


ArubaOS 7.2



User Guide

Copyright

© 2012 Aruba Networks, Inc. Aruba Networks trademarks include,  Aruba Networks®, Aruba Wireless Networks®, the registered Aruba the Mobile Edge Company logo, Aruba Mobility Management System®, Mobile Edge Architecture®, People Move. Networks Must Follow®, RFProtect®, Green Island®. All rights reserved. All other trademarks are the property of their respective owners.

Open Source Code

Certain Aruba products include Open Source software code developed by third parties, including software code subject to the GNU General Public License (GPL), GNU Lesser General Public License (LGPL), or other Open Source Licenses. The Open Source code used can be found at this site:

http://www.arubanetworks.com/open_source

Legal Notice

The use of Aruba Networks, Inc. switching platforms and software, by all individuals or corporations, to terminate other vendors' VPN client devices constitutes complete acceptance of liability by that individual or corporation for this action and indemnifies, in full, Aruba Networks, Inc. from any and all legal actions that might be taken against it with respect to infringement of copyright on behalf of those vendors.

Warranty

This hardware product is protected by the standard Aruba warranty of one year parts/labor. For more information, refer to the ARUBACARE SERVICE AND SUPPORT TERMS AND CONDITIONS.

Altering this device (such as painting it) voids the warranty.



www.arubanetworks.com

1344 Crossman Avenue
Sunnyvale, California 94089

Phone: 408.227.4500
Fax 408.227.4550

Chapter 1	About this Guide	19
	Audience.....	19
	Fundamentals.....	19
	WebUI	19
	CLI.....	19
	Related Documents	20
	Conventions.....	20
	Contacting Support	21
Chapter 2	System Overview	23
	System Information	23
	Factory Initial Configuration.....	24
	Spanning Tree Modes.....	24
	Auto Configuration.....	24
	Important Points to Remember.....	25
	Trace Options	25
	Understanding Interface Profiles	26
	Interface Numbering Convention	27
	Assigning an Interface Profile as an Access Port	27
	Assigning an Interface Profile as a Trunk.....	28
	Understanding Interface Group	28
	Configuring Interface Group	28
	Controller-ip.....	29
	Using the LCD	29
	LCD Management	29
	Using the LCD and USB Drive	30
	Upgrade an image	30
	Upload a pre-saved configuration	30
	LCD Functions with ArubaStack	30
	Disabling LCD Menu Functions	31
Chapter 3	Management Access.....	33
	Certificate Authentication Concepts.....	33
	Configuring Certificate Authentication	33
	In the CLI	33
	Public Key Authentication for SSH Access	33
	In the CLI	34
	Management Password Policy.....	34
	Defining a Management Password Policy	34
	Managing Certificates.....	35
	About Digital Certificates	36
	Obtaining a Server Certificate	36
	In the CLI	37
	Obtaining a Client Certificate	37
	Importing Certificates.....	37
	In the CLI	38

Viewing Certificate Information	38
Logging.....	38
Managing Files on the MAS.....	39
Transferring ArubaOS Image Files	40
In the WebUI	40
In the CLI	40
Backing Up and Restoring the Flash File System.....	40
Backup the Flash File System in the CLI.....	40
Restore the Flash File System in the WebUI	41
Restore the Flash File System in the CLI.....	41
Copying Log Files	41
In the WebUI.....	41
In the CLI	41
Copying Other Files	41
In the WebUI.....	41
In the CLI	42
Setting the System Clock.....	42
In the CLI	42
Clock Synchronization	42
Configuring NTP Authentication	42

Chapter 4	Profiles Management	45
	Profiles for Interfaces.....	45
	Profiles for VLANs.....	47
	Scope of the Profiles and Parameters.....	47
	Factory Initial vs Default vs Non-Default Profiles and Parameters	47
	Profiles and Parameters Assigned to the Interfaces and Groups.....	47
	AAA Profiles Assigned to the Interfaces, Groups, and VLANs	49
	Profiles and Parameters Assigned to the Port-Channel Members	50
	Creating a Profile.....	50
	Using the WebUI	50
	Using the CLI	50
	Viewing a Profile and its Parameters.....	51
	Using the CLI	51
	Displaying the List of Profiles Under Each Category.....	51
	Displaying the Parameters Assigned to Each Profile	52
	Applying and Activating a Profile.....	53
	Using the CLI	53
	Applying and Activating the Profiles for an Interface	53
	Applying and Activating the Profiles for an Interface Group	53
	Applying and Activating the Profiles for a Port-Channel	53
	Applying and Activating the Profiles for a VLAN	53
	Deleting a Profile	54
	Using the CLI	54
	Best Practices	54

Chapter 5	ArubaStack.....	55
	Important Points to Remember	55
	Stacking Topology.....	56
	Viewing the ArubaStack Information.....	56
	Dormant State	57
	Dynamic Election.....	57
	Configuring Priority	57
	Using the WebUI.....	58

Using the CLI	58
The Stacking Protocol.....	58
Auto Discovery	59
Primary Election	59
Election Anatomy	59
ArubaStack Pre-provisioning.....	59
Configuring ArubaStack Pre-provisioning	60
Using the WebUI.....	60
Using the CLI	60
ArubaStack Database.....	61
Removing an ArubaStack Database.....	61
Booting without an ArubaStack Database	61
Primary Switchover	62
ArubaStack Resiliency.....	62
Split Detect	63
Stack Join	63
Stack Merge—Dynamic Election	63
Stack Merge—Pre-Provisioning.....	64
Pre-provisioned and Dynamic ArubaStacks Merge	64
Pre-provisioned ArubaStacks Merge	65
Console Redirect.....	68
Management User Authentication	68
ArubaStack Member Replacement	69
Dynamic ArubaStack Configuration.....	69
Replacing a Linecard Member	69
Replacing a Secondary Member	71
Replacing a Primary Member	72
Preset ArubaStack Configuration	74
Replacing a Linecard Member	75
Replacing a Secondary Member	77
Replacing a Primary Member	79
Chapter 6 Ethernet Interfaces and PoE.....	83
Configuring the Management Port	83
Using the CLI	83
Sample Management Port Configuration	83
Gigabit Ethernet Network Interfaces	84
Time-Domain Reflectometer (TDR)	84
Configuring an Ethernet Interface.....	84
Using the CLI	84
Configuring Jumbo Frame Size.....	85
Verifying Jumbo Frame Size.....	85
Displaying Interface Counters and Statistics	85
Configuring an Interface Group	86
Using the CLI	86
Sample Interface Group Configuration.....	87
Verifying the Interface Group Configuration	87
Creating and Applying an Ethernet Link Profile to an Interface	89
Using the WebUI	90
Using the CLI	90
Ethernet Link Default Profile	90
Sample Ethernet Link Profile Configuration.....	90
Verifying Ethernet Link Profile Configuration	90
Ethernet Flow Control	91
Loopback Interfaces.....	91
Using the CLI	91

	Sample Loopback Interface Configuration	91
	Power Over Ethernet	92
	Power Management Modes	92
	Power Pools	92
	Mixed Mode PSUs	93
	PoE Priority	93
	PoE Guard-Band	93
	Configuring Power Over Ethernet	93
	Using the WebUI	93
	Using the CLI	94
	Sample PoE Configuration	94
	Creating and Applying a PoE Profile to an Interface	94
	Using the WebUI	94
	Using the CLI	94
	Sample PoE Profile Configuration	94
	Time Range Support for PoE	95
	PoE Factory-Initial and Default Profiles	95
	Monitoring Power-over-Ethernet	96
Chapter 7	Port-Channels	99
	Important Points to Remember	99
	Creating a Port-Channel	99
	Using the WebUI	100
	Using the CLI	100
	Default Enet-Link Profile for Port-Channels	100
	Sample Static Port-Channel Configuration	101
	Verifying the Port-Channel Configuration	101
	Link Aggregation Control Protocol	101
	LACP Port Modes	101
	LACP Session Timeout and Port Priority	102
	Creating and Applying a Dynamic Port-Channel Profile to an Interface	102
	Using the WebUI	102
	Using the CLI	102
	Sample Dynamic Port-Channel Configuration	103
	Verifying Port-Channel Configuration	103
	Verifying Port-Channel Neighbor Information	103
	Verifying Port-Channel Internal (Local) Information	103
	Verifying Port-Channel Counters Information	104
Chapter 8	Operations, Administration, and Maintenance	105
	Creating an OAM Profile	105
	Applying an OAM Profile	106
	Applying OAM to each Port Channel Member	106
	Related Show Commands	107
Chapter 9	VLANs	109
	VLANs Overview	109
	Creating VLANs	109
	Using the WebUI	109
	Using the CLI	109
	Sample VLAN Configuration	110
	Verifying VLAN Configuration	110
	Creating and Applying a Switching Profile to an Interface	111

	Using the WebUI	111
	Using the CLI	112
	Default Switching Profile	112
	Sample Access Port Configuration	112
	Verifying the Switching Profile Configuration for the Interface	113
	Sample Trunk Port Configuration	113
	Verifying the Trunk Configuration	113
	Managing the MAC Address Table	114
	Adding Static MAC Addresses	114
	Example Configuration	114
	Displaying the MAC Address Table	114
	Deleting the Static MACs	115
	Clearing the Learnt MACs	115
	Configuring the MAC Aging Time	115
	VLAN Profile	115
Chapter 10	GVRP	117
	GVRP Overview	117
	Enabling and Configuring GVRP Functionality	117
	Sample Configurations	118
Chapter 11	Link Layer Discovery Protocols	121
	Important Points to Remember	121
	LLDP	121
	Understanding LLDP	121
	LLDP Factory Initial and Default Profiles	122
	Configuring LLDP	122
	Configuring an LLDP Profile	122
	Applying LLDP Profile to an Interface	123
	Verifying LLDP Profile Configuration	123
	Monitoring LLDP	123
	Display LLDP Interface	123
	Display LLDP Interface <interface>	124
	Display LLDP Neighbor	124
	Display LLDP Neighbor Interface Detail	124
	Display LLDP Statistics	125
	Display LLDP Statistics Interface	125
	LLDP-MED	125
	Understanding LLDP-MED	125
	Configuring LLDP-MED	126
	LLDP-MED Usage	126
	Verifying the LLDP Profile Configuration to Check LLDP-MED Status	127
	Proprietary Link Layer Discovery Protocols	127
	Understanding Proprietary Link Layer Discovery Protocol	127
	CDP Receive Processing	127
	CDP Frame Information	127
	Configuring Propriety LLDP Receive Processing	128
	Verifying Propriety LLDP Receive Processing	128
	Monitoring the Proprietary Neighbor Discovery	129
Chapter 12	VoIP	131
	Voice VLANs	131
	Creating and Applying VoIP Profile to an Interface	132
	VoIP Auto-Discovery on Trusted Ports	132
	Enabling VoIP Auto-Discovery	132

	Verifying VoIP Mode Configuration	133
	Viewing Neighboring Phones.....	133
	VoIP Auto-discovery on Untrusted Ports	133
Chapter 13	MSTP	135
	Important Points to Remember	135
	Example MSTP Configuration	136
	Viewing Operational Information	137
	Loopguard and Rootguard	138
	Configuring Loopguard	138
	Configuring Rootguard.....	139
	Bridge Protocol Data Unit (BPDU) Guard.....	140
	Enabling and Configuring BPDU Guard Functionality	140
	Verifying the BPDU Guard Configuration:.....	141
	Sample Configuration	141
	Portfast	141
	Configuring Portfast	141
	Sample Topology and Configuration.....	142
Chapter 14	Rapid PVST+	153
	Important Points to Remember	153
	Configuring PVST+	153
	Configuring using the VLAN Profile.....	154
	Disable PVST+ on a VLAN.....	155
	Configuring using the Interface-based Profile	155
	Loopguard and Rootguard	155
	Configuring Loopguard	155
	Configuring Rootguard.....	156
	Verifying the Configuration.....	156
	Bridge Protocol Data Unit (BPDU) Guard.....	157
	Enabling and Configuring BPDU Guard Functionality	157
	Verifying the BPDU Guard Configuration:.....	157
	Sample Configuration	157
	Portfast	158
	Configuring Portfast	158
	Verify the Configuration.....	158
Chapter 15	Hot-Standby Link.....	159
	Important Point to Remember	159
	Configuration Steps.....	159
Chapter 16	Generic Router Encapsulation	161
	Overview	161
	Configuring an L2-GRE Tunnel.....	161
	Inter-tunnel flooding.....	161
	Understanding the VLAN Membership of Existing L2 GRE Tunnel.....	162
	Sample Configuration	163
Chapter 17	Layer 3 Routing.....	165
	Understanding Routed VLAN Interfaces	165
	Important Points to Remember.....	165
	Configuring Routed VLAN Interfaces	165
	Using the CLI	165

Sample Configuration	166
Verifying and Monitoring the VLAN Interfaces	166
Displaying VLAN Information	166
Displaying IP Interface Brief	167
Multinetting	167
Important Points to Remember	167
Configuring Secondary IP	167
Sample Configuration	167
Network Address Translation	168
IP Directed Broadcast	169
Configuring IP Directed Broadcast	169
Sample Configuration	169
Static Routes	170
Important Points to Remember	170
The Default Gateways	170
Configuring the Default Gateways and the Static Routes	170
Using the WebUI	170
Using the CLI	170
Sample Configuration	171
Verifying the IP Routes	171
Clearing the ARP Table	172
Route Configuration Limits	172
Route Metrics	172
Equal Cost Multipath	172
IP Prefix List	173

Chapter 18 DHCP Server & DHCP Relay 175

Important Points to Remember	175
Understanding DHCP Server and DHCP Relay	175
Configuring DHCP Server and DHCP Relay	176
Configuring DHCP Server	176
Configuring DHCP Relay	176
Applying DHCP Relay Profile to VLAN	177
Configuring a VLAN with a Relay Profile as DHCP Client	177
Points to Remember	177
Verifying DHCP Server and DHCP Relay	178
Verifying DHCP Relay Option 82 Logs	178
Network Log	178
System Log	178
Show Commands for IP DHCP	179
show interface-profile dhcp-relay-profile	179
show ip dhcp database	179
show ip dhcp binding	179
show ip dhcp statistics	180
show ip dhcp pool	180
show ip dhcp pool	180

Chapter 19 OSPFv2 183

OSPF Feature Overview	183
Key Features Supported by MAS	183
LSAs Originated by MAS	183
Configuring OSPF	184
Configuring OSPF	184
Configuring OSPF Area Types	184
Sample Configuration	185

	Configuring prefix-list with OSPF.....	185
	Sample Configuration.....	185
	Verifying the Configuration.....	185
	Enabling OSPF on a Loopback Interface.....	187
	OSPF MD5 Authentication	188
	Important Points to Remember.....	188
	Understanding OSPF MD5 Authentication	189
	Configuring OSPF MD5 Authentication	189
	Verifying OSPF MD5 Authentication	189
	Verifying OSPF MD5 Authentication Configuration from the Interface Profile.....	189
	Verifying the OSPF MD5 Authentication Configuration.....	190
	Verifying OSPF MD5 Authentication.....	190
Chapter 20	IPv6	191
	IPv6 Support for Mobility Access Switch	191
	Configure an IPv6 Interface Address	192
	Configure IPv6 Default Gateway	192
	Debug IPv6 Mobility Access Switch	192
Chapter 21	IGMP and PIM-SM	193
	Important Points to Remember	193
	Understanding IGMP and PIM-SM.....	193
	IGMP	193
	Basic IGMP Network Architecture	193
	PIM	194
	PIM Sparse Mode	194
	Configuring IGMP	194
	Configuring PIM Sparse Mode	194
	Configuring PIM-SM End to End.....	195
	Verifying PIM Sparse Mode	195
	Displaying PIM RPF Information.....	195
	Displaying PIM Neighbor Information.....	196
	Displaying PIM RP Information.....	196
	Displaying PIM Mroute Information	196
	Displaying PIM Statistical Information.....	196
Chapter 22	IGMP Snooping	199
	Important Points to Remember	199
	Multicast Support with IGMP Snooping	199
	Snooping Report and Query Support	200
	Mrouter	200
	Configuring a Static Mrouter Port	200
	Example Configuration	201
	Creating and Applying an IGMP Snooping Profile to a VLAN	201
	Using the CLI	201
	Sample Configuration	201
	IGMP Snooping Factory Initial and the Default Profiles	202
	Verifying IGMP Snooping Configuration.....	202
	Monitoring IGMP Snooping	203
	Clearing IGMP Counters and Membership	204
	Enabling IGMP Snooping Trace Options	204
Chapter 23	MLD Snooping	205

Important Points to Remember	205
Understanding MLD Snooping	205
Configuring MLD Snooping	205
Configuring MLD Snooping.....	205
Deleting an Mrouter Port on a VLAN.....	206
Verifying MLD Snooping	206
Verifying the MLD Snooping Profile.....	206
Verifying the Static and Dynamic Mrouter Port for MLD Snooping.....	206
Verifying the MLD Snooping Mrouter Detail	207
Verifying the Two Mrouter Entries with the Same IP Address.....	207
Verifying MLD Snooping Member Ports	208
Verifying the MLD Group	209
Verifying the MLD Snooping Group Count.....	209
Verifying the MLD Snooping Statistics	209
List of MLD Snooping Commands and Sample Outputs	210
Show Counters	210
Show VLAN Counters	211
Show MLD Mrouter Ports	211
Show MLD Mrouter Ports Detail.....	211
Show MLD Router Ports Per VLAN	211
.....Show Detected MLD Multicast Addresses	211
Show Detected MLD Multicast Addresses Per VLAN	212
Show Detected MLD Multicast Membership Information	212
Show Detected MLD Multicast Membership Information (Detailed Version)	212
Show Detected MLD Multicast Membership Information Per VLAN....	213
Show MLD-Snooping Profile	213
Show List of MLD-Snooping Profiles	213
.....Show List of References for MLD-Snooping Profile	213

Chapter 24 Port Security 215

Port Security Overview	215
Router Advertisement Guard	215
Points to remember	215
DHCP Trust	215
Loop Protect	216
Points to Remember.....	216
MAC Limit	216
Enabling Port Security Functionality.....	216
Enabling RA Guard Functionality	216
Enabling DHCP Trust Functionality.....	216
Enabling Loop Protect Functionality	217
Enabling MAC Limit Functionality	217
Attaching Port Security Profile to Interface.....	217
Viewing Port Errors	218
Recovering Ports Manually	218
Sample Configurations	218

Chapter 25 Storm Control..... 219

Important Points to Remember	219
Configuration Steps.....	219

Chapter 26 Access Control List 221

Types of ACLs	221
Port ACLs (PACLs)	221

Configuring the ACLs	222
Ethertype ACL	222
MAC ACL	222
Standard ACL	222
Extended ACL	222
Stateless ACL	223
Verifying the ACL configuration	224

Chapter 27 **Quality of Service..... 225**

QoS Concepts	225
Overview	225
Profiles and Queues	225
Classification	226
Trust Mode	226
Untrusted Mode	226
Profile	226
Policing	227
Configuring QoS	227
Configuring QoS Trust Mode	227
Configuring QoS-Profile under an Interface	228
Configuring QoS-Profile under a Stateless ACL	228
Configuring QoS-Profile under a User-Role	228
Configuring Policer under Policer-Profile	228
Configuring Policer-Profile under an Interface	228
Configuring Policer-Profile under a Stateless ACL	228
Configuring Policer-Profile under a User-role	229

Chapter 28 **Authentication Servers..... 231**

Important Points to Remember	231
Server and Server Group Concepts	231
Configuring Authentication Servers	232
Radius Server Authentication	232
RADIUS Server Username/Password Authentication	233
In the CLI	233
RADIUS Server Authentication with VSA	233
RADIUS Server Authentication with Server-Derivation Rule	233
In the CLI	233
Disabling Authentication of Local Management User Accounts	234
In the CLI	234
Verifying the configuration	234
Resetting the Admin or Enable Password	234
Bypassing the Enable Password Prompt	235
Setting an Administrator Session Timeout	235
Setting a CLI Session Timeout	236
Setting a WebUI Session Timeout	236
Configuring a RADIUS Server	236
Using the CLI	237
RADIUS Server Authentication Codes	237
RADIUS Change of Authorization	238
.....	238
Configuring an LDAP Server	238
Using the CLI	239
Configuring a TACACS+ Server	240
Using the CLI	240
Internal Database Concepts	240
Configuring the Internal Database	240

Using the CLI	241
Managing Internal Database Files	241
Using the CLI	241
Internal Database Utilities	242
Server Group Concepts	242
Configuring Server Groups	242
Using the CLI	242
Configuring Server List Order and Fail-Through	242
Using the CLI	242
Configuring Dynamic Server Selection	243
Using the CLI	244
Trimming Domain Information from Requests	244
Using the CLI	244
Configuring Server-Derivation Rules	244
Using the CLI	245
Configuring a Role Derivation Rule for the Internal Database	245
Using the CLI	245
Assigning Server Groups	245
User Authentication	246
Management Authentication	246
Using the CLI	246
Radius Accounting	246
Understanding Radius Accounting	246
Configuring RADIUS Accounting	248
TACACS+ Accounting	249
Authentication Timers	249
Using the CLI	250

Chapter 29

AAA Authentication	251
AAA Authentication Profile	251
Authentication Profile Concepts	251
Initial Role	251
MAC Auth Profile	251
MAC Default Role	251
802.1x Auth Profile	251
802.1x Default Role	251
User Derivation Rules	252
Authentication Schemes	252
MAC-Based Authentication	252
802.1x Authentication	252
Role/VLAN Derivation	252
Role Assignment Precedence	253
VLAN Assignment Precedence:	254
Current Limitations	254
Layer 2 Entry	254
Layer 3 Entry	255
User Roles	255
Authentication Roles	255
Access List	255
VLAN	255
User Derivation Rules	256
Configuring User Derivation Rules	256
Displaying User Derivation Rules	256
Configuring Authentication End to End	256
Configuring Authentication Server	256
Configuring a RADIUS Authentication Server	256
Configuring an Authentication Server Group	257

	Configuring a Server for Fail-Over with the Internal Database	257
	Configuring Internal Server Under a Server-Group	257
	Configuring a User Account with the Internal Database	257
	Displaying the Internal Database.....	257
	Maintaining Existing Accounts with the Internal Database	258
	Configuring Management Authentication	258
	Configuring AAA Timers.....	258
Chapter 30	Roles and Policies	259
	Policies	259
	Access Control Lists (ACLs).....	259
	Creating an Access Policy	260
	In the CLI	261
	Creating a Network Service Alias.....	261
	In the CLI	261
	User Roles	261
	Creating a User Role.....	262
	In the CLI	262
	User Role Assignments	262
	User Role in AAA Profile.....	263
	In the CLI	263
	User-Derived Roles or VLANs.....	263
	Configure a User-derived Role or VLAN in the CLI	263
	Default Role for Authentication Method.....	263
	In the CLI	264
	Server-Derived Role.....	264
	Sample configuration.....	264
	VSA-Derived Role	264
Chapter 31	MAC-Based Authentication	265
	MAC-Based Authentication Concepts	265
	Configuring MAC-Based Authentication	265
	Configuring the MAC Authentication Profile	265
	Using the CLI	266
	Configuring Clients	266
	Using the CLI to configure clients in the internal database	266
Chapter 32	802.1x Authentication	267
	802.1x Authentication Concepts	267
	Authentication with a RADIUS Server.....	267
	Authentication Terminated on the Mobility Access Switch	268
	Configuring 802.1x Authentication	269
	270
	Configuring a Server Rule Using the CLI	270
	LDAP Servers.....	270
	Configuring Certificates with Auth Termination	271
	Using the CLI	271
	Configuring 802.1x Authentication with Machine Authentication	271
	Role Assignment with Machine Authentication Enabled	271
	Authentication with an 802.1x RADIUS Server	273
	Creating an Alias for the Internal Network	274
	Using the CLI	274
	Creating the Student Role and Policy.....	274
	Using the CLI	274
	Creating the Faculty Role and Policy	274

	Using the CLI	274
	Creating the Guest Role and Policy	274
	Using the CLI	274
	Configuring the RADIUS Authentication Server	275
	Configuring 802.1x Authentication Profile	275
	Using the CLI	275
	Configuring AAA Profile	275
Chapter 33	Captive Portal	277
	Captive Portal Overview	277
	Configuring Captive Portal Authentication	277
	Captive Portal Configuration Parameters	278
	Captive Portal Configuration Example	279
	Configuring Captive Portal via the CLI	279
	Configuring Captive Portal via the WebUI	280
	Personalizing the Captive Portal Page	281
	Creating Walled Garden Access	283
	Creating Walled Garden Access	283
	Using the CLI to create walled garden access	283
	Mobility Access Switch Server Certificate	284
Chapter 34	Tunneled Nodes	285
	Important Points to Remember	285
	Tunneled Nodes Overview	286
	Support for Tunneled Node Back-up Server	287
	Creating and Configuring Tunneled Node Profile	287
	Path MTU Discovery	288
	Verifying and Monitoring Tunneled Nodes	288
	Verifying and Monitoring the Tunneled Nodes on the Controller	288
Chapter 35	Aruba Instant AP Integration	291
	Aruba Instant Overview	291
	Supported Devices	291
	IAP Integration with the MAS	292
	IAP Integration Features	292
	Rogue AP Containment	292
	PoE Prioritization	292
	GVRP Integration	292
	Viewing the Blacklisted MAC Address of the Rogue APs	292
	Recovering Ports Manually	293
Chapter 36	Aruba AirGroup Integration	295
	Overview	295
	Configuring mDNS packet forwarding	295
	Inter-tunnel flooding	296
	Sample Configuration	296
Chapter 37	Aruba ClearPass Policy Manager Integration	299
	Introduction	299
	Important Points to Remember	300
	Enabling Downloadable Role on MAS	300

	Using the WebUI	300
	Using the CLI	300
	Sample Configuration	301
	CPPM Server Configuration	301
	Adding a Device	301
	Adding Enforcement Profile	301
	Adding Enforcement Policy	303
	Adding Services	305
	MAS Configuration	307
	Configuring CPPM Server on MAS	307
	Configuring Server Group to include CPPM Server	307
	Configuring 802.1X Profile	307
	Configuring AAA Profile	307
	Show AAA Profile	307
Chapter 38	Virtual Private Networks	309
	Planning a Site-to-Site VPN Configuration	309
	Selecting an IKE protocol	309
	Supported IKE Modes	310
	VPN Topologies	310
	Configuring VPN	310
	Configuration Examples	312
	Main-Mode	312
	Aggressive-Mode with Tunnelled Node over VPN	312
	Static Route Support for VPN	313
Chapter 39	Port Mirroring	315
	Important Points to Remember	315
	The Source Port	315
	The Destination Port	315
	Mirroring Sampled Ratio	316
	Creating and Applying a Mirroring Profile to an Interface	316
	Using the CLI	316
	Sample Configuration	316
	Verifying Port Mirroring Configuration	316
Chapter 40	Remote Monitoring (RMON)	319
	Remote Monitoring (RMON) Overview	319
	Enabling RMON Service	319
	Configuring RMON Parameters	320
	Configuring the Alarm	320
	Configuring the Alarm Profile	320
	Configuring Ethernet Statistics Index	321
	Configuring History Group	321
	Configuring Event Entry	322
	Viewing RMON Active Configuration	322
	Viewing RMON Configuration	324
Chapter 41	MIB and SNMP	325
	SNMP Parameters for Mobility Access Switch	325
	Configuring SNMPv1/v2c Parameters	326
	Example	326
	Configuring SNMPv3 Parameters	326

Example.....	327
Viewing SNMP Configuration Parameters	327
Supported Standard MIBs	327
.....	330
Supported Enterprise MIBs.....	331
Supported Standard Traps	331
Supported Enterprise Traps.....	332

This guide describes the instructions and examples for configuring the ArubaOS Mobility Access Switch (Mobility Access Switch). This chapter covers:

- “Audience” on page 19
- “Fundamentals” on page 19
- “Related Documents” on page 20
- “Conventions” on page 20
- “Contacting Support” on page 21

Audience

This is intended for system administrators responsible for accessing networking infrastructures and assumes you are knowledgeable in Layer 2 and Layer 3 networking technologies.

Fundamentals

Throughout this document references are made to the Mobility Access Switch and configuring via the WebUI or command line interface (CLI).

WebUI

The WebUI is accessible through a standard Web browser from a remote management console or workstation. The WebUI includes a Quick Setup wizard that steps you through tasks that includes:

- Basic Information—Specify device name, domain name, password, and date and time
- Management—Specify switch management options, VLAN assignment, and static or DHCP IP address assignment
- Summary page with your settings and the ability to display your settings in a separate window for printing or saving.

The WebUI also includes post-setup Monitoring, Configuration, and Maintenance screens.

CLI

The CLI is a text-based interface accessible from a local console connected to the serial port on the S3500 or through a Telnet or Secure Shell (SSH) session.



By default, you access the CLI from the serial port or from an SSH session. You must explicitly enable Telnet on your Mobility Access Switch in order to access the CLI via a Telnet session.

When entering commands remember that:

- commands are not case sensitive
- the space bar will complete your partial keyword
- the backspace key will erase your entry one letter at a time
- the question mark (?) will list available commands and options

Related Documents

The following items are part of the complete documentation for the Aruba user-centric network:

- *Aruba S3500 Series Mobility Access Switch Installation Guide*
- *Aruba S2500 Series Mobility Access Switch Installation Guide*
- *ArubaOS Mobility Access Switch Command Line Reference Guide*
- *ArubaOS Mobility Access Switch Quick Start Guide*
- *Release Notes*

Conventions

The following conventions are used throughout this manual to emphasize important concepts:

Type Style	Description
<i>Italics</i>	This style is used to emphasize important terms and to mark the titles of books.
System items	This fixed-width font depicts the following: <ul style="list-style-type: none">• Sample screen output• System prompts• Filenames, software devices, and specific commands when mentioned in the text
Commands	In the command examples, this bold font depicts text that you must type exactly as shown.
<Arguments>	In the command examples, italicized text within angle brackets represents items that you should replace with information appropriate to your specific situation. For example: # send <text message> In this example, you would type “send” at the system prompt exactly as shown, followed by the text of the message you wish to send. Do not type the angle brackets.
[Optional]	In the command examples, items enclosed in brackets are optional. Do not type the brackets.
{Item A Item B}	In the command examples, items within curled braces and separated by a vertical bar represent the available choices. Enter only one choice. Do not type the braces or bars.

The following informational icons are used throughout this guide:



Indicates helpful suggestions, pertinent information, and important things to remember.



Indicates a risk of damage to your hardware or loss of data.



Indicates a risk of personal injury or death.

Contacting Support

Main Site	arubanetworks.com
Support Site	support.arubanetworks.com
North American Telephone	1-800-943-4526 (Toll Free) 1-408-754-1200
International Telephones	arubanetworks.com/support-services/aruba-support-program/contact-support/
Software Licensing Site	licensing.arubanetworks.com/login.php
Wireless Security Incident Response Team (WSIRT)	arubanetworks.com/support/wsirt.php
Support Emails	
Americas and APAC	support@arubanetworks.com
EMEA	emea_support@arubanetworks.com
WSIRT Email Please email details of any security problem found in an Aruba product.	wsirt@arubanetworks.com

This system overview is an introduction to the feature rich ArubaOS Mobility Access Switch and introduces functionality that is presented in greater detail in the rest of this document. This overview covers:

- “System Information” on page 23
- “Factory Initial Configuration” on page 24
- “Auto Configuration” on page 24
- “Trace Options” on page 25
- “Understanding Interface Profiles” on page 26
- “Understanding Interface Group” on page 28
- “Controller-ip” on page 29
- “Using the LCD” on page 29

System Information

- The Mobility Access Switch can be deployed as a standalone switch, an Aruba Stack (a group of up to 8 Mobility Access Switches) or with tunneled-node ports going back to a Mobility Controller. The Mobility Access Switch also supports role-based user access to provide enhanced endpoint security.

An ArubaStack is a set of interconnected Mobility Access Switches (up to 8) using stacking ports to form a stack. A stacking port is a physical port provisioned to run the stacking protocol. In factory default settings for Mobility Access Switches, 10 Gigabit uplink ports 2 and 3 are pre-provisioned to be stacking ports. Once a port is provisioned for stacking, it is no longer available to be managed as a network port. A stacking port can only be connected to other Mobility Access Switches running the Aruba Stacking Protocol (ASP).

- The default username/password to log in to the Mobility Access Switch is admin/admin123.
- The member interfaces belonging to the port-channel do not need to be contiguous; however, they must have the same physical parameters (speed/duplex).
- The maximum number of port-channel for the system is 64 (ID 0-63). Each port-channel may have up to 8 member interfaces.
- Port-channel interface(s) belong to VLAN 1 by default. If additional configurations are required, (such as configuring port-channel as a trunk) complete the port-channel configurations first, then apply the configurations to the port-channel.

Factory Initial Configuration

The Mobility Access Switch is pre-loaded with a factory initial configuration. To view the initial factory setting, execute the show running configuration command with the initial factory option.

```
(host) #show running-config | include factory-initial
Building Configuration...
interface-profile poe-profile "poe-factory-initial"
interface-profile lldp-profile "lldp-factory-initial"
vlan-profile igmp-snooping-profile "igmp-snooping-factory-initial"
    igmp-snooping-profile "igmp-snooping-factory-initial"
    lldp-profile "lldp-factory-initial"
    poe-profile "poe-factory-initial"
```



By default, MSTP is enabled in the factory setting.

Spanning Tree Modes

The spanning tree mode is set to MSTP in factory default.

```
(host) #show running-config | begin spanning-tree
Building Configuration...
spanning-tree
    mode mstp
```

To change spanning tree modes, use the spanning tree mode command. Once you change the spanning tree mode, the new spanning tree is automatically applied to all configured VLANs, including default VLAN 1.

```
(host)(config) #spanning-tree mode ?
mstp                Multiple spanning tree mode
pvst                Per-Vlan rapid spanning tree mode
(host)(config) #spanning-tree mode pvst
(host)(config) #
```

To verify the current spanning tree mode:

```
(host)(config) #show spanning-tree-profile

spanning-tree
-----
Parameter          Value
-----
spanning-tree-mode  pvst
```

For more detailed information on spanning tree, see [Chapter 13, “MSTP” on page 135](#) and [Chapter 14, “Rapid PVST+” on page 153](#).

Auto Configuration

The ArubaOS Mobility Access Switch supports an auto configuration through the DHCP server downloading a configuration file from the TFTP server. This process begins automatically when a Mobility Access Switch, with a factory default configuration, boots up. Basically, if the VLAN 1 (default VLAN) interface is in Auto config mode, we parse additional options in the DHCP message to obtain the TFTP server address and the configuration file name/path.



You can use any network port in stand-alone or stacking environments.

Important Points to Remember

- This process remains active for ten minutes. If the Mobility Access Switch is idle for 10 minutes and auto configuration is not complete, you must manually configure the Mobility Access Switch.
- During the auto configuration process, DHCP messages without auto configuration parameters are ignored.
- If quick-setup mode (WebUI or CLI) is started, auto configuration is disabled. If quick-setup mode is cancelled at any point, auto configuration remains disabled.
- Additionally, auto configuration is disabled when you attempt to configure an IP address for the VLAN interface or enable DHCP-client on the VLAN interface.
- If you do not choose to enter quick-setup and auto configuration is not disabled, the Mobility Access Switch reboots when the configuration is downloaded.



These options are processed only if the device is running with the factory-initial configuration. Otherwise these options in the incoming DHCP message are ignored.

The two options expected in the DHCP message are:

- TFTP server address must be included in one or both of the following:
 - siaddr
 - option 150

If the server address is included in both, the siaddr takes precedence.

- Configuration file path must be included in one or both of the following:
 - boot file option
 - option 67

The siaddr and the boot file option are part of the BOOTP parameters section of the DHCP message.



If the server IP address is provided but a configuration file is not provided, the DHCP message is ignored.

When these options are honored and processed, the Mobility Access Switch downloads the new configuration file, compares it with the configuration file in use, and if they differ, the new file is copied as default.cfg and the Mobility Access Switch reboots automatically.

Failures to download the configuration are logged. Successful download of the configuration are also logged. Reboot generates a message indicating that new configuration is being loaded. A syslog message is also logged.

Trace Options

The tracing feature is important for debugging the sequence of events that occur inside a process or protocol, for example message processing, state machine transitions, configuration change events, or timer events.

You can enable or disable trace options for various modules such as mstp, lldp, igmp, ospf, pim, rmon, layer2-forwarding, interface-manager, chassis-manager, and stack-manager via the `traceoptions` command.

You can use the following command to enable or disable the traceoptions for various modules:

```
(host) (config) #traceoptions
(host) (traceoptions) #?
```

chassis-manager	Control chassis manager trace options
igmp	Control igmp trace options
igmp-snooping	Control igmp-snooping trace options
interface-manager	Interface manager trace options
layer2-forwarding	Control Layer2 Forwarding trace options
lldp	Control LLDP trace options
mstp	Control MSTP trace options
no	Delete Command
ospf	Control ospf trace options
pim	Control pim sparse mode trace options
rmon	rmon trace options
routing	Control layer3 manager trace options
stack-manager	Control stack-manager trace options

The following command displays the enabled trace options:

```
(host) #show trace ?
```

chassis-manager	Show the contents of chassis manager trace file
igmp	Show the contents of igmp trace file
igmp-snooping	Show the contents of igmp-snooping trace file
interface-manager	Show the contents of interface manager trace file
layer2-forwarding	Show the contents of layer2-forwarding trace file
lldp	Show the contents of lldp trace file
mstp	Show the contents of mstp trace file
ospf	Show the contents of ospf trace file
pim	Show the contents of pim trace file
rmon	Show the contents of RMON trace file
stack-manager	Show the contents of stack-manager trace file

The following is an example configuration:

```
(host) (traceoptions) #layer2-forwarding flags fdb learning vlan
(host) (traceoptions) #show trace layer2-forwarding 10
```

For a complete listing of trace options commands, see the [ArubaOS 7.1 Command Line Reference Guide](#).

Understanding Interface Profiles

There are instances when multiple interfaces share the same characteristics; for example, physical interface characteristics, type of switch interface, and/or VLAN ID. Interface profiles are used when the same configuration is defined on a profile and applied to multiple interfaces.

The parameters are defined in the functional profile(s) and the name of the profile is referenced on the interfaces. The interface profile is particularly useful when a change is required. The change can be made on the profile without having to update the individual interfaces. [Table 1](#) list the profiles and their functions.

Table 1 *Interface Profiles*

Profile Type	Description
dhcp-relay-profile	Configure a dhcp relay profile
enet-link-profile	Configure an Ethernet Link
gvrp-profile	Configure a GVRP profile
igmp-profile	Configure an Interface IGMP profile

Table 1 *Interface Profiles*

Profile Type	Description
lACP-profile	Configure an LACP
lldp-profile	Configure an LLDP Profile
mirroring-profile	Configure a Mirroring profile
mstp-profile	Configure an Interface MSTP
oam-profile	Configure an OAM profile.
ospf-profile	Configure an Interface OSPF profile
pim-profile	Configure an Interface PIM profile
poe-profile	Configure a Power over Ethernet profile
port-security-profile	Configure a Port Security profile
pvst-port-profile	Configure an Interface PVST bridge
switching-profile	Configure a switching profile
tunneled-node-profile	Configure a Tunneled Node Server profile
voip-profile	Configure a VOIP profile

Interface Numbering Convention

The Mobility Access Switch numbering convention is three separate numbers:

- First number denotes slot number; in stacking mode, the first number is the stack member identification.
- Second number denotes the networking module; where 0 indicates front-facing networking interfaces and 1 indicates uplink module interfaces.
- Third number denotes the individual interface/port number.

For example, the interface gigabitethernet 0/0/20 denotes the slot number zero (0), module 0 and port number 20. Note that interface/port numbering starts at 0.

Assigning an Interface Profile as an Access Port

To assign an interface as an access port belonging to a particular VLAN, configure the switching profile to reference the VLAN (for example VLAN 200). Then apply the switching profile to the interface itself (for example gigabitethernet 0/0/10).

Configuring switching-profile that references VLAN 200:

```
(host) (config) #interface-profile switching-profile vlan_200
(host) (switching profile "vlan_200") #access-vlan 200
```

Applying the switching-profile to the gigabitethernet 0/0/10 interface:

```
(host) (config) #interface gigabitethernet 0/0/10
(host) (gigabitethernet "0/0/10") #switching-profile vlan_200
(host) (gigabitethernet "0/0/10") #exit
```

Assigning an Interface Profile as a Trunk

Similar to configuring an interface as an access port, assigning an interface profile as a trunk uses the trunk mode:

```
(host) (config) #interface-profile switching-profile TRUNK_PORTS
(host) (switching profile "TRUNK_PORTS") #switchport-mode trunk
```

Applying the switching-profile to the gigabitethernet 0/0/11 interface:

```
(host) (config) #interface gigabitethernet 0/0/11
(host) (gigabitethernet "0/0/11") #switching-profile TRUNK_PORTS
```

Native VLAN setting:

```
(host) (config) #interface-profile switching-profile TRUNK_PORTS
(host) (switching profile "TRUNK_PORTS") #native-vlan 100
```

By default, a trunk port allows all VLANs to be transported. This can be changed if necessary via the trunk parameter in the switching-profile:

```
(host) (config) #interface-profile switching-profile TRUNK_PORTS
(host) (switching profile "TRUNK_PORTS") #trunk allowed vlan all
```

Understanding Interface Group

It is often time consuming and tedious to configure multiple interfaces, which share the same configuration, via the command line. These interface can be grouped together so that any interface within the group can share the same configuration. When an interface is a member of an interface group, applying a specific profile to the interface will take precedence over interface group.

Configuring Interface Group

Define a group, for example First_Floor, which will contain the interfaces that share the same configuration. Apply valid interfaces members in ascending order; that is, from 0/0/0 through 0/0/30, and 0/0/32:

```
(host) (config) #interface-group gigabitethernet FIRST_FLOOR
(host) (gigabitethernet "FIRST_FLOOR") #apply-to 0/0/0-0/0/30,0/0/32
```

Notice there is no space in the list of interfaces.

Additionally, You can add or remove individual ports or ranges of ports without disrupting the existing port list using the following commands:

```
(host) (gigabitethernet "FIRST_FLOOR") #apply-to [add | remove] <interface-list>
```

Apply the switching-profile to the interface group:

```
(host) (gigabitethernet "FIRST_FLOOR") #switching-profile ACCESS_100
```

Verify your configuration or interface group using the **show interface-group-config** command.

```
(host) #show interface-group-config gigabitethernet FIRST_FLOOR

gigabitethernet "FIRST_FLOOR"
-----
Parameter                               Value
-----
Interface range members                 0/0/0-0/0/30,0/0/32
...
```

Controller-ip

The Mobility Access Switch automatically chooses the loopback IP or the first VLAN IP address as the controller-ip (also known as “Switch-IP”) address during the initial boot. If loopback does not exist, then the Mobility Access Switch automatically chooses the first VLAN IP as the controller-ip.

Aruba best practices recommends configuring the controller-ip to the loopback interface when using Ethernet and Mobility Access Switch functionalities.



If first the VLAN is automatically chosen (or configured) as the controller-id and the VLAN has no active member, then the controller-ip will be unreachable.

1. Set the loopback interface (0 in the example) address and mask:

```
(host)(config) #interface loopback 0
(host)(loopback "0") #ip address 10.10.10.1
```

2. Set the controller-ip loopback to interface 0.

```
(host)(config) #ip-profile
(host)(ip-profile) #controller-ip loopback 0
```

3. Verify your configuration with the **show switch ip** command.

```
(host)(loopback "0") #show switch ip

Switch IP Address: 10.10.10.1
Switch IP is from Loopback Interface: 0

(host)(loopback "0") #
```

Using the LCD

The Mobility Access Switch LCD panel is located at the front panel. The LCD displays:

- Boot status
- Hostname
- Alarm
- Interface LED modes: Admin, Speed/Duplex, PoE
- ArubaOS version
- Power supply, Fan status

LCD Management

In addition to displaying current status, LCD panel supports a user-interactive maintenance mode:

- ArubaOS software image upgrade
- Configuration file upload
- Erase configuration (write erase all)
- Factory default setting (restore factory-default stacking)
- Media (external USB) eject
- System reboot (reload)
- System Halt (halt)
- GUI Quick Setup

Using the LCD and USB Drive

You can upgrade your image or upload your pre-saved configuration by using your USB drive and your LCD commands.

Upgrade an image

1. Copy a new MAS image onto your USB drive (Windows on a PC) into a directory named **/arubaimage**.
2. Insert your USB drive into the Mobility Access Switch's USB slot. Wait for 30 seconds for MAS to mount the USB.
3. Navigate to **Upgrade Image** in the LCD's **Maintenance** menu. Select **partition** and confirm the upgrade (Y/N) and then wait for Mobility Access Switch to copy the image from USB to the system partition.
4. Execute a system reboot either from the LCD menu or from the command line to complete the upgrade.

Upload a pre-saved configuration

1. Copy your pre-saved configuration and name the copied file **aruba_usb.cfg**.
2. Move your pre-saved configuration file onto your USB drive (Windows on a PC) into a directory named **/arubaimage**.
3. Insert your USB drive into the Mobility Access Switch's USB slot. Wait for 30 seconds for MAS to mount the USB.
4. Navigate to the **Upload Config** in the LCD's Maintenance menu. Confirm the upload (Y/N) and then wait for the upload to complete.
5. Execute a system reboot either from the LCD menu or from the command line to reload from the uploaded configuration.

For detailed upgrade and upload instruction, see the *Upgrade Chapter in the Release Notes*.

LCD Functions with ArubaStack

Table 2 list the LED Stack mode and Maintenance mode along with each function. Some functions can be executed from any member in the ArubaStack (Primary, Secondary, or Line Card) to affect just that member. Other function are executed from the Primary only but affect all members of the ArubaStack. For example, system reboot can be executed on a member only to reboot just that member. Or, you can be executed system reboot on the Primary to reboot all members of the ArubaStack.

Table 2 list the LED Stack mode and Maintenance mode along with each function. Some functions can be executed from any member in the ArubaStack (Primary, Secondary, or Line Card) to affect just that member. Other function are executed from the Primary only but affect all members of the ArubaStack. For example, system reboot can be executed on a member only to reboot just that member. Or, you can execute system reboot on the Primary to reboot all members of the ArubaStack.

Table 2 LCD Functions Over Stacking

Mode	Any Stack Member (affects only local member)	Primary Only (affects all stack members)
LED Mode	Yes	
Status (display)		
Stack	Yes	

Table 2 LCD Functions Over Stacking

Mode	Any Stack Member (affects only local member)	Primary Only (affects all stack members)
AOS Version	Yes	
PS Status	Yes	
Fan Tray	Yes	
Maintenance		
Upgrade Image		Yes
Upload Configuration		Yes
Erase Config		Yes
Media Eject		Yes
Factory Default	Yes	
System Reboot	Yes	Yes
System Halt	Yes	Yes

Disabling LCD Menu Functions

For security purpose, you can disable all LCD menu functions by disabling the entire menu functionality using the following command:

```
host) (config) #lcd-menu
(host) (lcd-menu) #disable menu
```

To prevent inadvertent menu changes, you can disable LCD individual menu function using the following commands:

```
(host) (lcd-menu) #disable menu maintenance ?
erase-config Disable config erase menu
factory-default Disable factory default menu
gui-quick-setup Disable quick setup menu on LCD
media-eject Disable media eject menu on LCD
system-halt Disable system halt menu on LCD
system-reboot Disable system reboot menu on LCD
upload-config Disable config upload menu on LCD
upgrade-image Disable image upgrade menu on LCD
```

To display the current LCD functionality from the command line, use the following command:

```
(host) (config) #show lcd-menu

lcd-menu
-----
Menu Value
-----
menu maintenance upgrade-image partition0 enabled
menu maintenance upgrade-image partition1 enabled
menu maintenance system-reboot reboot-stack enabled
```

```
menu maintenance system-reboot reboot-local enabled
menu maintenance system-halt halt-stack enabled
menu maintenance system-halt halt-local enabled
menu maintenance upgrade-image enabled
menu maintenance upload-config enabled
menu maintenance erase-config enabled
menu maintenance factory-default enabled
menu maintenance media-eject enabled
menu maintenance system-reboot enabled
menu maintenance system-halt enabled
menu maintenance gui-quick-setup enabled
menu maintenance enabled
menu enabled
```


This chapter describes management access and tasks. It contains the following topics:

- “Certificate Authentication Concepts” on page 33
- “Public Key Authentication for SSH Access” on page 33
- “Management Password Policy” on page 34
- “Managing Certificates” on page 35
- “Logging” on page 38
- “Setting the System Clock” on page 42

Certificate Authentication Concepts

The Mobility Access Switch supports client certificate authentication for users accessing the Mobility Access Switch using the CLI. (The default is for username/password authentication.) You can use client certificate authentication only, or client certificate authentication with username/password (if certificate authentication fails, the user can log in with a configured username and password).



Each Mobility Access Switch can support a maximum of ten management users.

Configuring Certificate Authentication

To use client certificate authentication, you must do the following:

1. Obtain a client certificate and import the certificate into the Mobility Access Switch. Obtaining and importing a client certificate is described in “Managing Certificates” on page 35.
2. Configure certificate authentication for WebUI management. You can optionally also select username/password authentication.
3. Configure a user with a management role. Specify the client certificate for authentication of the user.

In the CLI

```
web-server
  mgmt-auth certificate
  switch-cert <certificate>
mgmt-user webui-cacert <ca> serial <number> <username> < role>
```

Public Key Authentication for SSH Access

The Mobility Access Switch supports public key authentication of users accessing the Mobility Access Switch using SSH. (The default is for username/password authentication.) When you import an X.509 client certificate into the Mobility Access Switch, the certificate is converted to SSH-RSA keys. When you enable public key authentication for SSH, the Mobility Access Switch validates the client’s credentials with the imported public keys. You can specify public key authentication only, or public key authentication with username/password (if the public key authentication fails, the user can login with a configured username and password).

To use public key authentication, you must do the following:

1. Import the X.509 client certificate into the Mobility Access Switch using the WebUI, as described in [“Importing Certificates” on page 37](#).
2. Configure SSH for client public key authentication. You can optionally also select username/password authentication.
3. Configure the username, role and client certificate.

In the CLI

```
ssh mgmt-auth public-key [username/password]
mgmt-user ssh-pubkey client-cert <certificate> <username> <role>
```

Management Password Policy

By default, the password for a new management user has no requirements other than a minimum length of 6 alphanumeric or special characters. However, if your company enforces a best practices password policy for management users with root access to network equipment, you may want to configure a password policy that sets requirements for management user passwords.

Defining a Management Password Policy

To define specific management password policy settings through the CLI, complete the following steps:

The table below describes the characters allowed in a management user password. The disallowed characters cannot be used by any management user password, even if the password policy is disabled.

Table 1 *Allowed Characters in a Management User Password*

Allowed Characters	Disallowed Characters
exclamation point: !	Parenthesis: ()
underscore: _	apostrophe: '
at symbol: @	semi-colon: ;
pound sign: #	dash: -
dollar sign: \$	equals sign: =
percent sign: %	slash: /
caret: ^	question mark: ?
ampersand: &	
star: *	
greater and less than symbols: < >	
curled braces: { }	
straight braces: []	
colon :	

Table 1 *Allowed Characters in a Management User Password*

Allowed Characters	Disallowed Characters
period: .	
pipe:	
plus sign: +	
tilde: ~	
comma: ,	
accent mark: `	

In the CLI

```
aaa password-policy mgmt
  enable
  no
  password-lock-out
  password-lock-out-time
  password-max-character-repeat.
  password-min-digit
  password-min-length
  password-min-lowercase-characters
  password-min-special-character
  password-min-uppercase-characters
  password-not-username
```

Managing Certificates

This section contains the following sections:

- [“About Digital Certificates” on page 36](#)
- [“Obtaining a Server Certificate” on page 36](#)
- [“Obtaining a Client Certificate” on page 37](#)
- [“Importing Certificates” on page 37](#)
- [“Viewing Certificate Information” on page 38](#)

The Aruba Mobility Access Switch is designed to provide secure services through the use of digital certificates. Certificates provide security when authenticating users and computers and eliminate the need for less secure password-based authentication.

There is a *default* server certificate installed in the Mobility Access Switch to demonstrate the authentication of the Mobility Access Switch for WebUI management access. However, this certificate does not guarantee security in production networks. Aruba *strongly* recommends that you replace the default certificate with a custom certificate issued for your site or domain by a trusted Certificate Authority (CA). This section describes how to generate a Certificate Signing Request (CSR) to submit to a CA and how to import the signed certificate received from the CA into the Mobility Access Switch.

The Mobility Access Switch supports client authentication using digital certificates for specific user-centric network services, such as AAA FastConnect. Each service can employ different sets of client and server certificates.

During certificate-based authentication, the Mobility Access Switch provides its server certificate to the client for authentication. After validating the Mobility Access Switch's server certificate, the client presents its own certificate to the Mobility Access Switch for authentication. After validating the client's certificate, the Mobility Access Switch can check the user name in the certificate with the configured authentication server (this action is optional and configurable).

About Digital Certificates

Clients and the servers to which they connect may hold authentication certificates that validate their identities. When a client connects to a server for the first time, or the first time since its previous certificate has expired or been revoked, the server requests that the client transmit its authentication certificate. The client's certificate is then verified against the CA which issued it. Clients can also request and verify the server's authentication certificate. For some applications, such as 802.1x authentication, clients do not need to validate the server certificate for the authentication to function.

Digital certificates are issued by a CA which can be either a commercial, third-party company or a private CA controlled by your organization. The CA is trusted to authenticate the owner of the certificate before issuing a certificate. A CA-signed certificate guarantees the identity of the certificate holder. This is done by comparing the digital signature on a client or server certificate to the signature on the certificate for the CA.

Digital certificates employ public key infrastructure (PKI), which requires a private-public key pair. A digital certificate is associated with a private key, known only to the certificate owner, and a public key. A certificate encrypted with a private key is decrypted with its public key. For example, party A encrypts its certificate with its private key and sends it to party B. Party B decrypts the certificate with party A's public key.

Obtaining a Server Certificate

Aruba strongly recommends that you replace the default server certificate in the Mobility Access Switch with a custom certificate issued for your site or domain by a trusted CA. To obtain a security certificate for the Mobility Access Switch from a CA:

1. Generate a Certificate Signing Request (CSR) on the Mobility Access Switch using the CLI.
2. Submit the CSR to a CA. Copy and paste the output of the CSR into an email and send it to the CA of your choice.
3. The CA returns a signed server certificate and the CA's certificate and public key.
4. Install the server certificate, as described in ["Importing Certificates" on page 37](#).



There can be only one outstanding CSR at a time in the **Mobility Access Switch**. Once you generate a CSR, you need to import the CA-signed certificate into the **Mobility Access Switch** before you can generate another CSR.

Table 2 *CSR Parameters*

Parameter	Description	Range
key	Length of private/public key.	1024/2048/4096
common_name	Typically, this is the host and domain name, as in <code>www.yourcompany.com</code> .	—
country	Two-letter ISO country code for the country in which your organization is located.	
state_or_province	State, province, region, or territory in which your organization is located.	

Table 2 CSR Parameters (Continued)

Parameter	Description	Range
city	City in which your organization is located.	
organization	Name of your organization.	
unit	Optional field to distinguish a department or other unit within your organization.	
email	Email address referenced in the CSR.	

In the CLI

1. Run the following command:

```
crypto pki csr {rsa key_len <key_val> || {ec curve-name <key_val>}} common-name <value>
country <country> state_or_province <state> city <city> organization <org> unit
<string> email <email>
```

2. Display the CSR output with the following command:

```
show crypto pki csr
```

3. Copy the CSR output between the BEGIN CERTIFICATE REQUEST and END CERTIFICATE REQUEST lines, paste it into an email and send it to the CA of your choice.

Obtaining a Client Certificate

You can use the CSR generated on the Mobility Access Switch to obtain a certificate for a client. However, since there may be a large number of clients in a network, you typically obtain client certificates from a corporate CA server. For example, in a browser window, enter `http://<ipaddr>/crtserv`, where <ipaddr> is the IP address of the CA server.

Importing Certificates

Use the WebUI or the CLI to import certificates into the Mobility Access Switch.



You cannot export certificates from the Mobility Access Switch.

You can import the following types of certificates into the Mobility Access Switch:

- Server certificate signed by a trusted CA. This includes a public and private key pair.
- CA certificate used to validate other server or client certificates. This includes only the public key for the certificate.
- Client certificate and client's public key. (The public key is used for applications such as SSH which does not support X509 certificates and requires the public key to verify an allowed certificate.)

Certificates can be in the following formats:

- X509 PEM unencrypted
- X509 PEM encrypted with a key
- DER
- PKCS7 encrypted
- PKCS12 encrypted

In the CLI

Use the following command to import CSR certificates:

```
crypto pki-import {der|pem|pfx|pkcs12|pkcs7} {PublicCert|ServerCert|TrustedCA} <name>
```

The following example imports a server certificate named **cert_20** in DER format:

```
crypto pki-import der ServerCert cert_20
```

Viewing Certificate Information

In the WebUI, the Certificate Lists section of the page lists the certificates that are currently installed in the Mobility Access Switch. Click **View** to display the contents of a certificate.

To view the contents of a certificate with the CLI, use the following commands:

Table 3 Certificate Show Commands

Command	Description
show crypto-local pki trustedCA [<name>][<attribute>]	Displays the contents of a trusted CA certificate. If a name is not specified, all CA certificates imported into the Mobility Access Switch are displayed. If name and attribute are specified, then only the attribute in the certificate are displayed. Attributes can be CN, validity, serial-number, issuer, subject, public-key.
show crypto-local pki serverCert [<name>][<attribute>]	Displays the contents of a server certificate. If a name is not specified, all server certificates imported into the Mobility Access Switch are displayed.
show crypto-local pki publiccert [<name>][<attribute>]	Displays the contents of a public certificate. If a name is not specified, all public certificates imported into the Mobility Access Switch are displayed.



All certificates on Primary node get synchronized with Secondary node only. Line Cards will not have these certificates synchronized. However, the certificates will get synchronized to the node when increasing the priority of the Line Card to make it primary.

Logging

For each category or subcategory of message, you can set the logging level or severity level of the messages to be logged. [Table 4](#) list the logging levels..

Table 4 Logging Levels

Logging Level	Description
Emergency	System is unusable
Alerts	Immediate action is needed.
Critical	Any critical conditions.
Errors	Error conditions.
Warning	Warning messages.
Notifications	Normal but signification conditions.

Table 4 *Logging Levels*

Logging Level	Description
Informational	Messages of general interest to system users.
Debug	Messages containing information useful for debugging.

The default logging level for all categories is Warning. Within each logging level are several log types you can select.

- network
- security
- system
- user
- user debug

Managing Files on the MAS

You can transfer the following types of files between the MAS and an external server or host:

- ArubaOS image file
- A specified file in the MAS's flash file system, or a compressed archive file that contains the entire content of the flash file system



You back up the entire content of the flash file system to a compressed archive file, which you can then copy from the flash system to another destination.

- Configuration file, either the active running configuration or a startup configuration
- Log files

You can use the following protocols to copy files to or from a MAS:

- File Transfer Protocol (FTP): Standard TCP/IP protocol for exchanging files between computers.
- Trivial File Transfer Protocol (TFTP): Software protocol that does not require user authentication and is simpler to implement and use than FTP.
- Secure Copy (SCP): Protocol for secure transfer of files between computers that relies on the underlying Secure Shell (SSH) protocol to provide authentication and security.



You can use SCP only for transferring image files to or from the MAS, or transferring files between the flash file system on the MAS and a remote host. The SCP server or remote host must support SSH version 2 protocol.

[Table 5](#) lists the parameters that you configure to copy files to or from a MAS.

Table 5 *File Transfer Configuration Parameters*

Server Type	Configuration
Trivial File Transfer Protocol (TFTP)	<ul style="list-style-type: none">• IP address of the server• filename

Table 5 File Transfer Configuration Parameters (Continued)

Server Type	Configuration
File Transfer Protocol (FTP)	<ul style="list-style-type: none"> • IP address of the server • username and password to log into server • filename
Secure Copy (SCP) You must use the CLI to transfer files with SCP.	<ul style="list-style-type: none"> • IP address of the server or remote host • username to log into server • absolute path of filename (otherwise, SCP searches for the file relative to the user's home directory)

For example, you can copy an ArubaOS image file from an SCP server to a system partition on a MAS or copy the startup configuration on a MAS to a file on a TFTP server. You can also store the contents of a MAS's flash file system to an archive file which you can then copy to an FTP server. You can use SCP to securely download system image files from a remote host to the MAS or securely transfer a configuration file from flash to a remote host.

Transferring ArubaOS Image Files

You can download an ArubaOS image file onto a MAS from a TFTP, FTP, or SCP server. In addition, the WebUI allows you to upload an ArubaOS image file from the local PC on which you are running the browser.

When you transfer an ArubaOS image file to a MAS, you must specify the system partition to which the file is copied. The WebUI shows the current content of the system partitions on the MAS. You have the option of rebooting the MAS with the transferred image file.

In the WebUI

1. Navigate to the **Maintenance > Image Management** page.
2. Select TFTP, FTP, SCP, or Local File.
3. Enter or select the appropriate values for the file transfer method.
4. Select the system partition to which the image file is copied.
5. Specify whether the MAS is to be rebooted after the image file is transferred, and whether the current configuration is saved before the MAS is rebooted.
6. Click **Upgrade**.

In the CLI

```
copy tftp: <tftphost> <filename> system: partition [0|1]
copy ftp: <ftphost> <user> <filename> system: partition {0|1}
copy scp: <scphost> <username> <filename> system: partition [0|1]
```

Backing Up and Restoring the Flash File System

You can store the entire content of the flash file system on a MAS to a compressed archive file. You can then copy the archive file to an external server for backup purposes. If necessary, you can restore the backup file from the server to the flash file system.

7. Click **Apply**.

Backup the Flash File System in the CLI

```
backup flash
copy flash: flashbackup.tar.gz tftp: <tftphost> <destfilename>
```



```
copy flash: flashback.tar.gz scp: <scphost> <username> <destfilename>
```

Restore the Flash File System in the WebUI

1. Navigate to the **Maintenance > Copy Files** page.
2. For Source Selection, specify the server to which the flashback.tar.gz file was previously copied.
3. For Destination Selection, select Flash File System.
4. Click **Apply**.

Restore the Flash File System in the CLI

```
copy tftp: <tftphost> <srcfilename> flash: flashback.tar.gz
copy scp: <scphost> <username> <srcfilename> flash: flashback.tar.gz
restore flash
```

Copying Log Files

You can store log files into a compressed archive file which you can then copy to an external TFTP or SCP server. The WebUI allows you to copy the log files to a WinZip folder which you can display or save on your local PC.

In the WebUI

1. Navigate to the **Maintenance > Copy Logs** page.
2. For Destination, specify the TFTP or FTP server to which log files are copied.
3. Select Download Logs to download the log files into a WinZip file on your local PC.
4. Click **Apply**.

In the CLI

```
tar logs
copy flash: logs.tar tftp: <tftphost> <destfilename>
copy flash: logs.tar scp: <scphost> <username> <destfilename>
```

Copying Other Files

The flash file system contains the following configuration files:

- **startup-config**: Contains the configuration options that are used the next time the MAS is rebooted. It contains all options saved by clicking the **Save Configuration** button in the WebUI or by entering the **write memory** CLI command. You can copy this file to a different file in the flash file system or to a TFTP server.
- **running-config**: Contains the current configuration, including changes which have yet to be saved. You can copy this file to a different file in the flash file system, to the startup-config file, or to a TFTP or FTP server.

You can copy a file in the flash file system or a configuration file between the MAS and an external server.

In the WebUI

1. Navigate to the **Maintenance > Copy Files** page.
2. Select the source where the file or image exists.
3. Select the destination to where the file or image is to be copied.
4. Click **Apply**.

In the CLI

```
copy startup-config flash: <filename>
copy startup-config tftp: <tftphost> <filename>

copy running-config flash: <filename>
copy running-config ftp: <ftphost> <user> <password> <filename> [<remote-dir>]
copy running-config startup-config
copy running-config tftp: <tftphost> <filename>
```

Setting the System Clock

You can set the clock on a Mobility Access Switch manually.

In the CLI

To set the date and time, enter the following command in privileged mode:

```
(host) #clock set <year> <month> <date> <hour> <minutes> <seconds>
```

To set the time zone and daylight savings time adjustment, enter the following commands in configure mode:

```
(host) (config) #clock timezone <WORD> <-23 - 23>

clock summer-time <zone> [recurring]
    <1-4> <start day> <start month> <hh:mm>
    first <start day> <start month> <hh:mm>
    last <start day> <start month> <hh:mm>
    <1-4> <end day> <end month> <hh:mm>
    first <end day> <end month> <hh:mm>
    last <end day> <end month> <hh:mm>
    [<-23 - 23>]
```

Clock Synchronization

You can use NTP to synchronize the Mobility Access Switch to a central time source. Configure the Mobility Access Switch to set its system clock using NTP by configuring one or more NTP servers. For each NTP server, you can optionally specify the NTP iburst mode for faster clock synchronization. The iburst mode sends up to ten queries within the first minute to the NTP server. (When iburst mode is not enabled, only one query is sent within the first minute to the NTP server.) After the first minute, the iburst mode typically synchronizes the clock so that queries need to be sent at intervals of 64 seconds or more.



The iburst mode is a configurable option and not the default behavior for the Mobility Access Switch, as this option is considered “aggressive” by some public NTP servers. If an NTP server is unresponsive, the iburst mode continues to send frequent queries until the server responds and time synchronization starts.

Configuring NTP Authentication

The Network Time Protocol adds security to an NTP client by authenticating the server before synchronizing the local clock. NTP authentication works by using a symmetric key which is configured by the user. The secret key is shared by both the Mobility Access Switch and an external NTP server. This helps identify secure servers from fraudulent servers.

The following example enables NTP authentication, adds authentication secret keys into the database, and specifies a subset of keys which are trusted. It also enables the iburst option.

```
(host) (config) #ntp authenticate
```

```
(host) (config) #ntp authentication-key <key-id> md5 <key-secret>  
(host) (config) #ntp trusted-key <key-id>  
(host) (config) #ntp <server IP> iburst key <key-id>
```


The Mobility Access Switch supports profile based configuration for interfaces, interface-groups, port-channels, and VLANs. You can use profiles to apply the same configuration to multiple interfaces and VLANs. It is often tedious to configure a lot of interfaces individually. For example, instead of setting the interface characteristics such as speed and duplex multiple times for multiple interfaces, you can define them in a profile and apply the profile to the interfaces. This is beneficial when you have many interfaces that share the same characteristics where you can define the parameters in a profile and then reference the name of the profile on the interfaces. When you need a change later, the change needs to be made only on the profiles and not on the individual interfaces. The profile-based configuration helps you to avoid having to manage large configurations on every interface and VLAN.

This chapter includes the following topics:

- “Profiles for Interfaces” on page 45
- “Profiles for VLANs” on page 47
- “Scope of the Profiles and Parameters” on page 47
- “Creating a Profile” on page 50
- “Viewing a Profile and its Parameters” on page 51
- “Applying and Activating a Profile” on page 53
- “Deleting a Profile” on page 54
- “Best Practices” on page 54

Profiles for Interfaces

The Mobility Access Switch uses profile-based configuration for the physical interfaces. You can apply the same profile to multiple interfaces that share the same characteristics such as physical specifications, type, and VLAN membership. You can also apply these profiles to an interface-group, or a port-channel.

You can create and apply the following profiles to an interface:

Table 1 *Interface Profiles*

Interface Profile	Description	Reference
dhcp-relay-profile	Specifies the dhcp relay profile for an interface.	See “Configuring DHCP Relay” on page 176.
enet-link-profile	Specifies the physical properties of an interface.	See “Creating and Applying an Ethernet Link Profile to an Interface” on page 89.
gvrp-profile	Specifies the gvrp profile parameters for an interface.	See “Enabling and Configuring GVRP Functionality” on page 117.
igmp-profile	Specifies the igmp profile parameters for an interface.	See “Configuring IGMP” on page 194.

Table 1 *Interface Profiles*

Interface Profile	Description	Reference
lACP-profile	Specifies the dynamic port-channel configuration parameters for an interface.	See “Creating and Applying a Dynamic Port-Channel Profile to an Interface” on page 102.
lldp-profile	Enables or disables the Link Level Discovery Protocol (LLDP) and LLDP MED extension.	See “Verifying the LLDP Profile Configuration to Check LLDP-MED Status” on page 127.
mirroring-in-profile	Specifies the ingress packet mirroring properties for an interface.	See Chapter 39, “Port Mirroring” on page 315
mirroring-out-profile	Specifies the egress packet mirroring properties for an interface.	See Chapter 39, “Port Mirroring” on page 315
mstp-profile	Specifies the MSTP configuration parameters for an interface.	See Chapter 13, “MSTP” on page 135
oam-profile	Specifies the OAM configuration parameters for an interface.	See “Operations, Administration, and Maintenance” on page 105
ospf-profile	Specifies the OSPF configuration parameters for an interface.	See “Configuring OSPF” on page 184.
pim-profile	Specifies the PIM configuration parameters for an interface.	See “Configuring PIM-SM End to End” on page 195.
poe-profile	Specifies the PoE configuration parameters for an interface.	See “Creating and Applying a PoE Profile to an Interface” on page 94.
port-security-profile	Specifies the port security parameters for an interface.	See “Enabling Port Security Functionality” on page 216.
pvst-port-profile	Specifies the parameters for PVST bridge.	See “Configuring using the Interface-based Profile” on page 155.
switching-profile	Specifies the switching parameters such as VLAN and port mode for an interface.	See “Creating and Applying a Switching Profile to an Interface” on page 111.
tunneled-node-profile	Specifies the controller information for a tunneled node interface.	See “Support for Tunneled Node Back-up Server” on page 287.
voip-profile	Specifies the VOIP configuration parameters for an interface that is connected to the VOIP devices and/or PCs and Laptops.	See “Creating and Applying VoIP Profile to an Interface” on page 132.

