setting the timestamp metadata from the arrival timestamp value.

Note: This version of OF-DPA does not support Ethernet Down MEPs at the local attachment circuit (UNI).

Figure 28 shows the corresponding pipeline fragment for termination. This same flow is used for both a provider level Up MEP or a client level MIP at the local attachment circuit. The processing is similar to the initiation case except using the corresponding egress tables.

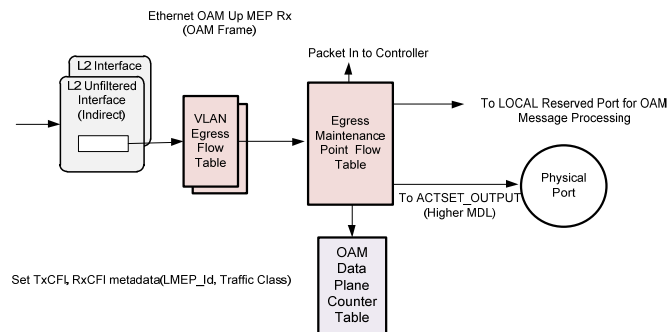


Figure 28 MPLS-TP Termination - Ethernet over MPLS-TP OAM PDU

While egress tables are not allowed to change the output port in the action set, note that they can still include output actions in an action list invoked with the Apply-Actions instruction. The Egress Maintenance Point Flow Table decides where the OAM frame will be processed by forwarding it either to CONTROLLER or LOCAL. It can output the frame or drop it by clearing the action set and not providing a Goto-Table instruction.

3.9.3.2 G.8113.1 OAM for MPLS-TP

MPLS-TP OAM methods are described in ITU-T G.8113.1/Y.1372.1 [15]. This standard describes a method for leveraging Ethernet OAM for MPLS-TP. In particular, it describes methods and procedures for applying ITU-T Y.1731 [16] Protocol Data Units (PDUs) for MPLS-TP OAM.

Packet formats for G.8113.1 are shown in Figure 29.

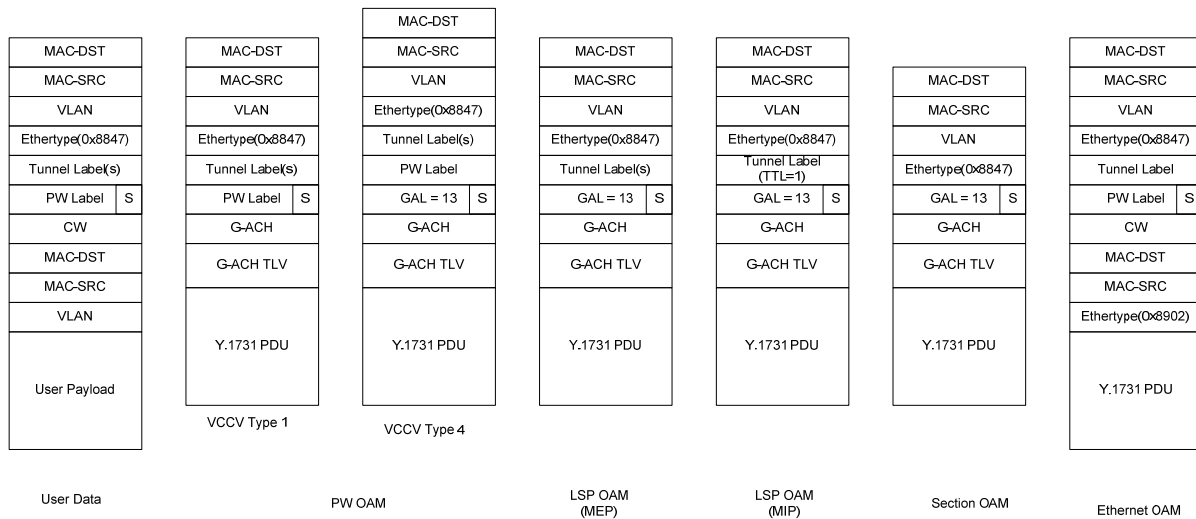


Figure 29 OAM MPLS-TP G.8113.1 Packet Formats

Note that G.8113.1 OAM only requires Down MEPs.

The OF-DPA objects for OAM processing for data frames at PW initiation are shown in Figure 30. The differences from the pipeline without OAM mainly consist of adding the Check Drop Status table action, and adding the LM counter action support to the MPLS L2 VPN and MPLS Tunnel Label group entries.

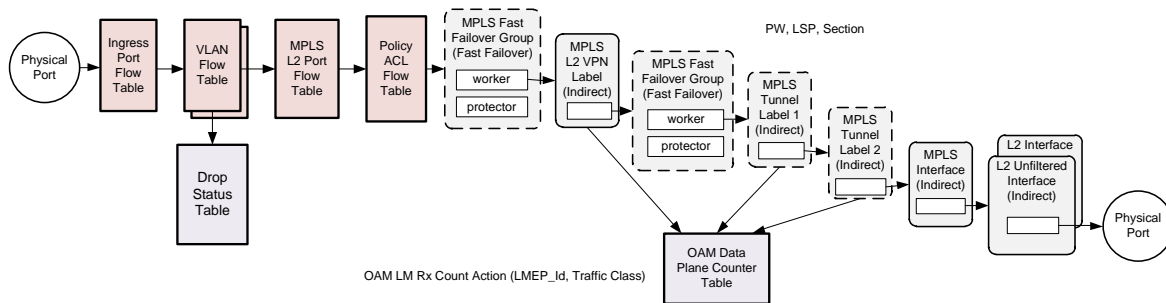


Figure 30 MPLS-TP Initiation - G.8113.1 OAM Data Frame

The corresponding OF-DPA objects for OAM data frame processing at pseudo-wire termination are shown in Figure 31. This mainly adds identifying the flow at classification time as a Down MEP, assigning LMEP id pipeline metadata, and adding the LM count action to the action set as required.

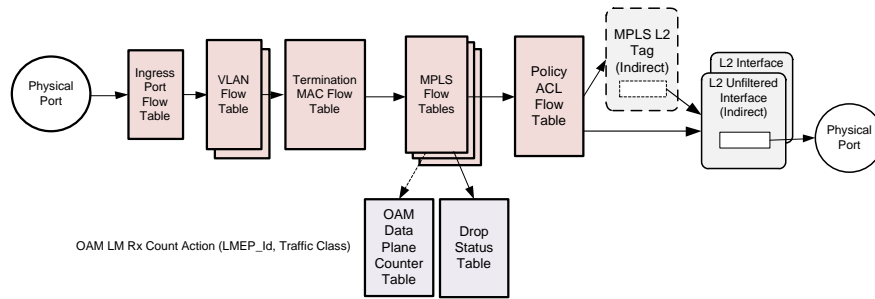


Figure 31 MPLS-TP Termination - G.8113.1 OAM Data Frame

Figure 32 shows the OF-DPA objects for LSR Down MEP or MIP processing for data frames. The receive counters are associated with MPLS Flow Table label matching as in Figure 31, and the transmit counters are associated with the MPLS Tunnel Label groups as in Figure 30.

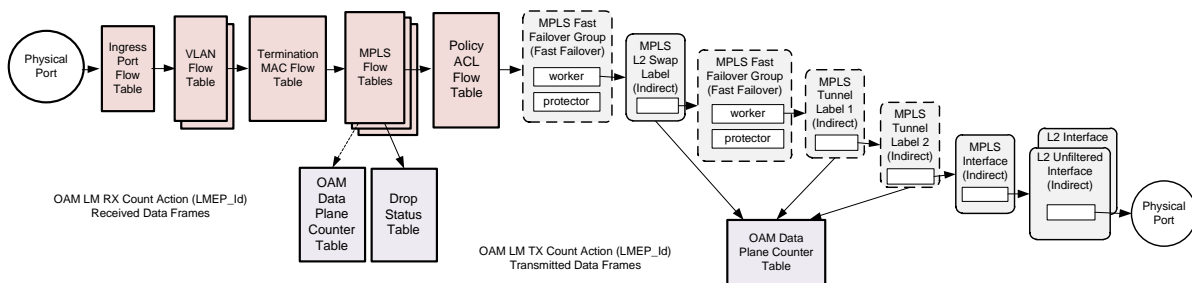


Figure 32 MPLS-TP - G.8113.1 OAM LSR Data Frame

The pipeline fragment in Figure 33 shows the OF-DPA objects for processing OAM data frames for PW termination, LSP, and Section OAM. For PW termination, OAM PDUs parsed according to the VCCV Type 1 packet format shown in Figure 29 are matched by PW label and G-ACH channel type. Similarly, VCCV Type 4 formats are recognized by the underlying bottom of stack GAL and similarly matched.




Figure 33 MPLS-TP - G.8113.1 OAM PDU

The OAM message packet formats for LSP and Section OAM are processed in almost the same way, the difference being in the match fields used in MPLS Flow Table rules. For example, LSP MIP detection requires matching on the MPLS-TTL.

Once matched as a frame requiring OAM processing by either the controller or local OAM Engine, the decision of whether to forward or drop the frame and to which destination is done in the Maintenance Point flow table on the basis of LMEP id and opcode, similar to the way it is done for Ethernet OAM.

3.10 Protection Switching

OF-DPA supports 1:1 linear and ring protection using Fast Failover group entries.

The OpenFlow 1.3.4 specification requires Fast Failover group types to support liveness monitoring to determine which bucket to use for forwarding. Fast Failover groups can be configured with `watch_port` and `watch_group` parameters, only one of which is used to determine bucket liveness. Liveness monitoring works as follows:

- A port is considered live if it has the `OFPPS_LIVE` flag set in its port state. Port liveness may be managed by code outside of the OpenFlow portion of a switch or defined outside of the OpenFlow specification, such as Spanning Tree or a KeepAlive mechanism. The port must not be considered live (and the `OFPPS_LIVE` flag must be unset) if one of the port liveness mechanisms enabled on the switch consider the port not live, or if the port config bit `OFPPC_PORT_DOWN` indicates the port is down, or if the port state bit `OFPPS_LINK_DOWN` indicates the link is down.
- A bucket is considered live if either `watch_port` is not `OFPP_ANY` and the port watched is live, or if `watch_group` is not `OFPG_ANY` and the group watched is live.
- A group is considered live if a least one of its buckets is live.

OF-DPA uses OAM Protection Liveness Logical Ports solely for the purpose of controlling the liveness property for OF-DPA MPLS Fast Failover group entry buckets. The controller or local OAM Engine can cause one or more OF-DPA MPLS Protection Fast Failover groups to switch buckets by changing the administrative state of an OAM Protection Liveness Logical Port to down.

The protection switching process and its linkage to the pipeline for switchover is illustrated in Figure 34.

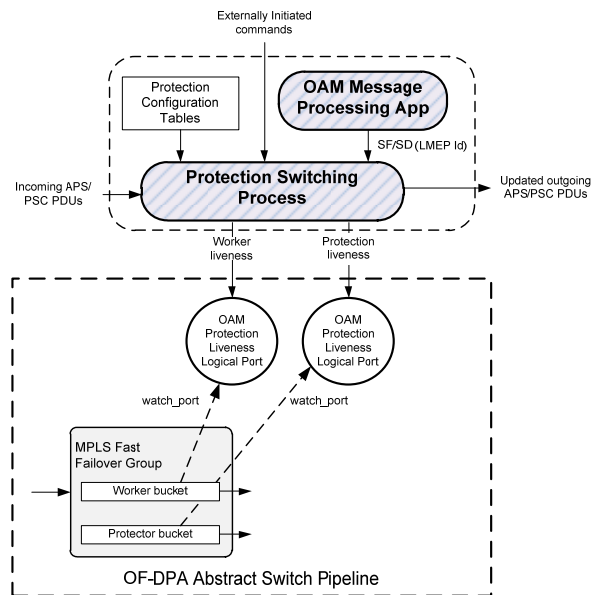


Figure 34 Protection Switching Process

OAM Protection Livness Logical Ports have the assigned number range shown in Table 147. OAM Protection Livness Logical Ports come pre-defined and require no configuration before being referenced. By default they are administratively up.

OF-DPA MPLS Fast Failover group entry buckets should be configured with watch_group OFPG_ANY and watch_port an OAM Protection Livness Logical Port.

The Protection Process can control the switchover from worker to protection path by changing the administrative state of an OAM Protection Livness logical port. The Protection Process is configured with the OAM Protection Livness logical port for use with a particular LMEP Id.

Note: Both the controller and a local OAM Engine can change the administrative status of an OAM Protection Livness Logical Port. The local mechanism is implementation dependent.

Note: An OAM Protection Livness Logical Port cannot be used in an output action.

4 OF-DPA OBJECT DESCRIPTIONS

OF-DPA 2.0 presents the application writer with a set of objects that can be programmed using OpenFlow 1.2.4. The programmable objects include flow tables, group table entries, and meter table entries. It also presents a set of objects that can be configured using an SDN configuration protocol such as OF-Config or OVSDB. The configurable objects include ports, queues, and OAM functions.

This section provides programming descriptions for these objects. For details consult the OF-DPA 2.0 TTP.

Flow tables have specific attributes, including entry types (rules) that have specific match fields, actions, and instructions. Flow entries can have "Goto-Table" instructions that determine the next table to process the packet. In other words, the flow entry programming determines the order in which packets traverse tables and accumulate actions in an action set. Actions in the action set are applied prior to the packet being forwarded when there is no next table specified. Specific forwarding actions, including egress packet edits, are for the most part included within the action sets of the group entries. OF-DPA 2.0 uses specific types of group entries to support different packet flow scenarios. Apply-actions instructions and action lists are also used for some VLAN tag and MPLS label packet editing, and to send packets to the controller.

In the general OpenFlow case packets pass from flow table to flow table and can be arbitrarily modified between tables. To take advantage of this generality each table stage would need to include a packet parser. In OF-DPA this kind of packet flow is conceptual - packets are parsed early in the pipeline and header fields are extracted. After that it is only these fields that are passed between tables and used for matching or modification by "apply actions" instructions. It is not expected that this distinction will matter to applications.

Section 3 showed the tables, group table entries, and ancillary objects used in the OF-DPA abstract pipeline for different packet processing flows.

Section 4.1 describes the OF-DPA flow tables in terms of their supported match fields, flow entry rule types, instructions, actions, expiration provisions, and statistics counters. Default miss actions are also specified for each table as applicable. Section 4.2 describes the OF-DPA group table entry types and action set constraints. Section 4.4 describes the meter table entry types.

Ingress packets always have an associated Tunnel Id metadata value. For packets from physical ports this value is always zero. For packets from tunnel logical ports the Tunnel Id metadata value is required and must be greater than zero to identify the tenant forwarding domain. Note that when tenant packets are forwarded the Tunnel Id is automatically supplied to the egress logical port.

4.1 Flow Tables

4.1.1 Ingress Port Flow Table

The Ingress Port Flow Table is the first table in the pipeline and, by convention, is numbered zero. The Ingress Port Flow Table decides whether to forward the packet using the main pipeline or in an isolated (e.g., tenant) forwarding domain.

Packets from certain types of logical ports can be processed differently than packets from physical ports or from other types of logical ports. OpenFlow uses a 32 bit value for ifNums, which allows OF_DPA to divide it into ranges and encode the port type in the high order 16 bits. In this version of OF-DPA, the high order 16 bits are zero for physical ports and one for overlay tunnel logical ports. Port numbering assignments can be found in Table 147, in the Ports section.

The Ingress Port Flow Table presents what is essentially a de-multiplexing logic function as an OpenFlow table that can be programmed from the controller. By default, packets from physical ports with null (zero) Tunnel Id metadata go to the VLAN Flow Table. Packets from logical ports must be accompanied by non-zero Tunnel Id metadata. Entries in this table must admit ingress packets by matching the ingress ifNum exactly, by matching Tunnel Id, or by some combination. For packets from logical ports there is an option to match Tunnel id type by masking the low order bits.

Note: OF-DPA may prevent certain types of rules from being added to other tables unless there is appropriate flow entry in the Ingress Port Flow Table.

The default on miss is for packets from physical ports to go to the VLAN Flow Table. There is no default rule for data center overlay tunnel packets from logical ports, which are dropped on miss.

4.1.1.1 Flow Entry Types and Match Fields

The Ingress Port Flow Table supports the flow entry types listed in Table 4.

Table 4 Ingress Port Flow Table Entry Types

Type	Description
Normal Ethernet Frames	Matches packets from local physical ports, identified by zero Tunnel Id. Normal Ethernet rules have Goto-Table instructions that specify the VLAN Flow Table.
Data Center Overlay Tunnel	Matches packets from a data center overlay tunnel logical port by Tunnel Id type. Overlay Tunnel rules have a Goto-Table instruction that specifies the Bridging Flow Table. The controller must add a rule or rules of this type

Type	Description
Frames	before adding rules in other tables that need to match Data Center Overlay Tunnel packets.

Note: Future versions of OF-DPA may support flow entry types for packets from other types of logical ports.

The Ingress Port Flow Table uses the match fields listed in Table 5. Note that QoS rules must have higher relative priority than overlapping non-QoS rules. DSCP rules must have higher relative priority than PCP rules that match the same flows.

Table 5 Ingress Port Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
IN_PORT	32	No	Yes	Ingress port. Depending on rule may be omitted to match any IN_PORT.
TUNNEL ID	32	Yes	No	Required in Data Center Overlay rules. Must be exact match in order to distinguish frames for different tenants. In Normal Ethernet rules must be zero.

4.1.1.2 Instruction Types

The Ingress Port Flow Table supports the instructions listed in Table 6.

Table 6 Ingress Port Flow Table Instructions

Name	Argument	Description
Goto-Table	Table	Depending on rule type, one of: VLAN Flow Table; or Bridging Flow Table.
Apply-Actions	Action list	Can contain at most one instance of each of the actions listed in Table 7.

4.1.1.3 Actions

The Ingress Port Flow Table actions can optionally set the packet VRF using an action list.

Table 7 Ingress Port Flow Table Action List

Name	Argument	Description
Set-Field	VRF	VRF for L3 lookups. Only applicable to Normal Ethernet Frame rules. Optional.

4.1.1.4 Counters and Flow Expiry

The Ingress Port Flow Table supports the basic table and flow entry counters listed in Table 8.

Table 8 Ingress Port Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported.

4.1.2 VLAN Flow Table

The VLAN Flow Table is used for IEEE 801.Q VLAN assignment and filtering to specify how VLANs are to be handled on a particular port.²² All packets must have an associated VLAN id in order to be processed by subsequent tables. Packets that do not match any entry in the VLAN table are filtered, that is, dropped by default. Note that IEEE defined BPDUs are always received untagged.²³

The VLAN Flow Table can optionally assign a non-zero VRF value to the packet based on the VLAN. OF-DPA defines VRF as a new pipeline metadata field. The VRF defaults to zero if not set.

4.1.2.1 Flow Entry Types and Match Fields

The VLAN Flow Table supports the Flow Entry Types listed in Table 9. Flow entries are differentiated based on IN_PORT, whether or not the packet was tagged, and the VLAN id in the tag. In addition OAM processing rules match packets by packet Ethertype and MAC-DST.

OpenFlow has traditionally used a 16-bit field for VLAN id. Since only the low order 12 bits are needed to express a VLAN id, OpenFlow has defined special values to indicate tagged and untagged packets. In particular, the VLAN id 0x0000 (OFPVID_NONE, defined in the OpenFlow specification) is used to

²² The VLAN Flow Table presents the hardware port and VLAN configuration bitmaps to the OpenFlow controller as a flow table.

²³ There are vendor specific BPDUs that are VLAN tagged.

represent an untagged packet, and 0x1000 (OFPVID_PRESENT) for a priority tagged packet. All tagged packets are represented by VLAN id values between 0x1001 and 0x1FFE²⁴ (OFPVID_PRESENT | VLAN id value). This convention must be followed in programming rules from the controller. For further explanation consult the OpenFlow 1.3.4 specification.

Note: OF-DPA 2.0 does not support matching packets just on whether or not they have a VLAN tag as described in Table 13 of OpenFlow 1.3.4.

Note: At most two tags are supported.

Entries in the OF-DPA VLAN Flow table are mutually exclusive. Any explicit rule priority assignments are ignored.

Table 9 VLAN Flow Table Flow Entry Types

Type	Description
VLAN Filtering	Exact match on IN_PORT and VLAN_VID parsed from the packet. For tagged packets with a VLAN tag containing a VLAN_VID greater than zero. Cannot be masked. VLAN_VID cannot be used in a Port VLAN Assignment rule for untagged packets. The only instruction is Goto-Table and must specify the Termination MAC Flow Table. Tagged packets that do not match any rule are treated as VLAN_VIDs that are not allowed on the port and are dropped. Can optionally assign a VRF for routed packets.
Untagged Packet Port VLAN Assignment	Exact match on IN_PORT and VLAN id == 0 (lower 12 bits of match field) value using a mask value of 0x0fff (masks off OFPVID_PRESENT). Action set must assign a VLAN_VID. The VLAN_VID value cannot be used in a VLAN Filtering rule. If the packet does not have a VLAN tag, one will be pushed if necessary at packet egress. Rule must have a Goto-Table instruction specifying the Termination MAC Flow Table. Untagged packets are dropped if there is no port VLAN assignment rule. Can optionally assign a VRF for routed packets.
Allow All VLANs	Wildcard VLAN match for a specific IN_PORT. Essentially turns off VLAN filtering and/or assignment for a physical port. Must be lower priority than any overlapping translation, filtering, MPLS, or VLAN assignment rule. Untagged packets that match this rule will be assigned an illegal VLAN and may be subsequently dropped. Should also define an L2 Unfiltered Interface group entry for the port.
VLAN Translate, Single Tag, or Single Tag to	Used to either modify the VLAN id on a single tagged packet, or to optionally modify the VLAN id and then push another tag onto a single tagged packet. Can also optionally assign a VRF for routed packets. By OpenFlow convention, the outermost VLAN tag is matched independent of

²⁴ Although accepted by OF-DPA, IEEE 802.1Q indicates VLAN id 4095 is a reserved value.

Type	Description
Double Tag	TPID.
VLAN Translate, Double Tag to Single Tag	Used in conjunction with the VLAN 1 Flow Table to pop the outer tag and optionally modify the remaining tag. Can also optionally assign a VRF for routed packets.
MPLS-TP Single Tagged Pseudo-wire Initiation	Used to classify a packet to an MPLS-TP pseudo-wire and assign a Tunnel_Id and MPLS L2 Port for VPLS or VPWS initiation. If the flow is also enabled for Ethernet OAM (i.e., there is an overlapping but higher relative priority Ethernet OAM frame for that IN_PORT and VLAN Id) must have an action to update LM counters.
MPLS-TP Double Tagged Pseudo-wire Initiation	Used in conjunction with the VLAN 1 Flow Table to classify a double tagged packet to a pseudo-wire. The outer tag must be popped in order to expose the inner tag for match by VLAN 1 Flow Table.
MPLS-TP Pseudo-wire Initiation, All Traffic on Port	Used to classify all packets on a port to a pseudo-wire and assign a Tunnel_Id and MPLS L2 Port for MPLS-TP VPLS or VPWS initiation. Must be higher relative priority than any overlapping rule. If enabled for Ethernet OAM (i.e., there is an overlapping but higher relative priority Ethernet OAM Frame rule) must have an action to update LM counters.
Ethernet OAM Frame	Enable for OAM. If packet is an IEEE 802.1ag CFM frame (ETH-TYPE is 0x8902), the Goto-Table instruction specifies the Maintenance Point Flow Table to determine where to forward the OAM frame for processing. MAC-DST is required, but might be the multicast group address defined for this purpose (01-80-C2-00-00-3x, where the last four bits can be masked).

Note: The untagged packet rule applies to both untagged packets, which match VLAN_VID = 0x1000, and IEEE 802.1P priority tagged packets, which match VLAN_VID = 0x0000. However the VLAN-PCP match field will be set from the value in a priority VLAN tag rather than default to zero in the case of a packet without a VLAN tag.

Note: A VLAN Flow Table rule cannot specify an IN_PORT and VLAN_VID combination that is used in a VXLAN Access Logical Port configuration. Conversely, it must include a rule to permit an IN_PORT and VLAN_VID combination used in a VXLAN Tunnel Next Hop configuration.

The VLAN Flow Table match fields are listed in Table 10.