Profiles for VLANs

You can configure the following profiles for a VLAN:

Table 2 VLAN Profiles

VLAN Profile	Description	Reference
igmp-snooping-profile	Specifies the IGMP snooping configuration parameters for a VLAN.	See “Creating and Applying an IGMP Snooping Profile to a VLAN” on page 201.
mld-snooping-profile	Specifies the MLD snooping configuration parameters for a VLAN.	See “Configuring MLD Snooping” on page 205.
pvst-profile	Specifies the PVST profile configuration parameters for a VLAN.	See “Configuring PVST+” on page 153.

Scope of the Profiles and Parameters

This section includes the following topics:

- “Factory Initial vs Default vs Non-Default Profiles and Parameters” on page 47
- “Profiles and Parameters Assigned to the Interfaces and Groups” on page 47
- “AAA Profiles Assigned to the Interfaces, Groups, and VLANs” on page 49
- “Profiles and Parameters Assigned to the Port-Channel Members” on page 50

Factory Initial vs Default vs Non-Default Profiles and Parameters

There are three factory initial profiles that are effective when you set the Mobility Access Switch to run on the factory initial setup. They are the following:

- `igmp-snooping-factory-initial` assigned to VLAN 1.
- `lldp-factory-initial` assigned to the default interface-group .
- `poe-factory-initial` assigned to the default interface-group.

The `lldp-factory-initial` and the `poe-factory-initial` profiles are also part of the default interface-group configuration and work as the default profiles for all the interfaces.

Any profile that has the `default` reserved keyword as the profile name is called the default profile. Similarly, any parameter assigned to the default interface-group is called the default value for the interface. Modifying any of the default parameters within the default profiles does not make the profile non-default. Similarly, modifying the default parameters for the default interface-group does not make the parameter non-default.

Profiles that you create with names other than `factory-initial` and `default` are called non-default profiles. Similarly, interface-groups that you create using other than the `default` keyword are called non-default interface-groups.

Profiles and Parameters Assigned to the Interfaces and Groups

The effective profile or the parameter for an interfaces is determined by the following concurrent rules:

1. A non-default profile or parameter takes precedence over the default profile or parameter irrespective of whether it is configured under the interface or the interface-group.
2. If the interface and the interface-group have a non-default profile or parameter, then an interface configuration takes precedence over interface-group configuration.

For example, the effective configuration is selected based on the rules in the following table:

Table 3 *Scope of the Interface Parameters and Profiles*

interface gigabitethernet <slot/module/port>	interface-group gigabitethernet <group- name>/default	Effective Profile/Parameter: show interface-config gigabitethernet <slot/module/ port>
default	default	default
default	A (non default)	A (non default)
B (non default)	default	B (non default)
C (non default)	D (non default)	C (non default)

By default, all the interfaces belong to a default interface-group. To view the configuration of the default interface-group, use the **show interface-group-config gigabitethernet default** command. When you create new interface-groups, the interfaces that do not belong to the new interface-groups continue to belong to the default interface-group. Note that overlapping ranges of interfaces among interface-groups is not supported.

You can view the default interface-group configuration using the following command:

```
(host)# show interface-group-config gigabitethernet default
gigabitethernet "default"
-----
Parameter                               Value
-----
Interface group members                  ALL
Interface MSTP profile                  default
Interface Tunneled Node profile         N/A
Interface VOIP profile                  N/A
Interface LLDP profile                  lldp-factory-initial
Interface PoE profile                  poe-factory-initial
Interface Ethernet link profile         default
Interface LACP profile                  N/A
QoS Profile                             N/A
Policer Profile                         N/A
Interface AAA profile                  N/A
Interface Ingress Mirroring profile     N/A
Interface Egress Mirroring profile     N/A
Interface shutdown                     Disabled
mtu                                     1514
Ingress ACL                             N/A
QoS Trust                               Disabled
Interface switching profile            default
Static Multicast Router port for the VLANs N/A
Interface Trusted/Untrusted            Trusted
MAC-Limit (Action)                     N/A
```

You can change the default interface-group using the following command:

```
(host)(config)# interface-group gigabitethernet default
```


For example, the following table determines the effective configuration of the `shutdown` parameter for an interface:

Table 4 *Scope of the Shutdown Parameter*

interface gigabitethernet <slot/module/port>	interface-group gigabitethernet <group- name>/default	Effective Parameter
no shutdown (default)	no shutdown (default)	no shutdown (default)
no shutdown (default)	shutdown (non default)	shutdown (non default)
shutdown (non default)	no shutdown (default)	shutdown (non default)
shutdown (non default)	shutdown (non default)	shutdown (non default)

For example, the following table determines the effective configuration of the `mtu` parameter for an interface:

Table 5 *Scope of the MTU Parameter*

interface gigabitethernet <slot/module/port>	interface-group gigabitethernet <group- name>/default	Effective Parameter
1514 (default)	1514 (default)	1514 (default)
1514 (default)	2000 (non default)	2000 (non default)
1000 (non default)	1514 (default)	1000 (non default)
2500 (non default)	3000 (non default)	2500 (non default)

AAA Profiles Assigned to the Interfaces, Groups, and VLANs

If no AAA profile is configured on the interface, interface-group, or VLAN, then, the default AAA profile is applied to the untrusted interfaces implicitly. If there are different non-default AAA profiles assigned to the interface, interface-group, and VLAN, the effective AAA profile is selected based on the rules in the following table:

Table 6 *Scope of a AAA Profile*

interface gigabitethernet <slot/module/ port>	interface-group gigabitethernet <group-name>/ default	vlan <vlan-id>	Effective AAA Profile
N/A	N/A	N/A	default
N/A	N/A	A (non default)	A (non default)
N/A	B (non default)	C (non default)	B (non default)
D (non default)	E (non default)	F (non default)	D (non default)

The default AAA profile is defined below:

```
(host) #show aaa profile default
```

```
AAA Profile "default"
```

```
-----
```

Parameter	Value
-----	-----
Initial role	logon
MAC Authentication Profile	N/A
MAC Authentication Default Role	guest
MAC Authentication Server Group	default
802.1X Authentication Profile	N/A
802.1X Authentication Default Role	guest
802.1X Authentication Server Group	N/A
Download Role from ClearPass	Enabled
L2 Authentication Fail Through	Disabled
RADIUS Accounting Server Group	N/A
RADIUS Interim Accounting	Disabled
XML API server	N/A
RFC 3576 server	N/A
User derivation rules	N/A
SIP authentication role	N/A
Enforce DHCP	Disabled
Authentication Failure Blacklist Time	3600 sec

You can modify the default AAA profile using the following command:

```
(host)(config)# aaa profile default
```

Profiles and Parameters Assigned to the Port-Channel Members

For port-channel members, apart from the following profiles and parameters, all the other profiles and parameters are inherited from the port-channel configuration:

- shutdown
- enet-link-profile
- lacp-profile
- lldp-profile

Creating a Profile

You can create the profiles using the WebUI or the CLI.

Using the WebUI

1. Navigate to the **Configuration > Ports** page.
2. Select the profile tab.
3. Click **New** under the Profile list.
4. Enter the details in the Profile Name column.
5. Complete the details of the Profile.
6. Click **Apply** and then **Save Configuration**.

Using the CLI

```
(host)(config)# aaa profile <profile-name>
{parameters}
exit
```

```

(host)(config)# vlan-profile igmp-snooping-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile enet-link-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile lacp-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile lldp-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile mirroring-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile mstp-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile poe-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile switching-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile tunneled-node-profile <profile-name>
    {parameters}
exit
(host)(config)# interface-profile voip-profile <profile-name>
    {parameters}
exit
(host)(config)# policer-profile <profile-name>
    {parameters}
exit
(host)(config)# qos-profile <profile-name>
    {parameters}
exit

```

Example:

```

(host) (config)# interface-profile enet-link-profile 10-HALF
(Ethernet Link "10-HALF") #duplex half
(Ethernet Link "10-HALF") #speed 10
(Ethernet Link "10-HALF") #exit

```

Viewing a Profile and its Parameters

Using the CLI

Displaying the List of Profiles Under Each Category

```

(host)# show aaa profile
(host)# show vlan-profile igmp-snooping-profile
(host)# show interface-profile enet-link-profile
(host)# show interface-profile lacp-profile
(host)# show interface-profile lldp-profile
(host)# show interface-profile mirroring-profile
(host)# show interface-profile mstp-profile
(host)# show interface-profile poe-profile

```

```
(host)# show interface-profile switching-profile
(host)# show interface-profile tunneled-node-profile
(host)# show interface-profile voip-profile
(host)# show policer-profile
(host)# show qos-profile
```

Example:

```
(host)# show aaa profile
```

```
AAA Profile List
```

```
-----
Name                References  Profile Status
-----
default             2
default-dot1x       0          Predefined (editable)
default-mac-auth    0          Predefined (editable)
profile-new         3
```

Displaying the Parameters Assigned to Each Profile

```
(host)# show aaa profile <profile-name>
(host)# show vlan-profile igmp-snooping-profile <profile-name>
(host)# show interface-profile enet-link-profile <profile-name>
(host)# show interface-profile lacp-profile <profile-name>
(host)# show interface-profile lldp-profile <profile-name>
(host)# show interface-profile mirroring-profile <profile-name>
(host)# show interface-profile mstp-profile <profile-name>
(host)# show interface-profile poe-profile <profile-name>
(host)# show interface-profile switching-profile <profile-name>
(host)# show interface-profile tunneled-node-profile <profile-name>
(host)# show interface-profile voip-profile <profile-name>
(host)# show policer-profile <profile-name> <profile-name>
(host)# show qos-profile <profile-name>
```

Example:

```
(host) #show aaa profile default
```

```
AAA Profile "default"
```

```
-----
Parameter                Value
-----
Initial role              logon
MAC Authentication Profile N/A
MAC Authentication Default Role guest
MAC Authentication Server Group default
802.1X Authentication Profile N/A
802.1X Authentication Default Role guest
802.1X Authentication Server Group N/A
Download Role from ClearPass Enabled
L2 Authentication Fail Through Disabled
RADIUS Accounting Server Group N/A
RADIUS Interim Accounting Disabled
XML API server            N/A
RFC 3576 server           N/A
User derivation rules     N/A
SIP authentication role   N/A
Enforce DHCP              Disabled
Authentication Failure Blacklist Time 3600 sec
```

Applying and Activating a Profile

Using the CLI

Applying and Activating the Profiles for an Interface

```
(host)(config)# interface gigabitethernet <slot/module/port>
  dhcp-relay-profile <profile-name>
  enet-link-profile <profile-name>
  gvrp-profile <profile-name>
  igmp-profile <profile-name>
  lacp-profile <profile-name>
  lldp-profile <profile-name>
  mirroring-in-profile <profile-name>
  mirroring-out-profile <profile-name>
  mstp-profile <profile-name>
  ospf-profile <profile-name>
  pim-profile <profile-name>
  poe-profile <profile-name>
  port-security-profile <profile-name>
  pvst-port-profile <profile-name>
  switching-profile <profile-name>
  tunneled-node-profile <profile-name>
  voip-profile <profile-name>
```

Applying and Activating the Profiles for an Interface Group

```
(host)(config)# interface-group gigabitethernet {default|<group-name>}
  dhcp-relay-profile <profile-name>
  enet-link-profile <profile-name>
  gvrp-profile <profile-name>
  igmp-profile <profile-name>
  lacp-profile <profile-name>
  lldp-profile <profile-name>
  mirroring-in-profile <profile-name>
  mirroring-out-profile <profile-name>
  mstp-profile <profile-name>
  ospf-profile <profile-name>
  pim-profile <profile-name>
  poe-profile <profile-name>
  port-security-profile <profile-name>
  pvst-port-profile <profile-name>
  switching-profile <profile-name>
  tunneled-node-profile <profile-name>
  voip-profile <profile-name>
```

Applying and Activating the Profiles for a Port-Channel

```
(host)(config)# interface port-channel <ID>
  enet-link-profile <profile-name>
  mirroring-in-profile <profile-name>
  mirroring-out-profile <profile-name>
  mstp-profile <profile-name>
  switching-profile <profile-name>
```

Applying and Activating the Profiles for a VLAN

```
(host)(config)# vlan <ID>
  pvst-profile <profile-name>
  mld-snooping-profile <profile-name>
```

```
igmp-snooping-profile <profile-name>
```

Deleting a Profile

Using the CLI

```
(host)(config)# no aaa profile <profile-name>
(host)(config)# no igmp-snooping-profile <profile-name>
(host)(config)# no interface-profile enet-link-profile <profile-name>
(host)(config)# no interface-profile lacp-profile <profile-name>
(host)(config)# no interface-profile lldp-profile <profile-name>
(host)(config)# no interface-profile mirroring-profile <profile-name>
(host)(config)# no interface-profile mstp-profile <profile-name>
(host)(config)# no interface-profile poe-profile <profile-name>
(host)(config)# no interface-profile switching-profile <profile-name>
(host)(config)# no interface-profile tunneled-node-profile <profile-name>
(host)(config)# no interface-profile voip-profile <profile-name>
(host)(config)# no interface-profile dhcp-relay-profile <profile-name>
(host)(config)# no interface-profile gvrp-profile <profile-name>
(host)(config)# no interface-profile igmp-profile <profile-name>
(host)(config)# no interface-profile ospf-profile <profile-name>
(host)(config)# no interface-profile pim-profile <profile-name>
(host)(config)# no interface-profile port-security-profile <profile-name>
(host)(config)# no interface-profile pvst-port-profile <profile-name>
```

Best Practices

You can manage the profiles efficiently by applying the following guidelines:

- You can use the following process to efficiently manage the profiles:
 - a. Identify the various interface-groups that you need such as Admin, Finance, Marketing, Customer Support, Engineering, and QA.
 - b. Identify the profiles that you need to create for each interface-group.
 - c. Create and apply those profiles to the appropriate interface-groups and port-channels.
 - d. Create and apply the non common profiles to the individual interfaces.
- Use the `show references` command to find out if the profile is used or not, and then, delete all the unused profiles to keep your configuration clean and easy to understand.

The ArubaStack feature enables simplified management by presenting a set of Mobility Access Switches as one entity, and reduces the operational complexity of managing multiple redundant links between access and distribution layer switches. Since the *ArubaStack* appears as one network node, loop prevention protocols are not required.

An ArubaStack is a set of interconnected Mobility Access Switches using stacking ports to form an ArubaStack. A stacking port is a physical port configured to run the stacking protocol. In factory default settings for Mobility Access Switches, uplink ports 2 and 3 are pre-provisioned to be stacking ports. Once a port is provisioned for stacking, it is no longer available to be managed as a network port. A stacking port can only be connected to other Mobility Access Switches running the Aruba Stacking Protocol (ASP). For detailed information on stacking commands see the [ArubaOS 7.2 Command Line Reference Guide](#).

This chapter contains the following sections:

- “Important Points to Remember” on page 55
- “Stacking Topology” on page 56
- “Dynamic Election” on page 57
- “ArubaStack Pre-provisioning” on page 59
- “ArubaStack Database” on page 61
- “ArubaStack Resiliency” on page 62
- “Management User Authentication” on page 68
- “ArubaStack Member Replacement” on page 69

Important Points to Remember

- Dynamic Election—An ArubaStack is formed and roles are assigned based on Auto Discovery.
- ArubaStack Pre-provisioning—ArubaStack members and roles are configured before the ArubaStack is formed.



Dynamic-election and Pre-provisioning cannot be configured together. You must choose one or the other for each ArubaStack.


- An ArubaStack can be formed using S3500s and S2500s.
- The ArubaStack members are Primary, Secondary and Line Card. A valid ArubaStack contains at least a Primary and a Secondary member.
 - Member—a collective term that includes Primary, Secondary, and Line Cards. All valid members run Aruba Stack Protocol (ASP) to discover each other.
 - Primary—runs all Layer2/Layer 3 functions and controls the ArubaStack. All configurations are performed on the Primary and then “pushed” to other members of the ArubaStack.
 - Secondary—back up for the Primary in the event of a hardware or software failure.
 - Line Card—a member of the ArubaStack that is neither a Primary or Secondary. The Line Card includes all interfaces required to *switch* traffic.
- The connection between the Mobility Access Switches cannot go over a Layer 2/Layer 3 cloud.

- One or more stacking ports might be connected between two Mobility Access Switches. The interconnection between the switches can form common topologies; chain, ring, hub-and-spoke etc.
- A port provisioned for stacking can not be managed as a network port.

Stacking Topology

Figure 1 displays an ArubaStack connected in a ring topology. After the election process (see “[Primary Election](#)” on page 59), member 0 is the Primary, member 1 is the Secondary, and member 2 is a Line Card.

Figure 1 ArubaStack Ring Topology



Viewing the ArubaStack Information

There are several commands available that allow you to view ArubaStack information such as topology, members, routes, interface and neighbors to name a few.

```
(host)#show stacking ?
asp-stats          Show asp stats on stacking interfaces
generated-preset-profile  Generate preset stack config from dynamic config
interface          Show configured stacking interfaces
internal           Show stacking internal details
location           Show stacking location
members            Show stacking members
neighbors          Show directly connected stacking neighbors
topology           Show stacking topology
```

For example, to view the ArubaStack topology, use the **show stacking topology** command.

```
(host)#show stacking topology
```

Member-id	Role	Mac Address	Interface	Neighbor	Member-id
0	*	Primary	000b.866a.f240	stack1/2	1
			stack1/3	2	
1		Secondary	000b.866b.0340	stack1/3	0
			stack1/2	2	


```

2          Linecard    000b.866b.3980    stack1/2    0
                                         stack1/3    1

```

Another example, to view the ArubaStack topology, use the **show stacking members** command.

```
(host) (config) #show stacking members
```

```
Member status: Active, Stack Id: 000b866af2404e339e0a
```

```
Stack uptime: 13 days 6 hours 3 minutes 52 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	000b.866a.f240	128	Active	ArubaS3500-24P	AU0000674
1	Secondary	000b.866b.0340	128	Active	ArubaS3500-24P	AU0000731
2	Linecard	000b.866b.3980	128	Active	ArubaS3500-24P	AU0000660

```

[S] - Split
[V] - Version Mismatch
[D] - Depleted Slots
[C] - Preset Configuration Mismatch
[I] - Preset Independent Stack

```



The member with the asterisk (*) indicates that you are logged onto that member (the Primary in the example above).

Dormant State

An ArubaStack member will enter the dormant state if it cannot contact a valid primary member. A member can become dormant for one of the following reasons:

- Split [S]—This member cannot connect to the primary member after an ArubaStack split.
- Version Mismatch [V]—This member's version of ArubaOS does not match that of the primary member.
- Depleted Slots [D]—The number of ArubaStack members has exceeded the maximum.
- Preset Configuration Mismatch [C]—This member's pre-provisioned configuration does not match the configuration of the primary member.
- Preset Independent Stack [I]—This member is part of a pre-provisioned ArubaStack that has not completely merged with another pre-provisioned ArubaStack.

Dynamic Election

Dynamic election is a stack-formation process that is completed automatically with only optional configuration (setting the priority value) done before the Mobility Access Switches are physically connected. The stacking protocol sends information between the ArubaStack members and the election process is completed to determine the primary and secondary members. The primary then assigns member-IDs and roles to the remaining members.

Configuring Priority

When adding a Mobility Access Switch to an ArubaStack, you may need to manually set the priority value so that the switch enters the ArubaStack as a Line Card (or a Primary or Secondary).

The switches priority value is one condition in the election process (see “[Primary Election](#)” on page 59). In the example below, the priority value (election-priority) is set to the default 128 assuring that the switch enters the ArubaStack as a Line Card.



In the example, the switch entering the ArubaStack has a previous member identification (member-id 2).

Using the WebUI

1. Navigate to the **Configuration > Stacking** page.
2. Click the **Add** button to add a MAS to the ArubaStack.
3. Enter the **Member ID**.
4. Enter the **Election Priority**.
5. Click **OK**.
6. Repeat this process until you have added all the necessary MAS's.
7. Set the **MAC persistence timeout** value.
8. Enable or disable **Split Detection** as required for your deployment.
9. Click **Apply** and the **Save Configuration**.

Using the CLI

```
(host)(stack-profile) #member-id 1 election-priority 128
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(host) (stack-profile) #member-id 1 location eng-building
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(host)(stack-profile) #write memory
Saving Configuration.....
```

```
Configuration Saved.
(host)(stack-profile) #
```



The command `member-id <member ID> location` is only available through CLI.

The Stacking Protocol

Each Mobility Access Switch runs an ArubaStack manager process that is responsible for running the Stacking protocol. The Stacking protocol is responsible for automatically:

- Identifying the ArubaStack neighbors and determining the ArubaStack topology.
- Assigning the switch's member ID to each member of the ArubaStack.
- Assigning each member of the ArubaStack a role; Primary, Secondary or Line Card.
- Setting up optimized communication path/channel between the ArubaStack members. This path/channel transports user data packets and the switch's own control packets.
- Converges the stacking topology during a ArubaStack link or ArubaStack member failure event; users and traffic are automatically re-routed via a different path.

Auto Discovery

The Stacking protocol exchanges information between Mobility Access Switches that are connected to each other and without any prior stacking related configuration. The protocol exchanges information between the different ArubaStack members, runs distributed election algorithm, and elects a Primary and Secondary members among the ArubaStack members. The Primary then assigns ArubaStack member IDs to all the members.

Primary Election

The ArubaStack manager discovers the ArubaStack topology. A Primary is elected based on the following in the order of priority.

1. Configured Priority (0-255). Priority is configured by administrator. Higher the priority, better the chances are for the MAS to become Primary. Default priority is 128.
2. Current Role (Primary, Secondary, LC). Weight associated with current role will be in descending order from Primary to LC. If the switch boots up in Dormant state it does not participate in election.
3. Uptime. Uptime for the switch in 100s of seconds.
4. Hardware Priority (0-31). Priority of becoming Primary if all of the above are same. This priority will be hardcoded based on the switch's hardware.
5. MAC Address of the switch. In Primary election, lower MAC wins.

Election Anatomy

The synchronization of the link state database also triggers a primary election task on all the ArubaStack members. This algorithm chooses one primary and one secondary amongst all the ArubaStack members based on the priority list in [“The Stacking Protocol” on page 58](#).

The system's MAC address of the ArubaStack members is the final tiebreaker. The ArubaStack member selected as a Primary asks for an explicit acknowledgment from the remaining ArubaStack members. Upon success, it assigns a ArubaStack unit ID and ArubaStack role for the remaining ArubaStack members and then conveys this information to each ArubaStack member. The ArubaStack unit ID and the chassis-role assigned by the Primary is persistent on a stacking database on all the ArubaStack members. Reboots, therefore, do not result in changes in ArubaStack unit IDs or roles.

Only a Mobility Access Switch that has an un-assigned ArubaStack ID or the same ArubaStack ID as the Primary is allowed to participate fully in the ArubaStack election. In addition, the ArubaStack members must be running the same software version. A Mobility Access Switch with a different software version is admitted into the ArubaStack for the purpose of administration but cannot participate in forwarding network traffic.

Interfaces for such a Mobility Access Switch is not created in the Primary. In the case of incompatible software versions, you can manually upgrade the ArubaStack members, or if configured, the Primary can automatically upgrade the ArubaStack members.

ArubaStack Pre-provisioning

The ArubaStack pre-provisioning feature allows you to configure the role and member-id of the members before the ArubaStack is created. In preset config the members are configured using their serial numbers, which can be found on the purchase order or can be located on the back of the Mobility Access Switch. Additionally, the CLI commands `show inventory` or `show stacking-profile` displays the serial number.

Configuring ArubaStack Pre-provisioning

All configuration for ArubaStack pre-provisioning is completed on a single Mobility Access Switch. Configuration consists of setting all parameters of all eventual members of the ArubaStack. This can be configured using the WebUI or the CLI. These parameters are:

- **Serial number:** The switch's serial number is used to identify the unit for ArubaStack formation. This is located on the purchase order, the rear of the unit, or the commands `show inventory` or `show stacking members` or `show stacking generated-preset-profile`.
- **ArubaStack-unit number:** The member-ID (or slot number) assigned to the switch.
- **Chassis-role:** The role assigned to the switch when configuring the ArubaStack. The roles are primary-capable or line card. Primary-capable switches can become a primary, secondary, or line card.



At least two MAS's in the ArubaStack must be assigned as primary-capable.

After the configuration has been saved, all Mobility Access Switches are physically connected. The ArubaStack then forms a chassis as specified in the configuration.

After the preset ArubaStack configuration is applied to the connected switches, primary-capable members choose one primary and one secondary by running the Primary-Election algorithm. The switches configured as line-card capable will become line cards and receive the configured slot number defined in the preset config after the primary election algorithm.

Using the WebUI

1. Navigate to the **Configuration > Stacking** page.
2. Click the **Enable pre-provisioning** check box.
3. Click the **Add** button to add a MAS to the ArubaStack.
4. Enter the **Member ID**.
5. Enter the **Serial Number**.
6. Select the device **Role** from the drop-down menu.
7. Click **OK**.
8. Repeat this process until you have added all the necessary MAS's.
9. Set the **MAC persistence timeout** value.
10. Enable or disable **Split Detection** as required for your deployment.
11. Click **Apply** and the **Save Configuration**.

Using the CLI

```
(host) (config) # stack-profile
(host) (stack-profile) #member-id 1
(host) (stack-profile) #member-id 1 serial-number AU00006600
(host) (stack-profile) #member-id 1 serial-number AU00006600 role line-card
(host) (stack-profile) #member-id 1 location eng-building
```



The command `member-id <member ID> location` is only available through CLI.

ArubaStack Database

Information related to the ArubaStack is kept in persistent storage so that the ArubaStack's Primary election procedure converges faster after subsequent reboots. This ArubaStack information includes:

- ArubaStack ID
- MAC address, role and member ID of all the members

When the switch boots using the ArubaStack database, it assumes the last role it had according to the ArubaStack database.

To accommodate any change in the ArubaStack topology since the last boot, the Mobility Access Switch uses a count down timer and then it verifies as follows:

- If I was the Primary and...
 - I see the Secondary which means that both the previous Primary and previous Secondary are present in the ArubaStack. I continue as Primary.
 - I do *not* see the Secondary, however, I can see *more* than half of the ArubaStack members in the database. I continue as Primary.
 - I do *not* see the Secondary and I can only see *less* than half of the ArubaStack members in the database. I transition into dormant state. The network interfaces of the switch will remain down.
- If I was the Secondary and...
 - I see the Primary which means that both the previous Primary and previous Secondary are present in the ArubaStack. I continue as Secondary.
 - I do *not* see the Primary, however, I can see *more* than half of the ArubaStack members in the database. I change to Primary.
 - I do *not* see the Primary and I can only see *less* than half of the ArubaStack members in the database. I transition into dormant state. The network interfaces of the switch will remain down.
- If I was a Line Card and...
 - I do *not* see Primary nor Secondary. I move to dormant state.
 - I do see both Primary and Secondary, The Primary will assign me my appropriate role and member-id.
 - I see either the Primary or the Secondary. I will wait for instructions from the member I see (Primary or Secondary).

Removing an ArubaStack Database

An ArubaStack database can be removed at each individual ArubaStack member to return the device to factory default settings. Use the command below to remove an ArubaStack database. Once removed, the device will be automatically reboot.

```
(host) #restore factory-default stacking
```

All configuration and stack settings will be restored to factory default on this member after reload.

Press 'y' to proceed with reload: [y/n]: y

System will now restart

.....

Booting without an ArubaStack Database

When Mobility Access Switches boot without the ArubaStack database, various timers are launched to assure that ArubaStack ports are brought up and RTMs (Routing Topology Messages) are exchanged with other members before deciding on its role. These timers are used to avoid unnecessary transition in roles

and changes in member-id. Because of these timers, the switch's boot up time is longer than with the ArubaStack database.

Primary Switchover

Best practices recommends executing the **database synchronize** command before attempting a system switch over. To view the switch over status, use the `show system switchover` command to verify synchronization before executing the `database synchronize` command.



Periodic synchronization is automatically executed every two minutes.

This command is successful only when both the Primary and Secondary are configured with the same stack-priority. Once this command is executed:

- the Secondary becomes the new Primary
- the old Primary becomes the new Secondary

The example below confirms that database synchronization to the secondary is current.

```
(host) #show system switchover

Secondary Switchover status
-----
System-state   :   synchronized to primary
Configuration  :   synchronized to primary
Database       :   synchronized to primary
```

ArubaStack Resiliency

When a member(s) of an ArubaStack exits the ArubaStack unexpectedly (due to hardware or software error for example) or members are removed from one ArubaStack to create another ArubaStack, it is known as a “stack split.” Keep-alive packets are exchanged among all the ArubaStack ports at regular intervals. When a member(s) of the ArubaStack exits the ArubaStack thereby isolating the remaining ArubaStack member(s), each ArubaStack member independently calculates the resultant state of the stack split.

Some rules governing the stack split are:

- After a stack split, members may transition to a dormant line card state regardless of their previous role.
- After a stack split, several members may form an inactive sub-stack of dormant line card switches.
- After a stack split if the Primary and Secondary members are within the same sub-stack, then that sub-stack is active and passing traffic.
- After a stack split if the Primary is in a different sub-stack than the Secondary, the active sub-stack is determined by the sub-stack with the most members.
- After a stack split if the Primary is in a different sub-stack than the Secondary *and* both sub-stacks contain the same number of members, the sub-stack with the Secondary becomes the active sub-stack. The Secondary rightly assumes that the Primary is completely offline.



An ArubaStack (or sub-stack) can never have two Primaries. The ArubaStack is designed to transition to an inactive state to avoid a collision of two Primaries.

Split Detect

The split detect feature, which detects if a split occurs in an ArubaStack, is enabled by default. When your ArubaStack has only two members, best practices recommends that you disable the split detection feature to ensure that the Primary does not transition to a dormant state if the Secondary is powered down. The command to disable split detections is shown below; note that you must save your configuration.

```
(host)(stack-profile) #no split-detection
WARNING!! This profile will not be applied till the configuration is saved.

(host)(stack-profile) #write memory
Saving Configuration.....
```

The **no split-detection** command is applied to a 2 member ArubaStack *only*. If you apply this command to an ArubaStack with more than 2 members, save the command, then execute the **show stack member** command, a warning notice is displayed.

```
(host)(stack-profile) #show stacking members

Member status: Active, Stack Id: 000b866af2404e339e0a
Id   Role      MAC Address      Priority  State   Model          Serial
--   ---      -
0   * Primary   000b.866a.f240   255      Active  ArubaS3500-24P AU0000674
1   Secondary  000b.866b.0340   200      Active  ArubaS3500-24P AU0000731
2   Linecard   000b.866b.3980   128      Active  ArubaS3500-24P AU0000660

Note: no-split-detect configured but not in effect
```



Split detect is not supported on pre-provisioned ArubaStacks.

Stack Join

Stack join occurs when a stack split creates two sub-stacks; an active sub-stack (includes the Primary and Secondary) and an inactive sub-stack with dormant Line Card members. The stack join pulls these two sub-stacks back together again as one active ArubaStack. The stack join is just resolving the broken connection between switches. There is no software command to issue. Once the connection is made, the stacking protocol will auto discover the ArubaStack topology. Original roles of the switches are maintained because all the switches in the ArubaStack know the identity of the ArubaStack Primary and Secondary and share the same ArubaStack ID.

Additionally, a stack join occurs when two or more MASs with factory default settings are connected via a stack port and then booted up. Those devices will join and the stack protocol will auto discover the stack topology. Each member's role is determined using the primary election algorithm ([“Primary Election” on page 59](#)).

Stack Merge—Dynamic Election

Stack merge takes place when two independently running ArubaStacks (with unique ArubaStack IDs) are connected to each other. Rules to determine which ArubaStack wins the merge are:

- A pre-provisioned ArubaStack wins over a dynamic-election ArubaStack
- An active ArubaStack wins over an inactive ArubaStack
- The ArubaStack with a higher stack priority (priority of the primary) wins
- The ArubaStack with more members wins over an ArubaStack with fewer members
- The ArubaStack with the lower ArubaStack uptime will merge into a higher uptime ArubaStack

- The tie breaker is the Stack ID; the ArubaStack with the lower Stack ID wins

The loosing ArubaStack members perform an automatic software reset to clear any previous software states and then those members join their place in the “winning” ArubaStack.

The following describes a merge scenario in which two MASs with less than 100 seconds of uptime are combined and the device with the lowest MAC becomes the primary. In this scenario, Device-A is the 48-port S3500 and Device-B is the 24-port S3500.

- On Device-A:

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a5ac04f7a3a6c
```

```
Stack uptime: 1 minutes 3 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	000b.866a.5ac0	128	Active	ArubaS3500-48P	AW0000155

- On Device-B:

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a75004f7a3a41
```

```
Stack uptime: 1 minutes 51 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	000b.866a.7500	128	Active	ArubaS3500-24T	AU0000229

- On Device-A, now acting as the primary for the ArubaStack:

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a5ac04f7a3a6c
```

```
Stack uptime: 22 minutes 20 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	000b.866a.5ac0	128	Active	ArubaS3500-48P	AW0000155
1	Secondary	000b.866a.7500	128	Active	ArubaS3500-24T	AU0000229

Stack Merge—Pre-Provisioning

Unlike ArubaStacks created by dynamic election, there is no automatic stack merge for deployments that include pre-provisioned ArubaStacks. If two ArubaStacks must be merged, the process of merging the members must be completed manually.

Pre-provisioned and Dynamic ArubaStacks Merge

In case of merge of one pre-provisioned ArubaStack and one dynamic-election ArubaStack, the pre-provisioned ArubaStack takes precedent. The two ArubaStacks will merge to form a single ArubaStack but the members from dynamic ArubaStack will become dormant if their config is not present in preset config. These members will remain dormant unless the pre-provisioned ArubaStack is modified to include members from dynamic ArubaStack. Complete the merge by taking the following steps.

1. The pre-provisioned ArubaStack will discover the new members and the members of the dynamic-election ArubaStack will become dormant.

After merge:

```
Member status: Active, Stack Id: 000b866b4a804f3f01c6
```

```
Stack uptime: 17 minutes 3 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
----	------	-------------	----------	-------	-------	--------

--	----	-----	-----	-----	-----	-----
0	* Primary	000b.866b.4a80	Preset	Active	ArubaS3500-48P	AW0000257
1	Secondary	000b.866c.2640	Preset	Active	ArubaS3500-48P	AW0000625
?	Linecard	000b.866a.6280	255	Dormant [C]	ArubaS3500-24T	AU0000183
?	Linecard	001a.1e08.7d80	255	Dormant [C]	ArubaS2500-48P	BL0000028

2. Add the former members of the dynamic-election ArubaStack to the stack-profile of the pre-provisioned ArubaStack.

After stack-profile update:

Member status: Active, Stack Id: 000b866b4a804f3f01c6

Stack uptime: 23 minutes 22 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	000b.866b.4a80	Preset	Active	ArubaS3500-48P	AW0000257
1	Secondary	000b.866c.2640	Preset	Active	ArubaS3500-48P	AW0000625
2	Linecard	000b.866a.6280	Preset	Active	ArubaS3500-24T	AU0000183
3	Linecard	001a.1e08.7d80	Preset	Active	ArubaS2500-48P	BL0000028

Pre-provisioned ArubaStacks Merge

If two pre-provisioned ArubaStacks are physically connected via a stack port, they will not merge automatically.



Aruba recommends that you remove the stack-profile configuration or execute `restore factory-default stacking` on each member of the joining ArubaStack before physical connection.

The following is an example of how to remove the pre-provisioned settings from a ArubaStack that will be merged with another pre-provisioned ArubaStack:

```
(Stack-B) #show stacking members
```

Member status: Active, Stack Id: 000b866a76c04f877710

Stack uptime: 1 minutes 56 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
1	Linecard	000b.866b.e300	Preset	Active	ArubaS3500-24P	AU0001357
4	* Primary	000b.866c.0ac0	Preset	Active	ArubaS3500-24P	AU0001517
7	Secondary	000b.866a.76c0	Preset	Active	ArubaS3500-24T	AU0000228

```
(Stack-B) #configure terminal
```

```
Enter Configuration commands, one per line. End with CNTL/Z
```

```
(Stack-B) (config) #stack-profile
```

```
(Stack-B) (stack-profile) #no member-id 1 serial-number AU0001357 role line-card
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #no member-id 4 serial-number AU0001517 role primary-capable
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #no member-id 7 serial-number AU0000228 role primary-capable
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #end
```

```
(Stack-B) #
```

```
(Stack-B) #write memory
Saving Configuration.....
```

```
(Stack-B) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a76c04f877710
```

```
Stack uptime: 16 minutes 3 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
1	Linecard	000b.866b.e300	128	Active	ArubaS3500-24P	AU0001357
4 *	Primary	000b.866c.0ac0	128	Active	ArubaS3500-24P	AU0001517
7	Secondary	000b.866a.76c0	128	Active	ArubaS3500-24T	AU0000228

In the case that two pre-provisioned ArubaStacks are physically connected before the stack-profile is removed from one of them, no merge will occur automatically. The following steps describe how to complete the merge without removing the physical connection:

Before Merge (primary ArubaStack, Stack-A):

```
(Stack-A) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a75004f846b14
```

```
Stack uptime: 15 hours 25 minutes 2 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
4	Linecard	001a.1e08.8140	Preset	Active	ArubaS2500-24P	BJ0000025
5	Secondary	000b.866a.7500	Preset	Active	ArubaS3500-24T	AU0000229
7 *	Primary	000b.866a.5ac0	Preset	Active	ArubaS3500-48P	AW0000155

Before Merge (joining ArubaStack, Stack-B):

```
(Stack-B) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a76c04f875627
```

```
Stack uptime: 22 minutes 51 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
1	Linecard	000b.866b.e300	Preset	Active	ArubaS3500-24P	AU0001357
4	Secondary	000b.866c.0ac0	Preset	Active	ArubaS3500-24P	AU0001517
7 *	Primary	000b.866a.76c0	Preset	Active	ArubaS3500-24T	AU0000228

1. The two ArubaStacks are physically connected using the stacking interfaces.



In this case, both ArubaStacks remain still independent, denoted by [I] but can see the members of the other ArubaStack.

After Physical Connection (primary ArubaStack, Stack-A):

```
(Stack-A) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a75004f846b14
```

```
Stack uptime: 15 hours 27 minutes 31 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
4	Linecard	001a.1e08.8140	Preset	Active	ArubaS2500-24P	BJ0000025
5	Secondary	000b.866a.7500	Preset	Active	ArubaS3500-24T	AU0000229
7 *	Primary	000b.866a.5ac0	Preset	Active	ArubaS3500-48P	AW0000155
?	Linecard	000b.866c.0ac0	Preset	Dormant [I]	ArubaS3500-24P	AU0001517
?	Linecard	000b.866a.76c0	Preset	Dormant [I]	ArubaS3500-24T	AU0000228

```
? Linecard 000b.866b.e300 Preset Dormant [I] ArubaS3500-24P AU0001357
```

After Physical Connection (joining ArubaStack, Stack-B):

```
(Stack-B) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a76c04f875627
```

```
Stack uptime: 26 minutes 59 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
1	Linecard	000b.866b.e300	Preset	Active	ArubaS3500-24P	AU0001357
4	Secondary	000b.866c.0ac0	Preset	Active	ArubaS3500-24P	AU0001517
7 *	Primary	000b.866a.76c0	Preset	Active	ArubaS3500-24T	AU0000228
?	Linecard	001a.1e08.8140	Preset	Dormant [I]	ArubaS2500-24P	BJ0000025
?	Primary	000b.866a.5ac0	Preset	Dormant [I]	ArubaS3500-48P	AW0000155
?	Linecard	000b.866a.7500	Preset	Dormant [I]	ArubaS3500-24T	AU0000229

2. Remove the configured stack-profile from the joining ArubaStack (Stack-B).

```
(Stack-B) #configure terminal
```

```
Enter Configuration commands, one per line. End with CNTL/Z
```

```
(Stack-B) (config) #stack-profile
```

```
(Stack-B) (stack-profile) #no member-id 1 serial-number AU0001357 role line-card  
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #no member-id 4 serial-number AU0001517 role primary-  
capable  
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #no member-id 7 serial-number AU0000228 role primary-  
capable  
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(Stack-B) (stack-profile) #end
```

```
(Stack-B) #write memory
```

3. The members of the joining ArubaStack now merge with the primary ArubaStack.

```
(Stack-A) #show stacking members
```

```
Member status: Active, Stack Id: 000b866a75004f846b14
```

```
Stack uptime: 15 hours 44 minutes 33 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	Linecard	000b.866a.76c0	Preset	Active	ArubaS3500-24T	AU0000228
1	Linecard	000b.866b.e300	Preset	Active	ArubaS3500-24P	AU0001357
2	Linecard	000b.866c.0ac0	Preset	Active	ArubaS3500-24P	AU0001517
4	Linecard	001a.1e08.8140	Preset	Active	ArubaS2500-24P	BJ0000025
5	Secondary	000b.866a.7500	Preset	Active	ArubaS3500-24T	AU0000229
7 *	Primary	000b.866a.5ac0	Preset	Active	ArubaS3500-48P	AW0000155

Console Redirect

Logging onto the ArubaStack using a console connection, from any member, redirects the session to the Primary. You can use a control sequence to redirect between the Primary command line and the ArubaStack's local member's (secondary or line card) command line.



If there is a disconnect between the Primary and its members, for example during an ArubaStack split or primary down, the console automatically redirects to a member command line until the new primary is elected.

Use the following control sequence to redirect console session:

- **Esc Ctrl-I** — redirects the console session from the Primary to a Secondary or Line Card member's command line.
- **Esc Ctrl-r** — redirects the Primary console session from a Secondary or Line Card member's session. This key sequence also enables the console redirect.

To verify the status of the console connection, execute the **show console status** command. In the example below, the ArubaStack has a Primary and a Secondary members only.

Management User Authentication

In an ArubaStack, management users are authenticated by a Primary member. The local user authentication credentials synchronize to all the members so that if the Primary becomes unreachable from other members, the authentication is performed locally. Apart from local admin users, you can configure an external authentication server.

From the Primary member console connection:

```
User:admin
Password: *****

(PPrimary) >enable
Password:*****

(PPrimary) #show console status

Redirect State: Idle
Member Id: 0
```

From a Non-primary member console connection:

```
User:admin
Password: *****

(PPrimary) >enable
Password:*****

(PPrimary) #show console status

Redirect State: Active
Member Id: 1
```

Enter **Esc Ctrl-I** to move to the local console. You will be required to login again.

```
*** CONNECTING TO LOCAL SLOT ***

(LC-1) #
User:admin
```

```

Password: *****

(LC-1) >enable
Password:*****

(LC-1) #show console status

Redirect State: Disabled
Member Id: 1

```

ArubaStack Member Replacement

The ArubaStack features allows the user to replace one or more members of a ArubaStack without bringing down the complete ArubaStack. Following are best practices, based on dynamic and preset ArubaStack configurations.



When replacing a unit with another unit that is not factory default, it is recommended to restore the unit to factory default as shown below.

```
(Aruba) #restore factory_default stacking
```

All configuration and stack settings will be restored to factory default on this member after reload.
 Press 'y' to proceed with reload: [y/n]: y
 System will now restart

Dynamic ArubaStack Configuration

The following section describes how to replace a member of a dynamic ArubaStack.

Replacing a Linecard Member

```
(host) #show stacking members
```

```

Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 3 minutes 55 seconds

```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	128	Active	ArubaS2500-48T	BK0000016
1	Linecard	001a.1e08.7b80	128	Active	ArubaS2500-48T	BK0000018
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

In the above ArubaStack of four members, if Linecard member 1 is down and to be replaced, complete the following steps:

1. Verify stacking members. Member 1 is down and the status will be displayed as Away and the role will be Unknown.

```
(host) #show stacking members
```

```

Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 11 minutes 16 seconds

```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	128	Active	ArubaS2500-48T	BK0000016
1	Unknown	001a.1e08.7b80	128	Away	ArubaS2500-48T	BK0000018
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015

```
3      Linecard    001a.1e08.7c80  128      Active    ArubaS2500-48T    BK0000014
```

2. To replace member 1, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 1
```

```
Member-id: 0
```

```
-----
```

```
Deleting Member-id: 1
```

```
Member-id: 2
```

```
-----
```

```
Deleting Member-id: 1
```

```
Member-id: 3
```

```
-----
```

```
Deleting Member-id: 1
```

3. Stacking database will be cleared and member 1 will not be visible in the show stacking command as shown below.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 18 minutes 29 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	128	Active	ArubaS2500-48T	BK0000016
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

4. Physically replace member with a new unit. The new unit will transition from an invalid unit Id shown by (?) and eventually be assigned the lowest stack-id available in the existing ArubaStack. In this case the new unit will be assigned unit ID 1.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 29 minutes 15 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	128	Active	ArubaS2500-48T	BK0000016
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014
?	Linecard	001a.1e08.7ac0	128	Active	ArubaS2500-48T	BK0000019

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 29 minutes 17 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	128	Active	ArubaS2500-48T	BK0000016
1	Linecard	001a.1e08.7ac0	128	Active	ArubaS2500-48T	BK0000019
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

Replacing a Secondary Member



The new member joining the ArubaStack will assume the role of Secondary only if the priority is configured to be higher than the Linecard members. If the priority is the same for all the members an existing member of the ArubaStack will be elected as the secondary and the new member joining the ArubaStack will be a Linecard.

In this scenario member-ID 1 is configured for a higher priority.

```
(host) #show stack-profile
```

```
stack-profile "default"
-----
Parameter                Value
-----
MAC persistence timeout   15 Minutes
Split Detection           Enabled
Election Priority:
  Member 0                250
  Member 1                250
```

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 42 minutes 40 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
---	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	250	Active	ArubaS2500-48T	BK0000016
1	Secondary	001a.1e08.7ac0	250	Active	ArubaS2500-48T	BK0000019
2	Linecard	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

In the above ArubaStack of four members, if the Secondary member 1 is down and needs to be replaced, here are the steps:

1. Verify stacking members. Secondary member 1 is down and the status will be displayed as Away and the role will be Unknown. An existing member will be elected as the secondary unless the secondary role is configured for a higher priority

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 43 minutes 50 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
---	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	250	Active	ArubaS2500-48T	BK0000016
1	Unknown	001a.1e08.7ac0	250	Away	ArubaS2500-48T	BK0000019
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

2. To replace member 1, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 1
```

```
Member-id: 0
-----
Deleting Member-id: 1
```

```
Member-id: 2
-----
Deleting Member-id: 1

Member-id: 3
-----
Deleting Member-id: 1
```

- Stacking database will be cleared and member 1 will not be visible in the show stacking command as shown below.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 44 minutes 46 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	250	Active	ArubaS2500-48T	BK0000016
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

- Physically replace member with a new unit. The new unit will transition from an invalid unit Id shown by (?) and eventually be assigned the lowest stack-id available in the existing ArubaStack. In this case the new unit will be assigned unit ID 1 and since member 1 is configured with higher priority it will be elected as secondary.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 47 minutes 6 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	250	Active	ArubaS2500-48T	BK0000016
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014
?	Unknown	001a.1e08.7a80	128	Away	ArubaS2500-48T	BK0000017

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 48 minutes 53 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	250	Active	ArubaS2500-48T	BK0000016
1	Secondary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK0000017
2	Linecard	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

Replacing a Primay Member

The new member joining the ArubaStack will assume the role of Primary only if the priority is configured to be higher than the Secondary member. If the priority of the primary and secondary are same, the existing Secondary member of the ArubaStack will be elected as the Primary and the new member joining the ArubaStack will be elected as Secondary.

If the priority is the same for all the members an existing secondary will take over the role of Primary member, and an existing Linecard member will assume the role of Secondary. The new member joining the ArubaStack will be a Linecard. In this scenario member-id 0 and 1 are configured for a higher priority


```
(host) #show stack-profile

stack-profile "default"
-----
Parameter          Value
-----
MAC persistence timeout 15 Minutes
Split Detection      Enabled
Election Priority:
  Member 0           255
  Member 1           250
```

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 1 hours 10 minutes 12 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7b00	255	Active	ArubaS2500-48T	BK00000016
1	Secondary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK00000017
2	Linecard	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK00000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK00000014

In the above stack of four members, if the Primary member 0 is down and needs to be replaced, here are the steps:

1. Verify stacking members. Primary member 0 is down and the status will be displayed as Away and the role will be Unknown. An existing Secondary member will be elected as the Primary and an existing Linecard member will be elected as Secondary.

```
(host) # show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	Unknown	001a.1e08.7b00	255	Away	ArubaS2500-48T	BK00000016
1	Primary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK00000017
2	* Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK00000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK00000014

2. To replace member 0, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 0
```

```
Member-id: 1
-----
Deleting Member-id: 0

Member-id: 2
-----
Deleting Member-id: 0

Member-id: 3
-----
Deleting Member-id: 0
```

3. Stacking database will be cleared and member 0 will not be visible in the show stacking command as shown below.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 1 hours 17 minutes 13 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
1	* Primary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK0000017
2	Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014

4. Physically replace member with a new unit. The new unit will transition from an invalid unit Id shown by (?) and eventually be assigned the lowest stack-id available in the existing ArubaStack. In this case the new unit will be assigned unit ID 0 and since member 0 is configured with highest priority it will be elected as Primary.

```
(host) # show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
1	Primary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK0000017
2	* Secondary	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014
?	Unknown	001a.1e08.7b00	255	Away	ArubaS2500-48T	BK0000016

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 47 minutes 6 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7b00	255	Active	ArubaS2500-48T	BK0000016
1	Secondary	001a.1e08.7a80	250	Active	ArubaS2500-48T	BK0000017
2	Linecard	001a.1e08.7c00	128	Active	ArubaS2500-48T	BK0000015
3	Linecard	001a.1e08.7c80	128	Active	ArubaS2500-48T	BK0000014



To avoid another switchover happened when the new unit becomes the primary, you may want to modify ArubaStack profile to keep member-1 as primary and new unit as secondary.

```
(host) #show stack-profile
```

```
stack-profile "default"
```

```
-----
```

Parameter	Value
-----	-----
MAC persistence timeout	15 Minutes
Split Detection	Enabled
Election Priority:	
Member 0	250
Member 1	255

Preset ArubaStack Configuration

The following section describes how to replace a member of a preset ArubaStack.

In a preset ArubaStack configuration, the units are assigned role and slot number using the stack-profile configuration. Here is a ArubaStack of four members configured as below

```
(host) #show stack-profile

stack-profile "default"
-----
Parameter          Value
-----
MAC persistence timeout 15 Minutes
Split Detection      Enabled
```

```
Preset-profile:
-----
```

```
Member-id      Serial-number  Role
0              BK0000020    Primary-capable
1              BK0000017    Primary-capable
2              BK0000015    Line-card
3              BK0000014    Line-card
```

Replacing a Linecard Member

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 2 hours 19 minutes 26 seconds

Id   Role      MAC Address      Priority  State   Model          Serial
--   -
0   *   Primary    001a.1e08.7bc0   Preset   Active   ArubaS2500-48T  BK0000020
1       Secondary 001a.1e08.7a80   Preset   Active   ArubaS2500-48T  BK0000017
2       Linecard   001a.1e08.7c00   Preset   Active   ArubaS2500-48T  BK0000015
3       Linecard   001a.1e08.7c80   Preset   Active   ArubaS2500-48T  BK0000014
```

In the above ArubaStack of four members, if Linecard member 2 is down and to be replaced, here are the steps:

1. Verify stacking members. Member 2 is down and the status will be displayed as Away and the role will be Unknown.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
Stack uptime: 2 hours 33 minutes 56 seconds

Id   Role      MAC Address      Priority  State   Model          Serial
--   -
0   *   Primary    001a.1e08.7bc0   Preset   Active   ArubaS2500-48T  BK0000020
1       Secondary 001a.1e08.7a80   Preset   Active   ArubaS2500-48T  BK0000017
2       Unknown   001a.1e08.7c00   Preset   Away     ArubaS2500-48T  BK0000015
3       Linecard   001a.1e08.7c80   Preset   Active   ArubaS2500-48T  BK0000014
```

2. To replace member 2, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 2
```

```
Member-id: 0
-----
Deleting Member-id: 2
```

```
Member-id: 1
-----
Deleting Member-id: 2
```

```
Member-id: 3
-----
Deleting Member-id: 3
```

3. Stacking database will be cleared and member 2 will not be visible in the show stacking command as shown below.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 2 hours 36 minutes 10 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial	
0	*	Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1		Secondary	001a.1e08.7a80	Preset	Active	ArubaS2500-48T	BK0000017
3		Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

4. Delete the serial number of member 2.

```
(host) (stack-profile) #no member-id 2 serial-number BK0000018 role line-card
```

5. Physically replace member with a new unit. The unit will not be an active part of the ArubaStack until the serial number is added to the stack-profile and will be displayed as Dormant

```
(host) (stack-profile) #show stacking members
```

```
Member status: Active, Stack Id: 001a1e087b004fcee152
Stack uptime: 4 hours 24 minutes 50 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial	
0	*	Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1		Secondary	001a.1e08.7a80	Preset	Active	ArubaS2500-48T	BK0000017
2		Linecard	001a.1e08.7b80	128	Dormant [C]	ArubaS2500-48T	BK0000018
3		Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

```
[S] - Split
[V] - Version Mismatch
[D] - Depleted Slots
[C] - Preset Configuration Mismatch
[I] - Preset Independent Stack
```

6. Add the serial number of the new unit to the ArubaStack using the following command and save the configuration.

```
(host) (stack-profile) #member-id 2 serial-number BK0000018 role line-card
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(host) (stack-profile) #write memory
Saving Configuration.....
```

```
Configuration Saved.
```

```
(host) #
```

7. The new unit will now be part of the ArubaStack

```
(host) #show stacking members
```

Member status: Active, Stack Id: 001ale087b004fcee152

Stack uptime: 3 hours 14 minutes 49 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1	Secondary	001a.1e08.7a80	Preset	Active	ArubaS2500-48T	BK0000017
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

Replacing a Secondary Member

In a stack-preset configuration at least two members in a ArubaStack must be configured as primary capable.

- An existing Linecard member will be elected as the Secondary if there is a unit that has a role as primary-capable
- If all other units are configured as Linecard, no Secondary member will be elected.
- If the Secondary unit needs to be replaced, the best practices are listed below.

```
(host) #show stack-profile
```

```
stack-profile "default"
```

Parameter	Value
-----	-----
MAC persistence timeout	14 Minutes
Split Detection	Enabled

```
Preset-profile:
```

Member-id	Serial-number	Role
0	BK0000020	Primary-capable
1	BK0000017	Primary-capable
2	BK0000018	Line-card
3	BK0000014	Line-card

In the above ArubaStack of four members, if the Secondary member 1 is down and needs to be replaced, here are the steps:

1. Verify stacking members. Secondary member 1 is down and the status will be displayed as Away and the role will be Unknown.

```
(host) #show stacking members
```

Member status: Active, Stack Id: 001ale087b004fcee152

Stack uptime: 4 hours 17 minutes 39 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1	Unknown	001a.1e08.7a80	Preset	Away	ArubaS2500-48T	BK0000017
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

2. To replace member 1, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 1
```

```
Member-id: 0
```

```
-----
```

Deleting Member-id: 1

Member-id: 2

Deleting Member-id: 1

Member-id: 3

Deleting Member-id: 1

3. Stacking database will be cleared and member 1 will not be visible in the show stacking command as shown below.

((host) #show stacking members

Member status: Active, Stack Id: 001ale087b004fcee152

Stack uptime: 4 hours 20 minutes 18 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

4. Delete the serial number of member 1 from the stack-profile.

(host) (stack-profile) #no member-id 1 serial-number BK0000017 role line-card

5. Physically replace member with a new unit.

6. The unit will not be an active part of the ArubaStack until the serial number is added to the stack-profile and will be displayed as Dormant.

(host) (stack-profile) #show stacking members

Member status: Active, Stack Id: 001ale087b004fcee152

Stack uptime: 4 hours 34 minutes 57 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	* Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1	Linecard	001a.1e08.7b00	128	Dormant [C]	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

[S] - Split

[V] - Version Mismatch

[D] - Depleted Slots

[C] - Preset Configuration Mismatch

[I] - Preset Independent Stack

7. Add the serial number of the new unit to the ArubaStack using the following command and save the configuration

(host) (config) #stack-profile member-id 1 serial-number BK0000016 role primary-capable
WARNING!! This profile will not be applied till the configuration is saved.

(host) (config) #write memory

Saving Configuration.....

Configuration Saved.

8. The new unit will now be part of the ArubaStack

```
(host) (config) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 4 hours 47 minutes 18 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	* Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1	Secondary	001a.1e08.7b00	Preset	Active	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

Replacing a Primary Member

In a stack-preset configuration at least two members in a ArubaStack must be configured as primary capable.

- The Secondary member will be elected as a Primary.
- An existing Linecard member will be elected as the Secondary if there is a unit that has a role as primary-capable
- If all other units are configured as Linecard, no Secondary member will be elected.
- If the Primary unit needs to be replaced, the best practices are listed below.

In this scenario member-id 0 and 1 are configured as primary capable

```
(host) #show stack-profile
```

```
stack-profile "default"
```

```
-----
```

Parameter	Value
MAC persistence timeout	14 Minutes
Split Detection	Enabled

```
Preset-profile:
```

```
-----
```

Member-id	Serial-number	Role
0	BK0000020	Primary-capable
1	BK0000016	Primary-capable
2	BK0000018	Line-card
3	BK0000014	Line-card

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	Primary	001a.1e08.7bc0	Preset	Active	ArubaS2500-48T	BK0000020
1	Secondary	001a.1e08.7a80	Preset	Active	ArubaS2500-48T	BK0000017
3	* Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014
4	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018

In the above ArubaStack of four members, if the Primary member 0 is down and needs to be replaced, here are the steps:

1. Verify stacking members. Primary member 0 is down and the status will be displayed as Away and the role will be Unknown. An existing Secondary member will be elected as the Primary.

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 4 hours 52 minutes 32 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	Unknown	001a.1e08.7bc0	Preset	Away	ArubaS2500-48T	BK0000020
1	* Primary	001a.1e08.7b00	Preset	Active	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

2. To replace member 0, clear the stacking database from the ArubaStack using the clear command as shown below.

```
(host) #clear stacking member-id 0
```

```
Member-id: 1
```

```
-----
```

```
Deleting Member-id: 0
```

```
Member-id: 2
```

```
-----
```

```
Deleting Member-id: 0
```

```
Member-id: 3
```

```
-----
```

```
Deleting Member-id: 0
```

3. Stacking database will be cleared and member 0 will not be visible in the show stacking command as shown below.

```
((host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 5 hours 12 minutes 55 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
1	* Primary	001a.1e08.7b00	Preset	Active	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

4. Delete the serial number of member 0 from the stack-profile.

```
(host) (stack-profile) #no member-id 0 serial-number BK0000020 role line-card
```

5. Physically replace member with a new unit.

6. The unit will not be an active part of the ArubaStack until the serial number is added to the stack-profile and will be displayed as Dormant

```
(host) #show stacking members
```

```
Member status: Active, Stack Id: 001ale087b004fcee152
```

```
Stack uptime: 5 hours 24 minutes 32 seconds
```

Id	Role	MAC Address	Priority	State	Model	Serial
0	Linecard	001a.1e08.7ac0	128	Dormant [C]	ArubaS2500-48T	BK0000019
1	* Primary	001a.1e08.7b00	Preset	Active	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

- [S] - Split
- [V] - Version Mismatch
- [D] - Depleted Slots
- [C] - Preset Configuration Mismatch
- [I] - Preset Independent Stack

7. Add the serial number of the new unit to the ArubaStack using the following command and save the configuration.

```
(host) (config) #stack-profile member-id 0 serial-number BK0000019 role primary-capable
WARNING!! This profile will not be applied till the configuration is saved.
```

```
(host) (config) #write memory
Saving Configuration.....
```

Configuration Saved.

```
(host) #
```

8. The new unit will now be part of the ArubaStack and be elected as Secondary

```
(host) #show stacking members
```

Member status: Active, Stack Id: 001a1e087b004fcee152

Stack uptime: 5 hours 29 minutes 51 seconds

Id	Role	MAC Address	Priority	State	Model	Serial
--	----	-----	-----	-----	-----	-----
0	Secondary	001a.1e08.7ac0	Preset	Active	ArubaS2500-48T	BK0000019
1	* Primary	001a.1e08.7b00	Preset	Active	ArubaS2500-48T	BK0000016
2	Linecard	001a.1e08.7b80	Preset	Active	ArubaS2500-48T	BK0000018
3	Linecard	001a.1e08.7c80	Preset	Active	ArubaS2500-48T	BK0000014

The Mobility Access Switch supports 24 or 48 gigabit ethernet network interfaces, up to four 10-gigabit ethernet uplink interfaces and a management port. You can configure power over ethernet only on the front panel ports of the Mobility Access Switch. The PoE cannot be configured for SFP/SFP+ port, management port, and any ports on the T models.

This chapter describes how to configure the ethernet interfaces.

This chapter includes the following topics:

- “Configuring the Management Port” on page 83
- “Gigabit Ethernet Network Interfaces” on page 84
- “Time-Domain Reflectometer (TDR)” on page 84
- “Configuring an Interface Group” on page 86
- “Creating and Applying an Ethernet Link Profile to an Interface” on page 89
- “Loopback Interfaces” on page 91
- “Power Over Ethernet” on page 92
- “Configuring Power Over Ethernet” on page 93
- “Creating and Applying a PoE Profile to an Interface” on page 94

Configuring the Management Port

The management interface is located above the console port on the rear panel of the Mobility Access Switch. It is labeled as *mgmt*. The management port is a dedicated interface for out-of-band management purpose. This interface is specifically available for the management of the system and cannot be used as a switching interface. You can configure only the IP address and description for this interface. The management port can be used to access the Mobility Access Switch from any location and configure the system.

You can configure the management port using the CLI.

Using the CLI

```
(host)(config)# interface mgmt
  description <name>
  ip address <ip-address> <mask>
  ipv6 [ <prefix> prefix_len <prefix_len> | link-local <link-local-address> ]
  no {...}
  shutdown
```

Sample Management Port Configuration

```
(host)(config)# interface mgmt
  description MGMT_PORT
  ip address 10.1.13.1 255.255.255.0
  no shutdown
```

Gigabit Ethernet Network Interfaces

The Mobility Access Switch supports 24 or 48 port gigabit ethernet interfaces of 10/100/1000 Mbps speeds.

A network gigabit ethernet interface is referred by its `<slot>/<module>/<port>`.

- Slot—The member ID of the stack.
- Module—There are two modules where the first one is the front-panel network module (0), while the other one is the uplink network module (1).
- Port—The individual port number.

For example, interface `gigabitethernet 0/0/20` refers to the first stack member (0) on the front-panel network module (0) at port number (20).



The Mobility Access Switch also supports four 10-Gigabit Ethernet interfaces for stacking and uplink purposes. See the Hardware Installation Guide for more information on the 10-Gigabit Ethernet uplink module.

Time-Domain Reflectometer (TDR)

Time-Domain Reflectometer (TDR) is a measurement technique used to characterize and locate faults in metallic cables such as twisted pair. TDR transmits a short rise electric pulse across the conducting cable and if the cable is properly terminated, the entire electric pulse is absorbed on the other end. If any faults exist in the cable, some of the incident signal is sent back towards the source. TDR also:

- Locates the position of faults within meters
- Detects and reports open circuits, short circuits, and impedance mismatches in a cable
- Detects pair swap (straight/crossover) on each pair of cable in twisted pair cable
- Detects pair polarity (positive/negative) on each channel pairs in a cable



TDR is not supported over management interfaces, Direct Attach Cables (DAC) or Fiber interfaces.

Use this command to execute a TDR diagnostic test on a specific `gigabitethernet` interface.

```
(host) (config)# run diagnostics interface gigabitethernet <slot/module/port> cable
```

Use the following command to view the test results for the Time-Domain Reflectometer (TDR) cable diagnostics:

```
(host)# show diagnostics interface gigabitethernet
```

Configuring an Ethernet Interface

To set up your network, you can configure the various parameters for each ethernet network and uplink interfaces individually.

Using the CLI

```
(host)(config)# interface gigabitethernet <slot/module/port>
  aaa-profile <profile_name>
  backup interface {gigabitethernet <slot/module/port> | port-channel <0-7>}
  clone <source>
  description <description>
  enet-link-profile <profile_name>
  igmp-snooping mrouter-vlan {add | delete} <vlan-id>
```

```

ip access-group in <in>
lACP-profile <profile_name>
lldp-profile <profile_name>
mac-limit <limit>
mirroring-in-profile <profile_name>
mirroring-out-profile <profile_name>
mstp-profile <profile_name>
mtu <64-9216>
no {...}
poe-profile <profile_name>
policer-profile <profile_name>
preemption delay <10-300>
preemption mode {forced | off}
qos trust
qos-profile <profile_name>
shutdown
switching-profile <profile_name>
trusted port
tunneled-node-profile <profile_name>
voip-profile <profile_name>
exit

```

Configuring Jumbo Frame Size

The Mobility Access Switch supports jumbo frames. You can enable jumbo frames on a per-interface basis with sizes from 64 to 9216 bytes. The default size is 1514 bytes.

```

(host)(config)# interface gigabitEthernet 0/0/6
mtu 9216
exit

```

Verifying Jumbo Frame Size

You can verify the jumbo frame size on an interface using the following command:

```

(host)# show interface gigabitEthernet 0/0/6
GE0/0/6 is administratively Up, Link is Down, Line protocol is Down
Hardware is Gigabit Ethernet, Address is 00:0b:86:6a:42:03
Encapsulation ARPA, Loopback not set
Configured: duplex (Auto), Speed (Auto), FC (Off), Autoneg (On)
Auto negotiation in progress
Interface index: 2
MTU 9216 bytes
Flags: Access, Trusted
Link status last changed:      0d 00:00:00 ago
Last update of counters:      0d 00:00:00 ago
Last clearing of counters:     0d 00:00:00 ago
<output truncated>

```

Displaying Interface Counters and Statistics

```

(host)# show interface gigabitEthernet 0/0/1 counters

```

Port	InOctets	InUcastPkts	InMcastPkts	InBcastPkts
GE0/0/1	0	0	0	0

Port	OutOctets	OutUcastPkts	OutMcastPkts	OutBcastPkts
GE0/0/1	0	0	0	0


```

(host)# show interface gigabitEthernet 0/0/1 statistics
Last update of counters:      0d 00:00:00 ago
Last clearing of counters:     0d 00:00:00 ago

```

```

Received Statistics:
  0 frames, 0 octets
  0 unicast, 0 multicast, 0 broadcast
  0 error frames, 0 error octets, 0 CRC events, 0 runts, 0 giants, 0 throttles
  0 drop events
Transmitted Statistics:
  0 frames, 0 octets
  0 unicast, 0 multicast, 0 broadcast
  0 throttles, 0 deferred
  0 collisions, 0 multiple collisions, 0 late collisions
Received and Transmitted Frame Size Statistics:
  0 64 octet, 0 65-127 octet, 0 128-255 octet, 0 256-511 octet, 0 512-1023 octet, 0 1024-
  max octet

```

Configuring an Interface Group

In the CLI configuration, it is often tedious to individually configure interfaces when there are multiple interfaces that have the same configuration. In such scenarios, you can group the interfaces together so that any interface within the group has the same configuration. When you configure an interface that is a member of an interface-group, applying a non-default profile or a parameter to the interface takes precedence over the interface-group configuration. By default, all the interfaces belong to a default interface-group.

To view the configuration of the default interface-group, use the **show interface-group-config gigabitethernet default** command. When you create non-default interface-groups, the excluded interfaces continue to belong to the default interface-group.



Interface-group and port-channel are not the same. Interface group assigns the configuration to individual interfaces whereas the port-channel makes a group of interfaces to work as a single logical interface.



You cannot have overlapping ranges of interfaces when you have multiple interface-groups. For more information about the scope of an interface and interface-group profiles, see [“Scope of the Profiles and Parameters” on page 47](#).

Using the CLI

```

(host)(config)# interface-group gigabitethernet {default|<group-name>}
  aaa-profile <profile_name>
  apply-to <interface range> add | remove
  clone <source>
  enet-link-profile <profile_name>
  igmp-snooping mrouter-vlan {add | delete} <vlan-id>
  ip access-group in <in>
  lacp-profile <profile_name>
  lldp-profile <profile_name>
  mac-limit <limit>
  mirroring-in-profile <profile_name>
  mirroring-out-profile <profile_name>
  mld-snooping mrouter-vlan {add | delete} <vlan-list>
  mstp-profile <profile_name>
  mtu <64-9216>
  tunneled-node-profile <profile-name>
  no {...}
  poe-profile <profile_name>

```

```

policer-profile <profile_name>
qos trust
qos-profile <profile_name>
shutdown
switching-profile <profile_name>
trusted port
voip-profile <profile_name>

```

Sample Interface Group Configuration

```

(host)(config)# interface-group gigabitethernet FINANCE
apply-to 0/0/0-0/0/20,0/0/32

```



Ensure that you do not add blank spaces between the ranges or multiple interfaces, and there must be three tuples in the individual, starting, and ending ranges. Also, the interface numbers should be in ascending order from start to finish of the range value. For example, 0/0, 0/1/0-1/1 is not a valid range because there is a space and the interface number format is not of slot/module/port in all the occurrences.

Verifying the Interface Group Configuration

You can use the following commands to view details about an interface-group.

```

(host)# show interface-group-config gigabitethernet default
gigabitethernet "default"
-----
Parameter                                         Value
-----
Interface group members                          ALL
Interface MSTP profile                          default
Interface Tunneled Node profile                 N/A
Interface VOIP profile                          N/A
Interface LLDP profile                          lldp-factory-initial
Interface PoE profile                           poe-factory-initial
Interface Ethernet link profile                 default
Interface LACP profile                         N/A
QoS Profile                                     N/A
Policer Profile                                N/A
Interface AAA profile                           N/A
Interface Ingress Mirroring profile             N/A
Interface Egress Mirroring profile             N/A
Interface shutdown                             Disabled
mtu                                              1514
Ingress ACL                                    N/A
QoS Trust                                       Disabled
Interface switching profile                    default
Static IGMP Multicast Router port for VLANs    N/A
Static MLD Multicast Router port for VLANs     N/A
Interface Trusted/Untrusted                    Trusted
MAC-Limit (Action)                             N/A

```

```

(host)# show interface-group-config gigabitethernet FINANCE
gigabitethernet "FINANCE"
-----
Parameter                                         Value
-----
Interface group members                          0/0/0-0/0/20,0/0/32

```

Interface MSTP profile	default
Interface Tunneled Node profile	N/A
Interface VOIP profile	N/A
Interface LLDP profile	default
Interface PoE profile	default
Interface Ethernet link profile	default
Interface LACP profile	N/A
QoS Profile	N/A
Policer Profile	N/A
Interface AAA profile	N/A
Interface Ingress Mirroring profile	N/A
Interface Egress Mirroring profile	N/A
Interface shutdown	Disabled
mtu	1514
Ingress ACL	N/A
QoS Trust	Disabled
Interface switching profile	default
Static Multicast Router port for the VLANs	N/A
Interface Trusted/Untrusted	Trusted
MAC-Limit (Action)	N/A

(host)# show interface-group-config gigabitethernet

gigabitethernet List

Name	References	Profile Status
-----	-----	-----
default	0	
FirstFloor	0	
SecondFloor	0	
Total:3		



In the case of LLDP and PoE profiles, the default interface-group has lldp-factory-initial and poe-factory-initial profiles applied, whereas a non-default interface-group that you create has the LLDP and PoE default profiles applied. The default LLDP and PoE profiles have LLDP and PoE disabled, while they are enabled in the factory-initial profiles.

You can view the differences in the LLDP and PoE factory-initial and default profiles using the following commands:

(host)# show interface-profile poe-profile poe-factory-initial

Power over Ethernet profile "poe-factory-initial"

Parameter	Value
-----	-----
Enable PoE interface	Enabled
Max Power on PoE port milliwatts	30000
PoE port priority	low
Power over Ethernet Cisco compatibility	Disabled

(host)# show interface-profile poe-profile default

Power over Ethernet profile "default"

Parameter	Value
-----	-----
Enable PoE interface	Disabled
Max Power on PoE port milliwatts	30000
PoE port priority	low


```
(host)# show interface-profile lldp-profile lldp-factory-initial
LLDP Profile "lldp-factory-initial"
-----
Parameter                                         Value
-----
LLDP pdu transmit                               Enabled
LLDP protocol receive processing                 Enabled
LLDP transmit interval (Secs)                   30
LLDP transmit hold multiplier                   4
LLDP-MED protocol                               Enabled

(host)# show interface-profile lldp-profile default
LLDP Profile "default"
-----
Parameter                                         Value
-----
LLDP pdu transmit                               Disabled
LLDP protocol receive processing                 Disabled
LLDP transmit interval (Secs)                   30
LLDP transmit hold multiplier                   4
LLDP-MED protocol                               Disabled
```

Creating and Applying an Ethernet Link Profile to an Interface

You can use the ethernet link profile to configure the gigabit ethernet switching and uplink ports. The ethernet interfaces support auto negotiation from 10BaseT to 1000BaseT as per IEEE 802.3u/z standards. When you enable auto negotiation, the device that is connected to the port is automatically configured to the highest speed supported by the device in the following order (highest to lowest):

- 10000 Mbps full duplex (supported only on the uplink interfaces)
- 1000 Mbps full duplex
- 100 Mbps full duplex
- 100 Mbps half duplex
- 10 Mbps full duplex
- 10 Mbps half duplex



The 10000 Mbps ports (10 gigabit uplink interfaces) cannot scale down to less than 1000 Mbps (1 gigabit speed).

Auto negotiation also supports the pause capabilities, automatic Media Detection Interface (MDI), and Media Detection Interface Crossover (MDIX) cable detection. The devices exchange information using the Fast link Pulse (FLP) bursts. The auto negotiation on the link is performed when you perform any of the following activities:

- Connect the device.
- Power on or reset the device at either end of the link.
- Make a negotiation request.

You can configure the ethernet link profile either using the CLI or the WebUI.

Using the WebUI

1. Navigate to the **Configuration > Ports > Ethernet** page.
2. Click **New** under the Profiles list, and enter a name for the Ethernet profile.
3. Click on the **Speed/Duplex** column and select the Speed and Duplex from the popup window.
4. Select a **Flow Control** option from the next column.
5. Select whether you need **Autonegotiation** enabled or disabled.
6. Click on the **Association** column and move the ports to the **Selected** list to apply this profile to selected ports.
7. Click **Apply**.

Using the CLI

```
(host)(config)# interface-profile enet-link-profile <profile-name>
  autonegotiation
  duplex {auto|full|half}
  speed {auto|10|100|10m_100m|1000|10000}
  flowcontrol {auto|on|off}
  no {...}
  exit
(host)(config)# interface gigabitethernet <slot/module/port>
  enet-link-profile <profile-name>
```



When the port speed is explicitly configured, the autonegotiation is disabled.

Ethernet Link Default Profile

```
(host)# show interface-profile enet-link-profile default
Ethernet Link "default"
```

```
-----
Parameter      Value
-----
Speed           auto
Duplex          auto
Autonegotiation Enabled
Flowcontrol     off
```

Sample Ethernet Link Profile Configuration

```
(host)(config)# interface-profile enet-link-profile intspd
  duplex full
  speed 1000
(host)(config)# interface gigabitethernet0/0/0
  enet-link-profile intspd
```

Verifying Ethernet Link Profile Configuration

```
(host)# show interface gigabitethernet 0/0/0
GE0/0/0 is administratively Up, Link is Down, Line protocol is Down
Hardware is Gigabit Ethernet, Address is 00:0b:86:6a:42:02
Encapsulation ARPA, Loopback not set
Configured: duplex (Auto), Speed (Auto), FC (Off), Autoneg (On)
Auto negotiation in progress
Interface index: 1
MTU 1514 bytes
```

```

Flags: Access, Trusted
Link status last changed:      0d 00:00:00 ago
Last update of counters:      0d 00:00:00 ago
Last clearing of counters:     0d 00:00:00 ago
Statistics:
    Received 0 frames, 0 octets
    0 broadcasts, 0 runts, 0 giants, 0 throttles
    0 error octets, 0 CRC frames
    0 multicast, 0 unicast
    Transmitted 0 frames, 0 octets
    0 broadcasts, 0 throttles
    0 errors octets, 0 deferred
    0 collisions, 0 late collisions
PoE Information:
    Interface: GE0/0/0, Administratively Disable, Port status: On
    Maximum power: 30000 mW, Power consumption: 0 mW
    Port voltage: 0 mV, Port current: 0 mA
    PD class: Class-0, Priority: Low, PSE port status: On

```

Ethernet Flow Control

Ethernet flow control prevents loss of frames by providing a back pressure. When an ethernet port receives frames faster than it can handle, it sends a PAUSE frame to stop the transmission from the sender for a specific period of time. The PAUSE frame has a destination group address of 01-80-c2-00-00-01.

Use the following command in the ethernet link profile to configure flow control for an ethernet port:

```
(host)(config)# [no] flow-control {on|off|auto}
```



When flow control frames are received, only pausing the transmit is supported. Sending flow control frames are not supported. This means that the system can only respond to PAUSE frames and cannot generate them. The flow-control can be enabled or disabled to respond to incoming PAUSE frames.

Loopback Interfaces

The Mobility Access Switch supports a maximum of 64 (0 to 63) loopback interfaces. You can configure the loopback interfaces using the CLI. Additionally, you can assign a secondary IP address to a loopback interface by using the **secondary** parameter.

Using the CLI

```

(host)(config)# interface loopback <0-63>
    clone <source>
    description <description>
    ip address <address> [secondary]
    no {...}
    ospf-profile
    exit

```

Sample Loopback Interface Configuration


```

(host)(config)# interface loopback 1
    description loopback01
    ip address 1.1.1.1
    exit

```

Power Over Ethernet

Power over Ethernet (PoE) as per IEEE 802.3at is a technology for wired Ethernet LANs to carry the electric-power required for the device in the data cables. You can use this technology to power IP phones, wireless LAN access points, cameras, embedded computers, thin clients, and LCDs.



The IEEE standard defined in IEEE 802.3af allows network equipment (power sourcing equipment) to provide up to 15.4 Watts of power at the output for powered devices (PDs). In addition, the IEEE 802.3at (PoE+) standard provides more power to PDs where up to 30.0 Watts of power on output is delivered on the standard copper cable. The Mobility Access Switch supports both PoE standards.

Power Management Modes

The Mobility Access Switch supports three PoE power management modes:

- **Static Mode**—The power deducted from the total power pool is the maximum power for that interface. This mode ensures that the maximum power specified by you for the interface is always reserved and cannot be shared by other PDs.
- **Dynamic Mode**—The power allocated from the total power pool for each port is the actual power consumed at that port. You can allocate any unused portion of power to the other PDs. This is the default mode.
- **Class-based Mode**—The power allocated for each port from the total power pool is the maximum power available for the class of PD connected to that port.

Power Pools

The Mobility Access Switch supports three models of power supply units (PSUs).

- **350W PSUs**—You can use the 350W PSUs for non-PoE models such as S3500-24T, S3500-24F, or S3500-48T. You can also install two 350W PSUs for system redundancy.
- **600W PSUs**—You can use the 600W PSUs for full PoE models such as S3500-24P or S3500-48P. You can also install two 600W PSUs to create a larger PoE pool.
- **1050W PSUs**—You can use the 1050W PSUs for full PoE models such as S3500-24P or S3500-48P.

Table 1 *Power Supply Pools*

Power Supply Capacity	System Power Redundancy	Power Available for PoE and PoE+Pool
350W	No	—
350W+350W	Yes	—
600W	No	400W

Table 1 *Power Supply Pools*

Power Supply Capacity	System Power Redundancy	Power Available for PoE and PoE+Pool
600W+600W	Yes	666W
1050W	No	850W
1050W+1050W	Yes	1440W

Mixed Mode PSUs

You can mix and match PSU models. The [Table 2](#) describes the various mixed mode PSU models.

Table 2 *Mixed Mode PSUs*

	350W	600W	1050W
350W	No PoE	PoE with 400W budget Not redundant for PoE	PoE with 850W budget Not redundant for PoE
600W	PoE with 400W budget Not redundant for PoE	PoE with 666W budget	PoE with 666W budget
1050W	PoE with 850W budget Not redundant for PoE	PoE with 666W budget	PoE with 1440W budget

PoE Priority

When you have power shortage in the PoE pool, you can configure PoE port priority to define which PoE ports should be provided with power while disabling power on other ports until enough power is available for all the PoE ports. Priority can be either low (default), high, or critical. When there is a power shortage, the Mobility Access Switch stops power to the low priority ports, then high priority ports, until there is enough PoE power available in the pool. If the PoE ports have the same priority, the power supply is stopped for ports with higher interface numbers, and then the lower interface numbers. For example, when there is an interface 0/0/4 and an interface 0/0/10 with the same priority, the Mobility Access Switch will stop power to the interface 0/0/10 before stopping power to the interface 0/0/4.

PoE Guard-Band

The PoE guard-band can provide protection when there is a sudden spike in the consumed power of PDs that could potentially impact other PoE enabled ports. When the guard-band is configured, the Mobility Access Switch reserves specified amount of power to prevent other PoE enabled ports from powering off and then on again. The default value for guard-band is 11,000mW. You can specify the guard-band value in steps of 1000 starting from 1000 to 30,000 milliwatts.

Configuring Power Over Ethernet

PoE/PoE+ is enabled on the Mobility Access Switch by default. It supports plug-and-play capability for PoE capable devices. You can configure PoE either using the CLI or the WebUI.

Using the WebUI

1. Navigate to the **Configuration > Ports > PoE** page.
2. Select a mode from the **Power Management Mode** drop-down list.

3. Click **Apply**, and then **Save Configuration**.



You can configure only one PoE management mode for the stack.

Using the CLI

```
(host)(config)# poe-management-profile slot <slot_num>
clone<source>
poe-powermanagement {class|dynamic|static}
poe-guardband <1000-30000 milliwatts>
no {...}
```



You can configure different PoE management modes (class/dynamic/static) on each stack member.

Sample PoE Configuration

```
(host)(config)# poe-management-profile slot 0
poe-powermanagement static
poe-guardband 15000
```

Creating and Applying a PoE Profile to an Interface

You can configure the PoE profile either using the CLI or the WebUI.

Using the WebUI

1. Navigate to the **Configuration > Ports > PoE** page.
2. Click **New** under the Profiles list, and enter a name for the PoE profile.
3. Click on the **Priority** column and select the priority from the drop-down list.
4. Enter the power in milliwatts in the **Power(/mW) Port** column.
5. Select whether the PoE state is enabled or disabled in the **State** column.
6. Select whether the Cisco compatibility is enabled or disabled in the **Cisco Legacy** column.
7. Click on the **Association** column and move the ports to the **Selected** list to apply this profile to the selected ports.
8. Click **Apply**, and then **Save Configuration**.

Using the CLI

```
(host)(config)# interface-profile poe-profile <profile-name>
cisco-compatibility
enable
poe-maxpower <milliwatts>
poe-priority {critical|high|low}
no {...}
exit
(host)(config)# interface gigabitethernet <slot/module/port>
poe-profile <profile-name>
```

Sample PoE Profile Configuration

```
(host)(config)# interface-profile poe-profile CAMERAS
poe-priority high
poe-maxpower 15000
```

```
enable
(host)(config)# interface gigabitethernet 0/0/15
poe-profile CAMERAS
```

Time Range Support for PoE

The PoE supports time range for controlling the mode of the PoE power (enable/disable) to the PoE port. The PoE port mode is enabled by the administrator.



By default, the time range profile is disabled in the poe-profile.

The PoE time range can be configured in two modes: **absolute** and **periodic**. In absolute mode, the time parameters correspond to a specific time range: start date, start time, end date, and the end time. The PoE port is enabled if the current system time is within this range. In periodic mode, the user can specify start day, start time, end day, and end time. The start day or end day can be daily, weekend, weekday, or any day of the week. The PoE port is enabled if the current day and time falls within the range.

The following are the invalid combinations for start and end values for the time range parameters in the periodic mode:

- **start-day:** daily, **end-day:** any other day other than daily
- **start-day:** weekend, **end-day:** any other day other than than weekend. (Here weekend refers to Saturday or Sunday)
- **start-day:** weekday, **end-day:** any other day other than weekday



Both the **start-time** and the **end- time** should not have identical time values if the **start-day** and the **end- day** are same.

You can configure the PoE profile using the following CLI :

```
(host)(config)# time-range-profile <profile_name>
```



As a best practice, avoid configuring the PoE time-of-day when the connected devices are in the process of being upgraded or when a power loss has rendered the connected device inoperable. In the case of an Aruba wireless Access Point, the PoE time-of-day should not be configured when an AP flash memory upgrade is in progress as it may result in potential corruption of the flash.

PoE Factory-Initial and Default Profiles

When the Mobility Access Switch is booted as factory-default and when it is booted for the first time, the poe-factory-initial profile is associated to all the ports.

```
(host)# show interface-profile poe-profile poe-factory-initial
Power over Ethernet profile "poe-factory-initial"
```

```
-----
Parameter                               Value
-----
Enable PoE interface                     Enabled
Max Power on PoE port milliwatts         30000
PoE port priority                         low
Power over Ethernet Cisco Compatibility   Disabled
time-range-profile                      N/A
```

```
(host)# show interface-profile poe-profile default
Power over Ethernet profile "default"
```

```
-----
Parameter                               Value
```

-----	-----
Disable PoE interface	Disabled
Max Power on PoE port milliwatts	30000
PoE port priority	low
Power over Ethernet Cisco Compatibility	Disabled
time-range-profile	N/A

Monitoring Power-over-Ethernet

You can use the following commands to verify the PoE configuration and monitor the PoE usage:

```
(host)# show poe interface gigabitethernet 0/0/13
GE0/0/13: Administratively Enable, Port status: Off
Maximum power: 30000 mW, Power consumption: 0 mW
Port voltage: 0 mV, Port current: 0 mA
PD class: Class-0, Priority: Low, PSE port status: Off, Time-range disable
Time-range: Periodic
```

```
Start: daily, 18:00:00 PST
```

```
End: daily, 09:00:00 PST
```

```
(host) #show poe interface brief
```

Interface	Admin	Consumption(mW)	Port Priority	Port Status
-----	-----	-----	-----	-----
GE0/0/0	Enable	0	Low	On
GE0/0/1	Enable	0	Low	On
GE0/0/2	Enable	0	Low	On
GE0/0/3	Enable	0	Low	On
GE0/0/4	Enable	0	Low	On
GE0/0/5	Enable	0	Low	On

```
(host)# show poe interface
```

```
GE3/0/0
```

```
-----
```

```
GE3/0/0: Administratively Enable, Port status: Off
```

```
Maximum power: 30000 mW, Power consumption: 0 mW
```

```
Port voltage: 0 mV, Port current: 0 mA
```

```
PD class: Class-0, Priority: Low, PSE port status: Off, PD detection in progress
```

```
GE3/0/1
```

```
-----
```

```
GE3/0/1: Administratively Enable, Port status: Off
```

```
Maximum power: 30000 mW, Power consumption: 0 mW
```

```
Port voltage: 0 mV, Port current: 0 mA
```

```
PD class: Class-0, Priority: Low, PSE port status: Off, PD detection in progress
```

```
GE3/0/2
```

```
-----
```

```
GE3/0/2: Administratively Disable, Port status: Off
```

```
Maximum power: 30000 mW, Power consumption: 0 mW
```

```
Port voltage: 0 mV, Port current: 0 mA
```

```
PD class: Class-0, Priority: Low, PSE port status: Off, CLI admin disabled
```

```
(host)# show poe
```

Port	Status	Voltage(mV)	Current(mA)	Power (mW)
----	-----	-----	-----	-----
GE0/0/0	On	0	0	0
GE0/0/1	On	0	0	0
GE0/0/2	On	0	0	0

GE0/0/3	On	0	0	0
GE0/0/4	On	0	0	0
GE0/0/5	On	0	0	0
GE0/0/6	On	0	0	0
GE0/0/7	On	0	0	0
GE0/0/8	On	0	0	0
GE0/0/9	On	0	0	0
GE0/0/10	On	0	0	0
GE0/0/11	On	0	0	0
GE0/0/12	On	0	0	0
GE0/0/13	On	0	0	0
GE0/0/14	On	0	0	0
GE0/0/15	On	0	0	0
GE0/0/16	On	0	0	0
GE0/0/17	On	0	0	0
GE0/0/18	On	0	0	0
GE0/0/19	On	0	0	0
GE0/0/20	On	0	0	0
GE0/0/21	On	0	0	0
GE0/0/22	On	0	0	0
GE0/0/23	On	0	0	0

(host)# show poe controller

Linecard	PowerBudget(W)	Power Consumption(W)	GuardBand(mW)	PoE Management
0	666	149	11000	Dynamic
1	400	130	11000	Dynamic
2	850	734	11000	Dynamic

(host)#show inventory

Show Inventory

System Card Slot : 2

SC Serial # : AW0000625 (Date: 07/15/11)

SC Model Name : ArubaS3500-48P

Mgmt Port HW MAC Addr : 00:0b:86:6c:26:41

HW MAC Addr : 00:0b:86:6c:26:40 to 00:0b:86:6c:26:7f

CPLD Version : (Rev: 11)

CPU Assembly # : 2010095E (Rev: 02.B0)

CPU Serial # : AB27010705 (Date: 07/08/11)

Fantray : Present (Version: 1)

Module 1 : Online

Module 1 Assembly # : 2010140B (Rev: 01.00)

Module 1 Serial # : UB26000012 (Date: 06/27/11)

Power Supply 0 : Present (1050W)

Power Supply 1 : Present (350W Not Supported)

System Temperature : CPU Temp : 39 C

: DPE-0 Temp : 47 C

: DPE-1 Temp : 47 C

: DDR Temp : 36 C

: USB Temp : 31 C

: Uplink Temp : 33 C

: Phy0 Temp : 23 C

: Phy1 Temp : 32 C

: Phy2 Temp : 30 C

System Voltages : 3300mV Sense : 3360.0 mV

: 2500mV Sense : 2480.0 mV

: 1800mV Sense : 1808.0 mV

```

: 1000mV Sense : 1004.0 mV
: 1100mV MAC : 1088.0 mV
: 1800mV VTT : 1840.0 mV
: 5000mV 5VDD : 5046.0 mV
: 5000mV 5VSB : 5022.0 mV
: 900mV VTT : 918.0 mV
Fantray Fan Tachometers : FAN 0 : 6367 RPM (Normal Speed)
: FAN 1 : 6585 RPM (Normal Speed)
: FAN 2 : 6459 RPM (Normal Speed)
: FAN 3 : 6490 RPM (Normal Speed)

```

(host)# show port status

```
Interface Admin Line Protocol Link PoE Trusted Mode
```

```

-----
GE0/0/0 Enable Down Down Enable Yes Access
GE0/0/1 Enable Down Down Enable Yes Access
GE0/0/2 Enable Down Down Enable Yes Access
GE0/0/3 Enable Down Down Enable Yes Access
GE0/0/4 Enable Down Down Enable Yes Access
GE0/0/5 Enable Down Down Enable Yes Access
GE0/0/6 Enable Down Down Enable Yes Access
GE0/0/7 Enable Down Down Enable Yes Access
GE0/0/8 Enable Down Down Enable Yes Access
GE0/0/9 Enable Down Down Enable Yes Access
GE0/0/10 Enable Down Down Enable Yes Access
GE0/0/11 Enable Down Down Enable Yes Access
GE0/0/12 Enable Down Down Enable Yes Access
GE0/0/13 Enable Down Down Enable Yes Access
GE0/0/14 Enable Down Down Enable Yes Access
GE0/0/15 Enable Down Down Enable Yes Access
GE0/0/16 Enable Down Down Enable Yes Access
GE0/0/17 Enable Down Down Enable Yes Access
GE0/0/18 Enable Down Down Enable Yes Access
GE0/0/19 Enable Down Down Enable Yes Access
GE0/0/20 Enable Down Down Enable Yes Access
GE0/0/21 Enable Down Down Enable Yes Access
GE0/0/22 Enable Down Down Enable Yes Access
GE0/0/23 Enable Down Down Enable Yes Access

```

A port-channel is a bundle of multiple physical interfaces that form a single logical interface. You can use port-channels to provide additional bandwidth or link redundancy between two devices. This chapter describes how to configure port-channels using the static Link Aggregation Group (LAG) and the dynamic Link Aggregation Control Protocol (LACP) methods.

This chapter includes the following topics:

- [“Important Points to Remember” on page 99](#)
- [“Creating a Port-Channel” on page 99](#)
- [“Link Aggregation Control Protocol” on page 101](#)
- [“Creating and Applying a Dynamic Port-Channel Profile to an Interface” on page 102](#)

Important Points to Remember

- A port-channel is always trusted. Any network that extends beyond the port-channel on the Mobility Access Switch must be a trusted network.
- The maximum port-channels supported per system is 64 groups; each group can be created statically or dynamically (via LACP).
- Each port-channel can have up to 8 member ports.
- The port-channel group identification (ID) range is 0 to 63 for both static and dynamic port-channels.
- The static and dynamic methods must use different group IDs and different port-channel members.
- When a port is added to a port-channel, it inherits the port-channel's properties such as VLAN membership and trunk status.
- Ports that are already assigned a feature profile cannot be part of a static or dynamic port-channel.
- All the ports in a port-channel must have the same port speed and duplex.
- There is no default LACP profile.
- For port-channel members, apart from the following profiles and parameters, all the other profiles and parameters are inherited from the port-channel configuration:
 - shutdown
 - lacp-profile
 - lldp-profile

Creating a Port-Channel

You can create port channels using the static method or the dynamic method.

- In the static method, you must first create the port-channel interface, and then add the physical interfaces to the port-channel.
- In the dynamic method, you must first create the lacp-profile and then apply the lacp-profile to the member interfaces.

Using the WebUI

1. Navigate to the **Configuration > Ports > Port Channel** page.
2. Select the **Group ID** for the port channel.
3. Select Static or LACP from the **Type** popup window and click **Ok**.
4. Click on the **Membership** column and move the ports to the **Selected** list to include the selected ports to the port channel.
5. Click **Apply** and then **Save Configuration**.

Using the CLI

```
(host) (config) #interface port-channel <0-63>
  backup [gigabitethernet <slot/module/port> | port-channel <0-63>]
  clone <source>
  description <description>
  enet-link-profile <profile-name>
  gvrp-profile <profile-name>
  igmp-snooping [ mrouter-vlan [ <vlan-list> | add <vlan-list> | delete <vlan-list>]]
  ip [access-group [in <ingress-acl> | out <egress-acl>]]
  mirroring-in-profile <profile-name>
  mirroring-out-profile <profile-name>
  mld-snooping [mrouter-vlan [<vlan-list> | add <vlan-list> | delete <vlan-list>]]
  mstp-profile <profile-name>
  mtu <64-9216>
  no
  policer-profile <profile-name>
  port-channel-members [<interface-list> | [add | delete] gigabitethernet <slot/
  module/port>]
  port-security-profile <profile-name>
  preemption [delay <10-300s> | mode [forced | off]]
  pvst-port-profile <profile-name>
  qos [trust [auto | dot1p | dscp | none]
  qos-profile <profile-name>
  shutdown
  switching-profile <profile-name>
```



For more information regarding Port-Channel profiles and configuration parameters, see the interface profiles and configuration parameters described in [“Profiles for Interfaces” on page 45](#).

Default Enet-Link Profile for Port-Channels

If you do not assign any enet-link-profile to the static or dynamic port-channel, the hidden **pc_default** profile is applied by default:

```
(show)# show interface-profile enet-link-profile pc_default
Ethernet Link "pc_default" (Predefined (editable))
-----
Parameter      Value
-----
Speed           1000
Duplex          full
Autonegotiation Enabled
Flowcontrol     off
```



Ethernet network ports and uplink ports can be combined to form a port-channel with members having different speed configuration.

Sample Static Port-Channel Configuration

```
(host)(config)# interface port-channel 1
port-channel-members gigabitethernet0/0/4,gigabitethernet0/0/5
[or]
port-channel-members add gigabitethernet 0/0/4
port-channel-members add gigabitethernet 0/0/5
exit
```

Verifying the Port-Channel Configuration

You can use the following command to verify the port-channel configuration:

```
(host)# show interface port-channel 1
port-channel 1 is administratively Up, Link is Down, Line protocol is Down
Hardware is Port-Channel, Address is 00:0b:86:6a:24:c0
Description: Link Aggregate
Member port(s):
    GE0/0/20 is administratively Up, Link is Down, Line protocol is Down
    GE0/0/21 is administratively Up, Link is Down, Line protocol is Down
    GE0/0/22 is administratively Up, Link is Down, Line protocol is Down
Speed: 0 Mbps
Interface index: 1442
MTU 1514 bytes
Flags: Access, Trusted
Link status last changed: 0d 00h:00m:00s ago
Last clearing of counters: 0d 00h:00m:00s ago
Statistics:
    Received 0 frames, 0 octets
    0 broadcasts, 0 runts, 0 giants, 0 throttles
    0 error octets, 0 CRC frames
    0 multicast, 0 unicast
    Transmitted 0 frames, 0 octets
    0 broadcasts, 0 throttles
    0 errors octets, 0 deferred
    0 collisions, 0 late collisions
```

Link Aggregation Control Protocol

The Mobility Access Switch supports Link Aggregation Control Protocol (LACP) based on the IEEE 802.3ad standard. LACP provides a standardized means for exchanging information with partner systems, to form a dynamic link aggregation group. LACP avoids port channel misconfiguration. You can define the LACP parameters in a lacp-profile, and then reference the profile in the ports to form a dynamic port-channel. A port-channel will be operationally down if all the ports in the port-channel are down.

LACP Port Modes

There are two modes in which the dynamic port-channel member interfaces can operate.

- Active mode—the interface is in active negotiating state. LACP runs on any link that is configured to be in the active state. The port in an active mode automatically initiates negotiations with other ports by initiating LACP packets.

- **Passive mode**—the interface is *not* in an active negotiating state and does not initiate negotiations. LACP runs on any link that is configured in a passive state. The port in a passive mode only responds to negotiations requests from other ports that are in an active state. .



A port in a passive state cannot set up a port-channel with another port in a passive state. Hence, to form a port-channel group between two ports, one port must be an active participant.

LACP Session Timeout and Port Priority

You can set the timeout for a LACP session. The timeout value is the amount of time that a port-channel interface waits for a LACPDU from the remote system before terminating the LACP session. The default time out value is long (90 seconds); short is 3 seconds. You can also set the port priority. The higher the value the lower the priority. The priority range is 1 to 65535 and the default is 255.

When a port in a port-channel is misconfigured (that is, the partner port is different from the other ports) or if the neighbor experiences time out or if it cannot exchange LACPDUs with the partner, then the port operational status is displayed as DOWN.



The port priority is used to select the ports that have the highest priority to form the port-channel when there are unspecified number of ports. However, only eight ports are supported in this release and hence the port priority is not useful in this release.

Creating and Applying a Dynamic Port-Channel Profile to an Interface

Using the WebUI

1. Navigate to the **Configuration > Ports > Port Channel** page.
2. Select the **Group ID** for the port channel.
3. Select LACP from the **Type** popup window.
4. Choose whether you want to select the LACP profile from a list of existing LACP profiles or you want to specify a new profile.
5. Select the LACP Profile name from the drop-down list or enter the name for the new LACP profile in the **Profile Name** text box.
6. Select the mode as passive or active from the **Mode** drop-down list.
7. Enter the priority in the **Priority** text box.
8. Select the timeout as long or short from the **Timeout** drop-down list.
9. Click on the **Membership** column and move the ports to the **Selected** list to include the selected ports to the port channel.
10. Click **Apply** and then **Save Configuration**.

Using the CLI

```
(host)(config)# interface-profile lacp-profile <profile-name>
  group-id <0-63>
  mode {active|passive}
  port-priority <1-65535>
  timeout {long|short}
  no {...}
  exit
(host)(config)# interface gigabitethernet <slot/module/port>
```

```
lACP-profile <profile-name>
```

Sample Dynamic Port-Channel Configuration

```
(host)(config)# interface-profile lACP-profile LACP_2
  group-id 2
  mode active
  exit
(host)(config)# interface gigabitEthernet 0/0/0
  lACP-profile LACP_2
  exit
(host)(config)# interface gigabitEthernet 0/0/1
  lACP-profile LACP_2
  exit
```

Verifying Port-Channel Configuration

```
(host)# show interface port-channel 2
port-channel 0 is administratively Up, Link is Down, Line protocol is Down
Hardware is Port-Channel, LACP enabled, Address is 00:0b:86:6a:25:40
Description: Link Aggregate
Member port(s):
  GE0/0/0 is administratively Up, Link is Down, Line protocol is Down
  GE0/0/1 is administratively Up, Link is Down, Line protocol is Down
Speed: 0 Mbps
Interface index: 1443
MTU 1514 bytes
Flags: Access, Trusted
Link status last changed: 0d 04h:10m:27s ago
Last clearing of counters: 0d 00h:00m:02s ago
Statistics:
  Received 0 frames, 0 octets
  0 broadcasts, 0 runts, 0 giants, 0 throttles
  0 error octets, 0 CRC frames
  0 multicast, 0 unicast
  Transmitted 0 frames, 0 octets
  0 broadcasts, 0 throttles
  0 errors octets, 0 deferred
  0 collisions, 0 late collisions
```

Verifying Port-Channel Neighbor Information

```
(host) #show lACP 2 neighbor

Flags: S - Device is requesting slow LACPDUs
       F - Device is requesting fast LACPDUs
       A - Device is in Active mode P - Device is in Passive mode
LACP Neighbor Table
-----
Port      Flags  Pri  OperKey  State  Num  Dev Id
----
GE0/0/0   SP     0    0x0      0x0    0x0  00:00:00:00:00:00
GE0/0/1   SP     0    0x0      0x0    0x0  00:00:00:00:00:00
```

Verifying Port-Channel Internal (Local) Information

```
(host) #show lACP 2 internal

Flags: S - Device is requesting slow LACPDUs
       F - Device is requesting fast LACPDUs
       A - Device is in Active mode P - Device is in Passive mode
```

LACP Internal Table

Port	Flags	Pri	AdminKey	OperKey	State	Num	Status
GE0/0/0	SA	255	0x3	0x3	0x5	0x7	down
GE0/0/1	SA	255	0x3	0x3	0x5	0x8	down

Verifying Port-Channel Counters Information

(host) #show lacp 2 counters

LACP Counter Table

Port	LACPDUTx	LACPDURx	MrkrTx	MrkrRx	MrkrRspTx	MrkrRspRx	ErrPktRx
GE0/0/0	0	0	0	0	0	0	0
GE0/0/1	0	0	0	0	0	0	0

Operations, Administration, and Maintenance (OAM) refers to the tools and utilities to install, monitor, and troubleshoot a network. This implementation of OAM complies with the IEEE 802.3ah standard and is able to report layer-2 network behavior. This helps network administrators monitor troubleshoot a network without sending technicians into the field to diagnose problems on location. OAM provides mechanisms to monitor link operation and health, and improve fault isolation.

The MAS OAM supports the following Link Fault Management Functionalities:

- **Discovery** – OAM-enabled local interface discovers remote interface enabled with OAM and notifies each other of own capabilities. After discovery, both sides send OAM PDUs periodically to monitor the link.
- **Remote fault detection** – Detection and handling of faulty link such as not receiving OAM PDU from the other peer within configured time-out or OAM PDU with “link-fault” flag.
- **Remote loopback** – Link segment testing controlled remotely using test frames. Usually remote loopback used during installation or for troubleshooting.

OAM is disabled by default. To enable OAM, you must create an OAM profile and apply it to a physical interface.

Creating an OAM Profile

OAM parameters are set by creating an OAM profile, which is a new type of interface profile.

```
(host) (config) # interface-profile oam-profile <oam-profile-name>
(host) (OAM profile "<oam-profile-name>") # ?
allow-loopback          Support OAM local loopback
clone                   Copy data from another OAM profile
discovery-mode          OAM discovery mode
link-fault-action       Action taken on link-fault detection
link-timeout            Timeout in seconds to declare link fault
no                       Delete Command
pdu-rate                Maximum OAM PDUs sent per second
remote-loopback         Put remote device into loopback mode
```

Table 1 OAM Profile Parameters Default Values

Parameter	Possible Values	Default Value
discovery-mode	Active, Passive	Active
remote-loopback	Enable, Disable	Disable
allow-loopback	Enable, Disable	Disable
pdu-rate	1 to10	5
link-timeout	2 to10	5
link-fault-action	Syslog, Error-disable	Error-disable

Sample Configuration

```
(host) (OAM profile "oam1") #allow-loopback
(host) (OAM profile "oam1") #link-fault-action syslog
(host) (OAM profile "oam1") #link-timeout 3
(host) (OAM profile "oam1") #pdu-rate 8

(host) (OAM profile "oam1") #show interface-profile oam-profile oam1

OAM profile "oam1"
-----
Parameter                                Value
-----
OAM discovery mode                        active
OAM remote-loopback                      Disabled
OAM local-loopback                       Enabled
OAM PDU rate (PDU per second)            8
OAM link-fault timeout (seconds)         3
OAM link-fault action                     syslog
```

Applying an OAM Profile

Once you've created an OAM profile, you must apply it to physical interfaces.

```
(host) (config) #interface gigabitethernet 0/0/1
(host) (gigabitethernet "0/0/1") #oam-profile <oam-profile-name>
(host) (config) #interface gigabitethernet 0/0/2
(host) (gigabitethernet "0/0/2") #oam-profile <oam-profile-name>
```



You cannot simultaneously apply both OAM and tunneled node settings to an interface.



An OAM profile must be applied to each port channel member interface.

Applying OAM to each Port Channel Member

In this first example, the output of the **show interface port channel** command identifies **GE0/0/12** and **GE0/0/13** as member ports of port channel 4:

```
(host) (config) #show interface port-channel 4
port-channel 4 is administratively Up, Link is Up, Line protocol is Up
Hardware is Port-Channel, LACP enabled, Address is 00:0b:86:6a:70:c0
Description: Link Aggregate
Member port(s):
    GE0/0/12 is administratively Up, Link is Up, Line protocol is Up
    GE0/0/13 is administratively Up, Link is Up, Line protocol is Up
Speed: 2 Gbps
Interface index: 1445
MTU 1514 bytes
Flags: Access, Trusted
Link status last changed: 0d 02h:25m:57s ago
Last clearing of counters: 0d 02h:25m:57s ago
Statistics:
    Received 4973595 frames, 1272848056 octets
    668 pps, 1.383 Mbps
```

```

    32 broadcasts, 0 runs, 0 giants, 0 throttles
0 error octets, 0 CRC frames
13602 multicast, 4959961 unicast
Transmitted 23674 frames, 6226872 octets
0 pps, 0 bps
39 broadcasts, 0 throttles
0 errors octets, 0 deferred
0 collisions, 0 late collisions

```

The commands in the example below apply an OAM profile to Port Channel Members GE0/0/12 and GE0/0/13:

```

(host) (config) #interface gigabitethernet 0/0/12
(host) (gigabitethernet "0/0/12") #oam-profile oam1
(host) (gigabitethernet "0/0/12") #interface gigabitethernet 0/0/13
(host) (gigabitethernet "0/0/13") #oam-profile oam1
(host) (gigabitethernet "0/0/13") #

```

Related Show Commands

The following show commands display the status of OAM on your Mobility Access Switches.

The **show oam brief** command displays a quick overview of the ports on which OAM is enabled.

Interface	OAM Mode	Link-fault Action	Loopback Local	Loopback Remote	Link State	Oper State	Remote MAC
GE0/0/1	Active	Syslog	Enable	Disable	Up	Up	00:0b:86:6a:4f:04
GE0/0/2	Active	Syslog	Enable	Disable	Up	Up	00:0b:86:6a:4f:03

The **show oam counters** command displays the total PDUs received and transmitted, as well as the number of errors, on OAM-enabled ports.

Interface	Total PDU Received	Error PDU Received	Unknown PDU Received	Total PDU Transmitted	Transmit Discarded
GE0/0/1	295	0	0	295	0
GE0/0/2	295	0	0	295	0

Use the **clear counters oam** command to clear any OAM counters:

```

(host) #clear counters oam

```

The **show oam interface gigabitethernet** command displays the OAM profile and status on a specific port:

```

show oam interface gigabitethernet <slot/port/module>
GE0/0/1 is operationally Up, Link is Up
  OAM link-fault action is syslog
  Local loopback is Enable, Remote loopback is Disable
  OAM PDU rate is 8, Link timeout is 3
Local:
  MAC address is 00:0b:86:6a:4f:03, PDU size is 64
  MUX state is Forward, Parser state is Forward
  Discovery mode is Active, Discovery state Completed
  Local is stable, Locat is satisfied
Remote:
  MAC address is 00:0b:86:6a:4f:04, PDU size is 64
  MUX state is Forward, Parser state is Forward
  Discovery mode is Active
  Remote is stable, Remote is valid

```


The Mobility Access Switch supports IEEE 802.1Q VLANs. It supports MAC-based VLANs, tag-based VLANs, port-based VLANs, and voice VLANs. You can optionally configure an IP address and netmask for a VLAN for inband management.

This chapter includes the following topics:

- “VLANs Overview” on page 109
- “Creating VLANs” on page 109
- “Creating and Applying a Switching Profile to an Interface” on page 111
- “Managing the MAC Address Table” on page 114
- “VLAN Profile” on page 115

VLANs Overview

The Mobility Access Switch supports the following types of VLANs:

- **MAC-based VLANs**—In the case of untrusted interfaces, you can associate a client to a VLAN based on the source MAC of the packet. Based on the MAC, you can assign a role to the user after authentication. For more information about how to assign MAC-based VLANs, see [Chapter 31, “MAC-Based Authentication”](#) on page 265.
- **Port-based VLANs**—In the case of trusted interfaces, all untagged traffic is assigned a VLAN based on the incoming port.
- **Tag-based VLANs**—In the case of trusted interfaces, all tagged traffic is assigned a VLAN based on the incoming tag.
- **Voice VLANs**—You can use the voice VLANs to separate voice traffic from data traffic when the voice and data traffic are carried over the same ethernet link. For more information on Voice VLANs, see [Chapter 12, “VoIP”](#) on page 131.

Creating VLANs

By default, all the ports in the Mobility Access Switch are assigned to VLAN 1. You can create VLANs and assign ports to them.

Using the WebUI

1. Navigate to the **Configuration > VLANs** page.
2. Click **New** under the VLANs list.
3. Enter the VLAN **ID**.
4. Enter a Description for the VLAN.
5. Click **Apply** and then **Save Configuration**.

Using the CLI

```
(host)(config)# vlan <id>
aaa-profile <profile-name>
```

```

clone <source>
description <name>
igmp-snooping-profile <profile-name>
mac-address-table static <mac-address> gigabitethernet <slot/module/port>
mac-aging-time <minutes>
mld-snooping-profile <profile-name>
no {...}
pvst-profile <profile-name>
exit

```

Sample VLAN Configuration

```

(host)(config)# vlan 100
    description Faculty
    exit
(host)(config)# vlan 200
    description Students
    exit

```

Verifying VLAN Configuration

You can verify the VLANs created and the ports assigned to the VLANs using the following commands:

```

(host)# show vlan
VLAN CONFIGURATION

```

```

-----
VLAN  Description                      Ports
----  -
1      All                             GE0/0/0-1 GE0/0/7 GE0/0/9-29 GE0/0/33
                                           GE0/0/35-41 GE0/0/44-47
100    Faculty                         GE0/0/0
101    Student                         GE0/0/0
102    Admin                           GE0/0/0
103    Finance                         GE0/0/0
104    HR                             GE0/0/0
105    Engineering                     GE0/0/0
106    QA                             GE0/0/0
107    Support                         GE0/0/0
108    Marketing                       GE0/0/0
109    Management                       GE0/0/0

```

```

(host)# show vlan detail

```

U - Untagged member, T - Tagged member

* - Active interface

Dot1q tag: 1, Description: VLAN0001

Number of interfaces: 36, Active: 5

VLAN membership:

Access:

```

GE0/0/1(U) GE0/0/7(U) GE0/0/9*(U) GE0/0/10*(U)
GE0/0/11(U) GE0/0/12(U) GE0/0/13(U) GE0/0/14(U)
GE0/0/15(U) GE0/0/16(U) GE0/0/17(U) GE0/0/18(U)
GE0/0/19(U) GE0/0/20(U) GE0/0/21(U) GE0/0/22(U)
GE0/0/23(U) GE0/0/24(U) GE0/0/25(U) GE0/0/26(U)
GE0/0/27(U) GE0/0/28(U) GE0/0/29(U) GE0/0/33(U)
GE0/0/35(U) GE0/0/36(U) GE0/0/37(U) GE0/0/38(U)
GE0/0/39(U) GE0/0/40(U) GE0/0/41(U) GE0/0/44(U)
GE0/0/45*(U) GE0/0/46*(U) GE0/0/47*(U)

```

Trunk:

```

GE0/0/0(U) GE0/0/0(T)

```

```

Dot1q tag: 100, Description: Faculty
Number of interfaces: 1, Active: 0
VLAN membership:
  Trunk:
    GE0/0/0(T)

(host)# show vlan extensive
Dot1q tag: 1, Description: VLAN0001
IGMP-snooping profile name: igmp-snooping-factory-initial
IGMP-snooping: Enabled
IGMP-snooping proxy: Disabled
MSTP instance: 0
MAC aging time: 5 minutes
Number of interfaces: 36, Active: 5
VLAN membership:
  GE0/0/0      Trunk  Trusted  Untagged
  GE0/0/0      Trunk  Trusted  Tagged
  GE0/0/1      Access Trusted  Untagged
  GE0/0/7      Access Trusted  Untagged
  GE0/0/9*     Access Trusted  Untagged
  ....
Dot1q tag: 100, Description: Faculty
MSTP instance: 0
MAC aging time: 300
Number of interfaces: 1, Active: 0
VLAN membership:
  GE0/0/0      Trunk  Trusted  Tagged

(host)#show vlan summary
Number of tunneled-node VLANs           :2
Number of operational VLANs             :10

```

Creating and Applying a Switching Profile to an Interface

You can assign VLAN membership to the interface using the switching profile. The switching profile has the following types of configurations for a port:

- **Switch-Port Mode**—Specifies whether the port is an access port connected to an end device or a trunk port for uplink connectivity.
- **Access VLAN**—Specifies the VLAN ID for the port, when the switch-port mode is access.
- **Native VLAN**—Specifies the VLAN for incoming untagged packets, when the switch-port mode is trunk. When a packet goes out of a trunk interface in native VLAN, it will be untagged. By default, VLAN 1 is the native VLAN. The native VLAN should be part of the trunk allowed VLANs.
- **Trunk Allowed VLANs**—Identifies the VLAN IDs for which the trunk carries the traffic.

Using the WebUI

1. Navigate to the **Configuration > Ports > Switching** tab.
2. Under the profiles list, click **New**.
3. Enter a name for the new switching profile under the **Name** column.
4. Select a mode from the drop-down list. It can be either trunk or access.
5. If you selected the mode as access, select the Access VLAN from the drop-down list. Only the VLANs created already are listed.

6. If you selected the mode as trunk, select the Native VLAN from the drop-down list. Only the VLANs created already are listed.
7. If you selected the mode as Trunk, select the trunk allowed VLANs from the Allowed VLAN column.
8. Select the interfaces that are part of this VLAN in the Association column.
9. Click **Apply** and then **Save Configuration**.

Using the CLI

```
(host)(config)# interface-profile switching-profile <profile-name>
    access-vlan <VLAN-ID>
    clone <source>
    native-vlan <VLAN-ID>
    switchport-mode {access|trunk}
    trunk allowed vlan [add|all|except|remove] <VLANs-List>
    storm-control-bandwidth <50-100>
    storm-control-broadcast
    storm-control-multicast
    storm-control-unknown-unicast
    no {...}
    exit
(host)(config)# interface gigabitethernet <slot/module/port>
    switching-profile <profile-name>
```



If you do not specify a switch-port mode, the port will be in switch-port mode access implicitly. In the case of switchport-mode trunk, the native vlan has to be in the allowed vlan list if you want the port to receive and transmit on the native vlan.

Default Switching Profile

```
(host)# show interface-profile switching-profile default
```

```
switching profile "default"
```

Parameter	Value
-----	-----
Switchport mode	access
Access mode VLAN	1
Trunk mode native VLAN	1
Enable broadcast traffic rate limiting	Enabled
Enable multicast traffic rate limiting	Disabled
Enable unknown unicast traffic rate limiting	Enabled
Max allowed rate limit traffic on port in percentage	50
Trunk mode allowed VLANs	1-4094

Sample Access Port Configuration

You can use the following steps to configure an interface as an access port that belongs to a particular VLAN:

1. Create a switching profile.
2. Apply the switching profile to the interface.

To configure a switching profile with access VLAN 200, use the following commands:

```
interface-profile switching-profile Student
    access-vlan 200
```

To apply the switching-profile to the interface (gigabitethernet 0/0/10), use the following commands:


```
interface gigabitethernet 0/0/10
  switching-profile Student
exit
```

Verifying the Switching Profile Configuration for the Interface

To verify the configuration, use one of the following commands:

```
(host) #show vlan
VLAN CONFIGURATION
-----
VLAN Description Ports
-----
1      VLAN0001      GE 0/0/0 GE 0/0/1 GE 0/0/11 GE 0/0/12
                        GE 0/0/13 GE 0/0/14 GE 0/0/15 GE 0/0/16
                        GE 0/0/17 GE 0/0/18 GE 0/0/19 GE 0/0/2
100    Faculty
200    Student GE 0/0/10

(host) #show interface gigabitethernet 0/0/0 switchport extensive
GE0/0/0
Link is Up
Flags: Access, Trusted

VLAN membership:

VLAN tag  Tagness  STP-State
-----  -
1          Untagged  FWD
```

Sample Trunk Port Configuration

To configure a trunk port, the switch-port mode should be set as trunk. To define the switching profile, use the following commands:

```
interface-profile switching-profile Upstream
  switchport-mode trunk
```

To apply the switching profile to the trunk ports, use the following commands:

```
interface gigabitethernet 0/0/11
  switching-profile Upstream
```

For trunk ports, there are times when the other side of the link requires traffic to be sent without any tags. This functionality is commonly referred as native VLAN. For this purpose, you can use the native-vlan parameter in the switching-profile:

```
interface-profile switching-profile Upstream
  native-vlan 100
```

By default, a trunk port allows all VLANs to be transported. You can change the allowed VLANs using the trunk allowed vlan parameter in the switching profile:

```
interface-profile switching-profile Upstream
  trunk allowed vlan all
```

Verifying the Trunk Configuration

You can use the following command to view the trunk configuration:

```
(host)# show trunk
Trunk Port Table
-----
Port          Vlans Allowed  Vlans Active   Native Vlan
```

----	-----	-----	-----
GE 0/0/11 ALL	1,100,200	100	
GE 0/0/12 2-45	2,30	45	

Managing the MAC Address Table

The Mobility Access Switch populates the MAC address table as a result of dynamic learning, static addition, and authentication process. These MACs are referred to as learnt, static, and auth MACs respectively. You can manage the MAC address table using the following tasks:

- “Adding Static MAC Addresses” on page 114
- “Displaying the MAC Address Table” on page 114
- “Deleting the Static MACs” on page 115
- “Clearing the Learnt MACs” on page 115
- “Configuring the MAC Aging Time” on page 115

Adding Static MAC Addresses

You can add static MAC addresses to a VLAN and thus to the MAC address table.

```
(host)(config)# vlan <vlan-id>
    mac-address-table static <mac-address> gigabitethernet <slot/module/port>
```

Example Configuration

```
(host)(config)# vlan 700
    description "vlan 700"
    aaa-profile default
    mac-aging-time 10
    mac-address-table static 00:01:02:03:04:05 gigabitethernet 0/0/14
    mac-address-table static 0a:0b:0c:0d:4e:0f gigabitethernet 0/0/16
(host)(config)# show vlan-config 700
VLAN "700"
-----
Parameter                Value
-----
Description              vlan 700
aaa-profile              default
igmp-snooping-profile    N/A
mld-snooping-profile     N/A
pvst-bridge-profile      predefinedprofile
MAC Aging time(Minutes)  10
Static mac address       00:01:02:03:04:05 gigabitethernet 0/0/14
Static mac address       0a:0b:0c:0d:4e:0f gigabitethernet 0/0/16
```

Displaying the MAC Address Table

```
(host)# show mac-address-table
Total MAC address: 2
Learnt: 1, Static: 1, Auth: 0
MAC Address Table
-----
Destination Address  Address Type  VLAN  Destination Port
-----
00:0b:86:0f:0a:80   Learnt       0226  GE0/0/42
00:10:db:00:00:11   Static       0201  GE0/0/0
```

```
(host)# show mac-address-table interface gigabitethernet 0/0/19
```

```
Total MAC address: 1
```

```
Learnt: 1, Static: 0, Auth: 0
```

```
MAC Address Table
```

```
-----
```

Destination Address	Address Type	VLAN	Destination Port
00:0c:34:46:f2:52	Learnt	0100	GE0/0/19

```
(host)#show mac-address-table summary
```

```
Total MAC address: 3
```

```
Learnt: 3, Static: 0, Auth: 0
```

```
(host)# show mac-address-table vlan 700
```

```
Total MAC address: 5
```

```
Learnt: 0, Static: 5, Auth: 0
```

```
MAC Address Table
```

```
-----
```

Destination Address	Address Type	VLAN	Destination Port
00:01:02:03:04:05	static	700	GE0/0/14
00:01:02:03:44:05	static	700	GE0/0/16
00:00:02:03:44:05	static	700	GE0/0/16
00:00:00:03:44:05	static	700	GE0/0/16
00:00:00:03:54:05	static	700	GE0/0/16

Deleting the Static MACs

You can use the following command to delete the static MAC addresses from the MAC address table:

```
(host)(config)# vlan <vlan-id>  
no mac-address-table static <mac-address>
```

Clearing the Learnt MACs

You can use the following commands to clear the learnt MACs from the MAC address table:

```
(host)(config)# clear mac-address-table  
(host)(config)# clear mac-address-table interface gigabitethernet 0/0/5  
(host)(config)# clear mac-address-table vlan 20  
(host)(config)# clear mac-address-table vlan 20 interface gigabitethernet 0/0/0
```

Configuring the MAC Aging Time

In the case of learnt MACs, you can configure the system to prune the MAC address if it does not get refreshed within the specified MAC aging time. The default value is 5 minutes. Use the following command to specify the MAC aging interval per VLAN:

```
(host)(config)# vlan <vlan-id>  
mac-aging-time <minutes>
```

VLAN Profile

A VLAN Profile (as opposed to interface profile) can be created to enable/modify IGMP-Snooping, MLD-Snooping and PVST settings. You can use the `vlan-profile` command followed by the particular feature.

```
(host) (config) #vlan-profile  
igmp-snooping-profile
```

```
mld-snooping-profile  
pvst-profile
```

For more information on configuring and applying IGMP Snooping profile to a VLAN, see [“Creating and Applying an IGMP Snooping Profile to a VLAN” on page 201](#).

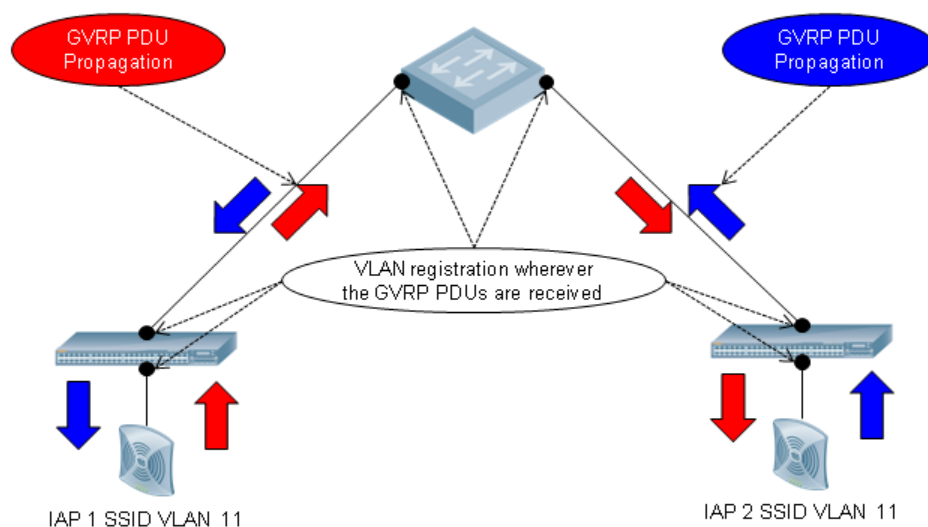
For more information on configuring and applying MLD Snooping profile to a VLAN, see [“Configuring MLD Snooping” on page 205](#)

For more information on configuring and applying PVST profile to a VLAN, see [“Configuring PVST+” on page 153](#).

This chapter includes the following topics:

- ## GVRP Overview

Figure 1 *GVRP Overview*



To enable GVRP in the Mobility Access Switch, you must configure the following two profiles and attach them to a trunk port:

- 
- NOTE**

You can enable GVRP only on trunk ports.

You can use the following CLI commands to define the GVRP global profile settings.

```
(host)(config)# gvrp
(host)(Global GVRP configuration)# enable
(host)(Global GVRP configuration)# join-time <milliseconds>
```

The join period timer controls the interval between the transmit PDU events that are applied to the applicant state machine. Default is 200 milliseconds.

```
(host)(Global GVRP configuration)# leave-time <milliseconds>
```

The leave period timer controls the period of time that the registrar state machine waits in the leaving state before transmitting to the empty state. Default is 600 milliseconds.

```
(host)(Global GVRP configuration)# leave-all-time <milliseconds>
```

The leave all period timer controls the frequency with which the leave all state machine generates LeaveAll PDUs. Default is 10000 milliseconds.

You can use the following CLI commands to define the interface specific gvrp-profile:

```
(host)(config)# interface-profile gvrp-profile <profile_name>
(host)(Interface GVRP profile <profile_name># registrar-mode [normal|forbidden]
```

In normal registrar mode, the Mobility Access Switch registers and de-registers VLANs to or from its connected switches and IAPs. In forbidden registrar mode, the Mobility Access Switch cannot register nor de-register VLANs to or from its connected switches and IAPs. Default is registrar-mode normal.

Sample Configurations

To enable and configure GVRP globally:

```
(host)(config)# gvrp
(host)(Global GVRP configuration)# enable
(host)(Global GVRP configuration)# join-time 200
(host)(Global GVRP configuration)# leave-time 600
(host)(Global GVRP configuration)# leave-all-time 10000
```

To enable and configure GVRP profile on an interface:

```
(host)(config)# interface-profile gvrp-profile Enable-GVRP
(host)(Interface GVRP profile "Enable-GVRP")# enable
(host)(Interface GVRP profile "Enable-GVRP")# registrar-mode normal
```

To attach GVRP profile to the interface:

```
(host) (config) # interface gigabitethernet 0/0/10
(host) (gigabitethernet "0/0/10") # gvrp-profile gvrp
```

The following example displays global GVRP status and current timer values:

```
(host) (config) #show gvrp-global-profile
```

Global GVRP configuration

```
-----
Parameter      Value
-----
GVRP status     Enabled
Join Time       200
Leave Time       600
Leave-all Time   10000
```

The following example displays the interfaces in which gvrp is enabled:

```
(host) (config) #show gvrp interfaces
```

Interface GVRP info

Interface	State	Registrar Mode
-----	-----	-----
gigabitethernet0/0/10	Enabled	Normal
gigabitethernet0/0/20	Disabled	N/A
port-channel1	Disabled	N/A

The Mobility Access Switch supports Link Layer Discovery Protocol (LLDP) to advertise identity information and capabilities to other nodes on the network, and store the information discovered about the neighbors. LLDP is also used to implement Voice VLAN configurations. For more information on Voice VLAN configuration, see [Chapter 12, “VoIP” on page 131](#).

This chapter contains the following major sections:

- [“Important Points to Remember” on page 121](#)
- [“LLDP” on page 121](#)
- [“LLDP-MED” on page 125](#)
- [“Proprietary Link Layer Discovery Protocols” on page 127](#)

Important Points to Remember

- Power via MDI and MAC/PHY Configuration/Status TLVs are not currently supported.
- Inventory-management, and extended power via MDI TLVs are not currently supported.
- LLDP-MED must be enabled to advertise a VOIP VLAN.

LLDP

This section contains the following sections:

- [“Understanding LLDP” on page 121](#)
- [“Configuring LLDP” on page 122](#)

Understanding LLDP

Link Layer Discovery Protocol (LLDP), defined in the IEEE 802.1AB standard, is a Layer 2 protocol that allows network devices to advertise their identity and capabilities on a LAN. The Mobility Access Switch supports a simple one-way neighbor discovery protocol with periodic transmissions of LLDP PDU.

- LLDP frames are constrained to a local link.
- LLDP frames are TLV (Type-Length-Value) form.
- LLDP Multicast address is 01-80-C2-00-00-0E.

LLDP Factory Initial and Default Profiles

This section contains the following sections:

- [“LLDP Factory Initial Profile” on page 122](#)
- [“Default LLDP Profile” on page 122](#)

LLDP Factory Initial Profile

When the Mobility Access Switch is booted as factory-default for the first time, the "lldp-factory-initial" profile is associated to all the ports.

To display this information, use the following command:

```
(host)# show interface-profile lldp-profile lldp-factory-initial
LLDP Profile "lldp-factory-initial"
-----
Parameter                                Value
-----
LLDP pdu transmit                        Enabled
LLDP protocol receive processing         Enabled
LLDP transmit interval (Secs)           30
LLDP transmit hold multiplier           4
LLDP fast transmit interval (Secs)      1
LLDP fast transmit counter              4
LLDP-MED protocol                       Enabled
Control proprietary neighbor discovery  Disabled
```

Default LLDP Profile

To display the default lldp profile information, use the following command:

```
(host)# show interface-profile lldp-profile default

LLDP Profile "default"
-----
Parameter                                Value
-----
LLDP pdu transmit                        Disabled
LLDP protocol receive processing         Disabled
LLDP transmit interval (Secs)           30
LLDP transmit hold multiplier           4
LLDP fast transmit interval (Secs)      1
LLDP fast transmit counter              4
LLDP-MED protocol                       Disabled
Control proprietary neighbor discovery  Disabled
```



When you use the default LLDP profile, the RX and TX parameters are disabled. You have to explicitly enable them for LLDP to work.

Configuring LLDP

- [“Configuring an LLDP Profile” on page 122](#)
- [“Applying LLDP Profile to an Interface” on page 123](#)

Configuring an LLDP Profile

To configure an LLDP profile, use the following command:

```
(host)(config)# interface-profile lldp-profile <profile-name>
```

```

clone <source>
lldp fast-transmit-counter <1-8>
lldp fast-transmit-interval <1-3600>
lldp receive
lldp transmit
lldp transmit-hold <1-100>
lldp transmit-interval <1-3600>}
no {...}
exit

```

Applying LLDP Profile to an Interface

To apply an LLDP profile to an interface, use the following command:

```

(host)(config)# interface gigabitethernet <slot/module/port>
    lldp-profile <profile-name>.

```



In the case of static and dynamic port-channels, the LLDP profile must be applied to the member interfaces.

Verifying LLDP Profile Configuration

```

(host)# show interface-profile lldp-profile <profile-name>
LLDP Profile " profile-name "
-----
Parameter                                         Value
-----
LLDP pdu transmit                               Enabled
LLDP protocol receive processing                 Enabled
LLDP transmit interval (Secs)                   30
LLDP transmit hold multiplier                   4
LLDP fast transmit interval (Secs)              30
LLDP fast transmit counter                      1
LLDP-MED protocol                              Disabled
Control proprietary neighbor discovery          Disabled

```

Monitoring LLDP

This section describes commands for monitoring LLDP. It contains the following sections:

- “Display LLDP Interface” on page 123
- “Display LLDP Interface <interface>” on page 124
- “Display LLDP Neighbor” on page 124
- “Display LLDP Neighbor Interface Detail” on page 124
- “Display LLDP Statistics” on page 125
- “Display LLDP Statistics Interface” on page 125

Display LLDP Interface

To display all LLDP information for all interfaces, use the following command:

```

(host)# show lldp interface
LLDP Interfaces Information
-----
Interface  LLDP TX  LLDP RX  LLDP-MED  TX interval  Hold Timer
-----
GE0/0/0    Enabled  Enabled  Enabled   30           120
GE0/0/1    Enabled  Enabled  Enabled   30           120
GE0/0/2    Enabled  Enabled  Enabled   30           120

```

```

GE0/0/3      Enabled  Enabled  Enabled  30          120
GE0/0/4      Enabled  Enabled  Enabled  30          120
GE0/0/5      Enabled  Enabled  Enabled  30          120
GE0/0/6      Enabled  Enabled  Enabled  30          120
GE0/0/7      Enabled  Enabled  Enabled  30          120
GE0/0/8      Enabled  Enabled  Enabled  30          120
GE0/0/9      Enabled  Enabled  Enabled  30          120
GE0/0/10     Enabled  Enabled  Enabled  30          120

```

<output truncated>

Display LLDP Interface <interface>

To display LLDP information for a specific interface, use the following command:

```
(host) #show lldp interface gigabitethernet 0/0/1
```

```

Interface: gigabitethernet0/0/1
LLDP Tx: Enabled, LLDP Rx: Enabled
Proprietary Neighbor Discovery: Disabled
LLDP-MED: Enabled
Fast Transmit interval: 1, Fast Transmit message counter: 4
Transmit interval: 30, Hold timer: 120

```

Display LLDP Neighbor

```

(host)#show lldp neighbor
Capability codes: (R)Router, (B)Bridge, (A)Access Point, (P)Phone, (O)Other
LLDP Neighbor Information

```

```

-----
Local Intf Chassis ID      Capability  Remote Intf  Expiry-Time (Secs)
-----
GE4/0/1    00:0b:86:6a:25:40    B:R         GE0/0/17     105
GE4/0/2    00:0b:86:6a:25:40    B:R         GE0/0/18     105

```

System name

```

-----
ArubaS3500
ArubaS3500

```

Number of neighbors: 2



NOTE

To view proprietary neighbors, use the **show neighbor-devices** command.

Display LLDP Neighbor Interface Detail

```
(host)# show lldp neighbor interface gigabitethernet 0/0/5 detail
```

```
Interface: GE0/0/5, Number of neighbors: 1
```

```

-----
Chassis id: 00:26:88:63:35:c0, Management address: 10.16.48.24
Interface description: ge 0/0/5.0, ID: 515, MTU: 1514
Last Update: Sat Apr 16 03:24:48 2011
Time to live: 120, Expires in: 117 Secs
System capabilities : Bridge,Router
Enabled capabilities: Bridge,Router
System name: EX-4200
System description:
  Juniper Networks, Inc. ex4200-24t , version 10.4R1.9 Build date: 2010-12-04 10:09:46
  UTC
Link aggregation id: 0

```

```

Auto negotiation: Supported, Enabled
Autoneg capability:
  10Base-T, HD: no, FD: yes
  100Base-T, HD: yes, FD: yes
  100Base-X, HD: no, FD: yes
  1000Base-T, HD: yes, FD: yes
Media attached unit type: Not advertised
LLDP-MED:
Device Type: Network Connectivity Device
Capability: LLDP-MED capabilities, Network policy, Location ID, Extended power via
MDI/PSE
Extended Power-over-Ethernet:
  Power Type & Source: PSE Device
  Power Source: Primary power source
  Power Priority: low
  Power Value: 0

```

Display LLDP Statistics

```
(host)# show lldp statistics
```

```
LLDP Statistics
```

```

-----
Interface  Received  Unknow TLVs  Malformed  Transmitted
-----
GE0/0/0    0          0             0           0
GE0/0/1    0          0             0           0
GE0/0/2    0          0             0           0
GE0/0/3    0          0             0           0
GE0/0/4    0          0             0           0
GE0/0/5    4          2             0           4
GE0/0/6    0          0             0           0
GE0/0/7    0          0             0           0
GE0/0/8    0          0             0           0
GE0/0/9    0          0             0           0
GE0/0/10   0          0             0           0
<output truncated>

```

Display LLDP Statistics Interface

```
(host)# show lldp statistics interface gigabitethernet 0/0/0
```

```
LLDP Statistics
```

```

-----
Interface                Received  Unknow TLVs  Malformed  Transmitted
-----
gigabitethernet0/0/0    0          0             0           0

```

LLDP-MED

This section contains the following sections:

- Understanding LLDP-MED
- Configuring LLDP-MED
- Verifying LLDP-MED

Understanding LLDP-MED

LLDP-MED (media endpoint devices) is an extension to LLDP developed by TIA (ANSI/TIA-1057) to support interoperability between VoIP end-point devices and other networking end-devices. LLDP-MED is focused mainly on discovery running between network devices and end-points such as IP phones.

Configuring LLDP-MED

LLDP-MED network policy discovery lets end-points and network devices advertise their VLAN IDs (e.g. voice VLAN), IEEE 802.1p, and DSCP values. The Mobility Access Switch switch can instruct end-devices to modify their settings to match VoIP requirements.