Table 10 VLAN Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
-------	------	----------	----------	-------------

Field	Bits	Maskable	Optional	Description
IN_PORT	32	No	No	Ingress port. Must be a physical port (high order 16 bits zero).
VLAN_VID	16	Yes	No	Outer VLAN id. The mask value can only be either 0x1fff for VLAN filtering rules and 0x0fff for untagged packet rules. Must be exact for other rules.
ETH_TYPE	16	No	Yes	Only allowed value is 0x8902, used in rules to match OAM frames. In non-OAM rules must be omitted.
MAC-DST	48	Yes	Yes	Required in rules that match OAM frames; must be omitted in other rule types.

4.1.2.2 Instruction Types

The VLAN table supports the instruction types listed in Table 11.

Table 11 VLAN Flow Table Instructions

Name	Argument	Description
Apply-Actions	Action List	The VLAN Flow Table supports the actions specified in Table 12.
Write-Actions	Action List	The VLAN Flow Table supports the actions specified in Table 13.
Goto-Table	Table	For VLAN filtering or Port VLAN assignment the next table should be the Termination MAC Flow Table. For VLAN translation or MPLS double tag classification the next table can be the VLAN 1 Flow Table. For OAM frames the next table must be the Maintenance Point Flow Table. A packet is dropped if it matches an entry that has no Goto-Table instruction.

4.1.2.3 Actions

The VLAN table uses Apply Actions for port VLAN tagging and assignment, and for VRF assignment. The action list can have at most one entry of each action type.

Table 12 VLAN Flow Table Action List Actions

Name	Argument	Description
Set Field	VLAN_VID, must be between 1 and 4094.	Sets the VLAN id on the outermost tag. If the packet is untagged then one is pushed with the specified VLAN id and priority zero.
Set Field	VRF	Optionally sets the VRF pipeline field. VRF must be the same in all rules for the same VLAN.
Set Field	OVID	Pipeline metadata field representing an outer tag VLAN Id that was popped, so that it can be used as a match field in the VLAN 1 Flow Table. Only required in flow entries with Goto-Table instructions specifying the VLAN 1 Flow Table.
Push VLAN	TPID	Used in translating single to double tag. TPID must be 0x8100 (inner VLAN tag) or 0x88a8 (outer VLAN tag).
Pop VLAN		Used in processing double tagged frames, where the Goto-Table instruction specifies the VLAN 1 Flow table.
Set Field	MPLS L2 Port	For pseudo-wire classification.
Set Field	Tunnel_Id	For pseudo-wire classification. Can be either a VPLS or a VPWS local port type.
Set-Field	LMEP_Id	Indicates MEP or MIP for OAM PDU processing.

Note: The untagged packet action is the same as in OpenFlow 1.0. The implicit addition of a tag to an untagged packet is tolerated but not condoned in OpenFlow 1.3.4.

Note: Untagged packet flows cannot be enabled for Ethernet OAM.

The VLAN table uses Write Actions for updating an OAM LM counter for this frame. The counter is updated at the end of the pipeline, after metering, in order to handle the case where the packet is dropped.

Table 13 VLAN Flow Table Action Set Actions

Name	Arguments	Description
------	-----------	-------------

Name	Arguments	Description
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Indicates MEP for which LM counters are to be incremented for data frames. Traffic Class defaults to zero if not set.

4.1.2.4 Counters and Flow Expiry

The VLAN Flow Table supports the table and flow entry counters listed in Table 14.

Table 14 VLAN Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported.

4.1.3 VLAN 1 Flow Table

The VLAN 1 Flow Table is used for double tag matching and actions. OpenFlow is defined such that it always matches the outermost VLAN tag. For OF-DPA, the VLAN table sets a pipeline metadata field (OVID), pops the outermost tag, and does a go to the VLAN 1 table. This second table matches what was the inner VLAN tag as the outermost tag but can also match on what was the outermost tag. Actions are similar to the single tag VLAN case.

4.1.3.1 Flow Entry Types and Match Fields

The VLAN 1 Flow Table supports the Flow Entry Types listed in Table 15.

Table 15 VLAN 1 Flow Table Flow Entry Types

Type	Description
VLAN Assignment	Exact match on IN_PORT, VLAN_VID, and OVID. Can optionally pop the tag, change the VLAN_VID, or push another tag and set the VLAN_VID for the pushed tag. Must have a Goto-Table instruction that specifies the Termination MAC Flow Table. Can also assign a VRF for MPLS L3 VPN.
MPLS L2 Stacked VLAN	Exact match on IN_PORT, VLAN_VID, and OVID. Can optionally pop the tag, change the VLAN_VID push another tag, or set the VLAN_VID for the pushed tag. Must assign an MPLS L2 Port value. Must have a Goto-Table instruction that specifies the MPLS L2 Port table. If enabled for Ethernet OAM (i.e., there is an overlapping but higher relative priority Ethernet OAM

Type	Description
	Frame rule) must specify an action to update LM counters.
Ethernet OAM Frame	Enabled for OAM. If packet is an IEEE 802.1ag CFM frame (ETH-TYPE is 0x8902), the rule must have a Goto-Table instruction that specifies the Maintenance Point Flow Table which will decide where to forward the OAM frame for processing.

The VLAN 1 Flow Table match fields are listed in Table 16.

Table 16 VLAN 1 Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
IN_PORT	32	No	No	Ingress port. Must be a physical port (high order 16 bits zero).
VLAN_VID	16	Yes	No	Inner VLAN id. Must be exact.
OVID	16	No	No	Outer VLAN id, set by a VLAN table flow entry.
ETH_TYPE	16	No	Yes	Only allowed value is 0x8902. In non-OAM rules must be omitted.
MAC-DST	48	Yes	Yes	Required in rules that match OAM frames, must be omitted in other rule types.

4.1.3.2 Instruction Types

The VLAN table supports the instruction types listed in Table 17.

Table 17 VLAN 1 Flow Table Instructions

Name	Argument	Description
Apply-Actions	Action List	The VLAN 1 Flow Table supports the actions specified in Table 18.
Write-Actions	Action List	The VLAN Flow Table supports the actions specified in Table 19.

Name	Argument	Description
Goto-Table	Table	For VLAN assignment must be the Termination MAC Flow Table. For OAM frames must be the Maintenance Point Flow Table. For MPLS-TP must be the MPLS L2 Port table.

4.1.3.3 Actions

The VLAN 1 Flow Table action list actions are as shown in Table 18.

Table 18 VLAN 1 Flow Table Action List Actions

Name	Argument	Description
Set Field	VLAN_VID	
Set Field	VRF	Optionally sets the VRF pipeline field. VRF must be the same in all rules for the same VLAN.
Push VLAN	TPID	Used in translating single to double tag. TPID must be one of 0x8100 or 0x88a8.
Pop VLAN		Used in processing double tagged frames.
Set Field	MPLS L2 Port	For pseudo-wire classification.
Set Field	Tunnel_Id	For pseudo-wire classification. Can be either VPLS or VPWS type.
Set-Field	LMEP_Id	Indicates MEP or MIP for OAM PDU processing.

The VLAN 1 Flow Table action list actions are as shown in Table 19.

Table 19 VLAN 1 Flow Table Action Set Actions

Name	Argument	Description
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Indicates MEP for which LM counters are to be incremented for data frames. Traffic Class defaults to zero if not set.

4.1.3.4 Counters and Flow Expiry

The VLAN 1 Flow Table supports the table and flow entry counters listed in Table 20.

Table 20 VLAN 1 Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported, as indicated in Table 21.

Table 21 VLAN 1 Flow Table Expiry

Name	Bits	Description
Hard Timeout	32	Number of seconds after which flow entry is removed. Optional, entry does not age out if zero or not specified.

4.1.4 MPLS L2 Port Flow Table

The MPLS L2 Port Flow Table is used for MPLS tunnel origination and to support per-customer counters. In future versions of OF-DPA it will be used for QoS classification for MPLS flows and for learning.

The MPLS L2 Port pipeline metadata encoding uses ranges distinguish different VPWS forwarding cases as shown in Table 22.

Table 22 MPLS L2 Port Metadata Naming Convention

Type	Numbering	Description
VPWS Local	0x0000nnnn	E-Line local client interface (UNI). Assigned by VLAN table.
VPWS Network	0x0002nnnn	E-Line network interface (NNI). Assigned by terminating MPLS label.

4.1.4.1 Flow Entry Types and Match Fields

The MPLS L2 Flow Table supports the Flow Entry Types listed in Table 23.

Table 23 MPLS L2 Port Flow Table Flow Entry Types

Type	Description
MPLS VPWS	Match on MPLS_L2_Port and Tunnel Id; Group (MPLS L2 VPN label or Fast Failover group). Must have a Goto-Table instruction that specifies the Policy ACL Flow Table.

The MPLS L2 Port Flow Table match fields are listed in Table 24.

Table 24 MPLS L2 Port Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
MPLS L2 Port	32	Yes	No	Either exact match or bit masked (0x00010000) to select whether a VPWS port.
Tunnel Id	32	No	No	Must be type MPLS-TP.

4.1.4.2 Instruction Types

The MPLS L2 Port table supports the instruction types listed in Table 25.

Table 25 MPLS L2 Port Flow Table Instructions

Name	Argument	Description
Write-Actions	Action Set	Only used for VPWS. Only action is Group, which must indicate one of: MPLS L2 VPN Label or Fast Failover.
Goto-Table	Table	Must be the Policy ACL Flow Table.

4.1.4.3 Counters and Flow Expiry

The MPLS L2 Port Flow Table supports the table and flow entry counters listed in Table 14.

Table 26 MPLS L2 Port Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed
Received Packets	64	Per-entry	Number of packets that hit this flow entry.
Received Bytes	64	Per-entry	Number of bytes that hit this flow entry.

Only hard interval timeout ageing per entry is supported.

4.1.5 Termination MAC Flow Table

The Termination MAC Flow Table determines whether to do bridging or routing on a packet. It identifies routed packets their destination MAC, VLAN, and Ethertype. Routed packet rule types use a Goto-Table instruction to indicate that the next table is one of the routing tables. The default on a miss is the Bridging Flow Table.

4.1.5.1 Flow Entry Types and Match Fields

The Termination MAC Flow Table implements the flow entry types listed in Table 27.

Table 27 Termination MAC Flow Table Entry Types

Name	Description
Unicast MAC	Used to identify an IPv4 or IPv6 router MAC address. Relative priority must be assigned so as to be lower than any multicast MAC rule. Must have a Goto-Table instruction specifying the Unicast Routing Flow Table.
IPv4 Multicast MAC	Wildcard rule that recognizes all IPv4 multicast MAC addresses specified in RFC 1112. If specified, this must be ETH_DST = 01-00-5e-00-00-00 with mask ff-ff-ff-80-00-00. There can only be one flow entry of this type. Must have a Goto-Table instruction specifying the Multicast Routing Flow Table.
IPv6 Multicast MAC	Wildcard rule that recognizes all IPv6 MAC addresses specified in RFC 2464. If specified, this must be ETH_DST = 33-33-00-00-00-00 with mask ff-ff-00-00-00-00. There can only be one flow entry of this type. Must have a Goto-Table instruction specifying the

Name	Description
	Multicast Routing Flow Table.
MPLS	Used to identify an MPLS label switch router or edge device MAC address. Must have a Goto-Table instruction specifying MPLS Flow Table 0 unless the platform does not support this table, in which it must specify MPLS Flow Table 1.

The Termination MAC Flow Table match fields are listed in Table 28. Strict rule priority must be assigned by the controller so that every flow entry has a unique priority.

Table 28 Termination MAC Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
IN_PORT	32	No	Yes	Physical (local) input port.
ETH_TYPE	16	No	No	Pre-requisite for IPv4 (0x0800), IPv6 (0x86dd), or MPLS (0x8847).
ETH_DST	48	No	No	Ethernet destination MAC. Prefix maskable for only the specific multicast IP flow entries in Table 27. Can only be field masked for unicast destination MACs.
VLAN_VID	16	Yes	Yes	Matches against the Outer VLAN id. Must be either omitted or exact.
IPv4_DST	32	Yes	Yes	Can only be used with 224/8 address and 224.0.0.0 mask values, otherwise must be omitted. Pre-requisite ETH_TYPE must be 0x0800.
IPv6_DST	128	Yes	Yes	Can only be used with FF00::/8 address and FF00:0:0:0:0:0 mask values, otherwise must be omitted. Pre-requisite ETH_TYPE must be 0x86dd.

4.1.5.2 Instruction Types

The Termination MAC Flow Table can have the instructions shown in Table 29.

Table 29 Termination MAC Flow Table Instruction Set

Name	Argument	Description
Goto-Table	Table	Unicast MAC rules with multicast IPV4_DST or IPV6-DST should specify the Multicast Routing Flow Table, otherwise they can only specify the Unicast Routing Flow Table. Multicast MAC rules can only specify the Multicast Routing Flow Table. MPLS rules must specify MPLS Flow Table 0 on platforms that support it, otherwise must specify MPLS Flow Table 1. The packet is dropped if the rule matches and there is no Goto-Table instruction.
Apply Actions	Action List	Optional. If supplied can only contain one action, output a copy to CONTROLLER.

4.1.5.3 Counters and Flow Expiry

The Termination MAC Flow Table counters are listed in Table 30.

Table 30 Termination MAC Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Number of active flow entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Termination MAC Flow Table only supports hard interval expiration.

4.1.6 Bridging Flow Table

The Bridging Flow Table supports Ethernet packet switching for potentially large numbers of flow entries using the hardware L2 tables. The default on a miss is to go to the Policy ACL Flow Table.

Note: The Policy ACL Flow Table is recommended for matching BPDUs.

The Bridging Flow Table forwards either based on VLAN (normal switched packets) or Tunnel id (isolated forwarding domain packets), with the Tunnel id metadata field used to distinguish different flow table entry types by range assignment. The naming convention for Tunnel id metadata is described in Table 31.

Table 31 Tunnel Id Metadata Naming Convention

Type	Numbering	Description
Data Center Overlay	0x0000nnnn	Identifies a data center overlay tenant isolated forwarding domain
MPLS-TP	0x0001nnnn	Identifies an MPLS-TP pseudo-wire isolated forwarding domain

The Bridging Flow Table flow entry types are listed in Table 32. VLAN and Tunnel id are mutually exclusive.

Table 32 Bridging Flow Table Flow Entry Types

Type	Description
Unicast VLAN Bridging	Matches switched unicast Ethernet frames by VLAN id and MAC_DST. MAC_DST must be unicast and cannot be masked. VLAN id must be present and non-zero. Tunnel id must be masked or omitted.
Multicast VLAN Bridging	Matches switched multicast Ethernet frames by VLAN id and MAC_DST. MAC_DST must be multicast and cannot be masked. VLAN id must be present and non-zero. Tunnel id must be masked or omitted.
DLF VLAN Bridging	Matches switched Ethernet frames by VLAN id only. MAC_DST must be field masked and match any destination. Must have lower relative priority than any unicast or multicast flow entries that specify this VLAN. VLAN id must be present and non-zero. Tunnel id must be masked or omitted.
Unicast Data Center Overlay Bridging	Matches switched unicast Ethernet frames by tunnel id and MAC_DST. MAC_DST must be unicast and cannot be masked. Tunnel id must be non-zero, type overlay tunnel (0x0000nnnn), and cannot be masked. VLAN id must be masked or omitted.
Multicast Data Center Overlay Bridging	Matches switched multicast Ethernet frames by tunnel id and MAC_DST. MAC_DST must be multicast and cannot be masked. Tunnel id must be non-zero, type overlay tunnel (0x0000nnnn), and cannot be masked. VLAN id must be masked or omitted.
DLF Data Center Overlay	Matches switched Ethernet frames by tunnel id only. MAC_DST is must be field masked and match any destination. Must have

Type	Description
Bridging	lower relative priority than any unicast or multicast flow entries that specify this tunnel id. Tunnel id must be non-zero, type overlay tunnel (0x0000nnnn), and cannot be masked. VLAN id must be masked or omitted.

Note: Exact match rules must be given higher relative priority assignments than any potentially overlapping wildcard rules. In any event, exact match rules are evaluated before any wildcard rules.

4.1.6.1 Flow Entry Types and Match Fields

Match fields for flow entry types are described in the following tables.

Table 33 Bridging Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
ETH_DST	48	Yes	Yes	Ethernet destination MAC, allowed values depend on flow entry type. Exact match only (mask must be all 1's if supplied).
VLAN_VID	16	Yes	Yes	VLAN id, allowed values depend on flow entry type. Exact match only (mask must be all 1's if supplied).
TUNNEL ID	32	Yes	Yes	Identifies isolated forwarding domain for data center overlay traffic. Allowed values depend on flow entry type. Exact match only.

4.1.6.2 Instruction Types

Default next table if no match is the ACL Policy Flow Table.

Table 34 Bridging Flow Table Instructions

Name	Argument	Description
Write-Actions	Action set	Only the actions in Section 4.1.6.3 can be specified.
Apply-Actions	Action list	Optional. If specified, can contain only a single output action to send a copy to CONTROLLER
Goto-Table	Table	Must be the ACL Policy Flow Table if specified. If packet matches and no next table is specified then the packet is

Name	Argument	Description
		dropped.

4.1.6.3 Action Set

The Bridging Flow Table supports the actions in Table 35 by flow entry type. The OF-DPA API validates consistency of flow entry type and OF-DPA group entry type references.

Table 35 Bridging Flow Table Action Set

Flow Entry Type	Argument	Description
Unicast VLAN Bridging	Group id	Must be an OF-DPA L2 Interface group entry for the forwarding VLAN.
Multicast VLAN Bridging	Group id	Must be an OF-DPA L2 Multicast group entry for the forwarding VLAN.
DLF VLAN Bridging	Group id	Must be an OF-DPA L2 Flood group entry for the forwarding VLAN.
Unicast Data Center Overlay Bridging	Output	Must be an overlay tunnel logical port for the tenant overlay forwarding domain tunnel id.
Multicast Data Center Overlay Bridging	Group id	Must be an OF-DPA L2 Overlay Multicast sub-type group entry for the tenant overlay forwarding domain tunnel id.
DLF Data Center Overlay Bridging	Group id	Must be an OF-DPA L2 Overlay Flood sub-type group entry tenant overlay forwarding domain tunnel id.

4.1.6.4 Counters and Flow Expiration

The Bridging Flow Table counters are listed in Table 36.

Table 36 Bridging Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Number of active entries in the table

Name	Bits	Type	Description
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Bridging Flow Table flow entry expiration is as shown in Table 37.

Table 37 Bridging Flow Table Flow Entry Expiration

Name	Bits	Description
Hard Timeout	32	Number of seconds after which flow entry is removed. Optional, entry does not age out if zero or not specified.
Idle Timeout	32	Number of seconds of inactivity, after which a flow entry is removed. Optional, flow entry does not age out if unspecified or zero.

4.1.7 Unicast Routing Flow Table

The Unicast Routing Flow Table supports routing for potentially large numbers of IPv4 and IPv6 flow entries using the hardware L3 tables.

The Unicast Routing Flow Table is a single table but organized as two mutually exclusive logical sub-tables by IP protocol, and supports the flow entry types listed in Table 38. A single table number is used for both logical tables.

Table 38 Unicast Routing Flow Table Entry Types

Type	Table	Prerequisite(s)	Description
IPv4 Unicast	Table 39	Ethertype=0x0800	Matches routed unicast IPv4 packets. The Goto-Table instruction specifies the Policy ACL Table.
IPv6 Unicast	Table 40	Ethertype=0x86dd	Matches routed unicast IPv6 packets. The Goto-Table instruction specifies the Policy ACL Table.

The Unicast Routing Flow Table can support multiple virtual routing tables. The VRF pipeline metadata match field value identifies the virtual routing table to use for a particular packet lookup.

4.1.7.1 Flow Entry Types and Match Fields

Match fields for flow entry types are described in the following tables.

Table 39 Unicast Routing Flow Table IPv4 Header Match Fields

Field	Bits	Maskable	Optional	Description
ETH_TYPE	16	No	No	Must be 0x0800
VRF	16	No	Yes	If omitted or zero indicates the default routing table.
IPv4 DST	32	Yes	No	Must be a unicast IPv4 address. Prefix maskable only, mask used for LPM forwarding. For the MPLS BFD rule, must have a value in 128/8 with mask 128.0.0.0.

Table 40 Unicast Routing Flow Table IPv6 Header Match Fields

Field	Bits	Maskable	Optional	Description
ETH_TYPE	16	No	No	Must be 0x86dd
VRF	16	No	Yes	If omitted or zero indicates the default routing table.
IPv6_DST	128	Yes	No	Must be a unicast IPv6 address. Prefix maskable only, used for LPM forwarding. For the MPLS BFD rule, must have a value in 0:0:0:0:FFFF:7F00/104 with mask 0:0:0:0:FFFF:7F00:0.

Note: Exact match rules must be given higher relative priority assignments than any LPM prefix match rules. In any event, the hardware evaluates exact match rules before any wildcard rules.

Note: Rules that specify a non-zero VRF must have higher relative priority than other overlapping rules. The wildcard rules are effectively “global” or “default” in that they are matched last, that is, if no specific VRF rule matches the packet. If the packet VRF is zero it can only match one of the wildcard rules.

4.1.7.2 Instruction Types

Default next table on a miss is the ACL Policy Flow Table.

Table 41 Unicast Routing Flow Table Instructions

Name	Argument	Description
Write-Actions	Action set	Only the actions in Table 42 can be specified.
Clear-Actions	-	Used to delete any forwarding decision so that the packet will be dropped.
Goto-Table	Table	Must specify the ACL Policy Flow Table. Required.

Other instruction types, specifically Apply Actions, are not supported.

4.1.7.3 Action Set

The actions in Table 42 are supported.

Table 42 Unicast Routing Flow Table Action Set

Name	Argument	Description
Group	Group id	Must be an OF-DPA L3 Unicast or L3 ECMP Group Entry.
Decrement TTL and do MTU check	-	MTU check is a vendor extension. An invalid TTL (zero before or after decrement) is always dropped and a copy sent to the CPU for forwarding to the CONTROLLER. Similarly, a packet that exceeds the MTU is dropped and a copy sent to the CONTROLLER. Required.

4.1.7.4 Counters and Flow Expiration

The Routing Flow Table counters are listed in Table 43.

Table 43 Unicast Routing Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Unicast Routing Flow Table flow entry expiration provisions are as shown in Table 44.

Table 44 Unicast Routing Flow Table Flow Entry Expiration

Name	Bits	Description
Hard Timeout	32	Number of seconds after which flow entry is removed. Optional, entry does not age out if zero or not specified.
Idle Timeout	32	Number of seconds of inactivity, after which a flow entry is removed. Optional, entry does not age out if zero or not specified.

4.1.8 Multicast Routing Flow Table

The Multicast Routing Flow Table supports routing for IPv4 and IPv6 multicast packets.

The Multicast Routing Flow Table can also support multiple virtual routing tables, matching on the packet VRF field value.

The Multicast Routing Flow Table is also organized as two mutually exclusive logical sub-tables by IP protocol, and supports the flow entry types listed in Table 45.

Table 45 Multicast Routing Flow Table Entry Types

Type	Table	Prerequisite(s)	Description
IPv4 Multicast	Table 46	Ethertype=0x0800	Matches routed multicast IPv4 packets.
IPv6 Multicast	Table 47	Ethertype=0x86dd	Matches routed multicast IPv6 packets.

4.1.8.1 Flow Entry Types and Match Fields

Match fields for flow entry types are described in the following tables.

Table 46 Multicast Routing Flow Table IPv4 Match Fields

Field	Bits	Maskable	Optional	Description
ETH_TYPE	16	No	No	Must be 0x0800. Required pre-requisite.
VLAN_VID	16	No	No	VLAN id

Field	Bits	Maskable	Optional	Description
VRF	16	No	Yes	VRF.
IPV4_SRC	32	Yes	Yes	Cannot be bit masked, but can be omitted.
IPV4_DST	32	Yes	No	Must be an IPv4 multicast group address.

Table 47 Multicast Routing Flow Table IPv6 Match Fields

Field	Bits	Maskable	Optional	Description
ETH_TYPE	16	No	No	Must be 0x86dd. Required prerequisite.
VLAN_VID	16	No	No	VLAN id
VRF	16	No	Yes	VRF.
IPV6_SRC	128	Yes	Yes	Cannot be bit masked, but can be omitted.
IPV6_DST	128	Yes	No	Must be an IPv6 multicast group address.

Note: Rules that specify a non-zero VRF are matched at higher relative priority than wildcard VRF rules.

4.1.8.2 Instruction Types

Default next table on miss is the ACL Policy Flow Table.