To configure the LLDP profile to enable LLDP-MED, use the following command:

```
(host)(config)# interface-profile lldp-profile <profile-name>
    lldp transmit
    lldp receive
    med enable
(host)(config)# interface gigabitethernet 0/0/18
lldp-profile <profile-name>
```

LLDP-MED Usage

In a converged network, LLDP-MED provides the following benefits:

- **Interoperability**
LLDP-MED offers vendor-independent management capabilities, enabling different convergence endpoints to inter-operate on one network.
- **Automatic deployment of network policies**
With LLDP-MED, administrators can automatically deploy voice VLAN.



The default transmit interval time is 30 seconds and the default transmit hold timer is 120 seconds. You can change the transmit-interval and transmit-hold timer in the lldp-profile.

- **Location services**
LLDP-MED allows deploying location services.
- **Detailed inventory management capabilities**
For each converged device, LLDP-MED can supply model, manufacturer, firmware and asset information.
- **Advanced PoE**
LLDP-MED enables advanced Power over Ethernet capabilities.
- **IP telephony network troubleshooting**
LLDP-MED enables detection of speed and duplex mismatches, and of improper static voice policy configurations.
- **More security**
LLDP-MED runs after 802.1X to prevent unauthenticated devices from gaining access to the network.
- **Hardware Information**
For each converged device, LLDP-MED can supply model, manufacturer and firmware.
- **IP Telephony Network Troubleshooting**
The information from the device attached and information from our own device is available for the user to take corrective action.

Verifying the LLDP Profile Configuration to Check LLDP-MED Status

To verify the LLDP profile configuration check LLDP-Med. status, use the following command:

```
(host) (config) #show interface-profile lldp-profile <profile-name>
```

```
LLDP Profile " profile-name"
-----
Parameter                                Value
-----
LLDP pdu transmit                        Enabled
LLDP protocol receive processing         Enabled
LLDP transmit interval (Secs)           30
LLDP transmit hold multiplier           4
LLDP fast transmit interval (Secs)      30
LLDP fast transmit counter               1
LLDP-MED protocol                       Enabled
Control proprietary neighbor discovery   Disabled
```

Propriety Link Layer Discovery Protocols

This section contains the following sections:

- [“Understanding Proprietary Link Layer Discovery Protocol” on page 127](#)
- [“Configuring Propriety LLDP Receive Processing” on page 128](#)
- [“Verifying Propriety LLDP Receive Processing” on page 128](#)
- [“Monitoring the Proprietary Neighbor Discovery” on page 129](#)

Understanding Proprietary Link Layer Discovery Protocol

Network companies can also define their proprietary data link layer discovery protocol. For instance, Cisco Discovery Protocol (CDP) is a proprietary data link layer discovery protocol. CDP is similar to LLDP and is used to share information about other directly connected vendor-specific equipment. CDP runs on many of vendor-specific devices including routers, switches, and VoIP phones.

When there are devices in the network that do not support LLDP, you can use the `proprietary-neighbor-discovery` knob in the LLDP interface profile to turn on the ability to receive proprietary discovery protocol packets and identify the neighbors. This release supports only CDP (Cisco Discovery Protocol). You can use the `show neighbor-devices` command to display the neighbors identified using LLDP and CDP protocols.

CDP Receive Processing

The Mobility Access Switch processes CDP frames that are received from CDP-supported devices. However, the Mobility Access Switch only receives CDP frames and does not forward CDP frames to other connected neighbors/devices. When new CDP information is received from an existing neighbor, the Mobility Access Switch updates the information and discards the existing information.

CDP Frame Information

The CDP frame contains the following information:

- Device ID
- IP Address
- Port ID
- Capabilities
- Software Version

- Platform
- Native VLAN

Configuring Propriety LLDP Receive Processing

Priority LLDP receive processing is configured under LLDP profile:

```
(host) (config) #interface-profile lldp-profile CDP-PROC
(host) (LLDP Profile "CDP-PROC") #proprietary-neighbor-discovery
(host) (LLDP Profile "CDP-PROC") #exit
```

The configured LLDP/CDP-PROC profile needs to be applied to the interface:

```
(host) (config) #interface gigabitethernet 2/0/23
(host) (gigabitethernet "2/0/23") #lldp-profile CDP-PROC
(host) (gigabitethernet "2/0/23") #exit
```

Verifying Propriety LLDP Receive Processing

Propriety LLDP receive processing configuration profile can be verified with the following command:

```
(host) #show interface-profile lldp-profile CDP-PROC
LLDP Profile "CDP-PROC"
```

```
-----
Parameter                                Value
-----
LLDP pdu transmit                        Disabled
LLDP protocol receive processing         Disabled
LLDP transmit interval (Secs)           30
LLDP transmit hold multiplier           4
LLDP fast transmit interval (Secs)      30
LLDP fast transmit counter               1
LLDP-MED protocol                       Disabled
Control proprietary neighbor discovery  Enabled
```

CDP-enabled neighboring devices can be viewed by following CLI command:

```
(host) #show neighbor-devices
Neighbor Devices Information
```

```
-----
Local Intf  Chassis ID      Protocol  Remote Intf      Expiry-Time (Secs)
-----
GE2/0/22    SEP002414B211B3  CDPv2    GigabitEthernet0/22  44
GE2/0/23    SEP00254593BFD8  CDPv2    Port 1             166
```

System name

```
-----
```

```
SEP002414B211B3.cisco.com
```

```
SEP00254593BFD8.cisco.com
```

Number of neighbors: 2

```
(host) #show neighbor-devices interface gigabitethernet 2/0/23
```

```
Neighbor Devices Information
```

```
-----
Local Intf  Chassis ID      Protocol  Remote Intf      Expiry-Time (Secs)
-----
GE2/0/23    SEP00254593BFD8  CDPv2    Port 1             137
```

System name

```
-----
```

```
SEP00254593BFD8.cisco.com
```



```
Number of neighbors: 1
(host) #show neighbor-devices interface gigabitethernet 2/0/23 detail
```

```
Interface: GE2/0/23, Number of neighbors: 1
```

```
-----
Chassis id: SEP00254593BFD8, Protocol: CDPv2
```

```
Management address: 5.5.5.21
```

```
Interface description: Port 1, ID: Port 1
```

```
Last Update: Sat Oct 1 14:24:43 2011
```

```
Time to live: 180, Expires in: 170 Secs
```

```
System capabilities :
```

```
Enabled capabilities:
```

```
System name: SEP00254593BFD8
```

```
System description:
```

```
  SCCP41.8-4-4S
```

```
Duplex: full
```

Monitoring the Proprietary Neighbor Discovery

You can use the following commands to display the neighbors discovered using the proprietary protocols such as CDP:

```
(host)# show neighbor-devices
```

```
(host)# show neighbor-devices interface gigabitethernet 0/0/1
```

```
(host)# show neighbor-devices interface gigabitethernet 0/0/1 detail
```


The Mobility Access Switch supports certain Voice functionalities.


This chapter includes the following topics:

- “Voice VLANs” on page 131
- “Creating and Applying VoIP Profile to an Interface” on page 132
- “VoIP Auto-Discovery on Trusted Ports” on page 132

Voice VLANs

The VoIP VLAN feature enables access ports to accept both untagged (data) and tagged (voice) traffic from IP phones connected directly to the Mobility Access Switch and separate these traffic into different VLANs (namely data VLAN and voice VLAN). You can configure a voice VLAN using the `voip-profile`.

The dot1p and DSCP values in the VoIP profile are communicated to the phone using LLDP. VoIP profile does not affect the QoS behavior on the switch. The QoS behavior depends on the QoS configuration on the port.



The following guidelines and limitations must be considered before creating a VoIP profile:

- If the port is configured as QoS trusted then the phone is expected to mark the DSCP and dot1p fields accordingly.
- To enable separate QoS treatment for the voice traffic ingressing an interface, you can either enable QoS Trust on the interface or apply the QoS-profile to the interface/access-list/user-role. For more information, see [Chapter 27, “Quality of Service” on page 225](#).
- Voice VLAN can be applied only to the access ports.
- Trunk ports and port-channels are not allowed to be part of a voice VLAN.
- You cannot assign a VoIP profile to untrusted interfaces. In the case of untrusted interfaces, the phone derives the `voip-vlan` from the role that is assigned to the phone after authentication.
- LLDP-MED instructs the attached VoIP phones to use the specified voice VLAN ID, 802.1p, and DSCP values. For details about configuring an LLDP profile, refer to [Chapter 11, “Link Layer Discovery Protocols” on page 121](#).

Creating and Applying VoIP Profile to an Interface

You can create and apply a VoIP profile to an interface using the following set of commands:

```
(host)(config)# interface-profile voip-profile <profile-name>
    clone <source>
    no{...}
    voip-dot1p <priority>
    voip-dscp <value>
    voip-vlan <VLAN-ID>
(host)(config)# interface gigabitethernet <slot/module/port>
    voip-profile <profile-name>
```

VoIP Auto-Discovery on Trusted Ports

ArubaOS provides support for VoIP Auto-discovery (also referred as CDP Fingerprinting) to discover the VoIP phones using neighbor discovery protocols (such as LLDP-MED and CDP) and assign Voice VLAN to the traffic originating from the phone. For more information on LLDP-MED, see [Chapter 11, “Link Layer Discovery Protocols”](#) on page 121.

You can configure VoIP either in static mode or auto-discover mode. By default, VoIP is configured in static mode. When VoIP operates in static mode, the phone is expected to know the Voice VLAN to be used and send the Voice traffic with the Voice VLAN tag. This is achieved, only if the Voice VLAN is configured statically on the phone or propagated to the phone using LLDP-MED.

In auto-discover mode, when LLDP-MED or CDP discovers a phone, the switch creates a rule to associate all the traffic originating from the phone to the Voice VLAN. Hence, the Voice VLAN need not be configured statically on the phone. The Voice VLAN can be tagged or untagged depending on the LLDP-MED configuration.

VoIP configured in auto-discover mode applies the Voice VLAN only to the first neighbor discovered in an interface. If both LLDP-MED and CDP neighbors are discovered, the preference is always given to the first LLDP-MED neighbor even if a CDP neighbor is already associated.

Enabling VoIP Auto-Discovery

You can use the following CLI command to enable VoIP in auto-discover mode:

```
(host) (config) #interface-profile voip-profile VOIP-1
(host) (VOIP profile "VOIP-1") #voip-mode auto-discover
(host) (VOIP profile "VOIP-1") #voip-vlan 5
```



You must enable the LLDP-profile with proprietary-neighbor-discovery/LLDP on the respective interface to identify the CDP/LLDP enabled phones.

You can enable proprietary-neighbor-discovery on an LLDP profile:

```
(host) (config) #interface-profile lldp-profile LLDP-1
(host) (LLDP Profile "LLDP-1") #lldp transmit
(host) (LLDP Profile "LLDP-1") #lldp receive
(host) (LLDP Profile "LLDP-1") #med enable
(host) (LLDP Profile "LLDP-1") #proprietary-neighbor-discovery
```

You can apply the configured LLDP-1 profile to an interface:

```
(host) (config) #interface gigabitethernet 0/0/0
(host) (gigabitethernet "0/0/0") #lldp-profile LLDP-1
(host) (gigabitethernet "0/0/0") # voip-profile VOIP-1
```

Verifying VoIP Mode Configuration

You can use the following command to verify the VoIP mode configuration on a VoIP profile:

```
(host) (config) #show interface-profile voip-profile VOIP-1
```

```
VOIP profile "VOIP-1"
-----
Parameter  Value
-----  -
VOIP VLAN  5
DSCP       46
802.1 UP   6
VOIP Mode  auto-discover
```

Viewing Neighboring Phones

You can use the following command to view the neighboring phones in the network and the Voice VLAN associated with the phones:

```
(host) #show neighbor-devices phones
Neighbor Phones
-----
Interface  Protocol  Phone MAC          Voice VLAN
-----  -
GE0/0/6    CDPv2     00:1b:54:c9:e9:fd  -
GE0/0/47   CDPv2     00:1b:54:c9:e9:fd  5
```

In the above output, "-" under the Voice VLAN column denotes that either Voice VLAN is not available or VoIP is not configured to run in auto-discover mode.

VoIP Auto-discovery on Untrusted Ports

This release of Mobility Access Switch automatically discovers the Cisco Discovery Protocol (CDP) phones on an untrusted interface and assigns a VoIP VLAN to the phone.

Complete the following steps to place a non-802.1x CDP phone in a VoIP VLAN by using a user derivation rule (UDR) to match **device-type**:

1. Create an LLDP profile.

```
(host) (config) #interface-profile lldp-profile ciscophones
(host) (LLDP Profile "ciscophones") #proprietary-neighbor-discovery
```

2. Create a VoIP profile.

```
(host) (config) #interface-profile voip-profile phone
(host) (VOIP profile "phone") #voip-vlan 100
```

3. Create a user-role and add the previously created VoIP profile to that role.

```
(host) (config-role) #user-role phonerole
(host) (config-role) #access-list stateless allowall-stateless
(host) (config-role) #voip-profile phone
```

4. Create a UDR and add the phone role.

```
(host) (config) #aaa derivation-rules user phoneudr
(host) (user-rule) #set role condition device-type equals "phone" set-value phonerole
```

5. Add the UDR to a AAA profile.

```
(host) (config) #aaa profile phone_client  
(host) (AAA Profile "phone_client") #user-derivation-rules phoneudr
```

6. Attach the LLDP profile and AAA profile to a port.

```
(host) (config) #interface gigabitethernet 0/0/2  
(host) (gigabitethernet "0/0/2") #lldp-profile ciscophones  
(host) (gigabitethernet "0/0/2") #aaa-profile phone_client
```

Alternatively, you can define the UDR for a VLAN assignment using the following command:

```
(host) (config) #aaa derivation-rules user <rule-name>  
(host) (user-rule) #set vlan condition device-type equals phone set-value <vlan-id>  
[position <priority> | description <descr>]
```



It is recommended to configure the UDR for the CDP phones that do not support LLDP or 802.1x authentication on an untrusted interface.

The implementation of Multiple Spanning Tree Protocol (MSTP) is based on the IEEE Standard 802.1D-2004 and 802.1Q-2005. In addition, MSTP supports the loopguard, rootguard, bpduguard, and portfast features.



To enable MSTP, use the spanning tree mode command.

MSTP maps a group of Virtual Local Area Networks (VLANs) to a reduced number of spanning tree instances. This allows VLAN bridges to use multiple spanning trees. This protocol enables network traffic from different VLANs to flow through different potential paths within a bridged VLAN. Because most networks do not need more than a few logical topologies, MSTP provides design flexibility as well as better overall network resource utilization.

Layer 2 networks typically use multiple paths and link redundancies to handle node and link failures. By definition, spanning tree uses a subset of the available physical links in its active logical topology to provide complete connectivity between any pair of end hosts. This chapter covers:

- [“Important Points to Remember” on page 135](#)
- [“Example MSTP Configuration” on page 136](#)
- [“Loopguard and Rootguard” on page 138](#)
- [“Bridge Protocol Data Unit \(BPDU\) Guard” on page 140](#)
- [“Sample Topology and Configuration” on page 142](#)

Important Points to Remember

- Configure MSTP using the command line only.
- Portfast, Loopguard, BPDUguard, and Rootguard are disabled by default.
- MSTP allows users to map a set of VLANs to a MSTP instance.
- MSTP allows formation of multiple spanning tree regions and each region can run multiple instances.
- For two switches to be in the same MSTP region, they must share the same name, the same version, and the same VLAN instance mapping.
- If a Mobility Access Switch receives RSTP/STP control packets from a neighbor, the neighbor is considered to be in a different region. For the RSTP/STP neighbor, the entire MSTP region looks like a single bridge.
- You can perform proper load balancing across redundant links using MSTP instances. The ability to configure the port cost and port priority values also provides you with the flexibility to determine the links that are chosen to carry the traffic.
- State machines (SM), as defined by the IEEE, get the port and instance information as input. As output, SMs provide the port-state for each port in every instance.

Example MSTP Configuration

Basic MSTP configuration includes setting the spanning tree mode to MSTP, entering the global MSTP mode, and assigning a region name.

1. Set the spanning tree mode:

```
(host)(config) #spanning-tree mode mstp
```

2. Verify the spanning tree mode:

```
(host)(config) #show spanning-tree-profile
```

```
spanning-tree
-----
Parameter          Value
-----
spanning-tree-mode mstp
```

3. Assign a region name:

```
(host) (Global MSTP) #region-name mstptechpubs
```

There are, of course, other MSTP options you can configure (such as forward delay, hello time). You can view the current MSTP configuration values using the **show mstp-global-profile** command.

```
(host) # show mstp-global-profile
```

```
Global MSTP
-----
Parameter          Value
-----
MSTP region name    mstptechpubs
MSTP revision        0
Instance bridge priority 1 4096
Instance vlan mapping 1 801-802
MSTP hello time      2
MSTP forward delay   15
MSTP maximum age     20
MSTP max hops        20
```

To view the interface MSTP configuration values, use the **show interface-profile mstp-profile** command:

```
(host) (config) #show interface-profile mstp-profile
```

```
Interface MSTP List
-----
Name           References  Profile Status
----
default        14
mstp_cost       3
techpubs       2
Total:4
```

To view the interface-profile named 'mstp_cost', use the **show interface-profile mstp_cost** command:

```
(config) #show interface-profile mstp-profile mstp_cost
```

```
Interface MSTP "mstp_cost"
-----
Parameter          Value
-----
Instance port cost  0 100
Instance port cost  1 200
Instance port cost  2 300
```


Instance port priority	N/A
Enable point-to-point	Disabled
Enable portfast	Disabled
Enable rootguard	Disabled
Enable loopguard	Disabled

Viewing Operational Information

To view MSTP operational information, use the **show spanning-tree interface all detail** command (the following is a partial output)

```
(host) #show spanning-tree mstp interface all detail

(GE0/0/23) of MST 0 is designated forwarding
Port path cost 20000, Port priority 128, Port identifier 128.24
Designated Root ID priority: 32768, Address: 000b.866a.f240
Designated Bridge ID priority: 32768, Address: 000b.866a.f240
Number of transitions to forwarding state: 1
Link type is point-to-point by default, Internal
BPDU sent: 108, Received: 9
Edge mode: Disabled
Root guard: Disabled
Loop guard: Disabled
(GE0/0/23) of MST 4 is designated forwarding
Port path cost 20000, Port priority 128, Port identifier 128.24
Designated Root ID priority: 32768, Address: 000b.866a.f240
Designated Bridge ID priority: 32768, Address: 000b.866a.f240
Number of transitions to forwarding state: 1
Link type is point-to-point by default, Internal
BPDU sent: 104, Received: 5

(GE1/0/22) of MST 0 is designated forwarding
Port path cost 20000, Port priority 128, Port identifier 128.167
Designated Root ID priority: 32768, Address: 000b.866a.f240
Designated Bridge ID priority: 32768, Address: 000b.866a.f240
Number of transitions to forwarding state: 1
Link type is point-to-point by default, Internal
BPDU sent: 107, Received: 8
Edge mode: Disabled
Root guard: Disabled
Loop guard: Disabled
(GE1/0/22) of MST 4 is designated forwarding
Port path cost 20000, Port priority 128, Port identifier 128.167
Designated Root ID priority: 32768, Address: 000b.866a.f240
Designated Bridge ID priority: 32768, Address: 000b.866a.f240
Number of transitions to forwarding state: 1
Link type is point-to-point by default, Internal
BPDU sent: 104, Received: 4
...
```

Or use the **show spanning-tree msti all detail** command (partial).

```
(host) #show spanning-tree mstp msti all detail

MST 0

vlans mapped          : 3,7
Configuration Digest : 0xED285086D33012C7D2B283FB89730D4D
```

```

Root ID                Address: 000b.866a.f240, Priority: 32768
Regional Root ID      Address: 000b.866a.f240, Priority: 32768
Bridge ID             Address: 000b.866a.f240, Priority: 32768
External root path cost 0, Internal root path cost 0

```

Interface	Role	State	Port Id	Cost	Type
GE0/0/23	Desg	FWD	128.24	20000	P2p
GE1/0/22	Desg	FWD	128.167	20000	P2p
GE1/0/23	Bkup	BLK	128.168	20000	P2p
GE2/0/23	Bkup	BLK	128.312	20000	P2p

MST 4

```

vlangs mapped          : 1
Root ID                Address: 000b.866a.f240, Priority: 32768
Bridge ID             Address: 000b.866a.f240, Priority: 32768
root path cost 0, remaining hops 20

```

Interface	Role	State	Port Id	Cost	Type
GE0/0/23	Desg	FWD	128.24	20000	P2p
GE1/0/22	Desg	FWD	128.167	20000	P2p
GE1/0/23	Bkup	BLK	128.168	20000	P2p
GE2/0/23	Bkup	BLK	128.312	20000	P2p

For a more complete listing of MSTP commands, see the [Command Line Reference Guide](#).

Loopguard and Rootguard

Loopguard provides additional protection against Layer 2 forwarding loops (spanning tree loops). A spanning tree loop is created when a spanning tree blocking port, in a redundant topology, erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant topology (not necessarily the spanning tree blocking port) is no longer receiving spanning tree BPDUs (Bridge Protocol Data Units).

If loopguard is enabled on a non-designated port and it stops receiving BPDUs, then that non-designated port is moved into the spanning tree loop-inconsistent blocking state.



Best practices is that loopguard be used on non-designated ports.

Configuring Loopguard

Below is a basic configuration for loopguard using the profile name *techpubs*.

```

(host) (config) #interface-profile mstp-profile techpubs
(host) (Interface MSTP "techpubs") #loopguard
(host) (Interface MSTP "techpubs") #

```

Associate the above mstp-profile to the interface:

```

(host) (config) #interface gigabitethernet 0/0/2
(host) (gigabitethernet "0/0/2") #mstp-profile techpubs
(host) (gigabitethernet "0/0/2") #

```

Verify the loopguard configuration:

```
(host) #show spanning-tree
```

```
MST 0
Root ID                Address: 0019.0655.3a80, Priority: 4097
Regional Root ID      Address: 000b.866c.3200, Priority: 16384
Bridge ID              Address: 000b.866c.3200, Priority: 16384
External root path cost 40000, Internal root path cost 0

Interface  Role      State  Port Id  Cost  Type
-----
GE0/0/1    Desg      FWD    128.2    20000 P2p
GE0/0/2    Loop-Inc  BLK    128.3    20000 P2p Bound <-- loopguard on GE0/0/2
GE0/0/22   Root      FWD    128.23   20000 P2p
```

Verify that loopguard is applied to the interface:

```
(host) #show spanning-tree mstp interface gigabitethernet 0/0/2 detail
```

```
(GE0/0/2) of MST 0 is loop inconsistent blocking
Port path cost 20000, Port priority 128, Port identifier 128.3
Designated Root ID priority: 4097, Address: 0019.0655.3a80
Designated Bridge ID priority: 16384, Address: 000b.866c.3200
Number of transitions to forwarding state: 1
Link type is point-to-point by default, Boundary
BPDU sent: 15, Received: 36
Edge mode: Disabled
Root guard: Disabled
Loop guard: Enabled <-- loopguard enabled
```

Configuring Rootguard

Rootguard provides a way to enforce the root bridge placement in the network. The rootguard feature guarantees that a port will not be selected as Root Port for the CIST or any MSTI. If a bridge receives superior spanning tree BPDUs on a rootguard-enabled port, the port is selected as an Alternate Port instead of Root Port and no traffic is forwarded across this port.

By selecting the port as an Alternate Port, the rootguard configuration prevents bridges, external to the region, from becoming the root bridge and influencing the active spanning tree topology.



Best practices is that rootguard be used on designated ports.

Below is a basic configuration for rootguard using the profile name *techpubs*.

```
(host) (config) #interface-profile mstp-profile techpubs
(host) (Interface MSTP "techpubs") #rootguard
(host) (Interface MSTP "techpubs") #
```

Associate the above mstp-profile to the interface:

```
(host) (config) #interface gigabitethernet 0/0/1
(host) (gigabitethernet "0/0/1") #mstp-profile techpubs
(host) (gigabitethernet "0/0/1") #
```

If a downstream bridge starts advertising itself as root without rootguard enabled, MSTP will accept that bridge as root. With rootguard enabled, it guards the root and prevents bridges from neighboring networks from becoming the root.

Verify the rootguard configuration:

```
(host) #show spanning-tree
```

```
MST 0
Root ID          Address: 0019.0655.3a80,  Priority: 4097
Regional Root ID Address: 000b.866c.3200,  Priority: 16384
Bridge ID        Address: 000b.866c.3200,  Priority: 16384
External root path cost 40000, Internal root path cost 0

Interface  Role          State  Port Id  Cost  Type
-----
GE0/0/1    Altn(Root-Inc)  BLK    128.22   20000  P2p    <---rootguard on GE0/0/1
GE0/0/2    Desg             FWD    128.301  20000  P2p
GE0/0/22   Root             FWD    128.23   20000  P2p
```

Use the show interface-profile *profile name* command to view the status of loopguard and rootguard.

```
(host) #show interface-profile mstp-profile techpubs
```

```
Interface MSTP "techpubs"
```

```
-----
Parameter          Value
-----
Instance port cost  N/A
Instance port priority N/A
Enable point-to-point Disabled
Enable portfast     Disabled
Enable rootguard     Enabled
Enable loopguard     Enabled
```

Bridge Protocol Data Unit (BPDU) Guard

BPDU guard functionality prevents malicious attacks on edge ports. When the malicious attacker sends a BPDU on the edge port, it triggers unnecessary STP calculation. To avoid this attack, use the BPDU guard on that edge port. The BPDU guard enabled port shuts down as soon as a BPDU is received.

Enabling and Configuring BPDU Guard Functionality

BPDU guard can be enabled or disabled at an interface level. By default, the BPDU is disabled. The BPDU guard functionality is configured as part of the mstp-profile configuration.

You can use the following command to configure the BPDU guard by using the MSTP profile:

```
(host) (config) #interface-profile mstp-profile <profile-name>
bpduguard
auto-recovery-time <recovery-time>
```

The following example shows how to enable and configure BPDU guard :

```
(host)(config)# interface-profile mst-profile BPDU-Guard1
bpduguard auto-recovery-time 60
```



You can configure BPDU guard with or without the auto-recovery-time option.

You can disable BPDU guard by using the following command:

```
(host) (config) #interface-profile <profile-name> no bpduguard
```

You can disable the auto recovery time by using the following command:

```
(host) (Interface MST "profile-name") #bpduguard no auto-recovery-time
```

Verifying the BPDU Guard Configuration:

```
(host) (config) #show interface-profile mstp-profile bpdu-guard
```

```
Interface MSTP "bpdu-guard"
-----
Parameter                                Value
-----
Instance port cost                       N/A
Instance port priority                   N/A
Enable point-to-point                   Disabled
Enable portfast                         Disabled
Enable rootguard                       Disabled
Enable loopguard                       Disabled
Enable bpduguard                       Enabled <——BPDU guard is enabled
Enable bpduguard auto recovery time    N/A
```

Sample Configuration

To enable and configure BPDU guard using the MSTP profile:

```
(host)(config)# interface-profile mst-profile BPDU-Guard1
bpduguard auto-recovery-time 60
```

To attach the MSTP profile to the interface:

```
(host) (config)# interface gigabitethernet <0/0/6>
mstp-profile BPDU-Guard1
```

Portfast

When the link on a bridge port goes up, MSTP runs its algorithm on that port. If the port is connected to a host that does not “speak” MSTP, it takes approximately 30 seconds for the port to transition to the forwarding state. During this time, no user data passes through this bridge port and some user applications may timeout.

Configuring Portfast

To immediately transition the bridge port into the forwarding state upon linkup, enable the MSTP Portfast feature.

```
(host) (config) #interface-profile mstp-profile portfast_techpubs
(host) (Interface MSTP "portfast_techpubs") #portfast
```

The bridge port still participates in MSTP; if a BPDU is received, it becomes a normal port.



The portfast is operational only on access ports.

Associate the above mstp-profile to the interface:

```
(host) (config) #interface gigabitethernet 0/0/1
(host) (gigabitethernet "0/0/1") #mstp-profile portfast_techpubs
(host) (gigabitethernet "0/0/1")
```

Use the **show interface-profile** command to view the status of Portfast.


```
(host) (config) #show interface-profile mstp-profile portfast_techpubs
```

Interface MSTP "portfast_techpubs"

Parameter	Value
Instance port cost	N/A
Instance port priority	N/A
Enable point-to-point	Disabled
Enable portfast	Enabled
Enable rootguard	Disabled
Enable loopguard	Disabled

Sample Topology and Configuration

Figure 1 MSTP Topology



Below is a the configuration for the topology in [Figure 1](#).

S3500 62 Configuration

```
!  
interface-profile switching-profile "access-port-509"
```

```

    access-vlan 509
!
interface-profile switching-profile "access-port-865"
    access-vlan 865
!
interface-profile switching-profile "access-vlan-2"
    access-vlan 2
!
interface-profile switching-profile "accessPortVlan100"
    access-vlan 100
!
interface-profile switching-profile "accessPortVlan120"
    access-vlan 120
!
interface-profile switching-profile "accessPortVlan150"
    access-vlan 150
!
interface-profile switching-profile "accessPortVlan200"
    access-vlan 200
!
interface-profile switching-profile "accessPortVlan40"
    access-vlan 40
!
interface-profile switching-profile "accessVlan12"
    access-vlan 12
!
interface-profile switching-profile "accessVlan6"
    access-vlan 6
!
interface-profile switching-profile "accessVlan9"
    access-vlan 9
!
interface-profile switching-profile "default"
!
interface-profile switching-profile "trunk-profile"
    switchport-mode trunk
!
interface-profile poe-profile "default"
!
interface-profile enet-link-profile "default"
!
interface-profile lACP-profile "pc0"
    group-id 0
    mode active
!
interface-profile lACP-profile "pc1"
    group-id 1
    mode active
!
interface-profile lldp-profile "default"
!
interface-profile lldp-profile "lldp-factory-initial"
    lldp transmit
    lldp receive
    med enable
!
interface-profile mstp-profile "default"
!

```

```

interface-profile mstp-profile "mstpPortfast"
    portfast
!
interface-profile mstp-profile "pathCost2000"
    instance 0 cost 2000
!
interface-profile mirroring-profile "toPort28"
!
spanning-tree
    mode mstp
!
mstp
    region-name "region1"
    instance 2 bridge-priority 4096
    instance 1 vlan 50-100
    instance 2 vlan 101-151
    instance 3 vlan 152-202
    instance 4 vlan 203-253
    instance 5 vlan 254-304
    instance 6 vlan 305-355
    instance 7 vlan 356-406
    instance 8 vlan 407-457
    instance 9 vlan 458-508
    instance 10 vlan 509-559
    instance 11 vlan 560-610
    instance 12 vlan 611-661
    instance 13 vlan 662-712
    instance 14 vlan 713-763
    instance 15 vlan 764-814
    instance 16 vlan 815-865
!
lacp
!
igmp-snooping-profile "default"
!
igmp-snooping-profile "igmp-snooping-factory-initial"
!
poemanagement member-id "default"
!
vlan "10"
!
vlan "100"
!
vlan "1000"
!
vlan "101"
!
vlan "102"
!
vlan "103"
!
vlan "104"
!
vlan "105"
!
vlan "106"
!
vlan "107"

```



```

!
vlan "108"
!
vlan "109"
!
vlan "11"
!
!
vlan "995"
!
vlan "996"
!
vlan "997"
!
vlan "998"
!
vlan "999"
!
interface gigabitethernet "0/0/0"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/12"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/2"
    lacp-profile "pc1"
!
interface gigabitethernet "0/0/20"
    mstp-profile "mstpPortfast"
!
interface gigabitethernet "0/0/24"
    shutdown
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/28"
    mstp-profile "mstpPortfast"
!
interface gigabitethernet "0/0/3"
    lacp-profile "pc1"
!
interface gigabitethernet "0/0/30"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/36"
    shutdown
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/42"
    lacp-profile "pc0"
!
interface gigabitethernet "0/0/43"
    lacp-profile "pc0"
!
interface gigabitethernet "0/0/46"
    shutdown
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/47"

```

```

        shutdown
        switching-profile "trunk-profile"
    !
interface vlan "4093"
    !
interface mgmt
    ip address 10.16.56.62 netmask 255.255.255.0
    !
interface port-channel "0"
    switching-profile "trunk-profile"
    !
interface port-channel "1"
    switching-profile "trunk-profile"
    !

snmp-server enable trap
end

(host62) #

```

S3500 63 Configuration

```

!
interface-profile switching-profile "access-poer-10"
    access-vlan 10
    !
interface-profile switching-profile "access-port-1000"
    access-vlan 1000
    !
interface-profile switching-profile "access-port-287"
    access-vlan 287
    !
interface-profile switching-profile "access-port-509"
    access-vlan 509
    !
interface-profile switching-profile "accessPortVlan100"
    access-vlan 100
    !
interface-profile switching-profile "accessPortVlan120"
    access-vlan 120
    !
interface-profile switching-profile "accessPortVlan150"
    access-vlan 150
    !
interface-profile switching-profile "accessPortVlan200"
    access-vlan 200
    !
interface-profile switching-profile "accessPortVlan40"
    access-vlan 40
    !
interface-profile switching-profile "accessVlan12"
    access-vlan 12
    !
interface-profile switching-profile "accessVlan6"
    access-vlan 6
    !
interface-profile switching-profile "accessVlan9"
    access-vlan 9
    !

```

```

interface-profile switching-profile "default"
!
interface-profile switching-profile "trunk-profile"
    switchport-mode trunk
!
interface-profile switching-profile "vlan-13-mgmt"
    access-vlan 13
!
interface-profile tunneled-node-profile "tunnuel-ip-10.10.1"
    controller-ip 10.10.10.2
    keepalive 5
!
interface-profile poe-profile "default"
!
interface-profile enet-link-profile "default"
!
interface-profile lacp-profile "pc1"
    group-id 1
    mode active
!
interface-profile lacp-profile "pc2"
    group-id 2
!
interface-profile lldp-profile "default"
!
interface-profile lldp-profile "lldp-factory-initial"
    lldp transmit
    lldp receive
    med enable
!
interface-profile mstp-profile "default"
!
interface-profile mstp-profile "mstpPortfast"
    portfast
!
interface-profile mirroring-profile "toPort31"
!
spanning-tree
    mode mstp
!
mstp
    region-name "region1"
    instance 3 bridge-priority 4096
    instance 0 bridge-priority 20480
    instance 1 vlan 50-100
    instance 2 vlan 101-151
    instance 3 vlan 152-202
    instance 4 vlan 203-253
    instance 5 vlan 254-304
    instance 6 vlan 305-355
    instance 7 vlan 356-406
    instance 8 vlan 407-457
    instance 9 vlan 458-508
    instance 10 vlan 509-559
    instance 11 vlan 560-610
    instance 12 vlan 611-661
    instance 13 vlan 662-712
    instance 14 vlan 713-763

```

```

        instance 15 vlan 764-814
        instance 16 vlan 815-865
    !
    lacp
    !
    igmp-snooping-profile "default"
    !
    igmp-snooping-profile "igmp-snooping-factory-initial"
    !
    poemanagement member-id "default"
    !
    vlan "10"
    !
    vlan "100"
    !
    vlan "1000"
    !
    vlan "101"
    !
    vlan "102"
    !
    vlan "103"
    !
    vlan "104"
    !
    vlan "105"
    !
    vlan "106"
    !
    vlan "107"
    !
    vlan "998"
    !
    vlan "999"
    !
    interface gigabitethernet "0/0/0"
        shutdown
    !
    interface gigabitethernet "0/0/12"
        lacp-profile "pc1"
    !
    interface gigabitethernet "0/0/13"
        lacp-profile "pc1"
    !
    interface gigabitethernet "0/0/16"
        switching-profile "trunk-profile"
    !
    interface gigabitethernet "0/0/17"
        shutdown
        switching-profile "trunk-profile"
    !
    interface gigabitethernet "0/0/31"
        mstp-profile "mstpPortfast"
        tunneled-node-profile "tunnuel-ip-10.10.1"
    !
    interface gigabitethernet "0/0/34"
        switching-profile "trunk-profile"
    !

```

```

interface gigabitethernet "0/0/36"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/42"
    lacp-profile "pc2"
!
interface gigabitethernet "0/0/43"
    lacp-profile "pc2"
!
interface gigabitethernet "0/0/44"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/45"
    switching-profile "vlan-13-mgmt"
!
interface gigabitethernet "0/0/46"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/47"
    mstp-profile "mstpPortfast"
!
interface gigabitethernet "0/0/6"
!
interface gigabitethernet "0/0/7"
!
interface mgmt
    ip address 10.16.56.63 netmask 255.255.255.0
!
interface port-channel "1"
    switching-profile "trunk-profile"
!
interface port-channel "2"
    switching-profile "trunk-profile"

```

(host 63) #

S3500 64 Configuration

```

!
interface-profile switching-profile "access-port-509"
    access-vlan 509
!
interface-profile switching-profile "access-port-865"
    access-vlan 865
!
interface-profile switching-profile "access-vlan-2"
    access-vlan 2
!
interface-profile switching-profile "accessPortVlan100"
    access-vlan 100
!
interface-profile switching-profile "accessPortVlan120"
    access-vlan 120
!
interface-profile switching-profile "accessPortVlan150"
    access-vlan 150
!
interface-profile switching-profile "accessPortVlan200"
    access-vlan 200

```

```

!
interface-profile switching-profile "accessPortVlan40"
    access-vlan 40
!
interface-profile switching-profile "accessVlan12"
    access-vlan 12
!
interface-profile switching-profile "accessVlan6"
    access-vlan 6
!
interface-profile switching-profile "accessVlan9"
    access-vlan 9
!
interface-profile switching-profile "default"
!
interface-profile switching-profile "trunk-profile"
    switchport-mode trunk
!
interface-profile poe-profile "default"
!
interface-profile enet-link-profile "default"
!
interface-profile lacp-profile "pc0"
    group-id 0
    mode active
!
interface-profile lacp-profile "pc2"
    group-id 1
    mode active
!
interface-profile lacp-profile "pc2"
    group-id 2
!
interface-profile lldp-profile "default"
!
interface-profile lldp-profile "lldp-factory-initial"
    lldp transmit
    lldp receive
    med enable
!
interface-profile mstp-profile "default"
!
interface-profile mstp-profile "mstpPortfast"
    portfast
!
interface-profile mstp-profile "pathCost2000"
    instance 0 cost 2000
!
interface-profile mirroring-profile "toPort28"
!
spanning-tree
    mode mstp
!
mstp
    region-name "region1"
    instance 2 bridge-priority 4096
    instance 0 bridge-priority 16384
    instance 1 vlan 50-100

```

```

instance 2 vlan 101-151
instance 3 vlan 152-202
instance 4 vlan 203-253
instance 5 vlan 254-304
instance 6 vlan 305-355
instance 7 vlan 356-406
instance 8 vlan 407-457
instance 9 vlan 458-508
instance 10 vlan 509-559
instance 11 vlan 560-610
instance 12 vlan 611-661
instance 13 vlan 662-712
instance 14 vlan 713-763
instance 15 vlan 764-814
instance 16 vlan 815-865
!
lacp
!
igmp-snooping-profile "default"
!
igmp-snooping-profile "igmp-snooping-factory-initial"
!
poemanagement member-id "default"
!
vlan "10"
!
vlan "100"
!
vlan "1000"
!
vlan "101"
!
vlan "102"
!
vlan "103"
!
vlan "104"
!
vlan "105"
!
vlan "106"
!
vlan "107"
!
vlan "108"
!
vlan "109"
!
vlan "11"
!
vlan "110"
!

interface gigabitethernet "0/0/0"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/12"
    switching-profile "trunk-profile"

```

```

!
interface gigabitethernet "0/0/2"
    lACP-profile "pc0"
!
interface gigabitethernet "0/0/20"
    mstp-profile "mstpPortfast"
!
interface gigabitethernet "0/0/24"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/28"
    mstp-profile "mstpPortfast"
!
interface gigabitethernet "0/0/3"
    lACP-profile "pc0"
!
interface gigabitethernet "0/0/36"
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/42"
    lACP-profile "pc2"
!
interface gigabitethernet "0/0/43"
    lACP-profile "pc2"
!
interface gigabitethernet "0/0/46"
    shutdown
    switching-profile "trunk-profile"
!
interface gigabitethernet "0/0/47"
    shutdown
    switching-profile "trunk-profile"
!
interface vlan "4093"
!
interface mgmt
    ip address 10.16.56.62 netmask 255.255.255.0
!
interface port-channel "0"
    switching-profile "trunk-profile"
!
interface port-channel "2"
    switching-profile "trunk-profile"
!

```


The implementation of Rapid PVST+ (Per-VLAN Spanning Tree Plus) is based on the IEEE Standards 802.1D-2004 and 802.1Q-2005 ensuring interoperability with industry accepted PVST+ protocols. In addition, Rapid PVST+ supports the loopguard, rootguard, bpduguard, and portfast features.



To enable PVST+ , use the spanning tree mode command.

Rapid PVST+ runs a separate spanning tree instance for each Virtual Local Area Network (VLAN). This allows the port to forward some VLANs while blocking other VLANs. PVST+ provides for load balancing of VLANs across multiple ports resulting in optimal usage of network resources.

Convergence occurs rapidly with Rapid PVST+. By default, each designated port in the spanning tree protocol sends out a BPDUs (Bridge Protocol Data Units) every 2 seconds. On a designated port in the topology, if hello messages are missed three consecutive times, or if the maximum age expires, the port immediately flushes all protocol information from the table. A port considers that it loses connectivity to its direct neighbor designated port when it misses three BPDUs or if the maximum age expires. This rapid aging of the protocol information allows for quick failure detection.

Rapid PVST+ provides for rapid recovery of connectivity following the failure of a device, a device port, or a LAN. It provides rapid convergence for edge ports, new root ports, and ports connected through point-to-point links.

This chapter covers:

- [“Important Points to Remember” on page 153](#)
- [“Configuring PVST+” on page 153](#)
- [“Loopguard and Rootguard” on page 155](#)
- [“Bridge Protocol Data Unit \(BPDU\) Guard” on page 157](#)

Important Points to Remember

- Configure Rapid PVST+ using the command line only.
- If your Mobility Access Switch is terminated on a router/switch spanning tree environment running PVST+, your Mobility Access Switch must be in PVST mode (**spanning-tree mode pvst** command).
- Once in Rapid PVST+ mode, a predefined *non-editable* PVST profile automatically associates all configured VLANs (including default VLAN 1) and PVST+ starts running on all configured VLANs.
- Rapid PVST+ inter-operates seamlessly with IEEE and PVST bridges when the Mobility Access Switch is placed in a network.

Configuring PVST+

You configure Rapid PVST+ via two profiles; the VLAN profile that enables you to configure the Rapid PVST+ properties and the interface-based profile that enables you to configure your Rapid PVST+ port properties.

Configuring using the VLAN Profile

Set the spanning tree mode to PVST+, assign a profile name, attach the profile to a VLAN, then configure PVST+ properties.

1. Set the spanning tree mode to PVST+.

```
(host)(config) #spanning-tree mode pvst
```

Verify the spanning tree mode:

```
(host)(config) #show spanning-tree-profile
```

```
spanning-tree
-----
Parameter          Value
-----
spanning-tree-mode  pvst
```

2. Assign a PVST+ profile name; in the example below the profile name is “techpubs”:

```
(host)(config) #vlan-profile pvst-profile techpubs
```

```
(host)(pvst-profile "techpubs") #
```

3. Attach the named profile to a VLAN; in the example below the profile name “techpubs” is attached to VLAN 1:

```
(host)(config) #vlan 1#
```

```
(host)(VLAN "1") #pvst-profile techpubs
```

4. View the other PVST+ options settings (such as forward delay, hello time and maximum age).

```
(host)(pvst-profile "techpubs") # ?
```

```
bridge-priority      Bridge-priority [0-61440 in steps of 4096]. Default:
                    32768
clone                Copy data from another pvst-profile
enable               Enable or disable PVST+ bridge.
forward-delay        Forward-delay in seconds [4-30]. Default: 15 seconds
hello-time           Hello-time in seconds [1-10]. Default: 2 seconds
max-age              Maximum age in seconds [6-40]. Default: 20 seconds
no                   Delete Command
```

5. To change one of the value, for example bridge hello time, execute the following command:

```
(host)(pvst-profile "techpubs") #hello-time 5
```

6. Then verify your change:

```
(host)(pvst-profile "techpubs") #show vlan-profile pvst-profile techpubs
```

```
pvst-profile "TechPubs"
-----
Parameter          Value
-----
Enable PVST+ bridge  Enabled
bridge priority      32768
bridge hello time    5      <—forward delay changed from 2 to 5 seconds
bridge forward delay 15
bridge maximum age   20
```

Disable PVST+ on a VLAN

The following example disables the PVST+ profile “techpubs” and then removes the PVST profile from VLAN 1.

```
(host)(config) #vlan-profile pvst-profile techpubs
(host)(pvst-profile "techpubs") #no enable
(host)(pvst-profile "techpubs") #exit
(host)(config) #vlan 1
(host)(VLAN "1") #pvst-profile techpubs
(host)(VLAN "1") #
```

Configuring using the Interface-based Profile

The interface-based Rapid PVST+ profile allows you to configure PVST+ port parameters.

1. Name the interface and view the configuration options.

```
(host) (config) #interface-profile pvst-port-profile techpubs
(host) (Interface PVST bridge "techpubs") #?
```

bpduguard	Enable or disable bpduguard
clone	Copy data from another Interface PVST bridge
loopguard	Enable or disable loopguard
no	Delete Command
point-to-point	Enable or disable point-to-point
portfast	Enable or disable portfast
rootguard	Enable or disable rootguard
vlan	spanning tree [1-4094]

2. Use any of the command options to further configure your interface-based profile.

```
(host)(Interface PVST bridge "techpubs") #vlan 3 cost 8
(host)(Interface PVST bridge "techpubs") #vlan 3 priority 240
```

Then verify your configuration. Notice that the cost and priority values include the original default value and the current value.

```
(host)(Interface PVST bridge "techpubs") #show interface-profile pvst-port-profile techpubs
```

```
Interface PVST bridge "techpubs"
-----
Parameter                               Value
-----
spanning tree port cost                  3 8 <—new value is displayed
spanning tree port priority              3 240 <—new value is displayed
Enable point-to-point                    Enabled
Enable portfast                          Disabled
Enable rootguard                         Disabled
Enable loopguard                         Disabled
```

Loopguard and Rootguard

Rapid PVST+ supports the loopguard and rootguard features.

Configuring Loopguard

Loopguard provides additional protection against Layer 2 forwarding loops (spanning tree loops). A spanning tree loop is created when a spanning tree blocking port, in a redundant topology, erroneously transitions to the forwarding state. This usually happens because one of the ports of a physically redundant

topology (not necessarily the spanning tree blocking port) is no longer receiving spanning tree BPDUs (Bridge Protocol Data Units).

If loopguard is enabled on a non-designated port receiving BPDUs, then that non-designated port is moved into the spanning tree loop-inconsistent blocking state.

Enable loopguard:

```
(host)(Interface PVST bridge "techpubs") #loopguard
```

Associate to the interface:

```
(host)(config) #interface gigabitethernet 0/0/2
(host)(gigabitethernet "0/0/2") #pvst-port-profile techpubs
```

Configuring Rootguard

Rootguard provides a way to enforce the root bridge placement in the network. The rootguard feature guarantees that a port will not be selected as Root Port. If a bridge receives superior spanning tree BPDUs on a rootguard-enabled port, the port is selected as an Alternate Port instead of Root Port and no traffic is forwarded across this port.

By selecting the port as an Alternate Port, the rootguard configuration prevents bridges, external to the region, from becoming the root bridge and influencing the active spanning tree topology.

Enable rootguard:

```
(host)(Interface PVST bridge "techpubs") #rootguard
```

Associate to the interface:

```
(host)(config) #interface gigabitethernet 0/0/2
(host)(gigabitethernet "0/0/2") #pvst-port-profile techpubs
```

Verifying the Configuration

Use the show interface-profile command to view the status of loopguard and rootguard.

```
(host) #show interface-profile pvst-port-profile techpubs

Interface PVST bridge "techpubs"
-----
Parameter                                Value
-----
Instance port cost                       3 8
Instance port priority                   3 240
Enable point-to-point                   Enabled
Enable portfast                         Enabled
Enable rootguard                        Enabled<—rootguard is enabled
Enable loopguard                        Enabled<—loopguard is enabled
Enable bpduguard                       Enabled
Enable bpduguard auto recovery time    60
```

Bridge Protocol Data Unit (BPDU) Guard

The BPDU guard functionality prevents malicious attacks on edge ports. When the malicious attacker sends a BPDU on the edge port, it triggers unnecessary STP calculation. To avoid this attack, use the BPDU guard on that edge port. The BPDU guard enabled port shuts down as soon as a BPDU is received.

Enabling and Configuring BPDU Guard Functionality

The BPDU Guard functionality can be enabled or disabled at an interface level. By default, the BPDU is disabled. The BPDU guard functionality can now be configured as part of the `pvst-port-profile` configuration.

You can use the following command to configure the BPDU guard by using the PVST profile:

```
(host) (config) #interface-profile pvst-port-profile <profile-name>
    bpduguard
    auto-recovery-time <recovery-time>
```

The following example shows how to enable and configure the BPDU guard functionality:

```
(host)(config)# interface-profile pvst-port-profile BPDU-Guard1
    bpduguard auto-recovery-time 60
```



You can configure BPDU guard with or without the `auto-recovery-time` option.

You can disable the BPDU guard functionality by using the following command:

```
(host) (config) #interface-profile <profile-name> no bpduguard
```

You can disable the auto recovery time by using the following command:

```
(host) (Interface PVST bridge "profile-name") #bpduguard no auto-recovery-time
```

Verifying the BPDU Guard Configuration:

```
(host) (config) #show interface-profile pvst-port-profile bpdu
```

```
Interface PVST bridge "bpdu"
```

```
-----
Parameter                                     Value
-----
Instance port cost                           N/A
Instance port priority                       N/A
Enable point-to-point                       Disabled
Enable portfast                             Disabled
Enable rootguard                            Disabled
Enable loopguard                            Disabled
Enable bpduguard                            Enabled <——BPDU guard is enabled
Enable bpduguard auto recovery time         N/A
```

Sample Configuration

To enable and configure BPDU guard using the PVST profile:

```
(host)(config)# interface-profile pvst-port-profile BPDU-Guard1
    bpduguard auto-recovery-time 60
```

To attach the PVST profile to the interface:

```
(host) (config)# interface gigabitethernet <0/0/6>
    pvst-port-profile BPDU-Guard1
```

Portfast

When the link on a bridge port goes up, PVST+ runs its algorithm on that port. If the port is connected to a host that does not “speak” PVST+, it takes approximately 30 seconds for the port to transition to the forwarding state. During this time, no user data passes through this bridge port and some user applications may time out.

Configuring Portfast

To immediately transition the bridge port into the forwarding state upon linkup, enable the PVST+ portfast feature.

```
(host)(config) #interface-profile pvst-port-profile techpubs
(host)(Interface PVST bridge "techpubs") #portfast
```

The bridge port still participates in PVST+; if a BPDU is received, it becomes a normal port.



NOTE

Portfast is operational only on access ports.

Verify the Configuration

Use the show interface-profile command to view the status of the portfast.

```
(host) (config) #show interface-profile pvst-port-profile bpdu
```

```
Interface PVST bridge "bpdu"
```

```
-----
```

Parameter	Value
-----	-----
Instance port cost	N/A
Instance port priority	N/A
Enable point-to-point	Disabled
Enable portfast	Disabled <— portfast is enabled
Enable rootguard	Disabled
Enable loopguard	Disabled
Enable bpduguard	Enabled
Enable bpduguard auto recovery time	N/A

The Hot-Standby Link (HSL) feature is a simplified failover mechanism. HSL enables a Layer 2 interface (or port-channel) to back-up another Layer 2 interface (or port-channel) so that these interfaces become mutual backups.

HSL consists of a pair of redundant links. One is the *primary* for traversing traffic, and the other is the *backup*. When the primary fails, a rapid traffic failover occurs to the awaiting backup.

One of the primary use cases for HSL is in an enterprise topology where each access switch is dual-homed to two distribution/core switches for redundancy purpose.

Important Point to Remember

- Spanning tree (MSTP and PVST+) must be disabled before configuring HSL. HSL and spanning tree can not be configured on the same system at the same time.
- HSL is a 1:1 ratio for primary and backup pairs. One backup interface can not be the backup of multiple primary interfaces. An interface can be part of only one HSL pair.
- HSL links are always trusted.
- Primary and backup interfaces must have the same switching profiles.
- Primary and backup interfaces cannot be members of the same port-channel.
- The interfaces cannot be Tunneled Node interfaces.

Configuration Steps

When a primary link goes down, the backup link becomes active. By default, when the link comes up it goes into the standby mode as the other interface is activated. You can force the primary interface to become active by enabling preemption.

Configure HSL directly in the interface. First, on the primary interface (for example 0/0/10), then specify the back-up interface (for example 0/0/11). Use the following steps, from the command line, to configure and verify HSL.

1. Configure the primary and backup interfaces.

```
(host) (config) #interface gigabitethernet 0/0/10
(host) (gigabitethernet "0/0/10") #backup interface gigabitethernet 0/0/11
```

2. Configure pre-emption if necessary (it is off by default).

```
(host) (gigabitethernet "0/0/10") #preemption mode forced
```

3. If pre-emption is configured, best practices recommends configuring *delay*. The range is 10 seconds to 5 minutes (300 seconds); default is 100 seconds.

```
(host) (gigabitethernet "0/0/10") #preemption delay 10
```

4. Verify the HSL configuration. The following show command is a partial output.

```
(host) #show interface-config gigabitethernet 0/0/10

gigabitethernet "0/0/10"
-----
Parameter                               Value
-----
Interface MSTP Profile                   disabled
...
Interface Trusted Mode                   Enabled
HSL backup interface                     gigabitethernet0/0/11
HSL preemption mode                      Forced
HSL preemption delay                     10
...
```

To view details of HSL on an interface, use the following show commands.

```
(host) #show hot-standby-link gigabitethernet 0/0/10

HSL Interface Info
-----
Primary Interface: GE-0/0/10 (Active)  Backup Interface:  GE-0/0/11 (Standby)
Preemption Mode: forced                  Preemption Delay: 10
Last Switchover Time: NEVER              Flap Count: 0
```

To view details of all HSL links, use the following show command.

```
(host) #show hot-standby-link

HSL Interfaces Info
-----
Primary    State    Backup    State    Last Switchover Time
-----
GE-0/0/10  Active  GE-0/0/11  Standby  Never
GE-0/0/3   Down    PC-4        Down     Never
PC-1       Down    GE-0/0/0    Active   Never
PC-2       Down    PC-3        Down     Never
```


This chapter describes the following topics related to Generic Router Encapsulation (GRE):

- “Overview” on page 161
- “Configuring an L2-GRE Tunnel” on page 161
- “Understanding the VLAN Membership of Existing L2 GRE Tunnel” on page 162
- “Sample Configuration” on page 163


Overview



Only Layer 2 (L2) GRE is supported in this release. GRE and L2-GRE may be used interchangeably in this document.

This release of ArubaOS Mobility Access Switch supports L2 connectivity through GRE tunnel. L2-GRE tunnel extends VLANs across Mobility Access Switches and Aruba controllers. GRE encapsulates Layer-2 frames with a GRE header and transmit through an IP tunnel over the cloud. Following figure shows how L2-GRE tunnel fits into network operations.

Figure 1 L2-GRE Tunnel Network Topology



Configuring an L2-GRE Tunnel

To configure an L2-GRE tunnel, see the following procedure.

```
(host) (config) #interface tunnel ethernet <tunnel-id>
(host) (Tunnel "tunnel-id") #description <interface-description>
(host) (Tunnel "tunnel-id") #source-ip <source-tunnel-ip>
(host) (Tunnel "tunnel-id") #destination-ip <destination-tunnel-ip>
(host) (Tunnel "tunnel-id") #switching-profile <profile-name>
(host) (Tunnel "tunnel-id") #keepalive <Tunnel heartbeat interval in seconds (1-86400)>
<Tunnel Heartbeat Retries (1-1024)>
```

Inter-tunnel flooding

There can be multiple L2-GRE tunnels terminating on the same device, either ArubaOS Mobility Access Switch or Mobility Controller. If the tunnels carry same VLANs, this may cause inter-tunnel flooding resulting in loops within the network. To avoid this scenario, disable inter-tunnel flooding in the switch and the controller.

```
(host) (config) #interface tunnel ethernet <tunnel-id>
(host) (Tunnel "tunnel-id") #no inter-tunnel-flooding
```

For additional parameters, see *ArubaOS 7.2 Command Line Interface* guide.

Understanding the VLAN Membership of Existing L2 GRE Tunnel

You can use the following commands to understand the VLAN membership of L2 GRE tunnel which is already configured.

Use the following command to check the VLAN membership of the existing L2 GRE tunnel:

```
(host) #show interface tunnel <tunnel-id>
tunnel 10 is administratively Up, Line protocol is Down
Description: GRE Interface
Internet address is unassigned
Source <source_IP>
Destination <destination_IP>
Protocol number 0
Tunnel mtu is set to 1100
Tunnel is an L2 GRE Tunnel
Tunnel is Trusted
Inter Tunnel Flooding is enabled
Tunnel keepalive is enabled
Tunnel keepalive interval is 3 seconds, retries 3
    Heartbeats sent 51347, Heartbeats lost 51346
    Tunnel is down 4 times
Switching-profile "100"
```

```
(host) #show interface-config tunnel <tunnel-id>
```

```
Tunnel "10"
-----
Parameter                Value
-----
Tunnel Description        N/A
Tunnel Source IP          <source_IP>
Tunnel Destination IP     <destination_IP>
Inter-Tunnel-Flooding     Enabled
Tunnel Mode               L2
Tunnel Protocol           0
Tunnel Keepalive          3/3
Tunnel MTU                1100
Tunnel Shutdown           Disabled
Tunnel Switching Profile  100
Tunnel Trusted            Enabled
```

This shows that Switching-Profile “100” is applied in L2 GRE tunnel interface. You can use the `show interface-profile switching-profile 100` command to view the VLAN configuration.

```
(host) #show interface-profile switching-profile 100
```

```
switching profile "100"
-----
Parameter                Value
-----
Switchport mode          access
Access mode VLAN         100
Trunk mode native VLAN   1
Enable broadcast traffic rate limiting Enabled
Enable multicast traffic rate limiting Disabled
Enable unknown unicast traffic rate limiting Enabled
```

```
Max allowed rate limit traffic on port in percentage 50
Trunk mode allowed VLANs 1-4094
```

You can use the `show vlan` command to view the port associated with the vlan:

```
(host) #show vlan
```

VLAN CONFIGURATION

```
-----
VLAN  Description  Ports
----  -
1      VLAN0001      GE0/0/1-19 GE0/0/21-26 GE0/0/28-33 GE0/0/35-36
                        GE0/0/38-47 GE0/1/0-3 GRE-TUN30
10     VLAN0010      GE0/0/34 Pc1
11     VLAN0011      GE0/0/34
20     VLAN0020      GE0/0/20
100    VLAN0100      GE0/0/0 GE0/0/27 GRE-TUN10 GRE-TUN20
```



MAC address learned on L2 GRE tunnel does not honor `mac-aging-timer` configuration , and ages out at 270 seconds.

Sample Configuration

To configure an L2-GRE tunnel and apply the switching profile:

```
(host) (config) #interface tunnel ethernet 1
(host) (Tunnel "1") #description L2-GRE_Interface
(host) (tunnel "1") #source-ip 10.0.0.1
(host) (tunnel "1") #destination-ip 10.0.1.2
(host) (tunnel "1") #switching-profile mDNS_vlan_200
(host) (tunnel "1") #keepalive 30 5
```



In the above example, **mDNS_vlan_200** was previously defined.

This chapter describes the Layer 3 Routing features available on the Mobility Access Switch. It contains the following sections:

- “Understanding Routed VLAN Interfaces” on page 165
- “Multinetting” on page 167
- “Network Address Translation” on page 168
- “IP Directed Broadcast” on page 169
- “Static Routes” on page 170
- “Route Metrics” on page 172
- “Equal Cost Multipath” on page 172
- “IP Prefix List” on page 173

Understanding Routed VLAN Interfaces

Routed VLAN Interfaces (RVI) are logical interfaces that enable routing and bridging between VLANs. You can route and bridge a protocol on the same interface. The traffic that remains in the bridge group (the bridged traffic) will be bridged among the bridged interfaces, and the traffic that needs to go out to another network (the routed traffic) will be routed internally to the appropriate output routed interface.

There can be an IPv4 address to each VLAN interface. You can also configure IGMP and PIM interface profiles to the VLAN interfaces. A total of 4094 routed VLAN interfaces can be configured in this release. VLAN interface 1 is configured by default.

Important Points to Remember

- The maximum number of VLAN interfaces supported are 4094.
- The Layer 2 VLAN must be configured before configuring the corresponding RVIs.
- The protocol status of a RVI is in up state only when the protocol status of at least one member port in the corresponding VLAN is in up state.

To assign member ports to a VLAN, create a switching profile with the corresponding VLAN, and assign the switching profile to the member interfaces.

Configuring Routed VLAN Interfaces

You can configure routed VLAN interfaces using the CLI.

Using the CLI

To configure routed VLAN interfaces, follow these steps:

1. Create the required VLANs.

```
(host)(config)# vlan <vlan-id>
exit
```

2. Create the switching profiles and reference the existing VLANs.

```
(host)(config)# interface-profile switching-profile <profile-name>
```

```

switchport-mode {access|trunk}
access-vlan <vlan-id>
trunk allowed vlan <vlan-list>
native-vlan <vlan-id>
exit

```

3. Apply the switching profiles to the physical interfaces.

```

(host)(config)# interface gigabitethernet <slot/module/port>
switching-profile <profile-name>
exit

```

4. Create the VLAN interfaces.

```

(host)(config)# interface vlan <vlan-id>
description <vlan-interface-description>
dhcp-relay-profile <profile-name>
igmp-profile <profile-name>
ip {address {{<ip-address> netmask <subnet-mask>}} | dhcp-client} | directed-
broadcast | nat {inside}}
ipv6 address {{<prefix> netmask <subnet-mask>}} | link-local <link-local>}
mtu <64-9216>
shutdown
no {...}
ospf-profile <profile-name>
pim-profile <profile-name>
exit

```

Sample Configuration

The following example shows how to configure a routed VLAN interface:

```

(host)(config) #interface vlan 10
(host)(vlan "10") #ip address 10.10.10.10 netmask 255.255.255.0
(host)(vlan "10") #description layer 3
(host)(vlan "10") #mtu 1500
(host)(vlan "10") #exit

```

Verifying and Monitoring the VLAN Interfaces

You can verify and monitor the VLAN interfaces using the following commands:

```

(host)#show interface vlan 10
VLAN10 is administratively Up, Line protocol is Up
Hardware is CPU Interface, Address is 00:0b:86:6a:f2:40
Description: layer3
Internet address is 10.10.10.10, Netmask is 255.255.255.0
IPv6 link-local address not assigned
Global Unicast address(es):
Routing interface is enable, Forwarding mode is enable
Directed broadcast is disabled, BCMC Optimization disabled
Encapsulation 802, Loopback not set
Interface index: 50331658
MTU 1500 bytes

```

Displaying VLAN Information

```

(host)# show vlan 10
VLAN CONFIGURATION
-----
VLAN  Description  Ports

```

```
-----
10      VLAN0010      GE0/0/2 GE0/0/16
```

Displaying IP Interface Brief

```
(host)#show ip interface brief
```

Interface	IP Address / IP Netmask	Admin	Protocol
vlan 1	unassigned / unassigned	Up	Up
vlan 7	20.2.1.1 / 255.255.255.0	Up	Down
vlan 10	10.10.10.10 / 255.255.255.0	Up	Down
loopback 3	1.1.1.1 / 255.255.255.0	Up	Up
mgmt	10.4.135.11 / 255.255.255.0	Up	Up

Multinetting

ArubaOS supports multiple IP addresses per VLAN and loopback interface. This allows the user to specify any number of secondary IP addresses. Secondary IP address can be used in a variety of situations, such as the following:

- If an insufficient number of host addresses are available on a particular network segment. Using secondary IP addresses on the routers or access devices allows you to have two logical subnets using one physical subnet
- If the an older network is built using Layer 2 bridges and has no subnetting. Secondary addresses can aid in the transition to a subnetted, router-based network.
- Two subnets of a single network might be otherwise seperated by another network. You can create a single network from subets that are physically seperated by another network using a secondary address.

Important Points to Remember

- OSPF advertises the secondary IP address in the router LSA but it does not form adjacency on the secondary IP address.
- PIM will not send hello packets on the secondary IP address.
- DHCP servers identify the subnets associated with secondary IP addresses used for allocation.

Configuring Secondary IP

To configure a secondary IP address, use the following command:

```
(host) (vlan "1") #ip address 1.1.1.1 255.255.255.0 ?
secondary          Make this IP address a secondary address
```

Sample Configuration

```
(host) (config) #interface vlan 2
(host) (vlan "2") #ip address 1.1.1.1 255.255.255.0 secondary

(host) (vlan "2") #show interface vlan 2

VLAN2 is administratively Up, Line protocol is Up
Hardware is CPU Interface, Address is 00:0b:86:6a:1c:c0
Description: 802.1Q VLAN
Internet address is 20.20.20.1, Netmask is 255.255.255.0
Internet address is 1.1.1.1, Netmask is 255.255.255.0 secondary
IPv6 link-local address is fe80::b:8600:26a:1cc0
Global Unicast address(es):
```

Routing interface is enable, Forwarding mode is enable
Directed broadcast is disabled, BCMC Optimization disabled

Network Address Translation

Aruba Mobility Access Switches support source Network Address Translation (NAT) with Port Address Translation (PAT) on VLAN interfaces. When source NAT is enabled on a VLAN interface, the IP address of the egress VLAN interface as determined by the routing table will be used as the source IP. For example, if "ip nat inside" is enabled on interface VLAN X and traffic will be routed out interface vlan Y, the IP address of interface VLAN Y will be used as the source IP for traffic from VLAN X

```
(host) (config) #interface vlan <vlan_id>
(host) (vlan "vlan_id") #ip nat inside
```



No packet fragmentation is supported by NATing.

To verify source NAT is enabled on a VLAN interface, use **show interface vlan <vlan-id>**. In the following example, source NAT has been enabled on interface VLAN 6. As a result, the output of **show interface vlan <vlan-id>** will included the bolded section below. If the bolded section is not displayed, source NAT has not been enabled.

```
(host) # show interface vlan 6

VLAN6 is administratively Up, Line protocol is Up
Hardware is CPU Interface, Address is 00:0b:86:6a:5d:c0
Description: 802.1Q VLAN
Internet address is 6.1.1.1, Netmask is 255.255.255.0
IPV6 link-local address is fe80::b:8600:66a:5dc0
Global Unicast address(es):
Routing interface is enabled, Forwarding mode is enabled
Interface is source NAT'ed
Directed broadcast is disabled, BCMC Optimization disabled
Encapsulation 802, Loopback not set
Interface index: 50331654
MTU 1700 bytes
```

Additionally, you can use the **show datapath vlan** command to verify that source NAT has been enabled.

```
(host) #show datapath vlan

Datapath VLAN Table Entries
-----
Flags: N - Nat Inside, M - Route Multicast, R - Routing
      S - Snoop MLD, G - Snoop IGMP, P - Proxy IGMP
      B - BCMC Optimization, A - Proxy ARP, U - Suppress ARP
      1(cert-id) - 8021X Term-PEAP, 2(cert-id) - 8021X Term-TLS
VLAN  Flags          Ports
----  -
6      NRU           1/0/14
100    RU            0/0/14
```

The **show datapath session** command can be used to to verify the packet flows that are being NAT'ed. This output however will not indicate the interface VLAN the flow(s) are using. To determine that information use the **show ip interface brief** command.


```
(host) #show datapath session
```

```
Datapath Session Table Entries
```

```
-----  
Flags: F - fast age, S - src NAT, N - dest NAT  
D - deny, R - redirect, Y - no syn  
H - high prio, P - set prio, T - set ToS  
C - client, M - mirror, V - VOIP  
Q - Real-Time Quality analysis  
I - Deep inspect, U - Locally destined  
E - Media Deep Inspect, G - media signal  
u - User Index
```

Source IP	Destination IP	Prot	SPort	DPort	Cntr	Prio	ToS	Age	Destination	TAge	UsrIdx	UsrVer	Flags
6.1.1.5	100.1.1.6	61	0	0	0/0		0 0	0	1/0/14	1	0	0	FSC
100.1.1.6	100.1.1.7	61	0	0	0/0		0 0	0	1/0/14	1	0	0	FNY

```
(host) #show ip interface brief
```

Interface	IP Address / IP Netmask	Admin	Protocol
vlan 100	100.1.1.7 / 255.255.255.0	Up	Up
vlan 6	6.1.1.1 / 255.255.255.0	Up	Up

IP Directed Broadcast

An IP directed broadcast is typically used by network management systems (NMS) for features like Wake On LAN to broadcast packets on a local subnet even though the source of that broadcast is located on a remote subnet. When the source device initiates this broadcast packet, it is routed through the network as a unicast packet until it reaches the target subnet. Other than the router directly attached to the target subnet, all routers across the network view it as a unicast packet. The router directly attached to the target subnet identifies the packet as a directed broadcast, converts it to a link-layer broadcast packet and propagates it across the target subnet.

This feature is disabled by default. When disabled, the directed broadcast packets are dropped unconditionally without generating an ICMP error packet. Due to the nature of propagating broadcast, Aruba does not recommend enabling this parameter as it can result in Denial of Service (DoS) attacks, if not used correctly. When absolutely necessary, you can enable this feature on a subnet by subnet basis. You can enable this feature on the Routed VLAN Interfaces (RVI) in the CLI.

Configuring IP Directed Broadcast

```
(host) (config) #interface vlan <id>  
(host)(vlan) #ip directed-broadcast
```

Sample Configuration

The following example shows how to configure a routed VLAN interface and enable IP directed broadcast:

```
(host)(config) #interface vlan 10  
(host)(vlan "10") #ip address 10.10.10.10 netmask 255.255.255.0  
(host)(vlan "10") #ip directed-broadcast  
(host)(vlan "10") #description layer 3  
(host)(vlan "10") #mtu 1500  
(host)(vlan "10") #exit
```

You can verify the preceding configuration using the following command:

```
(host)#show interface vlan 10  
VLAN10 is administratively Up, Line protocol is Up  
Hardware is CPU Interface, Address is 00:0b:86:6a:f2:40  
Description: layer3
```

```
Internet address is 10.10.10.10, Netmask is 255.255.255.0
IPv6 link-local address not assigned
Global Unicast address(es):
Routing interface is enable, Forwarding mode is enable
Directed broadcast is enabled, BCMC Optimization disabled
Encapsulation 802, Loopback not set
Interface index: 50331658
MTU 1500 bytes
```

Static Routes

The Mobility Access Switch supports static routes configuration. You can configure a default gateway and multiple static routes within the global IP-profile to route packets outside the local network. The static routes are active or added to the routing table only when the next hop is reachable, and can be removed from the static routes list only by using the `no` command.

Important Points to Remember

- You can have only one default gateway. However, you can have multiple static routes.
- You can have both an IPv4 and an IPv6 default gateway simultaneously.
- Static routes become active only when the nexthop is reachable.
- Nexthops have to be within the local network.

The Default Gateways

Default gateway is a special case of static route where the destination mask and prefix is 0/0. The next hop in a default gateway can be any valid IP address which can be reached through a routable or the management interface.

Configuring the Default Gateways and the Static Routes

You can configure the static routes within the global IP-profile. Each static route needs a destination, netmask and nexthop addresses.

The static routes are inserted in to the Forwarding Information Base (FIB), only when the nexthop matches the subnet of any of the RVI interfaces or the management interface. If the nexthop becomes unreachable, the Routing Information Base (RIB) gets purged but the static route is still retained. The static route can be completely removed from the system only by using the `no` command within the IP-profile.

You can configure the default gateways and the static routes using the CLI. You can also configure static routes using the WebUI.

Using the WebUI

1. Navigate to the **Configuration > Routing** page.
2. Click **New** under the static routes list.
3. Click on the **Destination IP** column and enter the destination IP address.
4. Click on the **Destination Mask** column and enter the destination netmask address.
5. Click on the **Next Hop** column and enter the nexthop IP address.
6. Click on the **Metric** column and enter the metric.
7. Press **Enter**.

Using the CLI

```
(host)(config) #ip-profile
```

controller-ip	Configure controller IP
default-gateway	Specify default gateway
no	Delete Command
prefix-list	Configure prefix list
route	Configure static route A.B.C.D

Sample Configuration

```
(host)(config) #ip-profile
(host)(ip-profile) #default-gateway 2.2.2.2
(host)(ip-profile) #no default gateway
(host)(ip-profile) #default-gateway import dhcp
(host)(ip-profile) #route 20.20.31.0 255.255.255.0 10.10.10.31
(host)(ip-profile) #route 20.20.32.0 255.255.255.0 10.10.10.32
(host)(ip-profile) #route 20.20.33.0 255.255.255.0 10.10.10.33
(host)(ip-profile) #no route 20.20.34.0 255.255.255.0 10.10.10.20
```

Verifying the IP Routes

```
(host) #show ip route
Codes: C - connected, O - OSPF, R - RIP, S - static
       M - mgmt, U - route usable, * - candidate default
Gateway of last resort is 10.18.7.254 to network 0.0.0.0 at cost 39
S    0.0.0.0/0 [39/0] via 10.18.7.254
C    10.10.10.0 is directly connected: vlan1
C    10.10.10.1 is directly connected: vlan1
C    10.10.10.20 is directly connected: vlan1
C    10.10.10.31 is directly connected: vlan1
C    10.10.10.32 is directly connected: vlan1
C    10.10.10.33 is directly connected: vlan1
M    10.18.7.0 is connected mgmt-intf: 10.18.7.125
M    10.18.7.125 is connected mgmt-intf: 10.18.7.125
M    10.18.7.254 is connected mgmt-intf: 10.18.7.125
S    20.20.31.0 [0] via 10.10.10.31
S    20.20.32.0 [0] via 10.10.10.32
S    20.20.33.0 [0] via 10.10.10.33
```

```
(host) #show ip route summary
```

```
Route Source  Total
-----
connected      6
static          5
ospf-intra      0
ospf-inter      0
ospf-ext1       0
ospf-ext2       0
ospf-nssa       0
```

```
(host) #show arp
```

```
IPV4 ARP Table
-----
Protocol  IP Address      Hardware Address  Interface
-----
Internet  40.40.40.252    00:0b:86:64:a8:c0  vlan40
```

Clearing the ARP Table

```
(host) #clear arp {<all>|<ip-address>}
```

Route Configuration Limits

The following table specifies the maximum number of routes and nexthops you can have in a Mobility Access Switch:

Table 1 *Route Configuration Limits*

Type of Route/Nexthop	Maximum Routes Supported
IPv4 Unicast + IPv4 Multicast Groups	6912
IPv4 Multicast Sources	1024
IPv6 Unicast + IPv6 Multicast Groups + IPv6 Multicast Sources	320
Address Resolution Protocol	4096 (3k distinct MACs)
Multicast downstream interface table	4096

Route Metrics

The Mobility Access Switch includes support for route metrics. For a given route destination, there can be multiple nexthops. A route metric enables the Mobility Access Switch to prefer one route over another or load balance when the metric is the same. For more details on load balancing across multiple nexthops, see [“Equal Cost Multipath” on page 172](#).

A route destination with a lower metric is added to the route manager. The higher metric routes are added only when the lower metric routes are removed.

The following example shows how to add a metric of 10 to a static route:

```
(host) (ip-profile) # route 192.168.1.0 255.255.255.0 192.168.2.1 10
```

Equal Cost Multipath



No commands are necessary to enable ECMP.

Equal Cost Multipath (ECMP) enables Mobility Access Switch to forward the data packets to any of the multiple nexthops of a routing destination. The route manager identifies the best routing destination based on the priority of the protocol. After the route manager identifies the best route, all the nexthops of that route are used for datapath forwarding. ECMP is auto-enabled and does not require any command to enable it.

ECMP provides flow-based load balancing for the chosen routing destination. For a given flow same nexthop is used to forward all the packets. For multiple flows, load balancing happens across multiple nexthops. ECMP uses the source IP and destination IP to define a flow. For TCP/UDP packets, it also uses the source and destination ports to define the flow. ECMP automatically load balances the traffic when multiple nexthops with equal cost exist

Apart from multiple nexthops, ECMP also enables addition of metric for a route. ECMP nexthops are per metric basis. For a given metric, there can be multiple nexthops (up to 4). A route with a lower metric is added to the route manager. The higher metric routes are added only when the lower metric routes are deleted.

IP Prefix List

The `ip prefix-list` command is used to configure IP prefix filtering. Prefix lists are used to either permit or deny the configured prefix based on the matching condition. The prefix list consists of an IP address and a bit mask. The IP address can be classful network, a subnet, or a single host route.



Any traffic that does not match any prefix-list entry is denied.

```
(host) (config) #ip-profile
(host) (ip-profile) #prefix-list <prefix-list-name>
    seq <sequence-number>
    deny|permit
    <network prefix A.B.C.D>
    <network mask A.B.C.D>
    ge <bit-length>|le <bit-length>
```

```
(host) (ip-profile) #prefix-list test seq 1 permit 5.5.5.0 255.255.255.0 ge 32
```

Parameter	Description
prefix-list	Prefix list name.
seq <sequence-number>	Sequence number. Prefix lists are evaluated starting with the lowest sequence number and continue down the list until a match is made. Once a match is made, the permit or deny statement is applied to that network and the rest of the list is ignored.
deny <network-prefix> <network mask>	Specify IPv4 packets to reject.
permit <network-prefix> <network mask>	Specify IPv4 packets to forward.
ge <bit-length>	Minimum prefix length to be matched.
le <bit-length>	Maximum prefix length to be matched.

If only a `ge` value is entered, the range is the value entered for `ge-length` argument to a full 32-bit length. If only the `le` value is entered, the range is from the value entered for `network-length` argument to `le-length` argument. If a `ge` or `le` value is not used, the prefix list is processed using an exact match. If both `ge` and `le` values are entered, the range falls between the values between the values used for the `ge-length` and `le-length` arguments. The behavior can be described as follows:

`network/length < ge-length <= le-length <= 32`



The `ge` and `le` values are optional parameters.

Once you have configured the desired prefix-list entries, you apply them to the global OSPF profile using the following command.

```
(host) (Global OSPF profile) #distribute-list prefix-list <prefix-list name>
```

The following is a sample configuration:

```
(host) (ip-profile) #prefix-list test seq 1 permit 5.5.5.0 255.255.255.0 ge 32
(host) (ip-profile) #prefix-list test seq 2 deny 6.6.6.0 255.255.255.0 ge 32
(host) (ip-profile) #prefix-list test seq 3 permit 10.10.0.0 255.255.255.0 ge 24 le 32
(host) (Global OSPF profile) #distribute-list test
```

Verify the IP Prefix List configuration by using the **show ip-profile** command.

```
(host) (ip-profile) #show ip-profile
```

```
ip-profile "default"
-----
Parameter          Value
-----
Default Gateway     10.18.7.254
Import DHCP Gateway Disabled
controller-ip       N/A
prefix-list test seq 1 permit 5.5.5.0 255.255.255.0 ge 32
prefix-list test seq 2 deny 6.6.6.0 255.255.255.0 ge 32
prefix-list test seq 3 permit 10.10.0.0 ge 24 le 32
```

This chapter describes the DHCP server and relay support on the Mobility Access Switch. It contains the following sections:

- [“Important Points to Remember” on page 175](#)
- [“Understanding DHCP Server and DHCP Relay” on page 175](#)
- [“Configuring DHCP Server and DHCP Relay” on page 176](#)
- [“Verifying DHCP Server and DHCP Relay” on page 178](#)

Important Points to Remember

- DHCP server identifier override sub-option is not supported in this release.

Understanding DHCP Server and DHCP Relay

Dynamic Host Configuration Protocol automates network-parameter assignment to network devices from one or more DHCP servers. Even in small networks, DHCP is useful because it makes it easy to add new machines to the network.

When a DHCP-configured client connects to a network, the DHCP client sends a broadcast query requesting necessary information from a DHCP server. The DHCP server manages a pool of IP addresses and information about client configuration parameters such as default gateway, domain name, the name servers, other servers such as time servers, and so forth.

On receiving a valid request, the server assigns the computer an IP address, a lease (length of time the allocation is valid), and other IP configuration parameters, such as the subnet mask and the default gateway. The query is typically initiated immediately after booting, and must complete before the client can initiate IP-based communication with other hosts.

During initialization, network clients try to dynamically obtain their IP addresses. In small networks, where all the systems are in the same IP subnet, the client and the server can communicate directly.

Clients on subnets that are not directly connected to a DHCP server must go through a "relay agent."

If DHCP relay is not enabled on the VLAN on which the request is received, but a pool is configured for that subnet, the IP is assigned from the internal DHCP server.

DHCP relay is enabled when a DHCP relay profile is attached to a VLAN interface. At this point, the relay agent receives the DHCP broadcast packets from the client and unicast them to one or more of the DHCP servers that are configured on the VLAN interface.

Configuring DHCP Server and DHCP Relay

This section contains the following sections:

- “Configuring DHCP Server” on page 176
- “Configuring DHCP Relay” on page 176
- “Applying DHCP Relay Profile to VLAN” on page 177

Configuring DHCP Server

DHCP server configuration is profile based. To configure the DHCP server, follow these steps:

1. Enable DHCP server configuration.

```
(host)(config) #service dhcp
```

2. Configure a DHCP server profile.

```
(host)(config) #ip dhcp pool pool-1
(host)(dhcp server profile "pool-1") #
```

3. Configure the domain name in the pool profile.

```
(host)(dhcp server profile "pool-1") #domain-name doc-domain
```

4. Configure the DNS servers. Up to 8 DNS servers can be configured.

```
(host)(dhcp server profile "pool-1") #dns-server 192.168.1.2
```

5. Configure the default router. Up to 8 routers can be configured.

```
(host)(dhcp server profile "pool-1") #default-router 192.168.1.1
```

6. Configure the Netbios name server. Up to 8 Netbios name servers can be configured.

```
(host)(dhcp server profile "pool-1") #netbios-name-server 192.168.1.3
```

7. Configure the lease time in days, hours, minutes, and seconds.

```
(host)(dhcp server profile "pool-1") #lease 30 24 60 60
```

8. Configure the network.

```
(host)(dhcp server profile "pool-1") #network 192.168.1.0 255.255.255.0
```

9. Configure the range between two IP addresses to be excluded.

```
(host)(dhcp server profile "pool-1") #exclude-address 192.168.1.1 192.168.1.3
```

10. Configure a vendor-class-identifier.

```
(host)(dhcp server profile "pool-1") #vendor-class-identifier testVendor
```

11. Configure server options.

```
(host)(dhcp server profile "pool-1") #option 50 ip 192.168.1.1
(host)(dhcp server profile "pool-1") #option 54 text server1
```

Configuring DHCP Relay

DHCP-Relay is supported with DHCP Option 82. DHCP Option 82 allows a DHCP relay agent to insert circuit specific information into a request that is being forwarded to a DHCP server.

DHCP Option 82 works by setting two sub-options:

- Circuit ID

The circuit ID includes information specific to the circuit on which the request arrives. Circuit identifier parameters can be interface-name, VLAN ID, or both.

- Remote ID

The remote ID carries information relating to the remote host end of the circuit. Remote identifier parameters can be the MAC address, the hostname of the relay agent, or a user defined string.

DHCP Relay Option 82 can be configured using DHCP Relay profile. To configure a DHCP Relay profile, follow these steps:

1. Configure a DHCP Relay profile under an interface profile.

```
(host)(config) #interface-profile dhcp-relay-profile relay1
```

2. Configure a helper address.

```
(host)(dhcp relay profile "relay1") #helper-address 172.16.30.1
```

3. Configure Option 82 circuit-identifier a VLAN only, an interface-name only or both VLAN and interface-name:

```
(host)(dhcp relay profile "relay1") #option82 circuit-identifier vlan
```

```
(host)(dhcp relay profile "relay1") #option82 circuit-identifier interface-name
```

```
(host)(dhcp relay profile "relay1") #option82 circuit-identifier interface-name vlan
```

4. Configure Option 82 remote-identifier with the host-name option.

```
(host)(dhcp relay profile "relay1") #option82 remote-identifier host-name
```

5. Configure Option 82 remote-identifier as MAC.

```
(host)(dhcp relay profile "relay1") #option82 remote-identifier mac
```

6. Configure Option 82 with the user defined option “myOwnString.”

```
(host)(dhcp relay profile "relay1") #option82 remote-identifier myOwnString
```

Applying DHCP Relay Profile to VLAN

The DHCP relay profile must be applied to the VLAN where DHCP clients connect. To configure a DHCP Relay profile to a VLAN, follow these steps:

1. Configure a VLAN interface.

```
(host)(config) #interface vlan 11
```

2. Configure an IP address on the VLAN interface.

```
(host)(vlan "11") #ip address 172.16.4.1 netmask 255.255.255.0
```

3. Configure DHCP Relay profile on the VLAN interface.

```
(host)(vlan "11") #dhcp-relay-profile relay1
```

Configuring a VLAN with a Relay Profile as DHCP Client

Keep the following points in mind before you configure a VLAN with a relay profile as DHCP client.

Points to Remember

- You can configure both static default gateway and default gateway import from DHCP.
- Static and OSPF routes have preference over DHCP and DHCP has preference over OSPF AS External routes.
- The DHCP routes will be installed only if `default gateway import dhcp` is specified in the `ip-profile`.
- If multiple VLANs act as DHCP clients with the `default-gw import dhcp` option, then the first valid DHCP gateway received in the response will be installed in the routing table.

To configure a VLAN with a relay profile to be a dhcp-client, follow these steps:

1. Configure a VLAN.

```
(host)(config) #interface vlan 4
```

2. Configure a DHCP relay profile.

```
(host)(vlan "4") #dhcp-relay-profile relay1
```

3. Set the IP address of an interface and use DHCP to obtain an IP address.

```
(host)(vlan "4") #ip address dhcp-client
(host)(vlan "4") #end
```

4. Display the VLAN Interface

```
(host)#show interface-config vlan 4
```

```
vlan "4"
-----
Parameter                                Value
-----
Interface OSPF profile                   N/A
Interface PIM profile                   N/A
Interface IGMP profile                   N/A
Interface shutdown                       Disabled
mtu                                       1500
IP Address                              N/A
IPv6 Address                            2012::12/64
IPv6 link local Address                 fe80::b:8600:a6a:3300
DHCP client                             Enabled
DHCP relay profile                       relay1
Interface description                   N/A
```

Verifying DHCP Server and DHCP Relay

This section contains the following sections:

- [“Verifying DHCP Relay Option 82 Logs” on page 178](#)
- [“Show Commands for IP DHCP” on page 179](#)

Verifying DHCP Relay Option 82 Logs

The debug level can be configured to log the DHCP relay messages. It can be configured in network or system logs.

Network Log

```
(host)(config) #logging level debugging network process dhcpd subcat dhcp
```

System Log

```
(host)(config) #logging level debugging system process dhcpd subcat all
```

The DHCP relay functionality can be verified by checking network or system logs as has been configured:

```
Sep 27 07:30:43 dhcpdwrap[1497]: <202523> <DEBUG> |dhcpdwrap| |dhcp| dhcprelay:
dev=eth1, length=341, from_port=67, op=2, giaddr=172.16.4.1
Sep 27 07:30:43 dhcpdwrap[1497]: <202527> <DEBUG> |dhcpdwrap| |dhcp| RelayToClient:
OFFER dest=172.16.4.2 client yiaddr=172.16.4.1 MAC=1c:75:08:9e:60:c8
```

```
Sep 27 07:30:43 dhcpdwrap[1497]: <202541> <DEBUG> |dhcpdwrap| |dhcp| Received DHCP
packet from Datpath, sos msg hdr flags 0x42 opcode 0x5a ingress 0x0 vlan 11 egress 0xb
src mac 00:0b:86:6a:41:40
Sep 27 07:30:43 dhcpdwrap[1497]: <202544> <DEBUG> |dhcpdwrap| |dhcp| Datapath vlan11:
ACK 1c:75:08:9e:60:c8 clientIP=172.16.4.2
```

Show Commands for IP DHCP

This section describes the following commands:

- “[show interface-profile dhcp-relay-profile](#)” on page 179
- “[show ip dhcp database](#)” on page 179
- “[show ip dhcp binding](#)” on page 179
- “[show ip dhcp statistics](#)” on page 180

show interface-profile dhcp-relay-profile

To display an IP DHCP Relay profile, use the following command:

```
(host)#show interface-profile dhcp-relay-profile relay1

dhcp relay profile "relay1"
-----
Parameter                                Value
-----
DHCP helper address                      172.16.30.1
Option82 Circuit-Id option               vlan interface-name
Option82 Remote-Id option                myOwnString
Giaddr as Source IP                      Disabled
```

show ip dhcp database

To display the complete IP DHCP database, use the following command:

```
(host)#show ip dhcp database
DHCP enabled
# pool-1
subnet 172.16.1.0 netmask 255.255.255.0 {
default-lease-time 43200;
max-lease-time 43200;
option domain-name "www.test.com";
option vendor-class-identifier "testStr";
option vendor-encapsulated-options "172.16.0.254";
option routers 172.16.1.254;
option user-option-43 code 43 = ip-address;
option user-option-43 172.16.1.254;
range 172.16.1.1 172.16.1.254;
authoritative;
```

show ip dhcp binding

To display the DHCP binding table, use the following command:

```
(host) #show ip dhcp binding

lease 172.16.1.251 {
  starts Fri Oct 21 08:10:29 2011
  ends Fri Oct 21 20:10:29 2011
  binding state active;
```

```

    next binding state free;
    hardware ethernet 00:25:90:0a:95:e1;
    uid "\001\000%\220\012\225\341";
}
lease 172.16.1.254 {
    starts Fri Oct 21 09:21:30 2011
    ends Fri Oct 21 21:21:30 2011
    binding state active;
    next binding state free;
    hardware ethernet 00:25:90:0a:95:d2;
    uid "\001\000%\220\012\225\322";
}
lease 172.16.1.253 {
    starts Fri Oct 21 13:09:32 2011
    ends Sat Oct 22 01:09:32 2011
    binding state active;
    next binding state free;
    hardware ethernet 00:25:90:0a:96:42;
    uid "\001\000%\220\012\226B";
}

```



The DHCP server assigns the abandoned leases only after all the free entries are exhausted.

show ip dhcp statistics

Displays the statistics in the pools stating the number of active leases, free leases etc

```
(host)#show ip dhcp statistics
```

```

Network Name 172.16.1.0/24
  Free leases      249
  Active leases    3
  Expired leases   0
  Abandoned leases 0

```

show ip dhcp pool

Displays the list of the dhcp pools configured and information about their references:

```
(host)#show ip dhcp pool
```

```

dhcp server profile List
-----
Name      References  Profile Status
----      -
pool-1    0
pool-2    0
pool-3    0
pool-4    0
Total:4

```

show ip dhcp pool

```
(host)#show ip dhcp pool <pool_name>
```

This command displays the details of the pool

```
(host)#show ip dhcp pool pool-1
```

```

dhcp server profile "pool-1"
-----
Parameter                                Value
-----
Domain name for the pool                  www.test.com
DHCP server pool                          192.168.1.0/255.255.255.0
DHCP pool lease time                     0 12 0 0
Vendor Class Identifier                   testStr
DHCP default router address              192.168.1.253
Configure DNS servers                     N/A
Configure netbios name servers            N/A
DHCP Option                              43 ip 192.168.1.254
Exclude address                          192.168.1.254
Exclude address                          192.168.1.253

```


This chapter contains the following sections:

- [“OSPF Feature Overview” on page 183](#)
- [“Configuring OSPF” on page 184](#)
- [“OSPF MD5 Authentication” on page 188](#)

OSPF Feature Overview

Open shortest path first (OSPFv2) is a dynamic interior gateway routing protocol (IGP) based on IETF RFC 2328. Aruba’s implementation of OSPFv2 allows the Mobility Access Switch to be effectively deployed in a Layer 3 topology.

This section contains the following sections:

- [“Key Features Supported by MAS” on page 183](#)
- [“LSAs Originated by MAS” on page 183](#)

Key Features Supported by MAS

- All stub area types
- Area border router (ABR)
- OSPF on VLAN and loopback interfaces
- OSPF MD5 authentication
- One OSPF instance
- Redistribute VLANs
- OSPF interface can belong to only one area

LSAs Originated by MAS

With current implementation, the following Link State Advertisement (LSA) types are generated by MAS:

- Type 1 Router LSA
- Type 2 Network LSA
- Type 3 Summary LSA
- Type 4 ASBR Summary LSA

Notes:

- Routes learned from VLAN-based access interfaces are distributed to OSPF as Router LSAs (Type 1).
- MAS can process Type 5 AS External LSA.
- MAS can process Type 7 NSSA External LSA.

Configuring OSPF

This section contains the following sections:

- “Configuring OSPF” on page 184
- “Configuring OSPF Area Types” on page 184
- “Configuring prefix-list with OSPF” on page 185
- “Verifying the Configuration” on page 185
- “Enabling OSPF on a Loopback Interface” on page 187

Configuring OSPF



The **router ospf** command must be configured to start the OSPF process.

To configure OSPF, follow these steps:

1. Enter the global OSPF configuration mode.

```
(host) (config) #router ospf
(host) (Global OSPF profile)
```

2. Assign the router identification.

```
(host) (Global OSPF profile) router-id 5.5.5.5
```

3. Assign areas.

```
(host) (Global OSPF profile) area 0.0.2.0
(host) (Global OSPF profile) area 0.0.0.1 stub
```

4. Create the interface OSPF profile “techpubs.”

```
(host) (config) #interface-profile ospf-profile techpubs
(host) (Interface OSPF profile "techpubs") #
```

5. Assign an area and cost to the profile “techpubs.”

```
(host) (Interface OSPF profile "techpubs") #area 0.0.2.0
(host) (Interface OSPF profile "techpubs") #cost 10
```

6. Attach the OSPF profile “techpubs” to a VLAN.

```
(host) (config) #interface vlan 2
(host) (vlan "2") #ospf-profile techpubs
(host) (vlan "2") #ip address 172.0.10.254 255.255.255.0
(host) (vlan "2")
```

Configuring OSPF Area Types

This release of ArubaOS Mobility Access Switch supports all Open Shortest Path First (OSPF) area types including Totally Stubby Area (TSA) and Not-So-Stubby-Area (NSSA). The following new commands are added to the Command Line Interface (CLI).

In the configuration mode, type **router ospf** to enter global OSPF profile mode.

To set an area as NSSA:

```
(host)(Global OSPF profile) #area <areaid> nssa
```

To set an area as Totally NSSA:

```
(host)(Global OSPF profile) #area <areaid> nssa no-summary
```

To set an area as TSA:

```
(host)(Global OSPF profile) #area <areaid> stub no-summary
```


To enable sending default route in NSSA:

```
(host)(Global OSPF profile) #area <areaid> nssa default-info-originate metric <cost>
metric-type <mtype>
```

To generate default Link State Advertisement (LSA) in normal area:

```
(host)(Global OSPF profile) #default-info-originate [always][metric <cost> metric-type
<mtype>]
```

For additional parameters, see *ArubaOS 7.2 Command Line Interface* guide.

Sample Configuration

```
(host)(config) #router ospf
(host)(Global OSPF profile) #area 0.0.0.1 nssa
(host)(Global OSPF profile) #area 0.0.0.2 nssa no-summary
(host)(Global OSPF profile) #area 0.0.1.0 stub no-summary
(host)(Global OSPF profile) #area 0.0.2.0 nssa default-info-originate metric 1 metric-
type 1
(host)(Global OSPF profile) #default-info-originate always
```

Configuring prefix-list with OSPF

You can filter networks received from LSA updates. The **prefix-list** command is used to configure IP prefix filtering. Prefix lists are used to either permit or deny the configured prefix based on a matching condition.



For a detailed description of the IP Prefix-list feature, see [“IP Prefix List” on page 173](#).

The **distribute-list** command filter networks received in updates. This command references to the user-defined prefix-list.

```
(host) (config) #router ospf
(host) (Global OSPF profile) #distribute-list <prefix-list name>
```

The **show router ospf** command verifies the distribute-list configuration.

```
(host) (config) #show router ospf
```

For **show router ospf** sample configuration, see [“Verifying the Configuration” on page 185](#).

Sample Configuration



This example assumes that a prefix-list called **aruba** has already been created.

```
(host) (config) #router ospf
(host) (Global OSPF profile) #distribute-list aruba
```

Verifying the Configuration

View the global OSPF profile values.

```
(host) (config) #show router ospf
```

```
Global OSPF profile "default"
-----
Parameter          Value
-----
State              Enabled
Area               0.0.0.0
```

```

Area          0.0.1.0 (stub)
Area          0.0.0.1 (nssa)
Area          0.0.0.2 (nssa)
Area          0.0.2.0 (nssa)
Area          0.0.0.4 (totally-stubby)
Router-id     10.10.10.10
Redistribute  vlan 2
Distribute-list  aruba

```

View the parameters and values for the interface OSPF profile “techpubs”.

```
(host) (vlan "2") #show interface-profile ospf-profile techpubs
```

```

Interface OSPF profile "techpubs"
-----
Parameter          Value
-----
Area               0.0.2.0
Cost               10
Dead-interval      Auto
Hello-interval     10
Retransmit-interval 5
Transmit-delay     1
Priority           1
State              Enabled

```

View the interface configuration for VLAN 2.

```
(host) (vlan "2") #show interface-config vlan 2
```

```

vlan "2"
-----
Parameter          Value
-----
Interface OSPF profile  techpubs
Interface PIM profile   N/A
Interface IGMP profile  N/A
Interface shutdown      Disabled
mtu                    1500
IP Address             172.0.10.254/255.255.255.0
IPv6 Address           N/A
IPv6 link local Address N/A
DHCP client            Disabled
DHCP relay profile      N/A
Interface description   N/A

```

Verify that the OSPF interface is running on VLAN 2.

```
(host) #show ip ospf interface vlan 2
```

```

Interface is vlan2, line protocol is up
Internet Address 172.0.10.254, Mask 255.255.255.0, Area 0.0.2.0
Router ID 5.5.5.5, Network Type BROADCAST, Cost: 10
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router id 0.0.0.0, Interface Address 0.0.0.0
Backup designated Router id 0.0.0.0, Interface Address 0.0.0.0
Timer intervals configured, Hello 10, Dead 40, Retransmit 5
Neighbor Count is 0
Tx Stat: Hellos 0 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 0
Rx Stat: Hellos 0 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 0

```

```
BadCksum 0 BadVer 0 BadNet 0 BadArea 0 BadDstAdr 0 BadAuType 0
BadAuth 0 BadNeigh 0 BadMTU 0 BadVirtLink 0
```

Verify the IP Routes

```
(host) #show ip route
```

```
Codes: C - connected, R - RIP
        O - OSPF, O(IA) - Ospf inter Area
        O(E1) - OSPF Ext Type 1, O(E2) - Ospf Ext Type 2
        M - mgmt, S - static, * - candidate default
        D - DHCP
```

```
Gateway of last resort is 10.232.10.1 to network 0.0.0.0 at cost 17
O(IA) * 0.0.0.0 /0 [17] via 10.232.10.1
O(IA) 1.0.0.99 /32 [2] via 10.232.10.1
O(IA) 1.0.0.103/32 [2] via 10.232.20.1
O(IA) 1.0.0.104/32 [3] via 10.232.10.1
O(IA) 1.0.0.105/32 [3] via 10.232.10.1
O(IA) 1.0.0.106/32 [3] via 10.232.10.1
O(IA) 1.0.0.108/32 [3] via 10.232.10.1
S 10.0.0.0 /8 [0] via 10.4.135.254
M 10.4.135.0/24 is directly connected: mgmt
M 10.4.135.91/32 is directly connected: mgmt
C 10.64.8.0/24 is directly connected: vlan66
C 10.64.8.1/32 is directly connected: vlan66
C 10.65.8.0/24 is directly connected: vlan21
C 10.65.8.1/32 is directly connected: vlan21
C 10.69.8.0/24 is directly connected: vlan61
C 10.69.8.1/32 is directly connected: vlan61
C 10.70.8.0/24 is directly connected: vlan81
C 10.70.8.1/32 is directly connected: vlan81
C 10.128.63.1/32 is directly connected: loopback0
C 10.128.64.0/24 is directly connected: vlan64
<omitted>
```

```
(host) #show ip route summary
```

Route Source	Total
-----	-----
connected	419
static	1
ospf-intra	400
ospf-inter	820
ospf-ext1	0
ospf-ext2	0
ospf-nssa	0

Enabling OSPF on a Loopback Interface

1. Create the loopback interface (3 in the example).

```
(host) (config) #interface loopback 3
(host) (loopback "3") #
```

2. Configure an IP address and Mask for the loopback.

```
(host) (loopback "3") #ip address 172.0.25.254
```

3. Attach the ospf-profile “techpubs” to the loopback interface.

```
(host) (loopback "3") #ospf-profile techpubs
```

4. Verify the loopback configuration:

```
(host) (loopback "3") #show interface loopback 3
```

```
loopback3 is administratively Up, Line protocol is Up
Hardware is Ethernet, Address is 00:0b:86:6a:f2:40
Description: Loopback
Internet address is 172.0.25.254, Netmask is 255.255.255.255
Interface index: 100663299
MTU 1514 bytes
```

Verify the interface configuration:

```
(host) (config) #show interface-config loopback 3
```

```
loopback "3"
-----
Parameter              Value
-----
Interface OSPF profile  techpubs
IP Address              172.0.25.254
Interface description   N/A
```

Verify that the OSPF is enabled on a Loopback interface:

```
(host) #show ip ospf interface loopback 3
```

```
Interface is loopback3, line protocol is up
Internet Address 172.0.25.254, Mask 255.255.255.255, Area 0.0.2.0
Router ID 5.5.5.5, Network Type LOOPBACK, Cost: 10
Transmit Delay is 1 sec, State LOOP, Priority 1
Timer intervals configured, Hello 10, Dead 40, Retransmit 5
Neighbor Count is 0
Tx Stat: Hellos 0 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 0
Rx Stat: Hellos 0 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 0
          BadCksum 0 BadVer 0 BadNet 0 BadArea 0 BadDstAdr 0 BadAuType 0
          BadAuth 0 BadNeigh 0 BadMTU 0 BadVirtLink 0
```

OSPF MD5 Authentication

This section contains the following sections:

- [“Important Points to Remember” on page 188](#)
- [“Understanding OSPF MD5 Authentication” on page 189](#)
- [“Configuring OSPF MD5 Authentication” on page 189](#)
- [“Verifying OSPF MD5 Authentication” on page 189](#)

Important Points to Remember

- This release only supports OSPF MD5 authentication on a per-interface basis.
- This release only supports one OSPF MD5 authentication key.
- This release does not support “simple” OSPF authentication.

Understanding OSPF MD5 Authentication

To protect Open Shortest Path First (OSPF) connections from spoofing attacks, the Mobility Access Switch supports MD5 authentication. MD5 is a message-digest algorithm that is specified in RFC 1321 and considered to be the most secure OSPF authentication mode.

Without MD5 authentication, a remote attacker can spoof an OSPF packet so that it appears to come from a trusted source, but can then change the routing tables of the unprotected device or exploit other vulnerabilities in the AOS OSPF network.

Note that you must configure the same MD5 key and password on both OSPF neighbors. The neighbor-ship only forms when both devices have the matching key and password.

This release only supports MD5 OSPF authentication, it does not support “simple” OSPF authentication. With simple authentication, the password traverses the network in clear-text. With MD5 OSPF authentication, the password does not traverse the network.

Configuring OSPF MD5 Authentication

To configure OSPF MD5 authentication, follow these steps:

1. Configure an OSPF profile in an interface profile:

```
(host)(config) #interface-profile ospf-profile ospf1
```

2. Configure an MD5 key and password.

```
(host)(Interface OSPF profile "ospf1") #message-digest-key 1 md5-passwd Aruba
```

3. Attach the interface OSPF profile to the vlan interface:

```
(host) (config) #interface vlan 1
(host) (vlan "1") #ospf-profile ospf1
```

Verifying OSPF MD5 Authentication

This section contains the following sections:

- [“Verifying OSPF MD5 Authentication Configuration from the Interface Profile” on page 189](#)
- [“Verifying the OSPF MD5 Authentication Configuration” on page 190](#)
- [“Verifying OSPF MD5 Authentication” on page 190](#)

Verifying OSPF MD5 Authentication Configuration from the Interface Profile

To verify the OSPF MD5 Authentication configuration from the Interface Profile, use the following show command:

```
(host)(config) #show interface-profile ospf-profile ospf1
```

```
Interface OSPF profile "ospf1"
```

```
-----
```

Parameter	Value
-----	-----
Area	0.0.0.0
Cost	1
Dead-interval	Auto
Hello-interval	10
Retransmit-interval	5
Transmit-delay	1
Priority	1
md5-key	1
md5-passwd	*****

State Enabled

Verifying the OSPF MD5 Authentication Configuration

To verify the OSPF MD5 Authentication configuration, use the following show command:

```
(host)(config) #show running-config

Building Configuration...
router ospf
  area 0.0.0.0
  interface-profile ospf-profile "ospf1"
    message-digest-key 1 md5-passwd 2aa9fdf39271f7779771543efd658fd0
  area 0.0.0.0
```

Verifying OSPF MD5 Authentication

To verify the OSPF MD5 Authentication, use the following show command:

```
(host)(config) #show ip ospf interface vlan 1

Interface is vlan1, line protocol is up
Internet Address 10.10.10.2, Mask 255.255.255.0, Area 0.0.0.0
Router ID 10.10.10.2, Network Type BROADCAST, Cost: 1
Transmit Delay is 1 sec, State DR, Priority 1
Designated Router id 10.10.10.2, Interface Address 10.10.10.2
Backup designated Router id 0.0.0.0, Interface Address 0.0.0.0
Timer intervals configured, Hello 10, Dead 40, Retransmit 5
Message digest authentication enabled key id:1
Neighbor Count is 0
Tx Stat: Hellos 19 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 19
Rx Stat: Hellos 0 DbDescr 0 LsReq 0 LsUpdate 0 LsAck 0 Pkts 0
          BadCksum 0 BadVer 0 BadNet 0 BadArea 0 BadDstAdr 0 BadAuType 0
          BadAuth 0 BadNeigh 0 BadMTU 0 BadVirtLink 0
```

The IPv6 protocol enables the next generation of large-scale IP networks by supporting addresses that are 128 bits long. This allows 2^{128} possible addresses (versus 2^{32} possible IPv4 addresses).

IPv6 addresses are represented as eight colon-separated fields of up to four hexadecimal digits each. The following are examples of IPv6 addresses:

```
FEDC:BA98:7654:3210:FEDC:BA98:7654:3210
1080:0:0:0:0:800:200C:417A
```

The use of the “::” symbol is a special syntax that you can use to compress one or more 16-bit groups of zeros or to compress leading or trailing zeros in an address. The “::” can appear only once in an address. For example, the address, 1080:0:0:0:0:800:200C:417A can also be represented as 1080::800:200C:417A.

IPv6 uses subnet identifiers to identify subnetworks to which nodes are attached. The subnet mask is a bitmask that specifies the prefix length. For example, 1080::800:200C:417A ffff:ffff:ffff:ffff:: represents all IPv6 addresses with the subnet identifier 1080:0:0:0.

IPv6 Support for Mobility Access Switch

ArubaOS provides IPv6 support on the Mobility Access Switch.



IPv6 support is currently limited to management functionality.

Following are the key points about IPv6 support on the Mobility Access Switch:

- Default IPv6 support on all RVI interfaces and Management interface.
- Auto-configured link local address on all IPv6 interfaces based on the MAC address and VLAN Id combination.
- Ability to override the auto configured link local address with another link local address.
- Ability to configure multiple global unicast addresses.
- Ability to ping other v6 hosts.
- Telnet support.
- Default gateway configuration support.

You can perform the following IPv6 operations on the Mobility Access Switch:

- [“Configure an IPv6 Interface Address” on page 192](#)
- [“Configure IPv6 Default Gateway” on page 192](#)
- [“Debug IPv6 Mobility Access Switch” on page 192](#)

You can also view the IPv6 related information on the Mobility Access Switch using the following commands:

- `show interface <intf name>`: View the IPv6 auto configured link local address and global unicast address of a VLAN interface
- `show ipv6 neighbors`: View the IPv6 neighbors
- `show ipv6 route`: View the IPv6 routes

- `show ipv6 interface brief`: View the IPv6 interfaces
- `show ipv6 interface`: View the IPv6 interface information in detail

Configure an IPv6 Interface Address

You can configure an IPv6 address for the management interface and VLAN interface of the Mobility Access Switch. The Mobility Access Switch can have multiple IPv6 addresses for each VLAN interface. You can configure IPv6 interface addresses using the following CLI commands.

To modify the auto-configured link local address of the VLAN interface:

```
(host)(config)#interface vlan <vlan#>
(host)(vlan "#")#ipv6 address link-local <X:X:X:X::X>
```

To configure global unicast address

```
(host)(config)#interface vlan <vlan#>
(host)(vlan "#")#ipv6 address <X:X:X:X::X> prefix_len <prefix_length>
```

To configure global unicast address on management interface:

```
(host)(config)#interface mgmt
(host)(mgmt)#ipv6 address <X:X:X:X::X> prefix_len <prefix_length>
```

To modify the auto-configured link local address of the management interface:

```
(host)(config)#interface mgmt
(host)(mgmt)#ipv6 address link-local <X:X:X:X::X>
```

Configure IPv6 Default Gateway

You can configure IPv6 default gateway using the following CLI command:

```
(host)(config)#ipv6-profile
(host)(ipv6-profile)#default-gateway <X:X:X:X::X>
```

Debug IPv6 Mobility Access Switch

You can now use the Ping command to debug IPv6 hosts.

To ping the global unicast address:

```
(host) #ping ipv6 <X:X:X:X::X>
```

To ping the link-local address of the host connected to the VLAN interface:

```
(host) #ping ipv6 interface vlan <interface-name> <X:X:X:X::X>
```

To ping the link-local address of the host connected to the management interface:

```
(host) #ping ipv6 interface mgmt <X:X:X:X::X>
```


This chapter contains the following major sections:

- “Important Points to Remember” on page 193
- “Understanding IGMP and PIM-SM” on page 193
- “Configuring IGMP” on page 194
- “Configuring PIM Sparse Mode” on page 194

Important Points to Remember

- PIM-SM runs on top of IGMP and needs an IGMP profile for the VLAN interface.
- IGMP must be enabled to run PIM-SM.
- IGMP is enabled by default and cannot be disabled.

Understanding IGMP and PIM-SM

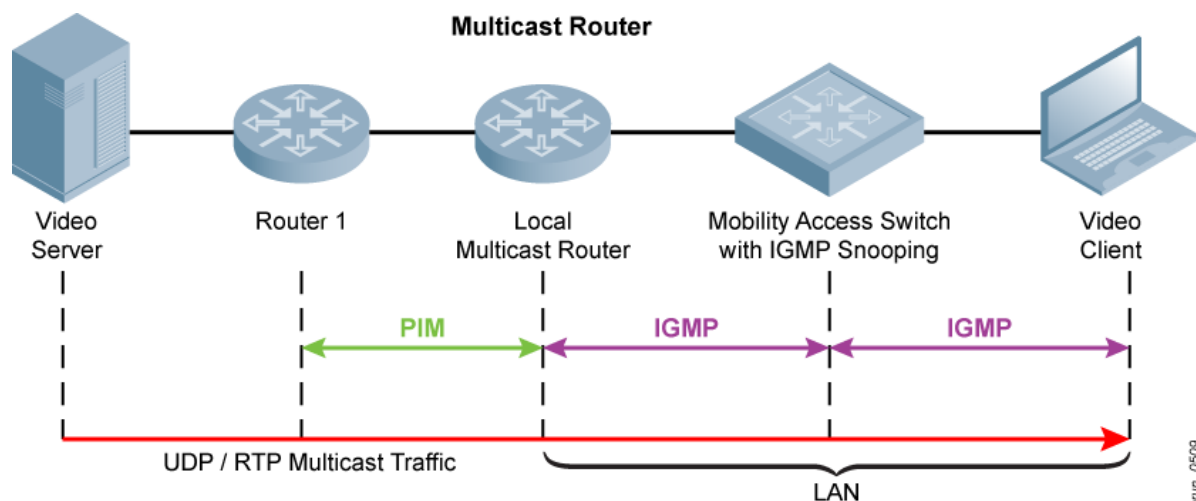
This section contains the following sections:

- “IGMP” on page 193
- “PIM” on page 194
- “PIM Sparse Mode” on page 194

IGMP

The Mobility Access Switch supports Internet Group Management Protocol (IGMP) as defined in IETF RFC 1112 (IGMPv1) and RFC 2236 (IGMPv2). IGMP allows hosts and adjacent routers on IP networks to establish multicast group memberships.

Basic IGMP Network Architecture



PIM

Protocol-Independent Multicast (PIM) is a family of multicast routing protocols for Internet Protocol (IP) networks that provide one-to-many and many-to-many distribution of data over a LAN, WAN or the Internet. It is termed protocol-independent because PIM does not include its own topology discovery mechanism, but instead uses routing information supplied by other traditional routing protocols such as the Border Gateway Protocol (BGP).

There are four variants of PIM, of which the Mobility Access Switch supports PIM Sparse Mode (PIM-SM).

PIM Sparse Mode

PIM-SM explicitly builds unidirectional shared trees rooted at a rendezvous point (RP) per group, and optionally creates shortest-path trees per source. PIM-SM generally scales fairly well for wide-area usage. PIM-SM is useful for routing multicast streams between VLANs, subnets, or local area networks (LANs) in applications such as IPTV.

Configuring IGMP

To configure an IGMP profile, follow these steps:

1. Configure an IGMP profile under an interface profile.

```
(host) (config) #interface-profile igmp-profile igmp1
(host) (Interface IGMP profile "igmp1") #
```

2. Enable IGMP profile (default is enabled).

```
(host) (Interface IGMP profile "igmp1") #no disable
```

3. Assign IGMP profile to a VLAN interface.

```
(host) (Interface IGMP profile "igmp1") #interface vlan 2
(host) (vlan "2") #igmp-profile igmp1
```

4. Verify the VLAN interface.

```
(host) (vlan "2") #show interface-config vlan 2
```

```
vlan "2"
-----
Parameter                                Value
-----
Interface OSPF profile                   ospf-a0
Interface PIM profile                    default
Interface IGMP profile                   igmp1
Interface shutdown                       Disabled
mtu                                      1500
IP Address                               20.1.1.4/255.255.255.0
IPv6 Address                             N/A
IPv6 link local Address                  N/A
DHCP client                             Disabled
DHCP relay profile                       N/A
Interface description                    N/A
```

Configuring PIM Sparse Mode

This section contains the following sections:

- [“Configuring PIM-SM End to End” on page 195](#)

- [“Verifying PIM Sparse Mode” on page 195](#)

Configuring PIM-SM End to End

To configure PIM-SM end to end, follow these steps:

1. Create a VLAN.

```
(host)(config) #vlan 7
(host)(VLAN "7") #exit
```

2. Create an interface-profile switching-profile profile to associate with a physical interface.

```
(host)(config) #interface-profile switching-profile ip-sp-profile
```

3. Add an access-vlan to set the VLAN when interface is in access mode.

```
(host)(switching profile "ip-sp-profile") #access-vlan 7
(host)(switching profile "ip-sp-profile") #exit
```

4. Associate the interface-profile switching-profile with a physical interface profile.

```
(host)(config) #interface gigabitethernet 0/0/0
(host)(gigabitethernet "0/0/0") #switching-profile ip-sp-profile
(host)(gigabitethernet "0/0/0") #exit
```

5. Create the routed VLAN interface (RVI).

```
(host)(config) #interface vlan 7
(host)(vlan "7") #
```

6. Assign an IP address to the routed VLAN interface (RVI).

```
(host)(vlan "7") #ip address 20.2.1.1 netmask 255.255.255.0
```

7. Associate the "default" PIM profile with the routed VLAN interface (RVI).

```
(host)(vlan "7") #pim-profile default
(host)(vlan "7") #exit
```

8. Use the router pim command to enter Global PIM profile mode and define the RP address and group range.

```
(host)(config) #router pim
(host)(Global PIM profile) #rp-address 224.0.0.1 group-range 225.0.0.0 255.0.0.0
```



When configuring static RP, please ensure the RP is active and reachable. If the RP is not reachable, multicast traffic fails.

Verifying PIM Sparse Mode

This section contains the following sections:

- [“Displaying PIM RPF Information” on page 195](#)
- [“Displaying PIM Neighbor Information” on page 196](#)
- [“Displaying PIM RP Information” on page 196](#)
- [“Displaying PIM Mroute Information” on page 196](#)
- [“Displaying PIM Statistical Information” on page 196](#)

Displaying PIM RPF Information

```
(co5) #show ip pim rpf 10.10.10.10
```

```
PIM RPF Information
-----
Address      Nexthop      RPF Interface
-----
10.10.10.10  20.20.1.9   vlan20
```

Displaying PIM Neighbor Information

To display PIM neighbor information, use the following command:

```
(host)# show ip pim neighbor
PIM Neighbor Information
-----
Interface  Neighbor IP  Expiry    UpTime
-----
vlan11     11.11.22.22  08:00:20  07:58:51
```

Displaying PIM RP Information

To display PIM RP information, use the following command:

```
(host)# show ip pim rp
PIM RP-Group Mapping
-----
Group/Prefix  RP
-----
224.0.0.0/4   11.11.22.22
```

Displaying PIM Mroute Information

To display PIM Mroute information, use the following command:

```
(host)# show ip pim mroute
IP Multicast Route Table
Flags:  D - Dense, S - Sparse, C - Connected, L - Local,
        P - Pruned, R - RP-bit set, T - SPT bit set, F - Register Flag
        J - Join SPT, A - Assert Winner
(*,224.1.1.6) 14:20:11 RP 11.11.22.22 flags:
    Incoming Interface: vlan11 RPF nbr: 11.11.22.22
    Outgoing Interface List:
        vlan22, 14:20:11
(22.22.99.99,230.1.1.1) 14:17:20 RP 11.11.22.22 flags: T
    Incoming Interface: vlan22 RPF nbr: 22.22.99.99
    Outgoing Interface List:
        vlan11, 14:17:20
```

Displaying PIM Statistical Information

To display PIM statistical information, use the following command:

```
(host)# show ip pim stats
PIM Statistics
-----
Interface  Counter      Value
-----
vlan11 Rx  Hellos      0056
           Rx Join/Prune  0000
           Rx Join      0000
           Rx Prune     0000
           Rx Register-Stop 0000
           Tx Hellos    0057
           Tx Join/Prune 0016
           Tx Join      0000
```

Tx Prunes	0000
Tx Register	0000
Invalid Hellos	0000
Invalid Join/Prune	0000
Invalid Join	0000
Invalid Prune	0000
Invalid Register	0000
Invalid Register-Stop	0000

You can enable multicast support on the Mobility Access Switch with IGMP snooping. You can enable the Mobility Access Switch to listen in on the IGMP conversation between hosts and network devices, and create a mapping table of which links need which IP multicast streams and which multicasts can be filtered from the links which do not need them.

This chapter includes the following topics:

- “Important Points to Remember” on page 199
- “Multicast Support with IGMP Snooping” on page 199
- “Mrouter” on page 200
- “Creating and Applying an IGMP Snooping Profile to a VLAN” on page 201
- “Sample Configuration” on page 201
- “IGMP Snooping Factory Initial and the Default Profiles” on page 202
- “Verifying IGMP Snooping Configuration” on page 202
- “Monitoring IGMP Snooping” on page 203

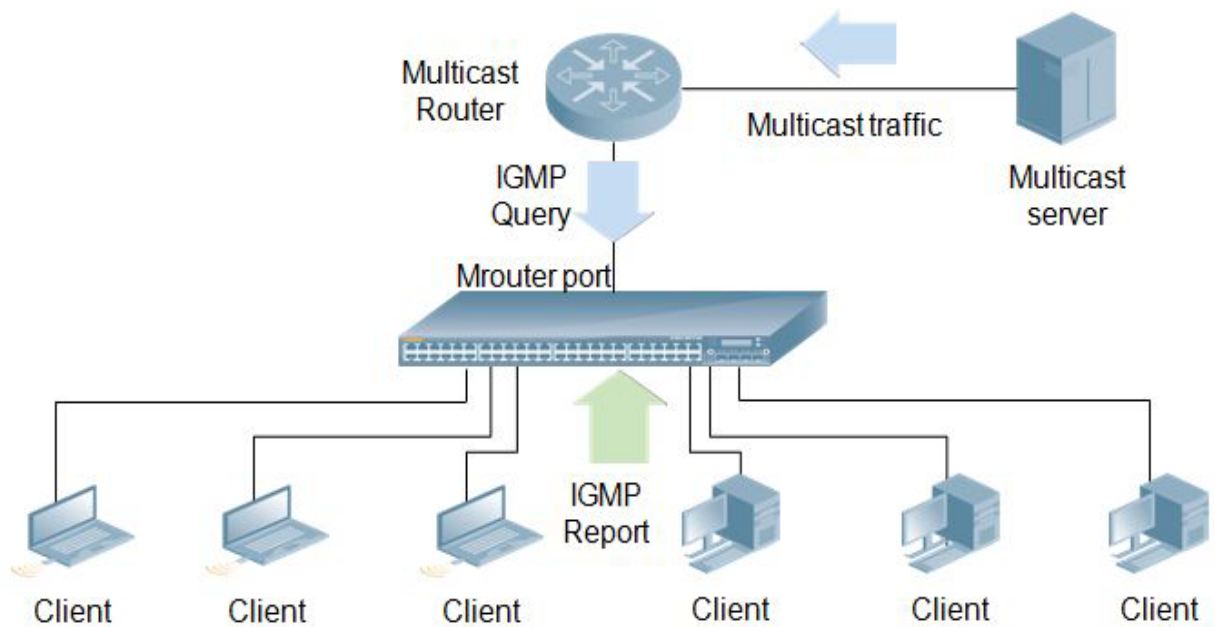
Important Points to Remember

- IGMP snooping is enabled by default.
- IGMP snooping is enabled on per-VLAN basis.
- IGMP snooping profile must be referenced in the VLAN and not on the interface.
- IGMP versions 1 and 2 are supported for snooping.

Multicast Support with IGMP Snooping

The Mobility Access Switch supports IGMP snooping, which prevents multicast flooding on Layer 2 network treating multicast traffic as broadcast traffic. All streams could be flooded to all ports on that VLAN. When multicast flooding occurs, end-hosts that happen to be in the same VLAN would be receiving all the streams only to be discarded without snooping.

When you enable IGMP snooping, the switch becomes IGMP-aware and processes the IGMP control messages as received. You must do this to correctly process all IGMP membership reports and IGMP leave messages. IGMP snooping is handled by the hardware for performance. Multicast routers and multicast receivers associated with each IP multicast group are learnt dynamically.



Snooping Report and Query Support

The Mobility Access Switch relays IGMP report from all receiver per group to the multicast router. In IGMP snooping proxy mode, reports to multicast router ports are suppressed. Query from multicast router is relayed to all ports in the VLAN. When snooping proxy is enabled, the switch queries hosts for interested receivers and it floods the query message received from a multicast router. When IGMP query message is seen, it becomes a mrouter port in IGMP snooping table. This port is used for forwarding multicast frames that are sourced from a VLAN to a multicast router for further processing.

Mrouter

VLANs in a Layer 2 switch needs to know the path to the PIM router that connects Layer 2 domain to the Layer 3 Network. When the multicast source is present on the Layer 2 switch, the traffic that originates from the Layer 2 switches need to know a port through which multicast traffic can reach the Layer 3 PIM router. For this reason, the VLAN in the Layer 2 switch on which IGMP snooping is enabled will designate a port as Mrouter port. The mrouter port can be detected dynamically or statically. The dynamic detection is based on IGMP query message or PIM hello messages. You can also configure static mrouter ports.

When multicast receivers are present on the VLAN in a Layer 2 switch, the IGMP report message from the host is forwarded out of the mrouter port towards the PIM router to let the PIM router know that there are receivers interested in receiving multicast traffic, so that, PIM routers can add the VLAN interface to the outgoing list in the multicast route on a multicast router.

Configuring a Static Mrouter Port

To configure a static mrouter port, follow these steps:

```
(host)(config)# interface gigabitethernet <slot/module/port>
  igmp-snooping mrouter-vlan <vlan-id|vlan-list>
  igmp-snooping mrouter-vlan {add | delete} <vlan-id>
```


Example Configuration

```
(host)(config)# interface gigabitethernet 0/0/9
  igmp-snooping mrouter-vlan 1
```

```
(host)# show igmp-snooping mrouter vlan 1
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

IGMP Snooping Multicast Router Ports

VLAN	Elected-Querier	Ports (Flags)	Expiry	UpTime	Src-IP
0001	10.10.10.6	GE0/0/9 (DM)	00:03:25	04:35:30	10.10.10.6
		GE0/0/9 (DP)	00:04:14	04:35:09	10.10.10.6

```
(host)# show igmp-snooping mrouter vlan 1 detail
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

```
Vlan:0001 Elected-Querier:10.10.10.6
  GE0/0/9   (DM) Expiry Time: 00:03:45  Uptime: 04:36:10
              Router IP: 10.10.10.6
              Router MAC: 00:19:06:55:15:40
  GE0/0/9   (DP) Expiry Time: 00:04:04  Uptime: 04:35:49
              Router IP: 10.10.10.6
              Router MAC: 00:19:06:55:15:40
```

Creating and Applying an IGMP Snooping Profile to a VLAN

Using the CLI

```
(host)(config)# vlan-profile igmp-snooping-profile <profile-name>
  clone <source>
  fast-leave
  last-member-query-count <1-5>
  last-member-query-interval <1-25 seconds>
  no {...}
  query-interval <1-18000 seconds>
  query-response-interval <1-25 seconds>
  robustness-variable <1-7>
  snooping
  snooping-proxy
  startup-query-count <1-10>
  startup-query-interval <1-18000 seconds>
(host)(config)# vlan <vlan-id>
  vlan-profile igmp-snooping-profile <profile-name>
```

Sample Configuration

```
(host)(config)# vlan-profile igmp-snooping-profile IGMP_SNOOP
  fast-leave
  last-member-query-count 2
  last-member-query-interval 15
  query-interval 6000
  query-response-interval 5
  robustness-variable 2
  snooping
```

```
snooping-proxy
startup-query-count 5
startup-query-interval 6000
(host)(config)# vlan 200
vlan-profile igmp-snooping-profile IGMP_SNOOP
```

IGMP Snooping Factory Initial and the Default Profiles

```
(host)# show vlan-profile igmp-snooping-profile igmp-snooping-factory-initial
igmp-snooping-profile "igmp-snooping-factory-initial"
```

Parameter	Value
-----	-----
Enable igmp snooping	Enabled
Enable igmp snooping proxy	Disabled
Enable fast leave	Disabled
startup-query-count	2
startup-query-interval(secs)	31
query-interval(secs)	125
query-response-interval(secs)	10
last-member-query-count	2
last-member-query-interval(secs)	1
robustness-variable	2

```
(host)# show vlan-profile igmp-snooping-profile default
igmp-snooping-profile "default"
```

Parameter	Value
-----	-----
Enable igmp snooping	Enabled
Enable igmp snooping proxy	Disabled
Enable fast leave	Disabled
startup-query-count	2
startup-query-interval	31
query-interval	125
query-response-interval	10
last-member-query-count	2
last-member-query-interval	1
robustness-variable	2

Verifying IGMP Snooping Configuration

```
(host)# show vlan-profile igmp-snooping-profile IGMP_SNOOP
igmp-snooping-profile "IGMP_SNOOP"
```

Parameter	Value
-----	-----
Enable igmp snooping	Enabled
Enable igmp snooping proxy	Disabled
Enable fast leave	Disabled
startup-query-count	2
startup-query-interval	31
query-interval	125
query-response-interval	10
<output truncated>	

Monitoring IGMP Snooping

```
(host)# show igmp-snooping counters vlan 2
```

IGMP Snooping Multicast Counters

Name	Value
received-total	0000
received-queries	0000
received-v1-reports	0000
received-v2-reports	0000
received-v3-reports	0000
received-pimv1-hello	0000
received-pimv2-hello	0000
received-leaves	0000
received-unknown-types	0000
len-errors	0000
checksum-errors	0000
transmitted-queries	0000
transmitted-joins	0000
transmitted-leaves	0000
transmitted-errors	0000
forwarded-queries	0000
forwarded-joins	0000
forwarded-leaves	0000

```
(host)# show igmp-snooping groups
```

IGMP Snooping Multicast Route Table

VLAN	Group	Port List
0100	224.0.1.40	GE 0/0/11
0100	239.255.255.250	GE 0/0/11

```
(host)# show igmp-snooping membership
```

IGMP Snooping Multicast Membership

VLAN	Group	Port	Expiry	UpTime
0001	224.0.1.40	GE0/0/9	00:03:36	04:47:27
0001	225.0.1.1	GE0/0/9	00:00:00	00:01:25
1900	225.0.1.1	GE0/0/3	00:03:49	04:47:32
0003	225.0.1.1	GE0/0/9	00:00:00	04:46:30
0003	239.0.0.1	GE0/0/9	00:00:00	04:44:42

```
(host)# show igmp-snooping mrouter
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

IGMP Snooping Multicast Router Ports

VLAN	Elected-Querier	Ports (Flags)	Expiry	UpTime	Src-Ip
0001	10.10.10.6	GE0/0/9 (DM)	00:04:07	04:45:55	10.10.10.6
		GE0/0/9 (DP)	00:04:09	04:45:34	10.10.10.6
0003	3.3.3.10	GE0/0/9 (DM)	00:04:15	04:45:25	3.3.3.10

		GE0/0/9 (DP)	00:04:06	04:44:56	3.3.3.10
0300	20.20.20.1	GE0/0/9 (DM)	00:04:15	04:45:25	20.20.20.1
		GE0/0/9 (DP)	00:04:05	04:45:13	20.20.20.1

You can also use the following commands:

```
(host)# show igmp-snooping counters vlan <vlan-id>
(host)# show igmp-snooping groups vlan <vlan-id>
(host)# show igmp-snooping membership vlan <vlan-id> | detail
(host)# show igmp-snooping mrouter vlan <vlan-id> | detail
```

Clearing IGMP Counters and Membership

```
(host)(config)# clear igmp-snooping counters
(host)(config)# clear igmp-snooping counters vlan <vlan-id>
(host)(config)# clear igmp-snooping membership
(host)(config)# clear igmp-snooping membership vlan <vlan-id>
(host)(config)# clear igmp-snooping mrouter
(host)(config)# clear igmp-snooping mrouter vlan <vlan-id>
```

Enabling IGMP Snooping Trace Options

```
(host)(config)# traceoptions
    igmp-snooping flags {all|config|errors|receive|transmit}
```

This chapter contains the following major sections:

- [“Important Points to Remember” on page 205](#)
- [“Understanding MLD Snooping” on page 205](#)
- [“Configuring MLD Snooping” on page 205](#)
- [“Verifying MLD Snooping” on page 206](#)

Important Points to Remember

- This release supports MLDv1 (RFC 2710), so MLDv2 specific packets are not processed.
- MLD snooping prevents multicast flooding on an Ethernet link, but it requires complex processing for each of the interfaces on switches that were not initially designed for this kind of task.
- MLD is embedded in ICMPv6, unlike IGMP, which uses a separate protocol. MLDv1 is similar to IGMPv2 and MLDv2 is similar to IGMPv3.

Understanding MLD Snooping

Multicast Listener Discovery (MLD) is a component of the Internet Protocol Version 6 (IPv6) suite. It is used by IPv6 routers for discovering multicast listeners on a directly attached link. When multicast is supported at the IPv6 level, it often broadcasts at lower levels. So, for example, an Ethernet switch broadcasts multicast traffic on all ports, even if only one host wants to receive it.

To prevent entire Ethernet segments from being flooded, MLD snooping can be implemented on Ethernet switches. The MLD snooping solution is similar to the IGMP snooping solution for IPv4. When MLD snooping is implemented on a switch, it detects all MLD version 1 messages that are exchanged on the link. It also maintains a table that indicates which IPv6 multicast groups should be forwarded for each of the interfaces.

Configuring MLD Snooping

This section contains the following sections:

- [“Configuring MLD Snooping” on page 205](#)
- [“Deleting an Mrouter Port on a VLAN” on page 206](#)

Configuring MLD Snooping

To configure MLD snooping, follow these steps:

1. Configure an MLD snooping profile in a VLAN profile.

```
(host) (config) #vlan-profile mld-snooping-profile MLD_Doc
(host) (mld-snooping-profile "MLD_Doc") #snooping
(host) (mld-snooping-profile "MLD_Doc") #
```

2. Apply the MLD snooping profile to the VLAN.

```
(host) (config) #vlan 10
```

```
(host) (VLAN "10") #mld-snooping-profile MLD_Doc
(host) (VLAN "10") #
```

3. Configure a static mrouter port.

```
(host) (config) #interface gigabitethernet 0/0/46
(host) (gigabitethernet "0/0/46") #mld-snooping mrouter-vlan 10
```

Deleting an Mrouter Port on a VLAN

To delete an mrouter port on a VLAN, use the following command:

```
(host) (gigabitethernet "0/0/4") #mld-snooping mrouter-vlan delete 2
```

Verifying MLD Snooping

This section contains the following sections:

- “Verifying the MLD Snooping Profile” on page 206
- “Verifying the Static and Dynamic Mrouter Port for MLD Snooping” on page 206
- “Verifying the MLD Snooping Mrouter Detail” on page 207
- “Verifying MLD Snooping Member Ports” on page 208
- “Verifying the MLD Group” on page 209
- “Verifying the MLD Snooping Group Count” on page 209
- “Verifying the MLD Snooping Statistics” on page 209

Verifying the MLD Snooping Profile

To verify an MLD snooping profile, use the following command:

```
(host) #show vlan-profile mld-snooping-profile MLD_Doc
```

```
mld-snooping-profile "MLD_Doc"
-----
Parameter                                Value
-----
robustness-variable                       2
last-member-query-interval(secs)         1
query-interval(secs)                     125
query-response-interval(secs)            10
Enable fast leave                         Disabled
Enable mld snooping                       Enabled
```

Verifying the Static and Dynamic Mrouter Port for MLD Snooping

To verify the static and dynamic mrouter port for MLD snooping, use the following command:

```
(host) #show mld-snooping mrouter vlan 1
```

Flags: D - Dnyamic, S - Static, P - PIM, M - IGMP/MLD

MLD Snooping Multicast Router Ports

VLAN	Elected-Querier	Ports (Flags)	Expiry	UpTime
0001	3555:5555:6666:6666:7777:7777:8888:8888	GE0/0/0 (S)	00:00:00	00:10:35
		GE0/0/3 (DM)	00:04:20	00:10:33
		GE0/0/3 (DP)	00:04:19	00:10:33

Verifying the MLD Snooping Mrouter Detail

To verify the mld-snooping mrouter detail and show identifiers for each field, use the following command:

```
(host) (VLAN "1") #show mld-snooping mrouter detail

Flags: D - Dnyamic, S - Static, P - PIM, M - IGMP/MLD

Vlan:0001 Elected-Querier:3555:5555:6666:6666:7777:7777:8888:8888
  GE0/0/0    (S)  Expiry Time: 00:00:00  Uptime: 00:03:54
              Router IP: N/A
              Router MAC: 00:00:00:00:00:00
  GE0/0/3    (DM) Expiry Time: 00:01:32  Uptime: 00:03:52
              Router IP: 3555:5555:6666:6666:7777:7777:8888:8888
              Router MAC: 00:00:00:00:02:00
  GE0/0/3    (DP) Expiry Time: 00:01:31  Uptime: 00:03:52
              Router IP: fe80::200:24ff:fef9:7ccd
              Router MAC: 00:00:24:f9:7c:cd
(host) (VLAN "1") #show igmp-snooping mrouter detail

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD

Vlan:0001 Elected-Querier:111.1.0.12
  GE0/0/0    (DM) Expiry Time: 00:04:12  Uptime: 00:00:08
              Router IP: 111.1.0.12
              Router MAC: 00:00:33:00:05:00
Vlan:0004 Elected-Querier:11.11.11.3
  GE0/0/4    (S)  Expiry Time: 00:00:00  Uptime: 00:19:54
              Router IP: N/A
              Router MAC: 00:00:00:00:00:00
  GE0/0/4    (DM) Expiry Time: 00:04:09  Uptime: 00:00:11
              Router IP: 11.11.11.3
              Router MAC: 00:00:09:0b:91:6d
```

Verifying the Two Mrouter Entries with the Same IP Address

Two mrouter entries with the same router IP address can be created if the PIM router is also the IGMP querier based on both protocol packets. To distinguish between the two IP addresses, flags are displayed in the commands **show igmp-snooping mrouter** and **show mld-snooping mrouter**.

show igmp-snooping mrouter

```
(host) (VLAN "1") #show igmp-snooping mrouter

Flags: D - Dnyamic, S - Static, P - PIM, M - IGMP/MLD

IGMP Snooping Multicast Router Ports
-----
VLAN  Elected-Querier  Ports (Flags)  Expiry    UpTime    Src-Ip
-----
0004  11.11.11.3           GE0/0/4 (S)    00:00:00   00:26:26   -
                                GE0/0/4 (DM)   00:03:52   00:06:43   11.11.11.3
                                GE0/0/4 (DP)   00:04:19   00:00:02   11.11.11.3
                                GE0/0/3 (DM)   00:03:52   00:06:43   11.11.11.11
```

show mld-snooping mrouter

If the 80 column limit is exceeded when displaying the **src-ip** and the elected querier in the same row of the **show mld-snooping mrouter** output, the **src-ip** is not shown. To find the **src-ip**, use the **show mld-snooping mrouter detail** command:

```
(host) (VLAN "1") #show mld-snooping mrouter

Flags: D - Dnyamic, S - Static, P - PIM, M - IGMP/MLD

MLD Snooping Multicast Router Ports
-----
VLAN  Elected-Querier                Ports (Flags)  Expiry  UpTime
----  -
0001  3555:5555:6666:6666:7777:7777:8888:8888  GE0/0/0 (S)   00:00:00  00:10:35
                                           GE0/0/3 (DM)   00:04:20  00:10:33
                                           GE0/0/3 (DP)   00:04:19  00:10:33
```

Similar to the output of "show mld-snooping mrouter detail", the output the show mld-snooping membership detail now includes labels for each field to enhance readability.