Table 48 Multicast Routing Flow Table Instructions

Name	Argument	Description
Write Actions	Action set	Only the actions in Table 48 can be specified.
Goto-Table	Table	Must be the Policy ACL Flow Table. In the event that there is no group entry referenced and no next table specified, the packet will be dropped.

Other instruction types, specifically Apply Actions, are not supported.

4.1.8.3 Action Set

The Multicast Routing Table supports the actions in Table 49.

Table 49 Multicast Routing Flow Table Action Set

Name	Argument	Description
Group	Group id	Must be an OF-DPA L3 Multicast group entry with the forwarding VLAN id as a name component.
Decrement TTL and do MTU check	-	MTU check is a vendor extension. An invalid TTL (zero before or after decrement) is always dropped and a copy sent to the CPU for forwarding to the CONTROLLER. Similarly, a packet that exceeds the MTU is dropped and a copy sent to the CONTROLLER. Required.

4.1.8.4 Counters and Flow Expiration

The Multicast Routing Flow Table counters are as shown in Table 50.

Table 50 Multicast Routing Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Multicast Routing Flow Table Flow entry expiration provisions are as shown in Table 51.

Table 51 Multicast Routing Flow Table Flow Entry Expiration

Name	Bits	Description
Hard Timeout	32	Number of seconds after which flow entry is removed. Optional, entry does not age out if zero or not specified.
Idle Timeout	32	Number of seconds of inactivity, after which a flow entry is removed. Optional, entry does not age out if zero or not specified.

4.1.9 MPLS Flow Tables

The MPLS pipeline can support three MPLS Flow Tables, MPLS Table 0, MPLS Table 1 and MPLS Table 2. An MPLS Flow Table lookup matches the label in the outermost MPLS shim header in the packets. MPLS Table 0 is only used to pop a protection label on platforms that support this table, or to detect an MPLS-TP Section OAM PDU. MPLS Table 1 and MPLS Table 2 can be used for all label operations. MPLS Table 1 and MPLS Table 2 are synchronized flow tables and updating one updates the other.

MPLS Table 0 only matches an outermost label that is either GAL (13) and bottom of stack (for LSP Section OAM), or else not bottom of stack. MPLS Table 0 has a built-in default rule with a Goto-Table instruction specifying MPLS Table 1 on a miss. An MPLS Table 0 rule can pop the outermost label and have a Goto-Table instruction specifying the MPLS Table 1. If the pop exposes an OAM frame (GAL or RAL) and the label is enabled for OAM, the frame can be forwarded to an OAM engine for further processing. If the pop is otherwise enabled for OAM, and LM counters are defined for this LMEP, then the LM counters are updated.

MPLS Table 1 matches the outermost label. If the label is not bottom of stack, it can pop the outermost label with a Goto-Table instruction specifying MPLS Table 2, for matching the next label. Up to three labels can be explicitly matched in this way, depending on platform. An entry that matches bottom of stack set can only be added or modified to MPLS Table 1 or MPLS Table 2.

L3 VPN rules for IPv6 unicast, IPv4 multicast, and IPv6 multicast are automatically added for the same label when an IPv4 L3 VPN rule is added.

4.1.9.1 Flow Entry Types and Match Fields

MPLS Table 0 Flow supports the flow entry types in Table 52.

Table 52 MPLS Flow Table 0 Flow Entry Types

Type	Prerequisite(s)	Actions	Use Case
Pop Tunnel Label	MPLS_BOS = 0	Pop outermost tunnel label; Set Ethertype to 0x8847; Decrement/check TTL; Optionally Copy TTL in; Optionally Copy EXP in; Goto-Table instruction specifies MPLS Table 1.	Pop Protection Label and forward based on an inner label
Pop Tunnel Label (MEP Data Frame)	MPLS_BOS = 0	Pop outermost tunnel label; Set Ethertype to 0x8847; Decrement/check TTL; Optionally Copy TTL in; Optionally Copy EXP in; Goto-Table instruction specifies MPLS Table 1. Increment OAM LM Counters for this LMEP.	MPLS-TP LSP OAM

Type	Prerequisite(s)	Actions	Use Case
Pop Tunnel Label (MEP OAM Frame)	MPLS_BOS = 0	Set the LMEP Id, Goto-Table instruction specifies the Maintenance Point Flow Table.	MPLS-TP LSP OAM
Section MEP OAM Frame	MPLS_BOS = 1	Exact match on GAL (13). Goto-Table instruction specifies the Maintenance Point Flow Table.	MPLS-TP Section OAM

MPLS Tables 1 and 2 support the flow entry types in Table 53. These match the same fields but actions differ depending on the table and packet flow use case.

Table 53 MPLS Flow Table 1 and 2 Flow Table Entry Types

Type	Prerequisite(s)	Actions	Use Case
Pop Tunnel Label	MPLS_BOS = 0	Pop outermost tunnel label, Ethertype stays 0x8847; Check TTL; Optionally Copy TTL in; Optionally Copy EXP in; Goto-Table instruction specifies MPLS Table 2.	Pop Protection Label and forward based on an inner label
Penultimate Hop Pop	MPLS_BOS = 0	Pop outermost tunnel label, Ethertype stays 0x8847; Decrement/check TTL; Optionally Copy TTL in; Optionally Copy EXP in; Group: MPLS Interface; Goto-Table instruction specifies the Policy ACL Flow Table.	Pop and forward based on this label to a next hop router. (PHP)
Swap Tunnel Label	MPLS_BOS = 0	Decrement/check TTL; Group: MPLS Swap, MPLS ECMP, or MPLS Fast Failover; Goto-Table instruction specifies the Policy ACL Flow Table.	Swap and forward based on this label (LSR)
Swap Pseudo-wire Label	MPLS_BOS = 1	Decrement/check TTL; Group: MPLS Swap or MPLS Fast Failover; Goto-Table instruction specifies the Policy ACL Flow Table.	Swap and forward based on this label (MS-PW, LSR)
Pop Tunnel Label (MEP Data Frame)	MPLS_BOS = 0	Pop outermost tunnel label, Ethertype stays 0x8847; Decrement/check TTL; Optionally Copy TTL in; Optionally Copy EXP in; Increment OAM LM Counters for this LMEP; Goto-Table	MPLS-TP LSP OAM MEP

Type	Prerequisite(s)	Actions	Use Case
		instruction specifies MPLS Table 1.	
Pop Tunnel Label (MEP OAM Frame)	MPLS_BOS = 0, next label is GAL.	Set the LMEP Id; Goto-Table instruction specifies the Maintenance Point Flow Table.	MPLS-TP LSP OAM MEP
Swap Tunnel Label (MEP OAM Frame)	MPLS_BOS = 0, TTL=1 and next label is GAL.	Set the LMEP Id; Goto-Table instruction specifies the Maintenance Point Flow Table.	MPLS-TP LSP MIP
Swap Pseudo-wire Label (MEP OAM Frame)	MPLS_BOS=1, next label is GAL	Set the LMEP Id; Goto-Table instruction specifies the Maintenance Point Flow Table.	MPLS-TP MEP (MS-PW, LSR)
L3 VPN Route (Unicast)	MPLS_BOS = 1	Pop outermost (VRF) label; Set Ethertype to either 0x0800 or 0x86dd; decrement/check TTL and optionally copy in; optionally Copy EXP in; Set Field VRF; Goto-Table instruction specifies the Unicast Routing Flow Table. IPv6 rule automatically added when IPv4 rule added, as are multicast rules.	Pop and do routing lookup on unicast inner IP (LER)
L3 VPN Route (Multicast)	MPLS_BOS = 1	Pop outermost (VRF) label; Set Ethertype to 0x0800 or 0x86dd (depending on IP protocol); Decrement/check TTL and optionally copy in; Optionally Copy EXP in; Set Field VRF; Goto-Table instruction specifies the Multicast Routing Flow Table. Automatically added for both IPv4 and IPv6 when IPv4 L3 VPN Route rule added.	Pop and do routing lookup on multicast inner IP (LER)
L2 Switch VPWS	MPLS_BOS = 1	Pop outermost (pseudo-wire) label; pop outer L2 header; Decrement/check TTL. If BFD frame optionally copy TTL in; Goto-Table instruction specifies the BFD Flow Table. If G.8113.1 frame, output to the Controller or LOCAL. Else optionally pop CW; Set Field MPLS L2 Port; Set Field Tunnel_Id; Group: MPLS L2 Tag, L2 Interface, or L2 Unfiltered Interface; Goto-Table instruction	Pop, decap, and L2 forward (MPLS-TP VPWS PW termination)

Type	Prerequisite(s)	Actions	Use Case
		specifies the Policy ACL Flow Table.	
L2 Switch VPWS (MEP Data Frame)	MPLS_BOS = 1	As with L2 Switch VPWS rule. If G.8113.1, Set Field LMEP_Id; increment OAM LM Counters for this LMEP.	Pop, decap, and L2 forward (MPLS-TP VPWS PW termination)
PW VCCV 1 (MEP OAM Frame)	MPLS_BOS = 1, ACH control word	Set the LMEP Id; Goto-Table Instruction specifies the Maintenance Point Flow Table.	
PW VCCV 3 (MEP OAM Frame)	MPLS_BOS = 1, TTL=1, ACH control word	Set the LMEP Id; Goto-Table Instruction specifies the Maintenance Point Flow Table.	
PW VCCV 4 (MEP OAM Frame)	MPLS_BOS = 0, next label is GAL, ACH control word	Set the LMEP Id; Goto-Table Instruction specifies the Maintenance Point Flow Table.	

The MPLS Flow Table match fields for all flow entry types are shown in Table 54.

Table 54 MPLS Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
ETH_TYPE	16	No	No	Must be 0x8847. This is the OpenFlow required pre-requisite for MPLS matching.
MPLS_BOS	1	No	No	Bottom of stack
MPLS_LABEL	20	No	No	Outermost label
IN_PORT	32	Yes	Yes	Physical (local) input port. Field maskable only.
MPLS_TTL	8	No	Yes	TTL=1 is the only value matched, drop unless an OAM PDU (indicated by a match on NEXT_LABEL_IS_GAL)

Field	Bits	Maskable	Optional	Description
MPLS_DATA_FIRST_NIBBLE	4	No	Yes	High order 4 bits of control word. 0000b for CW, 0001b for ACH. For L3 VPN, matches the IP version for setting the Ethertype.
MPLS_ACH_CHANNEL	16	No	Yes	Parsed from ACH Channel Type field if MPLS_DATA_FIRST_NIBBLE is 1. Only value is 0x8902, which is for G.8113.1.
NEXT_LABEL_IS_GAL	1	No	Yes	Parser peeks at next label, detects GAL and MPLS_BOS, and sets this pipeline match field.

4.1.9.2 Instruction Types

The MPLS Flow Table can have the instructions shown in Table 55.

Table 55 MPLS Flow Table Instructions

Name	Argument	Description
Goto-Table	Table	Depends on rule type as detailed in Table 52 and Table 53.
Apply Actions	Action List	Allowed actions are listed in Table 56 and depend on rule type.
Write Actions	Action Set	Allowed actions are listed in Table 57 and depend on rule type.

4.1.9.3 Actions

The MPLS Flow Table action list supports the actions in Table 57. These are applied immediately to the frame.

Table 56 MPLS Flow Table Action List

Name	Argument	Description
Pop Label	Ethertype	If BOS rules set the packet Ethertype based on matching the value of MPLS_DATA_FIRST_NIBBLE. Otherwise argument should be 0x8847, although it may not always be relevant for

Name	Argument	Description
		setting the packet Ethertype.
Decrement TTL		This action must always be performed. If TTL is invalid after decrement packet is sent to Controller. This does not apply to VCCV Type 3 frames, which should be output for OAM processing.
Copy TTL in		Optional. Only applies if there is an inner label or IP packet and TTL is valid after decrement.
Copy TC in		Optional. Only applies if there is an inner label. Primarily used to set the inner label EXP for PHP.
Set-Field	VRF	Required for L3 Route rules, otherwise not used.
Set-Field	MPLS L2 Port	Required for VPLS or VPWS pseudo-wire termination.
Set-Field	Tunnel Id	Required for VPLS or VPWS pseudo-wire termination.
Set-Field	Traffic Class	If specified, overrides Traffic Class from MPLS QoS table. Only used in conjunction with Set-Field QoS Index.
Set-Field	VLAN_Id	Only needed for L3 Multicast forwarding lookup, otherwise should be omitted.
Pop CW		Control word expected, pop without checking. Only used in conjunction with popping a bottom of stack pseudo wire label for MPLS-TP termination for data frames.
Pop VLAN		Pop outermost VLAN tag, only used in conjunction with popping an outermost Ethernet header.
Pop L2 Header		Pop outermost Ethernet header. Header cannot have a VLAN tag (must have already been popped). Only used for MPLS-TP termination.
Set-Field	LMEP_Id	Indicates MEP or MIP for OAM PDU processing.
OAM_LM_RX_Count	LMEP_Id, Traffic Class	Indicates MEP or MIP for which LM counters are to be incremented.

Name	Argument	Description
Check-Drop-Status	LMEP_Id, 0	Check that the path is not administratively locked.

Notes:

- For MPLS_TP PW termination actions must be listed in the following order: Pop Label, Pop CW, Pop VLAN (if any), Pop L2 Header.
- The parser must recognize the required OAM frame formats based on the next underlying label (GAL or RAL) or TTL. It must also recognize the payload G.8113.1 frame based on ACh, TTL, etc. These must be made available as match fields.
- Control word sequence number insertion, increment, or validation is not supported in this version of OF-DPA.

The MPLS Flow Table Write-Actions instruction can update the actions listed in Table 57 to the action set.

Table 57 MPLS Flow Table Action Set Actions

Name	Argument	Description
Group	Group id	Depends on rule type as detailed in Table 52 and Table 53.

4.1.9.4 Counters and Flow Expiry

The MPLS Flow Table counters are listed in Table 30.

Table 58 MPLS Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Number of active flow entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed
Received Packets	64	Per-entry	Number of packets that hit this flow entry.
Transmitted Packets	64	Per-entry	Number of bytes that hit this flow entry.

The MPLS Flow Table only supports hard flow entry expiration.

4.1.10 Maintenance Point Flow Table

The Maintenance Point Flow Table determines where to forward an OAM control frame for processing based on the LMEP Id and the MDL and opcode parsed from the Y.1731 PDU. Options are to send the packet to the CONTROLLER or LOCAL reserved port. The Controller can handle OAM message processing where latency is not critical.

The Maintenance Point Flow Table is used for G.8113.1 and Ethernet OAM.

4.1.10.1 Flow Entry Types and Match Fields

The Maintenance Point Flow Table implements the single flow entry type listed in Table 59.

Table 59 Maintenance Point Flow Table Entry Types

Name	Description
LMEP PDU	The only rule type.

The Maintenance Point Flow Table match fields are listed in Table 60.

Table 60 Maintenance Point Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
LMEP_Id	32	No	No	Local identifier for the MEP or MIP
OAM_Y1731_OPCODE	8	No	No	Parsed from the IEEE 802.1ag/Y.1731 header.
OAM_Y1731_MDL	3	No	No	Parsed from the IEEE 802.1ag/Y.1731 header.

The default on a miss is to drop (clear actions, no next table).

4.1.10.2 Instruction Types

The Maintenance Point Flow Table can have the instructions shown in Table 61. At least one must be specified. Since there is no next table the packet is dropped.

Table 61 Maintenance Point Flow Table Instructions

Name	Argument	Description
------	----------	-------------

Name	Argument	Description
Clear Actions		Used to drop and stop processing the OAM frame.
Apply Actions	Action List	Optional, Actions are listed in Table 62.
Goto Table		Optional, handles MIPs or when higher MDL OAM frames are to be treated as data, in which case it must be the MPLS L2 Port Flow Table.

4.1.10.3 Actions

The Maintenance Point Flow Table action list can include the actions shown in Table 62.

Table 62 Maintenance Point Flow Table Actions

Name	Argument	Description
Check-Drop-Status	LMEP_Id, 0	Check the Drop Status table for a Lock condition for this LMEP Id.
Output	Port	Can be CONTROLLER (sent to Controller using a Packet_In message) or LOCAL (for processing by the local OAM engine), or possibly both. For OAM frames that need to be processed at this Maintenance Point.
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Optional, handles MIPs or when higher MDL OAM frames are to be treated as data.
Set-Counter-Fields	LMEP Id, Traffic Class	Reads counters from the OAM Data Plane Counter Table and sets pipeline metadata from them.

4.1.10.4 Counters and Flow Expiry

The Maintenance Point Flow Table counters are listed in Table 63.

Table 63 Maintenance Point Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Number of active flow entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Maintenance Point Flow Table only supports hard expiration timers.

4.1.11 Policy ACL Flow Table

The Policy ACL Flow Table supports wide, multi-field matching. Most fields can be wildcard matched, and relative priority must be specified in all flow entry modification API calls. This is the preferred table for matching BPDU and ARP packets. It also provides the Metering instruction.

The Policy ACL Flow Table is organized as mutually exclusive logical sub-tables. Flow entries in the IPv6 logical tables match only packets that require matching on IPv6 header fields. The non-IPv6 logical table matches any packet that does not require matching on IPv6 header fields. Entries can optionally supply either a VLAN id or a Tunnel id match field, but not both. Following the OpenFlow single entry match semantics, since the Policy ACL Flow Table is considered a single table, a packet can match at most one rule in the entire table.

Note: If Ethertype is needed as a pre-requisite for other match fields it must be explicitly provided.

The Policy ACL Flow Table can optionally have a Goto-Table instruction specifying the Color Based Actions Flow Table. Otherwise it is the last table in the pipeline before performing the forwarding actions in the action set.

The default on table miss is to do nothing. The packet will be forwarded using the output or group in the action set, if any. If the action set does not have a group or output action the packet is dropped.

The Policy ACL Flow Table supports the flow entry types listed in Table 64.

Table 64 Policy ACL Flow Table Flow Entry Types

Type	Table	Prerequisite	Description
IPv4 VLAN	Table 65	Ethertype != 0x86dd, IN_PORT is a physical port	Matches packets by VLAN id except for IPv6. VLAN id is optional but must be non-zero if supplied. Tunnel id must not be supplied.
IPv6 VLAN	Table 66	Ethertype=0x86dd, IN_PORT is a physical port	Matches only IPv6 packets by VLAN id. VLAN id is optional but must be non-zero if supplied. Tunnel id must not be supplied.
IPv4 Data Center Overlay	Table 65	Ethertype != 0x86dd, IN_PORT is a tunnel logical port	Matches packets by tunnel id except for IPv6. A non-zero Tunnel id is required and is not maskable. VLAN id must not be supplied.
IPv6 Data Center Overlay	Table 66	Ethertype=0x86dd, IN_PORT is a tunnel logical port	Matches only IPv6 packets by tunnel id. A non-zero Tunnel id is required and is not maskable. VLAN id must not be supplied.

Type	Table	Prerequisite	Description
IPv4 MPLS-TP	Table 65	Ethertype != 0x86dd, IN_PORT is a tunnel logical port	Matches packets by tunnel id except for IPv6. A non-zero Tunnel id is required and is not maskable. VLAN id must not be supplied.
IPv6 MPLS-TP	Table 66	Ethertype=0x86dd, IN_PORT is a tunnel logical port	Matches only IPv6 packets by tunnel id. A non-zero Tunnel id is required and is not maskable. VLAN id must not be supplied.

4.1.11.1 Flow Entry Types and Match Fields

The available match fields for Policy ACL Flow Table flow entry types are as described in the following tables.

Table 65 Policy ACL Flow Table IPv4 Match Fields

Field	Bits	Maskable	Optional	Description or Prerequisite
IN_PORT	32	No	Yes	Physical or logical ingress port.
ETH_SRC	48	Yes	Yes	Ethernet source MAC
ETH_DST	48	Yes	Yes	Ethernet destination MAC
ETH_TYPE	16	No	Yes	Any value except 0x86dd. Explicit prerequisite must be 0x800 if IP fields are to be matched.
VLAN_VID	16	Yes	Yes	VLAN id. Cannot be masked for a VLAN bridging rule that redirects to a different L2 output group. Only applicable to VLAN flow entry types.
VLAN_PCP	3	No	Yes	802.1p priority field from VLAN tag. Always has a value, will be zero if packet did not have a VLAN tag.
VLAN_DEI	1	No	Yes	802.1p drop eligibility indicator field from VLAN tag. Always has a value, will be zero if packet did not have a VLAN tag.
TUNNEL ID	32	No	Yes	Tunnel forwarding domain. Applicable to data center overlay and MPLS-TP bridged flow entry types. IN_PORT must be a data center overlay tunnel or MPLS-TP logical port consistent with the tunnel-id range.

Field	Bits	Maskable	Optional	Description or Prerequisite
VRF	16	No	Yes	VRF.
IPV4_SRC	32	Yes	Yes	Matches SIP if Ethertype = 0x0800
ARP_SPA	32	Yes	Yes	Matches ARP source protocol address if Ethertype = 0x0806
IPV4_DST	32	Yes	Yes	Matches DIP if Ethertype = 0x0800
IP_PROTO	8	No	Yes	IP protocol field from IP header if Ethertype = 0x0800
IP_DSCP	6	No	Yes	Bits 0 through 5 of the IP ToS Field as defined in RFC 2474 if Ethertype = 0x0800
IP_ECN	2	No	Yes	Bits 6 through 7 of the IP ToS Field as defined in RFC 3168 if Ethertype = 0x0800
TCP_SRC	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 6
UDP_SRC	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 17
SCTP_SRC	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 132
ICMPV4_TYPE	8	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 1
TCP_DST	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 6
UDP_DST	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 17
SCTP_DST	16	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 132
ICMPv4_CODE	8	No	Yes	If Ethertype = 0x0800 and IP_PROTO = 1
MPLS L2 PORT	16	No	Yes	MPLS L2 Port

Table 66 Policy ACL Flow Table IPv6 Match Fields

Field	Bits	Maskable	Optional	Description
-------	------	----------	----------	-------------

Field	Bits	Maskable	Optional	Description
IN_PORT	32	No	Yes	Physical or logical ingress port.
ETH_SRC	48	Yes	Yes	Ethernet source MAC
ETH_DST	48	Yes	Yes	Ethernet destination MAC
ETH_TYPE	16	No	Yes	Must be 0x86dd
VLAN_VID	16	Yes	Yes	VLAN id. Cannot be masked for a VLAN bridging rule that redirects to a different L2 output group. Only applicable to VLAN flow entry types.
VLAN_PCP	3	No	Yes	802.1p priority field from VLAN tag. Always has a value, will be zero if packet did not have a VLAN tag.
VLAN_DEI	1	No	Yes	802.1p drop eligibility indicator field from VLAN tag. Always has a value, will be zero if packet did not have a VLAN tag.
TUNNEL ID	32	No	Yes	Tunnel forwarding domain. Applicable to data center overlay and MPLS-TP bridged flow entry types. IN_PORT must be a data center overlay tunnel or MPLS-TP logical port consistent with the tunnel-id range. VLAN_VID must not be supplied.
VRF	16	No	Yes	VRF
IPV6_SRC	128	Yes	Yes	Matches IPv6 SIP
IPV6_DST	128	Yes	Yes	Matches IPv6 DIP
IP_PROTO	8	No	Yes	Matches IPv6 Next header
IPV6_FLABEL	20	No	Yes	Matches IPv6 flow label
IP_DSCP	6	No	Yes	Bits 0 through 5 of the IP ToS Field as defined in RFC 2474 if Ethertype = 0x86dd
IP_ECN	2	No	Yes	Bits 6 through 7 of the IP ToS Field as defined in RFC 3168 if Ethertype = 0x86dd

Field	Bits	Maskable	Optional	Description
TCP_SRC	16	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 6
UDP_SRC	16	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 17
SCTP_SRC	16	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 132
ICMPV6_TYPE	8	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 58
TCP_DST	16	No	Yes	If Ethertype = 0x86dd 00 and IP_PROTO = 6
UDP_DST	16	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 17
SCTP_DST	16	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 132
ICMPv6_CODE	8	No	Yes	If Ethertype = 0x86dd and IP_PROTO = 58
MPLS L2 PORT	16	No	Yes	MPLS L2 Port

Notes:

- IPv6 Neighbor Discovery field matching is not supported in this version of OF-DPA.
- Not all IPv6 match fields are supported on all platforms.
- OF-DPA 2.0 permits bit masking L4 source and destination ports, and ICMP code. OpenFlow does not require these to be maskable.

4.1.11.2 Instruction Types

The Policy ACL Flow Table instructions are shown in Table 67.

Table 67 Policy ACL Flow Table Instructions

Name	Argument	Description
Meter	Meter identifier	Optional. Apply the meter indicated. Meter entry must exist prior to installing the flow.
Goto-Table	Color Based Actions	Optional. If not supplied pipeline processing is terminated and the action set is applied.
Apply Actions	Action list	Optional. Only the actions in Table 68 can be specified.

Name	Argument	Description
Clear Actions		Used to clear the action set for dropping the packet. Cannot be combined with write actions.
Write Actions	Action set	Only the actions in Table 69 or Table 70 can be specified, depending on rule type.

The packet is dropped if there is no group or output action since there is no next table.

Note: Apply-actions to CONTROLLER would be used if it is desired to output the packet to the CONTROLLER reserved port, rather than an output action in the write-actions action set.

Note: On some platforms, larger numbers of meters may be made accessible to flow entries that only match on: IN_PORT; VLAN_VID; MPLS_L2_PORT; or IN_PORT and VLAN_VID.

4.1.11.3 Action List Actions

The Policy ACL Flow Table action lists support the actions listed in Table 68.

Table 68 Policy ACL Flow Table Action List Actions

Name	Argument	Description
Set-Field	Color	New packet drop precedence. Optional. Overrides the color set by the meter.
Set-Field	Color Actions Index	Index into Color Based Actions Flow Table. Required for Color Based Actions.
Set-Field	Traffic Class	

4.1.11.4 Action Set Actions

The Policy ACL Flow Table action set supports the actions listed in Table 69 for VLAN match rule types, and the actions in Table 70 for tunnel match rule types.

Table 69 Policy ACL Flow Table VLAN Flow Entry Action Set

Name	Argument	Description
Group	Group	Sets output group entry for processing the packet after this table. Group must exist, be consistent with the type of rule

Name	Argument	Description
		and packet; and can be any of: L2 Interface, L2 Rewrite, L2 Multicast, L3 Unicast, L3 Multicast, or L3 ECMP; must respect VLAN id naming conventions. In particular, if the output is an L2 Rewrite group that does not set the VLAN id, the L2 Interface group it references must be consistent with the VLAN id in the matched flow entry.
Set-Queue	Queue-id	Determines queue to be used when packet is finally forwarded. Zero indicates the default queue. Cannot be used together with Set Traffic Class in the action list.

As with Unicast and Multicast Routing Flow Table actions, the decrement TTL and MTU checks are encoded by referencing an L3 Unicast or Multicast group entry. Note that if the group entry type is L2 Interface, L2 Rewrite, or L2 Multicast then these checks will not be done.

Table 70 Policy ACL Flow Table Tunnel Flow Entry Action Set

Name	Argument	Description
Group	Group	Sets output group entry for multicast forwarding or flooding. Group entry must exist, and must be one of OF-DPA L2 Overlay Multicast or L2 Overlay Flood sub-type with a tunnel id for the tenant forwarding domain.
Output	ifNum	Sets output port for unicast forwarding. Must be a tunnel logical port consistent with the rule forwarding domain.
Set-Queue	Queue-id	Determines queue to be used when packet is finally forwarded. Zero indicates the default queue. Cannot be used together with Set Traffic Class in action list.

4.1.11.5 Counters and Flow Expiration

The Policy ACL Flow Table counters are listed in Table 71. These are applicable to both VLAN and Tenant flow entries.

Table 71 Policy ACL Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table

Name	Bits	Type	Description
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed
Received Packets	64	Per-entry	Number of packets that hit this flow entry.
Received Bytes	64	Per-entry	Number of bytes that hit this flow entry.

Policy ACL Flow Table expiry provisions are shown in Table 72. Each flow entry can have its own timeout values.

Table 72 Policy ACL Flow Table Expiry

Name	Bits	Description
Hard Timeout	32	Number of seconds after which flow entry is removed. Optional, entry does not age out if zero or not specified.
Idle Timeout	32	Number of seconds of inactivity, after which a flow entry is removed. Optional, entry does not age out if zero or not specified.

4.1.12 Color Based Actions Flow Table

The Color Based Actions Flow Table provides packet editing actions based on the packet color. It is mainly used to set packet QoS fields such as DSCP or PCP.

Note: Actions in this table will override interface re-mark actions of the same type applied from group entries.

4.1.12.1 Flow Entry Types and Match Fields

The Maintenance Point Flow Table implements the single flow entry type listed in Table 73.

Table 73 Color Based Actions Flow Table Entry Types

Name	Description
Color Actions	The only rule type.

The Color Based Actions Flow Table match fields are listed in Table 74.

Table 74 Color Based Actions Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
Color	2	No	No	Packet color determined from prior stages
Color Actions Index	32	No	No	Set by Policy ACL Flow Table.

The default on a miss is to do nothing.

4.1.12.2 Instruction Types

The Color Based Actions Flow Table can have the instructions shown in Table 75. Since there is no next table, there are no Goto-Table or Write Metadata instructions.

Table 75 Color Based Actions Flow Table Instructions

Name	Argument	Description
Clear-Actions	none	Used to drop the packet.
Apply Actions	Action List	Allowed actions are listed in Table 62.

4.1.12.3 Actions

The Color Based Actions Flow Table can apply the actions listed in Table 76.

Table 76 Color Based Actions Flow Table Actions

Name	Argument	Description
Set-Field	Traffic Class	New packet traffic class. Optional
Set-Field	VLAN PCP	New outer VLAN priority marking. Optional.
Set-Field	VLAN DEI	Set or clear outer VLAN drop eligibility indicator. Optional.
Set-Field	IP_ECN	New ECN field marking. Applicable to IP packets. Optional.
Set-Field	IP_DSCP	New IP DSCP marking. Applicable to IP packets. Optional.

Name	Argument	Description
Set-Queue	Queue-id	Determines queue to be used when packet is finally forwarded. Cannot be used together with Set Traffic Class.
Output	Controller	Send a copy to the controller.

4.1.12.4 Counters and Flow Expiry

The Color Based Actions Flow Table counters are listed in Table 77.

Table 77 Color Based Actions Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Number of active flow entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

The Color Based Actions Flow Table only supports hard timer expiration for flow entries.

4.2 Egress Flow Tables

Egress flow tables permit matching in the context of the egress port after group entry processing. They are restricted from changing the port in the action set. However they can use the action list to send a copy to another port. They can also clear actions with no next table and cause the packet to be dropped.²⁵

OF-DPA 2.0 has three egress flow tables.

4.2.1 Egress VLAN Flow Table

The Egress VLAN Flow Table is used for VLAN translation and for OAM Maintenance Point processing.

4.2.1.1 Flow Entry Types and Match Fields

The Egress VLAN Flow Table supports the Flow Entry Types listed in Table 78.

Table 78 Egress VLAN Flow Table Flow Entry Types

Type	Description
------	-------------

²⁵ Egress Flow Tables are a planned feature of OpenFlow 1.5.

Type	Description
VLAN Translate Single Tag	Used to modify a single tagged packet. Can be used to remove the tag, change the VLAN Id, or to modify the VLAN id and push another tag.
VLAN Translate, Double Tag	Set the OVID metadata value to the outer VLAN Id and pop the outer tag. A Goto-Table instruction specifies the Egress VLAN 1 Flow table for further processing.
VLAN Translate Single Tag (Ethernet OAM MP)	Same actions as VLAN Translate Single Tag and also includes LM counter actions for an OAM UpMEP.
Ethernet OAM (Unicast)	Used for UpMEP receive OAM frame processing. Goto-Table instruction specifies the Egress Maintenance Point Flow table.
Ethernet OAM (Multicast)	Used for UpMEP receive OAM frame processing. Goto-Table instruction specifies the Egress Maintenance Point Flow table.

The Egress VLAN Flow Table match fields are listed in Table 10.

Table 79 Egress VLAN Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
ACTSET_OUTPUT	32	No	No	Egress physical port.
VLAN_VID	16	No	No	Outer VLAN id. Exact match.
ETH_TYPE	16	No	Yes	Required in rules that match OAM frames where the only allowed value is 0x8902; must be omitted in other rule types.
MAC-DST	48	Yes	Yes	Required in rules that match OAM frames; must be omitted in other rule types.

4.2.1.2 Instruction Types

The Egress VLAN Flow table supports the instruction types listed in Table 80.

Table 80 Egress VLAN Flow Table Instructions

Name	Argument	Description
Apply-Actions	Action List	The Egress VLAN Flow Table supports the actions specified in Table 81.
Clear-Actions		Used to clear the action set in preparation for dropping the packet.
Goto-Table	Table	One of: Egress VLAN 1 Flow Table, Egress Maintenance Point Flow Table, or none. If none the packet is output to the output port unless a Clear Actions instruction has been executed, in which case the packet is dropped.

4.2.1.3 Actions

The Egress VLAN Flow table uses Apply Actions for port VLAN tagging and assignment, as shown in Table 81.

Table 81 Egress VLAN Flow Table Action List

Name	Argument	Description
Set Field	VLAN_VID	
Set Field	OVID	Pipeline metadata field representing an outer tag VLAN Id that was popped, so that it can be used as a match field in the VLAN 1 Flow Table for double tag processing.
Push VLAN	TPID	Used in translating single to double tag. Value must be one of 0x8100 or 0x88a8.
Pop VLAN		Used in processing double tagged frames, where the Goto-Table instruction specifies the Egress VLAN 1 Flow table.
Set-Field	LMEP_Id	Indicates MEP or MIP for OAM PDU processing.
OAM_LM_RX_Count	LMEP_Id, Traffic Class	Indicates Up MEP or MIP for which LM counters are to be incremented.

4.2.1.4 Counters and Flow Expiry

The Egress VLAN Flow Table supports the table and flow entry counters listed in Table 82.

Table 82 Egress VLAN Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported.

4.2.2 Egress VLAN 1 Flow Table

The Egress VLAN 1 Flow Table is used for double tag VLAN translation and matching. As with the VLAN table, the OVID pipeline metadata field is used so that the Egress VLAN 1 Flow Table can match on two VLAN tags.

4.2.2.1 Flow Entry Types and Match Fields

The VLAN 1 Flow Table supports the Flow Entry Types listed in Table 83.

Table 83 Egress VLAN 1 Flow Table Flow Entry Types

Type	Description
VLAN Assignment	Exact match on IN_PORT, VLAN_VID, and OVID. Can optionally: pop the tag (packet becomes untagged); set the VLAN Id (single tag); or set the VLAN Id, push a tag, and set the VLAN Id for the pushed tag (double tagged). No next table.
VLAN Assignment (Ethernet OAM)	As above except increment LM counters for the LMEP Id and Traffic Class.
Ethernet OAM (Unicast)	Used for UpMEP receive OAM frame processing. Goto-Table instruction specifies the Egress Maintenance Point Flow table.
Ethernet OAM (Multicast)	Used for UpMEP receive OAM frame processing. Goto-Table instruction specifies the Egress Maintenance Point Flow table.

The Egress VLAN 1 Flow Table match fields are listed in Table 84.

Table 84 Egress VLAN 1 Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
ACTSET_OUTPUT	32	No	No	Egress port. Must be a physical port (high order 16 bits zero).
VLAN_VID	16	No	No	Inner VLAN id.
OVID	16	No	No	Outer VLAN id, set by a VLAN table flow entry.
ETH_TYPE	16	No	Yes	Only allowed value is 0x8902. In non-OAM rules must be omitted.
MAC-DST	48	Yes	Yes	Required in rules that match OAM frames, must be omitted in other rule types.

4.2.2.2 Instruction Types

The VLAN table supports the instruction types listed in Table 17.

Table 85 Egress VLAN 1 Flow Table Instructions

Name	Argument	Description
Apply-Actions	Action List	The Egress VLAN 1 Flow Table supports the actions specified in Table 86.
Clear-Actions		Used to drop the packet
Goto-Table	Table	For OAM frames must be the Maintenance Point Flow Table. Otherwise, there is no next table and the packet is forwarded to its output port.

4.2.2.3 Actions

The Egress VLAN 1 Flow Table action list is as shown in Table 86.

Table 86 Egress VLAN 1 Flow Table Action List

Name	Argument	Description
------	----------	-------------

Name	Argument	Description
Set Field	VLAN_VID	Used to set VLAN Id in inner tag or after pushing a new outer tag.
Pop VLAN		Used to remove inner tag from a double tagged frame leaving it untagged.
Push VLAN	TPID	Used push an outer tag so the final frame is double tagged. Value must be one of 0x8100 or 0x88a8.
Set-Field	LMEP_Id	Indicates MEP or MIP for OAM PDU processing.
OAM_LM_RX_Count	LMEP_Id, Traffic Class	Indicates Up MEP or MIP for which LM counters are to be incremented.

4.2.2.4 Counters and Flow Expiry

The Egress VLAN 1 Flow Table supports the table and flow entry counters listed in Table 87.

Table 87 Egress VLAN 1 Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported.

4.2.3 Egress Maintenance Point Flow Table

The Egress Maintenance Point Flow Table determines where to forward an OAM control frame for processing based on the LMEP Id and the opcode parsed from the Y.1731 PDU. Essentially duplicates the Maintenance Point Flow Table as an egress table.

The Egress Maintenance Point Flow Table is only used for Ethernet OAM.

4.2.3.1 Flow Entry Types and Match Fields

The Egress Maintenance Point Flow Table implements the single flow entry type listed in Table 88.

Table 88 Egress Maintenance Point Flow Table Entry Types

Name	Description
LMEP PDU	The only rule type.

The Egress Maintenance Point Flow Table match fields are listed in Table 89.

Table 89 Egress Maintenance Point Flow Table Match Fields

Field	Bits	Maskable	Optional	Description
LMEP_Id	32	No	No	Local identifier for the MEP or MIP
OAM_Y1731_OPCODE	16	No	No	Parsed from the IEEE 802.1ag/Y.1731 header.
OAM_Y1731_MDL	3	No	No	Parsed from the IEEE 802.1ag/Y.1731 header.

The built-in default on a miss is to drop (clear actions, no next table).

4.2.3.2 Instruction Types

The Egress Maintenance Point Flow Table can have the instructions shown in Table 90. At least one must be specified. Since there is no next table the packet is dropped.

Table 90 Egress Maintenance Point Flow Table Instructions

Name	Argument	Description
Write Actions	Action Set	
Clear Actions		Only used by built-in default rule.
Apply Actions	Action List	Actions are listed in Table 91.

4.2.3.3 Actions

Table 91 Egress Maintenance Point Flow Table Actions

Name	Argument	Description
------	----------	-------------

Name	Argument	Description
Output		Can be CONTROLLER (sent to Controller using a Packet_In message) or LOCAL (for processing by the local OAM engine), or possibly both.
OAM_LM_RX_Count	LMEP_Id, Traffic Class	Optional, handles MIPs or when higher MDL OAM frames are to be treated as data.
Set-Counter-Fields	LMEP_Id, Traffic Class	Reads counter values and uses them to set pipeline metadata to accompany the packet.

4.2.3.4 Counters and Flow Expiry

Egress Maintenance Point Flow Table supports the table and flow entry counters listed in Table 92.

Table 92 Egress Maintenance Point Flow Table Counters

Name	Bits	Type	Description
Active Entries	32	Table	Reference count of number of active entries in the table
Duration (sec)	32	Per-entry	Seconds since this flow entry was installed

Only hard interval timeout ageing per entry is supported.

4.3 Group Table

Most forwarding actions are embodied in group table entries. OF-DPA supports a defined set of group table entry types and enforces type checking consistency.

Each group entry has an identifier, type, counters, and one or more action buckets. OpenFlow has a single monolithic group table, but OF-DPA differentiates among types of group entries. For this purpose, OF-DPA encodes the group entry type in a group entry identifier field, effectively partitioning the group table identity name space to create logical sub-tables. The naming convention is shown in Table 93.

Table 93 OF-DPA Group Table Entry Identifier Naming Convention

Field	Bits	Description
Index	[27:0]	28 bit field, used to uniquely identify a group entry of the indicated type. May be used to further encode properties of the group entry,

Field	Bits	Description
		such as VLAN_VID.
Type	[31:28]	4 bit field that encodes the entry type, one of: 0: OF-DPA L2 Interface 1: OF-DPA L2 Rewrite 2: OF-DPA L3 Unicast 3: OF-DPA L2 Multicast 4: OF-DPA L2 Flood 5: OF-DPA L3 Interface 6: OF-DPA L3 Multicast 7: OF-DPA L3 ECMP 8: OF-DPA L2 Data Center Overlay 9: OF-DPA MPLS Label 10: OF-DPA MPLS Forwarding 11: OF-DPA L2 Unfiltered Interface

The OF-DPA 2.0 API validates the consistency checks on the group entry type when a group action is used in a flow or group entry action set. OF-DPA group entries must be defined before being used. OF-DPA 2.0 maintains reference counts for used entries, and an entry cannot be deleted if it is referenced by a flow entry or another group.

The index scheme varies by OF-DPA 2.0 group entry type and is described in the following sections.

4.3.1 OF-DPA L2 Interface Group Entries

L2 Interface Group entries are of OpenFlow indirect type, with a single action bucket.

OF-DPA L2 Interface group entries are used for egress VLAN filtering and tagging. The identifier convention is shown in Table 93. If a specific set of VLANs is allowed on a port, appropriate group entries must be defined for the VLAN and port combinations.

Note: OF-DPA uses the L2 Interface group declaration to configure the port VLAN filtering behavior. This approach was taken since OpenFlow does not support configuring VLANs on physical ports.

4.3.1.1 Naming Convention

Table 94 details the OF-DPA L2 Interface group entry identifier sub-fields that encode combinations of egress port and VLAN id.

Table 94 OF-DPA L2 Interface Group Entry Type Naming Convention

Field	Bits	Description
-------	------	-------------

Field	Bits	Description
Port identifier	[15:0]	Identifies a physical port (ifNum)
VLAN Id	[27:16]	VLAN id
Type	[31:28]	0 (L2 Interface)

4.3.1.2 Action Buckets

The single action bucket specifies the output port and whether or not the packet is egressed tagged. Although the pop action is a NOP if the packet has no VLAN tag, packets should always have a VLAN tag when the actions in the output group table are applied.

Note: If the packet came in untagged and a port VLAN was assigned, a VLAN tag was pushed as a VLAN Flow Table action.

Table 95 OF-DPA L2 Interface Group Entry Bucket Actions

Field	Argument	Description
Output	Port	Physical output port.
Pop VLAN	None	Pop the VLAN tag before sending the packet.
Set Field	DSCP	Static DSCP value for IP packets
Set Field	VLAN PCP	Static 802.1p value
Set-Field	VLAN DEI	Static 802.1p value

4.3.1.3 Counters

OF-DPA L2 Interface group entry counters are as shown in Table 96.

Table 96 OF-DPA L2 Interface Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entries currently referencing this group entry.

Name	Bits	Type	Description
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.2 OF-DPA L2 Unfiltered Interface Group Entries

L2 Unfiltered Interface Group entries are of OpenFlow indirect type, with a single action bucket. OF-DPA L2 Unfiltered Interface group entries are similar to L2 Interface group entries, but are used for forwarding to ports where egress VLAN filtering and tagging is not desired.

As with L2 Interface group entries, OF-DPA uses the L2 Unfiltered Interface group declaration to configure the port to not do VLAN filtering. Thus, a port cannot have both L2 Interface and L2 Unfiltered Interface groups defined for it.

4.3.2.1 Naming Convention

Table 97 details the OF-DPA L2 Unfiltered Interface group entry identifier sub-field encodings.

Table 97 OF-DPA L2 Unfiltered Interface Group Naming Convention

Field	Bits	Description
Port identifier	[15:0]	Identifies a physical port (ifNum)
Reserved	[27:16]	Must be zero
Type	[31:28]	11 (L2 Unfiltered Interface)

4.3.2.2 Action Buckets

The single action bucket, detailed in Table 98, specifies the output port. This is essentially the same as for the L2 Interface group action bucket but without the VLAN tag Pop action.

Table 98 OF-DPA L2 Unfiltered Interface Group Bucket Actions

Field	Argument	Description
Output	Port	Physical output port.
Set Field	DSCP	Static DSCP value for IP packets
Set Field	VLAN PCP	Static 802.1p value, only applies if packet has an outer VLAN tag.

Field	Argument	Description
Set-Field	VLAN DEI	Static 802.1p value, only applies if packet has an outer VLAN tag.

4.3.2.3 Counters

OF-DPA L2 Interface group entry counters are as shown in Table 99.

Table 99 OF-DPA L2 Unfiltered Interface Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entries currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.3 OF-DPA L2 Rewrite Group Entries

OF-DPA L2 Rewrite group entries are of indirect type and have a single action bucket. They are used when it is desired to modify Ethernet header fields for bridged packets. Use of an OF-DPA L2 Rewrite group entry is optional and only the Policy ACL Flow Table has the ability to use it in its Write-Actions instruction.

OF-DPA L2 Rewrite actions are optional with the exception of group. This permits an OF-DPA L2 Rewrite group entry to selectively modify the source MAC, destination MAC, and/or VLAN id.

If a Set Field action sets the VLAN id, the VLAN id must be the same as in a chained L2 Interface group entry. Note that if the VLAN id is not rewritten, the VLAN id in the L2 Interface group entry must be the same as the VLAN id matched in the Policy ACL Flow Table flow entry that forwarded to the rewrite group.

4.3.3.1 Naming Convention

Table 100 details the OF-DPA L2 Rewrite group entry identifier sub-fields that encode the type and VLAN id.

Table 100 OF-DPA L2 Rewrite Group Entry Type Naming Convention

Field	Bits	Description
Id	[27:0]	Index to differentiate group entries of this type

Field	Bits	Description
Type	[31:28]	1 (OF-DPA L2 Rewrite)

4.3.3.2 Action Buckets

The single action bucket references the output group for forwarding the packet and optional Ethernet header modifications.

Table 101 OF-DPA L2 Rewrite Group Entry Bucket Actions

Field	Argument	Description
Group	Group entry	Must chain to a L2 Interface group entry. Required.
Set Field	MAC_SRC	Re-write the source MAC. Optional.
Set Field	MAC_DST	Re-write the destination MAC. Optional.
Set Field	VLAN-id	Re-write the VLAN id. Optional.

Chained group entries must be defined before being used. OF-DPA maintains reference counts for used entries, and a group entry cannot be deleted if it is referenced by a flow entry or another group.

4.3.3.3 Counters

OF-DPA L2 Rewrite group entry counters are as shown in Table 102 for completeness.

Table 102 OF-DPA L2 Rewrite Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entries currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.4 OF-DPA L3 Unicast Group Entries

OF-DPA L3 Unicast group entries are used to supply the routing next hop and output interface for packet forwarding. To properly route a packet from either the Routing Flow Table or the Policy ACL Flow Table, the forwarding flow entry must reference an OF-DPA L3 Unicast Group entry.

OF-DPA L3 Unicast automatically includes the ALLOW-IN_PORT vendor extension property to allow packets to be sent out IN_PORT. This property overrides the OpenFlow default behavior, which is to not forward a packet to IN_PORT, and is inherited by chained group entries. It is not visible to the controller and hence cannot be modified or read.

All packets must have a VLAN tag. A chained L2 Interface group entry must be in the same VLAN as assigned by the OF-DPA L3 Unicast Group entry.

4.3.4.1 Naming Convention

The naming convention for OF-DPA L3 Unicast Group entries is shown in Table 103.

Table 103 OF-DPA L3 Unicast Group Entry Naming Conventioin

Field	Bits	Description
Id	[27:0]	Index to differentiate group entries of this type
Type	[31:28]	2 (OF-DPA L3 Unicast)

4.3.4.2 Action Buckets

The single action bucket is as shown in Table 104.

Table 104 OF-DPA L3 Unicast Bucket Actions

Field	Argument	Description
Group	Group-id	Must chain to a L2 Interface group entry. ALLOW-IN_PORT permits the chained group entry output action to include the packet IN_PORT. Required.
Set Field	MAC_DST	Write the next hop destination MAC. Required.
Set Field	MAC_SRC	Write the source MAC corresponding to the L3 output interface. Required.
Set Field	VLAN-id	Write the VLAN id corresponding to the L3 output interface. Required.

4.3.4.3 Counters

The OF-DPA L3 Unicast group entry counters are as shown in Table 105.

Table 105 OF-DPA L3 Unicast Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.5 OF-DPA L2 Multicast Group Entries

OF-DPA L2 multicast group entries are of OpenFlow ALL type. There can be multiple action buckets, each referencing an output port by chaining to an OF-DPA L2 Interface Group entry.