```
(host) (VLAN "1") #show igmp-snooping membership detail

Flags: H - IGMP/MLD listener, M - Multicast Router

Group:225.0.0.9 Vlan:0001
  Port: GE0/0/2    Expiry: 00:00:00 Uptime: 00:01:21
        (M) IP: 0.0.0.0          MAC: 00:0b:86:6a:20:80
  Port: GE0/0/4    Expiry: 00:02:59 Uptime: 00:01:21
        (H) IP: 11.11.11.1       MAC: 00:00:09:0b:91:6c
Group:225.0.0.10 Vlan:0001
  Port: GE0/0/2    Expiry: 00:00:00 Uptime: 00:01:21
        (M) IP: 0.0.0.0          MAC: 00:0b:86:6a:20:80
  Port: GE0/0/4    Expiry: 00:02:59 Uptime: 00:01:21
        (H) IP: 11.11.11.1       MAC: 00:00:09:0b:91:6c

(host) #show mld-snooping membership detail

Flags: H - IGMP/MLD listener, M - Multicast Router

Group:ff03::3 Vlan:0001
  Port: GE0/0/0    Expiry: 00:04:08 Uptime: 00:00:12
        (H) IP: fe80::5001       MAC: 00:00:02:00:05:00
  Port: GE0/0/4    Expiry: 00:00:00 Uptime: 00:00:12
        (M) IP: fe80::5002       MAC: 00:00:00:00:03:00
```

Verifying MLD Snooping Member Ports

To verify the MLD snooping member ports, use the following command:

```
(host) #show mld-snooping membership vlan 10

MLD Snooping Multicast Membership
-----
VLAN  Group  Port      Expiry  UpTime
----  -
0010  ff03::1  GE0/0/22  00:04:11  00:00:15
                                GE0/0/47  00:00:00  00:00:15
0010  ff03::2  GE0/0/22  00:04:11  00:00:15
                                GE0/0/47  00:00:00  00:00:15
```



```

0010 ff03::3 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::4 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::5 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::6 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::7 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::8 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::9 GE0/0/22 00:04:11 00:00:15
           GE0/0/47 00:00:00 00:00:15
0010 ff03::a GE0/0/22 00:04:11 00:00:15

```

MLD Snooping Multicast Membership

```

-----
VLAN  Group  Port      Expiry    UpTime
----  -
           GE0/0/47  00:00:00  00:00:15

```

Verifying the MLD Group

To verify the MLD group, use the following command:

```
(host) # show mld-snooping groups vlan 10
```

MLD Snooping Multicast Route Table

```

-----
VLAN  Group      Port List
----  -
0010  ff03::1  GE0/0/47 GE0/0/22
0010  ff03::2  GE0/0/47 GE0/0/22
0010  ff03::3  GE0/0/47 GE0/0/22
0010  ff03::4  GE0/0/47 GE0/0/22
0010  ff03::5  GE0/0/47 GE0/0/22
0010  ff03::6  GE0/0/47 GE0/0/22
0010  ff03::7  GE0/0/47 GE0/0/22
0010  ff03::8  GE0/0/47 GE0/0/22
0010  ff03::9  GE0/0/47 GE0/0/22
0010  ff03::a  GE0/0/47 GE0/0/22

```

Verifying the MLD Snooping Group Count

To verify the MLD snooping group count, use the following command:

```
(host) # show mld-snooping groups vlan 10 count
```

MLD Snooping Multicast Route Count

```

-----
VLAN  Count
----  -
0010  0010

```

Verifying the MLD Snooping Statistics

To verify the MLD snooping statistics, use the following command:

```
(host) #show mld-snooping counters vlan 10
```

MLD Snooping Counters

```

-----
Name                               Value
-----
received-total                     1110
received-queries                   0036
received-vl-reports                1074
received-leaves                    0000
received-unknown-types             0000
len-errors                         0000
checksum-errors                    0000
forwarded                          0930

```

List of MLD Snooping Commands and Sample Outputs

This section contains the following commands:

- [“Show Counters” on page 210](#)
- [“Show VLAN Counters” on page 211](#)
- [“Show MLD Mrouter Ports” on page 211](#)
- [“Show MLD Mrouter Ports Detail” on page 211](#)
- [“Show MLD Router Ports Per VLAN” on page 211](#)
- [“Show Detected MLD Multicast Addresses” on page 211](#)
- [“Show Detected MLD Multicast Addresses Per VLAN” on page 212](#)
- [“Show Detected MLD Multicast Membership Information” on page 212](#)
- [“Show Detected MLD Multicast Membership Information \(Detailed Version\)” on page 212](#)
- [“Show Detected MLD Multicast Membership Information Per VLAN” on page 213](#)
- [“Show MLD-Snooping Profile” on page 213](#)
- [“Show List of MLD-Snooping Profiles” on page 213](#)
- [“Show List of References for MLD-Snooping Profile” on page 213](#)

Show Counters

```

(host) #show mld-snooping counters
MLD Snooping Counters

```

```

-----
Name                               Value
-----
received-total                     0005
received-queries                   0001
received-vl-reports                0004
received-leaves                    0000
received-pim-v6                    0000
received-unknown-types             0000
len-errors                         0000
checksum-errors                    0000
forwarded                          0000

```

Show VLAN Counters

```
(host) #show mld-snooping counters vlan 1
```

MLD Snooping Counters

Name	Value
-----	-----
received-total	0005
received-queries	0001
received-vl-reports	0004
received-leaves	0000
received-pim-v6	0000
received-unknown-types	0000
len-errors	0000
checksum-errors	0000
forwarded	0000

Show MLD Mrouter Ports

```
(host) #show mld-snooping mrouter
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

MLD Snooping Multicast Router Ports

VLAN	Elected-Querier	Ports (Flags)	Expiry	UpTime
----	-----	-----	-----	-----
0001	fe11::d0d0	GE0/0/4 (DM)	00:04:12	00:00:08

Show MLD Mrouter Ports Detail

```
(host) #show mld-snooping mrouter detail
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

```
Vlan:0001 Elected-Querier:fe11::d0d0
      GE0/0/4      (DM) Expiry Time: 00:04:06 Uptime: 00:00:14
      Router IP: fe11::d0d0
      Router MAC: 00:00:00:00:03:00
```

Show MLD Router Ports Per VLAN

```
(host) #show mld-snooping mrouter vlan 1
```

Flags: D - Dynamic, S - Static, P - PIM, M - IGMP/MLD query

MLD Snooping Multicast Router Ports

VLAN	Elected-Querier	Ports (Flags)	Expiry	UpTime
----	-----	-----	-----	-----
0001	fe11::d0d0	GE0/0/4 (DM)	00:04:11	00:00:09

Show Detected MLD Multicast Addresses

```
(host) #show mld-snooping groups
```

MLD Snooping Multicast Route Table

VLAN	Group	Port List
0001	ff03::1	GE0/0/0 GE0/0/4
0001	ff03::2	GE0/0/0 GE0/0/4
0001	ff03::3	GE0/0/0 GE0/0/4
0001	ff03::4	GE0/0/0 GE0/0/4

Show Detected MLD Multicast Addresses Per VLAN

```
(host) #show mld-snooping groups vlan 1
```

MLD Snooping Multicast Route Table

VLAN	Group	Port List
0001	ff03::1	GE0/0/0 GE0/0/4
0001	ff03::2	GE0/0/0 GE0/0/4
0001	ff03::3	GE0/0/0 GE0/0/4
0001	ff03::4	GE0/0/0 GE0/0/4
0001	ff03::5	GE0/0/0 GE0/0/4

Show Detected MLD Multicast Membership Information

```
(host) #show mld-snooping membership
```

MLD Snooping Multicast Membership

VLAN	Group	Port	Expiry	UpTime
0001	ff03::1	GE0/0/0	00:02:12	00:02:08
0001	ff03::2	GE0/0/0	00:02:13	00:02:07
0001	ff03::3	GE0/0/0	00:02:14	00:02:06
0001	ff03::4	GE0/0/0	00:02:15	00:02:05
0001	ff03::5	GE0/0/0	00:02:16	00:02:04

Show Detected MLD Multicast Membership Information (Detailed Version)

```
(host) #show mld-snooping membership detail
```

Flags: H - IGMP/MLD listener, M - Multicast Router

```
Group:ff03::1 Vlan:0001
  Port: GE0/0/0    Expiry: 00:00:30 Uptime: 00:03:50
    (H) IP: fe80::200:24ff:fef9:7ccf MAC: 00:00:24:f9:7c:cf
Group:ff03::2 Vlan:0001
  Port: GE0/0/0    Expiry: 00:00:31 Uptime: 00:03:49
    (H) IP: fe80::200:24ff:fef9:7ccf MAC: 00:00:24:f9:7c:cf
Group:ff03::3 Vlan:0001
  Port: GE0/0/0    Expiry: 00:00:32 Uptime: 00:03:48
    (H) IP: fe80::200:24ff:fef9:7ccf MAC: 00:00:24:f9:7c:cf
Group:ff03::4 Vlan:0001
  Port: GE0/0/0    Expiry: 00:00:33 Uptime: 00:03:47
    (H) IP: fe80::200:24ff:fef9:7ccf MAC: 00:00:24:f9:7c:cf
Group:ff03::5 Vlan:0001
  Port: GE0/0/0    Expiry: 00:00:34 Uptime: 00:03:46
    (H) IP: fe80::200:24ff:fef9:7ccf MAC: 00:00:24:f9:7c:cf
```

Show Detected MLD Multicast Membership Information Per VLAN

```
(host) #show mld-snooping membership vlan 1
```

MLD Snooping Multicast Membership

VLAN	Group	Port	Expiry	UpTime
0001	ff03::1	GE0/0/0	00:02:12	00:02:08
0001	ff03::2	GE0/0/0	00:02:13	00:02:07
0001	ff03::3	GE0/0/0	00:02:14	00:02:06
0001	ff03::4	GE0/0/0	00:02:15	00:02:05
0001	ff03::5	GE0/0/0	00:02:16	00:02:04

Show MLD-Snooping Profile

```
(host) #show VLAN-profile mld-snooping-profile default
```

mld-snooping-profile "default"

Parameter	Value
robustness-variable	2
last-member-query-interval(secs)	10
query-interval(secs)	125
query-response-interval(secs)	10
Enable fast leave	Enabled
Enable mld snooping	Enabled

Show List of MLD-Snooping Profiles

```
(host) #show VLAN-profile mld-snooping-profile
```

mld-snooping-profile List

Name	References	Profile Status
default	2	
Total:1		

Show List of References for MLD-Snooping Profile

```
(host) #show references vlan-profile mld-snooping-profile default
```

References to mld-snooping-profile "default"

Referrer	Count
vlan "1" mld-snooping-profile	1
vlan "1111" mld-snooping-profile	1
Total References:2	

This chapter describes the following topics:

- “Port Security Overview” on page 215
- “Enabling Port Security Functionality” on page 216
- “Sample Configurations” on page 218

Port Security Overview

This release of ArubaOS Mobility Access Switch supports Port Security functionality which provides network security at Layer 2. You can now filter the unauthorized devices to send the control packets, restrict the number of MACs allowed on the interface, and detect unwanted loops in the network when not running spanning-tree protocol.

You can enable or disable this functionality at an interface level.

Router Advertisement Guard

The Router Advertisement (RA) Guard functionality analyzes the RAs and filters out RA packets sent by unauthorized devices. The RA guard feature is disabled by default. By enabling, the RA packets received on the interface are dropped and the port can be shutdown based on the interface configuration. The port can be re-activated after the configured time by configuring the `auto-recovery` option.

Points to remember

- The following RA messages are filtered by enabling the RA guard:
 - RA message with no extension header
 - RA message with multiple extension headers
 - RA message fragmented
- The following Unicast RA messages are not filtered by enabling the RA guard:
 - Unicast RA messages with multiple extension headers.
 - Unicast RA messages fragmented

DHCP Trust

The DHCP trust functionality provides support to filter the IPv4 DHCP packets from the unauthorized devices. The following IPv4 DHCP messages are filtered on an interface configured not to trust DHCP.

- DHCP offer messages
- DHCP Ack messages

You can enable the DHCP trust on any interface. By default the DHCP packets are trusted on the interface. When the DHCP Trust is disabled, the aforementioned DHCP messages that are received on the interface are dropped.

Loop Protect

The Loop Protect functionality detects the unwanted physical loops in your network. You can enable or disable this functionality at an interface level. A proprietary protocol data unit (PDU) is used to detect the physical loops in the network. When the system detects a loop, it disables the port that sends the PDU. You can re-enable the port automatically or manually.

Points to Remember

- It is recommended that you enable Loop Protect on all the Layer 2 interfaces when the spanning tree is disabled on the switch.
- The Loop Protect functionality will not detect any loops when MSTP or PVST (on any VLAN) is enabled on the switch.
- The Loop Protect functionality will work only on non-HSL interfaces. An error will be displayed when you try to enable this functionality on HSL interfaces.

MAC Limit

The MAC limit feature restricts the maximum number of MACs that can be learnt on the interface. When the MAC limit is enabled, it provides support to log the excess MACs or drop the new MAC learning requests or shuts down the port.

Enabling Port Security Functionality

The port security functionality will be configured as part of the port level security configuration. This profile can be attached to the interface.

Enabling RA Guard Functionality

RA Guard functionality can be enabled at the port level. You can configure the RA guard as part of the port level security configuration and can be attached to the interface.

```
(host)(config)# interface-profile port-security-profile <profile-name>
    ipv6-ra-guard action {drop|shutdown}auto-recovery-time <recovery-time>
```

The following example shows how to enable the RA Guard functionality:

```
(host)(config)# interface-profile port-security-profile RA-Guard1
    ipv6-ra-guard action shutdown auto-recovery-time 60
```

Enabling DHCP Trust Functionality

The DHCP trust functionality will be configured as part of the port level security configuration. This profile can be attached to the interface.

DHCP Trust can be enabled on any interface. By default, the DHCP Trust setting in a port-security-profile is to filter (block) these OFFER and ACK messages. You must explicitly enable DHCP Trust (trust dhcp) in the port-security-profile (if applied to a port) to allow these DHCP messages from valid devices.

```
(host)(config)# interface-profile port-security-profile <profile-name>
    trust dhcp
```

When **no trust dhcp** is configured the DHCP packets are dropped and a message is logged.

The following example shows how to enable the DHCP Trust functionality:

```
(host)(config)# interface-profile port-security-profile ps1
    trust dhcp
```


Enabling Loop Protect Functionality

Port Loop Protect functionality is configured as part of the port level security configuration. You can attach the port-security profile to any Layer 2 interface. Enabling Loop Protect will disable a port when it detects a loop. You can automatically re-enable the port by setting the auto-recovery option. Otherwise, you can recover the port manually using the `clear` command.

Use the following CLI commands to enable Loop Protect and the auto-recovery option:

```
(host) (config) #interface-profile port-security-profile <profile-name>
(host) (Port security profile "<profile-name>") #loop-protect auto-recovery-time <time in seconds>
```

Set a value for `auto-recovery-time` to enable the auto-recovery option. The port will then be automatically re-enabled and recovered from the error after the specified time. By default, auto-recovery is disabled. Auto-recovery remains disabled, if you enable `loop-protect` without setting the `auto-recovery-time` option or by setting the value to 0.

You can also disable the auto-recovery option using the following command:

```
(host) (Port security profile "<profile-name>") #no loop-protect auto-recovery-time
```

Use the following command to disable the Loop Protect functionality:

```
(host) (Port security profile "<profile-name>") #no loop-protect
```



It is recommended that you disable Spanning Tree using the following command before enabling Loop Protect on an interface:

```
(host) (config) #spanning-tree no mode
Otherwise, you will see the following warning message:
```

Warning: Port Loop Protect configured in the `port-security-profile`, will be inactive. It becomes active when MSTP/PVST is disabled.

Enabling MAC Limit Functionality

The MAC Limit functionality will be configured as part of the port level security configuration. This profile can be attached to the interface.

You can use the following command to configure the MAC Limit:

```
(host)(config)# interface-profile port-security-profile <profile-name>
    mac-limit <limit> action {drop|log|shutdown}
    auto-recovery-time <time in seconds>
```

The following example shows how to enable the MAC Limit functionality:

```
(host)(config)# interface-profile port-security-profile MAC_Limit
    mac-limit 30 action drop
    auto-recovery-time 50
```



The maximum value for `auto-recovery-time` for all the port security functionalities is 65535 seconds. You can apply `auto-recovery-time` option only if the action is shutdown.

Attaching Port Security Profile to Interface

To enable the Port Security functionality on an interface, you must attach a port-security profile to it. You can use the following commands to associate a port-security profile with an interface:

For GigabitEthernet:

```
(host) (config) #interface gigabitethernet <slot/mod/port>
```

```
(host) (gigabitethernet "<slot/mod/port>") #port-security-profile <profile-name>
```

For Port-channel:

```
(host) (config) #interface port-channel <id>
```

```
(host) (port-channel "<id>") #port-security-profile <profile-name>
```

Viewing Port Errors

You can use the following command to view the list of ports that are detected with port errors and the time at which they will be recovered automatically, if auto-recovery is enabled:

```
(host) #show port-error-recovery
```

Layer-2 Interface Error Information

Interface	Error	Recovery Time
-----	-----	-----
Pc5	Shutdown (Loop Detected)	2012-02-08 16:42:45 (PST)
GE0/0/42	Shutdown (Loop Detected)	No Auto recovery
Pc1	Shutdown (Loop Detected)	2012-02-07 16:45:40 (PST)
Pc2	Shutdown (RA Guard)	2012-02-08 16:42:45 (PST)
GE0/0/14	Log (Mac Limit Exceeded)	No Auto recovery
GE0/0/2	Drop (DHCP Trust Error)	2012-02-07 16:45:40 (PST)
GE0/0/5	Log (MAC Limit exceed)	No Auto recovery
	Drop (RA guard)	No Auto recovery
GE1/0/24	Shutdown (BPDU received)	2012-10-18 11:25:17 (PST)
		No Auto Recovery

Recovering Ports Manually

You can use the CLI to manually recover the port errors. To recover the ports on a specific interface execute the following command:

```
(host) #clear port-error-recovery interface <interface-name>
```

The following command clears the errors on gigabitethernet 0/0/42:

```
(host) #clear port-error-recovery interface gigabitethernet 0/0/42
```

To clear the port errors on all interfaces execute the following command:

```
(host) #clear port-error-recovery
```

Sample Configurations

To configure the port security profile:

```
(host) (config) # interface-profile port-security-profile port-security-1
(host (port security profile port-security-1)#
    ipv6-ra-guard action drop auto-recovery-time 60
    no trust dhcp
    loop-protect auto-recovery-time 10
    mac-limit 30 action drop auto-recovery-time 50
```

To attach the port security profile to the interface:

```
(host)(config)# interface gigabitethernet 0/0/6
    port-security-profile port-security-1

(host) (config) #interface port-channel 3
    port-security-profile port-security-1
```

Some protocols or features prevent bridge loops in a Layer 2 network, rogue switches, or end hosts can degrade the network by creating and propagating traffic storms.

Storm control prevents interfaces from disruptions by providing protection against excessive ingress rates of unknown-unicast, multicast, and broadcast traffic.

Important Points to Remember

- The configured storm control bandwidth percentage applies to all types of traffic.
- If the rate is 100%, no traffic is rate limited. If the rate is 50% then 50% of configured traffic is rate limited.
- Individual levels of storm control per traffic type is not supported. All types are set to single percentage.
- By default, storm control is enabled for unknown-unicast and broadcast traffic.
- Storm Control is configured from the command line only. You configure it under the switching-profile.

Configuration Steps

Use the following steps, from the command line, to configure and verify Storm Control.

1. Define the level of storm-control based on percentage of interface speed. Range is 50 to 100%.

```
(host) (config) #interface-profile switching-profile STORM_CONTROL
(host) (switching profile "STORM_CONTROL") #storm-control-bandwidth 80
```

2. Enable the type(s) of traffic you want controlled.

```
(host) (switching profile "STORM_CONTROL") #storm-control-unknown-unicast
(host) (switching profile "STORM_CONTROL") #storm-control-multicast
(host) (switching profile "STORM_CONTROL") #storm-control-broadcast
```

3. Apply the configured switching-profile to the interface.

```
(host) (config) #interface gigabitethernet 0/0/20
(host) (gigabitethernet "0/0/20") #switching-profile STORM_CONTROL
```

4. Verify the configuration.

```
(host) #show interface-profile switching-profile STORM_CONTROL
```

```
switching profile "STORM_CONTROL"
```

Parameter	Value
Switchport mode	access
Access mode VLAN	1
Trunk mode native VLAN	1
Enable broadcast traffic rate limiting	Enabled
Enable multicast traffic rate limiting	Enabled
Enable unknown unicast traffic rate limiting	Enabled
Max allowed rate limit traffic on port in percentage	80
Trunk mode allowed VLANs	1-4094

Access control lists (ACLs) are a common way of restricting certain types of traffic on a physical port. The Mobility Access Switch supports multiple types of access control lists to provide flexibility to control the traffic. This chapter describes the different types of ACLs supported and how to configure them on the Mobility Access Switch.

This chapter includes the following topics:

- “Types of ACLs” on page 221
- “Configuring the ACLs” on page 222
- “Verifying the ACL configuration” on page 224

Types of ACLs

- Ethertype ACLs are used to filter based on the Ethertype field in the frame header. Ethertype ACLs can be either named or numbered, with valid numbers in the range of 200-299. These ACLs can be used to permit IP while blocking other non-IP protocols, such as IPX or AppleTalk.
- MAC ACLs are used to filter traffic on a specific source MAC address or range of MAC addresses. MAC ACLs can be either named or numbered, with valid numbers in the range of 700-799 and 1200-1299.
- Standard ACLs permit or deny traffic based on the source IP address of the packet. Standard ACLs can be either named or numbered, with valid numbers in the range of 1-99 and 1300-1399. Standard ACLs use a bitwise mask to specify the portion of the source IP address to be matched.
- Extended ACLs permit or deny traffic based on source or destination IP address, or IP protocol. Extended ACLs can be named or numbered, with valid numbers in the range 100-199 and 2000-2699.
- Stateless ACLs are named, used to define stateless packet filtering and quality of service (QoS). A stateless ACL statically evaluates packet contents. .



In a stateless ACL the traffic in the reverse direction will be allowed unconditionally.

Port ACLs (PACLs)

ACLs provide the ability to filter ingress traffic based on conditions specified in the ACL. Port ACLs perform access control on all traffic entering or leaving the specified Layer 2 port. PACLs provides access control based on the Layer 3 addresses (for IP protocols), Layer 2 MAC addresses (for non-IP protocols), or Layer 4 port information and ranges. A Layer 2 port is a physical LAN or trunk port that belongs to a VLAN. The PACLs are applied on both the ingress and egress traffic with the following exceptions for egress traffic:

- Egress ACLs are applied only on interfaces and not on user roles.
- When QoS-profile is applied on egress ACL, only the dot1p and dscp values are applicable. The trafficclass, drop-precedence are not applicable.



You can apply all the types of ACLs to a port and only the MAC, Ethertype and Stateless ACLs can be applied to a user role. The MAC and Ethertype ACLs only apply to non-IP traffic and the Stateless ACL to IP traffic from the user.

Configuring the ACLs

ACL is order dependent. ACLs are executed in the sequential order in which access control entries (ACE) are defined. The Mobility Access Switch process the ACEs in the order in which it is configured. Usually the deny ACEs are configured before permit ACEs. There is an implicit deny at the end of every ACL. Therefore, if there are no matching ACEs for a given packet, then that packet will be dropped.

This section describes the CLIs to configure the different ACLs:

Ethertype ACL

The below command configures an Ethertype access control list (ACL).

```
(host)(config) #ip access-list eth ETHER_TYPE
(host)(config-eth-ETHER_TYPE) #deny 0x880
(host)(config-eth-ETHER_TYPE) #permit any
(host)(config-stateless-ETHER_TYPE) #exit
```

To configure the ACL when a particular access control entry(ACE) is changed in a particular ACL:

```
(host)(config) #ip access-list eth ETHER_TYPE
(host)(config-eth-ETHER_TYPE) #deny 0x0806
(host)(config-eth-ETHER_TYPE) #permit any
(host)(config-eth-ETHER_TYPE) #exit
```

MAC ACL

A range of MAC address can be matched by using a wildcard mask or a particular host using the host keyword:

```
(host)(config) #ip access-list mac MAC_LIST
(host)(config-mac-MAC_LIST) #deny 00:11:22:00:00:00 00:00:00:FF:FF:FF
(host)(config-mac-MAC_LIST) #deny host 00:66:77:88:99:AA
(host)(config-mac-MAC_LIST) #permit any
(host)(config-mac-MAC_LIST) #exit
```

Standard ACL

The Standard ACL match the source IP address of the packet. The IP address to be matched can be either a range of IP Addresses using wildcard mask or a particular host:

```
(host)(config) #ip access-list standard STANDARD
(host)(config-standard-STANDARD) #deny 1.1.1.0 0.0.0.255
(host)(config-standard-STANDARD) #deny host 192.168.10.100
(host)(config-standard-STANDARD) #permit any
(host)(config-standard-STANDARD) #exit
```

Extended ACL

The Extended ACL extends the standard ACL by matching IP address of the source and destination, port number of the source and destination, and the protocol:

```
(host)(config) #ip access-list extended EXTENDED
(host)(config-extended-EXTENDED) #deny icmp 1.1.1.0 0.0.0.255 2.2.2.0 0.0.0.255 echo-reply
(host)(config-extended-EXTENDED) #deny tcp host 192.168.1.1 eq 53 host 20.1.1.1 range 20 30 established
(host)(config-extended-EXTENDED) #permit any any any
(host)(config-extended-EXTENDED) #exit
```

Stateless ACL

Stateless ACL provides userlevel access control on statically configured ACL.

```
(host)(config) #ip access-list stateless STATELESS
(host)(config-stateless-STATELESS) #network 10.100.100.0 255.255.255.0 any tcp 8888
(host)(config-stateless-STATELESS) #deny log
(host)(config-stateless-STATELESS) #any host 10.100.100.200 any deny log
(host)(config-stateless-STATELESS) #any any any permit
(host)(config-stateless-STATELESS) #exit
```

Stateless ACL provides additional options that can be specified on matching the traffic. [Table 1](#) describes the parameters you configure for a stateless ACL.

Table 1 Stateless ACL Configuration Parameters

Parameter	Description
blacklist	Configure the ACL blacklist user when the ACL rule is matched. If the ACE entry is matched, the traffic from that particular user is denied and the user is blacklisted for 3600 seconds
log	Configure to display the log information when the ACL is applied.
policer-profile	To attach the policer-profile to the ACL
position	Defines or redefines the position of an ACE in an ACL.
qos-profile	QoS profile can be configured to assign specific TC/DP, DSCP, and 802.1p values. This option attaches the qos-profile to the ACL
time-range	Associate a time-range to an ACL. This configures the ACL to filter traffic during the specified time-range

The following ACL actions are not supported for Egress ACLs (For Stateless ACL applied in egress direction):

- Blacklist
- Log

For the policer profile attached to the egress ACL, only the following are permitted:

- Action: drop/permit
- counters

To apply ACL to a port in ingress direction, use the following CLI:

```
(host)(config) #interface gigabitethernet 0/0/0
(host)(gigabitethernet "0/0/0") #ip access-group in <acl_name>
(host)(gigabitethernet "0/0/0") #exit
```

To apply ACL to a port in egress direction, use the following CLI:

```
(host)(config) #interface gigabitethernet 0/0/0
(host)(gigabitethernet "0/0/0") #ip access-group out <acl_name>
(host)(gigabitethernet "0/0/0") #exit
```

Verifying the ACL configuration

```
(host)(config) #show ip access-list ETHER_TYPE
(host)(config) #ip access-list eth ETHER_TYPE
ETHER_TYPE
-----
Priority Action EtherType Mirror
-----
1          deny    0x8800
2          permit  any
```

You can use the same command to verify the ACL configuration after changing the ACE:

```
(host)(config) #show ip access-list ETHER_TYPE
ip access-list eth ETHER_TYPE
ETHER_TYPE
-----
Priority Action EtherType Mirror
-----
1          deny    0x8800
2          deny    0x8100 <-ACE has been edited
3          permit  any

(host)(config) #show ip access-list MAC_LIST
(host)(config-mac-MAC_LIST) #ip access-list mac MAC_LIST
(host)(config-mac-MAC_LIST) #deny 00:11:22:00:00:00 00:00:00:ff:ff:ff
(host)(config-mac-MAC_LIST) #deny host 00:66:77:88:99:aa
(host)(config-mac-MAC_LIST) #permit any
(host)(config-mac-MAC_LIST) #exit

(host)(config) #show ip access-list STANDARD
(host)(config-std-STANDARD) #ip access-list standard STANDARD
(host)(config-std-STANDARD) #deny 1.1.1.0 0.0.0.255
(host)(config-std-STANDARD) #deny host 192.168.10.100
(host)(config-std-STANDARD) #permit any
(host)(config-std-STANDARD) #exit

(host)(config) #show ip access-list EXTENDED
(host)(config-ext-EXTENDED) #ip access-list extended EXTENDED
(host)(config-ext-EXTENDED) #deny icmp 1.1.1.0 0.0.0.255 2.2.2.0 0.0.0.255 echo-reply
(host)(config-ext-EXTENDED) #deny udp 6.6.6.0 0.0.0.255 any eq 53
(host)(config-ext-EXTENDED) #permit 0 any any
(host)(config-ext-EXTENDED) #exit

(host)(config) #show ip access-list STATELESS
ip access-list stateless STATELESS
STATELESS
-----
Priority Source Destination Service Action TimeRange Log Expired QoS Policar Blacklist
Mirror IPv4/6
-----
1 10.100.100.0 255.255.255.0 any tcp 8888 deny Yes 4
2 any 10.100.100.200 any deny Yes 4
3 any any any permit 4
```


This chapter describes how to configure quality of service (QoS) on the Mobility Access Switch. This chapter contains the following major sections:

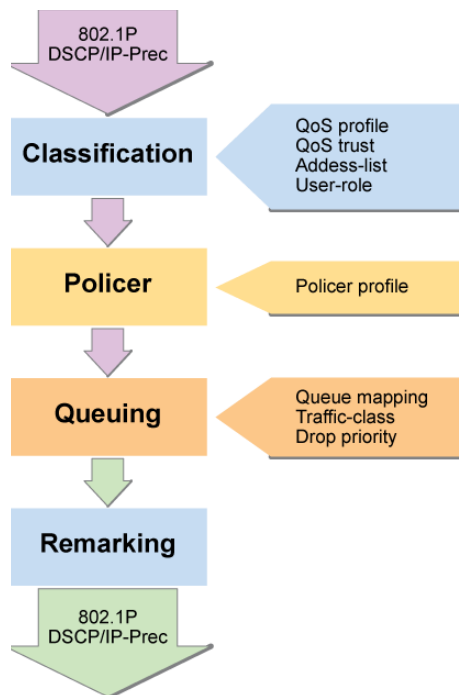
- “QoS Concepts” on page 225
- “Configuring QoS” on page 227

QoS Concepts

This section contains the following sections:

- “Overview” on page 225
- “Profiles and Queues” on page 225
- “Classification” on page 226
- “Policing” on page 227

Overview



Profiles and Queues

The Mobility Access Switch supports:

- A QoS profile that can be applied to an interface, user role, and traffic flow.
- Eight queues per interface in hardware.
- Eight traffic classes (TC), which map to the corresponding queue (0 – 7).
- Drop-precedence for controlling tail-drop.

Classification

This section contains the following sections:

- [“Trust Mode” on page 226](#)
- [“Untrusted Mode” on page 226](#)

Trust Mode

When the QoS mode on a port is set to be trusted, the received 802.1P/DSCP is considered trustworthy and the frame is allowed to exit with those values intact. The received DSCP or 802.1P value is used to index predefined QoS profiles to determine traffic class and drop precedence. These QoS profiles cannot be edited at this time.

The Mobility Access Switch supports several modes:

- Layer 2 QoS Trust Mode - Port is configured to trust the IEEE 802.1P user priority. This is relevant for 802.1Q packets
- Layer 3 QoS Trust Mode - Port is configured to trust the received DSCP value of the frame.
- Auto (L2+L3) trust mode prioritizes DSCP over 802.1P. If the received frame is IP, the DSCP value is used for indexing the QoS profile. If the received tagged frame is non-IP, then the 802.1P value is used for indexing the QoS profile.

The following table shows DSCP-Queue mapping:

Table 1 *DSCP-Queue Mapping*

DSCP	802.1p	Queue
0-7	0	0
8-15	1	1
16-23	2	2
24-31	3	3
32-39	4	4
40-47	5	5
48-55	6	6
56-63	7	7

- DP is defined as low for first 4 values (0-3) and high for last 4 values (4-7) for each DSCP range.
- For 802.1p, DP is defined low for all values.

Untrusted Mode

- The default is “untrust” for all interfaces where all incoming traffic are mapped to TC “0” and are then subsequently mapped to egress queue 0.

Profile

- QoS profile can be configured to assign specific TC/DP, DSCP, and 802.1p values.
- The QoS profile can be then applied to:
 - Interface (interface-profile)

- Stateless access-list
- User-role
- Policer profile

Policing

- Limits inbound transmission rate of a class of traffic on the basis of user-defined criteria.
- Policer can be applied to stateless ACL, interface, and user-role.
- 1-rate 3-color policer is supported at FCS.
 - Traffic rate below CIR or burst below CBS limit is considered “conforming” and is allowed to pass through the policer.
 - Traffic rate exceeding CIR, and bursting below EBS limit is considered “exceeding” and is allowed to pass through the policer by default.
 - Traffic rate exceeding CIR, and bursting above EBS limit is considered “violating” and is dropped at the policer by default.

Configuring QoS

This section contains the following sections:

- [“Configuring QoS Trust Mode” on page 227](#)
- [“Configuring QoS-Profile under an Interface” on page 228](#)
- [“Configuring QoS-Profile under a Stateless ACL” on page 228](#)
- [“Configuring QoS-Profile under a User-Role” on page 228](#)
- [“Configuring Policer under Policer-Profile” on page 228](#)
- [“Configuring Policer-Profile under an Interface” on page 228](#)
- [“Configuring Policer-Profile under a Stateless ACL” on page 228](#)
- [“Configuring QoS-Profile under a User-Role” on page 228](#)

Configuring QoS Trust Mode

To configure QoS trust mode, follow these steps:

1. In the configuration mode, configure the appropriate interface:


```
(host)(config) #interface gigabitethernet 0/0/6
```
2. In the interface mode, you can configure the following options:

To configure QoS trust auto, use the following command:

```
(host)(gigabitethernet "0/0/6") #qos trust auto
```

To configure QoS trust dot1p, use the following command:

```
(host)(gigabitethernet "0/0/6") #qos trust dot1p
```

To configure QoS trust dscp, use the following command:

```
(host)(gigabitethernet "0/0/6") #qos trust dscp
```

To disable QoS trust, use the following command:

```
(host)(gigabitethernet "0/0/6") #qos trust none
```

To display the predefined QoS profiles, use the following command.

```
(host)(config) #show qos-profile trusted
```

Configuring QoS-Profile

To configure a QoS under a QoS profile, use the following commands:

```
(host) (config) #qos-profile QOS1
(host) (QoS Profile "QOS1") #dot1p <value>
(host) (QoS Profile "QOS1") #drop-precedence <low/high>
(host) (QoS Profile "QOS1") #dscp <value>
(host) (QoS Profile "QOS1") #traffic-class <value>
```

Configuring QoS-Profile under an Interface

To configure a QoS profile on an Interface, use the following commands:

```
(host) (config) #interface gigabitethernet 0/0/19
(host) (gigabitethernet "0/0/19") #qos-profile QOS1
```

Configuring QoS-Profile under a Stateless ACL

To configure QoS Profile under a Stateless ACL, use the following commands:

```
(host) (config) #ip access-list stateless STATELESS
(host) (config-stateless-STATELESS)#any any any permit qos-profile QOS1
```

Configuring QoS-Profile under a User-Role

To configure QoS Profile under a user-role, use the following commands:

```
(host) (config) #user-role EMPLOYEE_1
(host) (config-role) #qos-profile QOS1
```

Configuring Policer under Policer-Profile

To configure Policer under a Policer profile, use the following commands:

```
(host) (config) #policer-profile 100MBPS
(host) (Policer Profile "100MBPS") #cir 100000 (100m)
(host) (Policer Profile "100MBPS") #cbs 100000 (100m)
(host) (Policer Profile "100MBPS") #pbs 110000 (110m)
(host) (Policer Profile "100MBPS") #exceed-action <permit | remark | drop>
(host) (Policer Profile "100MBPS") #exceed-profile <QoS profile for remark>
(host) (Policer Profile "100MBPS") #violate-action <permit | remark | drop>
```



When remark action is configured, a corresponding QoS profile must be configured also.

Configuring Policer-Profile under an Interface

To configure a policer profile on an interface, use the following commands:

```
(host) (config) #interface gigabitethernet 0/0/19
(host) (gigabitethernet "0/0/19") #policer-profile 100MBPS
```

Configuring Policer-Profile under a Stateless ACL

To configure a policer profile on an interface, use the following commands:

```
(host) (config) #ip access-list stateless STATELESS
(host) (config-stateless-STATELESS)#any any any permit policer-profile 100MBPS
```

Configuring Policer-Profile under a User-role

```
(host) (config) #user-role EMPLOYEE_1  
(host) (config-role) #policer-profile 100MBPS
```


This chapter describes how to configure authentication servers. It contains the following sections:

- “Important Points to Remember” on page 231
- “Server and Server Group Concepts” on page 231
- “Configuring Authentication Servers” on page 232
- “Internal Database Concepts” on page 240
- “Configuring the Internal Database” on page 240
- “Server Group Concepts” on page 242
- “Assigning Server Groups” on page 245
- “Authentication Timers” on page 249

Important Points to Remember

The Mobility Access Switch allows you to use an external authentication server or the internal user database to authenticate clients who need to access the wired network.

For an external authentication server to process requests from the Mobility Access Switch, you must configure the server to recognize the switch. Refer to the vendor documentation for information on configuring the authentication server.

Server and Server Group Concepts

The Mobility Access Switch supports the following external authentication servers:


- RADIUS (Remote Authentication Dial-In User Service)
- LDAP (Lightweight Directory Access Protocol)
- TACACS+ (Terminal Access Mobility Access Switch Access Control System)

Additionally, you can use the Mobility Access Switch’s internal database to authenticate users. You create entries in the database for users and their passwords and default role.

You can create *groups* of servers for specific types of authentication. For example, you can specify one or more RADIUS servers to be used for 802.1x authentication. The list of servers in a server group is an ordered list. This means that the first server in the list is always used unless it is unavailable, in which case the next server in the list is used. You can configure servers of different types in one group — for example, you can include the internal database as a backup to a RADIUS server.

Figure 1 shows a server group named Radii that contains two RADIUS servers, Radius-1 and Radius-2. The Radii server group is assigned to the server group for 802.1x authentication.

Figure 1 Server Group



Server names must be unique. You can configure the same server in multiple server groups, and you must configure the server before you can add it to a server group.



If you are using the **Mobility Access Switch's** internal database for user authentication, use the predefined "Internal" server group.

You can also include conditions for server-derived user roles or VLANs in the server group configuration. The server derivation rules apply to all servers in the group.

Configuring Authentication Servers

This section describes how to configure authentication servers on the Mobility Access Switch. It contains the following sections:

- "Configuring a RADIUS Server" on page 236
- "Configuring an LDAP Server" on page 238
- "Configuring a TACACS+ Server" on page 240
- "Configuring the Internal Database" on page 240
- "Using the CLI" on page 241
- "User Authentication" on page 246
- "Management Authentication" on page 246
- "Radius Accounting" on page 246
- "TACACS+ Accounting" on page 249
- "Using the CLI" on page 250

Radius Server Authentication

This section contains the following sections:

- "RADIUS Server Username/Password Authentication" on page 233
- "RADIUS Server Authentication with VSA" on page 233
- "RADIUS Server Authentication with Server-Derivation Rule" on page 233
- "Disabling Authentication of Local Management User Accounts" on page 234

- [“Verifying the configuration” on page 234](#)
- [“Resetting the Admin or Enable Password” on page 234](#)
- [“Bypassing the Enable Password Prompt” on page 235](#)
- [“Setting an Administrator Session Timeout” on page 235](#)

RADIUS Server Username/Password Authentication

In this example, an external RADIUS server is used to authenticate management users. Upon authentication, users are assigned the default role root.

In the CLI

```
aaa authentication-server radius rad1
    host <ipaddr>
    key <string>
aaa server-group corp_rad
    auth-server rad1

aaa authentication mgmt
    default-role root
    enable
    server-group corp_rad
```

RADIUS Server Authentication with VSA

In this scenario, an external RADIUS server authenticates management users and returns to the Mobility Access Switch the Aruba vendor-specific attribute (VSA) called Aruba-Admin-Role that contains the name of the management role for the user. The authenticated user is placed into the management role specified by the VSA.

The Mobility Access Switch configuration is identical to the [“RADIUS Server Username/Password Authentication” on page 233](#). The only difference is the configuration of the VSA on the RADIUS server. Ensure that the value of the VSA returned by the RADIUS server is one of the predefined management roles. Otherwise, the user will have *no* access to the Mobility Access Switch.

RADIUS Server Authentication with Server-Derivation Rule

A RADIUS server can return to the Mobility Access Switch a standard RADIUS attribute that contains one of the following values:

- The name of the management role for the user
- A value from which a management role can be derived

For either situation, configure a server-derivation rule for the server group.

In the following example, the RADIUS server returns the attribute Class to the Mobility Access Switch. The value of the attribute can be either “root” or “network-operations” depending upon the user; the returned value is the role granted to the user.



Ensure that the value of the attribute returned by the RADIUS server is one of the predefined management roles. Otherwise, the management user will not be granted access to the Mobility Access Switch.

In the CLI

```
aaa authentication-server radius rad1
    host <ipaddr>
    key <string>
```

```

aaa server-group corp_rad
  auth-server rad1
  set role condition Class value-of

aaa authentication mgmt
  default-role read-only
  enable
  server-group corp_radIn the CLI
aaa authentication-server radius rad1
  host <ipaddr>
  enable

aaa server-group corp_rad
  auth-server rad1
  set role condition Class equals it set-value root

aaa authentication mgmt
  default-role read-only
  enable
  server-group corp_rad

```

Disabling Authentication of Local Management User Accounts

You can disable authentication of management user accounts in local switches if the configured authentication server(s) (RADIUS or TACACS+) are not available.

You can disable authentication of management users based on the results returned by the authentication server. When configured, locally-defined management accounts (for example, admin) are not allowed to log in if the server(s) are reachable and the user entry is not found in the authentication server. In this situation, if the RADIUS or TACACS+ server is unreachable, meaning it does not receive a response during authentication, or fails to authenticate a user because of a timeout, local authentication is used and you can log in with a locally-defined management account.

In the CLI

```
mgmt-user localauth-disable
```

Verifying the configuration

To verify if authentication of local management user accounts is enabled or disabled, use the following command:

```
show mgmt-user local-authentication-mode
```

Resetting the Admin or Enable Password

This section describes how to reset the password for the default administrator user account (**admin**) on the Mobility Access Switch. The default password is admin123.

Use this procedure if the administrator user account password is lost or forgotten.

1. Connect a local console to the serial port on the Mobility Access Switch.
2. From the console, login in the Mobility Access Switch using the username **password** and the password **forgetme!**.
3. Enter enable mode by typing in **enable**, followed by the password **enable**.
4. Enter configuration mode by typing in **configure terminal**.
5. To configure the administrator user account, enter **mgmt-user admin root**. Enter a new password for this account. Retype the same password to confirm.

6. Exit from the configuration mode, enable mode, and user mode.

This procedure also resets the enable mode password to **enable**. If you have defined a management user password policy, make sure that the new password conforms to this policy.

Figure 2 is an example of how to reset the password. The commands in bold type are what you enter.

Figure 2 *Resetting the Password*

```
(host)
User: password
Password: forgetme!
(host) >enable
Password: enable
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(host) (config) #mgmt-user admin root
Password: *****
Re-Type password: *****
(host) (config) #exit
(host) #exit
(host) >exit
```

After you reset the administrator user account and password, you can login to the Mobility Access Switch and reconfigure the enable mode password. To do this, enter configuration mode and type the **enable secret** command. You are prompted to enter a new password and retype it to confirm. Save the configuration by entering **write memory**.

Figure 3 details an example reconfigure the enable mode password. Again, the command you enter displays in bold type.

Figure 3 *Reconfigure the enable mode password*

```
User: admin
Password: *****
(host) >enable
Password: *****
(host) #configure terminal
Enter Configuration commands, one per line. End with CNTL/Z

(host) (config) #enable secret
Password: *****
Re-Type password: *****
(host) (config) #write memory
```

Bypassing the Enable Password Prompt

The bypass enable feature lets you bypass the enable password prompt and go directly to the privileged commands (config mode) after logging on to the Mobility Access Switch. This is useful if you want to avoid changing the enable password due to company policy.

Use the `enable bypass` CLI command to bypass the enable prompt and go directly to the privileged commands (config mode). Use the `no enable bypass` CLI command to restore the enable password prompt.

Setting an Administrator Session Timeout

You can configure the number of seconds after which an Administrator's WebUI or CLI session times out.

Setting a CLI Session Timeout

To define a timeout interval for a CLI session, use the command:

```
login session timeout <value>
```

In the above command, <val> can be any number of minutes from 5 to 60 or seconds from 1 to 3600, inclusive. You can also specify a timeout value of 0 to disable CLI session timeouts.

Setting a WebUI Session Timeout

To define a timeout interval for a WebUI session, use the command:

```
web-server session timeout <session-timeout>
```

In the above command, <session-timeout> can be any number of seconds from 30 to 3600, inclusive

Configuring a RADIUS Server

[Table 1](#) describes the parameters you configure for a RADIUS server.

Table 1 RADIUS Server Configuration Parameters

Parameter	Description
Host	IP address of the authentication server. Default: N/A
Key	Shared secret between the Mobility Access Switch and the authentication server. The maximum length is 128 characters. Default: N/A
Authentication Port	Authentication port on the server. Default: 1812
Accounting Port	Accounting port on the server Default: 1813
Retransmits	Maximum number of retries sent to the server by the Mobility Access Switch before the server is marked as down. Default: 3
Timeout	Maximum time, in seconds, that the Mobility Access Switch waits before timing out the request and resending it. Default: 5 seconds
NAS ID	Network Access Server (NAS) identifier to use in RADIUS packets. Default: N/A
NAS IP	NAS IP address to send in RADIUS packets. You can configure a “global” NAS IP address that the Mobility Access Switch uses for communications with all RADIUS servers. If you do not configure a server-specific NAS IP, the global NAS IP is used. To set the global NAS IP in the CLI, enter the ip radius nas-ip <i>ipaddr</i> command. Default: N/A

Table 1 RADIUS Server Configuration Parameters (Continued)

Parameter	Description
Source Interface	<p>Enter a VLAN number ID.</p> <p>Allows you to use source IP addresses to differentiate RADIUS requests.</p> <p>Associates a VLAN interface with the RADIUS server to allow the group-specific source interface to override the global configuration.</p> <ul style="list-style-type: none"> • If you associate a Source Interface (by entering a VLAN number) with a configured server, then the source IP address of the packet will be that interface's IP address. • If you do not associate the Source Interface with a configured server (leave the field blank), then the IP address of the global Source Interface will be used.
Use MD5	<p>Use MD5 hash of cleartext password.</p> <p>Default: disabled</p>
Mode	<p>Enables or disables the server.</p> <p>Default: enabled</p>

Using the CLI

```

aaa authentication-server radius <name>
  host <ipaddr>
  key <key>
  enable

```

RADIUS Server Authentication Codes

A configured RADIUS server will return the following standard response codes.

Table 2 RADIUS Authentication Response Codes

Code	Description
0	Authentication OK.
1	Authentication failed—user/password combination not correct.
2	Authentication request timed out—No response from server.
3	Internal authentication error.
4	Bad Response from RADIUS server. Verify shared secret is correct.
5	No RADIUS authentication server is configured.
6	Challenge from server. (This does not necessarily indicate an error condition.)

RADIUS Change of Authorization

The following command configures a RADIUS server that can send user disconnect and change-of-authorization messages, as described in RFC 3576, “Dynamic Authorization Extensions to Remote Dial In User Service (RADIUS)”.

```
aaa rfc-3576-server <server-ip-addr>
    key <psk>
    no
```

The following command configures an RFC 3576 server:

```
(host) #aaa rfc-3576-server 10.1.1.245
(host) #key asdfjkl;
```

Configuring an LDAP Server

[Table 3](#) describes the parameters you configure for an LDAP server.

Table 3 *LDAP Server Configuration Parameters*

Parameter	Description
Host	IP address of the LDAP server. Default: N/A
Admin-DN	Distinguished name for the admin user who has read/search privileges across all the entries in the LDAP database (the user need not have write privileges but the user should be able to search the database, and read attributes of other users in the database).
Admin Password	Password for the admin user. Default: N/A
Allow Clear-Text	Allows clear-text (unencrypted) communication with the LDAP server. Default: disabled
Authentication Port	Port number used for authentication. Default: 389
Base-DN	Distinguished Name of the node which contains the entire user database to use. Default: N/A
Filter	Filter that should be applied to search of the user in the LDAP database: Default: (objectclass=*)
Key Attribute	Attribute that should be used as a key in search for the LDAP server. For Active Directory, the value is sAMAccountName. Default: sAMAccountName
Timeout	Timeout period of a LDAP request, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled

Table 3 LDAP Server Configuration Parameters (Continued)

Parameter	Description
Preferred Connection Type	<p>Preferred type of connection between the Mobility Access Switch and the LDAP server. The default order of connection type is:</p> <ol style="list-style-type: none"> 1. ldap-s 2. start-tls 3. clear-text <p>The Mobility Access Switch will first try to contact the LDAP server using the preferred connection type, and will only attempt to use a lower-priority connection type if the first attempt is not successful.</p> <p>Note: If you select clear-text as the preferred connection type, you must also enable the allow-cleartext option.</p>

Using the CLI

aaa authentication-server ldap <name>	
admin-dn	The Distinguished Name for the Admin user who can search for the LDAP user. E.g. (cn=Admin-Name,cn=Users,dc=department-name,dc=domain-name,dc=com)
admin-passwd	The password for the Admin user who can search for the LDAP user
allow-cleartext	Allow unencrypted communication with LDAP server
authport	Specify port number used for authentication. Range: 1-65535. Default : 389. Port 636 will be attempted for LDAP over SSL - LDAPS, 389 will be attempted for SSL over LDAP - Start TLS and for clear text.
base-dn	The Base Distinguished Name of search for the LDAP server. E.g. (cn=Users,dc=qa,dc=domain,dc=com)
clone	Copy data from another LDAP Server
enable	Enable LDAP server
filter	The filter that should be used as a key in a search for the LDAP server
host	IP address of LDAP server
key-attribute	The attribute that should be used as a key in search for the LDAP server. For PAP, the value is sAMAccountName. For EAP-TLS termination the value is userPrincipalName.
no	Delete Command
preferred-conn-type	Preferred connection type
timeout	Timeout period for LDAP request. Range: 1-30. Default: 20.

Configuring a TACACS+ Server

Table 4 defines the TACACS+ server parameters.

Table 4 TACACS+ Server Configuration Parameters

Parameter	Description
Host	IP address of the server. Default: N/A
Key	Shared secret to authenticate communication between the TACACS+ client and server. Default: N/A
TCP Port	TCP port used by server. Default: 49
Retransmits	Maximum number of times a request is retried. Default: 3
Timeout	Timeout period for TACACS+ requests, in seconds. Default: 20 seconds
Mode	Enables or disables the server. Default: enabled
Session Authorization	Enables or disables session authorization. Session authorization turns on the optional authorization session for admin users. Default: disabled

Using the CLI

The following command configures, enables a TACACS+ server and enables session authorization:

```
aaa authentication-server tacacs <name>
  clone default
  host <ipaddr>
  key <key>
  enable
  session-authorization
```

Internal Database Concepts

You can create entries, in the Mobility Access Switch's internal database, to use to authenticate clients. The internal database contains a list of clients along with the password and default role for each client. When you configure the internal database as an authentication server, client information in incoming authentication requests is checked against the internal database.

Configuring the Internal Database

The default server-group (aaa server-group "default") has the internal user database defined as the first authentication server by default. You must first add users if you want to effectively use the internal user database in the Mobility Access Switch.

Table 5 defines the required and optional parameters used in the internal database.

Table 5 *Internal Database Configuration Parameters*

Parameters	Description
User Name	(Required) Enter a user name or select Generate to automatically generate a user name. An entered username can be up to 64 characters in length.
Password	(Required) Enter a password or select Generate to automatically generate a password string. An entered password must be a minimum of 6 characters and can be up to 128 characters in length.
Role	Role for the client. In order for this role to be assigned to a client, you need to configure a server derivation rule, as described in “Configuring Server-Derivation Rules” on page 244 . (A user role assigned through a server-derivation rule takes precedence over the default role configured for an authentication method.)
E-mail	(Optional) E-mail address of the client.
Enabled	Select this checkbox to enable the user as soon as the user entry is created.
Expiration	Select one of the following options: <ul style="list-style-type: none"> Entry does not expire: No expiration on user entry Set Expiry time (mins): Enter the number of minutes the user will be authenticated before their user entry expires. Set Expiry Date (mm/dd/yyyy) Expiry Time (hh:mm): To select a specific expiration date and time, enter the expiration date in mm/dd/yyyy format, and the expiration time in hh:mm format.

Using the CLI

```
local-userdb add {generate-username|username <name>} {generate-password|password
<password>} {remote-ip<remote-ip>}
local-userdb modify {username < name>} {remote-ip<remote-ip>}
```

The output of **show local-userdb** command:

User Summary

Name IP Grantor-Name	Password	Role	E-Mail	Enabled	Expiry	Status	Sponsor-Name	Remote-
----	-----	----	-----	-----	-----	-----	-----	-----
68:b5:99:d7:ff:bc	68:b5:99:d7:ff:bc	mac-authenticated	Yes		Active		0.0.0.0	admin
00:1a:1e:01:11:0d	00:1a:1e:01:11:0d	mac-auth-101	Yes		Active		0.0.0.0	admin
00:1a:1e:01:11:0e	00:1a:1e:01:11:0e	mac-auth-102	Yes		Active		0.0.0.0	admin
wireless1	*****	authenticated	Yes		Active		0.0.0.0	admin

Managing Internal Database Files

ArubaOS allows you to import and export tables of user information to and from the internal database. These files should not be edited once they are exported. ArubaOS only supports the importing of database files that were created during the export process. Note that importing a file into the internal database overwrite and removes all existing entries.

Using the CLI

Enter the following command in enable mode:

```
local-userdb export <filename>
local-userdb import <filename>
```

Internal Database Utilities

The local internal database also includes utilities to clear all users from the database and to restart the internal database to repair internal errors. Under normal circumstances, neither of these utilities are necessary.

Server Group Concepts

You can create *groups* of servers for specific types of authentication — for example, you can specify one or more RADIUS servers to be used for 802.1x authentication. You can configure servers of different types in one group — for example, you can include the internal database as a backup to a RADIUS server.

Configuring Server Groups

Server names are unique. You can configure the same server in more than one server group. The server must be configured before you can include it in a server group.

Using the CLI

```
aaa server-group <name>
auth-server <name>
```

Configuring Server List Order and Fail-Through

The list of servers in a server group is an ordered list. By default, the first server in the list is always used unless it is unavailable, in which case the next server in the list is used. You can configure the order of servers in the server group. In the CLI, use the **position** parameter to specify the relative order of servers in the list (the lowest value denotes the first server in the list).

As mentioned previously, the first available server in the list is used for authentication. If the server responds with an authentication failure, there is no further processing for the user or client for which the authentication request failed. You can optionally enable *fail-through* authentication for the server group so that if the first server in the list returns an authentication deny, the Mobility Access Switch attempts authentication with the next server in the ordered list. The Mobility Access Switch attempts authentication with each server in the list until either there is a successful authentication or the list of servers in the group is exhausted. This feature is useful in environments where there are multiple, independent authentication servers; users may fail authentication on one server but can be authenticated on another server.

Before enabling fail-through authentication, note the following:

- This feature is not supported for 802.1x authentication with a server group that consists of external EAP-compliant RADIUS servers. You can, however, use fail-through authentication when the 802.1x authentication is terminated on the Mobility Access Switch (AAA FastConnect).
- Enabling this feature for a large server group list may cause excess processing load on the Mobility Access Switch. Aruba recommends that you use server selection based on domain matching whenever possible (see [“Configuring Dynamic Server Selection”](#) on page 243).
- Certain servers, such as the RSA RADIUS server, lock out the Mobility Access Switch if there are multiple authentication failures. Therefore you should not enable fail-through authentication with these servers.

In the following example, you create a server group ‘corp-serv’ with two LDAP servers (ldap-1 and ldap-2), each of which contains a subset of the usernames and passwords used in the network. When fail-through authentication is enabled, users that fail authentication on the first server in the server list should be authenticated with the second server.

Using the CLI

```
aaa authentication-server ldap ldap-1
```

```

host 10.1.1.234
aaa authentication-server ldap ldap-2
host 10.2.2.234

aaa server-group corp-serv
  auth-server ldap-1 position 1
  auth-server ldap-2 position 2
  allow-fail-through

```

Configuring Dynamic Server Selection

The Mobility Access Switch can dynamically select an authentication server from a server group based on the user information sent by the client in an authentication request. For example, an authentication request can include client or user information in one of the following formats:

- <domain>\<user> — for example, corpnet.com\darwin
- <user>@<domain> — for example, darwin@corpnet.com
- host/<pc-name>.<domain> — for example, host/darwin-g.finance.corpnet.com (this format is used with 802.1x machine authentication in Windows environments)

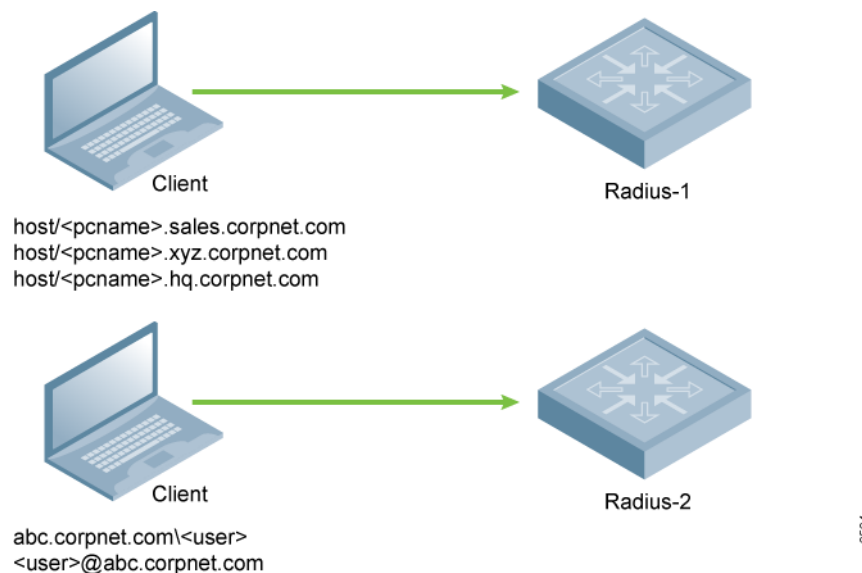
When you configure a server in a server group, you can optionally associate the server with one or more match rules. A match rule for a server can be one of the following:

- The server is selected if the client/user information *contains* a specified string.
- The server is selected if the client/user information *begins* with a specified string.
- The server is selected if the client/user information *exactly* matches a specified string.

You can configure multiple match rules for the same server. The Mobility Access Switch compares the client/user information with the match rules configured for each server, starting with the first server in the server group. If a match is found, the Mobility Access Switch sends the authentication request to the server with the matching rule. If no match is found before the end of the server list is reached, an error is returned and no authentication request for the client/user is sent.

For example, [Figure 4](#) depicts a network consisting of several subdomains in corpnet.com. The server radius-1 provides 802.1x machine authentication to PC clients in xyz.corpnet.com, sales.corpnet.com, and hq.corpnet.com. The server radius-2 provides authentication for users in abc.corpnet.com.

Figure 4 Domain-Based Server Selection Example



You configure the following rules for servers in the corp-serv server group:

- radius-1 will be selected if the client information starts with “host/”.
- radius-2 will be selected if the client information contains “abc.corpnet.com”.

Using the CLI

```
aaa server-group corp-serv
  auth-server radius-1 match-authstring starts-with host/ position 1
  auth-server radius-2 match-authstring contains abc.corpnet.com position 2
```

Trimming Domain Information from Requests

Before the Mobility Access Switch forwards an authentication request to a specified server, it can truncate the domain-specific portion of the user information. This is useful when user entries on the authenticating server do not include domain information. You can specify this option with any server match rule. This option is only applicable when the user information is sent to the Mobility Access Switch in the following formats:

- <domain>\<user> — the <domain>\ portion is truncated
- <user>@<domain> — the @<domain> portion is truncated



This option does not support client information sent in the format host/<pc-name>.<domain>

Using the CLI

```
aaa server-group corp-serv
  auth-server radius-2 match-authstring contains abc.corpnet.com trim-fqdn
```

Configuring Server-Derivation Rules

When you configure a server group, you can set the VLAN or role for clients based on attributes returned for the client by the server during authentication. The server derivation rules apply to all servers in the group. The user role or VLAN assigned through server derivation rules takes precedence over the default role and VLAN configured for the authentication method.



The authentication servers must be configured to return the attributes for the clients during authentication. For instructions on configuring the authentication attributes in a Windows environment using IAS, refer to the documentation at <http://technet2.microsoft.com/windowsserver/en/technologies/ias.mspx>.

The server rules are applied based on the first match principle. The first rule that is applicable for the server and the attribute returned is applied to the client and would be the only rule applied from the server rules. These rules are applied uniformly across all servers in the server group.

Table 6 describes the server rule parameters you can configure.

Table 6 Server Rule Configuration Parameters

Parameter	Description
Role or VLAN	The server derivation rules can be for either user role or VLAN assignment. With Role assignment, a client can be assigned a specific role based on the attributes returned. In case of VLAN assignment, the client can be placed in a specific VLAN based on the attributes returned.

Table 6 Server Rule Configuration Parameters (Continued)

Parameter	Description
Attribute	This is the attribute returned by the authentication server that is examined for <i>Operation</i> and <i>Operand</i> match.
Operation	<p>This is the match method by which the string in <i>Operand</i> is matched with the attribute value returned by the authentication server.</p> <ul style="list-style-type: none"> contains – The rule is applied if and only if the attribute value contains the string in parameter <i>Operand</i>. starts-with – The rule is applied if and only if the attribute value returned starts with the string in parameter <i>Operand</i>. ends-with – The rule is applied if and only if the attribute value returned ends with the string in parameter <i>Operand</i>. equals – The rule is applied if and only if the attribute value returned equals the string in parameter <i>Operand</i>. not-equals – The rule is applied if and only if the attribute value returned is not equal to the string in parameter <i>Operand</i>. value-of – This is a special condition. What this implies is that the role or VLAN is set to the value of the attribute returned. For this to be successful, the role and the VLAN ID returned as the value of the attribute selected must be already configured on the Mobility Access Switch when the rule is applied.
Operand	This is the string to which the value of the returned attribute is matched.
Value	The user role or the VLAN applied to the client when the rule is matched.
position	<p>Position of the condition rule. Rules are applied based on the first match principle. 1 is the top.</p> <p>Default: bottom</p>

Using the CLI

```

aaa server-group <name>
  auth-server <name>
  set {role|vlan} condition <condition> set-value {<role>|<vlan>}
    [position number]

```

Configuring a Role Derivation Rule for the Internal Database

When you add a user entry in the Mobility Access Switch's internal database, you can optionally specify a user role (see [“Internal Database Concepts” on page 240](#)). In order for the role specified in the internal database entry to be assigned to the authenticated client, you must configure a server derivation rule as shown in the following sections:

Using the CLI

```

aaa server-group internal
  set role condition Role value-of

```

Assigning Server Groups

You can create server groups for the following purposes:

- user authentication
- management authentication
- accounting

You can configure all types of servers for user and management authentication. However, TACACS+ is not supported for 802.1x authentication. For Accounting only RADIUS and TACACS+ servers are supported (see [Table 7](#)).

Table 7 *Server Types and Purposes*

	RADIUS	TACACS+	LDAP	Internal Database
User authentication	Yes	Yes (for MAC Authentication only)	Yes	Yes
Management authentication	Yes	Yes	Yes	Yes
Accounting	Yes	Yes	No	No

User Authentication

For information about assigning a server group for user authentication, see the configuration chapter for the authentication method.

Management Authentication

Users who need to access the Mobility Access Switch to monitor, manage, or configure the Aruba user-centric network can be authenticated with RADIUS, TACACS+, or LDAP servers or the internal database.



Only user record attributes are returned upon a successful authentication. Therefore, to derive a different management role other than the default mgmt auth role, set the server derivation rule based on the user attributes.

Using the CLI

```
aaa authentication mgmt
  server-group <group>
```

Radius Accounting

This section describes how user statistics are maintained and made available for RADIUS accounting. It contains the following sections:

- [“Understanding Radius Accounting” on page 246](#)
- [“Configuring RADIUS Accounting” on page 248](#)

Understanding Radius Accounting

RADIUS accounting supports sending user statistics in radius accounting stop and interim records. This document describes how user statistics are maintained and made available for RADIUS accounting.

When RADIUS accounting is enabled in the AAA profile, RADIUS accounting start and stop records are sent to the server. RADIUS accounting stop records contain received bytes and packet counters. The accounting start record is sent when a user authenticates. The stop record is sent when a user logs out or is deleted from the system. If interim accounting is enabled, updates are sent out at a fixed interval. Each interim record includes cumulative user statistics.

Currently, only received packets and bytes in accounting records are transmitted to the radius server.

User Activity and Statistics

RADIUS accounting allows user activity and statistics to be reported from the Mobility Access Switch to RADIUS servers. RADIUS accounting works as follows:

- The Mobility Access Switch generates an Accounting Start packet when a user logs in. The code field of transmitted RADIUS packet is set to 4 (Accounting-Request). Note that sensitive information, such as user passwords, are not sent to the accounting server. The RADIUS server sends an acknowledgement of the packet.
- The Mobility Access Switch sends an Accounting Stop packet when a user logs off; the packet information includes various statistics such as elapsed time, input and output bytes and packets. The RADIUS server sends an acknowledgement of the packet. The following is the list of attributes that the Mobility Access Switch can send to a RADIUS accounting server:
 - Acct-Status-Type:
This attribute marks the beginning or end of accounting record for a user. Currently, possible values include Start and Stop.
 - User-Name:
Name of user.
 - Acct-Session-Id:
A unique identifier to facilitate matching of accounting records for a user. It is derived from the user name, IP address and MAC address. This is set in all accounting packets.
 - Acct-Authentic:
This indicates how the user was authenticated. Current values are 1 (RADIUS), 2 (Local) and 3 (LDAP).
 - Acct-Session-Time:
The elapsed time, in seconds, that the client was logged in to the Mobility Access Switch. This is only sent in Accounting-Request records where the Acct-Status-Type is Stop.
 - Acct-Terminate-Cause:
Indicates how the session was terminated and is sent in Accounting-Request records where the Acct-Status-Type is Stop. Possible values are:
 - 1: User logged off
 - 4: Idle Timeout
 - 5: Session Timeout. Maximum session length timer expired.
 - 7: Admin Reboot: Administrator is ending service, for example prior to rebooting the Mobility Access Switch.
 - NAS-Identifier:
This is set in the RADIUS server configuration.
NAS-IP-Address: IP address of the master Mobility Access Switch. You can configure a “global” NAS IP address: in the WebUI, navigate to the Configuration > Security > Authentication > Advanced page; in the CLI, use the `ip radius nas-ip` command.
 - NAS-Port:
Physical or virtual port (tunnel) number through which the user traffic is entering the Mobility Access Switch.
 - NAS-Port-Type:
Type of port used in the connection. This is set to one of the following:
 - 5: admin login
 - 15: wired user type
 - 19: wireless user
 - Framed-IP-Address: IP address of the user.
 - Calling-Station-ID: MAC address of the user.

- Called-station-ID: MAC address of the Mobility Access Switch.

The following attributes are sent in Accounting-Request packets when Acct-Status-Type value is Start:

- Acct-Status-Type
- User-Name
- NAS-IP-Address
- NAS-Port
- NAS-Port-Type
- NAS-Identifier
- Framed-IP-Address
- Calling-Station-ID
- Called-station-ID
- Acct-Session-Id
- Acct-Authentic

The following attributes are sent in Accounting-Request packets when Acct-Status-Type value is Stop:

- Acct-Status-Type
- User-Name
- NAS-IP-Address
- NAS-Port
- NAS-Port-Type
- NAS-Identifier
- Framed-IP-Address
- Calling-Station-ID
- Called-station-ID
- Acct-Session-Id
- Acct-Authentic
- Terminate-Cause
- Acct-Session-Time

The following attributes are sent only in Accounting Stop packets (they are not sent in Accounting Start packets):

- Acct-Input-Octets
- Acct-Output-Octets
- Acct-Input-Packets
- Acct-Output-Packets

Configuring RADIUS Accounting

Radius accounting support is enabled and disabled in the AAA profile. By default, it is disabled.

To enable radius-accounting, use the command radius-accounting:

```
(host) #configure terminal
(host) (config)#aaa profile default
(host) (AAA Profile "default") #radius-accounting foobar
(host) (AAA Profile "default") #show aaa profile test

AAA Profile "TEST"
```


Parameter	Value
Initial role	logon
MAC Authentication Profile	N/A
MAC Authentication Default Role	guest
MAC Authentication Server Group	default
802.1X Authentication Profile	N/A
802.1X Authentication Default Role	guest
802.1X Authentication Server Group	N/A
Download Role from ClearPass	Enabled
L2 Authentication Fail Through	Disabled
RADIUS Accounting Server Group	foobar
RADIUS Interim Accounting	Disabled
XML API server	N/A
RFC 3576 server	N/A
User derivation rules	N/A
SIP authentication role	N/A
Enforce DHCP	Disabled
Authentication Failure Blacklist Time	3600 sec

To disable the feature, use the command `no radius-accounting`:

```
(host) (AAA Profile "default") #no radius-accounting
```

TACACS+ Accounting

TACACS+ accounting allows commands issued on the Mobility Access Switch to be reported to TACACS+ servers. You can specify the types of commands that are reported (action, configuration, or show commands) or have all commands reported.

Using the CLI

```
aaa tacacs-accounting server-group <group> command {action|all|configuration|show} mode {enable|disable}
```

Authentication Timers

[Table 8](#) describes the timers you can configure that apply to all clients and servers. These timers can be left at their default values for most implementations.

Table 8 *Authentication Timers*

Timer	Description
User Idle Timeout	<p>Maximum period after which a client is considered idle if there is no user traffic from the client.</p> <p>The timeout period is reset if there is a user traffic. If Mobility Access Switch does not see traffic from the user for more than the timeout period, then that user entry will be deleted from the system. If the keyword seconds is not specified, the value defaults to minutes at the command line.</p> <p>Range: 1 to 255 minutes (30 to 15300 seconds)</p> <p>Default: 5 minutes (300 seconds)</p>

Table 8 *Authentication Timers (Continued)*

Timer	Description
Authentication Server Dead Time	<p>Maximum period, in minutes, that the Mobility Access Switch considers an unresponsive authentication server to be “out of service”.</p> <p>This timer is only applicable if there are two or more authentication servers configured on the Mobility Access Switch. If there is only one authentication server configured, the server is never considered out of service and all requests are sent to the server.</p> <p>If one or more backup servers are configured and a server is unresponsive, it is marked as out of service for the dead time; subsequent requests are sent to the next server on the priority list for the duration of the dead time. If the server is responsive after the dead time has elapsed, it can take over servicing requests from a lower-priority server; if the server continues to be unresponsive, it is marked as down for the dead time.</p> <p>Range: 0–50 Default: 10 minutes</p>
Logon User Lifetime	<p>Maximum time, in minutes, unauthenticated clients are allowed to remain logged on.</p> <p>Range: 0–255 Default: 5 minutes</p>

Using the CLI

To set an authentication timer, use the following command:

```
aaa timers {dead-time <minutes>|idle-timeout <number>|logon-lifetime <minutes>}
```

This chapter describes AAA authentication. It contains the following major sections:

- “AAA Authentication Profile” on page 251
- “Configuring Authentication End to End” on page 256

AAA Authentication Profile

- “Authentication Profile Concepts” on page 251
- “Authentication Schemes” on page 252
- “Role/VLAN Derivation” on page 252
- “User Roles” on page 255
- “Authentication Roles” on page 255
- “User Derivation Rules” on page 256

Authentication Profile Concepts

The AAA profile can be applied on a global or per port or per VLAN basis, but only if the port is marked as un-trusted. If no AAA profile is configured on a port or a VLAN that the port is part of, the AAA profile configured under the wired authentication profile (aaa authentication wired) is applied globally by default.

AAA profile cannot be attached to an interface that is configured with a Tunneled Node profile.

If the port is marked as trusted, no authentication can be applied to traffic to the port.

The global AAA profile has limited ability to perform granular access control. The ability to apply an AAA profile on a per port/VLAN basis provides the administrator with greater flexibility and more granular access control. With per-port AAA profile, users can specify a unique AAA profile for each un-trusted port.

The AAA profile can be configured with the following parameters:

Initial Role

- The Initial Role is applied to all packets before a Layer 3 user entry is created.

MAC Auth Profile

- The MAC Auth Profile contains the MAC authentication profile parameters.

MAC Default Role

- The MAC Default Role is the default role a user receives upon successful MAC authentication.

802.1x Auth Profile

- The 802.1x Auth Profile contains the 802.1x authentication profile parameters.

802.1x Default Role

- The 802.1x Default Role is the default role a user receives upon successful 802.1x authentication.

User Derivation Rules

- The User Derivation Rules provide the means to derive a new VLAN or role, based on user attributes.

Authentication Schemes

The Mobility Access Switch supports the following authentication schemes:

- MAC Based Authentication
- 802.1X Authentication

MAC-Based Authentication

MAC-Based Authentication is a simple authentication method that is used more often as a filtering mechanism than as an actual authentication method. MAC-Based Authentication is frequently used when devices such as phones, printers, and scanners do not support 802.1x. It is also used in conjunction with 802.1x, so that the 802.1x authenticator and the back-end authentication server do not have to handle the load of authenticating users or devices that are not part of the back end database.

802.1x Authentication

802.1x authentication is a sophisticated method of network authentication that is widely supported across client OS and networking devices. This scheme provides a number of authentication methods, including PEAP and TLS. Both of these methods rely on TLS protocol to establish a secure tunnel to exchange user credentials, and authenticate the user. User validation can be done using a password or a certificate. The Mobility Access Switch supports using 802.1x authentications in the following modes:

- Authenticator Mode
- Authentication (EAP-Termination) Mode

Authenticator Mode

The authenticator mode is a generic method where the EAP frames from the user are packaged and sent to a RADIUS server. In the authentication server mode, also known as eap-termination mode, the controller can terminate the EAP frames to provide crypto hardware acceleration support to terminate the TLS tunnel. The controller dataplane terminates the phase 1 of the 802.1x authentication and provides with the TLS keys to the control plane to terminate the TLS tunnel. The phase 2 continues in the control plane with the user validation done using MSChapV2, PAP or Certification verification depending on the EAP mode the user was configured.

Authentication Server (EAP-Termination) Mode

In the authentication server mode, or eap-termination mode, the controller can terminate the EAP frames to provide crypto hardware acceleration support to terminate the TLS tunnel.

802.1x also supports key exchange in data encryption for wireless users. For wired users that are deployed today there is no key exchange and the security is limited to authenticating the user.

Role/VLAN Derivation

A user can be assigned a role/VLAN at different stages in its life cycle and the derivation can be done on various parameters. The precedence of the assignment is from 1 to 5 with 1 being the lowest and 5 being the highest. A user can be assigned a different role/VLAN in the following stages:

1. Initial Role/VLAN

This role is applied to the ingress on which the user traffic arrives. For wireless and tunneled-mode users, the ingress is a GRE tunnel and for wired users it is a port or VLAN. This role provides the means to control what kind of initial traffic is allowed, which is predominantly determined based on the allowed modes of authentication. There are cases where initial role is configured to deny all DHCP traffic so that the creation of the user happens after MAC based or 802.1x authentication is completed.

2. User Derived Role/VLAN

This role is only assigned based on the user MAC address. For this role derivation, user-derivation-rules must be defined and applied under the AAA profile.

3. Default Authentication Role/VLAN

This role is assigned when a user successfully completes a specific authentication type. Each authentication type can have a different role and this provision is defined in the AAA profile for Layer 2 authentication types. A VLAN can be configured under the default authentication role. This VLAN is assigned to the user after successful authentication. If a VLAN is not present under the user role, the client gets a default port based VLAN or VLAN derived via user derivation rule, server derivation rule or Vendor Specific Attribute.

4. Server Derived Role/VLAN

This role is derived from the attributes sent by the back-end authentication server. For this role to be applied, a set of “server derivation rules” must be defined under the server-group. The server group contains both the server definitions and the rules that are applied to the attributes returned from the list of servers.

5. Aruba VSA

Aruba Vendor Specific Attributes (VSA) override any of the above rules and derivations. If the back-end authentication server sends an VSA like Aruba-User-Role or Aruba-User-VLAN, the value of these attributes are sent to the user.

There are no rules that must be configured for this derivation to happen.



Roles and VLANs can be derived using VSA, but neither user role nor VLAN derivation is possible using two separate entries of VSA attributes under an IAS profile of the Windows authentication server.

Role Assignment Precedence

The precedence of role assignment in reducing order is as follows:

1. Vendor specific attribute (VSA) derived via Captive Portal authentication
2. Server derived via Captive Portal authentication
3. Default Captive Portal authentication
4. VSA derived via 802.1x authentication
5. Server derived via 802.1x authentication
6. Default 802.1x authentication
 - 802.1X authentication Default Role—Users get this role after successful machine (if it is enabled) and user authentication (username/password or certificates).
 - Machine authentication-Default User Role—Users get this role after a successful user authentication (username/password or certificates) and a failed machine authentication.
 - Machine Authentication-Default Machine Role—Users get this role after a successful machine authentication and a failed user authentication.
7. MAC authentication default role

8. Role derived via UDR matching the MAC address
9. AAA Profile Initial Role



If the “dhcp-option” based UDR is configured to derive a role and when the matching options are found in DHCP packets, it overrides all the above precedence. The client will get a VLAN configured under dhcp-option based UDR. If a VLAN is not configured, then the client will either stay in current VLAN or follow the VLAN assignment precedence. For more details, see [“VLAN Assignment Precedence:” on page 254](#).

VLAN Assignment Precedence:

The precedence of VLAN assignment in reducing order is given below:



No VLAN will be derived if Captive Portal authentication is successful. Any VLAN derived will be ignored after a successful Captive Portal authentication.

1. Explicit VSA derived via 802.1x authentication
2. VLAN configured under VSA derived 802.1x authentication role
3. Explicit server derived via 802.1x authentication
4. VLAN configured under server derived 802.1x authentication role
5. VLAN defined under the respective default authentication role
 - 802.1X authentication default role
 - Machine authentication—default user role
 - Machine authentication—default machine role
 - MAC authentication default role
6. Explicit UDR based on MAC address match to derive a VLAN
7. VLAN defined under UDR based on matching MAC address
8. VLAN defined under AAA profile initial role
9. Default VLAN assigned to the port



If the dhcp-option based UDR is configured to derive a VLAN and if matching options are found in DHCP packets, it overrides all the above precedence.

Current Limitations

- If the MAC authenticated client has received a VLAN via SDR or VSA and going further for successful 802.1x authentication, its VLAN is overwritten and client is assigned a new VLAN (precedence is based on points 1 to 9 above).
- SDR and VSA are not available for machine authentication.

Layer 2 Entry

Layer 2 user entry is created when the wired station connects to the network or when a Layer 2 “miss trigger” is sent to the control plane for a wired user. The Layer 2 user entry with 0.0.0.0 and MAC address is created both in the control plane and dataplane. The user entry inherits the initial role or the user derived role from the AAA profile. This user entry controls the Layer 2 traffic the user can send prior to getting an IP address. It also maintains the statistics for a given MAC address, assuming a user can potentially get multiple IP addresses. Location based ACLs are applied using the Layer 2 user entry.

Layer 3 Entry

After getting an IP address, the user entry shows up in the user table as “Layer 3 Entry.”

User Roles

User roles are a key component for role based policy enforcement.

Fully authenticated Layer 2 roles are assigned when a user has successfully completed all configured Layer 2 authentication methods.

The following authentication command is available in all roles:

```
reauthentication-interval <minutes>
  policer-profile < policer profile name>
  qos-profile <qos profile name>
  voip-profile <voip profile name>
```



For more detail, see the Chapter 15 Roles and Policies.

Authentication Roles

After authentication, the station or user is given a role that defines the behavior of the user. The role can be defined with the following:

- Access List
- VLAN
- Reauthentication Interval

Access List

This ACL is applied to the user. Three types of ACLs can be applied:

- Ether ACL
These access rules can be applied to specific Ether types.
- MAC ACL
These access rules are applied based on MAC address
- Layer 2 - 4
These access rules are applied based on Layer 3 and Layer 4 information such as IP-Address, protocol, and port.

VLAN

The VLAN attribute is set on initial roles or Layer 2 authenticated roles, so that the user ends on a new VLAN.

- Reauthentication Interval

This is defined in terms of minutes and is sometimes used to re-trigger authentication after a specified interval.

User Derivation Rules

This section contains the following sections:

- [“Configuring User Derivation Rules” on page 256](#)
- [“Displaying User Derivation Rules” on page 256](#)



DHCP Signature (DHCP-Option) is supported in addition to MAC Address-based UDRs.

Configuring User Derivation Rules

To configure user derivation rules, use the following command:

```
aaa derivation-rules user student
    set role condition macaddr equals "00:25:90:0a:95:d2" set-value student-role
    set vlan condition macaddr equals "00:25:90:0a:95:d2" set-value 202
```

Displaying User Derivation Rules

To display user derivation rules, use the following command:

```
(host)(config) #show aaa derivation-rules user udr_rule1
```

User Rule Table

Pr	Attribute	Operation	Operand	Action	Value	Total Hits	New Hits	Desc
1	macaddr	equals	00:aa:bb:cc:dd:e1	set role	authentic	0	0	
2	macaddr	equals	00:aa:bb:cc:dd:e2	set vlan	3912	0	0	

Rule Entries: 2

Configuring Authentication End to End

This section describes how to configure authentication end-to-end using the command-line interface. This section contains the following sections:

- [“Configuring Authentication Server” on page 256](#)
- [“Configuring Management Authentication” on page 258](#)
- [“Configuring AAA Timers” on page 258](#)

Configuring Authentication Server

Prior to configuring authentication, an authentication server must be defined. The Mobility Access Switch supports the following authentication server types: RADIUS, TACACS+, LDAP, and the Internal Database.



TACACS+ is not supported for 802.1X authentication.

Configuring a RADIUS Authentication Server

To configure a RADIUS authentication server, use the following commands:

```
(host)(config) #aaa authentication-server radius RADIUS1
(host)(RADIUS Server "RADIUS1") #host 10.20.20.200
(host)(RADIUS Server "RADIUS1") #key <shared-secret>
(host)(RADIUS Server "RADIUS1") #exit
```


Displaying the Authentication Server Configuration

To display the authentication server configuration for verification, use the following command:

```
(host) #show aaa authentication-server all
```

Auth Server Table

Name	Type	IP addr	AuthPort	AcctPort	Status	Requests
Internal	Local	172.16.0.254	n/a	n/a	Enabled	0
RADIUS1	Radius	10.20.20.200	1812	1813	Enabled	0

Configuring an Authentication Server Group



Authentication servers are referenced in server groups.

To configure the server in a server group, use the following commands:

```
(host) (config) #aaa server-group AUTH_SERVER
(host) (Server Group "AUTH_SERVER") #auth-server RADIUS1
(host) (Server Group "AUTH_SERVER") #exit
```

Configuring a Server for Fail-Over with the Internal Database

You can define multiple authentication servers for fail-over purposes. When you define multiple authentication servers, reference the servers in a single server-group.

```
(host) (config) #aaa server-group AUTH_SERVER
(host) (Server Group "AUTH_SERVER") #auth-server Internal
(host) (Server Group "AUTH_SERVER") #auth-server RADIUS2
```

Configuring Internal Server Under a Server-Group

To configure the internal database server, use the Internal keyword for the authentication-server, and the following commands:

```
(host) (config) #aaa server-group INTERNAL_SERVER
(host) (Server Group "INTERNAL_SERVER") #auth-server Internal
(host) (Server Group "INTERNAL_SERVER") #exit
```

Configuring a User Account with the Internal Database

To use the Internal Server, create a user account with the following command:

```
(host) #local-userdb add username <username> password <password> role dot1x-
authenticated
```

Displaying the Internal Database

To display the user database, use the following commands:

```
(host) # show local-userdb
```

User Summary

Name	Password	Role	E-Mail	Enabled	Expiry	Status	Sponsor-Name	Remote-IP	Grantor-Name
USER1	*****	guest		Yes		Active		0.0.0.0	admin

User Entries: 1

Maintaining Existing Accounts with the Internal Database

To add an existing user account, use the following command:

```
(host) #local-userdb add username labuser1 password abcdef
```

To modify an existing user account, use the following command:

```
(host) #local-userdb modify username USER1 role <ROLE>
```

To delete an existing user account, use the following command:

```
(host) #local-userdb del username USER1
```

To delete all existing user accounts, use the following command:

```
(host) #local-userdb del-all
```

Configuring Management Authentication

Similar to user/port authentication, management user can also be authenticated by using the AAA profile, such as using central authentication server for authenticating access to the network devices.

Authentication server can be the same server used for user authentication, or a separate server can be created for management authentication purpose. Similar to AAA authentication server configuration, the server needs to be defined first, then referenced on the server-group:

```
(host) (config) #aaa authentication-server tacacs TACACS1
(host) (TACACS Server "TACACS1") #host 10.20.20.202
(host) (TACACS Server "TACACS1") #key <shared-secret>
(host) (TACACS Server "TACACS1") #exit

(host) (config) #aaa server-group MGMT_AUTH_SERVER
(host) (Server Group "MGMT_AUTH_SERVER") #auth-server TACACS1
(host) (Server Group "MGMT_AUTH_SERVER") #exit
```

Once the server-group is defined (or used existing server-group), the AAA profile for management can be configured:

```
(host) (config) #aaa authentication mgmt
(host) (Management Authentication Profile) #enable
(host) (Management Authentication Profile) #server-group MGMT_AUTH_SERVER
(host) (Management Authentication Profile) #exit
```

Configuring AAA Timers

AAA timers such as dead-time, timeout for idle, as well as logon-lifetime can be defined at global level:

```
(host) (config) #aaa timers dead-time 10
(host) (config) #aaa timers idle-timeout 300
(host) (config) #aaa timers logon-lifetime 5
(host) (config) #aaa timers stats-timeout 300 seconds
```



Logon-lifetime is not applicable for 802.1x and MAC authentication as the user entry is deleted and the session is terminated when the idle-timeout hits.

Timers can be viewed using the following CLI command:

```
(host) #show aaa timers
User idle timeout = 300 seconds
Auth Server dead time = 10 minutes
Logon user lifetime = 5 minutes
User Interim stats frequency = 300 seconds
```

The idle-timeout is set to 5 minutes, which is the default.

Every client is associated with a user role, which determines the client's network privileges and how often it must re-authenticate. A *policy* is a set of rules that applies to traffic that passes through the ArubaOS Mobility Access Switch. You specify one or more policies for a user role. Finally, you can assign a user role to clients before or after they authenticate to the system.

This chapter describes assigning and creating roles and policies using the ArubaOS command line. This chapter describes the following topics:

- “Policies” on page 259
- “User Roles” on page 261
- “User Role Assignments” on page 262

Policies

An access policy identifies specific characteristics about a data packet passing through the Mobility Access Switch and takes some action based on that identification. In a Mobility Access Switch, that action can be a access -type action such as permitting or denying the packet, an administrative action such as logging the packet, or a quality of service (QoS) action such as setting 802.1p bits or placing the packet into a priority queue. You can apply access policies to user roles to give differential treatment to different users on the same network, or to physical ports to apply the same policy to all traffic through the port.

Policies differ from access control lists (ACLs) in the following ways:

- Policies are *stateless*, meaning that they statically evaluates packet contents. The traffic in the reverse direction will be allowed unconditionally.
- Access policies are *uni-directional*, meaning that they keep track of data connections traveling into or out of the network. ACLs are normally applied to traffic inbound to an interface.
- Access policies are *dynamic*, meaning that address information in the policy rules can change as the policies are applied to users.

Access Control Lists (ACLs)

- Ethertype ACLs are used to filter based on the Ethertype field in the frame header. Ethertype ACLs can be either named or numbered, with valid numbers in the range of 200-299. These ACLs can be used to permit IP while blocking other non-IP protocols, such as IPX or AppleTalk.
- MAC ACLs are used to filter traffic on a specific source MAC address or range of MAC addresses. MAC ACLs can be either named or numbered, with valid numbers in the range of 700-799 and 1200-1299.
- Standard ACLs permit or deny traffic based on the source IP address of the packet. Standard ACLs can be either named or numbered, with valid numbers in the range of 1-99 and 1300-1399. Standard ACLs use a bitwise mask to specify the portion of the source IP address to be matched.
- Extended ACLs permit or deny traffic based on source or destination IP address, source or destination port number, or IP protocol. Extended ACLs can be named or numbered, with valid numbers in the range 100-199 and 2000-2699.
- Stateless ACLs are used to define stateless packet filtering and quality of service (QoS). A stateless ACL statically evaluates packet contents. The traffic in the reverse direction will be allowed unconditionally. Stateless ACLs are named ACLs.

Mobility Access Switch provides both standard and extended ACLs for compatibility with router software from popular vendors, however access policies provide equivalent and greater functionality than standard and extended ACLs and should be used instead.

You can apply only the MAC, Ethertype and Stateless ACLs to a user role. The MAC and Ethertype ACLs only apply to non-IP traffic and the Stateless ACL to IP traffic *from* the user.

Creating an Access Policy

This section describes how to configure the rules that constitute an access policy. An access policy can then be applied to a user role (until the policy is applied to a user role, it does not have any effect).

[Table 1](#) describes required and optional parameters for a rule.

Table 1 Access Policy Rule Parameters

Field	Description
Source (required)	Source of the traffic, which can be one of the following: <ul style="list-style-type: none">• any: Acts as a wildcard and applies to any source address.• host: This refers to traffic from a specific host. When this option is chosen, you must configure the IP address of the host.• network: This refers to a traffic that has a source IP from a subnet of IP addresses. When this option is chosen, you must configure the IP address and network mask of the subnet.• alias: This refers to using an alias for a host or network.
Destination (required)	Destination of the traffic, which can be configured in the same manner as Source.
Service (required)	Type of traffic, which can be one of the following: <ul style="list-style-type: none">• any: This option specifies that this rule applies to any type of traffic.• tcp: Using this option, you configure a range of TCP port(s) to match for the rule to be applied.• udp: Using this option, you configure a range of UDP port(s) to match for the rule to be applied.• service: Using this option, you use one of the pre-defined services (common protocols such as HTTPS, HTTP, and others) as the protocol to match for the rule to be applied.• protocol: Using this option, you specify a different layer 4 protocol (other than TCP/UDP) by configuring the IP protocol value.
Action (required)	The action that you want the controller to perform on a packet that matches the specified criteria. This can be one of the following: <ul style="list-style-type: none">• permit: Permits traffic matching this rule.• drop: Drops packets matching this rule without any notification.
Log (optional)	Logs a match to this rule. This is recommended when a rule indicates a security breach, such as a data packet on a policy that is meant only to be used for voice calls.
policer-profile	To attach the policer-profile to the ACL.
position	Configure the ACL to display the filter position.
qos-profile	QoS profile can be configured to assign specific TC/DP, DSCP, and 802.1p values. This option attaches the qos-profile to the ACL .
Time Range (optional)	Time range for which this rule is applicable.

Table 1 Access Policy Rule Parameters (Continued)

Field	Description
Black List (optional)	Automatically blacklists a client that is the source or destination of traffic matching this rule. This option is recommended for rules that indicate a security breach where the blacklisting option can be used to prevent access to clients that are attempting to breach the security.

In the CLI

```
(host)(config) #ip access-list stateless STATELESS
(host)(config-stateless-STATELESS) #network 10.100.100.0 255.255.255.0 any tcp 8888
deny log
(host)(config-stateless-STATELESS) #any host 10.100.100.200 any deny log
(host)(config-stateless-STATELESS) #any any permit
(host)(config-stateless-STATELESS) #exit
```

Creating a Network Service Alias

A network service alias defines a TCP, UDP or IP protocol and a list or range of ports supported by that service. When you create a network service alias, you can use that alias when specifying the network service for multiple session ACLs.

In the CLI

To define a service alias via the command-line interface, access the CLI in config mode and issue the following command:

```
netSERVICE <name> <protocol>|tcp|udp {list <port>,<port>}|{<port> [<port>]}
[ALG <service>]
```

User Roles

This section describes how to create a new user role. When you create a user role, you specify one or more policies for the role. [Table 2](#) list the parameters you can configure for the user role.

Table 2 User Role Parameters

Field	Description
Access Policies (required)	One or more policies that define the privileges of a wired client in this role. There are three ways to add a access policy to a user role: <ul style="list-style-type: none"> Use an existing policy via CLI Edit and use the existing policy via CLI Create a new policy CLI Note: For more information, see “Configuring the ACLs.”
Re-authentication Interval (optional)	Time, in minutes, after which the client is required to reauthenticate. Enter a value between 0-4096. 0 disables reauthentication. Default: 0 (disabled)
Role VLAN ID (optional)	By default, a client is assigned a VLAN on the basis of the ingress VLAN for the client to the Mobility Access Switch. You can override this assignment and configure the VLAN ID that is to be assigned to the user role. You configure a VLAN by navigating to the Configuration > VLANs page.

Table 2 *User Role Parameters (Continued)*

Field	Description
policer-profile (optional)	Specifies the policer activities configuration parameters for the user under this role.
qos-profile (optional)	Specifies the QoS configuration parameters for the user under this role.
voip-profile (optional)	Specifies the VOIP configuration parameters for an user connected to the interface (VOIP devices and/or PCs and Laptops).

Creating a User Role

The following example creates the user role ‘web-guest’ and assigns the previously-configured ‘web-only’ policy to this user role.



You cannot delete a user-role that is referenced in a **aaa-profile**. Remove all references to the role and then perform the delete operation. Deleting user-roles used by external authentication servers is also inadvisable without first modifying the external authentication server not to reference that role.

In the CLI

```
user-role web-guest
  access-list stateless web-only position 1
```

After assigning the user role, you can use the **show reference user-role <role>** command to see the profiles that reference this user role.

User Role Assignments

A client is assigned a user role by one of several methods. A role assigned by one method may take precedence over one assigned by a different method. The methods of assigning user roles are, from lowest to highest precedence:

1. The user role can be derived from user attributes upon the client’s association with an interface (this is known as a *user-derived role*). You can configure rules that assign a user role to clients that match the mac address. For example, you can configure a rule to assign the role “VoIP-Phone” to any client that has a MAC address that starts with bytes xx:yy:zz. User-derivation rules are executed before client authentication.
2. The user role can be the default user role configured for an authentication method, such as 802.1x or MAC authentication. For each authentication method, you can configure a default role for clients who are successfully authenticated using that method.
3. The user role can be derived from attributes returned by the authentication server (this is known as a *server-derived role*). If the client is authenticated via an authentication server, the user role for the client can be based on the attribute returned by the server during authentication. In case the attribute is not returned by the server, the client gets the default authentication role defined under aaa profile. Server-derivation rules are executed after client authentication.
4. The user role can be derived from Aruba Vendor-Specific Attributes (VSA) for RADIUS server authentication. A role derived from an Aruba VSA takes precedence over any other user roles.

The following sections describe the methods of assigning user roles.

User Role in AAA Profile

An AAA profile defines the user role for unauthenticated clients (initial role) as well as the default user role for MAC and 802.1x authentication. To configure user roles in the AAA profile:

In the CLI

```
aaa profile <profile>
  initial-role <role>
  dot1x-default-role <role>
  mac-default-role <role>
```

User-Derived Roles or VLANs

Attributes derived from the client's can be used to assign the client to a specific role or VLAN, as user-derivation rules are executed before the client is authenticated.

You configure the user role or VLAN to be assigned to the client by specifying condition rules; when a condition is met, the specified user role or VLAN is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can optionally add a description of the user rule.

[Table 3](#) describes the conditions for which you can specify a user role or VLAN.

Table 3 Conditions for a User-Derived Role or VLAN

Rule Type	Condition	Value
DHCP-Option	One of the following: <ul style="list-style-type: none">• equals• starts with	DHCP signature ID. Note: This string is not case sensitive.
MAC address of the client	One of the following: <ul style="list-style-type: none">• contains• ends with• equals• does not equal• starts with	MAC address (xx:xx:xx:xx:xx:xx)

Configure a User-derived Role or VLAN in the CLI

```
aaa derivation-rules user <name>
  set role|vlan
  condition macaddr
  contains|ends-with|equals|not-equals|starts-with <string>
  set-value <role>
  position <number>
```



There are many online tools available for converting ASCII text to a hexadecimal string.

Default Role for Authentication Method

For each authentication method, you can configure a default role for clients who are successfully authenticated using that method. To configure a default role for an authentication method:

In the CLI

To configure the default user role for MAC or 802.1x authentication:

```
aaa profile <profile>
  mac-default-role <role>
  dot1x-default-role <role>
```

Server-Derived Role

If the client is authenticated via an authentication server, the user role for the client can be based on one or more attributes returned by the server during authentication. You configure the user role to be derived by specifying condition rules; when a condition is met, the specified user role is assigned to the client. You can specify more than one condition rule; the order of rules is important as the first matching condition is applied. You can also define server rules based on client MAC address, even though the MAC address is not returned by the server as an attribute.



The roles and VLANs in the sample below are defined under the **aaa server-group <server-group-name>** configuration.

Sample configuration

```
set role|vlan
  condition <attribute name>
  contains|ends-with|equals|not-equals|starts-with <attribute value>
  set-value <role> | <vlan>
  position <number>
```

VSA-Derived Role

Many Network Address Server (NAS) vendors, including Aruba, use VSAs to provide features not supported in standard RADIUS attributes. For Aruba systems, VSAs can be employed to provide the user role and VLAN for RADIUS-authenticated clients, however the VSAs must be present on your RADIUS server. This involves defining the vendor (Aruba) and/or the vendor-specific code (14823), vendor-assigned attribute number, attribute format (such as string or integer), and attribute value in the RADIUS dictionary file. VSAs supported on Mobility Access Switches conform to the format recommended in RFC 2865, “Remote Authentication Dial In User Service (RADIUS)”.

This chapter describes the following topics:

- “MAC-Based Authentication Concepts” on page 265
- “Configuring MAC-Based Authentication” on page 265
- “Configuring Clients” on page 266

MAC-Based Authentication Concepts

MAC-based authentication is used to authenticate devices based on their physical media access control (MAC) address. While not the most secure and scalable method, MAC-based authentication implicitly provides an additional layer of security authentication devices. MAC-based authentication is often used to authenticate and allow network access through certain devices while denying access to the rest. For example, if clients are allowed access to the network via station A, then one method of authenticating station A is MAC-based. Clients may be required to authenticate themselves using other methods depending on the network privileges required.

Configuring MAC-Based Authentication

This section describes how to configure MAC-based authentication on the Mobility Access Switch. Before configuring MAC-based authentication, you must configure:

- The user role that will be assigned as the default role for the MAC-based authenticated clients.
- You configure the default user role for MAC-based authentication in the AAA profile. If derivation rules exist or if the client configuration in the internal database has a role assignment, these values take precedence over the default user role.
- The authentication server group that the Mobility Access Switch uses to validate the clients. The internal database can be used to define clients for MAC-based authentication.

Configuring the MAC Authentication Profile

Table 1 describes the MAC-based authentication parameters.

Table 1 MAC Authentication Profile Configuration Parameters

Parameter	Description
Delimiter	Delimiter used in the MAC string: <ul style="list-style-type: none">• colon specifies the format xx:xx:xx:xx:xx:xx• dash specifies the format xx-xx-xx-xx-xx-xx• none specifies the format xxxxxxxxxxxx• oui-nic specifies the format xxxxxx-xxxxxx Default: none
Case	The case (upper or lower) used in the MAC string. Default: lower

Table 1 MAC Authentication Profile Configuration Parameters (Continued)

Parameter	Description
Max Authentication failures	Number of times a station can fail to authenticate before it is blacklisted. A value of 0 disables blacklisting. Default: 0

Using the CLI

```
aaa authentication mac <profile>
  case {lower|upper}
  delimiter {colon|dash|none|oui-nic}
  max-authentication-failures <number>
```

Configuring Clients

You can create entries in the Mobility Access Switch's internal database that can be used to authenticate client MAC addresses. The internal database contains a list of clients along with the password and default role for each client. To configure entries in the internal database for MAC authentication, you enter the MAC address for both the user name and password for each client.



You must enter the MAC address using the delimiter format configured in the MAC authentication profile. The default delimiter is none, which means that MAC addresses should be in the format xxxxxxxxxxxx. If you specify colons for the delimiter, you can enter MAC addresses in the format xx:xx:xx:xx:xx:xx.

Using the CLI to configure clients in the internal database

Enter the following command in enable mode:

```
local-userdb add username <macaddr> password <macaddr>...
```

This chapter describes the following topics:

- “802.1x Authentication Concepts” on page 267
- “Configuring 802.1x Authentication” on page 269
- “Configuring 802.1x Authentication with Machine Authentication” on page 271

802.1x Authentication Concepts

IEEE 802.1x is an IEEE Standard for port-based Network Access Control (PNAC). It is part of the IEEE 802.1x group of networking protocols. It provides an authentication mechanism to devices wishing to attach to a LAN or WLAN.

802.1x authentication involves three parties:

- The *supplicant*, or client, is the device attempting to gain access to the network. You can configure the Aruba user-centric network to support 802.1x authentication for wired users.
- The *authenticator* is the gatekeeper to the network and permits or denies access to the supplicants. The Aruba Mobility Access Switch acts as the authenticator, relaying information between the authentication server and supplicant. The EAP type must be consistent between the authentication server and supplicant and is transparent to the Mobility Access Switch.
- The *authentication server* provides a database of information required for authentication and informs the authenticator to deny or permit access to the supplicant.

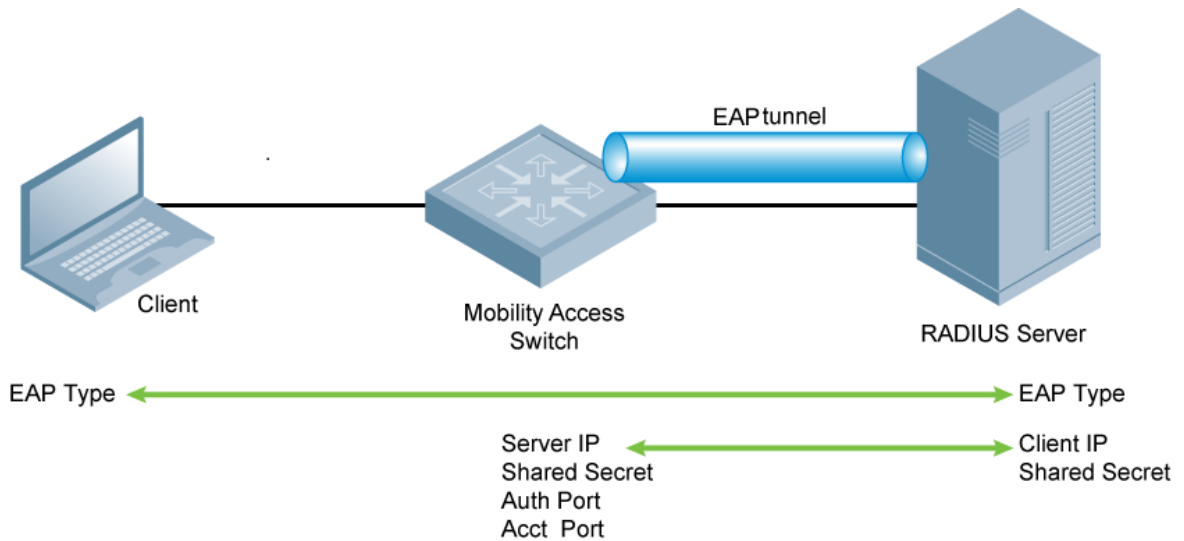
The 802.1x authentication server is typically an EAP-compliant Remote Access Dial-In User Service (RADIUS) server which can authenticate either users (through passwords or certificates) or the client computer.

In Aruba user-centric networks, you can terminate the 802.1x authentication on the Mobility Access Switch. The Mobility Access Switch passes user authentication to its internal database or to a “backend” non-802.1x server. This feature is useful for deployments where an 802.1x EAP-compliant RADIUS server is not available or required for authentication.

Authentication with a RADIUS Server

See [Table 1](#) below for an overview of the parameters that you need to configure on authentication components when the authentication server is an 802.1x EAP-compliant RADIUS server.

Figure 1 802.1x Authentication with RADIUS Server



The supplicant and authentication server must be configured to use the same EAP type. The Mobility Access Switch does not need to know the EAP type used between the supplicant and authentication server.

For the Mobility Access Switch to communicate with the authentication server, you must configure the IP address, authentication port, and accounting port of the server on the Mobility Access Switch. The authentication server must be configured with the IP address of the RADIUS client, which is the Mobility Access Switch in this case. Both the Mobility Access Switch and the authentication server must be configured to use the same shared secret.




Additional information on EAP types supported in a Windows environment, Microsoft supplicants, and authentication server, is available at [http://technet.microsoft.com/en-us/library/cc782851\(WS.10\).aspx](http://technet.microsoft.com/en-us/library/cc782851(WS.10).aspx).

The client communicates with the Mobility Access Switch through an EAP tunnel in order to authenticate to the network. Therefore, the network authentication and encryption configured must be the same on both the client and the Mobility Access Switch.

Authentication Terminated on the Mobility Access Switch

User authentication is performed either via the Mobility Access Switch's internal database or a non-802.1x server.

Figure 2 802.1x Authentication with Termination on Mobility Access Switch



In this scenario, the supplicant is configured for EAP-Transport Layer Security (TLS) or EAP-Protected EAP (PEAP).

- EAP-TLS is used with smart card user authentication. A smart card holds a digital certificate which, with the user-entered personal identification number (PIN), allows the user to be authenticated on the network. EAP-TLS relies on digital certificates to verify the identities of both the client and server.

EAP-TLS requires that you import server and certification authority (CA) certificates onto the Mobility Access Switch. The client certificate is verified on the Mobility Access Switch (the client certificate must be signed by a known CA) before the user name is checked on the authentication server.

- EAP-PEAP uses TLS to create an encrypted tunnel. Within the tunnel, one of the following “inner EAP” methods is used:
 - EAP-Generic Token Card (GTC): Described in RFC 2284, this EAP method permits the transfer of unencrypted usernames and passwords from client to server. The main uses for EAP-GTC are one-time token cards such as SecureID and the use of an LDAP or RADIUS server as the user authentication server. You can also enable caching of user credentials on the Mobility Access Switch as a backup to an external authentication server.
 - EAP-Microsoft Challenge Handshake Authentication Protocol version 2 (MS-CHAPv2): Described in RFC 2759, this EAP method is widely supported by Microsoft clients. A RADIUS server must be used as the backend authentication server.

If you are using the Mobility Access Switch’s internal database for user authentication, you need to add the names and passwords of the users to be authenticated. If you are using an LDAP server for user authentication, you need to configure the LDAP server on the Mobility Access Switch, and configure user IDs and passwords. If you are using a RADIUS server for user authentication, you need to configure the RADIUS server on the Mobility Access Switch.

Configuring 802.1x Authentication

The Mobility Access Switch supports 802.1x (dot1x) authentication including **termination**. For example, the list of termination options for the profile name *techpubsAuth* is shown below.

```
(host) (802.1X Authentication Profile "techpubsAuth") # termination ?
eap-type          Configure the EAP method.Default method is EAP-PEAP
enable            Enable Dot1x Termination.Default is disabled
enable-token-caching  Enable Token Caching.Default is disabled
inner-eap-type     Configure the inner EAP method.Default method is
                  EAP-MSCHAPV2
token-caching-period  Configure the Token Caching Period
```

The following example configures various options for the 802.1x Authentication profile *techpubsAuth*.

```
(host) (802.1X Authentication Profile "techpubsAuth") #termination enable
(host) (802.1X Authentication Profile "techpubsAuth") #termination eap-type eap-peap
(host) (802.1X Authentication Profile "techpubsAuth") #max-authentication-failures 2
(host) (802.1X Authentication Profile "techpubsAuth") #timer reauth-period 3600
(host) (802.1X Authentication Profile "techpubsAuth") #framed-mtu 1500
(host) (802.1X Authentication Profile "techpubsAuth") #reauth-max 2
(host) (802.1X Authentication Profile "techpubsAuth") #reauthentication
```

To verify the above configurations, execute the show command below:

```
(host) (config) #show aaa authentication dot1x techpubsAuth

802.1X Authentication Profile "techpubsAuth"
-----
Parameter                               Value
-----
Max authentication failures              2          <--
```

Enforce Machine Authentication	Disabled	
Machine Authentication: Default Machine Role	guest	
Machine Authentication Cache Timeout	24 hr(s)	
Blacklist on Machine Authentication Failure	Disabled	
Machine Authentication: Default User Role	guest	
Interval between Identity Requests	30 sec	
Quiet Period after Failed Authentication	30 sec	
Reauthentication Interval	3600 sec	<--
Use Server provided Reauthentication Interval	Disabled	
Authentication Server Retry Interval	30 sec	
Authentication Server Retry Count	2	
Framed MTU	1500 bytes	<--
Number of times ID-Requests are retried	3	
Maximum Number of Reauthentication Attempts	2	<--
Maximum number of times Held State can be bypassed	0	
Reauthentication	Enabled	<--
Termination	Enabled	<--
Termination EAP-Type	eap-peap	<--
Termination Inner EAP-Type	N/A	
Enforce Suite-B 128 bit or more security level Authentication	Disabled	
Enforce Suite-B 192 bit security level Authentication	Disabled	
Token Caching	Disabled	
Token Caching Period	24 hr(s)	
CA-Certificate	N/A	
Server-Certificate	N/A	
TLS Guest Access	Disabled	
TLS Guest Role	guest	
Ignore EAPOL-START after authentication	Disabled	
Handle EAPOL-Logoff	Disabled	
Ignore EAP ID during negotiation.	Disabled	
Check certificate common name against AAA server	Enabled	



Use the privileged mode in the CLI to configure users in the Mobility Access Switch's internal database.

To add users to the local database, use the following command:

```
local-userdb add username <user> password <password> role <user_role>
```

Configuring a Server Rule Using the CLI

```
aaa server-group dot1x_internal
set role condition Role value-of
```

LDAP Servers

If you are using a LDAP server for authentication, the following variables should be set.

- termination enabled
- EAP type of TLS or PEAP (with inner-EAP-type set to GTC)

Below is an example configuration for the profile *techpubsAuth* for an LDAP server:

```
(host) (802.1X Authentication Profile "techpubsAuth") #termination enable
(host) (802.1X Authentication Profile "techpubsAuth") #termination eap-type eap-peap
(host) (802.1X Authentication Profile "techpubsAuth") # termination inner-eap-type
eap-gtc
```

To verify the configuration, execute the **show aaa authentication dot1x profile_name** command.

Configuring Certificates with Auth Termination

The Mobility Access Switch supports 802.1x authentication using digital certificates for auth termination.

- **Server Certificate**—A server certificate installed in the Mobility Access Switch verifies the authenticity of the Mobility Access Switch for 802.1x authentication. Aruba Mobility Access Switches ship with a demonstration digital certificate. Until you install a customer-specific server certificate in the Mobility Access Switch, this demonstration certificate is used by default for all secure HTTP connections and auth termination. This certificate is included primarily for the purposes of feature demonstration and convenience and is not intended for long-term use in production networks. Users in a production environment are urged to obtain and install a certificate issued for their site or domain by a well-known certificate authority (CA). You can generate a Certificate Signing Request (CSR) on the Mobility Access Switch to submit to a CA. For information on how to generate a CSR and how to import the CA-signed certificate into the Mobility Access Switch, see [“Managing Certificates” on page 35](#)
- **Client Certificates**—Client certificates are verified on the Mobility Access Switch (the client certificate must be signed by a known CA) before the user name is checked on the authentication server. To use client certificate authentication for auth termination you need to import the following certificates into the Mobility Access Switch (see [“Importing Certificates” on page 37](#)):
 - Mobility Access Switch’s server certificate
 - CA certificate for the CA that signed the client certificates

Using the CLI

```
aaa authentication dot1x <profile>
    termination enable
    server-cert <certificate>
    ca-cert <certificate>
```

Configuring 802.1x Authentication with Machine Authentication

When a Windows device boots, it logs onto the network domain using a machine account. Within the domain, the device is authenticated before computer group policies and software settings can be executed; this process is known as *machine authentication*. Machine authentication ensures that only authorized devices are allowed on the network.

You can configure 802.1x for both user and machine authentication (select the **Enforce Machine Authentication** option described in [Table 1 on page 272](#)). This tightens the authentication process further since both the device and user need to be authenticated.

Role Assignment with Machine Authentication Enabled

When you enable machine authentication, there are two additional roles you can define in the 802.1x authentication profile:

- Machine authentication default machine role
- Machine authentication default user role

While you can select the same role for both options, you should define the roles as per the policies that need to be enforced. Also, these roles can be different from the 802.1x authentication default role configured in the AAA profile.

With machine authentication enabled, the assigned role depends upon the success or failure of the machine and user authentications. In certain cases, the role that is ultimately assigned to a client can also depend

upon attributes returned by the authentication server or server derivation rules configured on the Mobility Access Switch.

Table 1 describes role assignment based on the results of the machine and user authentications.

Table 1 *Role Assignment for User and Machine Authentication*

Machine Auth Status	User Auth Status	Description	Role Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	Initial role defined in the AAA profile will be assigned. If no initial role is explicitly defined, the default initial role (logon role) is assigned.
Failed	Passed	Machine authentication fails (for example, the machine information is not present on the server) and user authentication succeeds. Server-derived roles do not apply.	Machine authentication default user role configured in the 802.1x authentication profile.
Passed	Failed	Machine authentication succeeds and user authentication has not been initiated. Server-derived roles do not apply.	Machine authentication default machine role configured in the 802.1x authentication profile.
Passed	Passed	Both machine and user are successfully authenticated. If there are server-derived roles, the role assigned via the derivation take precedence. This is the <i>only</i> case where server-derived roles are applied.	A role derived from the authentication server takes precedence. Otherwise, the 802.1x authentication default role configured in the AAA profile is assigned.

For example, if the following roles are configured:

- 802.1x authentication default role (in AAA profile): dot1x_user
- Machine authentication default machine role (in 802.1x authentication profile): dot1x_mc
- Machine authentication default user role (in 802.1x authentication profile): guest

Role assignments would be as follows:

- If both machine and user authentication succeed, the role is dot1x_user. If there is a server-derived role, the server-derived role takes precedence.
- If only machine authentication succeeds, the role is dot1x_mc.
- If only user authentication succeeds, the role is guest.
- On failure of both machine and user authentication, the initial role defined in the AAA profile is assigned.

With machine authentication enabled, the VLAN to which a client is assigned (and from which the client obtains its IP address) depends upon the success or failure of the machine and user authentications. The VLAN that is ultimately assigned to a client can also depend upon attributes returned by the authentication server or server derivation rules configured on the Mobility Access Switch. If machine authentication is successful, the client is associated to the VLAN configured on the interface. However, the client can be assigned a derived VLAN upon successful user authentication.



You can optionally assign a VLAN as part of a user role configuration. It is recommended not to use VLAN derivation if user roles are configured with VLAN assignments.

Table 2 describes VLAN assignment based on the results of the machine and user authentications when VLAN derivation is used.

Table 2 VLAN Assignment for User and Machine Authentication

Machine Auth Status	User Auth Status	Description	VLAN Assigned
Failed	Failed	Both machine authentication and user authentication failed. L2 authentication failed.	VLAN configured on the interface or, VLAN configured under initial role
Failed	Passed	Machine authentication fails (for example, the machine information is not present on the server) and user authentication succeeds.	VLAN configured on the interface or, VLAN configured under Machine authentication default user role
Passed	Failed	Machine authentication succeeds and user authentication has not been initiated.	VLAN configured on the interface or, VLAN configured under Machine authentication default machine role
Passed	Passed	Both machine and user are successfully authenticated.	Derived VLAN or, VLAN configured on the interface

Authentication with an 802.1x RADIUS Server

- An EAP-compliant RADIUS server provides the 802.1x authentication. The RADIUS server administrator must configure the server to support this authentication. The administrator must also configure the server to all communications with the Aruba Mobility Access Switch.
- 802.1x authentication based on PEAP with MS-CHAPv2 provides both computer and user authentication. If a user attempts to log in without the computer being authenticated first, the user is placed into a more limited “guest” user role.

Windows domain credentials are used for computer authentication, and the user’s Windows login and password are used for user authentication. A single user sign-on facilitates both authentication to the network and access to the Windows server resources.

You can create the following policies and user roles for:

- Student
- Faculty
- Guest
- Sysadmin
- Computer

Creating an Alias for the Internal Network

Using the CLI

```
netdestination "Internal Network"  
  network 10.0.0.0 255.0.0.0  
  network 172.16.0.0 255.255.0.0
```

Creating the Student Role and Policy

The **student** policy prevents students from using telnet, POP3, FTP, SMTP, SNMP, or SSH to the wired portion of the network. The **student** policy is mapped to the **student** user role.

Using the CLI

```
ip access-list stateless student  
  any alias "Internal Network" svc-telnet deny  
  any alias "Internal Network" svc-pop3 deny  
  any alias "Internal Network" svc-ftp deny  
  any alias "Internal Network" svc-smtp deny  
  any alias "Internal Network" svc-snmp deny  
  any alias "Internal Network" svc-ssh deny  
  
user-role student  
access-list stateless student  
access-list stateless allowall
```

Creating the Faculty Role and Policy

The **faculty** policy is similar to the **student** policy. However, the faculty members are allowed to use POP3 and SMTP. The **faculty** policy is mapped to the **faculty** user role.

Using the CLI

```
ip access-list stateless faculty  
  any alias "Internal Network" svc-telnet deny  
  any alias "Internal Network" svc-ftp deny  
  any alias "Internal Network" svc-snmp deny  
  any alias "Internal Network" svc-ssh deny  
  
user-role faculty  
access-list stateless faculty  
access-list stateless allowall
```

Creating the Guest Role and Policy

The **guest** policy permits only access to the Internet (via HTTP or HTTPS) and only during daytime working hours. The **guest** policy is mapped to the **guest** user role.

Using the CLI

```
time-range working-hours periodic  
  weekday 07:30 to 17:00  
  
ip access-list stateless guest  
  any host 10.1.1.25 svc-dhcp permit time-range working-hours  
  any host 10.1.1.25 svc-dns permit time-range working-hours  
  any alias "Internal Network" any deny  
  any any svc-http permit time-range working-hours  
  any any svc-https permit time-range working-hours  
  any any any deny
```

```
user-role guest
access-list stateless guest
```

Configuring the RADIUS Authentication Server

You can set the role condition to identify the user's group. The Mobility Access Switch uses the literal value of this attribute to determine the role name. The following example uses the RADIUS server name *radiusTechPubs* to configure the Radius server.

```
(host) (config) #aaa authentication-server radius radiusTechPubs
(host) (RADIUS Server "radiusTechPubs") #host 10.41.255.30
(host) (RADIUS Server "radiusTechPubs") #key hometown
(host) (RADIUS Server "radiusTechPubs") #exit

(host) (config) #aaa server-group radiusTechpubs
(host) (Server Group "radiusTechpubs") #auth-server radiusTechpubs
(host) (Server Group "radiusTechpubs") #set role condition Class Value-of
```

Configuring 802.1x Authentication Profile

In the 802.1x authentication profile, configure enforcement of machine authentication before user authentication. If a user attempts to log in without machine authentication taking place first, the user is placed in the limited guest role.

Using the CLI

```
aaa authentication dot1x dot1x
  machine-authentication enable
  machine-authentication machine-default-role student
  machine-authentication user-default-role guest
```

Configuring AAA Profile

A AAA profile specifies the 802.1x authentication profile and 802.1x server group to be used for authenticating clients. The AAA profile also specifies the default user roles for 802.1x authentication.

Using the CLI

```
aaa profile aaa_dot1x
  dot1x-default-role faculty
  authentication-dot1x dot1x
  dot1x-server-group radiusTechpubs
```


Captive portal is an L3 authentication method supported by Mobility Access Switch. A captive portal presents a web page which requires user action before network access is granted. The required action can be simply viewing and agreeing to an acceptable use policy, entering Email ID, or entering a user ID and password which must be validated against a database of authorized users. The Mobility Access Switch supports both internal and external captive portals.

This chapter describes the following topics:

- “Captive Portal Overview” on page 277
- “Configuring Captive Portal Authentication” on page 277
- “Captive Portal Configuration Example” on page 279
- “Personalizing the Captive Portal Page” on page 281
- “Creating Walled Garden Access” on page 283
- “Mobility Access Switch Server Certificate” on page 284

Captive Portal Overview

You can configure captive portal for guest users where no authentication is required, or for registered users who must be authenticated against an external authentication server or the Mobility Access Switch’s internal user database.



Captive portal is most often used for guest access, access to open systems (such as public hot spots), or as a way to connect to a VPN.

You can use captive portal for guest and registered users at the same time. The default captive portal web page provided with ArubaOS Mobility Access Switch displays login prompts only for registered users. The Mobility Access Switch supports the creation of 16 different customer login pages. The login page displayed is based on the AAA Profile applied to the port that the user is connected.

Configuring Captive Portal Authentication

This section describes how to configure Captive Portal authentication on the Mobility Access Switch. Before configuring Captive Portal authentication, you must configure the following:

- The user role that will be assigned as the initial role. This initial role does not require any Captive Portal specific ACLs because once Captive Portal is added to the user-role, the necessary ACLs will automatically be added.
- The authentication server group that the Mobility Access Switch uses to validate the guest or registered users. The internal user database or an external authentication server may be used.



A read-only ACL using the same name defined in `captive-portal <name>` is automatically generated upon adding `captive-portal <name>` to a user-role. This ACL is configured to redirect http/https traffic and permit DNS and DHCP traffic. You can use the `show rights <user-role>` command to verify this ACL.

Captive Portal Configuration Parameters

Table 3 describes configuration parameters for Captive Portal Authentication profile page in the WebUI. In the CLI, you configure these options with the `aaa authentication captive-portal` commands.

Table 3 *Captive Portal Authentication Profile Parameters*

Parameter	Description
default-guest-role	Role assigned to guest. Default: guest
default-role	Role assigned to the Captive Portal user upon login. When both user and guest logon are enabled, the default role applies to the user logon; users logging in using the guest interface are assigned the guest role. Default: guest
enable-welcome-page	Displays the configured welcome page before the user is redirected to their original URL. If this option is disabled, redirection to the web URL happens immediately after the user logs in. Default: Enabled
guest-logon	Enables Captive Portal logon without authentication. Default: Disabled
ip-addr-in-redirection-url	Sends IP address of one of the interface in the redirection URL when external captive portal servers are used. Default: Disabled
login-page	URL of the page that appears for the user logon. This can be set to any URL. Default: /auth/index.html
logon-wait	Configure parameters for the logon wait interval Default: 10 seconds
Logon wait CPU utilization threshold	CPU utilization percentage above which the Logon wait interval is applied when presenting the user with the logon page. Default: 60%
Logon wait minimum wait	Minimum time, in seconds, the user will have to wait for the logon page to pop up if the CPU load is high. This works in conjunction with the Logon wait CPU utilization threshold parameter. Default: 5 seconds
logout-popup-window	Enables a pop-up window with the Logout link for the user to logout after logon. If this is disabled, the user remains logged in until the user timeout period has elapsed or the station reloads. Default: Enabled
max-authentication-failures	The number of authentication failures before the user is blacklisted. Default : 0, Range: 0-10
protocol-http	Use HTTP protocol on redirection to the Captive Portal page. If you use this option, modify the captive portal policy to allow HTTP traffic. Default: disabled (HTTPS is used)
redirect-pause	Time, in seconds, that the system remains in the initial welcome page before redirecting the user to the final web URL. If set to 0, the welcome page displays until the user clicks on the indicated link. Default: 10 seconds
server-group	Name of the group of servers used to authenticate Captive Portal users.

Table 3 *Captive Portal Authentication Profile Parameters (Continued)*

Parameter	Description
show-fqdn	Allows the user to see and select the fully-qualified domain name (FQDN) on the login page. The FQDNs shown are specified when configuring individual servers for the server group used with captive portal authentication. Default: Disabled
show-acceptable-use-policy	Show the acceptable use policy page before the logon page. Default: Disabled
single-session	Allows only one active user session at a time. Default: Disabled
switchip-in-redirection-url	Sends the Mobility Access Switch's IP address in the redirection URL when external captive portal servers are used. An external captive portal server can determine the Mobility Access Switch from which a request originated by parsing the 'switchip' variable in the URL. Default: Disabled
use-chap	Use CHAP protocol. You should not use this option unless instructed to do so by an Aruba representative. Default: Disabled
user-logon	Enables Captive Portal with authentication of user credentials. Default: Enabled
user-vlan-in-redirection-url	Sends VLAN ID of the user in the redirection URL when external captive portal servers are used.
welcome-page	URL of the page that appears after logon and before redirection to the web URL. This can be set to any URL. Default: /auth/welcome.html
white-list	Name of an existing white list on an IPv4 or IPv6 network destination. The white list contains authenticated websites that a guest can access.

Captive Portal Configuration Example

Configuring Captive Portal via the CLI

To configure Captive Portal via the command-line interface, access the CLI configuration mode and issue the following commands:

1. Create a Captive Portal profile

```
(host)(config)#aaa authentication captive-portal cp-profile
(host)(Captive Portal Authentication Profile "cp-profile") #default-role guest
(host)(Captive Portal Authentication Profile "cp-profile") #server-group cp-srv
```



It is assumed that a AAA server-group named "cp-srv" was previously created. To create a AAA server-group, refer the procedure mentioned in ["Configuring Server Groups" on page 242](#).

You can use the following URL to configure an external captive portal authentication on an external server:

```
(host)(config)#aaa authentication captive-portal cp-profile
```

```
(host) (Captive Portal Authentication Profile "cp-profile") #login-page https://<external_server_IP>/<login_page_path>
```

You can use the following URLs to configure an external captive portal authentication on CPPM:

For pre-6.0 ClearPass Policy Manager (Onboard, Legacy Captive Portal Capability):

```
(host)(Captive Portal Authentication Profile "cp-profile") #login-page https://<clearpass-server>/agent/portal/
```

For pre-6.0 ClearPass Guest:

```
(host)(Captive Portal Authentication Profile "cp-profile") #login-page https://<clearpass-guest-server>/<admin-defined-name>.php
```

For 6.0 ClearPass Policy Manager and ClearPass Guest (Integrated Platform):

```
(host)(Captive Portal Authentication Profile "cp-profile") #login-page https://<clearpass-server>/agent/portal/ (Onboard, Legacy Captive Portal Capability)
(host)(Captive Portal Authentication Profile "cp-profile") #login-page https://<clearpass-server>/guest/ (ClearPass Guest)
```

Please refer to ClearPass Policy Manager and ClearPass Guest documentation for more details.

2. Attach a Captive Portal profile to a user role

```
(host)(config) #user-role cp-first
(host)(config-role) #captive-portal cp-profile
```

3. Designate the **cp-first** user-role as the initial role of the AAA profile **cp_aaa**

```
(host)(config) #aaa profile cp_aaa
(host)(AAA Profile "cp_aaa") #initial-role cp-first
```

4. Apply the configured AAA profile to the interface

```
(host)(config) #interface gigabitethernet 0/0/0
    aaa-profile cp_aaa
    no trusted port
```



By default, the authenticated Captive Portal users will be assigned the **guest** user-role.

Configuring Captive Portal via the WebUI



This release of ArubaOS supports creating a user role only using the CLI. To create the user role using the CLI, refer the procedure mentioned in [“Configuring Captive Portal Authentication” on page 277](#).

1. Navigate to the **Configuration>Authentication** page.
2. Select initial role as **cp-first** from the **Initial-Role** drop-down list.
3. Click the **New** button to create a new AAA profile, enter the name of the profile (for example, **profile1**) in the **Name** textbox.
4. Select the authentication method as **captive-portal** from the **Authentication Method** drop-down list.
5. Select the **specify new profile** radio button and enter the captive portal profile name (for example, **c-portal**) in the **Profile Name** textbox.
6. Select the server-group as **cp-srv** from the **Auth Server** drop-down list.



It is assumed that a AAA server-group named "cp-srv" was previously created. To create a AAA server-group, refer the procedure mentioned in [“Configuring Server Groups” on page 242](#).

7. Click **Ok** and **Apply**.
8. To assign AAA profile to the port, select the port from the **Ports Assign** list.

9. Click **Ok** and **Apply**.
10. To make the port untrusted, navigate to **Configuration>Ports** page and select the port from the **Ports** list.
11. Select the **Disabled** radio button from the **Trusted** list.
12. Click **Ok** and **Apply**.




By default, authenticated Captive Portal users will be assigned the **guest** user-role.

Personalizing the Captive Portal Page

The first screen displayed before the captive portal login page informs the user about the authentication requirement and a link (here) is provided . By clicking on this link, the user can access the captive portal login page.

Figure 3 displays the screen that appears before the captive portal login page.

Figure 3 *Authentication Request Page*



The following can be personalized on the default captive portal page:

- Captive portal background
- Page text

The background image and text should be visible to users with a browser window on a 1024 by 768 pixel screen. The background should not clash if viewed on a much larger monitor. A good option is to have the background image at 800 by 600 pixels, and set the background color to be compatible. The maximum image size for the background can be around 960 by 720 pixels, as long as the image can be cropped at the bottom and right edges. Leave space on the left side for the login box.

1. Navigate to the **Configuration > Captive Portal** page.
2. Select the captive portal profile that you want to customize from the **Profile** drop-down list.
3. Select the image that you want to customize from the **Background** drop-down list.

The default page design is as shown below:

Figure 4 *Personalizing the Captive Portal - Default Image*

Configuration > Captive Portal

Profile:

Template:

Page Design:

Background:

Logo:

[Preview current settings](#)



4. To add the policy text:
 - a. Click on the **policy text** tab and enter the acceptable use policy for guests in HTML format.
 - b. Click Apply.
 - c. To view the changes, click on the **Preview current settings** link which displays the Captive Portal page as it will be seen by users



You can configure policy text from the WebUI. To enable it from the CLI, use `show-acceptable-use-policy` command.

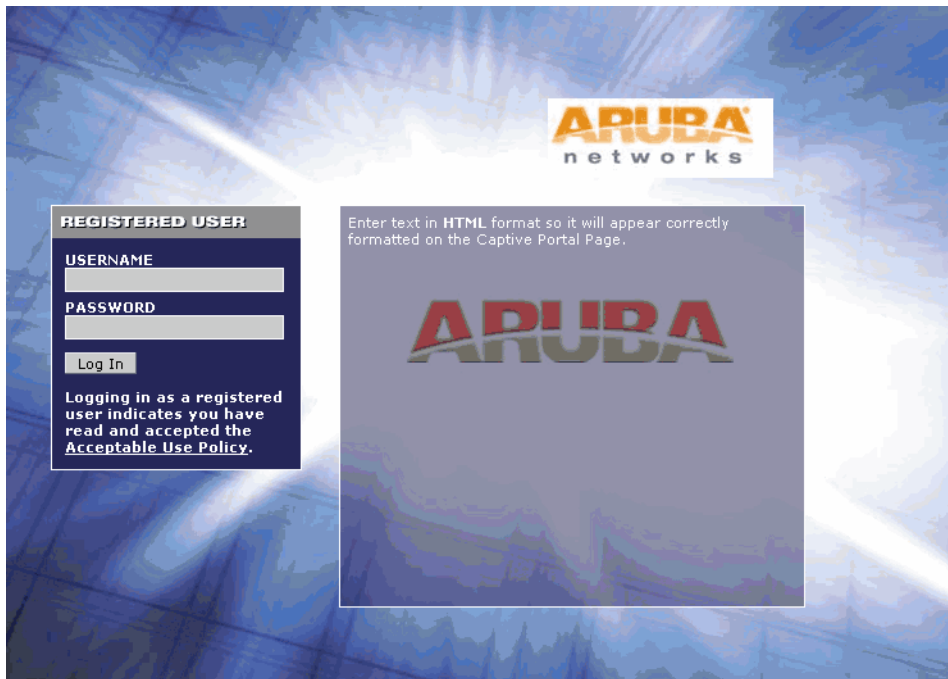
5. To customize the page background:
 - a. Select the **CUSTOM Image** from the **Background** drop-down list.
 - b. Set the background color in the Custom page background color field. The color code must a hexadecimal value in the format #hhhhhh.
 - c. To view the page background changes, click on the **Preview current settings** link and displays the Captive Portal page as it will be seen by users

Figure 5 *Customizing the Captive Portal Background Page*



6. To customize the captive portal background text:
 - a. Enter the text that needs to be displayed in the **Welcome Text (in HTML format)** message box.
 - b. To view the background text changes, click **Preview current settings link** at the bottom on the page. This displays the Captive Portal page as it will be seen by users.

Figure 6 *Customizing the Captive Portal Background Text*



Creating Walled Garden Access

On the Internet, a walled garden typically controls a user's access to web content and services. The walled garden directs the user's navigation within particular areas to allow access to a selection of websites or prevent access to other websites.

Creating Walled Garden Access

Walled garden access is needed when an external or internal captive portal is used. A common example could be a hotel environment where unauthenticated users are allowed to navigate to a designated login page (for example, a hotel website) and all its contents.

Users who do not sign up for Internet service can view "allowed" websites (typically hotel property websites). The website names must be DNS-based (not IP address based) and support the option to define wildcards. This works for client devices with or without HTTP proxy settings.

When a user attempts to navigate to other websites not configured in the white list walled garden profile, the user is redirected back to the login page. In addition, the black listed walled garden profile is configured to explicitly block navigation to websites from unauthenticated users.

Using the CLI to create walled garden access

This example configures a destination named Mywhite-list and adds the domain names, google.com and cnn.com to that destination. It then adds the destination name Mywhite-list (which contains the allowed domain names google.com and cnn.com) to the white list.

```
(host)(config)#netdestination "Mywhite-list"
(host)(config)#name www.google.com
(host)(config)#name www.cnn.com

(host)(config) #aaa authentication captive-portal default
(host)(Captive Portal Authentication Profile "default")#white-list Mywhite-list
```



Ensure not to prefix named netdestination with "http://" or "https://".

Mobility Access Switch Server Certificate

The Mobility Access Switch is designed to provide secure services through the use of digital certificates. A server certificate installed in the Mobility Access Switch verifies the authenticity of the Mobility Access Switch for captive portal.

ArubaOS Mobility Access Switch ships with a demonstration digital certificate. Until you install a customer-specific server certificate in the Mobility Access Switch, this demonstration certificate is used by default for all secure HTTP connections such as captive portal. This certificate is included primarily for the purposes of feature demonstration and convenience and is not intended for long-term use in production networks.

Users in a production environment are urged to obtain and install a certificate issued for their site or domain by a well-known certificate authority (CA). You can generate a Certificate Signing Request (CSR) on the Mobility Access Switch to submit to a CA.

You can use the following command to assign a customized captive portal certificate:

```
(host)(config) #web-server
(host)(Web Server Configuration) #captive-portal-cert
(host)(Web Server Configuration) #captive-portal-cert <captive-portal-cert-name>
```



NOTE

For information on how to generate a CSR and to import a certificate into the Mobility Access Switch, see [“Obtaining a Server Certificate” on page 36](#).

Tunneled Node (previously known as Mux) provides the ability to tunnel the ingress packets (via GRE) from an interface on the Mobility Access Switch (Tunneled Node port) to an Mobility Controller (Tunneled Node server). You can use the Tunneled Nodes to allow the Mobility Controller to provide centralized security policy, authentication, and access-control.

This chapter includes the following topics:

- “Important Points to Remember” on page 285
- “Tunneled Nodes Overview” on page 286
- “Support for Tunneled Node Back-up Server” on page 287
- “Creating and Configuring Tunneled Node Profile” on page 287
- “Verifying and Monitoring Tunneled Nodes” on page 288
- “Verifying and Monitoring the Tunneled Nodes on the Controller” on page 288

Important Points to Remember

- The minimum required version of Mobility Controller ArubaOS is 6.1.2.4.
- Multiple VLAN interfaces are supported in ArubaOS and the GRE tunnel is sourced with the “Switch IP” of the switch.
- Only the following Aruba Mobility Controllers support Tunneled Nodes:
 - 7200 Series Controllers
 - 6000 Series Chassis (M3 module).
 - 3000 Series Controllers
 - 600 Series Controllers
- Ensure that there is an IP reachability between the Mobility Access Switch and the Mobility Controller.
- The Tunneled Node is configured on per-port basis.
- The Tunneled Node is not supported on port-channels. However, Tunneled Node traffic can traverse port-channels.
- The GRE tunnel is created when the interface state transitions to *up* state and the controller is reachable.
- If the interface is up but the Mobility Controller is not reachable, the Mobility Access Switch will retry at every 60 seconds to form a GRE tunnel.
- The Mobility Access Switch allocates an internal VLAN for every Tunneled Node interface. This VLAN is used only for Tunneled Node internal processing. An available internal VLAN ID with the highest number (starting with 4094) is used by default. If you create a new VLAN with the ID that is already assigned to a Tunneled Node, then that VLAN ID is released and then the system allocates the next available VLAN ID. There can be traffic disruption in the mean time.
- Ensure that the VLANs specified in the switching profile and assigned to the Tunneled Node interface is present on the Mobility Controller.
- Only one Tunneled Node profile is supported on the Mobility Access Switch and hence only one Mobility Controller can be used as the Tunneled Node server.

- Spanning tree processing does not take place on the Tunneled Node interface.
- A policer-profile and qos-profile may be applied to a Tunneled Node interface.
- To support Tunneled Node, the Mobility Controller must have an AP and Security bundle license per Mobility Access Switch or ArubaStack.


Tunneled Nodes Overview

This section provides detailed information on the Tunneled Node, also known as a wired Tunneled Node. The Tunneled Node provides access and security using an overlay architecture.

The Tunneled Node connects to one or more client devices at the edge of the network and then establishes a GRE tunnel to the controller. This approach allows the controller to support all the centralized security features, such as IEEE 802.1x authentication, captive-portal authentication, and stateful firewall.

To configure the Tunneled Node, you must specify the IP address of the controller and identify the ports that are to be used as Tunneled Node ports. A tunnel is established between the controller and the Mobility Access Switch for each active Tunneled Node port. [Figure 1](#) shows how the Tunneled Node fits into network operations. Traffic moves through GRE tunnels between the active Tunneled Node ports and the controller. Policies are configured and enforced on the controller. On the controller, you can assign the same policy to Tunneled Node user traffic as you would to any untrusted wired traffic.

Figure 1 Tunneled Node configuration operation



arun_109



The Tunneled Node port can also be configured as a trunk port. This allows you to have multiple clients on different VLANs on the trunk port.

Support for Tunneled Node Back-up Server

ArubaOS provides support for Tunneled Node back-up server by allowing you to configure primary and back-up controllers in the Tunneled Node profile. The Mobility Access Switch keeps checking for the reachability of both primary and the back-up servers configured on the Tunneled Node profile. When the primary controller goes down and if the back-up controller is reachable, the Mobility Access Switch automatically establishes a Tunneled Node between the back-up controller. This ensures that the ports on the Mobility Access Switch do not lose connectivity at any point. The Mobility Access Switch switches back to the primary controller as soon as it finds the primary controller reachable.

Creating and Configuring Tunneled Node Profile

You can create, configure, view, and apply a Tunneled Node profile to an interface using the following commands:

To create a Tunneled Node Profile:

```
(host)(config)# interface-profile tunneled-node-profile <profile-name>
```

To configure the primary and the back-up server for a Tunneled Node:

```
(host)(config)(Tunneled Node Server profile "<profile-name>")#
  backup-controller-ip <IP-address>
  controller-ip <IP-address>
  keepalive <0-40seconds>
  mtu <1024-1500>
no {...}
```

To view a Tunneled Node profile configuration, execute the following command:

```
(host)# show interface-profile tunneled-node-profile tunnell
Tunneled Node Server profile "tunnell"
Parameter                               Value
-----
Controller IP Address                    1.1.1.1
Backup Controller IP Address             2.2.2.1
Keepalive timeout in seconds             10
MTU on path to controller                 1400
```

To apply the Tunneled Node profile to an interface:

```
(host)(config)# interface gigabitethernet <slot/module/port>
  tunneled-node-profile <profile-name>
```



Tunneled Node profile must be applied to the interface along with the switching profile.



For information about how to configure the Tunneled Node server (controller) to use the appropriate Tunneled Node clients, see the appropriate version of the controller User Guide.

Path MTU Discovery

The MTU specified in the Tunneled Node profile must match the path MTU on your network. To determine the correct path MTU between the Tunneled Node client and the controller, use the `ping <ip-address> mtu discovery do size <size>` command. For example, see the following output:

```
(host)# ping 10.13.6.44 mtu_discovery do size 16508
Press 'q' to abort.
PING 10.13.6.44 (10.13.6.44)
From 10.16.48.21 icmp_seq=1 Frag needed and DF set (mtu = 1500)
From 10.16.48.21 icmp_seq=1 Frag needed and DF set (mtu = 1500)
From 10.16.48.21 icmp_seq=1 Frag needed and DF set (mtu = 1500)
From 10.16.48.21 icmp_seq=1 Frag needed and DF set (mtu = 1500)
From 10.16.48.21 icmp_seq=1 Frag needed and DF set (mtu = 1500)
```

Verifying and Monitoring Tunneled Nodes

```
(host)# show tunneled-node state
```

Tunneled Node State

```
-----
IP             MAC             Port state      vlan  tunnel  inactive-time
--            ---
172.16.30.2    00:0b:86:6a:23:80  GE0/0/11  complete  0400  4088    0000
172.16.30.2    00:0b:86:6a:23:80  GE0/0/34  complete  0400  4091    0000
```

```
(host)# show tunneled-node config
```

Tunneled Node Client: Enabled

Tunneled Node Server: 172.16.30.2

Tunneled Node Loop Prevention: Disabled



The `show tunneled-node config` command displays the Tunneled Node server IP address of the controller to which Mobility Access Switch is connected at that moment.

```
(host)# show vlan
```

VLAN CONFIGURATION

```
-----
VLAN  Description      Ports
----  -
4088  MUX Internal VLAN    GE 0/0/11 TUNNEL-0
<output truncated>
```

Verifying and Monitoring the Tunneled Nodes on the Controller

```
(host)# show tunneled-node state
```

Tunneled Node State

```
-----
IP             MAC             s/p state      vlan  tunnel  inactive-time
--            ---
172.16.50.2    00:0b:86:6a:23:80  gigabitethernet0/0/34  complete  400    9        1
172.16.50.2    00:0b:86:6a:23:80  gigabitethernet0/0/11  complete  400    10       1
```

```
(host)# show user-table
```

Users

```
-----
IP             MAC             Name           Role           Age(d:h:m)  Auth  VPN link
AP name        Roaming  ESSID/BSSID/Phy  Profile        Forward mode  Type
-----
```


This chapter describes the following topics:

- “Aruba Instant Overview” on page 291
- “IAP Integration with the MAS” on page 292
- “Viewing the Blacklisted MAC Address of the Rogue APs” on page 292

Aruba Instant Overview

Aruba Instant virtualizes Aruba Mobility Controller capabilities on 802.11n access points (APs), creating a feature-rich enterprise-grade wireless LAN (WLAN) that combines affordability and configuration simplicity.

Aruba Instant is a simple, easy to deploy turn-key WLAN solution consisting of one or more access points. An Ethernet port with routable connectivity to the Internet or a self-enclosed network, is used to deploy an Instant Wireless Network. An Instant Access Point (IAP) can be installed at a single site or deployed across multiple geographically-dispersed locations. Designed specifically for easy deployment, and proactive management of networks, Instant is ideal for small customers or remote locations without any on-site IT administrator.

Aruba Instant consists of an Instant Access Point (IAP) and a Virtual Controller (VC). The Virtual Controller resides within one of the access points. In an Aruba Instant deployment only the first IAP needs to be configured. After the first IAP is deployed, the subsequent IAPs will inherit all the required information from the Virtual Controller.

Supported Devices

The following is a list of Instant devices supported by Aruba:

- IAP-92
- IAP-93
- IAP-104
- IAP-105
- IAP-134
- IAP-135
- IAP-175P/175AC
- RAP-3WN/3WN-US/3WNP/3WNP-US



IAP-104, IAP-105, IAP-134, IAP-135, and IAP-175 support an unlimited number of IAPs on Layer 2 networks. IAP - 92/93 supports 16 IAPs.

For more information on IAP, see the *Instant Access Point 6.2.0.0-3.2 User Guide*.

IAP Integration with the MAS

This release of ArubaOS Mobility Access Switch includes new integration features with Aruba Instant AP (IAP) 3.1 software.

IAP Integration Features

The IAP integration features saves the wastage of power and bandwidth consumed by the rogue APs on the wired network.

This section describes two main IAP integration features:

- Rogue AP containment
- PoE prioritization
- GVRP Integration



Ensure that LLDP is enabled on ports where IAPs are connected.

Rogue AP Containment

When a rogue AP is detected by IAP, the IAP sends out the MAC Address of the rogue AP to the MAS using the Aruba's proprietary LLDP TLV protocol (MAC information TLV with action as Blacklist). The MAS blacklists the MAC address of the rogue AP and turns off the PoE on the port or the MAS installs a bridge entry with the source MAC command as `DROP` to discard the packets originating from or carried to the Rogue AP.



- To enable the rogue AP containment feature, connect the IAPs to the LLDP enabled MAS ports.
- The rogue AP containment functionality is supported only on trusted ports.

PoE Prioritization

When an IAP is plugged into a PoE enabled port on the Mobility Access Switch, the Mobility Access Switch automatically increases the PoE priority from low (default) to high. This only occurs if the `poe-profile` associated with the given port is using the `poe-factory-initial` profile and the default `poe-priority` has not been manually changed.

For information on PoE and configuring the PoE on MAS, see [“Power Over Ethernet” on page 92](#).

GVRP Integration

Configuring GVRP in Mobility Access Switch enables the switch to register/de-register the dynamic VLAN information received from a GVRP applicant such as an IAP in the network. GVRP support also enables the switch to propagate the registered VLAN information to the neighboring bridges in the network.



When VLANs are added on WLAN or wired profiles, the VLANs are advertised to the upstream switch using GVRP messages.

For information on enabling and configuring GVRP on Mobility Access Switch, see [“Enabling and Configuring GVRP Functionality” on page 117](#).

Viewing the Blacklisted MAC Address of the Rogue APs

You can use the following command to view details on the blacklisted MAC addresses received from the IAPs:

```
(host) #show lldp neighbor interface gigabitethernet 1/0/40 detail

Interface: gigabitethernet1/0/40, Number of neighbors: 1
-----
Chassis id: d8:c7:c8:ce:0d:63, Management address: 192.168.0.252
Interface description: bond0, ID: d8:c7:c8:ce:0d:63, MTU: 1522
Device MAC: d8:c7:c8:ce:0d:63
Last Update: Thu Sep 27 10:59:37 2012
Time to live: 120, Expires in: 103 Secs
System capabilities : Bridge,Access point
Enabled capabilities: Access point
System name: d8:c7:c8:ce:0d:63
System description:
  ArubaOS (MODEL: 105), Version 6.1.3.4-3.1.0.0 (35380)
Auto negotiation: Supported, Enabled
Autoneg capability:
  10Base-T, HD: yes, FD: yes
  100Base-T, HD: yes, FD: yes
  1000Base-T, HD: no, FD: yes
Media attached unit type: 1000BaseTFD - Four-pair Category 5 UTP, full duplex mode (30)
MAC:          7c:d1:c3:c7:e9:72: Blacklist
MAC:          9c:b7:0d:7d:0b:72: Blacklist
MAC:          7c:d1:c3:d1:02:c8: Blacklist
```

Viewing Port Errors

The following command displays the state of the interface due to the detection of the blacklisted rogue AP by the MAS:

```
(host) # show port-error-recovery

Layer-2 Interface Error Information
-----
Interface  Error                               Error seen time                Recovery time
-----
GE0/0/47   Blacklisted device detected  2012-05-09 20:37:10 (PST)  2012-05-09 20:42:10 (PST)
```

Recovering Ports Manually

You can use the following command to manually recover the state of the interface:

```
(host) (config) #clear port-error-recovery interface <interface-name>
```

The following command clears the errors on gigabitethernet 0/0/42:

```
(host) (config) #clear port-error-recovery interface gigabitethernet 0/0/42
```

To clear the port errors on all interfaces execute the following command:

```
(host) (config) #clear port-error-recovery
```



The interface recovers from the port error state automatically after five minutes and can be re-activated.

This chapter describes the following topics:

- “Overview” on page 295
- “Configuring mDNS packet forwarding” on page 295
- “Sample Configuration” on page 296

Overview

Aruba AirGroup is a unique enterprise-class capability that leverages zero configuration networking to allow mobile devices to use services like the Apple AirPrint wireless printer service and the Apple AirPlay streaming service. These services use multicast DNS (mDNS) packets to locate devices and the services that those devices offer.

To ensure Wired and Wireless AirPrint/AirPlay devices can communicate with one another previously required all devices to be on the same Layer-2 network which may not be desirable. Airgroup, which was introduced in ArubaOS 7.2 for the Mobility Access Switch and ArubaOS 6.1.3.4-AirGroup for the Mobility Controller, avoids that need by enabling the ability to just redirect mDNS traffic to a Mobility Controller regardless of VLAN. A simple rule on the MAS is used to redirect all incoming mDNS packets on a port to an L2-GRE tunnel which is then terminated on a Mobility Controller. This allows the Mobility Controller to handle the rest of the AirGroup functionality.

Aruba AirGroup is available in two deployment models; Integrated and Overlay. The location of the mDNS proxy function primarily differentiates the two deployment models. The Mobility Access Switch can interoperate in either deployment model but uses the same underlying features like L2-GRE tunnels used in the Overlay Deployment Model between Mobility Controller.

For more information about Aruba AirGroup, Overlay Deployment Model, and configuration, see the *Aruba AirGroup Deployment Guide*.

Configuring mDNS packet forwarding

To configure mDNS packet forwarding to an AirGroup Mobility Controller, see the following procedures.

1. Create a switching profile and add VLAN for mDNS traffic.

```
(host) (config) #interface-profile switching-profile <profile-name>
(host) (switching profile) #switchport-mode trunk
(host) (switching profile) #trunk allowed vlan <vlan-list>
```



Both ends of an L2-GRE tunnel must carry the same user VLANs.

2. Configure an L2-GRE tunnel and apply the switching profile.

This release of ArubaOS Mobility Access Switch supports L2 connectivity through GRE tunnel. L2-GRE tunnel extends VLANs across switches and Aruba controllers.



If the MAS and AirGroup controller are on the same L2 network, L2-GRE tunnel is not required.

```
(host) (config) #interface tunnel ethernet <tunnel-id>
(host) (Tunnel "tunnel-id") #description <interface-description>
(host) (Tunnel "tunnel-id") #source-ip <source-tunnel-ip>
(host) (Tunnel "tunnel-id") #destination-ip <destination-tunnel-ip>
(host) (Tunnel "tunnel-id") #switching-profile <profile-name>
(host) (Tunnel "tunnel-id") #keepalive <Tunnel heartbeat interval in seconds (1-86400)> <Tunnel Heartbeat Retries (1-1024)>
```

3. Configure a stateless ACL with mDNS UDP port 5353 redirect rule.

```
(host) (config) #ip access-list stateless <name of the access-list>
(host) (config-stateless)#any any udp 5353 redirect tunnel <L2-GRE-tunnel-ID>
```



The Extended-action options appearing in a stateless ACL after `redirect tunnel <ID>` are unsupported.

4. Apply redirect ACL to either a port or user role.

a. Apply redirect ACL to a port.



Before you apply redirect ACL to a port, you must create explicit allow rules while configuring mDNS redirect ACL to permit non-mDNS traffic.

```
(host) (config) #interface gigabitethernet <slot/module/port>
(host) (gigabitethernet) #ip access-group in <ingress-access-control-list>
```

b. Apply redirect ACL to a user role.



Add the mDNS redirect ACL to position one of the user-role.

```
(host) (config) #user-role <role-name>
(host) (config-role) #access-list stateless <name-of-access-list> position 1
```

Inter-tunnel flooding

There can be multiple switches from the same L2 network having L2-GRE tunnel terminating at a single controller. This may generate inter-tunnel flooding resulting in loops within the switch network. To avoid this scenario, disable inter-tunnel flooding in the switch and the controller.

```
(host) (config) #interface tunnel ethernet <tunnel-id>
(host) (Tunnel "tunnel-id") #no inter-tunnel-flooding
```

Sample Configuration

To create a switching profile and add VLAN for mDNS traffic:

```
(host) (config) #interface-profile switching-profile mDNS_vlan_200
(host) (switching profile "mDNS_vlan_200") #switchport-mode trunk
(host) (switching profile "mDNS_vlan_200") #trunk allowed vlan 200
```


To configure an L2-GRE tunnel and apply the switching profile:

```
(host) (config) #interface tunnel ethernet 1
(host) (Tunnel "1") #description L2-GRE_Interface
(host) (tunnel "1") #source-ip 10.0.0.1
(host) (tunnel "1") #destination-ip 10.0.1.2
(host) (tunnel "1") #switching-profile mDNS_vlan_200
(host) (tunnel "1") #keepalive 30 5
```

To configure stateless ACL with mDNS redirect rule:

```
(host) (config) #ip access-list stateless mDNS_redirect
(host) (config-stateless-mDNS_redirect)#any any udp 5353 redirect tunnel 1
```

To apply redirect ACL to a port:

```
(host) (config) #interface gigabitethernet 0/0/1
(host) (gigabitethernet "0/0/1") #ip access-group in mDNS_redirect
```

To apply redirect ACL to a user role:

```
(host) (config) #user-role employee
(host) (config-role) #access-list stateless mDNS_redirect position 1
```


ArubaOS 7.2 for the Mobility Access Switch and ClearPass Policy Manager (CPPM) 6.0 includes support for centralized policy definition and distribution. With this new release, ArubaOS Mobility Access Switch introduces downloadable roles. By using this feature, when CPPM successfully authenticates a user, the user is assigned a role by CPPM and if the role is not defined on the MAS, the role attributes can also be automatically downloaded.

This chapter contains the following sections:

- “Introduction” on page 299
- “Important Points to Remember” on page 300
- “Enabling Downloadable Role on MAS” on page 300
- “Sample Configuration” on page 301

Introduction

In order to provide highly granular per-user level access, user roles can be created when a user has been successfully authenticated. During the configuration of a policy enforcement profile at CPPM, the administrator can define a role that should be assigned to the user after successful authentication. In RADIUS authentication, when CPPM successfully authenticates a user, the user is assigned a role by CPPM and if the role is not defined on the MAS, the role attributes can also be automatically downloaded.



This feature works for 802.1X and/or MAC-based authenticated users. Downloading a role for Captive Portal authenticated user is not supported in this release.

Important Points to Remember

- CPPM does not perform any error checking to confirm accuracy in the configuration. Therefore, it is recommended that you review your CPPM policy prior to enabling this feature.
- Only attributes that are listed below, herein referred to as whitelist role attributes, can be defined in CPPM. Other attributes under user-role such as policer-profile, qos-profile, vlan, and voip-profile may be referenced, but cannot be defined in CPPM.
 - `netdestination`
 - `netservice`
 - `ip access-list stateless`
 - `user-role`
- Any other attributes that are referred to by a CPPM policy must either be defined within the CPPM policy itself or configured on the MAS before the policy is downloaded.
- An instance name (name of a whitelist role attribute as stated above) is case-sensitive. Attributes must adhere to the following rules:
 - Should not match any CLI option nested under a command from the whitelist.
 - Should not contain a number or a combination of numbers.
 - Should not contain any periods '.'.

Example below are considered as invalid configurations and will fail CPPM role download on MAS:

```
netservice 'tcp' tcp 443
```

The first instance of **tcp** is a user-defined field while the second is an operator of the **netservice** command. This violates the first rule.

```
netdestination 'alias'
```

The user-defined name **alias** is also a valid operator of the **netdestination** command. This violates the first rule.

```
netdestination '10.1.5'
```

This user-defined name uses both numbers and periods. This violates the second and third rule.

```
ip access-list stateless '100'
```

This user-defined name uses numbers. This violates the second rule.

It is recommended that some naming convention similar to the CamelCase (mixture of upper and lower case letters in a single word) be used to avoid collisions with the CLI options in the role description.

Enabling Downloadable Role on MAS

You can enable role download using the CLI or WebUI.

Using the WebUI

1. Navigate to the **Configuration > Authentication > Profiles** tab.
2. Select an AAA profile.
3. Select **Enabled** from the **Role Download** drop-down list.

Using the CLI

```
(host) (config) #aaa profile <profile-name>
(host) (AAA profile) #download-role
```

Sample Configuration

The following example shows the configuration details to integrate CPPM server with MAS to automatically download roles.

CPPM Server Configuration

Adding a Device

1. From the **Configuration > Network > Devices** page, click the **Add Device** link.
2. On the **Device** tab, enter the **Name**, **IP or Subnet Address**, and **RADIUS Shared Secret** fields. Keep the rest of the fields as default.
3. Click **Save**.

The fields are described in the following figure and table.

Figure 1 *Device Tab*

Add Device			
Device	SNMP Read Settings	SNMP Write Settings	CLI Settings
Name:	Mobility Access Switch		
IP or Subnet Address:	10.1.1.1 (e.g., 192.168.1.10 or 192.168.1.1/24)		
Description:			
RADIUS Shared Secret:	Verify:
TACACS+ Shared Secret:		Verify:	
Vendor Name:	IETF		
Enable RADIUS CoA:	<input type="checkbox"/>		
Attributes			
Attribute	Value		
1. Click to add...			
Add Cancel			

Table 1 *Device Tab*

Container	Description
Name	Specify the name or identity of the device.
IP or Subnet Address	Specify the IP address or subnet (example 10.1.1.1/24) of the device.
RADIUS Shared Secret	Enter and confirm a Shared Secret for each of the two supported request protocols.

Adding Enforcement Profile

1. From **Configuration > Enforcement > Profiles** page, click **Add Enforcement Profile**.

2. On the **Profile** tab, select **Aruba Downloadable Role Enforcement** from the **Template** drop-down list.
3. Enter the **Name** of the enforcement profile.
Keep the rest of the fields as default.
4. Click **Save**.

The fields are described in the following figure and table.

Figure 2 Enforcement Profiles Page

Configuration » Enforcement » Profiles » Add Enforcement Profile

Enforcement Profiles

Profile	Attributes	Summary
Template:	Aruba Downloadable Role Enforcement ▼	
Name:	Enforcement_Profile_1	
Description:		
Type:	RADIUS	
Action:	<input checked="" type="radio"/> Accept <input type="radio"/> Reject <input type="radio"/> Drop	
Device Group List:	<div> <div>Remove</div> <div>View Details</div> <div>Modify</div> </div> Add new Device Group	
	<div>--Select-- ▼</div>	

Table 2 Enforcement Profiles Page

Container	Description
Template	Policy Manager comes pre-packaged with several enforcement profile templates. In this example, select Aruba Downloadable Role Enforcement - RADIUS template that can be filled with user role definition to create roles that can be assigned to users after successful authentication.
Name	Specify the name of the enforcement profile.

5. On the **Attributes** tab, select **Radius:Aruba** from the **Type** drop-down list.
6. From the **Name** drop-down list, select **Aruba-CPPM-Role**.
7. In the **Value** field, enter the attribute for the downloadable-role.
8. Click **Save**.

The fields are described in the following figure and table.

Figure 3 *Enforcement Profiles Attributes Tab*

Configuration » Enforcement » Profiles » Add Enforcement Profile

Enforcement Profiles

Profile	Attributes	Summary
Type	Name	Value
1. Radius:Aruba	Aruba-CPPM-Role	= ip access-list stateless icmp_permit any any svc-icmp permit ! user-role cppm-D-Role access-list stateless icmp_permit vlan 101 !
2. Click to add...		

Table 3 *Enforcement Profiles Attributes Tab*

Container	Description
Type	Type is any RADIUS vendor dictionary that is pre-packaged with Policy Manager, or imported by the Administrator. This field is prepopulated with the dictionary names.
Name	Name is the name of the attribute from the dictionary selected in the Type field. The attribute names are prepopulated from the dictionary.
Value	Value is attribute for the downloadable role. You can enter freeform text to define the role and policy. Note: The maximum limit for freeform text is 16,000 bytes.

Adding Enforcement Policy

1. From **Configuration > Enforcement > Policies** page, click **Add Enforcement Policy**.
2. On the **Enforcement** tab, enter the name of the enforcement policy.
3. From the **Default Profile** drop-down list, select **[Deny Access Profile]**.
Keep the rest of the fields as default.
4. Click **Save**.

The fields are described in the following figure and table.

Figure 4 *Enforcement Policies Enforcement Tab*

Configuration » Enforcement » Policies » Add

Enforcement Policies

Enforcement	Rules	Summary
Name:	<input type="text" value="Enforcement_Policy_1"/>	
Description:	<div><div></div><div>...</div></div>	
Enforcement Type:	<input checked="" type="radio"/> RADIUS <input type="radio"/> TACACS+ <input type="radio"/> WEBAUTH (SNMP/Agent/CLI/CoA) <input type="radio"/> Application	
Default Profile:	<input type="text" value="[Deny Access Profile]"/>	View Details Modify

Table 4 *Enforcement Policies Enforcement Tab*

Container	Description
Name	Specify the name of the enforcement policy.
Default Profile	An Enforcement Policy applies Conditions (roles, health, and time attributes) against specific values associated with those attributes to determine the Enforcement Profile. If none of the rules matches, Policy Manager applies the Default Profile. See “Adding Enforcement Profile” on page 301 to add a new profile.

- On the **Rules** tab, click **Add Rule**.
- On the **Rules Editor** popup, select the appropriate values in the **Conditions** section.
- In the **Enforcement Profiles** section, select the RADIUS enforcement profile that you created in step [“Adding Enforcement Profile” on page 301](#) from the **Profile Names** drop-down list.
- Click **Save**.

The fields are described in the following figure and table.

Figure 5 *Enforcement Policies Rules Editor*

Rules Editor

Conditions

Match ALL of the following conditions:

	Type	Name	Operator	Value
1.	Authentication	Source	EQUALS	[Local User Repository]
2.	Click to add...			

Enforcement Profiles

Profile Names:

- [RADIUS] Enforcement_Profile_1

Buttons: Move Up, Move Down, Remove

Dropdown: -Select to Add-

Buttons: Save, Cancel

Table 5 *Enforcement Policies Rules Editor*

Container	Description
Type	The rules editor appears throughout the Policy Manager interface. It exposes different namespace dictionaries depending on Service type. When working with service rules, you can select Authentication namespace dictionary
Name	Drop-down list of attributes present in the selected namespace. In this example, select Source .
Operator	Drop-down list of context-appropriate (with respect to the attribute) operators. In this example, select EQUALS .

Table 5 Enforcement Policies Rules Editor

Container	Description
Value	Drop-down list of the Authentication source database. In this example, select [Local User Repository] .
Profile Names	Name of the RADIUS enforcement profile.

Adding Services

1. From the **Configuration > Services** page, click the **Add Service** link.
2. On the **Service** tab, select **802.1X Wired** from the **Type** drop-down-list.
3. In the **Name** field, enter the name of the service.

Keep the rest of the fields as default.

4. Click **Save**.

The fields are described in the following figure and table.

Figure 6 Service Tab

Configuration » Services » Add

Services

Service	Authentication	Roles	Enforcement	Summary
Type:	802.1X Wired			
Name:	Service_1			
Description:	802.1X Wired Access Service			
Monitor Mode:	<input type="checkbox"/> Enable to monitor network access without enforcement			
More Options:	<input type="checkbox"/> Authorization <input type="checkbox"/> Posture Compliance <input type="checkbox"/> Audit End-hosts <input type="checkbox"/> Profile Endpoints			
Service Rule				
Matches <input type="radio"/> ANY or <input checked="" type="radio"/> ALL of the following conditions:				
Type	Name	Operator	Value	
1. Radius:IETF	NAS-Port-Type	EQUALS	Ethernet (15)	
2. Radius:IETF	Service-Type	BELONGS_TO	Login-User (1), Framed-User (2), Authenticate-Only (8)	
3. Click to add...				

Table 6 Service Tab

Container	Description
Type	Select the desired service type from the drop down menu. In this example, select 802.1X Wired .
Name	Specify the name of the service.

5. On the **Authentication** tab, select **[Local User Repository]** **[Local SQL DB]** from the **Authentication Sources** drop-down list.

Keep the rest of the fields as default.

6. Click **Save**.

The fields are displayed in the following figure.

Figure 7 *Authentication Tab*

Configuration » Services » Add

Services

Service	Authentication	Roles	Enforcement	Summary
Authentication Methods:				
<div>[EAP PEAP] [EAP FAST] [EAP TLS] [EAP TTLS] [EAP MSCHAPv2]</div>		<div>Move Up Move Down Remove View Details Modify</div>		
--Select to Add--				
Authentication Sources:				
<div>[Local User Repository] [Local SQL DB]</div>		<div>Move Up Move Down Remove View Details Modify</div>		
--Select to Add--				
Strip Username Rules:				
<input type="checkbox"/> Enable to specify a comma-separated list of rules to strip username prefixes or suffixes				

7. On the **Enforcement** tab, select the enforcement policy that you created in step “[Adding Enforcement Policy](#)” on page 303 from the **Enforcement Policy** drop-down list.

Keep the rest of the fields as default.

8. Click **Save**.

The fields are displayed in the following figure.

Figure 8 *Enforcement Tab*

Configuration » Services » Add

Services

Service	Authentication	Roles	Enforcement	Summary
Use Cached Results:		<input type="checkbox"/> Use cached Roles and Posture attributes from previous sessions		
Enforcement Policy:		<div>Enforcement_Policy_1</div> <div>Modify</div> Add new Enforcement Policy		
Enforcement Policy Details				
Description:				
Default Profile:		[Deny Access Profile]		
Rules Evaluation Algorithm:		first-applicable		
Conditions		Enforcement Profiles		
1. (Authentication:Source EQUALS [Local User Repository])		Enforcement_Profile_1		

For more configuration details on CPPM, see the *ClearPass Policy Manager 6.0 User Guide*.

MAS Configuration

Configuring CPPM Server on MAS

```
(host) (config) #aaa authentication-server radius cppm_server
(host) (RADIUS Server "cppm_server") #host <ip_address_of_cppm_server>
(host) (RADIUS Server "cppm_server") #key <shared_secret>
```

Configuring Server Group to include CPPM Server

```
(host) (config) #aaa server-group cppm_grp
(host) (Server Group "cppm_grp") #auth-server cppm_server
```

Configuring 802.1X Profile

```
(host) (config) #aaa authentication dot1x cppm_dot1x_prof
```

Configuring AAA Profile

```
(host) (config) #aaa profile cppm_aaa_prof
(host) (AAA Profile "cppm_aaa_prof") #authentication-dot1x cppm_dot1x_prof
(host) (AAA Profile "cppm_aaa_prof") #dot1x-server-group cppm_grp
(host) (AAA Profile "cppm_aaa_prof") #download-role
```

Show AAA Profile

```
(host) #show aaa profile cppm_aaa_prof
```

```
AAA Profile "cppm_aaa_prof"
-----
Parameter                                Value
-----
Initial role                             logon
MAC Authentication Profile                N/A
MAC Authentication Default Role           guest
MAC Authentication Server Group           default
802.1X Authentication Profile              cppm_dot1x_prof
802.1X Authentication Default Role         guest
802.1X Authentication Server Group         cppm_grp
Download Role from ClearPass             Enabled
L2 Authentication Fail Through            Disabled
RADIUS Accounting Server Group             N/A
RADIUS Interim Accounting                 Disabled
XML API server                           N/A
RFC 3576 server                           N/A
User derivation rules                     N/A
SIP authentication role                   N/A
Enforce DHCP                             Disabled
Authentication Failure Blacklist Time     3600 sec
```


Wireless networks can use virtual private network (VPN) connections to further secure wireless data from attackers.



The Mobility Access Switch only supports Site-to-Site VPN configurations in tunnel mode. IPSec transport mode is not supported in this release.



There is no Equal Cost Multiple Path (ECMP) support over VPN.

Planning a Site-to-Site VPN Configuration

Site-to-site VPNs allow networks (for example, a branch office network) to connect to other networks (for example, a corporate network). Unlike a remote access VPN, hosts in a site-to-site VPN do not run VPN client software. All traffic for the other network is sent and received through a VPN gateway which encapsulates and encrypts the traffic.

The following IKE authentication methods are supported for site-to-site VPNs:

- Preshared Key authentication
- Certificate authentication. You can configure a RSA server certificate and a CA certificate for each site-to-site VPN IPsec map configuration. If you are using certificate-based authentication, the peer must be identified by its certificate subject-name distinguished name (for deployments using IKEv2) or by the peer's IP address (for IKEv1).



Certificate-based authentication is supported for site-to-site VPN between two Aruba devices with static IP addresses. Additionally, Certificate-based authentication is also supported with dynamic IP addresses when IKEv2 is used.

Selecting an IKE protocol

Mobility Access Switches running ArubaOS version 7.2 and later support both IKEv1 and the newer IKEv2 protocol to establish IPsec tunnels. IKEv2 is simpler, faster, and a more reliable protocol than IKEv1.

If your IKE policy uses IKEv2, you should be aware of the following caveats when you configure your VPN:

- ArubaOS does not support separate pre-shared keys for both directions of an exchange; the same pre-shared key must be used by both peers. ArubaOS does not support mixed authentication with both pre-shared keys and certificates; each authentication exchange requires a single authentication type. (For example, if a Site-to-Site peer authenticates with a pre-shared key, the other peer must also authenticate with a pre-shared key.)
- ArubaOS does not support IKEv2 mobility (MOBIKE), Authentication Headers (AH) or IP Payload Compression Protocol (IPComp).

Supported IKE Modes

ArubaOS supports site-to-site VPNs using IKEv2 or IKEv1 Main-mode/Aggressive-mode. By default, site-to-site VPN uses IKEv1 Main-mode with Pre-Shared-Keys to authenticate the IKE security association (SA). This method requires static IP addresses between the peers and therefore will not work for dynamically addressed peers.


To support site-site VPN with dynamically addressed devices, you must use IKEv1 Aggressive-mode or IKEv2 with certificates. The VPN endpoint with a dynamic IP address must be configured to be the initiator and the endpoint with the static IP address must be configured as the responder.

Aruba MAS and Mobility Controllers can use IKEv1 or IKEv2 to establish a site-to-site VPN between another MAS or Mobility Controller or between that MAS and third-party device. Note, however, that only Aruba devices (Mobility Access Switches or Mobility Controllers) and devices running Windows 2008 Server or Strongswan 4.3 support IKEv2 authentication.

VPN Topologies

You must configure VPN settings on the devices at both the local and remote sites. In the following figure, a VPN tunnel connects Network A to Network B across the Internet.

Figure 1 Site-to-Site VPN Configuration Components



To configure the VPN tunnel on Mobility Access Switch, you need to configure the following:

- The source network (Network A).
- The destination network (Network B).
- The VLAN or loopback interface on the MAS connected to the Layer-3 network (Interface A in the [Figure 1](#)).
- The peer gateway address, which is the IP address of the Mobility Controller's interface connected to the Layer-3 network (Interface B in the [Figure 1](#)).

Configuring VPN

To configure a site-to-site VPN with a static IP MAS device and static IP Mobility Controller using IKEv1, issue the following commands:

```
crypto-local ipsec-map <name> <priority>
  src-net <ipaddr> <mask>
  dst-net <ipaddr> <mask>
  peer-ip <ipaddr>
  interface [loopback <loopback-number>|vlan <vlan-id>]
  version vl
  pre-connect enable|disable
```

For certificate authentication:

```
set ca-certificate <cacert-name>
set server-certificate <cert-name>
```

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  version v1
  authentication rsa-sig
  group 1|2
  hash {md5|sha|sha1-96}
  lifetime <seconds>
```

For preshared key authentication:

```
crypto-local isakmp key <key> address <ipaddr> netmask <mask>
```

```
crypto isakmp policy <priority>
  encryption {3des|aes128|aes192|aes256|des}
  version v1
  authentication pre-share
  group {1|2}
  hash {md5|sha|sha1-96}
  lifetime <seconds>
```

To configure site-to-site VPN with a static MAS and a dynamically addressed Mobility Controller that initiates IKE Aggressive-mode for Site-Site VPN:

```
crypto-local ipsec-map <name> <priority>
  src-net <ipaddr> <mask>
  dst-net <ipaddr> <mask>
  peer-ip <ipaddr>
  local-fqdn <local_id_fqdn>
  interface [loopback <loopback-number>|vlan <vlan-id>]
  pre-connect [enable|disable]
```

For the Pre-shared-key:

```
crypto-local isakmp key <key> address <ipaddr> netmask 255.255.255.255
```

For a static IP Mobility Controller that responds to IKE Aggressive-mode for Site-Site VPN:

```
crypto-local ipsec-map <name2> <priority>
  src-net <ipaddr> <mask>
  dst-net <ipaddr> <mask>
  peer-ip 0.0.0.0
  peer-fqdn fqdn-id <peer_id_fqdn>
  vlan <id>
```

For the Pre-shared-key:

```
crypto-local isakmp key <key> fqdn <fqdn-id>
```

For a static IP MAS that responds to IKE Aggressive-mode for Site-Site VPN with One PSK for All FQDNs:

```
crypto-local ipsec-map <name2> <priority>
  src-net <ipaddr> <mask>
  peer-ip 0.0.0.0
  peer-fqdn any-fqdn
  vlan <id>
```

For the Pre-shared-key for All FQDNs:

```
crypto-local isakmp key <key> fqdn-any
```

Configuration Examples

Main-Mode

The following example shows a MASs a with static IP address and Mobility Controller with a static IP address.

MAS:

```
crypto-local ipsec-map map1 10
src-net 1.1.1.1 255.255.255.0
dst-net 2.2.2.2 255.255.255.0
peer-ip 3.3.3.3
interface vlan 50
version v1
pre-connect enable

crypto-local isakmp key secret address 3.3.3.3 netmask 255.255.255.255
```

Controller:

```
(host) (config) #crypto-local ipsec-map map2 10
src-net 2.2.2.2 255.255.255.0
dst-net 1.1.1.1 255.255.255.0
peer-ip 4.4.4.4
vlan 50
version v1
trusted enabled

crypto-local isakmp key secret address 4.4.4.4 netmask 255.255.255.255
```

Aggressive-Mode with Tunneled Node over VPN

This release of ArubaOS also adds support for Tunneled Node over VPN. This allows you to provide all the centralized security policy, authentication, and access-control from a tunneled node over a VPN connection.

The following example shows site-to-site VPN configured between MAS with a dynamic IP address and Mobility Controller with a static IP address. In this example, the MAS is configured to be the initiator of IKE Aggressive-mode and the Mobility Controller is the responder of IKE Aggressive-mode.

1. Establish a VPN connection between the MAS and the Mobility Controller.

MAS:

```
(host) (config) #crypto-local ipsec-map here-there-vpn 100
src-net 101.1.1.1 255.255.255.0
dst-net 100.1.1.1 255.255.255.0
peer-ip 2.2.2.2
local-fqdn test@abc.com
interface vlan 2

crypto-local isakmp key secret address 2.2.2.2 netmask 255.255.255.255
```

Mobility Controller:

```
(host) (config) #crypto-local ipsec-map there-here-vpn 100
src-net 100.1.1.0 255.255.255.0
dst-net 101.1.1.0 255.255.255.0
peer-ip 0.0.0.0
peer-fqdn fqdn-id test@abc.com
vlan 2
```



```
crypto-local isakmp key secret fqdn test@abc.com
```

2. Establish a Tunneled Node connection between the MAS and Mobility Controller. Ensure that the MAS's switch IP is in the IPSec source network and the Mobility Controller's IP address is in the IPSec destination network.

```
(host)(config)(Tunneled Node Server profile "tunnel1")#  
controller-ip 100.1.1.1
```

```
(host)# show interface-profile tunneled-node-profile tunnel1  
Tunneled Node Server profile "tunnel1"  
Parameter                               Value  
-----  
Controller IP Address                    100.1.1.1  
Keepalive timeout in seconds             10  
MTU on path to controller                 1400
```

3. Apply the tunneled node profile to an interface.

Static Route Support for VPN

You can also configure a static route to be used with VPN to and from your Mobility Access Switch. Use the following command to configure a static route using an IPSec map.

```
(host) (config) #ip-profile  
(host) (ip-profile) #route <destip> <netmask> ipsec <mapname> metric <metric>
```

The value **metric** is used to enable IPSec route redundancy. **Metric** is cost assigned to the IPSec map that determines which map should be used first and which map should be used if the first map is unavailable.

```
(host) (ip-profile) #route 5.5.5.0 255.255.255.0 ipsec map1 metric 10  
(host) (ip-profile) #route 5.5.5.0 255.255.255.0 ipsec map2 metric 20
```

In the above example, map1 would be used over map2. However, if map1 was unavailable, map2 would be used.



Pre-connect must be enabled on the IPSec maps for IPSec route redundancy.

The static route to IPSec map can be configured before or after the crypto map. If the static route is configured before the IPSec map, the static route is kept in the configuration; however, the route is not pushed to the routing table.

You can use port mirroring to send copies of all or sampled packets seen on specific port(s) or port-channel to a destination. You can use this method for appliances such as sniffers that monitor network traffic for further analysis.

This chapter includes the following topics:

- [“Important Points to Remember” on page 315](#)
- [“The Source Port” on page 315](#)
- [“The Destination Port” on page 315](#)
- [“Mirroring Sampled Ratio” on page 316](#)
- [“Creating and Applying a Mirroring Profile to an Interface” on page 316](#)
- [“Sample Configuration” on page 316](#)
- [“Verifying Port Mirroring Configuration” on page 316](#)

Important Points to Remember

- The destination port must be a local interface.
- A VLAN cannot be configured as the destination.
- The Mobility Access Switch mirroring session limit is one.

The Source Port

You can use port mirroring to take a copy of the ingress and egress packets on one or more ports. Packets are sent to the destination without modification at Layer 2. Any number of network ports can be configured for monitoring. Port-channel can also be the source for mirroring. If the bandwidth for source is greater than the destination, packets loss can occur. The Mobility Access Switch does not distinguish whether the source port is a Layer 2 access or trunk interface.

The Destination Port

One port can be the destination interface; Port-channels and VLANs cannot be a destination. Normal traffic forwarding will not be performed on the destination port. Only the mirrored packets can be received on the destination port. A destination port cannot be a port mirroring source port at the same time. The destination port does not participate in any Layer 2 protocol, including Spanning-tree. Switching profile such as access or trunk profile cannot be applied on the destination port.

Mirroring Sampled Ratio

You can configure the Mobility Access Switch to mirror at a ratio of one out of X packets (1:X) to the destination. The value of X can be between 0 and 2,047.

Table 1 *Sampled Ratio Values*

Ratio (X value)	Description
0	Does not mirror any packet to the destination.
1	Mirrors all packets to the destination (1:1). This is the default.
100	Mirrors 1 out of 100 packets to the destination.
...	...
2047	Mirrors 1 out of 2,047 packets to the destination.

Creating and Applying a Mirroring Profile to an Interface

Using the CLI

```
(host)(config)# interface-profile mirroring-profile <profile-name>
  destination gigabitethernet <slot/module/port>
  ratio <0-2047>
  clone <source>
  no {...}
(host)(config)# interface gigabitethernet <slot/module/port>
  mirroring-in-profile <profile-name>
  mirroring-out-profile <profile-name>
```

The mirroring-in-profile is used for ingress traffic and the mirroring-out-profile is used for egress traffic.

Sample Configuration

```
(host)(config)# interface-profile mirroring-profile MIRROR
  destination gigabitethernet 0/0/40
  ratio 10
  exit
(host)(config)# interface gigabitethernet 0/0/30
  mirroring-in-profile MIRROR
  mirroring-out-profile MIRROR
```

Verifying Port Mirroring Configuration

```
(host) (config) #show mirroring
```

```
Mirroring Profile Name : MIRROR
Mirroring Ratio       : 10
Mirroring Destination : GE0/0/40
Ingress mirrored ports : GE0/0/30
Egress mirrored ports  : GE0/0/30
```

```
(host)# show interface-config gigabitethernet 0/0/30
```

```
gigabitethernet "0/0/30"
```

```
-----
```

Parameter	Value
-----------	-------

-----	-----
-------	-------

```
<output truncated>
```

```
Ingress Port Mirroring Profile    MIRROR
```

```
Egress Port Mirroring Profile    MIRROR
```

```
<output truncated>
```

```
(host)# show interface-profile mirroring-profile MIRROR
```

```
Mirroring profile "MIRROR"
```

```
-----
```

Parameter	Value
-----------	-------

-----	-----
-------	-------

```
gigabitethernet    0/0/30
```

```
Port mirroring ratio  10
```


This chapter describes the following topics:

- “Remote Monitoring (RMON) Overview” on page 319
- “Enabling RMON Service” on page 319
- “Configuring RMON Parameters” on page 320
- “Viewing RMON Active Configuration” on page 322

Remote Monitoring (RMON) Overview

This release of ArubaOS Mobility Access Switch supports RMON, which provides standard information that a network administrator can use to monitor, analyze, and troubleshoot a group of distributed local area networks (LANs). Monitoring devices (commonly called "probes") contain RMON software agents that collect information and analyze packets. These probes act as servers and the Network Management applications that communicate with them act as clients. While both agent configuration and data collection use SNMP, RMON is designed to operate differently than other SNMP-based systems:

- Probes have more responsibility for data collection and processing, which reduces SNMP traffic and the processing load of the clients.
- Information is only transmitted to the management application when required, instead of continuous polling.

This release of ArubaOS supports the following RMON groups:

- ethernet statistics
- history control
- ethernet history
- alarm
- event

Enabling RMON Service

You can use the following command to enable RMON service on the Mobility Access Switch:

```
(host)(config)# service rmon
```

The `service rmon` command is disabled by default. When the `service rmon` command is disabled, the `rmon` data is not populated in the CLI display command but all the other configurations can be done. When the `service rmon` command is enabled, all the configurations done before would be applied.

Configuring RMON Parameters

Configuring the Alarm

Table 1 describes the alarm parameters

Table 1 Alarm Configuration Parameters

Parameter	Description
alarm-profile	To associate an alarm profile.
monitor	Configures an OID to monitor.
owner	Configures an owner of this alarm entry.

You can use the following command to associate the alarm profile with the alarm entry:

```
(host)(config)#rmon alarm <alarm_index>
(host)(alarm_index)#alarm-profile <alarm-profile-name>
```

You can use the following command to monitor an interface or OID:

```
(host)(alarm_index)#monitor <oid>
```

You can use the following command to monitor OID on gigabitethernet interface:

```
(host)(alarm_index)#monitor gigabitethernet <slot/module/port> oid-type <oid_types>
```

You can use the following command to monitor OID on port-channel interface:

```
(host)(alarm_index)#monitor port-channel <port-channel id> oid-type <oid_types>
```

Configuring the Alarm Profile

Table 2 describes the alarm-profile parameters

Table 2 Alarm Profile Configuration Parameters

Parameter	Description
falling-event	Associate an event index or profile to the falling event.
falling-threshold-value	Specifies the value at which the event is generated.
rising-event	Associate an event profile or index to the rising event.
rising-threshold-value	Specifies the value at which the event is generated.
sample-type	Specifies whether the sample type is either delta or absolute <ul style="list-style-type: none">When the sample-type is delta, the value of the selected variable at the last sample will be subtracted from the current value, and the difference is compared with the thresholds.When the sample-type is absolute, the value of the selected variable will be compared directly with the thresholds at the end of the sampling interval.
startup-alarm	Configures initial alarm (rising, falling, or either)

To configure the alarm variable, first you have to create an alarm profile. You can use the following command to create the alarm profile:

```
(host)(config)#rmon alarm-profile <profile-name>
```



```

falling-event<event-index>
falling-threshold-value <value>
interval<interval>
rising-event <event-index>
rising-threshold-value <value>
sample-type <absolute|delta>
startup-alarm {falling|rising|rising-or-falling}

```

Configuring Ethernet Statistics Index

Table 3 describes the ethernet statistics index parameters.

Table 3 Ethernet Statistics Index Configuration Parameters

Parameter	Description
monitor	Configures an OID to monitor.
owner	Configure the owner of the etherstat entry.

You can use the following command to configure ethernet statistics collection on an interface:

```
(host)(config)# rmon etherstat <etherstat-index>
```

You can use the following command to monitor an OID:

```
(host) (etherstat_index)#monitor <oid>
```

You can use the following command to monitor OID on gigabitethernet interface:

```
(host) (etherstat_index)#monitor gigabitethernet <slot/module/port>
```

You can use the following command to monitor OID on port-channel interface:

```
(host) (etherstat_index)#monitor port-channel <port-channel id>
```

Configuring History Group

Table 4 describes the history group parameters.

Table 4 History Group Configuration Parameters

Parameter	Description
monitor	Configures the OID to monitor.
owner	Configures the owner of the history entry.
samples	Number of samples
sampling-interval	Interval of each sample

You can use the following command to create the history group profile:

```

(host)(config)#rmon history <history-index>
    samples <number>
    sampling-interval <interval>
    owner <owner>

```

You can use the following command to monitor an OID:

```
(host) (history_index)#monitor <oid>
```

You can use the following command to monitor OID on gigabitethernet interface:

```
(host) (history_index)#monitor gigabitethernet <slot/module/port>
```

You can use the following command to monitor OID on port-channel interface:

```
(host) (history_index)#monitor port-channel <port-channel id>
```

Configuring Event Entry

Table 5 describes the event entry parameters.

Table 5 Event Entry Configuration Parameters

Parameter	Description
description	Configures description of the event.
owner	Configures owner of the event.
Type	<p>Specifies whether to send SNMPtrap or create log entry when the event occurs.</p> <ul style="list-style-type: none">When type is log or log-and-trap, an RMON log entry is created when the event is triggered and sets the eventType in the RMON MIB to log or log-and-trap.When type is trap or log-and-trap, SNMP trap is generated.When type is none, no action is taken for this event.

You can use the following command to configure the event entry:

```
(host)(config)#rmon event <event-index>
```

You can use the following command to configure the event type:

```
(host) (event-index)#type
```

You can use the following command to clear the RMON log entries:

```
(host)# clear rmon log-table
```

Viewing RMON Active Configuration

You can use the following command to list the alarm-oids supported on device to use it as an alarm variable.

```
(host)#show rmon alarm-oid
```

Supported OID List

Object Name	Object Identifier
-----	-----
ifOutOctets	1.3.6.1.2.1.2.2.1.16
ifInUcastPkts	1.3.6.1.2.1.2.2.1.11
ifOutUcastPkts	1.3.6.1.2.1.2.2.1.17
ifOutBroadcastPkts	1.3.6.1.2.1.31.1.1.1.5
ifInErrors	1.3.6.1.2.1.2.2.1.14
ifHCInOctets	1.3.6.1.2.1.31.1.1.1.6
ifHCInUcastPkts	1.3.6.1.2.1.31.1.1.1.7
ifHCInMulticastPkts	1.3.6.1.2.1.31.1.1.1.8
ifHCOutMulticastPkts	1.3.6.1.2.1.31.1.1.1.12
ifHCOutBroadcastPkts	1.3.6.1.2.1.31.1.1.1.13

You can use the following command to display the RMON event table information:

```
(host)#show rmon event-table
```

RMON Event Table:

Event Index	Type	Last Seen	Description	Owner
1	log and Trap	10-25-2011@19-28-16	desc_log_1	admin
4	log	-	desc_log_2	guest

You can use the following command to display the log table information. The latest log entry will be displayed as the first one:

```
(host) #show rmon log-table
```

RMON Log Table:

Log Id	Event Id	Creation Time	Description
1	3	3-22-2012@23-39-43	Rising threshold log: ifHCInOctets.455

You can use the following command to display the log table based on an event index:

```
(host)#show rmon log-table event <event-id> log <log-id>
```

You can use the following command to display the alarms on the device either briefly or detailed on alarm entry index basis:

```
(host)# show rmon alarms {brief | entry <index>}
```

The following command displays the details on the alarm on the device:

```
(host)#show rmon alarms brief
```

Total: 1 entry

RMON Alarm Table:

RMON Alarm Table

Alarm Index	Variable	Rising Threshold Value	Falling Threshold Value	Owner
1	ifInErrors.8	10	0	config

```
(host) #show rmon alarms entry 1
```

```
Alarm 1 is active, owned by config
  Monitors ifHCInMulticastPkts.1 every 10 seconds
  Taking delta sample, last value was 0
  Rising threshold value is 300, assigned to event 1
  Falling threshold value is 100, assigned to event 1
```

You can use the following command to display the history table either briefly or detailed on history entry index basis:

```
(host)# show rmon history {brief | entry <index>}
```

The following example displays the history table information:

```
(host)#show rmon history brief
```

Total: 1 entry

RMON History Table

History Index	Interface	Octets	Pkts	Bcast Pkts	MCast Pkts	Utilization
---------------	-----------	--------	------	------------	------------	-------------

```
1          gigabitethernet0/0/1  1323196  19594  0          19554          17
```

```
(host) #show rmon history entry 1
```

```
Entry 1 is active, and owned by config
  Monitors gigabitethernet0/0/0 every 1800 seconds
  Buckets requested 50, Buckets granted 50
  0 sample(s) created
```

Viewing RMON Configuration

You can use the following list of commands to display the RMON configurations which may or may not get applied. For active configuration, see [“Viewing RMON Active Configuration” on page 322](#).

You can use the following command to display the configuration done for a specific alarm-profile:

```
(host)#show rmon-config alarm-profile [profile-name]
```

You can use the following command to display the configuration for a specific alarm entry:

```
(host)#show rmon-config alarm [index]
```

You can use the following command to display the configuration done for a specific etherstat index:

```
(host)#show rmon-config etherstat [index]
```

You can use the following command to display the configuration done for a specific event index.

```
(host)#show rmon-config event [index]
```

You can use the following command to display the configuration done for a specific history index:

```
(host)#show rmon-config history [index]
```

ArubaOS Mobility Access Switch supports versions 1, 2c, and 3 of Simple Network Management Protocol (SNMP) for reporting purposes only. In other words, SNMP cannot be used for setting values in an Aruba system in the current Mobility Access Switch.



Aruba-specific management information bases (MIBs) describe the objects that can be managed using SNMP.

SNMP Parameters for Mobility Access Switch

You can configure the following SNMP parameters for the Mobility Access Switch.

Table 1 *SNMP Parameters for the Mobility Access Switch*

Parameter	Description
Read Community Strings	Community strings used to authenticate requests for SNMP versions lower than version 3.
Enable Trap Generation	Activates the SNMP trap generation functionality. The configured SNMP trap receivers will receive the generated traps when this option is enabled.
Trap/Inform receivers	Host information about a trap receiver. This host needs to be running a trap receiver to receive and interpret the traps sent by the Mobility Access Switch. Configure the following for each host/trap receiver: <ul style="list-style-type: none"> IP address SNMP version: can be 1, 2c, or 3. Community string UDP port on which the trap receiver is listening for traps. The default is the UDP port number 162. This is optional, and will use the default port number if not modified by the user.
If you are using SNMPv3 to obtain values from the ArubaOS Mobility Access Switch, you can configure the following parameters:	
User name	Name of the user.
Authentication protocol	An indication of whether messages sent on behalf of this user can be authenticated, and if so, the type of authentication protocol used. This can take one of the two values: <ul style="list-style-type: none"> MD5: HMAC-MD5-96 Digest Authentication Protocol SHA: HMAC-SHA-96 Digest Authentication Protocol
Authentication protocol password	The (private) authentication key for use with the authentication protocol, if messages sent on behalf of this user can be authenticated. This is a string password for MD5 or SHA depending on the choice above.

Table 1 *SNMP Parameters for the Mobility Access Switch*

Parameter	Description
Privacy protocol	An indication of whether messages sent on behalf of this user can be protected from disclosure, and if so, the type of privacy protocol which is used. This can take one of the following values: <ul style="list-style-type: none"> • DES (Data Encryption Standard) • AES (Advanced Encryption Standard) NOTE: Under DES, only CBC-DES Symmetric Encryption Protocol is supported.
Privacy protocol password	The (private) privacy key for use with the privacy protocol, if messages sent on behalf of this user can be encrypted/decrypted with DES.
Context	SNMP v3 context information used in SNMP agent.
Engine ID	Agent engine ID for SNMPv3.
SNMP Server Group	View access group entry for SNMPv3
View	SNMP view entry. The view entry is associated with an OID. This is used for configuring groups and community strings.

Configuring SNMPv1/v2c Parameters

Execute the following commands to configure the basic SNMP v1/v2c parameters:

```
(host)(config) #snmp-server community <string> view <view-name>
(host)(config) #snmp-server enable trap
(host)(config) #snmp-server host <ipaddr> version {1 <security-string>} | {2c
<security-string> [inform] [interval <seconds>] [retrycount <number>]} udp-port <port>
all auth generic ptopo rmon snmp stacking system vlan
(host)(config) #snmp-server inform queue-length <size>
(host)(config) #snmp-server trap source <ipaddr>
```

Example

The following is a sample SNMP v2c configuration:

```
(host)(config) #snmp-server community public view V2c_View
(host)(config) #snmp-server enable trap
(host)(config) #snmp-server host 10.13.6.70 version 2c public rmon stacking udp-port
4050
(host)(config) #snmp-server inform queue-length 250
(host)(config) #snmp-server trap source 10.13.7.80
```

Configuring SNMPv3 Parameters

Execute the following commands to configure the basic SNMP v3 parameters:

```
(host)(config) # snmp-server context <context-name>
(host)(config) #snmp-server view <view-name> oid-tree <OID> {included | excluded}
(host)(config) #snmp-server group <group-name> {v1 | v2c | [v3 {auth|no-auth|priv}]}
[context-prefix <name> context-match {exact|prefix}] notify <notify-view-name> read
<read-view-name>}
(host)(config) #snmp-server engine-id <engineid>
(host)(config) #snmp-server user <user-name> group <name> {v1 | v2c | {v3[auth-prot
{md5|sha} <password>] [priv-prot {AES|DES} <password>]}}
(host)(config) #snmp-server host <ipaddr> version 3 <user-name> [engine-id <engineid>]
[inform] [interval <seconds>] [retrycount <number>] udp-port <port> all auth generic
ptopo rmon snmp stacking system vlan
```

Example

You can use the following sample commands to configure SNMP v3:

To do SNMPv3 Get/GetNext operation:

```
(host) (config) #snmp-server view V3-View oid-tree ifTable included
(host) (config) #snmp-server view V3-View oid-tree ifName.0 excluded
(host) (config) #snmp-server community public view V3-View
```

To send SNMPv3 Traps:

```
(host) (config) #snmp-server context V3-Context
(host) (config) #snmp-server view V3-View oid-tree ifTable included
(host) (config) #snmp-server view V3-View oid-tree ifName.0 excluded
(host) (config) #snmp-server group V3-Group v3 auth notify ALL read V3-View context-
prefix V3-Context context-match exact
(host) (config) # snmp-server user V3-User group V3-Group v3 auth-prot md5 abcd1234
(host) (config) #snmp-server host 10.13.6.66 version 3 V3-User engine-id
8000052301A9FEA484 vlan
```

Viewing SNMP Configuration Parameters

You can use the following show commands to view the SNMP configuration details on the Mobility Access Switch:

- `show snmp group-snmpp`: View the View Access Group information populated from the snmpd process.
- `show snmp group-trap`: View the View Access Group information populated from the trapd process.
- `show snmp view`: View the View information with the included and excluded OID details.
- `show snmp context`: View the list of context names configured on the Mobility Access Switch.
- `show snmp community`: View the SNMP community table.
- `show snmp user-table`: View the user-table entries.
- `show snmp trap-hosts`: View the target trap host entries.
- `show snmp trap-group`: View the list of trap filter groups that can be applied while configuring trap hosts. You can also view the traps associated with a specific trap filter.
- `show snmp notify filter profile-name`: View the SNMP Target profile names.
- `show snmp engine-id`: View the SNMP engine ID.
- `show snmp inform stats`: View the SNMP inform statistics.
- `show snmp trap-list`: View the list of SNMP traps supported and their status.
- `show snmp trap-queue`: View the list of SNMP traps in queue.

Supported Standard MIBs

The following table gives the list of supported standard MIBs, supported tables in each MIB, and the scalars that are not supported in each MIB:

Table 2 *Supported MIBs*

MIB Name	Supported Tables	Scalars Not Supported
RFC1213-MIB	<ul style="list-style-type: none">• ipNetToMediaTable• tcp Globals• tcpConnTable• udp Globals• udpConnTable• sysinfo	—

Table 2 *Supported MIBs*

MIB Name	Supported Tables	Scalars Not Supported
IF-MIB(RFC 1213, ifXTable RFC 2233, RFC 2863)	<ul style="list-style-type: none"> ifTable ifXtable ifTableLastChange 	<ul style="list-style-type: none"> ifOutDiscards ifOutErrors ifInUnknownProtos ifInNUcastPkts ifOutNUcastPkts
EtherLike-MIB(RFC 3635)	<ul style="list-style-type: none"> dot3StatsTable 	<ul style="list-style-type: none"> dot3StatsSQETestErrors dot3StatsSymbolErrors dot3StatsEtherChipSet dot3StatsCarrierSenseErrors dot3StatsInternalMacTransmitErrors dot3StatsRateControlAbility dot3StatsRateControlStatus dot3StatsAlignmentErrors dot3StatsSingleCollisionFrames
ALARM-MIB-1(RFC 3877)	<ul style="list-style-type: none"> alarmModelTable alarmActiveStatsTable alarmClearTable 	—
NOTIFICATION-LOG(RFC3014)	<ul style="list-style-type: none"> Notification MIB(Globals) nImConfigLogTable 	—
SNMP-MPD-MIB(RFC 2572)	<ul style="list-style-type: none"> snmpMPDStats 	—
SNMP-FRAMEWORK-MIB(RFC 2571)	<ul style="list-style-type: none"> snmpEngine 	—
SNMPv2-MIB(RFC 1907)	—	<ul style="list-style-type: none"> snmpInTooBig snmpInNoSuchNames snmpInBadValues snmpInReadOnly snmpInGenErrs snmpInTotalReqVars snmpInTotalSetVars snmpInGetRequests snmpInGetNexts snmpInSetRequests snmpInGetResponses snmpInTraps snmpOutTooBig snmpOutNoSuchNames snmpOutBadValues snmpOutGenErrs snmpOutGetRequests snmpOutGetNexts snmpOutSetRequests snmpOutGetResponses snmpOutTraps
SNMP-TARGET-MIB(RFC 2573)	<ul style="list-style-type: none"> snmpTargetObjects snmpTargetAddrTable snmpTargetParamsTable 	—

Table 2 *Supported MIBs*

MIB Name	Supported Tables	Scalars Not Supported
SNMP-NOTIFICATION-MIB(RFC 2573)	<ul style="list-style-type: none"> • snmpNotifyTable • snmpNotifyFilterProfileTable • snmpNotifyFilterTable 	—
Q-BRIDGE-MIB(RFC 4363)	<ul style="list-style-type: none"> • dot1qBase • dot1qFdbTable • dot1qTpFdbTable • dot1qStaticUnicastTable • dot1qVlanStaticTable 	—
BRIDGE-MIB(RFC 4188)	<ul style="list-style-type: none"> • dot1dBase • dot1dTpFdbTable • dot1dStaticTable • dot1dBasePortTable 	—
PTOPO-MIB(RFC 2922)	<ul style="list-style-type: none"> • ptopoConnTable 	—
LLDP-MIB	<ul style="list-style-type: none"> • lldpPortConfigTable • lldpConfigManAddrTable • lldpStatsTxPortTable • lldpStatsRxPortTable • lldpLocPortTable • lldpLocManAddrTable • lldpRemTable • lldpRemManAddrTable 	—
RMON-MIB(RFC 2819)	<ul style="list-style-type: none"> • etherStatsTable • historyControlTable • etherHistoryTable • alarmTable • eventTable • logTable 	—
RMON2-MIB (RFC 4502)	<ul style="list-style-type: none"> • probeConfig 	—

Table 2 *Supported MIBs*

MIB Name	Supported Tables	Scalars Not Supported
HC-RMON-MIB (RFC 3273)	<ul style="list-style-type: none"> etherStatsHighCapacityGroup etherHistoryHighCapacityGroup 	<ul style="list-style-type: none"> etherStatsHighCapacityOverflowPkts64Octets etherStatsHighCapacityPkts64Octets etherStatsHighCapacityOverflowPkts65to127Octets etherStatsHighCapacityPkts65to127Octets etherStatsHighCapacityOverflowPkts128to255Octets etherStatsHighCapacityPkts128to255Octets etherStatsHighCapacityOverflowPkts256to511Octets etherStatsHighCapacityPkts256to511Octets etherStatsHighCapacityOverflowPkts512to1023Octets etherStatsHighCapacityPkts512to1023Octets etherStatsHighCapacityOverflowPkts1024to1518Octets etherStatsHighCapacityPkts1024to1518Octets
OSPF-MIB	<ul style="list-style-type: none"> ospfGeneralGroup ospfAreaTable ospfStubAreaTable ospfIfTable ospfNbrTable ospfLsdbTable ospfExtLsdbTable 	<ul style="list-style-type: none"> ospfDemandExtensions ospfIfDemand ospfNbmaNbrPermanence ospfNbrHelloSuppressed ospfStubMetric ospfImportAsExtern ospfNbmaNbrPermanence ospfNbrHelloSuppressed ospfIfAuthKey ospfExtLsdbAdvertisement ospfLsdbAdvertisement
ENTITY-MIB	<ul style="list-style-type: none"> entityGeneral entPhysicalTable entLogicalTable entAliasMappingTable entPhysicalContainsTable 	<ul style="list-style-type: none"> entPhysicalMfgName entPhysicalAssetID entPhysicalUri entPhysicalHardwareRev entPhysicalAlias entPhysicalMfgDate entLPMappingTable



To get OID for ENTITY-MIB, a new MIB called ARUBA-VENDORTYPE has been added.

Supported Enterprise MIBs

The following table gives the list of supported enterprise MIBs, supported tables in each MIB, and the scalars that are not supported in each MIB:

Table 3 *Supported Enterprise MIBs*

MIB Name	Supported Tables	Scalars Not Supported
ARUBA-SYSTEMEXT	<ul style="list-style-type: none">• wlsxSysExtProcessorTable• wlsxSysExtStorageTable• wlsxSysExtMemoryTable• wlsxSysExtCardTable• wlsxSysExtFanTable• wlsxSysExtPowerSupplyTable	<ul style="list-style-type: none">• wlsxSysExtSwitchMasterIp• wlsxSysExtSwitchRole
ARUBA-SWITCH	<ul style="list-style-type: none">• wlsxSysXProcessorTable• wlsxSysXStorageTable• wlsxSysXMemoryTable	<ul style="list-style-type: none">• wlsxSwitchMasterIP• wlsxSwitchRole
ARUBA-USER	<ul style="list-style-type: none">• wlsxUserTable• wlsxUserSessionTimeTable	—
ARUBA-IFEXT	<ul style="list-style-type: none">• wlsxIfExtNPortTable	—
ARUBA-POE	<ul style="list-style-type: none">• wlsxPsePortTable• wlsxPseSlotTable	—
ARUBA-STACKING	<ul style="list-style-type: none">• wlsxStackMemberTable• wlsxStackProtoIfTable• wlsxStackTopoTable	—

Supported Standard Traps

The following table gives the list of supported standard traps:

Table 4 *Standard Traps*

Supported Traps
<ul style="list-style-type: none">• authenticationFailure• coldStart• linkDown• linkUp• warmStart• ptopoConfigChange• lldpRemTablesChange• risingAlarm• fallingAlarm• ospfIfStateChange• ospfNbrStateChange• entConfigChange

Supported Enterprise Traps

The following table gives the list of supported enterprise traps:

Table 5 *Supported Enterprise Traps*

Supported Traps
<ul style="list-style-type: none">• wlsxAuthMaxAclEntries• wlsxAuthServerReqTimedOut• wlsxColdStart• wlsxFanFailure• wlsxFanOK• wlsxFanTrayInsertedTrap• wlsxFanTrayRemovedTrap• wlsxFlashSpaceOK• wlsxInRangeVoltage• wlsxInformQueueOverFlow• wlsxLowMemory• wlsxLowOnFlashSpace• wlsxMemoryUsageOK• wlsxNAuthMaxAclEntries• wlsxNAuthServerIsDown• wlsxNAuthServerIsUp• wlsxNAuthServerReqTimedOut• wlsxNFanFailure• wlsxNGBICInserted• wlsxNLowMemory• wlsxNLowOnFlashSpace• wlsxNOutOfRangeTemperature• wlsxNOutOfRangeVoltage• wlsxNProcessDied• wlsxNUserEntryAuthenticated• wlsxNUserEntryCreated• wlsxNUserEntryDeAuthenticated• wlsxNUserEntryDeleted• wlsxNormalTemperature

Table 5 *Supported Enterprise Traps*

Supported Traps
<ul style="list-style-type: none">• wlsxOutOfRangeTemperature• wlsxOutOfRangeVoltage• wlsxPowerSupplyFailureTrap• wlsxPowerSupplyMissingTrap• wlsxPowerSupplyOK• wlsxPowerSupplyOKTrap• wlsxProcessDied• wlsxProcessRestart• wlsxStackIfStateChangeTrap• wlsxStackTopologyChangeTrap• wlsxUserAuthenticationFailed• wlsxUserEntryAuthenticated• wlsxUserEntryChanged• wlsxUserEntryCreated• wlsxUserEntryDeAuthenticated• wlsxUserEntryDeleted• wlsxVlanLinkDown• wlsxVlanLinkUp• wlsxWarmStart• wlsxIfStateChangeTrap (Enhanced for BPDU guard feature)