Note: By OpenFlow default, a packet cannot be forwarded back to the IN_PORT from which it came in. An action bucket that specifies the particular packet's ingress port is not evaluated.

All of the OF-DPA L2 Interface Group entries referenced by the OF-DPA Multicast Group entry, and the OF-DPA Multicast Group entry itself, must be in the same VLAN.

4.3.5.1 Naming Convention

OF-DPA L2 Multicast group entries use the naming convention in Table 106.

Table 106 OF-DPA L2 Multicast Group Entry Type Naming Convention

Field	Bits	Description
Id	[15:0]	Index to differentiate group entries of this type
VLAN Id	[27:16]	VLAN id
Type	[31:28]	3 (L2 Multicast)

4.3.5.2 Action Buckets

The contents of OF-DPA L2 Multicast Group entry buckets can contain only the value shown in Table 107.

Table 107 OF-DPA L2 Multicast Bucket Actions

Field	Argument	Description
-------	----------	-------------

Field	Argument	Description
Group	Group-id	Must chain to a L2 Interface group entry whose VLAN id name component matches the VLAN id component of this group entry's name.

4.3.5.3 Counters

The VL2 Multicast group entry counters are as shown in Table 108.

Table 108 OF-DPA L2 Multicast Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.6 OF-DPA L2 Flood Group Entries

The OF-DPA L2 Flood Group entries are used by VLAN Flow Table wildcard (destination location forwarding, or DLF) rules. Like OF-DPA L2 Multicast group entry types they are of OpenFlow ALL type. The action buckets each encode an output port. Each OF-DPA L2 Flood Group entry bucket forwards a replica to an output port, except for packet IN_PORT.

The main difference from OF-DPA L2 Multicast Group entries is how they are processed in the hardware.

All of the OF-DPA L2 Interface Group entries referenced by the OF-DPA Flood Group entry, and the OF-DPA Flood Group entry itself, must be in the same VLAN.

Note: There can only be one OF-DPA L2 Flood Group entry defined per VLAN.

4.3.6.1 Naming Convention

OF-DPA L2 Flood group entries follow the naming convention shown in Table 109.

Table 109 OF-DPA L2 Flood Group Entry Naming Convention

Field	Bits	Description
Id	[15:0]	Index to differentiate group entries of this type

Field	Bits	Description
VLAN Id	[27:16]	VLAN id
Type	[31:28]	4 (OF-DPA L2 Flood)

4.3.6.2 Action Buckets

The contents of the OF-DPA L2 Flood Group Entry action buckets can contain only the values shown in Table 110.

Table 110 OF-DPA L2 Flood Bucket Actions

Field	Argument	Description
Group	Group-id	Must chain to L2 Interface group entry whose VLAN id name component is the same as the VLAN id in this entry's name.

4.3.6.3 Counters

The OF-DPA L2 Multicast group entry counters are as shown in Table 111.

Table 111 OF-DPA L2 Flood Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.7 OF-DPA L3 Interface Group Entries

OF-DPA L3 interface group entries are of indirect type and have a single action bucket. They are used to supply outgoing routing interface properties for multicast forwarding. For unicast forwarding, use of OF-DPA L3 Unicast group entries is recommended.

OF-DPA L3 Interface uses the ALLOW-IN_PORT vendor extension that permits packets to be sent out IN_PORT.

The VLAN id in the name must be the same as the VLAN_VID assigned in the Set Field action and the VLAN id in the name of the chained OF-DPA L2 Interface group.

4.3.7.1 Naming Convention

Table 112 details the OF-DPA L3 Interface group entry identifier sub-fields.

Table 112 OF-DPA L3 Interface Group Entry Type Naming Convention

Field	Bits	Description
Id	[27:0]	Index to differentiate group entries of this type
Type	[31:28]	5 (OF-DPA L3 Interface)

4.3.7.2 Action Buckets

The single action bucket specifies the MAC_SRC, VLAN_VID, TTL decrement action, and an output group for forwarding the packet. All actions are required.

Table 113 OF-DPA L3 Interface Group Entry Bucket Actions

Field	Argument	Description
Group	Group entry	Must chain to a L2 Interface group entry. This group entry can output the packet to IN_PORT. The VLAN id component of the chained group entry's name must match the Set Field value for VLAN id.
Set Field	MAC_SRC	Write the source MAC corresponding to the L3 output interface.
Set Field	VLAN-id	Write the VLAN id corresponding to the L3 output interface.

Referenced group entries must be defined before being used. OF-DPA maintains reference counts for used entries, and an entry cannot be deleted if it is referenced by a flow entry or another group.

4.3.7.3 Counters

OF-DPA L3 Interface group entry counters are as shown in Table 102 for completeness.

Table 114 OF-DPA L3 Interface Group Entry Counters

Name	Bits	Type	Description
------	------	------	-------------

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.8 OF-DPA L3 Multicast Group Entries

OF-DPA L3 Multicast group entries are of OpenFlow ALL type. The action buckets describe the interfaces to which multicast packet replicas are forwarded.

IP multicast packets are forwarded differently depending on whether they are switched or routed. Packets must be switched in the VLAN in which they came in, and cannot be output to IN_PORT. Packets that are multicast in other VLANs or MPLS L3 VPNs must be routed and must be allowed to egress via IN_PORT. This difference is reflected in the actions that are programmed in the action buckets.

Note that any chained OF-DPA L2 Interface Group entries must be in the same VLAN as the OF-DPA L3 Multicast group entry. However chained OF-DPA L3 Interface Group entries must be in different VLANs from the OF-DPA L3 Multicast Group entry, and from each other.

4.3.8.1 Naming Convention

The naming convention for OF-DPA L3 Multicast Group entries is shown in Table 115.

Table 115 OF-DPA L3 Multicast Group Entry Naming Convention

Field	Bits	Description
Index	[15:0]	Used to differentiate between OF-DPA L3 multicast group entries.
VLAN Id	[27:16]	VLAN id
Type	[31:28]	6 (OF-DPA L3 Multicast)

4.3.8.2 Action Buckets

The action buckets contain the values shown in Table 116.

Table 116 OF-DPA L3 Multicast Bucket Actions

Field	Argument	Description
-------	----------	-------------

Field	Argument	Description
Group	Group-id	Can chain to one of: L3 Interface; L2 Interface; L3 Fast Failover; MPLS L3 VPN Label group entry types. Chained group entry names must conform to the VLAN id requirements above.

4.3.8.3 Counters

The OF-DPA L3 Multicast group entry counters are as shown in Table 117.

Table 117 OF-DPA L3 Multicast Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.9 OF-DPA L3 ECMP Group Entries

OF-DPA L3 ECMP group entries are of OpenFlow type SELECT. For IP routing the action buckets reference the OF-DPA L3 Unicast group entries that are members of the multipath group for ECMP forwarding.

An OF-DPA L3 ECMP Group entry can also be used in a Provider Edge Router. In this packet flow it can chain to either an MPLS L3 Label group entry or to an MPLS Fast Failover group entry.

An OF-DPA L3 ECMP Group entry can be specified as a routing target instead of an OF-DPA L3 Unicast Group entry. Selection of an action bucket for forwarding a particular packet is hardware specific.

4.3.9.1 Naming Convention

The naming convention for OF-DPA L3 ECMP Group entries is as shown in Table 118.

Table 118 OF-DPA L3 ECMP Group Entry Naming Convention

Field	Bits	Description
Id	[27:0]	Used to differentiate OF-DPA L3 ECMP group entries.

Field	Bits	Description
Type	[31:28]	7 (OF-DPA L3 ECMP)

4.3.9.2 Action Buckets

The action buckets contain the single value listed in Table 119.

Table 119 OF-DPA L3 ECMP Group Entry Bucket Actions

Field	Argument	Description
Group	Group-id	May chain to an OF-DPA L3 Unicast, MPLS Fast Failover, or MPLS L3 VPN Label group entry.

4.3.9.3 Counters

The OF-DPA L3 ECMP group entry counters are as shown in Table 120.

Table 120 OF-DPA L3 ECMP Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.10 OF-DPA L2 Overlay Group Entries

OF-DPA L2 Overlay Group Entries are of OpenFlow all type. The action buckets describe the tenant access logical ports and/or tunnel endpoint logical ports to which packets are to be replicated by this group. Note that all tenant logical ports must be for the same tenant as the tunnel id in the group name.

Tenant access and tunnel endpoint logical port configuration is described in Section 5.1.3.

4.3.10.1 OF-DPA L2 Overlay Group Sub-Types

There are four OF-DPA L2 Overlay Group sub-types. These can be considered OF-DPA group entries in their own right, but are described together here since they perform similar functions. The differences relate to usage (whether in DLF or multicast flows) and to the underlay remote tunnel endpoint type (whether unicast or multicast). Note that regardless of whether forwarded (overlay) packets are

themselves unicast or multicast, they will be replicated using the underlay tunnel type corresponding to the OF-DPA L2 Overlay Group sub-type name component.

Figure 35 shows an OF-DPA L2 Overlay Flood Over Unicast Tunnels group entry. Buckets can specify multiple access and/or tunnel logical ports. OF-DPA will use unicast underlay tunnels to forward packets for the specified logical ports. OF-DPA L2 Overlay Flood Over Unicast Tunnels group entries can only be referenced by tunnel DLF rule types.

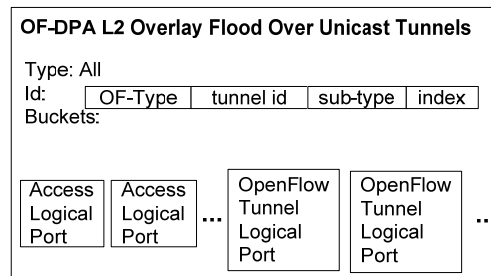


Figure 35 OF-DPA L2 Overlay Flood Over Unicast Tunnels

Figure 36 illustrates an OF-DPA L2 Overlay Flood Over Multicast Tunnels group entry. There can be at most one bucket specifying a tunnel logical port. OF-DPA will forward packets over the tenant multicast underlay tunnel configured on the tunnel logical port. A multicast IP group address must have been configured for the tenant on that logical port.

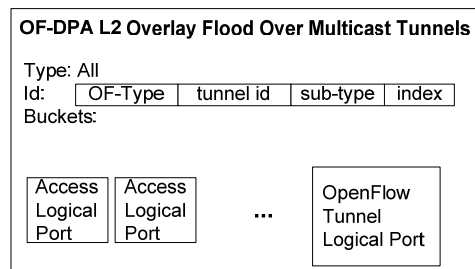


Figure 36 OF-DPA L2 Overlay Flood Over Multicast Tunnels

OF-DPA L2 Overlay Flood Over Multicast Tunnels group entries can only be referenced by a tunnel DLF rule.

Figure 37 shows an OF-DPA L2 Overlay Multicast Over Unicast Tunnels group entry. Multiple tunnel logical port buckets can be specified. OF-DPA will use unicast underlay tunnels to forward packets for the specified logical ports. OF-DPA L2 Overlay Multicast Over Unicast Tunnels group entries cannot be referenced by tunnel DLF rule types.

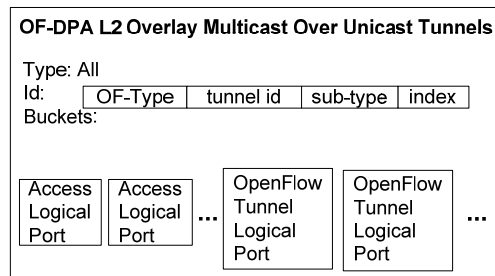


Figure 37 OF-DPA L2 Overlay Multicast Over Unicast Tunnels

Figure 38 illustrates the OF-DPA L2 Overlay Multicast Over Multicast Tunnels group entry. There can be at most one bucket specifying a tunnel logical port configured with a multicast IP group address for the tenant. OF-DPA will use unicast underlay tunnels to forward packets for the specified logical ports. OF-DPA L2 Overlay Multicast Over Unicast Tunnels group entries cannot be referenced by tunnel DLF rules.

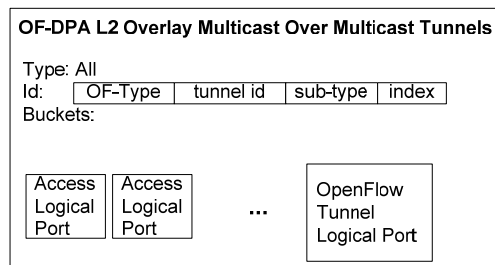


Figure 38 OF-DPA L2 Overlay Multicast Over Multicast Tunnels

4.3.10.2 Naming Convention

The naming convention for OF-DPA L2 Overlay group entries is shown in Table 121.

Table 121 OF-DPA L2 Overlay Group Entry Naming Convention

Field	Bits	Description
Index	[9:0]	Used to differentiate L2 Overlay group entries of the same sub-type.
Sub-Type	[11:10]	Identifies the type of forwarding and undelay tunnel used: 0: OF-DPA L2 Overlay Flood Over Unicast Tunnels 1: OF-DPA L2 Overlay Flood Over Multicast Tunnels 2: OF-DPA L2 Overlay Multicast Over Unicast Tunnels 3: OF-DPA L2 Overlay Multicast Over Multicast Tunnels
Tunnel Id	[27:12]	Low order 16 bits of the tenant forwarding domain identifier. Must uniquely identify the tenant.

Field	Bits	Description
Type	[31:28]	8 (OF-DPA L2 Overlay)

4.3.10.3 Action Buckets

The action buckets for all OF-DPA L2 Overlay Group Entry sub-types contain the values shown in Table 122.

Table 122 OF-DPA L2 Overlay Group Sub-Type Entry Bucket Actions

Field	Argument	Description
Output	Logical port	Must be a logical port in the tenant forwarding domain. Can be either an access or tunnel logical port.

4.3.10.4 Counters

The OF-DPA L2 Overlay Flood group entry counters are as shown in Table 123. These counters are individually maintained by sub-type.

Table 123 OF-DPA L2 Overlay Group Sub-Type Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.11 OF-DPA MPLS Interface Group Entry

An OF-DPA Interface Group Entry is of OpenFlow type INDIRECT. It is used to set the outgoing L2 header to reach the next hop label switch router or provider edge router.

4.3.11.1 Naming Convention

The naming convention for OF-DPA MPLS group entries is shown in Table 124.

Table 124 OF-DPA MPLS Interface Group Entry Naming Convention

Field	Bits	Description
-------	------	-------------

Field	Bits	Description
Index	[23:0]	Used to differentiate MPLS group entries of the same sub-type
Sub-Type	[27:24]	Identifies the type of the entry: 0: OF-DPA MPLS Interface
Type	[31:28]	9 (OF-DPA MPLS Label)

4.3.11.2 Action Buckets

The action bucket for the OF-DPA MPLS Interface Group entry type contains the actions listed in Table 125.

Table 125 OF-DPA MPLS Interface Group Entry Bucket Actions

Field	Argument	Description
Set-Field	MAC-DST	Destination MAC address. Required.
Set-Field	MAC-SRC	Source MAC address. Required.
Set-Field	VLAN_VID	VLAN. Required.
OAM_LM_TX_Count	LMEP_Id	Used for Section OAM Loss Measurement. Optional.
Group	Group entry	May chain to either an OF-DPA L2 Interface or an L2 Unfiltered Interface group entry type. Required.

4.3.11.3 Counters

The OF-DPA MPLS Interface Group entry type counters are as shown in Table 126.

Table 126 OF-DPA MPLS Interface Group Type Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.12 OF-DPA MPLS Label Group Entries

OF-DPA Label Group entries are of OpenFlow INDIRECT type. There are four MPLS label Group entry sub-types, all with similar structure. These can be used in different configurations to push up to three labels for tunnel initiation or LSR swap.

4.3.12.1 Naming Convention

The naming convention for OF-DPA MPLS Label group entries is shown in Table 121. Different sub-types are used for different action sets.

Table 127 OF-DPA MPLS Label Group Entry Naming Convention

Field	Bits	Description
Index	[23:0]	Used to differentiate MPLS group entries of the same sub-type
Sub-Type	[27:24]	Identifies the type of the entry: 1: OF-DPA MPLS L2 VPN Label 2: OPF-DPA MPLS L3 VPN Label 3: OF-DPA MPLS Tunnel Label 1 4: OF-DPA MPLS Tunnel Label 2 5: OF-DPA MPLS Swap Label
Type	[31:28]	9 (OF-DPA MPLS Label)

The contents of the action buckets vary according to the type of label.

4.3.12.2 MPLS VPN Label Action Buckets

OF-DPA MPLS VPN Label groups are used for VPWS or L3 VPN initiation.

The action bucket for OF-DPA MPLS VPN Label sub-type can contain the actions shown in Table 128.

Table 128 OF-DPA MPLS L2 VPN Label Group Bucket Actions

Field	Argument	Description
Push L2 Header		Pushes an outermost Ethernet header on the packet. All fields are initialized to zero. If the outermost header has a VLAN tag it must be pushed using a Push VLAN action. Used only in MPLS L2 VPN group entries for MPLS-TP VPLS and VPWS pseudo-wire initiation.

Field	Argument	Description
Push VLAN	Ethertype = 0x8100	Used only in MPLS L2 VPN group entries to push a single VLAN tag on a newly pushed Ethernet header.
Push MPLS Header	Ethertype = 0x8847	Used in all MPLS Label group entry types to push a new MPLS shim header on the frame immediately after the Ethernet header, including any VLAN tags. The new header becomes the outermost MPLS label.
Push CW		Used only in MPLS L2 VPN group entries to push a 32-bit Control Word after the PW label. The Control Word is initialized to all zeros and is pushed between the MPLS shim label and the packet payload. This action can only be executed once and only after a Push MPLS Header for the bottom of stack (PW) label.
Set-Field	MPLS Label	20-bit value, set on the outermost MPLS header label field.
Set-Field	BOS	1-bit outermost MPLS header bottom of stack field.
Set-Field	TC	3-bit value for outermost MPLS header EXP field. Cannot be used in the same action set as Set TC from table action. Overrides any value set by an MPLS Flow Table action.
Set TC From Table	QoS Index	Index into the MPLS VPN Label Remark Action to lookup EXP value based on packet Traffic Class and Color pipeline match field values.
Set-Field	TTL	Eight-bit value for outermost MPLS TTL field. Cannot be used in the same action set as a Copy TTL outwards action. Overrides any value set by an MPLS Flow Table action.
Copy TTL outwards		Used only in MPLS L3 VPN Label group entries to copy the TTL from the IP header.
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Used to increment the OAM pseudo-wire loss management counter for the label at an MPLS-TP PW or LSP Down MEP. Optional.
Set PRI From Table	QoS Index	Sets the 802.1p priority field in the L2 header. Index into the MPLS VPN Label Remark Action table to lookup VLAN_PRI value based on packet Traffic Class and Color

Field	Argument	Description
		values. Not used if final packet does not have a VLAN tag.
Group	Group entry	Possible values: OF-DPA MPLS Interface; OF-DPA MPLS Tunnel Label 1 group entry

4.3.12.3 MPLS Tunnel Label 1 Action Buckets

The action bucket for OF-DPA MPLS Tunnel Label 1 sub-type can contain the actions shown in Table 129.

Table 129 OF-DPA MPLS Tunnel Label 1 Group Bucket Actions

Field	Argument	Description
Push MPLS Header	Ethertype = 0x8847	Push a new MPLS shim header on the frame immediately after the Ethernet header, including any VLAN tags. The new header becomes the outermost MPLS label.
Set-Field	MPLS Label	20-bit value for outermost MPLS label field.
Set-Field	TC	3-bit value for outermost MPLS header EXP field. Cannot be used in the same action set as Set TC from table action. Overrides any value set by an MPLS Flow Table action.
Copy TC outwards		Copy EXP field from the previously outermost shim header which is now immediately under this tunnel label.
Set TC From Table	QoS Index	Index into the MPLS Tunnel Label Remark Action to lookup EXP value based on packet traffic class and color values.
Set PRI From Table	QoS Index	Sets the 802.1p priority field if outermost label. Index into the MPLS Tunnel Label Remark Action to lookup VLAN_PRI value based on packet Traffic Class and Color values. Not used if final packet does not have a VLAN tag.
Set-Field	TTL	Eight-bit value for outermost MPLS TTL field. Overrides a copy out action if both are supplied.
Copy TTL outwards		Copy TTL from next MPLS shim or IP header.

Field	Argument	Description
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Used to increment the OAM LSP loss management counter for the label at an MPLS-TP LSP Down MEP. Optional.
Group	Group entry	Possible values: OF-DPA MPLS Interface; OF-DPA MPLS Tunnel Label 2 group entry

4.3.12.4 MPLS Tunnel Label 2 Action Buckets

The action bucket for OF-DPA MPLS Tunnel Label 2 sub-type can contain the values shown in Table 130.

Table 130 OF-DPA MPLS Tunnel Label 2 Actions

Field	Argument	Description
Push MPLS Header	Ethertype = 0x8847	Push a new MPLS shim header on the frame immediately after the Ethernet header, including any VLAN tags. The new header becomes the outermost MPLS label.
Set-Field	MPLS Label	20-bit value for outermost MPLS label field.
Set-Field	TC	3-bit value for outermost MPLS header EXP field. Cannot be used in the same action set as Set TC from table action. Overrides any value set by an MPLS Flow Table action.
Copy TC outwards		Copy EXP field from the previously outermost shim header which is now immediately under this tunnel label.
Set TC from table	QoS Index	Index into the MPLS Tunnel Label Remark Action table to lookup EXP value based on packet traffic class and color values.
Set PRI From Table	QoS Index	Sets the 802.1p priority field if outermost label. Index into the MPLS Tunnel Label Remark Action to lookup VLAN_PRI value based on packet Traffic Class and Color values. Not used if final packet does not have a VLAN tag.
Set-Field	TTL	Eight-bit value for outermost MPLS TTL field.
Copy TTL outwards		Copy TTL from next MPLS shim.

Field	Argument	Description
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Used to increment the OAM LSP loss management counter for the label at an MPLS-TP LSP Down MEP. Optional.
Group	Group entry	Only possible value is an OF-DPA MPLS Interface group entry.

4.3.12.5 MPLS Swap Label Action Buckets

The action bucket for the OF-DPA MPLS Swap Label sub-type can contain the values shown in Table 131.

Table 131 OF-DPA MPLS Swap Label Actions

Field	Argument	Description
Set-Field	MPLS Label	20-bit value to replace that in the outermost MPLS label field.
Set-Field	TC	3-bit value for outermost MPLS header EXP field. Cannot be used in the same action set as Set TC from table action. Overrides any value set by an MPLS Flow Table action.
Copy TC outwards		Copy EXP field from next shim. If next header is not an MPLS shim, operation is a NOP and TC is unchanged.
Set TC From Table	QoS Index	Index into the MPLS VPN Label Remark Action table to lookup EXP value based on packet traffic class and color values.
Set-Field	TTL	Eight-bit value for outermost MPLS TTL field.
Copy TTL outwards		Copy TTL from next MPLS shim or IP header. If next header is neither MPLS nor IP, operation is a NOP.
Set PRI From Table	QoS Index	Sets the 802.1p priority field if outermost label. Index into the MPLS VPN Label Remark Action to lookup VLAN_PRI value based on packet Traffic Class and Color values. Not used if final packet does not have a VLAN tag.
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Used to increment the OAM LSP loss management counter for the label at an MPLS-TP LSP Down MEP.

Field	Argument	Description
		Optional.
Group	Group entry	Possible values: OF-DPA MPLS Interface; OF-DPA MPLS Tunnel Label 1 group entry.

4.3.12.6 Counters

The counters supported for all OF-DPA MPLS Label sub-type group entries are as shown in Table 132.

Table 132 OF-DPA MPLS Label Group Sub-Type Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.13 OF-DPA MPLS Fast Failover Group Entry

OF-DPA MPLS Fast Failover Group Entries are of OpenFlow FAST FAILOVER type. The action buckets chain to an MPLS label group. Fast Failover group entry types have two buckets, for working and protection paths. By convention the bucket selection is ordered, with the first bucket always being the working bucket. The working bucket is always selected if its liveness condition is satisfied (i.e., the path working is functioning). If the liveness condition is not valid, then the group entry fails over to next bucket in order, until one is found that is “live.”

Note: The bucket liveness “watch_port” should be programmed with an OAM Protection Liveness Logical Port, and the “watch_group” should be programmed to ANY.

4.3.13.1 Naming Convention

The naming convention for OF-DPA MPLS group entries is shown in Table 133.

Table 133 OF-DPA MPLS Fast Failover Group Entry Naming Convention

Field	Bits	Description
Index	[23:0]	Used to differentiate MPLS group entries of the same sub-type
Sub-Type	[27:24]	Identifies the type of the entry:

Field	Bits	Description
		6: OF-DPA MPLS Fast Failover
Type	[31:28]	10: OF-DPA MPLS Forwarding

4.3.13.2 Action Buckets

The actions for the working and protection path buckets are shown in Table 134. Both buckets should be programmed with the same types of groups. If this convention is not followed, results may be unpredictable.

MPLS-TP L2 VPN flows should use MPLS L2 VPN Label groups. Likewise, MPLS L3 VPN flows should use MPLS L3 VPN Label groups.

Table 134 OF-DPA MPLS Fast Failover Group Entry Bucket Actions

Field	Argument	Description
Group	Group entry	Depending on the packet flow can be one of: MPLS L2 VPN Label; MPLS L3 Label; MPLS Swap Label; or MPLS Tunnel Label 1.
Watch_Port	OAM Protection Liveness Logical Port	Determines liveness of bucket. The bucket is live if the logical port is operational. The first live bucket is used.

4.3.13.3 Counters

The counters for the OF-DPA MPLS Fast Failover Group entry are as shown in Table 150.

Table 135 OF-DPA MPLS Fast Failover Tag Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.14 OF-DPA MPLS ECMP Group Entry

OF-DPA MPLS ECMP Group Entries are of OpenFlow SELECT type and used for MPLS ECMP multipath forwarding.

All buckets must reference the same type of OF-DPA MPLS Group. If this convention is not followed results may be unpredictable.

4.3.14.1 Naming Convention

The naming convention for OF-DPA MPLS ECMP group entries is shown in Table 136.

Table 136 OF-DPA MPLS ECMP Group Entry Naming Convention

Field	Bits	Description
Index	[23:0]	Used to differentiate MPLS group entries of the same sub-type
Sub-Type	[23:24]	Identifies the type of the entry: 8: OF-DPA MPLS ECMP
Type	[31:28]	10: OF-DPA MPLS Forwarding

4.3.14.2 Action Buckets

OF-DPA MPLS ECMP Group type entry buckets each have a single group as shown in Table 137.

Table 137 OF-DPA MPLS ECMP Group Entry Bucket Actions

Field	Argument	Description
Group	Group entry	Depending on the packet flow can be one of: MPLS Fast Failover, MPLS L2 VPN Label; MPLS L3 Label; MPLS Swap Label; MPLS Tunnel Label 1, or MPLS Interface.

Note: The maximum number of buckets available in any OF-DPA MPLS ECMP Group entry is platform dependent.

4.3.14.3 Counters

The OF-DPA MPLS ECMP Group entry type counters are as shown in Table 138.

Table 138 OF-DPA MPLS ECMP Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.3.15 OF-DPA MPLS L2 Tag Group Entry

OF-DPA MPLS L2 Tag Group Entries are of OpenFlow Indirect type. They are optionally used to perform service delimiting (SD) tag actions on tunnel termination required by RFC4448 [30]. The action bucket chains to an L2 Interface or L2 Unfiltered Interface Group entry for forwarding the packet to a local attachment circuit interface.

4.3.15.1 Naming Convention

The naming convention for OF-DPA MPLS L2 Tag group entries is shown in Table 139.

Table 139 OF-DPA MPLS L2 Tag Group Entry Naming Convention

Field	Bits	Description
Index	[23:0]	Used to differentiate MPLS group entries of the same sub-type
Sub-Type	[27:24]	Identifies the type of the entry: 10: OF-DPA MPLS L2 Tag
Type	[31:28]	10: OF-DPA MPLS Forwarding

4.3.15.2 Action Buckets

The action bucket actions for the OF-DPA MPLS L2 Tag Group type are shown in Table 140.

Table 140 OF-DPA L2 Tag Group Entry Bucket Actions

Field	Argument	Description
Push VLAN		Pushes a new outermost VLAN tag
Pop VLAN		Pops the outermost VLAN tag

Field	Argument	Description
Set Field	VLAN	Sets a new value in the VLAN tag. Ignored if the packet does not have a VLAN tag.
Group	Group entry	Must be one of: L2 Interface; L2 Unfiltered Interface group entry

4.3.15.3 Counters

The OF-DPA MPLS L2 Tag Group entry type counters are as shown in Table 141.

Table 141 OF-DPA MPLS L2 Tag Group Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow or group entities currently referencing this group entry.
Duration (sec)	32	Per-entry	Seconds since this group entry was installed

4.4 Meters

This section describes the OF-DPA Meter Table. OF-DPA metering contains extensions to support RFC 2698 and RFC 2698 token bucket meters. They are based on the Meter Table definitions in OpenFlow 1.3, but extended to allow specifying the policer mode (TrTCM or SrTCM) and for changing the packet Color field.

This version of OF-DPA supports a single meter per flow using a Policy ACL Flow Table instruction. OpenFlow 1.3 specifies that a flow entry can have a single meter instruction that is evaluated before any other instructions, in particular, before an Apply-Actions instruction. OF-DPA meter bands are only type “color set,” that is, the only action is to set the packet color. A Policy ACL Flow Table rule that has a meter instruction cannot specify a Set-Field Color action using an Apply-Actions or Write-Actions instruction.

The meter mode may be configured as color-aware or color-blind. Color-blind operation effectively treats all incoming packets as Green, regardless of color. The default is color-blind. OF-DPA Meter Bands are programmed using OpenFlow Meter Band Experimenter fields.

Note: Color-based metering has been proposed as a future OpenFlow feature [28]. Although an effort has been made to align OF-DPA meters with this proposal as much as is feasible, there are a number of differences, including how the configured parameters are used to measure rate and burst using the token bucket algorithm.

4.4.1 Meter Table Entries

A Meter Table entry contains the parameters shown in Table 142. OF-DPA supports Meter Table entries with two meter bands and a default “Green” band. This default has a rate and burst of zero. Furthermore the Red band rates must be greater than the Yellow band rates. In OpenFlow terms the range determines which band applies for a particular packet.




Figure 39 Meter Entry Example (TrTCM)

Figure 39 shows an example two rate three color meter (TrTCM) entry. In this example, both Yellow and Red bands are defined as Color Set, the only option. Packets that exceed the Committed Information Rate (CIR) or the Committed Burst Size (CBS) are considered out of profile. If a packet is out of profile but does not exceed the Peak Information Rate (PIR) or Peak Burst Size (PBS) it falls in the Yellow band. The Yellow band in this example sets the packet color to yellow. A subsequent queuing function would typically use the packet color to determine the packet drop precedence. Packets that exceed and the PIR or PBS fall in the Red band and are colored Red. These are usually dropped.

Table 142 lists the configuration parameters for OF-DPA Meters. All OF-DPA meters require a burst parameter.

Table 142 OF-DPA Meter Table Entry Parameters

Field	Description
Meter Id	Meter instance. Meter Id encodes the mode and color awareness of the meter according to the convention in Table 143.
Flags	Bit position: 0: Kbps (only one of Kbps or Packets set) 1: Packets 2: Burst (must always be set) 3: Stats (per-band counters, platform dependent)

Field	Description
Meter Bands	Two meter bands, Yellow (01), and Red (10). Configuration is as described in Table 145.
Counters	Per-meter entry counters are specified in Table 144

The meter identifier is structured so that it encodes the color-awareness and mode of the meter²⁶. Table 143 shows identifier details.

Table 143 OF-DPA Meter Entry Naming Convention

Field	Bits	Description
Index	[16:0]	Unique identifier for this meter entry
Reserved	[27:17]	
Color-Aware	[28]	0: Color-Blind 1: Color-Aware
Mode	[31:29]	1: TrTCM Mode 2: SrTCM Mode 4: Modified TrTCM Mode All other values reserved.

Table 144 Meter Entry Counters

Name	Bits	Type	Description
Reference Count	32	Per-entry	Number of flow entities currently referencing this meter table entry.
Duration (sec)	32	Per-entry	Seconds since this meter table entry was installed

²⁶ Adding flags to the meter header would have required an entirely new experimenter meter message type, since OpenFlow does not have Experimenter provisions to add flags. Instead OF-DPA incorporates the flags into the meter identifier. Experimenter Color Set meter bands are defined for the color actions.

4.4.2 Meter Bands

Meter bands are configured using the entries in Table 145. The OpenFlow Drop type can only be used for Red bands.

OF-DPA meter operation is described in detail in Table 1 and Table 2.

Table 145 Meter Band Configuration Parameters

Field	Description
Type	Meter bands must be Color Set. A Yellow band sets the Color to 0x01. A Red band sets the Color to 0x10.
Rate	Minimum rate for applying this meter. This value programs the token bucket CIR or PIR in the units specified by the meter entry.
Burst	Burst size for packets to be in profile. This value programs the token bucket CBS, PBS, or EBS depending on mode.
Counters	OF-DPA meter band counters are as in Table 146.

Note: All meters must support both rate and burst, and all bands must be configured with a burst size greater than zero.

Note: For single rate meters, both bands must be programmed with the same rate value. The programmed rate and burst values are interpreted as in Table 3.

Table 146 Meter Band Counters

Name	Bits	Type	Description
Packet Count	64	Per-band	In-band packet count.
Byte Count	64	Per-band	In-band byte count

Note: Per-band counters are not supported on all platforms.

5 CONFIGURATION

This section describes the properties of OF-DPA features that must be configured by some mechanism outside of the OpenFlow protocol. These are presented requirements and do not prescribe any particular configuration approach. Configuration information is modeled using a combination of UML diagrams and tables. This information could be accessed using OF-Config [10], OVSDB [11], or some local CLI.

Section 5.1 describes the OF-DPA physical and logical port configuration, and Section 5.2 describes queue configuration. OAM message processing configuration diagrams are included in Section 5.3, and protection configuration is in Section 5.4

5.1 Ports

This section lists the OF-DPA supported properties for physical, reserved, and logical ports. Ports are identified using a 32-bit ifNum value. The most significant two bytes indicate the type of port. Port types with their numbering conventions are listed in Table 147.

Table 147 Port Type Numbering Conventions

Type	Numbering	Description
Physical	0x0000 xxxx	Physical (front panel) port
Reserved	0xFFFF xxxx	Reserved ports as defined in the OpenFlow specification.
Overlay Tunnel Logical Port	0x0001 xxxx	Logical port used for VXLAN or NVGRE overlay tunnels. Can be a VXLAN Tunnel Endpoint Network Port, an NVGRE Tunnel Endpoint Network Port, a VXLAN Access Port, or an NVGRE Access Port. Port number assigned by configuration.
OAM Protection Liveness Logical	0xF000 xxxx	Logical port used for OAM protection liveness testing. These ports are pre-defined and do not need to be

Type	Numbering	Description
Port		configured in order to be used. The default operational state is up (LIVE). Can be administratively configured down, in which case the operational state will change to 1 (down).

5.1.1 Physical Ports

OF-DPA supports physical ports that are available on specific target platforms. This section is unchanged from OF-DPA 1.0.

Physical ports are front panel ports on the Abstract Switch. Figure 40, from OF-Config 1.1.1, shows the UML port configuration data model.

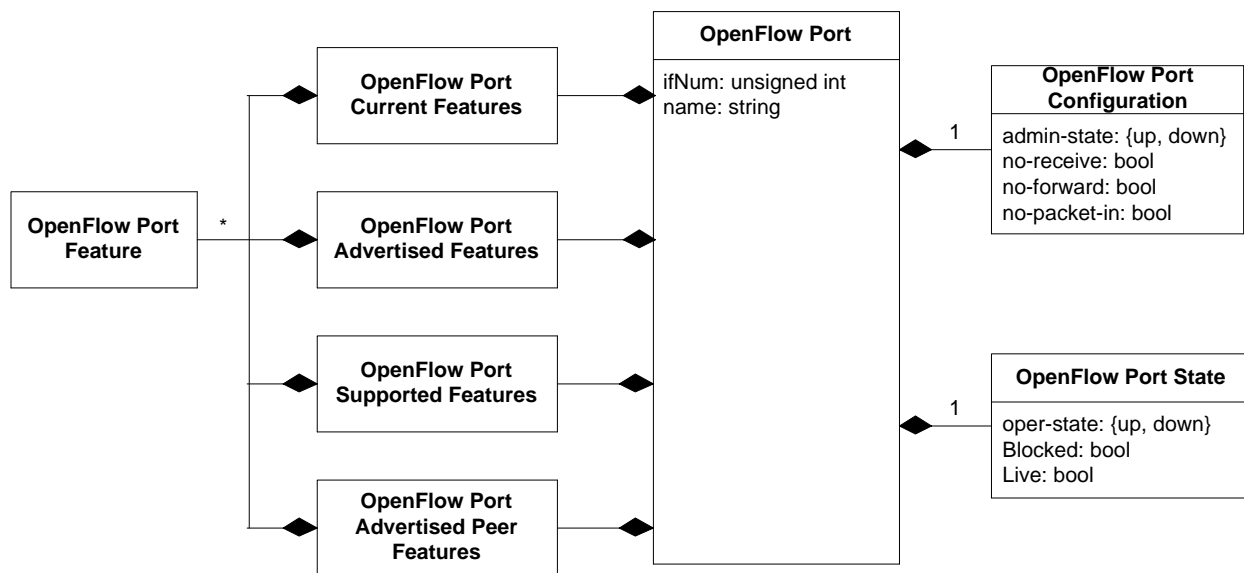


Figure 40 Port Properties Configuration

5.1.1.1 Features

OpenFlow Port Features are further modeled in terms of the sub-classes shown in Figure 41, also from OF-Config 1.2.

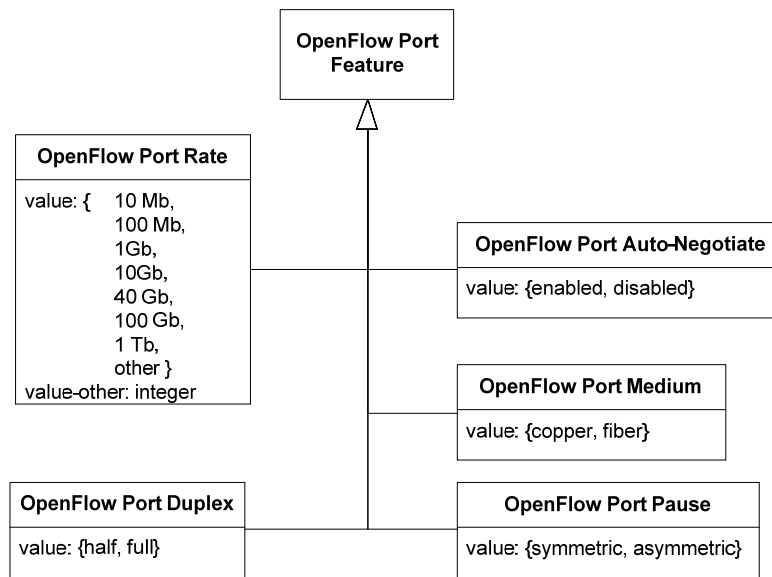


Figure 41 OpenFlow Feature Sub-Classes

For OF-Config, which is based on NetConf and YANG, UML models such as these are used to directly derive XML style sheets for representing configuration protocol messages.

OF-DPA supports the physical port features listed in Table 148.

Table 148 OF-DPA Port Features

Name	Bits	Configurable?	Description
Number	32	No	ifNum (should be the same as in interface MIB)
Hardware Address	48	No	MAC address assigned to port.
Name	128	Yes	16-byte string name (should be the same as in interface MIB)
Configured State	32	Yes	Port is administratively up (0) or down (1)
Current State	32	No	Port link (operational) state is up (0), live (4), or down (1). Generally a port is live if operationally up.
Current Features	32	No	OF-DPA supports the feature bitmap in Table 149. A one indicates the feature is currently active.
Advertised	32	No	OF-DPA supports the feature bitmap in Table 149. A zero bit

Name	Bits	Configurable?	Description
Features			indicates the feature is not available.
Supported Features	32	No	OF-DPA supports the features in Table 149. A zero bit indicates the feature is not supported.
Peer Features	32	No	Bitmap indicating capabilities advertised by the peer from Table 149.
Current Speed	32	No	Current port bitrate in kbps
Max Speed	32	No	Maximum port bitrate in kbps

Note: Not all of the above may be applicable to the LOCAL or CONTROLLER reserved port.

Table 149 shows the port features bitmap referenced from the table above and the OpenFlow Port Features sub-classes in Figure 41.

Table 149 Port Features Bitmap

Feature	Bit	Description
10 Mbps HD	0	10 Mbps half-duplex
10 Mbps FD	1	10 Mbps full-duplex
100 Mbps HD	2	100 Mbps half-duplex
100 Mbps FD	3	100 Mbps full-duplex
1GB HD	4	1 Gbps half-duplex
1GB FD	5	1 Gbps full-duplex
10GB FD	6	10 Gbps full-duplex
40GB FD	7	40 Gbps full-duplex
100GB FD	8	100 Gbps full-duplex

Feature	Bit	Description
1TB FD	9	1 Tbps full-duplex
Other	10	Other rate, not in the above list
Copper	11	Copper medium
Fiber	12	Fiber medium
Autoneg	13	Auto-negotiation
Pause	14	Pause enabled
Pause_Asym	15	Asymmetric pause supported

5.1.1.2 Counters

OF-DPA supports the port counters listed in Table 150.

Table 150 OF-DPA Physical Port Counters

Name	Bits	Description
Received Packets	64	Total packets received
Transmitted Packets	64	Total packets transmitted
Received Bytes	64	Total bytes received
Transmitted Bytes	64	Total bytes transmitted
Receive Drops	64	Received packets dropped for any reason
Transmit Drops	64	Transmitted packets dropped for any reason
Receive Errors	64	Received packet errors
Transmit Errors	64	Transmit packets errors
Receive Frame Alignment	64	Received packets with frame alignment errors

Name	Bits	Description
Errors		
Receive Overrun Errors	64	Received packet overruns
Receive CRC Errors	64	Received packet CRC errors
Collisions	64	Collisions
Duration (sec)	32	Time in seconds since configured

5.1.2 Reserved Ports

OF-DPA supports the reserved ports listed in Table 151. These do not required configuration and are listed for general information.

Table 151 OF-DPA Reserved Ports

Name	Required	Description	Use	Supported?
ALL	Yes	Required but not supported in OF-DPA.	Output	No
IN_PORT	Yes	Used to send packets to the ingress port to override OpenFlow default behavior. OF-DPA uses group ALLOW-IN_PORT property instead. Not to be confused with the IN_PORT match field.	Output	No
CONTROLLER	Yes	The OpenFlow controller. Output destination for sending packets to the Agent which, in turn, sends to the OpenFlow Controller in a Packet_In message. Also can optionally be used to indicate the source of packets received by the Agent in a Packet_Out message.	Input or output	Yes
TABLE	Yes	Used in Packet_Out messages to indicate that the packet must be recirculated through the pipeline. Must always be the first table in the pipeline if specified.	Output	Yes
ANY	Yes	Special value used in some requests.	Neither	Yes

Name	Required	Description	Use	Supported?
LOCAL	No	Used to send and receive packets with the local Network Protection App. Analogous to Controller but the destination is a local OAM engine rather than the Agent. The exact mechanism is implementation dependent.	Input or output	Yes, for OAM
NORMAL	No	Not supported in OF-DPA	Output	No
FLOOD	No	Not supported in OF-DPA	Output	No

5.1.3 Logical Ports

Logical Ports are used to model functionality as external of the OpenFlow pipeline, such as adding or removing tunnel encapsulation headers. The Abstract Switch receives packets from an ingress logical port after they have been processed by an external function, and forwards packets to a destination logical port for processing by an external function. Thus, packets received from logical ports may not be the same as were transferred on the wire.

Logical port external function parameters are configured on the Logical Port. OF-DPA uses Logical ports for overlay tunnels and as objects for MPLS Protection Group entry buckets to watch.

Except for Section 5.1.3.3 on OAM Protection Liveness logical ports, this section is unchanged from OF-DPA 1.0.

5.1.3.1 Overlay Tunnels

Tunnel packets enter the OF-DPA data path from Tunnel Logical Ports, along with tunnel id metadata. The tunnel id identifies the tenant forwarding domain. Tunnel Logical Ports are modeled according the UML data model in Figure 42, which shows Tunnel Logical Ports as abstract classes.

Note that there are two Tunnel Logical Port sub-types shown. Access ports connect local servers in the tenant forwarding domain. Tunnel Endpoints connect to remote switches.

Tunnel Logical Ports must have a specified protocol in order to be instantiated. OF-DPA supports VXLAN [21] overlays.

Note: VXLAN support is hardware platform and version dependent.

The Tunnel Endpoint abstract class provides necessary configuration parameters common to different protocol sub-classes. This includes the local and remote endpoint addresses, the TTL for packet origination, and the multipath properties for forwarding tunnel initiation packets.

Similarly, the Tunnel Access Port abstract class provides necessary parameters for locally attached servers. Three methods are supported: all traffic on a port, all traffic with a particular VLAN id on a port, or packets tagged with an IEEE 802.1BR [20] port extension tag (ETAG) on that port.

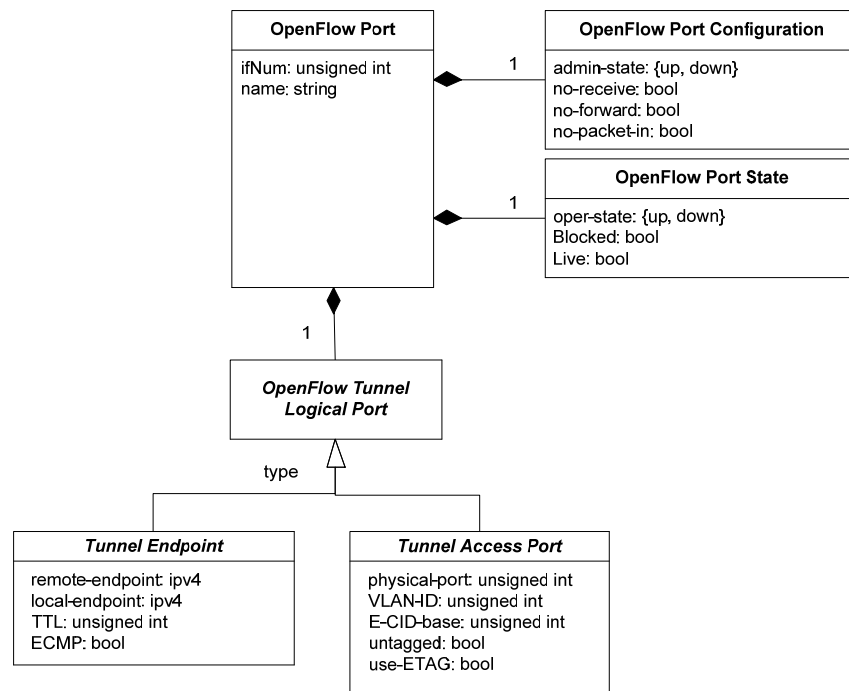


Figure 42 Tunnel Logical Port Configuration

5.1.3.2 VXLAN Tunnel Logical Port Configuration

VXLAN Logical Port configuration is shown in the UML diagram in Figure 43. Two types of VXLAN Logical Port can be configured: VXLAN Tunnel Endpoints, and VXLAN Access Ports.

VXLAN Tunnel Endpoint Logical Ports are used to forward packets to a remote tunnel endpoint, or VTEP. The VXLAN Tunnel Endpoint class is configured with protocol specific header properties as well as tunnel initiation forwarding properties. The use-entropy flag indicates that a hash value is to be inserted instead of the configured udp-src-port-if-no-entropy setting.

Note: This version of OF-DPA supports hardware with a single system-wide configuration for the terminator-udp-dest-port and use-entropy settings. As a result, all configured VXLAN Tunnel Endpoints must specify the same values for these parameters.

VXLAN Tunnel Endpoints must be configured with forwarding state for tunnel initiation packets. This can be specified in terms of a VXLAN Unicast Tunnel Next Hop or an ECMP VXLAN Next Hop Group multipath object, if the ECMP flag is set. The ECMP class aggregates one or more VXLAN Unicast Tunnel Next Hop objects.

Traffic for multiple tenants can be multiplexed over a VXLAN Tunnel Endpoint. The VXLAN Tenant class provides the protocol header information (VN_ID) for distinguishing a particular Tenant's traffic. It also

identifies the Tenant's isolated forwarding domain for ingress and egress packets. The VXLAN Tunnel class can also provide an IP multicast group address for this Tenant's traffic. Note that multiple VXLAN Tenants can share an IP multicast address.

The VXLAN Access Port class configures a logical port for a VXLAN Tenant's locally attached endpoint. The Access Port configuration specifies how traffic is classified to a particular VXLAN Tenant isolated forwarding domain. This can be one of: all traffic on a port; traffic on a port with a particular VLAN Id; and traffic with a particular E-Tag [20].

Note that there is some interaction between the overlay tunnel configuration, VLAN Flow Table entries, and L2 Interface Groups. For Access Ports, configuration must be mutually exclusive in order to isolate overlay tenant traffic. This means that a VLAN Flow Table entry must not specify filtering for local tenant traffic configured via an Access Port, and an L2 Interface Group must not call out the same port and VLAN properties as for local tenant traffic. Tunnel Endpoint operation, however, depends upon L2 Interface Group settings for forwarding underlay VXLAN packets initiated by a Tunnel Next Hop, and upon VLAN Flow Table entries to permit receiving underlay terminated tunnel packets.

OF-DPA provides configuration APIs for VXLAN Tunnel Endpoints, VXLAN Tenants, VXLAN Access Ports, ECMP VXLAN Next Hop Groups, and VXLAN Unicast Tunnel Next Hops. Configuration parameters are listed in Table 152 thru Table 155.

Note: OF-DPA derives header fields for multicast tunnel origination packets from related configuration values. For example, if multiple local endpoint IP addresses are configured OF-DPA will select one.

Table 152 VXLAN Tunnel Endpoint Logical Port Configuration Parameters

Name	Bits	Description
Remote Endpoint	32	IPv4 SIP in termination packets, or DIP for unicast origination packets.
Local Endpoint	32	IPv4 DIP for termination packets, or SIP for origination packets.
TTL	8	TTL value for use in origination packets.
ECMP	1	Use multipath forwarding for origination packets.
Terminator UDP Dest Port	16	Destination UDP port for recognizing termination VXLAN frames.
Initiator UDP Dest Port	16	Destination UDP port to put in originating VXLAN frames.
UDP Source Port	16	Default source port to use if entropy option is not used
Use Entropy	1	Insert hash value in place of UDP source port

Table 153 VXLAN Access Logical Port Configuration Parameters

Name	Bits	Description
Port	32	Local port
VLAN id	16	VLAN id to match or use if VLAN tagged
E-CID	16	E-CID value to match or use if IEEE 802.1BR tagged
Untagged	1	All traffic on port is for the same tenant
Use ETAG	1	Use IEEE 802.1BR tagging rather than VLAN id

Table 154 VXLAN Tenant Configuration Parameters

Name	Bits	Description
Tunnel id	32	Value to associate with packets for this tenant. Identifies the tenant forwarding domain.
VN_ID	24	Segment identifier in the VXLAN header that identifies this tenant.
Multicast IP	32	Multicast group IP address associated with this tenant

Table 155 VXLAN Next Hop Configuration Parameters

Name	Bits	Description
Next Hop Id	32	Identifier used to reference next hop objects
MACSA	64	Underlay source MAC address
MACDA	64	Underlay destination MAC address
VLAN id	16	VLAN id to use if tagged
Port	32	Egress port for forwarding

5.1.3.3 OAM Protection Liveness Logical Ports

OAM Protection Liveness Logical Ports have no configuration. They are pre-defined in a specific range and default to administratively up. The operational state is always the same as the administrative state.

OAM Protection Liveness Logical Ports are used as a data object to control Fast Failover bucket liveness state.

5.2 Queues

OF-DPA supports eight queues per standard port. On some platforms, OF-DPA supports a larger number of service based queues.

This section is unchanged from OF-DPA 1.0.

5.2.1 Configuration

OF-DPA queue configuration parameters are listed in Table 156. Queue_Id is always relative to the port to which the queue is attached. Queue_Id values must be a value between zero and seven.

Table 156 OF-DPA Queue Configuration Parameters

Name	Bits	Description
Queue_Id	32	Identifier for this specific queue. Must be a value between 0 and 7.
Port	32	Port to which this queue is attached
Max Rate	16	Maximum rate in terms of a percentage of the port rate, specified in increments of .1%. A value of 1000 means no maximum rate.
Min Rate	16	Minimum rate in terms of a percentage of the port rate, specified in increments of 0.1%. A value of 1000 means no minimum rate.

5.2.2 Counters

OF-DPA queues counters are as shown in Table 157.

Table 157 OF-DPA Queue Counters

Name	Bits	Description
Transmit Packets	64	Total packets transmitted
Transmit Bytes	64	Total bytes transmitted
Duration (sec)	64	Duration in seconds

5.3 OAM Message Processing

This section describes configuration for a local Network Protection App. Similar configuration could be used for Controller hosted Network Protection App. The Configuration model is intended to closely follow the parameters defined in OAM standards [15][16][27].

5.3.1 MPLS-TP Ethernet OAM Configuration

OAM Engine configuration for Ethernet OAM follows IEEE 802.1ag [27]. The configuration information model for Fault Management is diagrammed in Figure 44, and for Performance Monitoring in Figure 45.

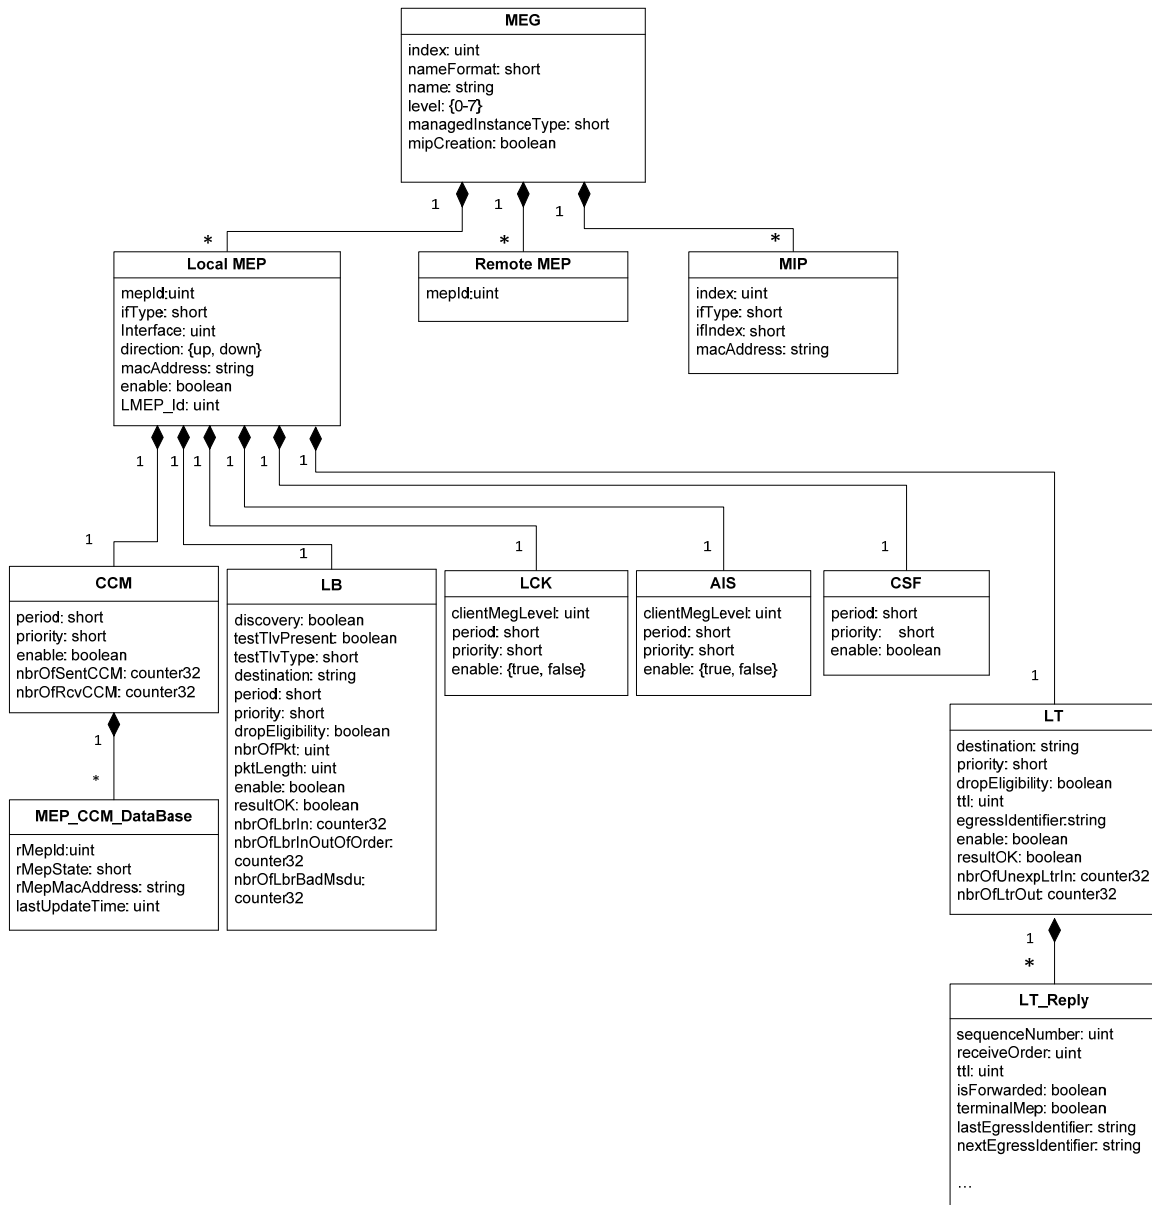


Figure 44 OAM Ethernet Fault Management Configuration

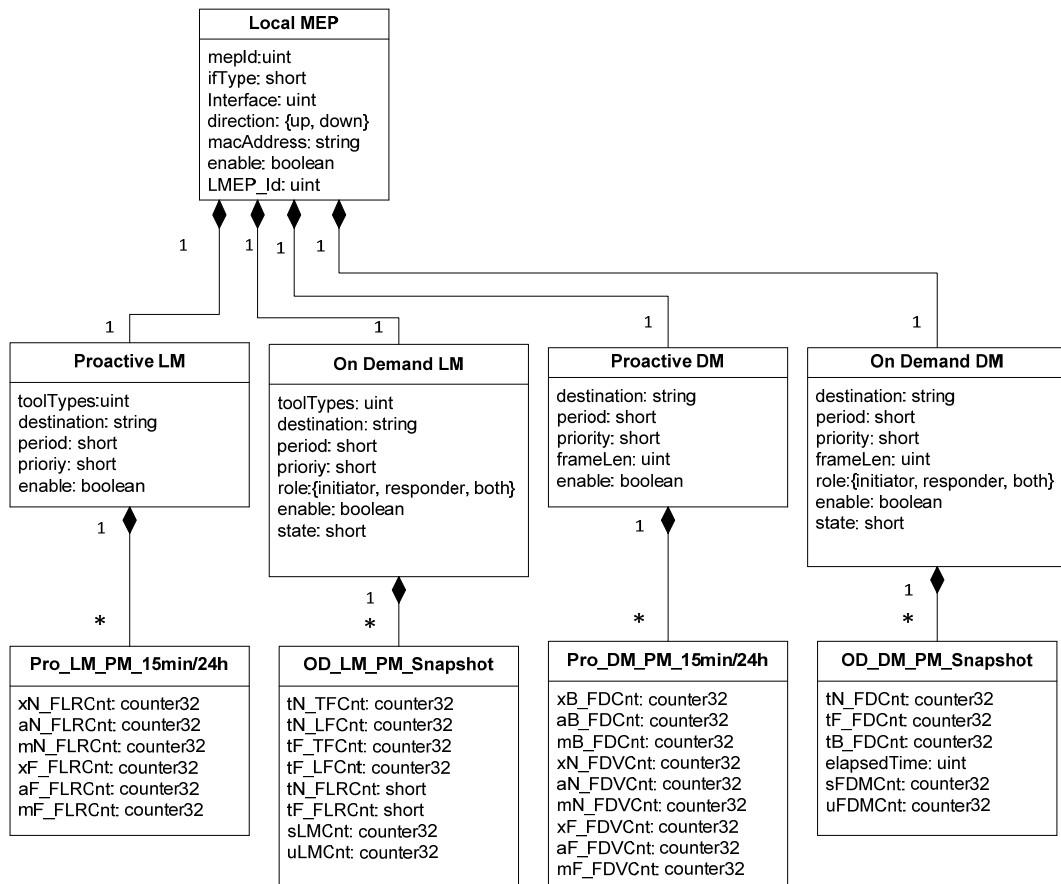


Figure 45 OAM Ethernet Performance Monitoring Configuration

5.3.2 MPLS-TP G.8113.1 OAM Configuration

OAM Engine configuration for G.8113.1 OAM follows [15]. The configuration information for Fault Management is shown in Figure 46, and for Performance Monitoring in Figure 47.

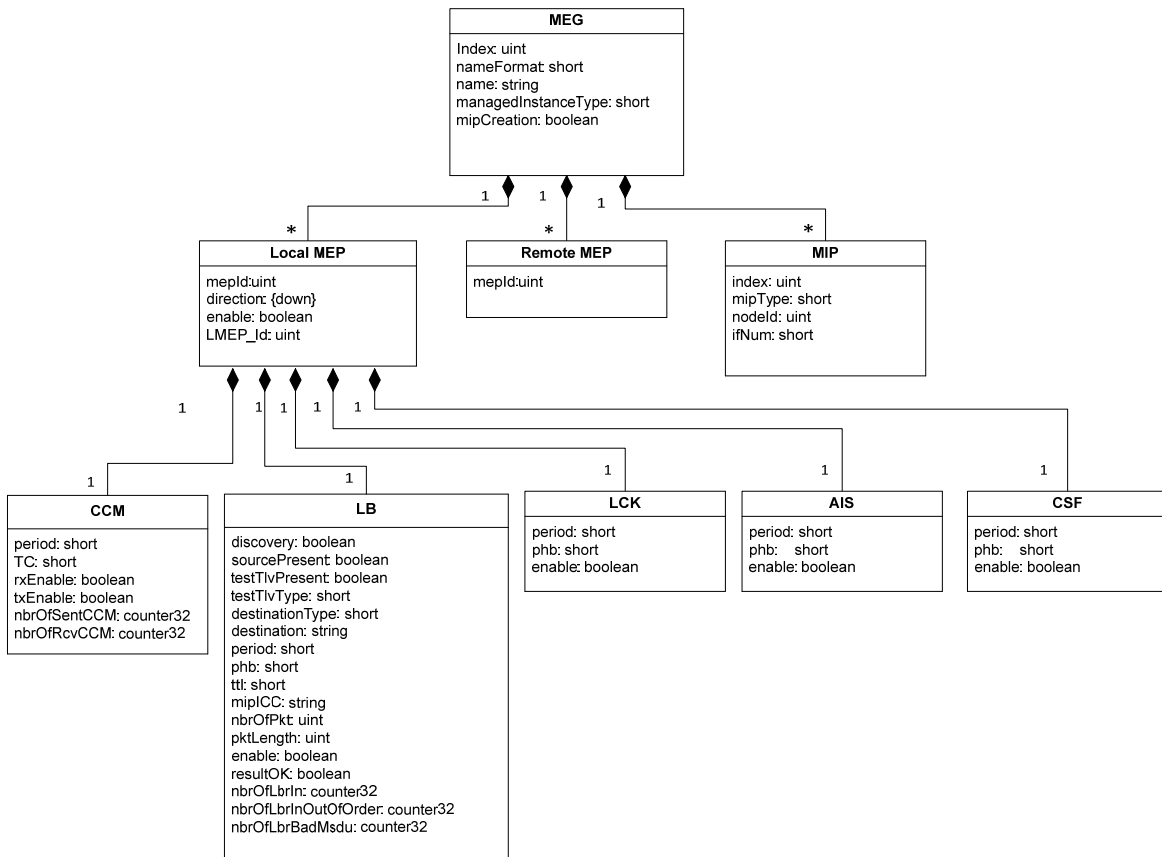


Figure 46 OAM G.8113.1 Fault Management Configuration

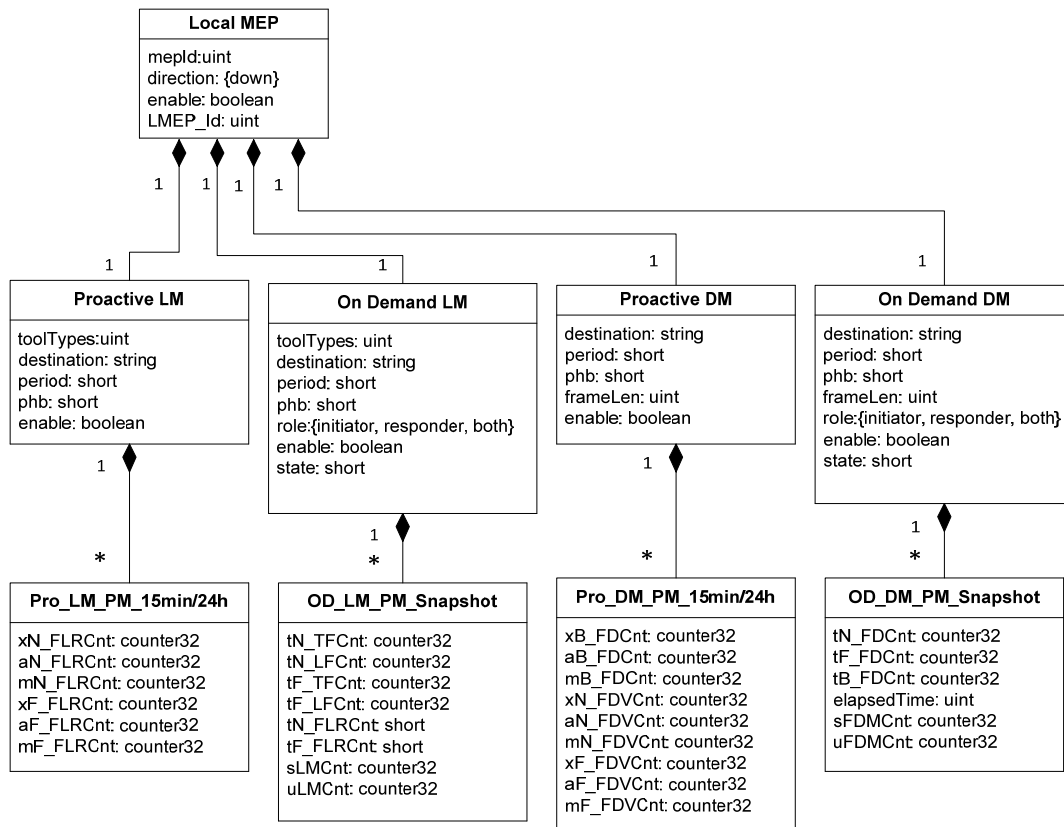


Figure 47 OAM G.8113.1 Performance Monitoring Configuration

5.4 Protection

The Protection Process can control the switchover from worker to protection path by changing the administrative state of an OAM Protection Liveness logical Port. The Protection Process uses the information in the Fast Failover Configuration Table to determine which port to update for a particular path.

5.4.1 MPLS-TP Linear Protection

The Protection Process can control the switchover from worker to protection path by changing the administrative state of an OAM Protection Liveness logical Port. The Protection Process uses the information in the MPLS-TP linear protection configuration tables to determine which port to update for a particular path.

The MPLS-TP Linear Protection data model is diagrammed below. Note that Tail End Protection is for future use.

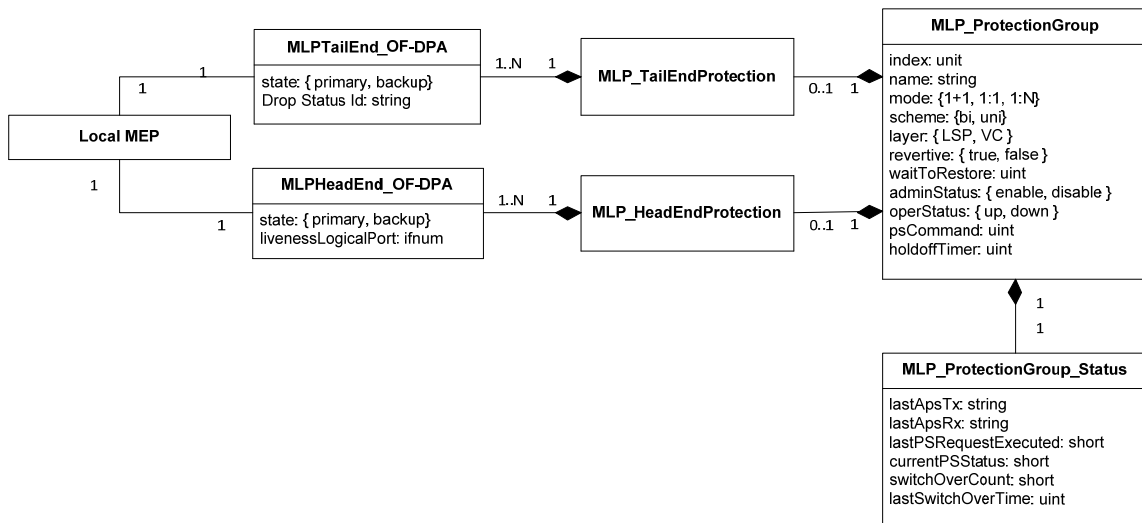


Figure 48 MPLS-TP Linear Protection Configuration

6 VENDOR EXTENSION FEATURES

OF-DPA provides vendor extensions for source MAC learning, L3 forwarding IN_PORT control, MPLS and OAM actions and pipeline match fields, and new ancillary object types.

In many cases the vendor extension features only affect the OpenFlow abstract switch and can be accommodated by the existing OpenFlow 1.3.4 protocol. In others, an OpenFlow 1.3.4 agent and compatible controller can be extended using the OpenFlow Experimental facility to add new protocol elements as needed.

6.1 Source MAC Learning

OpenFlow 1.3.4 does not provide for flow tables that provide different views into the same database, e.g., using different lookup keys for different purposes.²⁷

The Bridging Flow Table contains MAC forwarding entries and is looked up by MAC_DST and either VLAN or Tunnel_id. An exact match hit in the table sets the id of the group entry for forwarding the packet. However if there is no exact match, a flow entry that matches the VLAN or Tunnel_id but wildcards the ingress port can provide a flooding forwarding group entry for destination location forwarding (DLF).

Source MAC learning is typically used to discover the MAC-to-port binding for populating the MAC table. A second lookup is done in the same table using the MAC_SRC and VLAN or Tunnel_id. If there is a hit, the output interface is compared against the IN_PORT. If there is a mismatch, an entry for this MAC and VLAN or Tunnel_id is added to the table along with the interface. If there is a hit but the interface values are different, it means the end station has moved and the entry needs to be updated accordingly.²⁸

OF-DPA implements optional logic for identifying when a MAC-to-port binding needs to be learned as a vendor extension. This function looks up all packets, regardless of whether they will be processed using the Bridging Flow Table or the Routing Flow Table. If the MAC_SRC and VLAN or Tunnel_id miss, or if the source has moved, the logic does one of two things, depending on the configuration.

Note: Network Virtualization SDN use cases, especially in data center and enterprise networks, centrally manage L2 forwarding and VLAN tables based on network discovery and do not rely on learning and flooding.

Note: The learning port depends on the packet flow. Bridging and routing learns physical ports, while overlay tunnels learn tunnel endpoint logical ports.

²⁷ This is addressed to some degree by the table synchronization features in OpenFlow 1.4.

²⁸ This is subject, of course, to security policy

6.1.1 Controller Managed Learning

If all learning is to be managed by the OpenFlow Controller, OF-DPA will send a PACKET_IN message to the Controller with a reason code (no match), special table id that indicates the learning lookup, and the MAC_SRC, VLAN/Tunnel_Id, and IN_PORT match fields that missed. The controller can then send a flow mod message to the switch to add an appropriate entry in the Bridging Flow Table or update Group entries. Since a PACKET_OUT is not expected as well, there is no need to buffer the miss packet, which would have already been forwarded normally.

To prevent multiple PACKET_IN learning messages, OF-DPA adds a pending (disabled) entry in the Bridging Flow Table. This entry will be removed after a configured interval if the controller does not come back with a flow mod.

For Controller managed learning the feature would be configured with destination CONTROLLER.

6.1.1.1 Configuration

The configuration options are listed in Table 158.

Table 158 Source MAC Learning Feature Configuration

Name	Description
Enable	Enable the source MAC learning feature.
Destination	CONTROLLER.
Duration	If the destination is CONTROLLER, indicates the time interval after which the pending entry is removed if the Controller does not issue a Flow Mod to keep it.

6.2 Additional Group Properties

OF-DPA adds the vendor extension property "ALLOW-IN_PORT" to OF-DPA L3 Interface group entries. This property applies to the group entry and to any referenced group entries. All L3 Interface Group entries automatically come with the property set, and it cannot be overridden. This obviates the need for special protocol support in OpenFlow 1.3.4.

6.3 MTU Check

OF-DPA adds an MTU check with the TTL check described in the OpenFlow specification. The same error code is used for both TTL and MTU check.

MTU check is required in order to implement an IP router and enable it to set the appropriate ICMP destination unreachable reason code.

This approach does not require special protocol support in OpenFlow 1.3.1.

6.4 Table Numbering

OF-DPA table number assignments are shown in Table 159.

Table 159 Flow Table Number Assignments

Flow Table Name	Number
Ingress Port Flow Table	0
VLAN Flow Table	10
VLAN 1 Flow Table	11
MPLS L2 Port Flow Table	13
Termination MAC Flow Table	20
MPLS 0 Flow Table	23
MPLS 1 Flow Table	24
MPLS 2 Flow Table	25
Maintenance Point Flow Table	26
Unicast Routing Flow Table	30
Multicast Routing Flow Table	40
Bridging Flow Table	50
Policy ACL Flow Table	60
Color Based Actions Flow Table	65
Egress VLAN Flow Table	210
Egress VLAN 1 Flow Table	211

Flow Table Name	Number
Egress Maintenance Point Flow Table	226
Source MAC Learning Flow Table	254

6.5 Experimenter Features

This section describes extensions to OpenFlow required by OF-DPA 2.0 that require protocol support in order to be used with OpenFlow 1.3.4. These are encoded using the OpenFlow Experimenter facilities described.

Note that some of the extensions described in this section have already been proposed for OpenFlow 1.5. In many cases the experimenter versions will be able to be replaced with equivalent OpenFlow 1.5 protocol elements once agent and controller support becomes available. At the time of this writing OpenFlow 1.5 is still a work in progress.

The experimenter id used for these features in OF-DPA is the Broadcom OUI: 00-10-18.

6.5.1 OAM Data Plane Counter Table

The OAM Data Plane Counter Table is defined to maintain the LM counters needed for OAM processing. Entries in this table implement a globally accessible LM counter resource. They can be updated using actions (OAM_LM_TX_Count, OAM_LM_RX_Count) from flow tables and group entries and read by the controller or a local OAM processor. OF-DPA 2.0 includes actions to update LM counters from the VLAN Flow Tables, MPLS Flow Tables, Egress VLAN Flow Tables, and MPLS Label Group entries.

The OAM Data Plane Counter Table is indexed by LMEP_Id and Traffic Class. Each entry can hold a packet counter for transmit and receive directions. Furthermore each entry maintains a reference count.

Table 160 OAM Data Plane Counter Table Entry

Field	Bits	Description
LMEP Id	32	Index key
Traffic Class	4	Index key
Transmit Packets	64	TxFCL in [15]
Receive Packets	64	RxFCL in [15]
Reference Count	64	Number of objects referencing this entry.

An experimenter message type is used to support modifications to the OAM Data Plane Counter Table. This is only used as a Controller/Switch message. An entry in the table should be created for each configured LMEP Id.

```

struct ofdpa_OAM_DPC_mod_msg {
    struct ofp_header  header;      /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter; /* Experimenter ID: 00-00-10-18 */
    uint32_t          exp_type;     /* OAM Data Plane Counter Table (1). */
    uint32_t          command;      /* one of OFDPA_MSG_MOD_* */
    uint32_t          LMEP_id;      /* index into table entry */
    uint64_t          transmit_packets; /* clear only */
    uint64_t          receive_packets; /* clear only */
    uint8_t           traffic_class; /* sub-index into table entry */
    uint8_t           pad[3];       /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_OAM_DPC_mod_msg) == 60);

enum ofdpa_message_mod_command {
    OFDPA_MSG_MOD_ADD = 0; /* New counter entry */
    OFDPA_MSG_MOD_MODIFY = 1; /* Modify entry. Used to zero counters. */
    OFDPA_MSG_MOD_DELETE = 2; /* Delete entry */
};

```

An experimenter multipart message type is used to read status of the OAM Data Plane Counter Table. It is used for both request and reply messages.

```

struct ofdpa_OAM_DPC_multipart {
    struct ofp_header  header;      /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter; /* Experimenter ID: 00-00-10-18 */
    uint32_t          exp_type;     /* OAM Data Plane Counter Table (1). */
    uint32_t          LMEP_id;      /* Zero indicates all entries*/
    uint64_t          transmit_packets; /* In multipart reply only */
    uint64_t          receive_packets; /* In multipart reply only */
    uint64_t          reference_count; /* In multipart reply only */
    uint8_t           traffic_class; /* Zero indicates all traffic classes */
    uint8_t           pad[3];       /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_OAM_DPC_multipart) == 40);

```

6.5.2 Drop Status Action Table

The Drop Status Action Table performs a packet drop action when invoked by a Check Drop Status action with an Index and Type argument. It is used to implement the OAM LCK function in order to drop OAM and data frames during a lock condition. Entries in this table are indexed by LMEP_id.

The Drop Status Table can be considered an auxiliary object similar to a Meter entry with “drop” band except that the drop action is controlled by the Network Protection Process using an unspecified internal interface, similar to what is done for the switchover control on the OAM Protection Liveness Logical port.

Table 161 Drop Status Table Entry

Field	Bits	Description
Type	8	0: Lock Status 1-7: Reserved
Index	32	Index key. LMEP Id must be specified with Type 0.
Drop Action	1	0: Do not drop 1: Drop

Note: The default action on a Check Drop-Status action lookup miss is “Do not drop.”

An experimenter message type is used to support modifications to the Drop Status Table. This is only used as a Controller/Switch message. An entry in the table should be created for each configured LMEP Id.

```
struct ofdpa_OAM_DST_mod_msg {
    struct ofp_header  header;          /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter;    /* Experimenter ID: 00-00-10-18 */
    uint32_t          exp_type;        /* Drop Status Table (2). */
    uint32_t          command;         /* one of OFDPA_MSG_MOD_* */
    uint32_t          index;           /* index into table entry */
    uint8_t           type;            /* type of entry */
    uint8_t           drop-status;     /* 0=do not drop; 1= drop */
    uint8_t           pad[6];         /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_OAM_DST_mod_msg) == 24);
```

An experimenter multipart message type is used to read status of the Drop Status Table. It is used for both request and reply messages.

```
struct ofdpa_DROP_STATUS_multipart {
    struct ofp_header  header;          /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter;    /* Experimenter ID: 00-00-10-18 */
    uint32_t          exp_type;        /* Drop Status Table (2) */
    uint32_t          index;           /* Zero indicates all entries*/
    uint8_t           type;            /* Zero indicates all entries*/
    uint8_t           drop-status     /* 0=do not drop; 1= drop */
    uint8_t           pad[1];         /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_DROP_STATUS_multipart) == 16);
```

6.5.3 MPLS Label Remark Action Tables

The MPLS Label Remark Action tables can be invoked from the MPLS Label Group entries using optional Set-TC-From-Table and Set-PCP-DFI-From-Table actions. The Set-TC-From-Table action uses the packet Traffic Class and Color and a supplied QoS Index argument to perform a Set-Field(MPLS_TC) action on the label. If the outermost label, the Set-PCP-DFI-From-Table action uses the packet Traffic Class and Color along with a supplied QoS Index argument to perform Set-Field(PCP) and Set-Field(DEI) actions on the outermost VLAN tag.

While similar to match-action tables in some respects, these tables are invoked in the context of evaluating an Action Set to perform Set-Field actions. These tables are modified and read using new Experimenter messages.

There are two MPLS Label Remark Action tables:

- MPLS VPN Label Remark Action Table
- MPLS Tunnel Label Remark Action Table

Entries in both of these have the fields shown in Table 162.

Table 162 MPLS Label Remark Table Entry Fields

Field	Bits	Description
Index	4	Mapping profile index
Traffic Class	4	Traffic Class
Color	2	Color
MPLS_TC	3	Three-bit EXP value to use with the MPLS label.
VLAN_PCP	3	Three-bit PRI value to use in the outermost VLAN tag.
VLAN_DEI	1	DEI value to use in the outermost VLAN tag.

6.5.3.1 MPLS VPN Label Remark Action Table

An experimenter message type is used to support modifications to the MPLS VPN Label Remark Action Table. This is only used as a Controller/Switch message.

```
struct ofdpa_MPLS_VPN_LABEL_REMARK_ACTION_mod_msg {
    struct ofp_header  header;      /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter; /* Experimenter ID: 00-00-10-18 */
}
```

```

uint32_t    exp_type;           /* MPLS VPN Label Remark Action (3) */
uint32_t    command;           /* one of OFDPA_MSG_MOD_* */
uint32_t    index;             /* index */
uint8_t     traffic_class;     /* traffic class */
uint8_t     color;             /* color */
uint8_t     mpls_tc;           /* MPLS TC value to set */
uint8_t     vlan_pcp;          /* outer vlan PCP to set */
uint8_t     vlan_dei;          /* outer vlan DEI to set */
uint8_t     pad[3];            /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_MPLS_VPN_LABEL_REMARK_ACTION_mod_msg) == 24);

```

An experimenter multipart message type is used to read status of the MPLS VPN Label Remark Action Table. It is used for both request and reply messages.

```

struct ofdpa_MPLS_VPN_LABEL_REMARK_ACTION_multipart {
    struct ofp_header  header;    /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t           experimenter; /* Experimenter ID: 00-00-10-18 */
    uint32_t           exp_type;    /* MPLS VPN Label Remark Action (3) */
    uint32_t           index;       /* Zero indicates all values */
    uint8_t            traffic_class; /* Zero indicates all values */
    uint8_t            mpls_tc      /* Zero indicates all values */
    uint8_t            vlan_pcp     /* Zero indicates all values */
    uint8_t            vlan_dei     /* Zero indicates all values */
    uint8_t            pad[6];      /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_MPLS_VPN_LABEL_REMARK_ACTION_multipart) ==
16);

```

6.5.3.2 MPLS Tunnel Label Remark Action Table

An experimenter message type is used to support modifications to the MPLS Tunnel Label Remark Action Table. This is only used as a Controller/Switch message.

```

struct ofdpa_MPLS_TUNNEL_LABEL_REMARK_ACTION_mod_msg {
    struct ofp_header  header;    /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t           experimenter; /* Experimenter ID: 00-00-10-18 */
    uint32_t           exp_type;    /* MPLS VPN Label Remark Action (3) */
    uint32_t           command;     /* one of OFDPA_MSG_MOD_* */
    uint32_t           index;       /* index */
    uint8_t            traffic_class; /* traffic class */
    uint8_t            color;       /* color */
    uint8_t            mpls_tc;     /* MPLS TC value to set */
    uint8_t            vlan_pcp;    /* outer vlan PCP to set */
    uint8_t            vlan_dei;    /* outer vlan DEI to set */
    uint8_t            pad[3];      /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_MPLS_TUNNEL_LABEL_REMARK_ACTION_mod_msg) ==
24);

```

An experimenter multipart message type is used to read status of the MPLS Tunnel Label Remark Action Table. It is used for both request and reply messages.

```
struct ofdpa_MPLS_TUNNEL_LABEL_REMARK_ACTION_multipart {
    struct ofp_header  header;      /* Type OFPT_EXPERIMENTER, 16 bytes */
    uint32_t          experimenter; /* Experimenter ID: 00-00-10-18 */
    uint32_t          exp_type;     /* MPLS VPN Label Remark Action (3) */
    uint32_t          index;       /* Zero indicates all values */
    uint8_t           traffic_class; /* Zero indicates all values*/
    uint8_t           mpls_tc      /* Zero indicates all values */
    uint8_t           vlan_pcp     /* Zero indicates all values */
    uint8_t           vlan_dei     /* Zero indicates all values */
    uint8_t           pad[6];      /* align message on 64-bit boundary */
};
OFP_ASSERT(sizeof(struct ofdpa_MPLS_TUNNEL_LABEL_REMARK_ACTION_multipart) ==
16);
```

6.5.4 Actions

OF-DPA 2.0 requires a number of new actions which are described in Table 163.

Table 163 OF-DPA Experimenter Actions

Action	Argument	Description
Push L2 Header	None	Push a new outermost Ethernet header on the packet.
Pop L2 Header	None	Pop the outermost Ethernet header from the packet. The header cannot have a VLAN tag.
Push CW	None	Push MPLS-TP PW Control Word. Used in MPLS L2 VPN Label group type entries.
Pop CW	None	Pop MPLS-TP PW Control Word. Used in MPLS label match actions when bottom of stack.
Copy TC In	None	Copy EXP field value in. Used in conjunction with popping an MPLS label.
Copy TC Out	None	Copy EXP value out, used when pushing a label.
Set TC From Table	QoS Index	MPLS label EXP field re-mark action based on packet Traffic Class, Color, and the provided QoS index argument to look up the new value in an MPLS Label Remark Action Table.

Action	Argument	Description
Set PCP DFI From Table	QoS Index	Outer VLAN tag re-mark action for PCP and DEI fields based on packet Traffic Class, Color, and the provided QoS index argument to look up the new value in an MPLS Label Remark Action Table.
OAM_LM_TX_Count	LMEP_Id, Traffic Class	Update loss measurement transmit counter (TxFC1) in the OAM Data Plane Counter Table for a Maintenance Point.
OAM_LM_RX_Count	LMEP_Id, Traffic Class	Update loss measurement receive counter (RxFC1) in the OAM Data Plane Counter Table for a Maintenance Point.
Set-Counter-Fields	LMEP_Id, Traffic Class	Set pipeline metadata fields (TxFC1, RxFC1) from the indicated OAM Data Plane Counter Table entry. Set pipeline metadata field RxTime from the current timestamp. These along with the LMEP Id are sent with PDUs that are output to an Network Protection App.
Decrement TTL and do MTU check	None	Decrement TTL and do MTU check.
Check Drop-Status	Index, type	Drops the packet if the indexed entry action indicates drop.

The action types are programmed using the following assignments:

```

/* OF-DPA Experimenter Acton types */
enum ofdpa_action_exp_type {
    OFDPA_ACT_PUSH_L2_Header =          1,
    OFDPA_ACT_POP_L2_Header =          2,
    OFDPA_ACT_PUSH_CW =                 3,
    OFDPA_ACT_POP_CW =                  4,
    OFDPA_ACT_COPY_TC_IN =              5,
    OFDPA_ACT_COPY_TC_OUT =             6,
    OFDPA_ACT_SET_TC_FROM_TABLE =       7,
    OFDPA_ACT_SET_PCP_DFI_FROM_TABLE =  9,
    OFDPA_ACT_OAM_LM_RX_COUNT =         10,
    OFDPA_ACT_OAM_LM_TX_COUNT =         11,
    OFDPA_ACT_OAM_SET_COUNTER_FIELDS =  12,
    OFDPA_ACT_DEC_TTL_MTU =             13,
    OFDPA_ACT_CHECK_DROP_STATUS =       14,
};

```


6.5.5 Match Fields

Experimenter pipeline metadata match fields are described in Table 164.

Table 164 OF-DPA Experimenter Match Fields

Action	Bits	Masked	Pre-requisite	Description
VRF	16	No		Pipeline metadata. Virtual Router, used to select virtual routing table.
Traffic Class	4	No		Pipeline metadata. QoS traffic class.
Color	2	No		Pipeline metadata. Drop precedence. Values are: 00: Green 01: Yellow 10: Red 11: reserved
DEI	1	No	VLAN tag	Drop eligibility indication from 802.1Q.
LMEP Id	32	No	None	Pipeline metadata. Used to identify a local MEP or MIP instance.
MPLS L2 Port	32	Yes		Pipeline metadata. Used to identify an MPLS-TP pseudo wire endpoint.
OVID	16	No	VLAN tag	Pipeline metadata. Outer VLAN id, which has been popped in the VLAN Flow Table, to enable double tag matching in the VLAN 1 Flow Table.
MPLS_DATA_FIRST_NIBBLE	4	No	ETH-TYPE=0x8847 and MPLS_BOS=1	Determine if data (0000b) or control (0001b)

Action	Bits	Masked	Pre-requisite	Description
MPLS_ACH_CHANNEL	16	No	MPLS control frame determined by the value of MPLS_DATA_FIRST_NIBBLE	
MPLS_TTL	8	No	ETH-TYPE=0x8847	
MPLS_NEXT_LABEL_IS_GAL	1	No	ETH-TYPE=0x8847	Pipeline metadata derived from the packet parser "peeking" at the next label.
OAM_Y1731_MDL	3	No	ETH-TYPE=0x8902	OAM PDU Maintenance Domain Level
OAM_Y1731_OPCODE	8	No	ETH-TYPE=0x8902	OAM PDU opcode
COLOR_ACTIONS_INDEX	32	No	None	Pipeline metadata. Used to identify an entry in the Color Based Actions Flow Table.
ACTSET_OUTPUT	32	No		Output port assigned by match table or group output action. Only used in egress tables. This is the same as the similarly named pipeline metadata match field expected in OpenFlow 1.5. Read only from Controller.
TxFCl	64	No		OAM Data Counter Table value set by Set-Counter-Fields. Pipeline metadata field sent to Network Protection App. Read only from Controller.
RxFCl	64	No		OAM Data Counter Table value set by Set-Counter-Fields. Pipeline metadata field sent to Network Protection App. Read only from Controller.

Action	Bits	Masked	Pre-requisite	Description
RxTIME	64	No		Timestamp value for current OAM PDU. Pipeline metadata field sent to Network Protection App. Read only from Controller.

The match fields are programmed using the following assignments:

```

/* OF-DPA Experimenter Match Field types */
enum ofdpa_match_exp_type {
    OFDPA_OXM_VRF = 1,
    OFDPA_OXM_TRAFFIC_CLASS = 2,
    OFDPA_OXM_COLOR = 3,
    OFDPA_OXM_DEI = 4,
    OFDPA_OXM_QOS_INDEX = 5,
    OFDPA_OXM_LMEP_ID = 6,
    OFDPA_OXM_MPLS_TTL = 7,
    OFDPA_OXM_MPLS_L2_Port = 8,
    OFDPA_OXM_OVID = 10,
    OFDPA_OXM_MPLS_DATA_FIRST_NIBBLE = 11,
    OFDPA_OXM_MPLS_ACH_CHANNEL = 12,
    OFDPA_OXM_MPLS_NEXT_LABEL_IS_GAL = 13,
    OFDPA_OXM_OAM_Y1731_MDL = 14,
    OFDPA_OXM_OAM_Y1731_OPCODE = 15,
    OFDPA_OXM_COLOR_ACTIONS_INDEX = 16,
    OFDPA_OXM_TXFCL = 17,
    OFDPA_OXM_RXFCL = 18,
    OFDPA_OXM_RX_TIMESTAMP = 19,
    OFDPA_OXM_ACTSET_OUTPUT = 42,
}

```

6.5.6 Color Set Meter Band

The experimenter definition to support OF-DPA 2.0 color set meter bands is given below. This is the only meter band type supported and replaces the corresponding fields in Meter modification messages.

```

/* OFPMT_EXPERIMENTER band for Color Set */

struct ofp_meter_band_experimenter_color_set {
    uint16_t    type;        /* OFPMBT_EXPERIMENTER (0xFFFF) */
    uint16_t    len;        /* Length in bytes of this band. */
    uint32_t    rate;       /* Rate for this band. */
    uint32_t    burst_size; /* Size of bursts. */
    uint32_t    experimenter; /* Experimenter Id: (00-00-10-18) */
    uint16_t    exp_type;   /* OFPMBT_COLOR_SET (3) */
    uint8_t     color;      /* New color, one of: Yellow(1), Red(2) */
    uint8_t     pad[5];     /* Align on 64-bit boundary */
};

OFP_ASSERT(sizeof(struct ofp_meter_band_experimenter_color_set) == 24);

```

APPENDIX A :REFERENCES

- [1] "Software Defined Networking Definition," Open Networking Foundation, <https://www.opennetworking.org/sdn-resources/sdn-definition>
- [2] OpenFlow 1.0 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.0.0.pdf>
- [3] OpenFlow 1.3 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.3.0.pdf>
- [4] OpenFlow 1.3.1 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.3.1.pdf>
- [5] OpenFlow 1.3.2 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.3.2.pdf>
- [6] OpenFlow 1.3.3 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.3.3.pdf>
- [7] OpenFlow 1.3.4 Specification, <https://www.opennetworking.org/images/stories/downloads/specification/openflow-spec-v1.3.4.pdf>
- [8] OpenFlow Table Type Patterns 1.0, <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow/OpenFlow%20Table%20Type%20Patterns%20v1.0.pdf>
- [9] OpenFlow™ Data Plane Abstraction (OF-DPA): Abstract Switch Specification, Version 1.0, February 21, 2014, https://github.com/Broadcom-Switch/of-dpa/blob/master/doc/OFDPA_OASS-ETP101-R.PDF
- [10] OpenFlow Management and Configuration Protocol 1.2 (OF-Config 1.2), <https://www.opennetworking.org/images/stories/downloads/sdn-resources/onf-specifications/openflow-config/of-config-1.2.pdf>
- [11] Pffaf, B., and Davie, B, "The Open vSwitch Database Management Protocol," RFC7047, December 2013

- [12] Katz, D., and Ward, D., "Bidirectional Forwarding Detection," RFC 5880, June, 2010
- [13] Aggarwal, R., Kompella, K., Nadeau, T., Swallow, G., "Bidirectional Forwarding Detection (BFD)for MPLS Label Switched Paths (LSPs)," RFC 5884, June, 2010
- [14] Allan, D., Swallow, G., and Drake, J., "Proactive Connectivity Verification, Continuity Check, and Remote Defect Indication for the MPLS Transport Profile," RFC 6428, November, 2011
- [15] International Telecommunications Union, "Operations, administration and maintenance mechanism for MPLS-TP in packet transport networks", Recommendation ITU-T G.8113.1/Y.1372.1, November, 2012
- [16] International Telecommunications Union, "Operations, administration and maintenance mechanism for MPLS-TP networks using the tools defined for MPLS", Recommendation ITU-T G.8113.2/Y.1372.2, November, 2012
- [17] International Telecommunications Union, "OAM functions and mechanisms for Ethernet based networks," Recommendation ITU-T Y.1731, November, 2013
- [18] "Ryu, a Software Defined Network Framework," <http://osrg.github.io/ryu/>
- [19] "OpenDaylight Home Page," <http://www.opendaylight.org/>
- [20] IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks – Bridge Port Extension, IEEE Std 802.1BR™-2012
- [21] VXLAN: A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks - <http://datatracker.ietf.org/doc/draft-mahalingam-dutt-dcops-vxlan>
- [22] Blake, S., Black, D., Carlson, M., Davies, E., Wang, Z., and Weiss, W., "An Architecture for Differentiated Services," RFC2475, December 1998
- [23] Heinanen, J. and Guerin, R., "A Single Rate Three Color Marker," RFC 2697, September, 1999
- [24] Heinanen, J. and Guerin, R., "A Two Rate Three Color Marker," RFC 2698, September, 1999
- [25] Hanks, S., Li, T., Farinacci, D. and P. Traina, "Generic Routing Encapsulation", RFC 1701, October 1994
- [26] Vigoureux, M., Ward, D., and Betts, M., "Requirements for Operations, Administration, and Maintenance (OAM) in MPLS Transport Networks," RFC 5860, May, 2010
- [27] IEEE Standard for Local and metropolitan area networks – Virtual Bridged Local Area Networks – Virtual Bridged Local Area Networks – Amendment 5: Connectivity Fault Management, IEEE Std 802.1ag™-2007
- [28] Open Networking Foundation Extensibility Working Group JIRA EXT-407, available to ONF member companies at <https://rs.opennetworking.org/bugs/browse/EXT-407>.
- [29] International Telecommunications Union, "Characteristics of Ethernet transport network equipment functional blocks", Recommendation ITU-T G.8021/Y.1341, May, 2012
- [30] Martini, L., Rosen, E., El-Aawar, N., and Heron, G., "Encapsulation Methods for Transport of Ethernet over MPLS Networks," RFC 4448, April 2006